

An Explainable Hybrid Machine Learning Framework for Smart Irrigation Prediction and Anomaly-Aware Decision Support

MD Ashfi Ar Raihan¹, Sabrina Islam Tesa¹, Antu Chakrabarty¹, Mhamud Hasan Kwoshick¹, Tawfiqul Hasan¹, Sultanul Arifeen Hamim¹, Md. Mazid-Ul-Haque¹, Abhijit Bhowmik¹, and Md Manirul Islam¹

Department of Computer Science, American International University-Bangladesh (AIUB), 408/1, Kuratoli, Khilkhet, Dhaka-1229, Bangladesh
{23-53818-3, 23-53821-3, 23-53810-3, 23-53812-3, 23-54422-3}@student.aiub.edu, {hamim,mazid,abhijit,manirul}@aiub.edu

Abstract. Efficient irrigation management is important for sustainable agriculture because over-irrigation and under-irrigation can reduce water efficiency, soil quality, and crop productivity. Conventional irrigation practices often depend on fixed schedules, manual observation, or simple threshold rules, which are not sufficiently adaptive to changing soil and weather conditions. This study proposes an explainable hybrid machine learning framework for smart irrigation prediction and anomaly-aware decision support. A public smart farming sensor dataset was used, and selected soil, weather, crop, and irrigation-related features were considered. Since the dataset did not contain direct irrigation decision or failure labels, irrigation need and anomaly labels were constructed using soil moisture, temperature, humidity, and soil moisture change. SMOTE was applied to handle class imbalance, while CTGAN was used to evaluate the effect of synthetic tabular data augmentation through a CTGAN versus non-CTGAN comparison. Multiple machine learning classifiers were evaluated using accuracy, F1-score, ROC-AUC, and cross-validation. The CTGAN-enhanced Random Forest model achieved 0.9989 accuracy and 0.9898 F1-score for irrigation need prediction, while CatBoost achieved 0.9540 accuracy and 0.9541 F1-score for anomaly-aware warning. However, because the irrigation target was constructed from variables that were also available to the classifier, the near-perfect irrigation result primarily indicates successful recovery of the predefined pseudo-labeling rules rather than independently validated agronomic prediction. SHAP and LIME were used for explanation, and the final recommendation engine converted model outputs into practical irrigation actions.

Keywords: Smart irrigation, machine learning, anomaly detection, CTGAN, SMOTE, SHAP, LIME, decision support.

1 Introduction

Efficient irrigation management is a major challenge in modern agriculture because water availability, soil condition, weather variation, and crop demand

change continuously. Conventional irrigation practices often depend on fixed schedules, manual observation, or simple soil-moisture thresholds, which may cause over-irrigation, under-irrigation, water loss, soil degradation, and reduced crop productivity. Recent studies show that machine learning can support smart irrigation by analyzing soil moisture, temperature, humidity, rainfall, and related environmental variables [1]. Similar work has also shown that machine learning-based irrigation support can reduce water usage and improve soil moisture forecasting [2]. However, many existing studies mainly focus on irrigation prediction or scheduling, while anomaly-related conditions such as leakage, abnormal moisture change, and possible system failure are less frequently integrated into the same decision-support framework [3].

A practical solution is to combine irrigation need prediction, anomaly-aware warning, explainability, and recommendation generation within one data-driven framework. Sensor and IoT-based irrigation studies have emphasized the importance of real-time monitoring and multi-source agricultural data for better irrigation control [4]. Recent studies also show that machine learning methods can perform well in irrigation-related prediction and control tasks, but there is still a need for compact frameworks that connect prediction results with practical field-level decisions [5]. Therefore, a framework that can handle class imbalance, compare synthetic augmentation, explain model behavior, and generate recommendations can make irrigation decision support more useful.

To address these issues, this study proposes an explainable hybrid machine learning framework for smart irrigation prediction and anomaly-aware decision support. A public smart farming sensor dataset was used, and selected soil, weather, crop, and irrigation-related features were considered. Since the dataset did not contain direct irrigation decision or real failure labels, irrigation need and anomaly labels were constructed using soil moisture, temperature, humidity, and soil moisture change. Multiple machine learning models were evaluated under CTGAN and non-CTGAN settings, SMOTE was used for imbalance handling, SHAP and LIME were used for explanation, and the final outputs were converted into practical irrigation recommendations. The main contributions of this study are summarized as follows:

- A hybrid irrigation decision-support framework is proposed by integrating irrigation need prediction and anomaly-aware warning generation.
- Irrigation need and anomaly labels are constructed from sensor-based soil and environmental conditions because direct decision and failure labels are not available in the original dataset.
- SMOTE and CTGAN are used to address class imbalance and small dataset limitations, and a CTGAN versus non-CTGAN comparison is performed.
- Multiple machine learning models are evaluated using accuracy, precision, recall, F1-score, ROC-AUC, and cross-validation, with final model selection mainly based on F1-score.
- SHAP and LIME are incorporated to explain global and local model predictions.
- A recommendation engine is developed to convert prediction and anomaly outputs into practical irrigation actions.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the methodology. Section 4 discusses the experimental results. Section 5 discusses the findings. Section 6 concludes the paper and outlines future work. Finally, Section 7 provides details on code and data availability.

2 Literature Review

Machine learning (ML) and artificial intelligence (AI) have driven transformative advances across domains where high-stakes decision-making and trustworthiness are essential. A major share of this research focuses on clinical and cognitive healthcare, ranging from comprehensive frameworks for cognitive assessment [6] and comparative benchmarks for dementia identification [7] to interpretable fuzzy reasoning for automated clinical diagnosis [8]. In neurodegenerative care, deep transfer-learning ensembles have markedly enhanced Alzheimer's disease classification [9], complemented by integrated Internet-of-Healthcare-Things (IoHT) paradigms [10], pervasive sensor systems [11], and inclusive AI architectures tailored for broader neurological disorders [12]. Parallel developments in neurodevelopmental conditions highlight automated psychological screening [13], IoT-based assessment platforms [14], and assistive fuzzy interfaces for pediatric care [15]. Beyond cognitive health, diagnostic machine learning has proven effective across varied medical imaging and signal modalities. Recent breakthroughs include structural brain MRI segmentation [16], bio-inspired optimization for brain pathology classification [17], few-shot grading of diabetic retinopathy [18], convolutional thermal imaging for breast cancer screening [19], and acoustic analysis of respiratory patterns [20]. Critical state-of-the-art syntheses have mapped diagnostic avenues for epileptic seizure prediction [21] and polycystic ovary syndrome (PCOS) detection [22]. At the sensor level, learning algorithms increasingly process high-dimensional physiological signals, from raw extracellular neural recordings [23] and in-vivo artifact removal [24] to fuzzy-ensemble EEG emotion recognition [25] and wearable stress tracking [26]. In personal mobility and assistive environments, ambient intelligence supports continuous patient safety through emotion-aware fall monitoring [27], multimodal kinetic fusion [28], indoor navigation for patients with progressive neurodegeneration [29], and neuro-fuzzy controllers for autonomous wheelchairs [30]. Closely aligned human-centered pipelines extend into behavioral analytics, addressing continuous cognitive engagement [31], sentiment-driven academic performance quantification [32], and data-driven anomaly detection in daily human activities [33]. Underpinning these intelligent services is the need for robust, decentralized, and tamper-resistant infrastructure. Recent work addresses next-generation 6G medical connectivity [34], Bayesian-blockchain frameworks for smart healthcare security [35], and end-to-end protective protocols for medical IoT services [36]. Similar security-by-design strategies secure low-power edge nodes via lightweight S-MQTT architectures [37], software-defined fuzzy intrusion detection [38], distributed KYC verification in banking [39], dynamic trust evaluation

in ad-hoc mobile networks [40], and provenance tracking in agricultural supply chains [41]. Beyond clinical and security contexts, intelligent optimization has demonstrated success in cyber-physical automation—including adaptive fuzzy controllers for photovoltaic arrays [42], recurrent networks for commodity market forecasting [43], forensic camera source attribution [44], and autonomous aquabots for precision pond management [45]. Taken together, alongside the pedagogical shift toward Industry 5.0 educational frameworks [46] and formal explainability paradigms [47], these developments reflect a decisive transition toward multimodal, transparent, and context-aware intelligent systems.

In addition, the use of data-driven methods in irrigation management has increased because water scarcity, irregular rainfall, soil variability, and inefficient irrigation practices directly affect crop production. Younes et al. [1] reviewed machine learning applications for smart irrigation and reported that neural-network and ensemble-based approaches often achieved strong performance, with median accuracy values around 0.95 and median R^2 values around 0.87. Lephondo et al. [2] applied machine learning for irrigation support and reported soil moisture forecasting performance with a correlation value of about 96%. Umutoni and Samadi [3] reviewed ML-based irrigation decision-making and found that gradient boosting achieved an RMSE of 0.11 and accuracy of 92.7%, while a binary tree classifier reached about 95% success rate. These studies show that machine learning can improve irrigation scheduling compared with fixed-rule and manual irrigation practices.

Recent studies have also emphasized IoT, sensor integration, real-time irrigation management, and water-related agricultural prediction. Nsoh et al. [4] reported that some multi-source irrigation models achieved R^2 values between 0.7602 and 0.9181, with RMSE reduction between 32.505% and 57.125%. Del-Coco et al. [5] reported high performance in irrigation control tasks, including 0.98 pump activation accuracy and 0.97 soil moisture prediction accuracy. Vallejo-Gómez et al. [48] highlighted the importance of soil moisture, temperature, humidity, and wireless sensor networks in automated irrigation systems. Vij et al. [49] showed that IoT and machine learning can reduce manual intervention in farm irrigation. Mokhtar et al. [50] predicted irrigation water requirements using several ML and deep learning models, where DNN achieved RMSE = 0.055 mm and CNN/DNN models produced R^2 values above 0.90. El Bilali and Taleb [51] applied machine learning to irrigation water quality prediction in a semi-arid environment. Hussein et al. [52] used ML for irrigation water quality index prediction, where XGBoost achieved $R^2 = 0.971$, RMSE = 2.827, and MAE = 2.221.

Anomaly detection and explainability are important for making irrigation systems more reliable and transparent. Benameur et al. [53] proposed a low-cost irrigation anomaly detection system using deep learning and reported that the best autoencoder model achieved 97.60% accuracy and 97.60% F1-score. Adikisson et al. [54] showed that autoencoder-based methods can detect abnormal behavior in smart farming sensor environments. Cartolano et al. [55] analyzed explainable AI models for smart agriculture and showed the importance of expla-

nation methods in agricultural prediction. Abdul Razak et al. [56] emphasized that explainability is necessary for trust in agriculture-related decision-making. Bouni et al. [57] showed that SHAP and LIME can support global and local explanation in agricultural recommendation tasks. SMOTE [58] is widely used for minority-class oversampling, while CTGAN [59] is useful for synthetic tabular data generation. SHAP [60] and LIME [61] are also commonly used for interpreting global feature importance and local prediction behavior.

From the reviewed studies, it is clear that previous works have improved irrigation prediction, IoT-based automation, water requirement estimation, water quality prediction, anomaly detection, and explainable agricultural modeling. However, many studies focus on only one part of the problem, such as irrigation scheduling, water quality estimation, or sensor automation. Fewer studies combine irrigation need prediction, anomaly-aware warning, class imbalance handling, synthetic tabular augmentation, explainability, and recommendation generation in one compact framework. Moreover, many irrigation datasets do not contain direct irrigation decision labels or real failure labels. Therefore, this study addresses these gaps by constructing irrigation need and anomaly labels from sensor-based conditions, comparing CTGAN and non-CTGAN settings, applying SMOTE, using SHAP and LIME, and generating practical irrigation recommendations from model outputs.

3 Methodology

This study follows a hybrid irrigation decision-support pipeline that combines irrigation need prediction, anomaly-aware warning, class imbalance handling, CTGAN-based augmentation, explainability, and recommendation generation. No new machine learning algorithm was developed; instead, existing classifiers were evaluated for two connected tasks: binary irrigation need prediction and anomaly-aware classification. The overall workflow of the proposed framework is shown in Fig. 1.

3.1 Dataset and Preprocessing

The dataset used in this study was the publicly available *Smart Farming Sensor Data for Yield Prediction* dataset from Kaggle [62]. It contains soil, weather, crop, irrigation, and sensor-related farming records. Since the objective of this work is irrigation decision support, non-essential attributes such as farm ID, sensor ID, region, location, crop yield, and disease status were removed from the predictive feature set. The date field was not used as a model input. Soil-moisture change was calculated using consecutive records before the unnecessary identifier information was excluded from model training. The features used were soil pH, soil moisture, temperature, humidity, sunlight hours, rainfall, irrigation type, and crop type. Missing numerical values were handled using the median, while categorical missing values were handled using the mode. Categorical variables were encoded, and scaling was applied for scale-sensitive classifiers such as Logistic Regression, KNN, and SVM.




Fig. 1: Workflow of the proposed smart irrigation prediction and anomaly-aware decision-support framework.

3.2 Irrigation and Anomaly Label Construction

The original dataset did not contain direct irrigation decision labels. Therefore, a binary pseudo-label named *Irrigation_Need* was constructed using soil moisture, temperature, and humidity:

$$Y_i^{irr} = \begin{cases} 1, & \text{if } SM_i < 30 \wedge Temp_i > 25 \wedge Hum_i < 60, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The thresholds used for pseudo-label construction were treated as heuristic operating conditions for this experimental framework rather than universal agronomic limits. They were selected to represent comparatively dry soil together with warm and low-humidity conditions under which irrigation demand would be more likely. Because irrigation requirements vary by crop, soil type, growth stage, and local climate, these thresholds should not be interpreted as fixed recommendations for all agricultural environments. Here, $Y_i^{irr} = 1$ indicates that irrigation is required, while $Y_i^{irr} = 0$ indicates that irrigation is not required.

For anomaly-aware warning, soil moisture change was calculated as:

$$\Delta SM_i = SM_i - SM_{i-1} \quad (2)$$

Here, SM_i represents the soil-moisture value of the current record and SM_{i-1} represents the soil-moisture value of the preceding record in the available data

sequence. The date variable itself was not supplied to the machine-learning classifiers. Using soil moisture and moisture change, the anomaly label was constructed as:

$$Y_i^{anom} = \begin{cases} 1, & \text{if } SM_i > 80, \\ 2, & \text{if } SM_i < 15, \\ 3, & \text{if } \Delta SM_i < -20, \\ 4, & \text{if } \Delta SM_i = 0, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Similarly, the anomaly thresholds were used to create interpretable experimental warning categories from the available sensor variables. They provide a rule-based basis for evaluating the anomaly-aware pipeline but require agronomic and field validation before operational use. Class 0 represents normal condition, class 1 represents over-irrigation risk, class 2 represents under-irrigation risk, class 3 represents possible leakage, and class 4 represents possible system anomaly. In the available dataset, observed records were available only for the normal, under-irrigation, and possible-leakage categories. No observed samples satisfied the rules for over-irrigation or system anomaly. Therefore, classes 1 and 4 are retained only as rule-defined decision categories in the framework and were not empirically validated as observed anomaly classes.

3.3 SMOTE, CTGAN, and Model Evaluation

The available records were divided into training and test subsets before model evaluation. The test subset was kept separate from model fitting and was used to evaluate the trained classifiers. SMOTE was applied only to the training data so that synthetic minority samples were not directly introduced into the test set. The CTGAN experiment was conducted as a separate augmentation setting and was compared with the non-CTGAN baseline. Therefore, the reported CTGAN and non-CTGAN results should be interpreted as an experimental comparison on this dataset rather than as evidence of external generalization. A synthetic minority sample is generated as:

$$x_{new} = x_i + \lambda(x_{nn} - x_i) \quad (4)$$

where x_i is a minority-class sample, x_{nn} is one of its nearest minority-class neighbors, and λ is a random value between 0 and 1.

CTGAN was used to generate synthetic tabular farming records and evaluate whether generative augmentation could improve model performance under a small dataset setting. The synthetic records were labeled using the same irrigation and anomaly label rules. Two experimental settings were compared:

$$\begin{aligned} D_{baseline} &= SMOTE(D_{real}), \\ D_{CTGAN} &= SMOTE(D_{real} \cup D_{synthetic}). \end{aligned} \quad (5)$$

This CTGAN versus non-CTGAN comparison is a key novelty of the study because synthetic data were not used only for augmentation; their effect was experimentally compared against a non-CTGAN SMOTE baseline.

The evaluated classifiers included Logistic Regression, KNN, SVM, Naive Bayes, Decision Tree, Random Forest, XGBoost, LightGBM, CatBoost, and selected ensemble settings. Performance was evaluated using accuracy, precision, recall, F1-score, ROC-AUC, and 10-fold cross-validation. F1-score was used as the main model-selection metric because the irrigation and anomaly targets were imbalanced. Cross-validation was used as a supporting measure to examine the stability of model performance across different data partitions. Since the data were imbalanced, final model selection was mainly based on F1-score:

$$M^* = \arg \max_{M_j} F1(M_j) \quad (6)$$

where M^* represents the selected classifier and $F1(M_j)$ represents the F1-score of the j -th candidate classifier.

3.4 Explainability and Recommendation Engine

SHAP and LIME were used to support model transparency. SHAP was applied to identify global feature importance, while LIME was used to explain individual predictions. The final recommendation engine combines irrigation output, anomaly output, and input conditions to generate practical irrigation actions:

$$D_i = \begin{cases} \text{Irrigation needed,} & \hat{y}_i^{irr} = 1 \wedge \hat{y}_i^{anom} = 0, \\ \text{No irrigation needed,} & \hat{y}_i^{irr} = 0 \wedge \hat{y}_i^{anom} = 0, \\ \text{Potential over-irrigation warning,} & \hat{y}_i^{anom} = 1, \\ \text{Monitor under-irrigation,} & \hat{y}_i^{anom} = 2, \\ \text{Possible leakage detected,} & \hat{y}_i^{anom} = 3, \\ \text{Potential system-anomaly warning,} & \hat{y}_i^{anom} = 4. \end{cases} \quad (7)$$

This recommendation layer is another novelty of the study because the framework converts prediction and anomaly outputs into actionable irrigation decisions instead of stopping at classification results.

4 Results

Since the dataset was imbalanced, F1-score was considered the main evaluation metric, while accuracy, ROC-AUC, and cross-validation performance were used as supporting measures.

4.1 Class Distribution Analysis

The dataset contained 500 records after preprocessing and pseudo-label construction. Table 1 shows the original irrigation need distribution, the balanced distribution after SMOTE, and the CTGAN-generated synthetic distribution. The original target was imbalanced, with 438 non-irrigation samples and only

Table 1: Class Distribution of the Irrigation Need Target

Class	Meaning	Original	Original %	SMOTE	CTGAN
0	Irrigation not required	438	87.60%	350	384
1	Irrigation required	62	12.40%	350	116
Total	–	500	100%	700	500

62 irrigation-required samples. After SMOTE, both classes were balanced with 350 samples each. CTGAN generated 500 additional synthetic records for evaluating the effect of synthetic tabular augmentation.

The anomaly target was also imbalanced, as shown in Table 2. Normal records were dominant, while under-irrigation and leakage appeared as minority classes. No original samples were found for over-irrigation and system failure. This distribution shows that the available dataset does not provide observed examples for all predefined anomaly categories. Therefore, performance for the missing over-irrigation and system-anomaly categories cannot be evaluated from the current data.

Table 2: Class Distribution of the Anomaly Target

Class	Meaning	Samples	Percentage
0	Normal condition	399	79.80%
1	Over-irrigation risk	0	0.00%
2	Under-irrigation	55	11.00%
3	Possible leakage	46	9.20%
4	Possible system anomaly	0	0.00%
Total	–	500	100%

4.2 CTGAN Synthetic Data Quality Evaluation

CTGAN was used to generate 500 synthetic farming sensor records. The generated data were evaluated using distribution comparison, correlation comparison, Kolmogorov–Smirnov testing, and Train on Synthetic, Test on Real validation. Although the KS test showed distributional differences between real and synthetic numerical features, TSTR validation achieved 0.9840 accuracy and 0.9836 F1-score for irrigation need prediction. Therefore, CTGAN was used as an augmentation strategy rather than as a replacement for real sensor records.

4.3 CTGAN versus Non-CTGAN Model Comparison

Table 3 compares model performance under non-CTGAN and CTGAN-enhanced settings. CTGAN did not improve all classifiers equally. Linear and distance-based models showed lower F1-scores after CTGAN augmentation, while Random Forest showed the strongest improvement. This supports the novelty of evaluating CTGAN experimentally instead of using synthetic data without comparison.

Table 3: CTGAN vs Non-CTGAN Performance Comparison

Model	Framework	Accuracy	F1-Score
Logistic Regression	Non-GAN	0.9214	0.9436
Logistic Regression	CTGAN	0.9468	0.7992
KNN	Non-GAN	0.9229	0.9281
KNN	CTGAN	0.9290	0.8742
SVM	Non-GAN	0.9843	0.9509
SVM	CTGAN	0.9734	0.9376
Naive Bayes	Non-GAN	0.9786	0.9694
Naive Bayes	CTGAN	0.9789	0.9248
Random Forest	Non-GAN	0.9986	0.9681
Random Forest	CTGAN	0.9989	0.9898
XGBoost	Non-GAN	0.9957	0.9792
XGBoost	CTGAN	0.9989	0.9792
KNN + NB Ensemble	Non-GAN	0.9657	0.9436
KNN + NB Ensemble	CTGAN	0.9612	0.9248
XGB + SVM Ensemble	Non-GAN	0.9971	0.9681
XGB + SVM Ensemble	CTGAN	0.9989	0.9584

4.4 Best Model Selection

The final models were selected mainly based on F1-score because both irrigation and anomaly targets were imbalanced. This near-perfect score should be interpreted in the context of the pseudo-label construction. Soil moisture, temperature, and humidity were used to define the irrigation target and were also available among the predictive variables. Therefore, the result mainly demonstrates that the model can reproduce the deterministic relationships used to generate the pseudo-labels. It does not by itself demonstrate equivalent performance for independent field irrigation decisions. For anomaly-aware warning, CatBoost achieved 0.9540 accuracy, 0.9573 precision, 0.9540 recall, and 0.9541 F1-score using 10-fold cross-validation. The selected final models are shown in Table 4.

Table 4: Selected Final Models

Task	Model	Acc.	F1
Irrigation prediction	Random Forest	0.9989	0.9898
Anomaly detection	CatBoost	0.9540	0.9541

4.5 Explainability and Recommendation Output

SHAP and LIME were used to improve model transparency. SHAP identifies global feature importance, while LIME explains individual predictions. Fig. 2 and Fig. 3 show the SHAP-based global feature importance for irrigation prediction and anomaly-aware warning, respectively.




Fig. 2: SHAP summary plot showing the global importance of input features for irrigation prediction.




Fig. 3: SHAP summary plot showing the global importance of input features for anomaly-aware warning.

LIME was additionally used for instance-level interpretation. Unlike SHAP, which summarizes feature influence across multiple observations, LIME explains a single prediction by showing which input conditions locally supported or opposed the predicted class. In the tested cases, the LIME output provided a local explanation of the irrigation decision based on the input conditions used by the selected classifier. Therefore, LIME complements the global SHAP analysis by allowing individual recommendations to be inspected separately.

The final recommendation engine converts model outputs into practical irrigation actions. This is a key novelty because the framework does not stop at prediction; it produces decision-support recommendations by combining irrigation and anomaly outputs.

Table 5: Decision and Recommendation Mapping

Irr. Output	Anom. Output	Recommendation
1	0	Irrigation needed
0	0	No irrigation needed
–	1	Stop irrigation; over-irrigation risk
–	2	Monitor field; under-irrigation alert
–	3	Possible leakage detected
–	4	System anomaly detected
Final configuration		Random Forest (Prediction) + Cat-Boost (Anomaly Detection)

```

Test Case 1: Soil Moisture: 95%,Temperature_C: 30C ,Humidity: 60% → Over-Irrigation Risk (Stop Irrigation)
Test Case 2: Soil Moisture: 5%,Temperature_C: 30C ,Humidity: 60% → Irrigation Needed (Soil Too Dry)
Test Case 3: Soil Moisture: 25%,Temperature_C: 30C ,Humidity: 60% → Irrigation Needed (Soil Too Dry)
Test Case 4: Soil Moisture: 45%,Temperature_C: 30C ,Humidity: 60% → Optimal Moisture (No Irrigation Needed)
Test Case 5: Soil Moisture: 70%,Temperature_C: 30C ,Humidity: 60% → Slightly Wet (Monitor Only)
Test Case 6: Soil Moisture: 85%,Temperature_C: 30C ,Humidity: 60% → Over-Irrigation Risk (Stop Irrigation)

```

Fig. 4: Rule-based test case outputs showing recommendations under different soil moisture, temperature, and humidity conditions.

```

-----
Sample 33:
Moisture: 24.77
Prediction: 0, Anomaly: [0]
Decision: Irrigation Needed (Soil Too Dry)
-----
Sample 34:
Moisture: 10.31
Prediction: 0, Anomaly: [1]
Decision: System Anomaly Detected
-----
Sample 35:
Moisture: 39.35
Prediction: 0, Anomaly: [0]
Decision: Optimal Moisture (No Irrigation Needed)
-----
Sample 36:
Moisture: 22.48
Prediction: 0, Anomaly: [0]
Decision: Irrigation Needed (Soil Too Dry)
-----
Sample 37:
Moisture: 39.61
Prediction: 0, Anomaly: [0]
Decision: Optimal Moisture (No Irrigation Needed)
-----

```

Fig. 5: Sample recommendation outputs generated by combining irrigation prediction, anomaly warning, and suggested decision.

5 Discussion

The experiments show that the proposed framework can reproduce the constructed irrigation labels and provide anomaly-aware decision-support outputs from the available smart-farming data. Among the evaluated models, the CTGAN-enhanced Random Forest produced the highest irrigation prediction score, while CatBoost produced the strongest anomaly-classification result. SHAP and LIME provide global and instance-level interpretation, and the recommendation layer translates model outputs into understandable irrigation-related actions.

The near-perfect irrigation result must, however, be interpreted carefully. Irrigation labels were created using fixed conditions based on soil moisture, temperature, and humidity, and these variables were also available to the classifiers. Therefore, the 0.9989 accuracy mainly indicates that the Random Forest successfully learned the relationships used to construct the pseudo-labels. It does

not demonstrate near-perfect prediction of independent farmer decisions, crop responses, water savings, or other agronomic outcomes. Furthermore, only 500 original records were available, and no independent validation across different farms, seasons, or geographical locations was performed. The present results should therefore be treated as an evaluation of the proposed decision-support pipeline within the available dataset rather than evidence of broad agronomic generalization.

Several additional limitations should be considered. The pseudo-label thresholds were used as heuristic experimental rules, and their sensitivity was not systematically evaluated. The dataset contained no observed samples for the predefined over-irrigation and system-anomaly categories, so these remain unvalidated rule-based warning categories. Repeated experiments with multiple independent random seeds and formal confidence intervals were also not evaluated. In practical deployment, sensor drift, calibration errors, changing environmental conditions, and false alerts may affect system reliability. False irrigation or anomaly warnings could result in unnecessary water use, crop stress, or unnecessary system inspection. Future work should therefore use independently observed irrigation decisions and failure events, larger multi-farm and temporal datasets, threshold sensitivity analysis, repeated-run statistical evaluation, and real-time field testing.

6 Conclusion

This study presented an explainable hybrid machine learning framework for smart irrigation prediction and anomaly-aware decision support. The framework combined irrigation need prediction, anomaly-aware warning, SMOTE-based imbalance handling, CTGAN-based synthetic augmentation, SHAP–LIME explainability, and recommendation generation. Since the dataset did not contain direct irrigation decision or failure labels, rule-based pseudo-labels were constructed using soil moisture, temperature, humidity, and soil moisture change. The CTGAN-enhanced Random Forest model achieved 0.9989 accuracy and 0.9898 F1-score for irrigation need prediction, while CatBoost achieved 0.9540 accuracy and 0.9541 F1-score for anomaly-aware warning. The recommendation engine converted model outputs into irrigation-related decision-support actions, including irrigation need, under-irrigation risk, and possible anomaly warnings. However, the irrigation target was generated from predefined sensor-based rules rather than independently observed irrigation decisions. Therefore, the near-perfect irrigation performance mainly represents recovery of the pseudo-labeling rules and should not be interpreted as field-level agronomic generalization. In addition, no observed samples were available for the predefined over-irrigation and system-anomaly categories. Future work should validate the framework using independent field-labelled irrigation decisions, observed failure events, larger multi-farm datasets, temporal evaluation, threshold sensitivity analysis, and real-time IoT sensor streams.

7 Code and Data Availability

The source code and data for this study are publicly available on GitHub at: <https://github.com/ashfi-raihan/Research/blob/main/IrrigationAnomaly/IrrigationAnomaly.ipynb>.

References

1. A. Younes, Z. E. A. El Assad, O. El Meslouhi, D. E. A. El Assad, and E.-d. A. Majid. The application of machine learning techniques for smart irrigation systems: A systematic literature review. *Smart Agricultural Technology*, 7:100425, 2024.
2. I. Lephondo, A. Telukdarie, I. Munien, U. Onkonkwo, and A. Vermeulen. The outcomes of smart irrigation system using machine learning to minimize water usage within the agriculture sector. *Procedia Computer Science*, 237:525–532, 2024.
3. L. Umutoni and V. Samadi. Application of ml approaches in supporting irrigation decision making: A review. *Agricultural Water Management*, 294:108710, 2024.
4. B. Nsoh, A. Katimbo, M. A. S. Al-Ismaili, J. D. Luck, D. M. Heeren, and V. Singh. Internet of things-based automated solutions utilizing machine learning for smart and real-time irrigation management: A review. *Sensors*, 24(23):7480, 2024.
5. M. Del-Coco, M. Leo, and P. Carcagnì. Machine learning for smart irrigation in agriculture: How far along are we? *Information*, 15(6):306, 2024.
6. Abdul Rehman Javed, Ayesha Saadia, Huma Mughal, Thippa Reddy Gadekallu, Muhammad Rizwan, Praveen Kumar Reddy Maddikunta, Mufti Mahmud, Madhusanka Liyanage, and Amir Hussain. Artificial Intelligence for Cognitive Health Assessment: State-of-the-Art, Open Challenges and Future Directions. *Cogn. Comput.*, 15(6):1767 – 1812, 2023.
7. Yunus Miah, Chowdhury Nazia Enam Prima, Sharmeen Jahan Seema, Mufti Mahmud, and M. Shamim Kaiser. Performance Comparison of Machine Learning Techniques in Identifying Dementia from Open Access Clinical Datasets. *Adv. Intell. Syst. Comput.*, 1188:79 – 89, 2021.
8. Tianhua Chen, Pan Su, Yinghua Shen, Lu Chen, Mufti Mahmud, Yitian Zhao, and Grigoris Antoniou. A dominant set-informed interpretable fuzzy system for automated diagnosis of dementia. *Front. Neurosci.*, 16, 2022.
9. Tanjim Mahmud, Koushick Barua, Anik Barua, Sudhakar Das, Nanziba Basnin, Mohammad Shahadat Hossain, Karl Andersson, M. Shamim Kaiser, and Nahed Sharmen. Exploring Deep Transfer Learning Ensemble for Improved Diagnosis and Classification of Alzheimer’s Disease. *Proc. BI 2023*, 13974 LNAI:109 – 120, 2023.
10. Sabrina Jesmin, M. Shamim Kaiser, and Mufti Mahmud. Artificial and Internet of Healthcare Things Based Alzheimer Care During COVID 19. *Proc. BI 2020*, 12241 LNAI:263 – 274, 2020.
11. Tajim Md. Niamat Ullah Akhund, Md. Julkar Nayeem Mahi, A.N.M. Hasnat Tanvir, Mufti Mahmud, and M. Shamim Kaiser. ADEPTNESS: Alzheimer’s disease patient management system using pervasive sensors - early prototype and preliminary results. *Proc. BI 2018*, 11309 LNAI:413 – 422, 2018.
12. Shamim Al Mamun, M. Shamim Kaiser, and Mufti Mahmud. An Artificial Intelligence Based Approach Towards Inclusive Healthcare Provisioning in Society 5.0: A Perspective on Brain Disorder. *Proc. BI 2021*, 12960 LNAI:157 – 169, 2021.

13. Sabbir Ahmed, Md. Farhad Hossain, Silvia Binte Nur, M. Shamim Kaiser, and Mufti Mahmud. Toward Machine Learning-Based Psychological Assessment of Autism Spectrum Disorders in School and Community. *Proc. TEHI 2021*, 376:139 – 149, 2022.
14. Md. Eftekhair Alam, M. Shamim Kaiser, Mohammad Shahadat Hossain, and Karl Andersson. An IoT-belief rule base smart system to assess autism. In *Proc. iCEE-iCT 2018*, pages 672 – 676, 2018.
15. Anjum Ismail Sumi et al. fASSERT: A fuzzy assistive system for children with autism using internet of things. *Proc. BI 2018*, 11309 LNAI:403 – 412, 2018.
16. Hafsa Moontari Ali, M. Shamim Kaiser, and Mufti Mahmud. Application of Convolutional Neural Network in Segmenting Brain Regions from MRI Data. *Proc. BI 2019*, 11976 LNAI:136 – 146, 2019.
17. B. Deepa, M. Murugappan, M.G. Sumithra, Mufti Mahmud, and Mabrook S. Al-Rakhami. Pattern Descriptors Orientation and MAP Firefly Algorithm Based Brain Pathology Classification Using Hybridized Machine Learning Algorithm. *IEEE Access*, 10:3848 – 3863, 2022.
18. M. Murugappan, N.B. Prakash, R. Jeya, A. Mohanarathinam, G.R. Hemalakshmi, and Mufti Mahmud. A novel few-shot classification framework for diabetic retinopathy detection and grading. *Meas.: J. Int. Meas. Confed.*, 200, 2022.
19. Mathew Jose Mammoottil, Lloyd J. Kulangara, Anna Susan Cherian, Prabu Mohandas, Khairunnisa Hasikin, and Mufti Mahmud. Detection of Breast Cancer from Five-View Thermal Images Using Convolutional Neural Networks. *J. Healthc. Eng.*, 2022, 2022.
20. Himadri Mukherjee, Priyanka Sreerama, Ankita Dhar, Sk. Md. Obaidullah, Kaushik Roy, Mufti Mahmud, and K.C. Santosh. Automatic Lung Health Screening Using Respiratory Sounds. *J. Med. Syst.*, 45(2), 2021.
21. Sobhana Jahan, Farhana Newsheen, Mahathir Mahmud Antik, Md. Sazzadur Rahman, M. Shamim Kaiser, A. S. M. Sanwar Hosen, and In-Ho Ra. AI-Based Epileptic Seizure Detection and Prediction in Internet of Healthcare Things: A Systematic Review. *IEEE Access*, 11:30690 – 30725, 2023.
22. Samia Ahmed, Md. Sazzadur Rahman, Ismate Jahan, M. Shamim Kaiser, A. S. M. Sanwar Hosen, Deepak Ghimire, and Seong-Heum Kim. A Review on the Detection Techniques of Polycystic Ovary Syndrome Using Machine Learning. *IEEE Access*, 11:86522 – 86543, 2023.
23. Mufti Mahmud and Stefano Vassanelli. Processing and analysis of multichannel extracellular neuronal signals: State-of-the-art and challenges. *Front. Neurosci.*, 10(JUN), 2016.
24. Marcos Fabietti, Mufti Mahmud, Ahmad Lotfi, Alberto Avarua, David Gugganmos, Randolph Xudo, and Michela Chiappalone. Neural Network-based Artifact Detection in Local Field Potentials Recorded from Chronically Implanted Neural Probes. In *Proc. IJCNN 2020*, 2020.
25. Trishita Dhara, Pawan Kumar Singh, and Mufti Mahmud. A Fuzzy Ensemble-Based Deep learning Model for EEG-Based Emotion Recognition. *Cogn. Comput.*, 16(3):1364 – 1378, 2024.
26. Sayandeep Ghosh, SeongKi Kim, Muhammad Fazal Ijaz, Pawan Kumar Singh, and Mufti Mahmud. Classification of Mental Stress from Wearable Physiological Sensors Using Image-Encoding-Based Deep Neural Network. *Biosensors*, 12(12), 2022.
27. M. Jaber Al Nahian, Tapotosh Ghosh, Mohammed Nasir Uddin, Md. Maynul Islam, Mufti Mahmud, and M. Shamim Kaiser. Towards Artificial Intelligence Driven

- Emotion Aware Fall Monitoring Framework Suitable for Elderly People with Neurological Disorder. *Proc. BI 2020*, 12241 LNAI:275 – 286, 2020.
28. Md. Nahiduzzaman, Moumitu Tasnim, Nishat Tasnim Newaz, M. Shamim Kaiser, and Mufti Mahmud. Machine Learning Based Early Fall Detection for Elderly People with Neurological Disorder Using Multimodal Data Fusion. *Proc. BI 2020*, 12241 LNAI:204 – 214, 2020.
 29. Milon Biswas, Ashiqur Rahman, M. Shamim Kaiser, Shamim Al Mamun, K. Shayekh Ebne Mizan, Mohammad Shahidul Islam, and Mufti Mahmud. Indoor Navigation Support System for Patients with Neurodegenerative Diseases. *Proc. BI 2021*, 12960 LNAI:411 – 422, 2021.
 30. M. Shamim Kaiser, Zamshed Iqbal Chowdhury, Shamim Al Mamun, Amir Hussain, and Mufti Mahmud. A Neuro-Fuzzy Control System Based on Feature Extraction of Surface Electromyogram Signal for Solar-Powered Wheelchair. *Cogn. Comput.*, 8(5):946 – 954, 2016.
 31. Muhammad Arifur Rahman, David J. Brown, Nicholas Shopland, Andrew Burton, and Mufti Mahmud. Explainable Multimodal Machine Learning for Engagement Analysis by Continuous Performance Test. *Proc. HCII 2022*, 13309 LNCS:386 – 399, 2022.
 32. Johanna Watkins, Marcos Fabielli, and Mufti Mahmud. SENSE: A Student Performance Quantifier using Sentiment Analysis. In *Proc. IJCNN*, pages 1–6, 2020.
 33. Salisu Wada Yahaya, Ahmad Lotfi, and Mufti Mahmud. Towards a data-driven adaptive anomaly detection system for human activity. *Pattern Recognit. Lett.*, 145:200 – 207, 2021.
 34. M. Shamim Kaiser, Nusrat Zenia, Fariha Tabassum, Shamim Al Mamun, M. Arifur Rahman, Md. Shahidul Islam, and Mufti Mahmud. 6g access network for intelligent internet of healthcare things: Opportunity, challenges, and research directions. *Adv. Intell. Syst. Comput.*, 1309:317 – 328, 2021.
 35. Fahiba Farhin, M. Shamim Kaiser, and Mufti Mahmud. Secured smart healthcare system: Blockchain and bayesian inference based approach. *Adv. Intell. Syst. Comput.*, 1309:455 – 465, 2021.
 36. Fahiba Farhin, M. Shamim Kaiser, and Mufti Mahmud. Towards Secured Service Provisioning for the Internet of Healthcare Things. In *Proc. AICT 2020*, 2020.
 37. Abdur Rahman, Shanto Roy, M Shamim Kaiser, and Md. Shahidul Islam. A lightweight multi-tier S-MQTT framework to secure communication between low-end IoT nodes. In *Proc. NSysS 2018*, 2018.
 38. Fahiba Farhin, Ishrat Sultana, Nahida Islam, M Shamim Kaiser, Md. Sazzadur Rahman, and Mufti Mahmud. Attack Detection in Internet of Things using Software Defined Network and Fuzzy Neural Network. In *Proc. ICIEV icIVPR 2020*, 2020.
 39. Abdullah Al Mamun, Sheikh Riad Hasan, Md Salahuddin Bhuiyan, M. Shamim Kaiser, and Mohammad Abu Yousuf. Secure and Transparent KYC for Banking System Using IPFS and Blockchain Technology. In *Proc. TENSYP 2020*, pages 348 – 351, 2020.
 40. Antesar M. Shabut, M. Shamim Kaiser, Keshav P. Dahal, and Wenbing Chen. A multidimensional trust evaluation model for MANETs. *J. Netw. Comput. Appl.*, 123:32 – 41, 2018.
 41. Sm Al-Amin, Shipra Rani Sharkar, M. Shamim Kaiser, and Milon Biswas. Towards a blockchain-based supply chain management for e-agro business system. *Adv. Intell. Syst. Comput.*, 1309:329 – 339, 2021.

42. Lotfi Farah, Amir Hussain, Abdelfateh Kerrouche, Cosimo Ieracitano, Jamil Ahmad, and Mufti Mahmud. A highly-efficient fuzzy-based controller with high reduction inputs and membership functions for a grid-connected photovoltaic system. *IEEE Access*, 8:163225 – 163237, 2020.
43. Oluwatamilore Orojo, Jonathan Tepper, T.M. McGinnity, and Mufti Mahmud. A Multi-recurrent Network for Crude Oil Price Prediction. In *Proc. SSCI 2019*, pages 2940 – 2945, 2019.
44. Md Hasan Al Banna, Md Ali Haider, Md Jaber Al Nahian, Md Maynul Islam, Kazi Abu Taher, and M. Shamim Kaiser. Camera model identification using deep CNN and transfer learning approach. In *Proc. ICREST 2019*, pages 626 – 630, 2019.
45. Md. Moniruzzaman Hemal, Atiqur Rahman, Nurjahan, Farhana Islam, Samsuddin Ahmed, M. Shamim Kaiser, and Muhammad Raisuddin Ahmed. An Integrated Smart Pond Water Quality Monitoring and Fish Farming Recommendation Aquabot System. *Sensors*, 24(11), 2024.
46. Y. Supriya, Dasari Bhulakshmi, Sweta Bhattacharya, Thippa Reddy Gadekallu, Pratik Vyas, Rajesh Kaluri, S. Sumathy, Srinivas Koppu, David J. Brown, and Mufti Mahmud. Industry 5.0 in Smart Education: Concepts, Applications, Challenges, Opportunities, and Future Directions. *IEEE Access*, 12:81938 – 81967, 2024.
47. Suman Lata Tripathi and Mufti Mahmud. *Explainable Machine Learning Models and Architectures*. 2023.
48. D. Vallejo-Gómez, M. Osorio, and C. A. Hincapié. Smart irrigation systems in agriculture: A systematic review. *Agronomy*, 13(2):342, 2023.
49. A. Vij, S. Vijendra, A. Jain, S. Bajaj, A. Bassi, and A. Sharma. Iot and machine learning approaches for automation of farm irrigation system. *Procedia Computer Science*, 167:1250–1257, 2020.
50. A. Mokhtar, N. Al-Ansari, W. El-Ssawy, R. Graf, P. Aghelpour, H. He, S. M. Hafez, and M. Abuarab. Prediction of irrigation water requirements for green beans-based machine learning algorithm models in arid region. *Water Resources Management*, 37:1557–1580, 2023.
51. A. El Bilali and A. Taleb. Prediction of irrigation water quality parameters using machine learning models in a semi-arid environment. *Journal of the Saudi Society of Agricultural Sciences*, 19(7):439–451, 2020.
52. E. E. Hussein, A. Derdour, B. Zerouali, A. Almaliki, Y. J. Wong, M. Ballesta-de los Santos, P. M. Ngoc, M. A. Hashim, and A. Elbeltagi. Groundwater quality assessment and irrigation water quality index prediction using machine learning algorithms. *Water*, 16(2):264, 2024.
53. R. Benameur, M. Benyamina, A. Benamar, and A. Khaldi. An innovative smart and sustainable low-cost irrigation system for anomaly detection using deep learning. *Sensors*, 24(4):1162, 2024.
54. M. Adkisson, A. K. Sangaiah, and S. R. Nayak. Autoencoder-based anomaly detection in smart farming ecosystem. In *Proceedings of the IEEE International Conference on Big Data*, pages 1–6, 2021.
55. A. Cartolano, A. Cuzzocrea, G. Pilato, and G. M. Grasso. Analyzing and assessing explainable ai models for smart agriculture environments. *Multimedia Tools and Applications*, 83:1–24, 2024.
56. S. F. Abdul Razak, S. Yogarayan, M. S. Sayeed, and M. I. F. Mohd Derafi. Agriculture 5.0 and explainable ai for smart agriculture: A scoping review. *Emerging Science Journal*, 8(2):744–760, 2024.

57. M. Bouni, B. Hssina, K. Douzi, and S. Douzi. Interpretable machine learning techniques for an advanced crop recommendation model. *Journal of Electrical and Computer Engineering*, 2024:7405217, 2024.
58. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
59. L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni. Modeling tabular data using conditional gan. In *Advances in Neural Information Processing Systems*, volume 32, pages 7335–7345, 2019.
60. S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.
61. M. T. Ribeiro, S. Singh, and C. Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144, 2016.
62. Atharva Soundankar. Smart farming sensor data for yield prediction. <https://www.kaggle.com/datasets/atharvasoundankar/smart-farming-sensor-data-for-yield-prediction>, 2024. Accessed: 27-Jun-2026.