Water and Soil Management and Modeling

Online ISSN: 2783 - 2546

University of Mohaghegh Ardabili

Optimizing machine learning to reduce crop classification uncertainty in semi-arid Canal command areas

Mohansing Rajaput [1]*, Abhilash Ramadasa [2], Basavanand M. Dodamani [3]

[1] Ph.D. Scholar, Department of Water Resources and Ocean Engineering, National Institute of Technology Karnataka, Surathkal, Mangaluru-575025, Karnataka, India.
[2] Scientist 'C', National Institute of Hydrology, Hard Rock Regional Centre, Visvesvaraya Nagar, Belagavi-590019, Karnataka, India.
[3] Professor, Department of Water Resources and Ocean Engineering, National Institute of Technology Karnataka, Surathkal, Mangaluru-575025, Karnataka, India.

Abstract

Effective water resource management and canal performance analysis in semi-arid regions are fundamentally reliant on precise estimates of crop water demand, which are typically derived from accurate, up-to-date crop inventories. For command areas in semi-arid India, this vital information is a primary input for agro-hydrological models used to assess irrigation efficiency and plan water allocation. However, the inherent complexity of these landscapes characterized by small, fragmented landholdings introduces substantial uncertainty into remote sensing based crop classification, threatening the reliability of subsequent management decisions. This study systematically addresses this input uncertainty by performing a comprehensive, multi-factorial sensitivity analysis using multi-temporal Sentinel-1 (SAR) and Sentinel-2 (optical) data. We investigated the combined effects of four multi-sensor data fusion strategies, six Machine Learning (ML) classifiers, three feature selection techniques, and five training/testing data split ratios. The findings of the study provide crucial operational insights for modelers and managers. The synergistic fusion of Sentinel-1 and Sentinel-2 data was identified as the single most critical factor for achieving high accuracy. Furthermore, classification performance showed high sensitivity to training data volume, with an optimal threshold observed at an 80/20 train/test split. The Extreme Gradient Boosting (XGBoost) classifier, coupled with Backward Elimination feature selection, emerged as the superior strategy, achieving a maximum overall accuracy of 98.4%. By identifying this optimized workflow, this research provides a robust and scalable method for generating highly reliable spatial input data, thereby minimizing uncertainty in crop water requirement calculations and significantly enhancing the predictive capacity and practical utility of agro-hydrological models for sustainable water management.

Keywords: Agro-hydrological modeling, Canal performance analysis, Crop classification, Semi-arid agriculture, Water resource management.

Article Type: Research Article
Academic Editor: Raoof Mostafazadeh

*Corresponding Author, E-mail: mohansingr@gmail.com
Citation: Rajaput, M., Abhilash, R., & Dodamani, B. M. (2026). Optimizing machine learning to reduce crop classification uncertainty in semi-arid Canal command areas. Water and Soil Management and Modeling, 6(2) (Special Issue: New Approaches to Water and Soil Management and Modeling), 303-325.
doi: 10.22098/mmws.2026.18944.1745
Received: 05 December 2025, Received in revised form: 30 January 2026, Accepted: 09 May 2026, Published online: 03 June 2026
Water and Soil Management and Modeling, Year 2026, Vol. 6, No.2 (Special Issue), pp. 303-325
Publisher: University of Mohaghegh Ardabili

1. Introduction

Effective water resource management and sustainable food security are paramount in arid and semi-arid regions, where water is the scarcest and most critical resource (Chang et al., 2023). In India, large-scale irrigation systems are managed under canal command areas, and the persistent challenge lies in accurately assessing their performance and ensuring equitable water distribution (Amarasinghe et al., 2021; Kumar et al., 2021). Evaluating the efficiency of these systems is typically accomplished using agro-hydrological models, which are essential tools for simulating water balance, determining irrigation schedules, and forecasting total water demand across the command area (Rajaput et al., 2025).

A core limitation in running these models and achieving reliable canal performance analysis is the lack of precise, timely, and scalable data on crop coverage. In the semi-arid landscapes typical of the region, the complex mosaic of small, highly fragmented farm plots and dynamic cropping calendars makes traditional ground surveys unreliable. Without an accurate spatial inventory of the crops being grown, the calculation of actual crop evapotranspiration (ETc) and the total irrigation water demand become uncertain, which directly hinders effective water allocation and data-driven agricultural policy (Abhilash et al., 2022).

To overcome this critical information deficit, Earth Observation (EO) technologies offer a viable and scalable solution. Freely available data from the Sentinel-1 (S1 - radar) and Sentinel-2 (S2 - optical) missions provide the necessary temporal and spatial resolution for monitoring agricultural dynamics (GSARS, 2017; Qader et al., 2021; Yao et al., 2022). The fusion of optical and all-weather radar data has proven to be an effective strategy for capturing both biophysical characteristics and structural information of crops (Eisfelder et al., 2024). When integrated with advanced Machine Learning (ML) classifiers, raw satellite data can be transformed into high-fidelity, field-level crop maps that directly feed agro-hydrological models (Gao et al., 2023; Maponya et al., 2020; Vizzari et al., 2024). It is important to note that this optimization of spatial input data falls under the broader, well-established domain of applying machine learning for predictive water management (Ahmed et al., 2024; Solomatine & Ostfeld, 2008), akin to applications in demand forecasting or water quality modeling. By refining the input parameters for physical models, such data-driven approaches significantly enhance the reliability of water resource planning.

However, despite the technical sophistication of these tools, a significant applied challenge remains. The accuracy of the final crop map and therefore the reliability of the derived water demand calculations is highly sensitive to the initial methodological choices made by the user. Most existing research focuses on demonstrating a specific algorithm or data product, neglecting the crucial interactive effects of the entire processing chain. Furthermore, the challenge of transferring models across different regions remains a recognized issue in the literature; however, developing a standard, reproducible methodology as presented in this work, offers a robust pathway for addressing this challenge at larger scales (Hoppe et al., 2024; Rusňák et al., 2023).

Key methodological decisions, such as the strategy used for multi-sensor data fusion, the optimization of the feature space through selection (FS) techniques, and the volume of training data used, introduce inherent uncertainty into the final crop map (Gao et al., 2023; Kabolizadeh et al., 2023; Ramezan, 2022). Without systematically understanding and optimizing these workflow components, the resulting spatial data may be insufficient or unreliable for the rigorous requirements of water resource modeling and decision-making. While Deep Learning (DL) has emerged as a powerful tool in this domain, foundational studies such as this one are essential. By providing transparent performance benchmarks and insights into feature importance and data requirements, this work constitutes an essential knowledge base required for the effective design and training of more complex deep learning models in the future (Maxwell et al., 2018; Nigar et al., 2024).

This study moves beyond a simple technical demonstration to develop and validate an

optimal, end-to-end classification workflow specifically designed to minimize input uncertainty for water management applications in a semi-arid canal command area. We conduct a comprehensive, multi-factorial sensitivity analysis to disentangle the effects of these key methodological choices.

The primary objectives of this study are: (1) To systematically determine the impact of data fusion strategy, feature selection technique, classifier choice, and training data size on crop classification accuracy, with the explicit goal of identifying the most robust combination for generating reliable inputs for agro-hydrological modeling. (2) To establish the critical performance threshold for training data size that balances mapping accuracy with the operational costs of data collection in complex, fragmented landscapes. (3) To leverage these systematic findings to establish a highly accurate and reproducible workflow that can directly support the operational assessment of irrigation performance and facilitate more precise water budgeting within the canal command area. (4) To create a practical, rigorous performance benchmark that provides actionable guidance for water managers deploying ML-based crop classification systems in similar semi-arid environments.

In the context of this study, 'actionable guidance' is operationally defined by the balance between classification accuracy and resource constraints. Specifically, we evaluate: (1) Optimal Data Allocation: establishing evidence-based protocols for partitioning limited ground truth datasets to ensure complex classifiers receive sufficient training data; (2) Computational Feasibility: validating efficient machine learning architectures that can be deployed on standard workstations without requiring specialized high-performance computing resources; and (3) Data Accessibility: utilizing standardized, open-access Earth Observation products to ensure the workflow is reproducible by non-specialist water management agencies.

By pursuing these objectives, this work directly contributes to the advancement of Water and Soil Management and Modelling by providing a validated, uncertainty-reduced methodology for quantifying the spatial extent of crops, a foundational requirement for sustainable water resource planning.

2. Materials and Methods

2.1 Study Area and Ground Truth Data

The study was conducted within the command area of the Narayanapura Left Bank Canal (NLBC) irrigation project, located in the Yadgir district of Karnataka, India (refer Figure 1). This area is situated in a semi-arid region where efficient water management is of paramount importance for agricultural productivity. The specific focus was on the command area served by distributaries D2, D3, and D4 of the Hunasagi Branch Canal, covering a gross command area of 4639.61 ha (46.396sqkm). The region is characterized by a complex cropping pattern with small, fragmented farm holdings, making it a challenging yet representative environment for testing advanced crop classification methodologies. The primary crops grown during the Kharif season include paddy, cotton, and pulses, while the Rabi season, which is the focus of this study, is dominated by paddy, jowar, groundnut, and bajra.

The analysis specifically focuses on the Rabi (winter) cropping season, which represents the primary irrigated agricultural cycle in this semi-arid command area. Unlike the Kharif (monsoon) season, which is largely rain-fed and subject to persistent cloud cover that limits optical remote sensing, the Rabi season relies heavily on canal water distribution, making accurate crop mapping during this period critical for irrigation scheduling and water demand estimation.

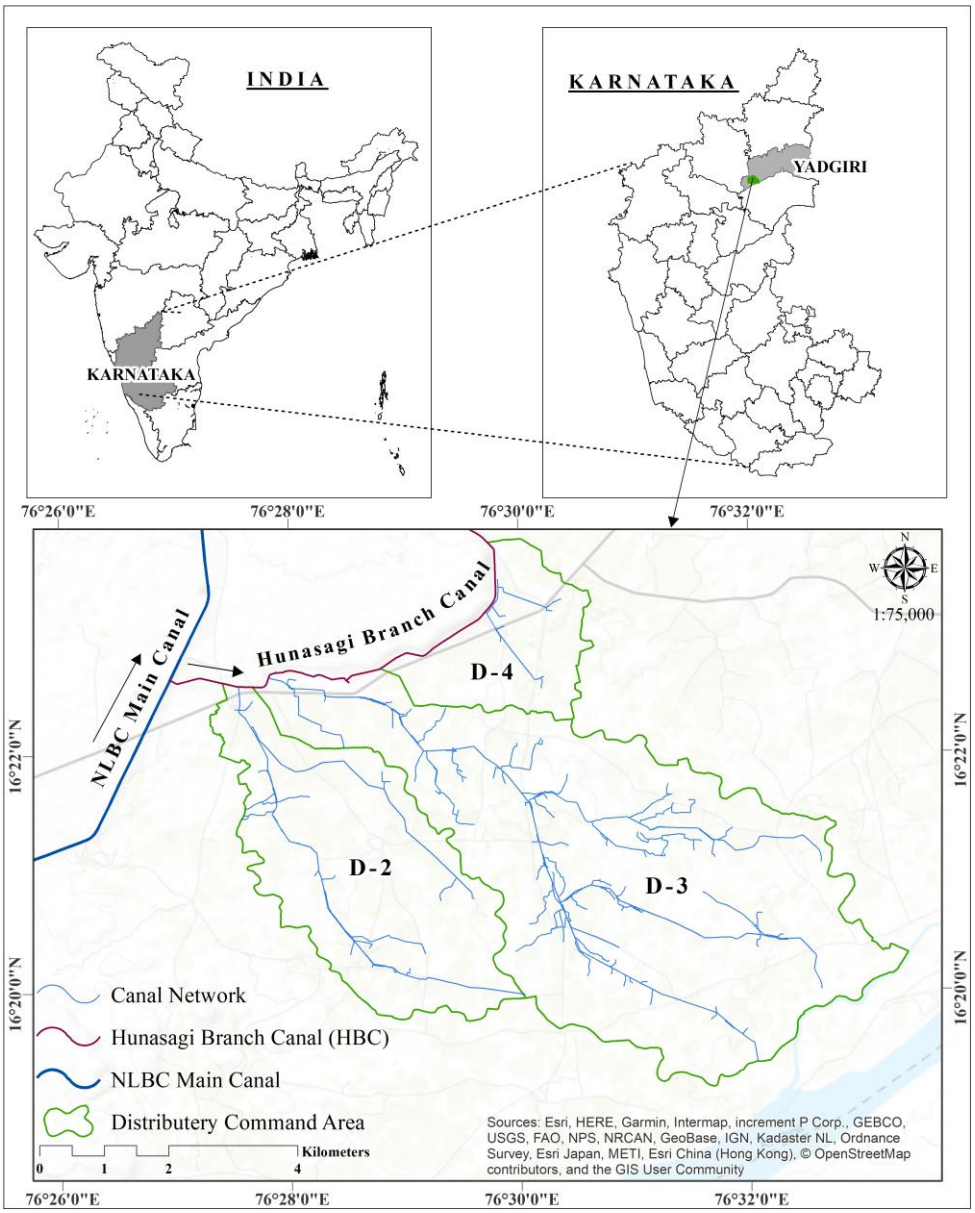

Figure 1. Location of the Study area.

Ground truth data, essential for training and validating the classification models, were collected through field surveys conducted during the 2020-2021 Rabi season. A total of 314 ground truth locations (polygons) were meticulously selected and geolocated. To ensure a robust sample size relative to the high-dimensional feature space, all valid pixels falling within these ground truth polygons were extracted for the analysis. This process generated a total dataset of over 20,000-pixel, ensuring a sampling density of approximately 7 samples per km² and satisfying the statistical requirement of having at least 10–30 times more samples than input features (Congalton & Green, 2019).

These locations represented the major crop types and other land cover classes present in the study area, including Paddy, Jowar, Groundnut, Bajra, Fallow Land, Built-up areas, and Waterbodies. The spatial distribution of these locations (refer Figure 2) was designed to capture the heterogeneity of the landscape and ensure representative sampling across the different distributary commands.

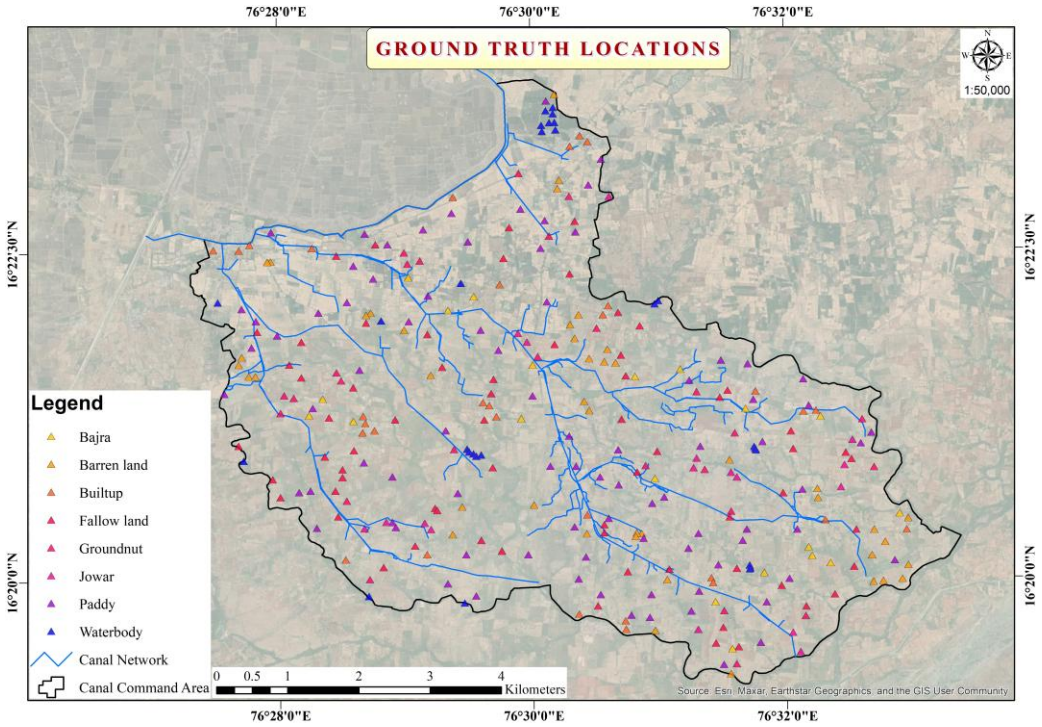

Figure 2. Spatial distribution of ground truth locations.

2.2 Satellite Data and Pre-processing

This study utilized a multi-temporal dataset of Sentinel-1 (S1) and Sentinel-2 (S2) imagery, acquired from the Copernicus Open Access Hub (ESA, 2024). The time series spanned from October 2020 to January 2021, coinciding with the main growing period of the Rabi season crops. This period was selected based on hydro-meteorological analysis, representing a normal water year (with 890mm rainfall) without significant cyclone impacts, thus ensuring that the observed spectral and backscatter signatures were primarily driven by crop phenology rather than extreme weather events.

Sentinel-1 Data: S1-A Ground Range Detected (GRD) products, acquired in the Interferometric Wide (IW) swath mode, were used. These data provide dual-polarization (VV and VH) backscatter information, which is sensitive to the canopy structure and moisture content of vegetation. The pre-processing workflow for S1 data was executed using the Sentinel Application Platform (SNAP) software and included the following steps: (1) application of the orbit file, (2) thermal noise removal, (3) radiometric calibration to convert digital numbers (DN) to sigma-naught (σ^0) backscatter coefficients, (4) speckle filtering using a Lee filter to reduce noise, and (5) range-Doppler terrain correction using the SRTM 30m DEM to correct for geometric distortions.

Sentinel-2 Data: Standard S2 Level-2A products were employed to utilize atmospherically corrected surface reflectance data. This selection ensures operational feasibility for water management applications by eliminating the need for complex, user-side atmospheric correction workflows. Images with less than 20% cloud cover were selected to ensure clear views of the study area. Ten spectral bands were chosen for the analysis, covering the visible (B2, B3, B4), red-edge (B5, B6, B7, B8, B8A), and short-wave infrared (B11, B12) regions, as these are most relevant for vegetation and land use classification.

To ensure spatial consistency and facilitate data fusion, all S1 and S2 images were resampled to a common spatial resolution of 10 m and re-projected to the UTM Zone 43N coordinate system.

2.3 Experimental Design: A Multi-Factorial Framework

To systematically address the research objectives, a multi-factorial experiment was designed to evaluate the influence and interaction of four primary factors on crop classification performance. This framework, illustrated in Figure 3, allows for a comprehensive analysis of the entire classification workflow, moving beyond a simple comparison of individual components. The total combinatorial experiment resulted in 360 unique model configurations, providing a rich dataset for a robust sensitivity analysis.

Factor 1. Data Source Synergy

To isolate the contributions of SAR data, optical data, derived vegetation indices, and their full synergy, four distinct datasets were prepared (refer Table 1). These datasets represent the four levels of this experimental factor.

Table 2 contains the equations of the extracted indices for preparing the feature datasets with references. The wide range of features and their possible correlations would lead to data redundancy in the methodology process. In order to optimize the feature set for optimum classification effectiveness, FS methods were used to find and keep the most pertinent features while removing redundant and irrelevant ones.

Table 1. Four datasets for evaluating crop classification accuracy.

Dataset	Description	No. of features	Data source
Dataset 1	Comprised the VV and VH polarization time series from Sentinel-1 images, along with the derived indices Simple Ratio-1 (SR1), Simple Ratio-2 (SR2), and Simple Ratio-3 (SR3), preparing total of 55-layer features based on available band layers.	VV - 11 VH - 11 SR1 - 11 SR2 - 11 SR3 - 11 from Oct 2020 to Jan 2021	Sentinel-1
Dataset 2	Consisted of the time series from Sentinel-2 Bands 2 to 8A, B11, and B12, producing 120-layer features.	12 images with total 10 bands from Oct 2020 to Jan 2021	Sentinel-2
Dataset 3	Included the time series of vegetation indices—Radar Vegetation Index (RVI) derived from Sentinel-1 with optical indices Enhanced Vegetation Index (EVI), Green Normalized Difference Vegetation Index (GNDVI), and Normalized Difference Vegetation Index (NDVI) extracted from Sentinel-2 images, producing a multi-sensor feature set of 47 layers.	RVI - 11 EVI - 12 GNDVI - 12 NDVI - 12 from Oct 2020 to Jan 2021	Sentinel-1 and Sentinel-2
Dataset 4	Combined the time series of Sentinel-1 images and their derived indices with the time series of Sentinel-2 images and their corresponding indices, producing 222-layer features.	DS1 and DS2 and DS3	Combination of all above.

Table 2. The Extracted Indices from Sentinel-1 and 2 satellite imagery in this study.

Feature	Formula	Reference
SR1	VV/VH	Kabolizadeh et al. (2023)
SR2	VH-VV	Kabolizadeh et al. (2023)
SR3	VH + VV	Kabolizadeh et al. (2023)
RVI	(4xVH)/(VV + VH)	Mandal et al. (2020)
NDVI	(NIR-R)/(NIR + R)	Rouse et al. (1974)
GNDVI	(NIR – Green) / (NIR + Green)	Huete et al. (2002)
EVI	2.5 * ((NIR – Red) / (NIR + 6 * Red – 7.5 * Blue + 1))	Gitelson et al. (1996)

Factor 2. Feature Space Optimization

To assess the impact of dimensionality reduction and the removal of redundant features, three widely used feature selection (FS) techniques were applied. These methods were chosen to represent different strategies for optimizing the feature space. •Sequential Forward Selection (SFS): An iterative wrapper method that starts with an empty feature set and sequentially adds the feature that provides the greatest improvement in model performance until a predefined criterion is met (Kabolizadeh et al., 2023; Montgomery et al., 2012). Recursive Feature Elimination (RFE): A backward wrapper method that begins with the full feature set, ranks features based on importance (derived from a base model), and recursively removes the least significant ones until the optimal subset is identified (He et al., 2022; Zhang et al., 2022). • Backward Elimination (BE): A statistical approach that starts with all features and iteratively removes the one whose removal results in the least significant loss of model performance, typically based on a statistical metric like p-value or accuracy change (Montgomery et al., 2012; Song et al., 2019).

Factor 3. Classifier Architecture

To evaluate performance across a spectrum of algorithmic complexity and learning paradigms, six popular and well-established machine learning classifiers were selected. These classifiers offer a balance between simple, interpretable models and advanced ensemble methods suitable for complex datasets. • Naïve Bayes (NB): A simple probabilistic classifier based on Bayes' theorem with a strong assumption of feature independence (Manning et al., 2008; Raja et al., 2022). • Classification and Regression Trees (CART): A decision tree-based model that provides an interpretable, rule-based classification but is prone to overfitting (Pal & Mather, 2003). • K-Nearest Neighbors (KNN): A non-parametric, instance-based learning algorithm that classifies a data point based on the majority class of its nearest neighbors (Saadatfar et al., 2020). • Support Vector Machine (SVM): A powerful classifier that seeks to find an optimal hyperplane that maximizes the margin between different classes in a high-dimensional space (Raja et al., 2022; Schölkopf & Smola, 2001). • Random Forest (RF): An ensemble method that constructs a multitude of decision trees on random subsets of the data and features, aggregating their predictions to improve accuracy and control overfitting (Biau & Scornet, 2016; Ok et al., 2012). • Extreme Gradient Boosting (XGBoost): An advanced and highly efficient implementation of the gradient boosting framework, which builds a strong model by sequentially adding weak learners that correct the errors of the previous ones (Chen & Guestrin, 2016; Zhang et al., 2022).

Factor 4. Training Data Dependency

To explicitly test the sensitivity of model performance to the volume of available training data, five different train-test splitting ratios were systematically evaluated. The 314 ground truth polygons were randomly partitioned into training and testing sets according to these ratios. Crucially, for each split, all pixel-level spectral signatures within the selected training polygons were extracted to train the classifier, while the testing polygons were strictly withheld for validation. The specific polygon splits were: • 50/50 (157 training, 157 testing) • 60/40 (188 training, 126 testing) • 70/30 (220 training, 94

testing) • 80/20 (251 training, 63 testing) • 90/10 (283 training, 31 testing).

This random partitioning was stratified to maintain a consistent class distribution across both subsets, minimizing potential biases caused by systematic variations(Uçar et al., 2020). Splits below 40–60% were excluded to avoid insufficient training data, which can compromise model learning and increase the risk of overfitting (Song et al., 2019). This approach ensured a comprehensive evaluation of how different data splits influence model performance.

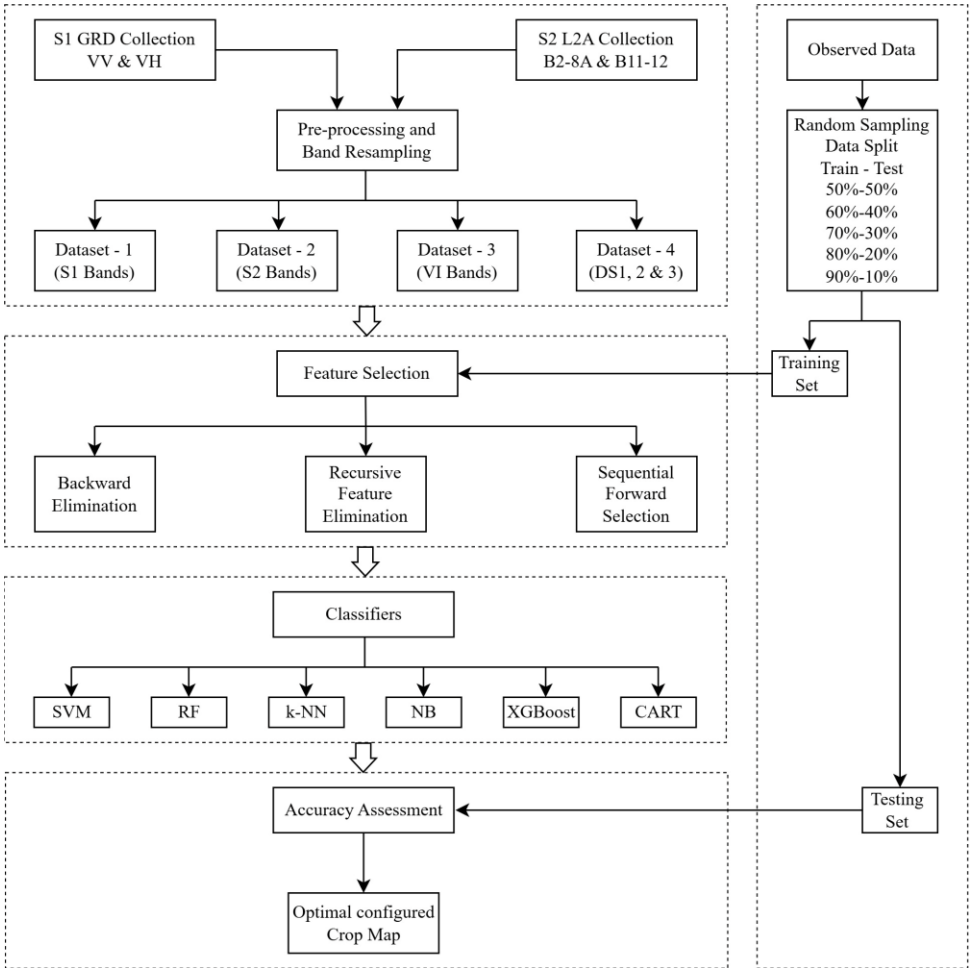

Figure 3. Methodology of the Study.

2.4 Classifier Configuration and Hyperparameters

To ensure a fair and reproducible comparison across the different machine learning architectures, effectively isolating the impact of input data fusion and feature selection, we adopted a standardized hyperparameter strategy. Rather than performing exhaustive grid-searching for each of the 360 model combinations (which could introduce bias by optimizing one classifier more successfully than others), we utilized nominal parameter settings consistent with best practices for baseline comparative studies (Table 3). A fixed random_state=42 was applied to all stochastic models (RF, SVM, XGBoost, CART) to ensure deterministic reproducibility.

2.5 Accuracy Assessment

The performance of each model configuration was rigorously evaluated using a confusion matrix derived from the independent test set. Three standard accuracy metrics were calculated to provide a comprehensive assessment of

classification performance: • Overall Accuracy (OA): The percentage of correctly classified samples across all classes, providing a single, intuitive measure of overall model performance. It is calculated as the sum of the diagonal elements of the confusion matrix divided by the total number of samples. • Kappa Coefficient (K): A statistical measure that compares the observed accuracy with the expected accuracy from random chance. It provides a more robust assessment than OA, especially in cases of class imbalance. Kappa values range from -1 to 1, where 1 indicates perfect agreement and 0 indicates agreement equivalent to chance. • F1-Score (Macro-Averaged): The harmonic mean of precision and recall. The macro-averaged F1-score calculates the F1-score for each class independently and then averages them, giving equal weight to each class. This metric is particularly useful for evaluating performance on multi-class problems and is sensitive to the performance on minority classes (He et al., 2022; Kabolizadeh et al., 2023).

Table 3. Hyperparameter settings for the evaluated classifiers.

Classifier	Library	Key Hyperparameters	Rationale
NB	sklearn.naive_bayes	GaussianNB() (Default)	Non-parametric baseline.
KNN	sklearn.neighbors	n_neighbors=5, metric='minkowski'	Standard default for local neighborhood classification.
SVM	sklearn.svm	kernel='linear', C=1.0	Linear kernel selected for computational efficiency in high-dimensional space.
RF	sklearn.ensemble	n_estimators=100, criterion='gini'	Lightweight ensemble baseline chosen to test performance with minimal computational complexity.
CART	sklearn.tree	criterion='gini', splitter='best'	Standard single decision tree configuration.
XGBoost	xgboost	booster='gbtree', learning_rate=0.3, max_depth=6	Standard gradient boosting defaults for multi-class classification.

3. Results

The multi-factorial experimental design yielded a comprehensive dataset of over 360 unique model runs, allowing for a detailed analysis of the factors influencing crop classification accuracy. The results are presented in a structured narrative, first identifying the single optimal configuration and then systematically dissecting the main and interactive effects of the four experimental factors. The performance of all combinatorial methods is summarized in the heat map in Figure 4.

3.1 Overview of Optimal Performance

The analysis revealed a clear optimal configuration that significantly outperformed all other combinations. The highest classification accuracy was achieved using the fully fused dataset (DS4), the Extreme Gradient Boosting (XGBoost) classifier, the Backward Elimination (BE) feature selection technique, and an 80/20 training-testing data split. This optimal configuration yielded an Overall Accuracy of 0.9835 and a macro-averaged F1-Score of 0.9834. The highest Kappa coefficient (0.9830) was achieved with a nearly identical setup, using the RFE feature selection method instead of BE, underscoring the robustness of this specific workflow configuration. The performance metrics for this top-performing model are summarized in Table 4.

To assess whether the class imbalance in the ground truth data biased the classification results, we examined the Producer's Accuracy (PA) and User's Accuracy (UA) for each class (Table 5). The results presented in Table 5 reflect the pooled confusion matrix derived from the k-fold cross-validation procedure. This represents the aggregated performance on the 314 validation polygons when they served as independent hold-out test sets. For this validation, we evaluated the classification accuracy based on the centroid of

each polygon, ensuring an object-oriented assessment.

Notably, this object-level accuracy calculated from the pooled matrix (98.4%) is nearly identical to the mean pixel-based accuracy derived from the cross-validation splits (98.35%), confirming the stability of the model across different spatial scales. Despite the significant disparity in sample sizes ranging from 84 samples for Paddy to only 17 for Jowar, the model demonstrated robust performance across all categories. Specifically, the minority crops Bajra (n=20) and Jowar (n=17) achieved exceptional accuracy, with both PA and UA exceeding 94%. These results confirm that the XGBoost classifier effectively resolved the spectral signatures of minority crops and that the high OA is not an artifact of masking poor performance in under-represented classes.

3.2 Analysis of Main Effects

To understand the individual contribution of each experimental factor, their main effects were analysed by aggregating performance metrics across all other variables. This approach isolates the general impact of data source, classifier choice, feature selection, and training data size, as visualized in the comparative strip plot in Figure 5.

The Critical Role of Data Fusion

The choice of input dataset had the most pronounced effect on classification accuracy. A visual comparison of the performance distributions for each dataset, aggregated across all classifiers, FS methods, and data splits (Figure 5a), clearly shows the superiority of the fused dataset. DS4, which synergistically combines S1 SAR, S2 optical, and derived vegetation indices, consistently yielded the highest accuracies. The median OA for DS4 was substantially higher than for DS1, DS2, or DS3. This demonstrates that the complementary information provided by the different sensors, structural and moisture information from SAR, and spectral-phenological information from optical is not merely additive but synergistic, enabling the models to better discriminate between spectrally similar crop types. DS1, relying solely on SAR data, exhibited the lowest overall performance, highlighting the limitations of using SAR alone for detailed crop type classification in this environment.

Classifier Performance Hierarchy

A distinct performance hierarchy emerged among the six tested classifiers. A box plot visualizing the OA distributions for each classifier (Figure 5b) would reveal that the ensemble methods, XGBoost, SVM and Random Forest, consistently outperformed the other models. XGBoost demonstrated the highest median accuracy and the tightest distribution of results, indicating both high performance and robustness across different experimental conditions. RF and SVM showed comparable, strong performance, ranking just below XGBoost. KNN performed moderately well, while the simpler models, CART and Naïve Bayes, consistently delivered the lowest accuracies. The poor performance of NB, in particular, can be attributed to its strong and often violated assumption of feature independence, which is ill-suited for the complex and correlated nature of multi-temporal remote sensing data.

While formal statistical significance tests (e.g., McNemar's test) were not conducted due to the aggregate nature of the multi-factorial analysis, the consistency of the performance hierarchy is evident in the distribution of results (Figure 6). XGBoost maintained the highest median accuracy across all data split ratios, with a notably tighter interquartile range compared to Random Forest and SVM. This empirical consistency across diverse experimental conditions suggests that the observed superiority of the gradient boosting architecture is a robust feature of the model rather than a product of random variation.

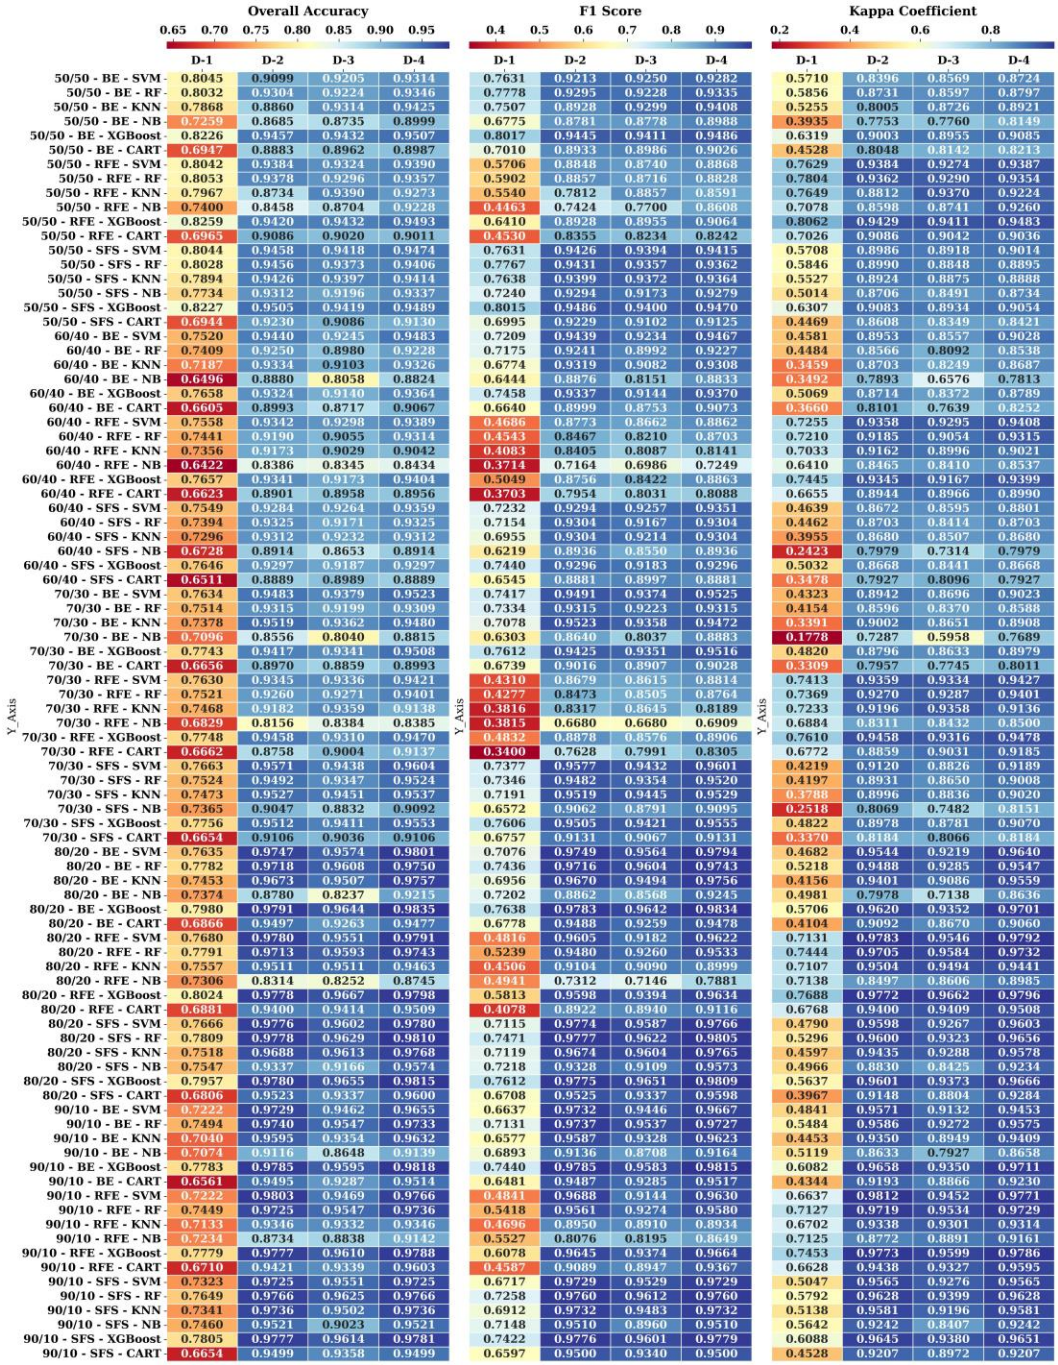

Figure 4. Heat map of crop classification accuracy for all combinations.

Table 4. Summary of top performing model.

Component	Optimal Choice	Metric	Value
Data Split	(S1 + S2 + VIs)	Overall Accuracy	0.9835
Classifier	XGBoost	Kappa Coefficient	0.9830*
Feature Selection	Backward Elimination (BE)	F1-Score	0.9834
Data Split	80/20	-	-

***Note: The highest Kappa was achieved with the RFE technique, but the performance with BE was nearly identical (0.9796 with an 80/20 split) and was optimal for both OA and F1-Score.**

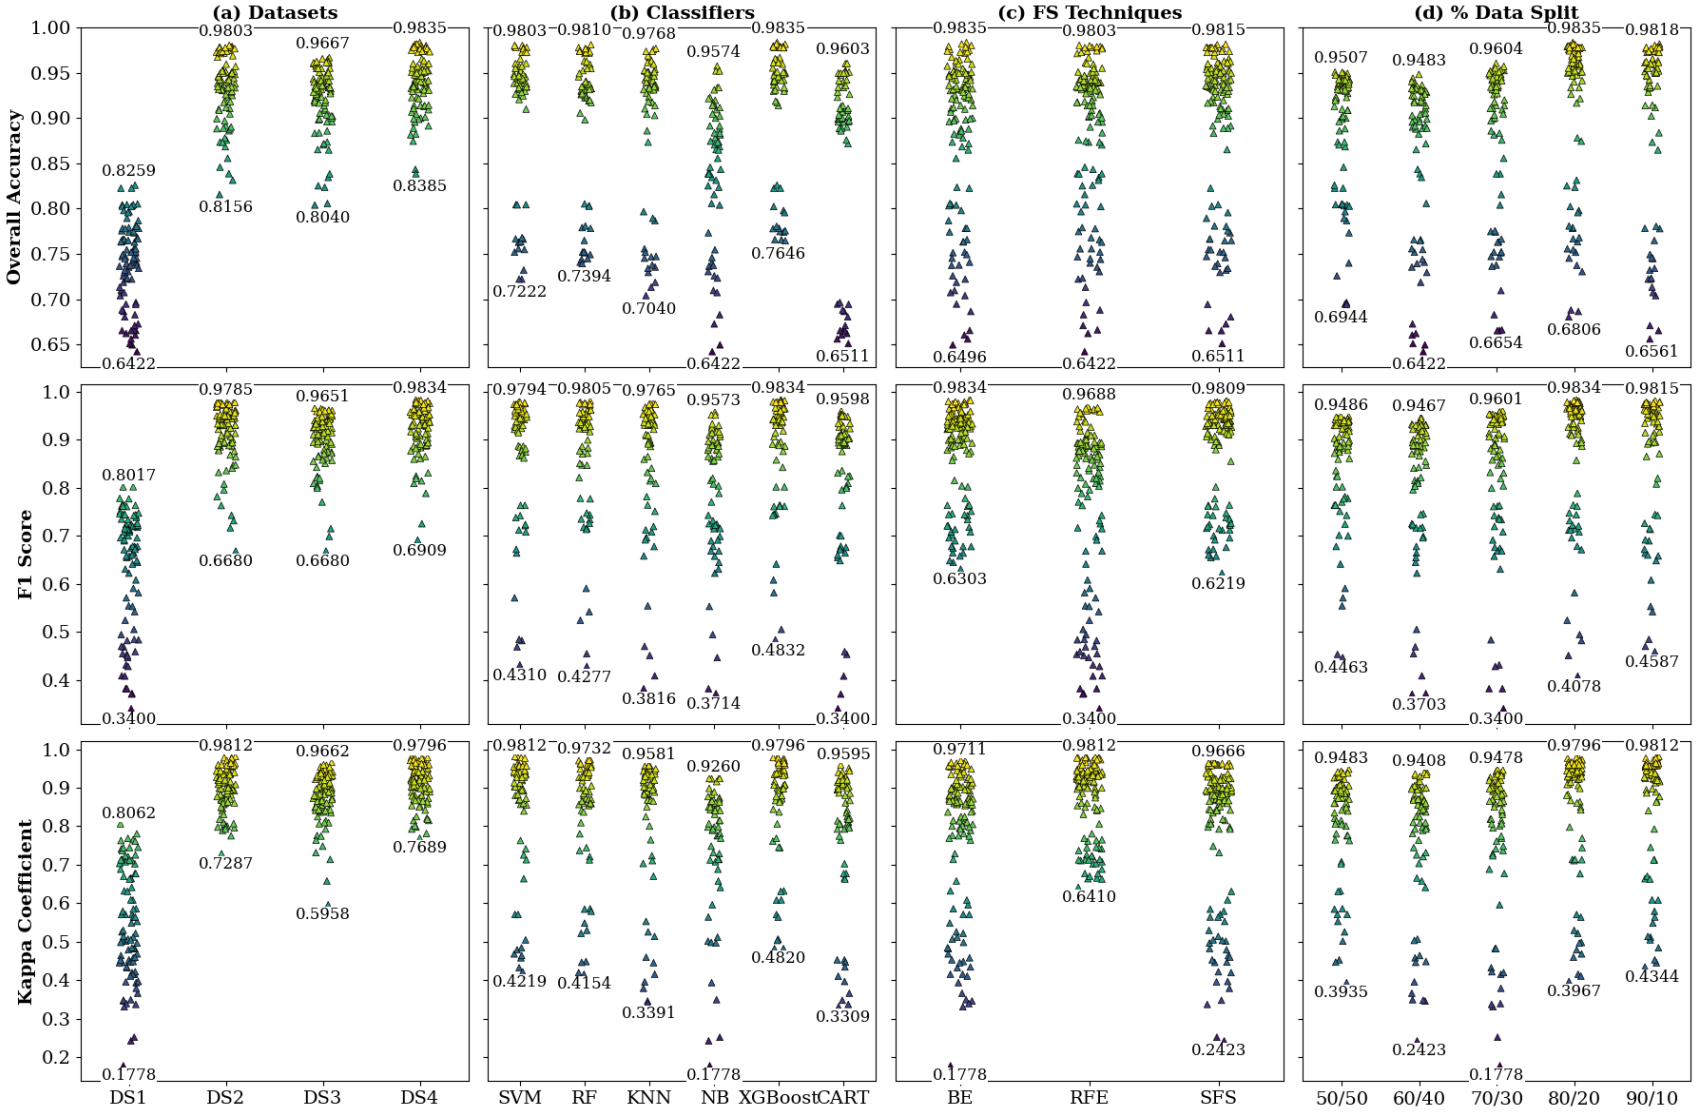

Figure 5. Strip plot for the OA, F1 score and kappa coefficient with maximum and minimum values.

Table 5. Confusion matrix for optimal combination DS4, XGboost, BE technique with 80/20 data split.

True \ Predicted	Bajra	Buildup	Fallowland	Groundnut	Jowar	Paddy	Barren land	Waterbody	Total	Prod. Accuracy %
Bajra (20)	19	0	0	0	1	0	0	0	20	**95.0**
Buildup (31)	0	31	0	0	0	0	0	0	31	**100.0**
Fallowland (67)	0	0	66	0	0	0	1	0	67	**98.5**
Groundnut (25)	0	0	0	25	0	0	0	0	25	**100.0**
Jowar (17)	1	0	0	0	16	0	0	0	17	**94.1**
Paddy (84)	0	0	0	0	0	83	0	1	84	**98.8**
Wasteland (43)	0	0	1	0	0	0	42	0	43	**97.7**
Waterbody (27)	0	0	0	0	0	0	0	27	27	**100.0**
Total Predicted	20	31	67	25	17	83	43	28	314	
User's Accuracy %	**95.0**	**100.0**	**98.5**	**100.0**	**94.1**	**100.0**	**97.7**	**96.4**		**OA: 98.4%**

The Consistent Benefit of Feature Selection

The application of feature selection techniques provided a consistent, albeit moderate, improvement in average classification accuracy. The process of removing redundant and irrelevant features helps to reduce model complexity, mitigate the risk of overfitting, and enhance predictive performance. When comparing the three FS techniques (BE, RFE, and SFS), their performance distributions were found to be very similar (Figure 5c). No single FS method demonstrated a consistent, statistically significant advantage over the others across all scenarios. This suggests that the primary benefit comes from the act of feature selection itself, rather than the specific algorithm used, at least for this set of well-established methods.

The Training Data Threshold

The analysis of data split ratios revealed a critical, non-linear relationship between the amount of training data and model performance. The strip plots of the performance metrics against the percentage of training data (Figure 5d) would illustrate this trend clearly. Performance increased as the training set size grew from 50% to 80%. However, the performance began to plateau between the 80/20 and 90/10 splits, with the accuracy becoming marginal. This finding suggests the existence of a 'training data threshold' around 80% for this specific problem and landscape complexity. Below this threshold, the models are likely under-trained, failing to capture the full spectral variability of the crop classes. Above this threshold, additional training data yields diminishing returns, indicating that the model has learned the salient patterns from

the available data. The lowest overall performance was observed with the 60/40 split, underscoring the risk of using insufficient training data in complex classification tasks.

3.3 Analysis of Interaction Effects and Feature Importance

Beyond the main effects, the factorial design allows for an investigation into the more nuanced interactions between the experimental factors, revealing how the optimal choice for one factor may depend on the level of another.

Interaction between Data Split and Classifier Performance

A significant interaction was observed between the training data split and the choice of classifier. A grouped box and whisker plot, with data split on the x-axis and separate box for each classifier, would effectively visualize this relationship (refer Figure 6). The more advanced ensemble classifiers, XGBoost, RF, SVM as well as KNN, showed a performance increase as the training set size grew from 50% to 80%. Their performance was highly dependent on having a sufficiently large training set to learn the complex decision boundaries. In contrast, the simpler models like NB and CART showed a less pronounced improvement, plateauing earlier. This indicates that while advanced classifiers have a higher performance ceiling, they require more data to reach it. Conversely, when training data is severely limited, the performance gap between simple and complex models narrows.

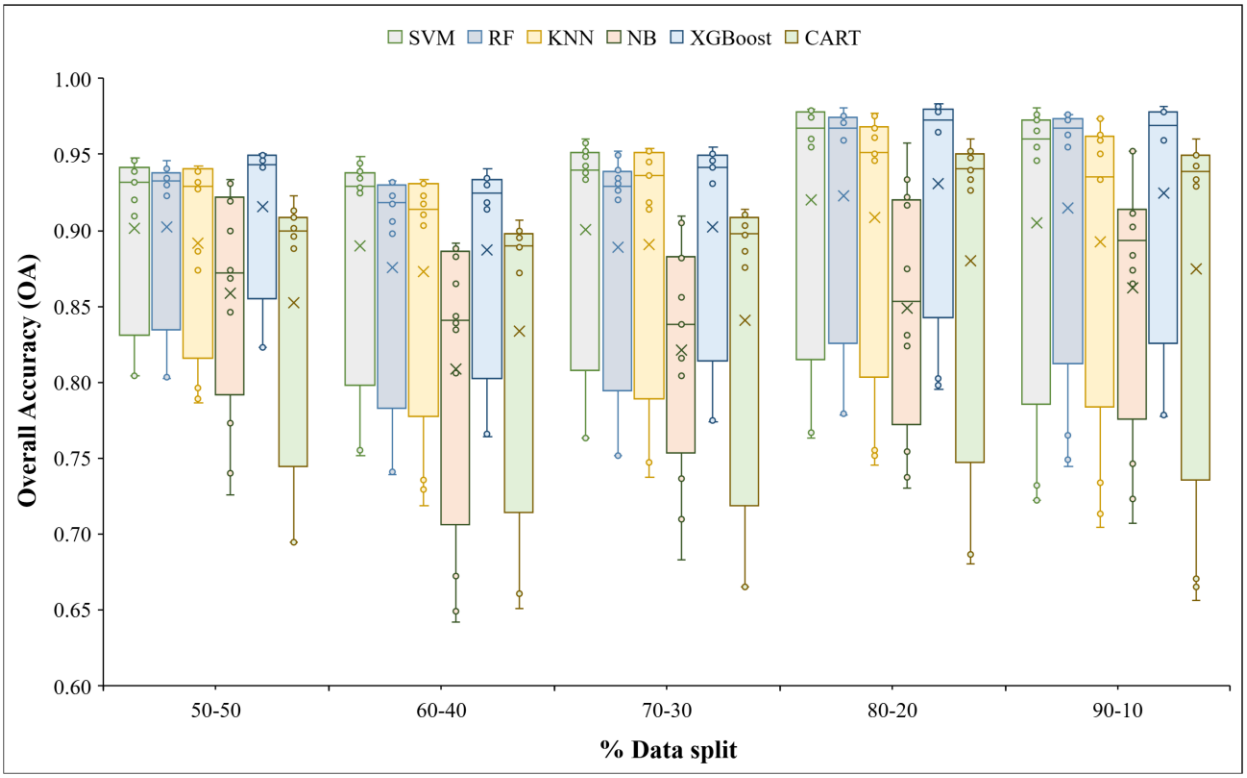

Figure 6. Box-and-whisker plot for the OA for ML classifier performance under different data split.

Feature Importance in the Optimal Model

An analysis of the feature importance for the best-performing model (DS4, XGBoost, BE, 80/20 split) provides critical insights into the biophysical drivers of classification accuracy. The analysis, as depicted in Figure 7a, reveals that a diverse set of features from all three data sources was selected as important, reinforcing the value of data fusion. The most influential features can be grouped into three key categories:

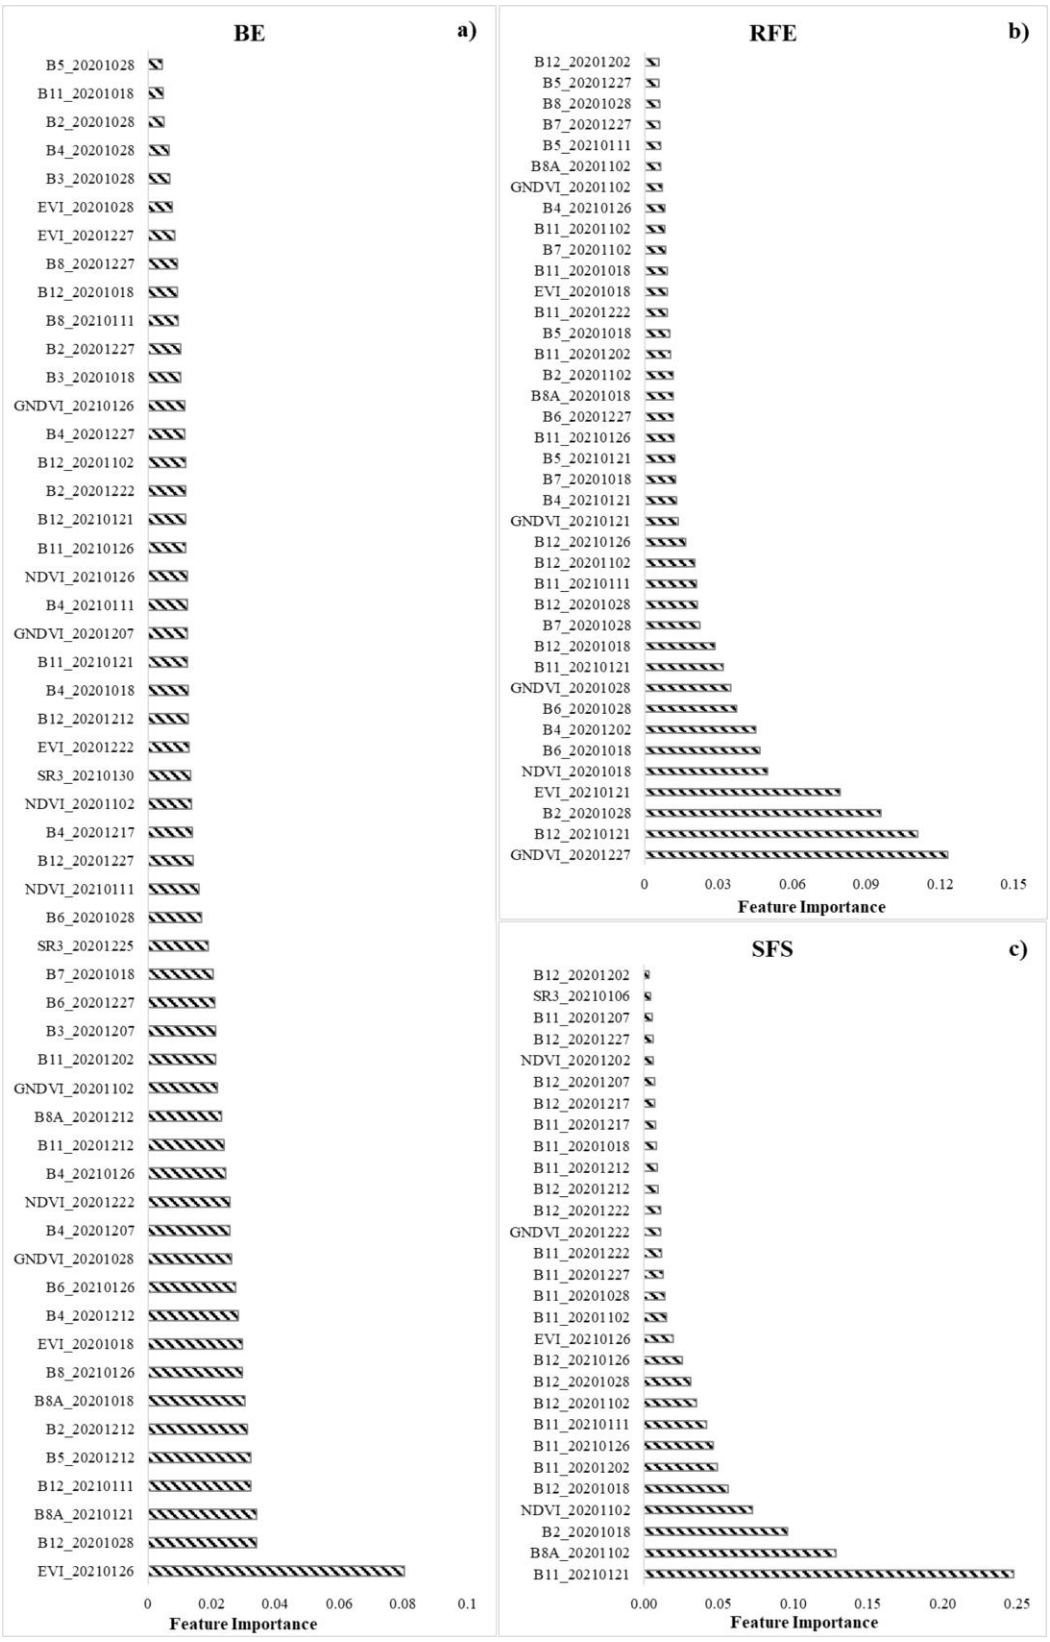

Figure 7. Feature importance for the maximum performance achieved for the dataset 4 by the 80/20 data split and XGBoost with FS technique (a) BE, (b) RFE, and (c) SFS.

1. **Sentinel-2 SWIR and Red-Edge Bands:** Bands B11 and B12 (Short-Wave Infrared) and bands B5 through B8A (Red-Edge) were consistently ranked as highly important. This is biophysically significant, as SWIR bands are sensitive to vegetation water content and canopy structure, while red-edge bands are highly correlated with chlorophyll content and are crucial for distinguishing subtle differences in vegetation health and phenological stage.
2. **Key Vegetation Indices:** Time-series of NDVI, EVI, and GNDVI were identified as critical predictors. These indices are designed to track photosynthetic activity and green biomass, and their temporal profiles provide a unique signature for different crop growth cycles.
3. **SAR-Derived Structural Information:** The simple ratio SR3 (VH+VV) from Sentinel-1 data emerged as an important feature. This index is related to the overall backscatter from the canopy and is sensitive to changes in biomass and vegetation structure, providing information that is complementary to the spectral data from S2

The RFE and SFS methods, when applied to the same dataset and classifier (Figure 7b and 7c), also highlighted the importance of S2 bands and VIs, though RFE tended to de-emphasize S1 bands, which may explain its slightly lower peak performance (OA: 0.9798) compared to the BE configuration (OA:0.9835) that retained a more diverse feature set.

3.4 Final Classified Map

The application of the optimal workflow (XGBoost classifier with BE feature selection on the DS4 dataset using an 80/20 data split) resulted in a high-quality, detailed crop classification map of the study area (Figure 8). The map exhibits high spatial coherence and accurately delineates the complex mosaic of agricultural fields. The area statistics derived from this map (Table 6) provide a quantitative inventory of the land cover distribution. Paddy was identified as the dominant crop, covering 49.66% of the area (2305.25 ha), followed by a significant portion of Fallow Land (35.10%) due to less sowing of non-cash crops. This output represents a tangible and actionable product for local water management authorities, providing the precise spatial data on cropping patterns required for improved irrigation performance assessment.

Table 6. Cropping area from DS4 using XGboost, BE technique with 80/20 data split.

Class	Area in Ha	% of Area
Bajra	132.54	2.86%
Buildup	114.92	2.48%
Fallow land	1627.80	35.08%
Groundnut	203.18	4.38%
Jowar	72.86	1.57%
Other	6.06	0.13%
Paddy	2309.84	49.79%
Barren land	154.72	3.33%
Waterbody	17.70	0.38%
Total	**4639.61**	**100.00%**

4. Discussion

This study's systematic, multi-factorial analysis of a crop classification workflow provides insights that extend beyond the identification of a single optimal model. By disentangling the effects of data fusion, classifier choice, feature selection, and training data partitioning, we can interpret the underlying reasons for the observed performance patterns and place these findings within the broader context of the remote sensing literature.

4.1 Interpretation of Key Findings

The Overwhelming Power of Synergistic Data: The finding that the fully fused dataset (DS4) consistently and significantly outperformed all others is the most critical outcome of this study. This result is not merely an artifact of 'more data is better'; it speaks to the fundamental complementarity of optical and SAR remote sensing for agricultural monitoring (Kabolizadeh et al., 2023). Sentinel-2's multispectral bands capture the biochemical properties of the crop's chlorophyll content, leaf structure, and water

content which manifest as distinct phenological curves over the growing season (Liu et al., 2024). However, these signals can be ambiguous between different crop types at certain growth stages. Sentinel-1's SAR data adds a crucial second dimension of information related to the physical structure of the canopy: plant height, density, biomass, and surface roughness (Vizzari et al., 2024). The combination of these two information streams allows the classifier to build a much more robust and unique spatio-temporal signature for each crop type, enabling it to resolve ambiguities that would be insurmountable with either data source alone. This synergy is the primary driver of the high accuracies achieved.

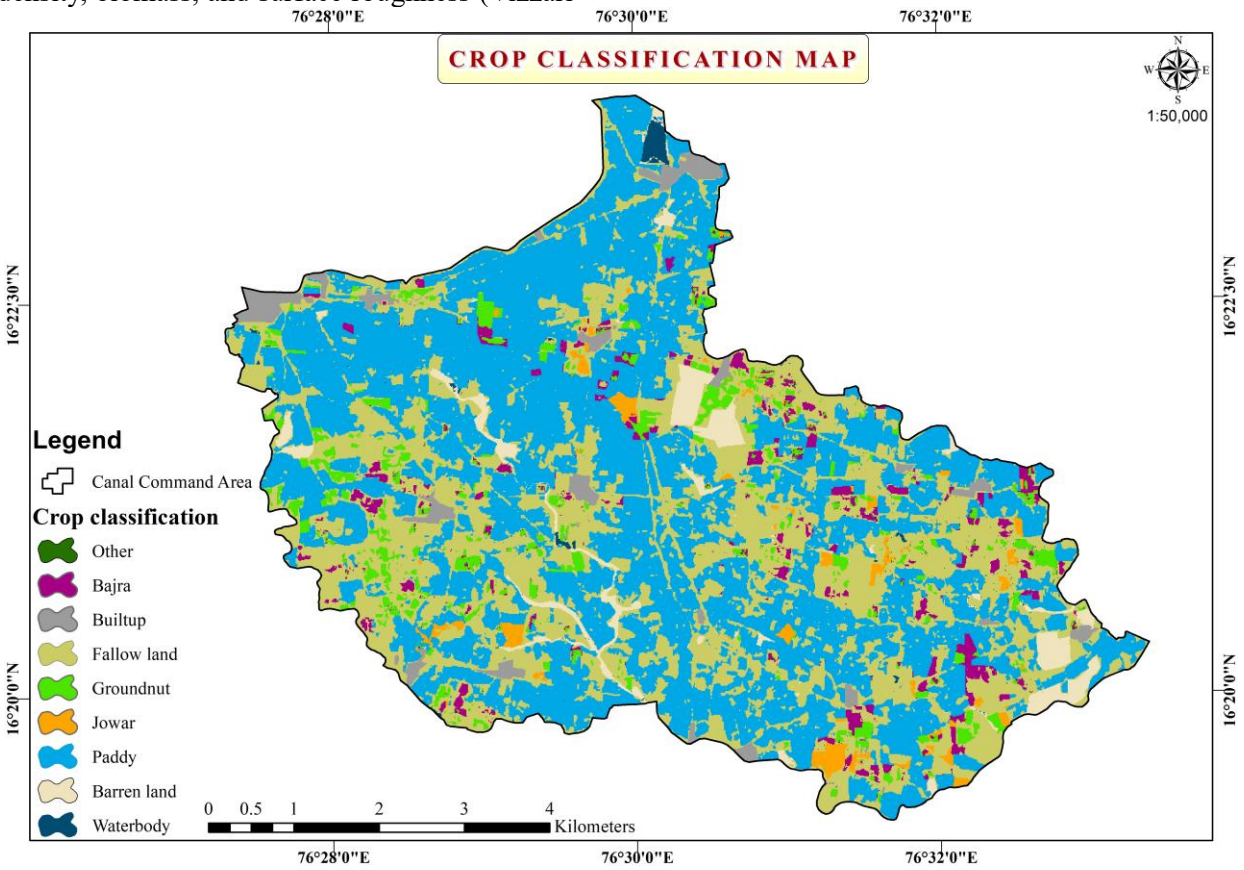

Figure 8. Final Crop Classification Map from DS4 using XGboost, BE technique with 80/20 data split.

Classifier-Complexity Mismatch and the Dominance of Gradient Boosting: The clear performance hierarchy among classifiers, with XGBoost at the top and Naïve Bayes at the bottom, highlights the concept of a 'classifier-complexity mismatch'. The input feature space, especially for DS4 with 222 variables, is high-dimensional, non-linear, and contains complex inter-dependencies. Simple models like Naïve Bayes, which assume feature independence, are fundamentally mismatched to this data structure and are thus incapable of learning the effective decision boundaries, resulting in poor performance (Banda et al., 2023). In contrast, XGBoost is exceptionally well-suited for this type of problem. Its foundation in gradient boosting allows it to iteratively build a strong model by correcting the errors of a sequence of weak learners (decision trees) (Mirzaei et al., 2023). This process is inherently adept at capturing complex, non-linear interactions between features. Furthermore, XGBoost's built-in regularization mechanisms (L1 and L2) are critical for preventing overfitting in a high-dimensional feature space, a key reason for its superior and more robust performance compared to even other powerful ensemble methods like Random Forest (Shao et al., 2024).

The 'Training Data Threshold' Hypothesis: The discovery of a performance plateau around an 80/20 data split has significant practical and

theoretical implications. We hypothesize that this 'training data threshold' is a function of landscape complexity. In a heterogeneous environment like the study command area, with small fields and high intra-class spectral variability (e.g., different planting dates or management practices for the same crop), a substantial number of training samples are required to adequately represent the full feature space of each class. Below this 80% threshold, the training set is likely not diverse enough, leading to a model that is under-fitted and generalizes poorly. Once this threshold is reached, the training set has sufficiently sampled the landscape's variability, and adding more samples provides only marginal new information, hence the performance plateau. This finding challenges the often-unquestioned adoption of standard splits like 70/30 and suggests that for complex agricultural mapping tasks, investing in a larger ground truth dataset for training is a critical prerequisite for achieving high accuracy. This has direct implications for the design and budgeting of field data collection campaigns.

4.2 Contextualizing within the State-of-the-Art

A potential critique of this study could be the exclusion of deep learning models. However, this was a deliberate methodological choice to establish a rigorous and, crucially, interpretable baseline for crop classification performance. While DL models like CNNs and LSTMs have shown great promise, their "black box" nature often makes it difficult to understand the precise drivers of their predictions (Yadav et al., 2024). Our study, by focusing on a range of traditional and ensemble ML models, provides a clear and detailed picture of the fundamental relationships between data, features, and model performance. The key findings such as the critical importance of SWIR and red-edge bands, the synergistic value of S1-S2 fusion, and the existence of a training data threshold are not specific to the ML models used but are fundamental principles of the data itself. These insights serve as an essential foundation that can and should inform the design of more complex DL architectures. For example, a future DL model could be made more efficient and effective by using the feature importance results from this study to pre-select the most informative input bands, or by ensuring the training dataset meets the 80% threshold identified here. Therefore, this work does not compete with the DL literature but rather complements it by providing the foundational benchmarks and data-driven principles upon which more advanced models can be built.

4.3 Quantifying the Reduction in Hydrological Uncertainty

A primary objective of this study was to move beyond classification metrics and quantify the reduction in input uncertainty for agro-hydrological modeling. To achieve this, we compared the potential volumetric water error of the optimized workflow (XGBoost + DS4) against a non-optimized baseline scenario typical of data-scarce regions (NB + DS1, OA: 64.2%). As shown in Table 7, the non-optimized model significantly overestimated the area of Paddy—the most water-intensive crop—predicting 3005.25 ha compared to the 2309.84 ha identified by the optimized model. This overestimation of 695.41 ha was largely due to the misclassification of Fallow Land as Paddy.

To translate this spatial error into hydrological terms, we calculated the Gross Irrigation Requirement (GIR) for Paddy during the 2020-21 Rabi season using the CROPWAT 8.0 model. Based on local meteorological data and a calculated canal system efficiency of 51% [Citation for 51% efficiency], the GIR was determined to be 2345.1 mm (23,451 m^3/ha).

$$\text{Volumetric Uncertainty} = \Delta\text{Area}_{\text{Paddy}} * \text{GIR}_{\text{Paddy}}$$

Applying this water duty to the area discrepancy, the non-optimized model yielded a volumetric overestimation of approximately 16.31 million Cubic Meters (MCM) of water (695.41 ha × 23,451 m^3/ha). In a water-scarce semi-arid command area, an error of this magnitude would lead to significant misallocation of resources, potentially depriving tail-end farmers of their equitable share. The optimized workflow effectively eliminates this uncertainty, ensuring that agro-hydrological models receive the precise spatial inputs required for accurate demand forecasting.

4.4 Practical Implications for Command Area Management

The ultimate goal of this research is to produce actionable information for real-world applications. The final, high-accuracy (OA > 98%) crop map is a direct input for enhancing Irrigation Performance Assessment (IPA) and command area management. Water resource managers can use this map to: • Improve Water Demand Forecasting: By knowing the precise area and location of water-intensive crops like paddy versus less-intensive crops like jowar or groundnut, authorities can create far more accurate models of aggregate water demand across the canal network. • Enable Equitable Water Distribution: The map allows for a data-driven approach to water allocation, ensuring that distributaries serving larger areas of thirsty crops receive adequate supply, promoting fairness and reducing conflict among water users. • Monitor Policy Compliance: If policies are in place to encourage crop diversification or limit the cultivation of certain crops in specific zones, this methodology provides a transparent and efficient tool for monitoring compliance over large areas. The demonstrated workflow, once operationalized, can provide this critical information on an annual basis, transforming command area management from a reactive to a proactive, data-driven process.

Table 7: Comparison of crop area estimation and deviation between the optimized (XGBoost + DS4) and non-optimized (Naïve Bayes + DS1) classification models.

Class	Optimized Model Area (ha)	Non-Optimized Model Area (ha)	Difference (ha)	% Deviation
Bajra	132.53	115.28	-17.25	-13.02%
Buildup	114.92	125.6	10.68	9.29%
Fallow land	1627.8	785.02	-842.78	-51.77%
Groundnut	203.18	218.5	15.32	7.54%
Jowar	72.86	88.39	15.53	21.31%
Other	6.06	4.85	-1.21	-19.97%
Paddy	2309.84	3005.25	695.41	30.11%
Barren land	154.72	286.51	131.79	85.18%
Waterbody	17.7	10.21	-7.49	-42.32%
Total	**4639.61**	**4639.61**		

(Note: The deviation shows that the non-optimized model effectively "stole" area from Fallow Land (-51%) and incorrectly assigned it to Paddy (+30%) and Barren Land (+85%). This is a very strong argument for why optimization is needed.)

4.5 Limitations and Future Work

While this study provides a robust and comprehensive analysis, it is important to acknowledge its limitations, which in turn point toward avenues for future research.

First, the scope of this analysis was limited to the 2020-2021 Rabi season within the NLBC command area. While the specific trained model weights are local to this spatiotemporal context, the methodological principles established here specifically the necessity of SAR-optical fusion, the superiority of gradient boosting (XGBoost) over standard classifiers, and the identification of the 80% training data threshold are expected to be highly transferable to other semi-arid agricultural regions with similar fragmented landholdings. Future research should focus on validating this optimized workflow across multi-year datasets to assess its robustness under varying climatic conditions (e.g., drought vs. flood years) and extending it to larger river basins. Second, while the pixel-based training approach provided sufficient statistical power (>20,000 spectral observations), the number of independent ground truth polygons (n=314) was constrained by operational logistics. A larger set of field polygons in future studies would further strengthen the validation of object-level accuracy. Third, although training data was derived from field polygons, the classification inference was performed at the pixel level. While highly

accurate, this approach can produce characteristic 'salt-and-pepper' noise within homogenous fields. Future research could explore Object-Based Image Analysis (OBIA) to mitigate this effect and create more coherent classification maps that better represent field boundaries (Vizzari et al., 2024). Finally, the principles established in this study provide a crucial foundation to inform the design of more advanced deep learning models, such as hybrid CNN-LSTM architectures (Liu et al., 2024).

5. Conclusion

The current research successfully addressed the critical challenge of input uncertainty in agro-hydrological modeling by validating an optimized Machine Learning workflow for crop classification in a semi-arid canal command area. While the primary output is a high-precision crop map, the broader contribution is a robust, reproducible methodology that enhances the reliability of sustainable water management decisions.

The multi-factorial analysis yielded unequivocal evidence that the synergistic fusion of Sentinel-1 (SAR) and Sentinel-2 (optical) data is the single most critical factor for minimizing classification error in these complex landscapes. Furthermore, the identification of a "training data threshold" at an 80/20 split provides a crucial operational insight. It offers water managers a clear cost-benefit guideline for planning field data campaigns, balancing the need for robust model training with the practical constraints of survey budgeting. Among the evaluated strategies, the Extreme Gradient Boosting (XGBoost) classifier, coupled with Backward Elimination feature selection, emerged as the superior architecture for handling the non-linear complexities of the data.

The optimal workflow achieved an exceptional Overall Accuracy of 98.4%. From an operational perspective, this accuracy effectively reduces the spatial area estimation error to less than 2%. By minimizing this error particularly for high-water-demand crops like Paddy, water managers can significantly reduce the volumetric uncertainty in seasonal water budgets, preventing the potential misallocation of resources. This high fidelity is not merely a statistical benchmark but a fundamental requirement for reducing uncertainty in crop water requirement calculations. By providing a validated, low-error input layer, this research significantly enhances the predictive capacity of agro-hydrological models, enabling more precise canal performance analysis and supporting data-driven policy for equitable water allocation in water-scarce regions.

Finally, it is important to note that these findings are derived from a semi-arid landscape characterized by specific crop phonologies. While the methodological framework (XGBoost optimization with 80/20 splitting) is designed to be robust, the transferability of these specific model weights to humid or highly heterogeneous Agro-climatic zones requires further validation.

Acknowledgements:
We are grateful to Krishna Bhagya Jala Nigam Limited (KBJNL) for their support in collecting field data for this analysis. We also acknowledge the use of Google Gemini (https://gemini.google.com/app) for improving the vocabulary and grammar of the article during the drafting process.

Author Contributions:
Mohansing Rajaput: Conceptualization, coding, formal analysis, writing of the original draft, and editing.
Abhilash Ramadasa: Writing, reviewing, and editing.
Basavanand M. Dodamani: Review, editing, and supervision.

Conflicts of interest:
The authors of this article declared no conflict of interest regarding the authorship or publication of this article.

Data availability statement:
The Sentinel-1 and Sentinel-2 satellite imagery that support the findings of this study are openly available from the Copernicus Open Access Hub at https://scihub.copernicus.eu/. The ground truth data collected for this research are not publicly available due to their nature but are

available from the corresponding author upon reasonable request.

References

Abhilash, R., Basappa, V., Chakragiri, S. V., & Patankar, D. B. (2022). Geospatial Approach for Integrated Command Area Management. Journal of Irrigation and Drainage Engineering, 148(4), 1–11. doi:10.1061/(asce)ir.1943-4774.0001659

Ahmed, A. A., Sayed, S., Abdoulhalik, A., Moutari, S., & Oyedele, L. (2024). Applications of machine learning to water resources management: A review of present status and future opportunities. Journal of Cleaner Production, 441, 140715. doi: 10.1016/j.jclepro.2024.140715

Amarasinghe, U. A., Sikka, A., Mandave, V., Panda, R. K., Gorantiware, S., Chandrasekharana, K., & Ambast, S. K. (2021). A re-look at canal irrigation system performance: a pilot study of the Sina irrigation system in Maharashtra, India. Water Policy, 23, 114–129. doi:10.2166/wp.2020.291

Banda, L., Rai, A., Kansal, A., & Vashisth, A. K. (2023). Suitable Crop Prediction based on affecting parameters using Naïve Bayes Classification Machine Learning Technique. 2023 International Conference on Disruptive Technologies, ICDT 2023, 43–46. doi:10.1109/ICDT57929.2023.10150814

Biau, G., & Scornet, E. (2016). A random forest guided tour. Test, 25(2), 197–227. doi:10.1007/s11749-016-0481-7

Chang, Z., Li, H., Chen, D., Liu, Y., Zou, C., Chen, J., Han, W., Liu, S., & Zhang, N. (2023). Crop Type Identification Using High-Resolution Remote Sensing Images Based on an Improved DeepLabV3+ Network. Remote Sensing, 15(21), 5088. doi:10.3390/rs15215088

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 13-17-August-2016, 785–794. doi:10.1145/2939672.2939785

Congalton, R. G., & Green, K. (2019). Assessing the Accuracy of Remotely Sensed Data. CRC Press. doi:10.1201/9780429052729

Eisfelder, C., Boemke, B., Gessner, U., Sogno, P., Alemu, G., Hailu, R., Mesmer, C., & Huth, J. (2024). Cropland and Crop Type Classification with Sentinel-1 and Sentinel-2 Time Series Using Google Earth Engine for Agricultural Monitoring in Ethiopia. Remote Sensing, 16(5), 866. doi:10.3390/rs16050866

ESA. (2024). Copernicus open access hub. https://scihub.copernicus.eu/

Gao, S., Tang, B. H., Huang, L., & Chen, G. (2023). Identification of tea plantations in typical plateau areas with the combination of Sentinel-1/2 optical and radar remote sensing data based on feature selection algorithm. International Journal of Remote Sensing, 7033–7053. doi:10.1080/01431161.2023.2198655

Gitelson, A. A., Kaufman, Y. J., & Merzlyak, M. N. (1996). Use of a green channel in remote sensing of global vegetation from EOS-MODIS. Remote Sensing of Environment, 58(3), 289–298. doi:10.1016/S0034-4257(96)00072-7

GSARS. (2017). Handbook on remote sensing for agricultural statistics. In Handbook on Remote Sensing for Agricultural Statistics. GSARS Handbook. https://openknowledge.fao.org/server/api/core/bitstreams/dcadd248-cccb-43ad-9e3c-7aa0cf2bb577/content

He, S., Peng, P., Chen, Y., & Wang, X. (2022). Multi-Crop Classification Using Feature Selection-Coupled Machine Learning Classifiers Based on Spectral, Textural and Environmental Features. Remote Sensing, 14. doi:10.3390/rs14133153

Hoppe, H., Dietrich, P., Marzahn, P., Weiß, T., Nitzsche, C., Freiherr von Lukas, U., Wengerek, T., & Borg, E. (2024). Transferability of Machine Learning Models for Crop Classification in Remote Sensing Imagery Using a New Test Methodology: A Study on Phenological, Temporal, and Spatial Influences. Remote Sensing, 16(9), 1493. doi:10.3390/rs16091493

Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. Remote Sensing of Environment,

83(1–2), 195–213. doi:10.1016/S0034-4257(02)00096-2

Kabolizadeh, M., Rangzan, K., & Habashi, K. (2023). Improving classification accuracy for separation of area under crops based on feature selection from multi-temporal images and machine learning algorithms. Advances in Space Research, 72(11), 4809–4824. doi: 10.1016/j.asr.2023.09.044

Kumar, M. D., Sahasranaman, M., Verma, M. S., Kumar, S., & Narayanamoorthy, A. (2021). Getting the irrigation statistics right. International Journal of Water Resources Development. doi:10.1080/07900627.2021.1921711

Liu, N., Zhao, Q., Williams, R., & Barrett, B. (2024). Enhanced crop classification through integrated optical and SAR data: a deep learning approach for multi-source image fusion. International Journal of Remote Sensing, 45(19–20), 7605–7633. doi:10.1080/01431161.2023.2232552

Mandal, D., Kumar, V., Ratha, D., Dey, S., Bhattacharya, A., Lopez-Sanchez, J. M., McNairn, H., & Rao, Y. S. (2020). Dual polarimetric radar vegetation index for crop growth monitoring using sentinel-1 SAR data. Remote Sensing of Environment, 247. doi: 10.1016/j.rse.2020.111954

Manning, C. D., Raghavan, P., & Schütze, H. (2008). Introduction to Information Retrieval. Cambridge University Press. doi:10.1017/CBO9780511809071

Maponya, M. G., van Niekerk, A., & Mashimbye, Z. E. (2020). Pre-harvest classification of crop types using a Sentinel-2 time-series and machine learning. Computers and Electronics in Agriculture, 169. doi: 10.1016/j.compag.2019.105164

Maxwell, A. E., Warner, T. A., & Fang, F. (2018). Implementation of machine-learning classification in remote sensing: an applied review. International Journal of Remote Sensing, 39(9), 2784–2817. doi:10.1080/01431161.2018.1433343

Mirzaei, A., Bagheri, H., & Khosravi, I. (2023). Enhancing Crop Classification Accuracy through Synthetic SAR-Optical Data Generation Using Deep Learning. ISPRS International Journal of Geo-Information, 12(11). doi:10.3390/ijgi12110450

Montgomery, D. C., Peck, E. A., & Vining, G. G. (2012). Introduction to Linear Regression Analysis (5th ed.). John Wiley & Sons. https://www.kwcsangli.in/uploads/3--Introduction_to_Linear_Regression_Analysis__5th_ed._Douglas_C._Montgomery__Elizabeth_A._Peck__and_G._.pdf

Nigar, A., Li, Y., Jat Baloch, M. Y., Alrefaei, A. F., & Almutairi, M. H. (2024). Comparison of machine and deep learning algorithms using Google Earth Engine and Python for land classifications. Frontiers in Environmental Science, 12. doi:10.3389/fenvs.2024.1378443

Ok, A. O., Akar, O., & Gungor, O. (2012). Evaluation of random forest method for agricultural crop classification. European Journal of Remote Sensing, 45(1), 421–432. doi:10.5721/EuJRS20124535

Pal, M., & Mather, P. M. (2003). An assessment of the effectiveness of decision tree methods for land cover classification. Remote Sensing of Environment, 86(4), 554–565. doi:10.1016/S0034-4257(03)00132-9

Qader, S. H., Dash, J., Alegana, V. A., Khwarahm, N. R., Tatem, A. J., & Atkinson, P. M. (2021). The Role of Earth Observation in Achieving Sustainable Agricultural Production in Arid and Semi-Arid Regions of the World. Remote Sensing, 13(17), 3382. doi:10.3390/rs13173382

Raja, S. P., Sawicka, B., Stamenkovic, Z., & Mariammal, G. (2022). Crop Prediction Based on Characteristics of the Agricultural Environment Using Various Feature Selection Techniques and Classifiers. IEEE Access, 10, 23625–23641. doi:10.1109/ACCESS.2022.3154350

Rajaput, M., Abhilash, R., & Dodamani, B. M. (2025). A systematic review of performance assessment in canal irrigation systems: Integrating socio-technical, remote sensing, and AI-driven approaches for a climate-resilient future. Water and Soil Management and Modeling, 5(4), 254–276. doi:10.22098/mmws.2025.18343.1683

Ramezan, C. A. (2022). Transferability of Recursive Feature Elimination (RFE)-Derived Feature Sets for Support Vector Machine Land Cover Classification. Remote Sensing, 14. doi:10.3390/rs14246218

Rouse, J. W., Haas, R. H., Schell, J. A., Deering, D. W., & Harlan, J. C. (1974). Monitoring the Vernal Advancement of Retrogradation (Greenwave Effect) of Natural Vegetation. https://ntrs.nasa.gov/api/citations/19750020419/downloads/19750020419.pdf

Rusňák, T., Kasanický, T., Malík, P., Mojžiš, J., Zelenka, J., Sviček, M., Abrahám, D., & Halabuk, A. (2023). Crop Mapping without Labels: Investigating Temporal and Spatial Transferability of Crop Classification Models Using a 5-Year Sentinel-2 Series and Machine Learning. Remote Sensing, 15(13), 3414. doi:10.3390/rs15133414

Saadatfar, H., Khosravi, S., Joloudari, J. H., Mosavi, A., & Shamshirband, S. (2020). A new k-nearest neighbors' classifier for big data based on efficient data pruning. Mathematics, 8(2). doi:10.3390/math8020286

Schölkopf, B., & Smola, A. J. (2001). Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. The MIT Press. doi:10.7551/mitpress/4175.001.0001

Shao, Z., Ahmad, M. N., & Javed, A. (2024). Comparison of Random Forest and XGBoost Classifiers Using Integrated Optical and SAR Features for Mapping Urban Impervious Surface. Remote Sensing, 16(4). doi:10.3390/rs16040665

Solomatine, D. P., & Ostfeld, A. (2008). Data-driven modelling: some past experiences and new approaches. Journal of Hydroinformatics, 10(1), 3–22. doi:10.2166/hydro.2008.015

Song, Q., Xiang, M., Hovis, C., Zhou, Q., Lu, M., Tang, H., & Wu, W. (2019). Object-based feature selection for crop classification using multi-temporal high-resolution imagery. International Journal of Remote Sensing, 40(5–6), 2053–2068. doi:10.1080/01431161.2018.1475779

Uçar, M. K., Nour, M., Sindi, H., & Polat, K. (2020). The Effect of Training and Testing Process on Machine Learning in Biomedical Datasets. Mathematical Problems in Engineering, 2020. doi:10.1155/2020/2836236

Vizzari, M., Lesti, G., & Acharki, S. (2024). Crop classification in Google Earth Engine: leveraging Sentinel-1, Sentinel-2, European CAP data, and object-based machine-learning approaches. Geo-Spatial Information Science. doi:10.1080/10095020.2024.2341748

Yadav, S., Sharma, S., & Chaudhary, P. (2024). Review on Advancement in Machine Learning and Deep Learning Techniques for Crop Classification. In Advancements in Communication and Systems (pp. 593–610). Soft Computing Research Society. doi:10.56155/978-81-955020-7-3-53

Yao, J., Wu, J., Xiao, C., Zhang, Z., & Li, J. (2022). The Classification Method Study of Crops Remote Sensing with Deep Learning, Machine Learning, and Google Earth Engine. Remote Sensing, 14(12), 2758. doi:10.3390/rs14122758

Zhang, N., Chen, M., Yang, F., Yang, C., Yang, P., Gao, Y., Shang, Y., & Peng, D. (2022). Forest Height Mapping Using Feature Selection and Machine Learning by Integrating Multi-Source Satellite Data in Baoding City, North China. Remote Sensing, 14(18). doi:10.3390/rs14184434