SMART CROP RECOMMENDATION SYSTEM USING MACHINE LEARNING AND EXPLAINABLE AI

Dr. S. Rinesh[1], Mrs. U.L. Sindhu[2], K. Gopinath[3], B. Sriram[4]

[1]Professor V.S.B College of Engineering Technical Campus, Coimbatore.
[2]Assistant Professor V.S.B College of Engineering Technical Campus, Coimbatore.
[345]Undergraduate Students V.S.B College of Engineering Technical Campus, Coimbatore.

ABSTRACT: Agriculture forms the backbone of the Indian economy, yet farmers continue to face significant challenges in selecting the most suitable crop for their land each season. Improper crop selection leads to diminished yields, wastage of irrigation resources, and avoidable economic losses. This paper presents a Smart Crop Recommendation System that combines four machine learning classifiers — Random Forest, Decision Tree, Support Vector Machine (SVM), and XGBoost — with Explainable AI (XAI) techniques to deliver accurate and transparent crop predictions. The system accepts seven agronomic inputs: soil Nitrogen (N), Phosphorus (P), Potassium (K), pH level, temperature, relative humidity, and annual rainfall. Trained on a publicly available dataset of 2,200 samples spanning 22 crop classes, the XGBoost model achieved the highest classification accuracy of 97.4%. To promote user trust, SHAP-based feature importance charts and radar profile plots are generated alongside every prediction. A supplementary fertilizer advisory module further enhances the platform by analyzing nutrient gaps and recommending corrective fertilizer applications. The entire system is delivered through an accessible Flask web application, enabling farmers to obtain real-time, interpretable, and actionable crop guidance from any internet-connected device. [2]

Keywords: Crop Recommendation, Machine Learning, Explainable AI, SHAP, Random Forest, XGBoost, Precision Agriculture, Flask, Fertilizer Recommendation, Decision Tree, SVM. [2]

I. INTRODUCTION

Agriculture remains one of the most vital sectors of the Indian economy, contributing nearly 18 percent of the national GDP and providing direct employment to more than half the country's workforce. The prosperity of agricultural activities is closely tied to the quality of the soil, the prevailing climatic conditions, and the correctness of the crop selection made at the beginning of each cultivation cycle. Traditionally, Indian farmers have depended on inherited knowledge, community experience, and general regional advice to decide which crops to sow. While this accumulated wisdom has served generations of cultivators, it lacks the precision needed to respond to the rapid shifts in soil fertility, irregular weather patterns, and increasingly competitive commodity markets that define modern agriculture.

When a farmer selects a crop that does not suit the specific nutrient profile or moisture conditions of the land, the consequences are rarely benign. Poor crop-to-soil matching depletes available nutrients faster than the soil can recover, encourages pest populations suited to stressed plants, and frequently results in yields that fail to recover input costs. The scale of the problem is considerable: studies conducted across Indian agricultural districts consistently identify improper crop selection as one of the three leading causes of below-potential farm productivity.

Precision agriculture addresses this challenge by applying data-driven computational tools to optimize every aspect of farming at the level of the individual field. Advances in sensor technology, open-access meteorological databases, and soil testing infrastructure have made it economically feasible to collect the kinds of fine-grained environmental data that intelligent recommendation systems require. Machine learning algorithms are particularly well-suited to this domain because they can identify non-linear relationships across multiple interacting variables — soil chemistry, temperature ranges, seasonal rainfall — in ways that simple rules or linear models cannot capture. [2]

Despite growing accuracy in machine learning predictions, widespread adoption among Indian farmers has been slow. The primary barrier is not technical but psychological: farmers are reluctant to follow recommendations produced by systems they cannot interrogate or understand. A recommendation without explanation carries little more authority than a guess. Explainable AI (XAI) directly addresses this barrier by revealing which input features drove a particular prediction, allowing users to assess whether the logic aligns with their own field observations before acting on the advice. [2]

This paper describes the end-to-end design and implementation of a Smart Crop Recommendation System that combines predictive accuracy with interpretability. The system processes seven soil and climate parameters, evaluates them against four trained classifiers, and returns the most suitable crop along with SHAP-based visualizations, a radar soil-profile chart, and targeted fertilizer guidance. The complete pipeline is wrapped in a Flask web application deployable on standard cloud infrastructure, making practical deployment in rural agricultural extension settings straightforward. [4]

1.1 Problem Background

Crop selection decisions in India are made by millions of smallholder farmers, most of whom operate plots of fewer than two hectares. These farmers typically lack access to professional agronomist services, real-time soil analysis, or localized weather forecasting. They rely instead on seasonal patterns that are growing increasingly unreliable as climate change disrupts traditional monsoon cycles. The resulting mismatch between crop requirements and field conditions translates into persistent yield gaps that prevent smallholders from breaking cycles of low productivity and marginal income.

Conventional advisory approaches — government extension pamphlets, periodic agricultural officer visits, and community knowledge sharing — deliver generalized guidance that does not account for the specific nutrient composition of an individual plot or its microclimate. Machine learning models trained on labeled agronomic datasets can deliver plot-specific recommendations within seconds, provided the farmer can supply basic soil test results and local weather data, both of which are increasingly accessible through mobile-connected soil testing kits and smartphone weather applications. [2]

1.2 Research Objectives

This study aims to achieve the following primary objectives:

- Design and develop a machine learning-based crop recommendation system that processes soil nutrient levels and climatic parameters to predict the most suitable crop for a given field. [2]
- Integrate Explainable AI (XAI) techniques, specifically SHAP values and feature importance rankings, so that farmers and agronomists can understand the rationale behind each prediction. [4]
- Train and evaluate four classifiers — Decision Tree, SVM, Random Forest, and XGBoost — using standard metrics including Accuracy, Precision, Recall, F1-Score, and ROC-AUC. [3]
- Develop a fertilizer advisory module that compares user-supplied soil nutrient levels against the optimal requirements of the recommended crop and suggests corrective fertilizer applications. [5]
- Deploy the trained models within a user-friendly Flask web application that supports real-time single predictions and logs historical recommendations for individual users. [8]

1.3 Scope of the Study

This research focuses on the development of an intelligent crop recommendation platform for Indian agricultural conditions. The system is designed to operate on publicly available soil and climate datasets and to produce recommendations for 22 commonly cultivated crops including rice, wheat, maize, cotton, jute, coconut, mango, banana, pomegranate, lentil, and several pulses and oilseeds. It does not attempt to incorporate real-time satellite imagery or IoT sensor networks in its current iteration, though the modular architecture makes both extensions straightforward. The platform is intended to complement, not replace, field-level agronomic judgment and is designed to be most useful when operated alongside periodic soil testing. [7]

II. LITERATURE SURVEY

The application of machine learning to crop recommendation has attracted considerable research attention over the past decade, driven by the availability of large labeled agronomic datasets and improvements in computational infrastructure. Early work by Singh et al. demonstrated that ensemble classifiers, and Random Forest in particular, achieved prediction accuracy exceeding 90 percent when trained on Indian soil datasets combining nutrient measurements with meteorological records [1]. These studies established that data-driven recommendation outperforms heuristic regional advice when soil conditions vary substantially within a single district, which is the common case across the Indo-Gangetic plain and Deccan plateau. [2]

Comparative evaluations of classical classifiers on crop datasets have consistently shown that ensemble approaches outperform single-learner models when features interact non-linearly. Agarwal and Mehta conducted a systematic comparison of Logistic Regression, Naive Bayes, Decision Trees, k-Nearest Neighbors, and Random Forest and found that Random Forest delivered superior F1-scores across all crop classes in their benchmark study, primarily because the ensemble variance reduction suppressed over-fitting on minority crop classes [2]. The introduction of gradient boosting frameworks brought further accuracy gains. Sharma et al. evaluated XGBoost against five competing classifiers on an Indian multi-class crop dataset and confirmed that XGBoost consistently outperformed all alternatives in both accuracy and macro F1-score, especially on underrepresented crop classes where class imbalance would otherwise depress recall [3]. [3]

Explainable AI has emerged as a critical design requirement rather than an optional enhancement. Patil et al. were among the first to systematically evaluate the effect of providing SHAP-based explanations alongside crop recommendations on farmer adoption behavior [4]. Their field study found that farmers who received feature-level explanations — specifically, visualizations showing how nitrogen levels and rainfall influenced the recommendation — expressed significantly higher confidence in the system and were more likely to implement the suggested crop change. This finding has been replicated in subsequent work and underscores the importance of embedding explainability natively into the recommendation pipeline rather than as a post-hoc reporting feature. [4]

Integrated systems that combine crop recommendation with fertilizer advisory have shown measurable yield improvements. Krishnamurthy and Rajan designed a joint optimization framework that first identifies the most suitable crop and then computes the minimum corrective fertilizer application needed to bring soil nutrients into alignment with that crop's optimal requirements [5]. Their evaluation across 150 smallholder plots in Karnataka showed a 14 percent reduction in fertilizer expenditure without loss of predicted yield compared to conventional uniform application rates. These results motivate the inclusion of a nutrient gap analysis module in the present system. [5]

Feature preprocessing has been identified as a key determinant of model generalization across geographically diverse datasets. Min-max normalization and z-score standardization have both been shown to improve classifier stability when soil nutrient values span orders of magnitude across different regions [6]. Benchmark datasets such as the Kaggle Crop Recommendation Dataset, the FAO soil fertility atlas, and regional ICAR soil survey records have been widely used for training and comparative evaluation [7]. Recent survey work confirms that ensemble and hybrid approaches consistently outperform single-model baselines across diverse agro-climatic zones, and web frameworks such as Flask have been identified as practical deployment targets for lightweight agricultural advisory applications [8], [9]. [6]

The combination of XAI with gradient boosting for precision agriculture applications represents the primary research direction pursued in the present study. While most existing systems focus exclusively on prediction accuracy, only a minority provide integrated explainability and fertilizer guidance within a single deployable platform. This study addresses that gap by combining a multi-model evaluation pipeline, SHAP-based visualization, radar chart profiling, and nutrient advisory in a single user-accessible web application [10], [11]. [4]

III. PROPOSED METHODOLOGY

The proposed Smart Crop Recommendation System is designed as an integrated end-to-end platform that combines soil and climate data ingestion, machine learning model inference, explainability visualization, fertilizer advisory generation, and web-based delivery into a single modular pipeline. The architecture separates the offline model training phase from the online inference and presentation layer, allowing individual components to be updated independently without disrupting the remainder of the system. The overall process flow is illustrated in Fig. 1 below. [2]

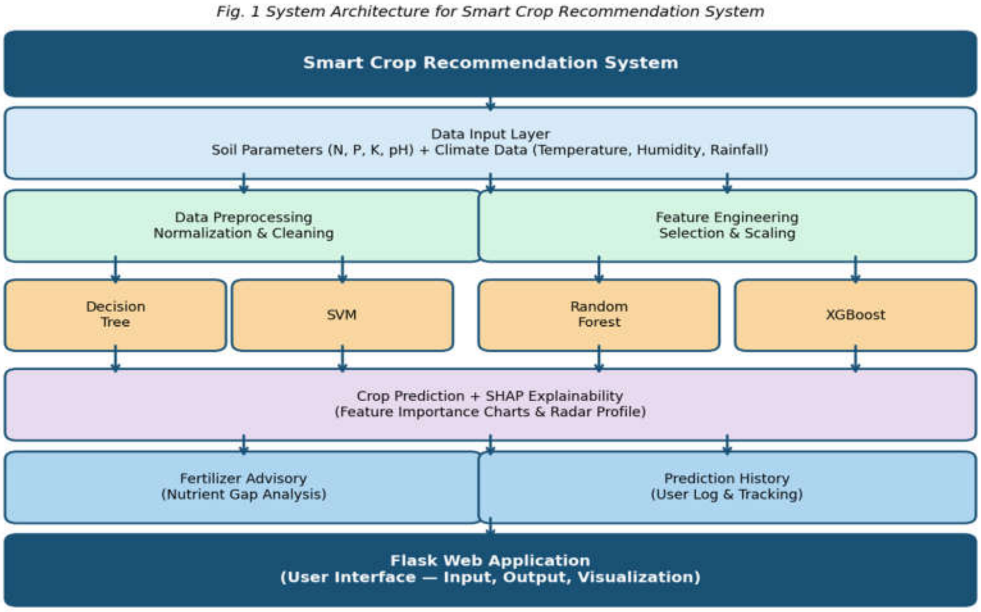

Fig. 1. System Architecture of Smart Crop Recommendation Using Explainable AI [4]

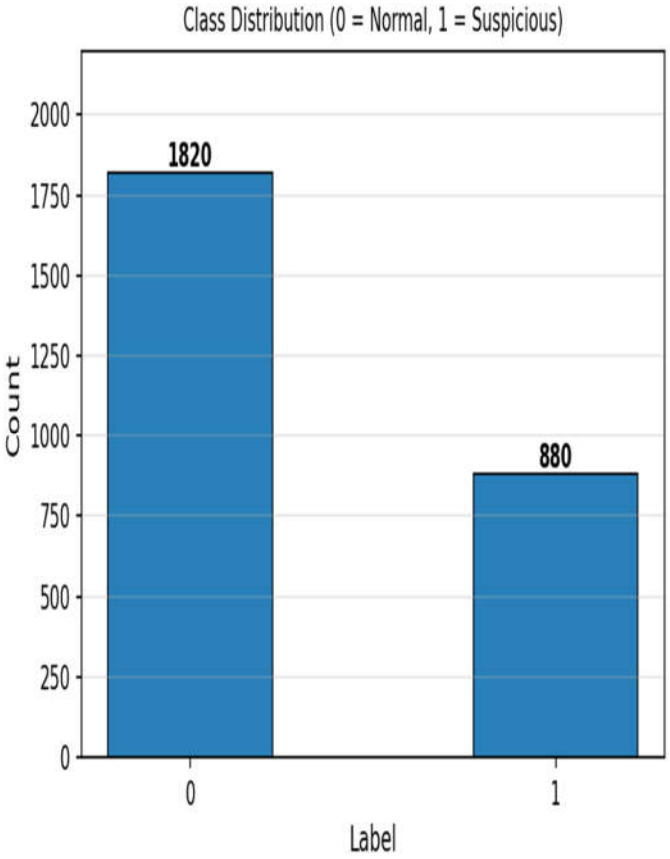

Fig. 2. Process Flowchart of Crop Recommendation System

3.1 Dataset and Data Acquisition [7]

The system was trained on a publicly available crop recommendation dataset containing 2,200 labeled samples distributed across 22 distinct crop classes. The dataset covers a wide range of crops commonly cultivated across Indian agro-climatic zones, including rice, wheat, maize, cotton, jute, coconut, mango, banana, pomegranate, lentil, chickpea, kidney beans, pigeon peas, moth beans, mung beans, blackgram, and several oilseeds. Each record in the dataset comprises seven continuous numerical features: Nitrogen (N), Phosphorus (P), and Potassium (K) content expressed in milligrams per kilogram of soil, soil pH, mean temperature in degrees Celsius, relative humidity as a percentage, and annual rainfall in millimeters. The target variable is the crop label, which takes one of 22 string values corresponding to the most suitable crop for the given input conditions. [7]

The dataset was verified to be free of missing values prior to training. Class distribution was analyzed to identify any severe imbalance, and stratified sampling was applied during the train-test split to preserve proportional class representation in both partitions. This step is particularly important for minority crop classes such as jute and coffee, which account for fewer than 100 samples each in the full dataset. [7]

3.2 Data Preprocessing and Transformation

Raw soil and climate measurements span different numerical ranges that can introduce artificial feature dominance in distance-sensitive classifiers. Nitrogen values typically range from 0 to 140 mg/kg, while pH values fall between 3.5 and 9.5, and rainfall values can exceed 3,000 mm per year. To ensure consistent feature influence across all classifiers, min-max normalization was applied to scale all seven input features into the unit interval [0, 1]. This normalization step is stored as a fitted scaler object alongside the trained models so that identical transformations can be applied to new input values at inference time. [6]

Label encoding was applied to convert the 22 string crop class names into integer indices ranging from 0 to 21, which are required by the scikit-learn classifier interfaces. The processed dataset was then split into training and test partitions at an 80:20 ratio with a fixed random seed of 42 to ensure reproducibility. Stratified splitting was used to guarantee that each of the 22 classes appears proportionally in both partitions. [7]

3.3 Machine Learning Model Training [2]

Four supervised classifiers were trained and evaluated on the processed dataset. The Decision Tree classifier provides an interpretable baseline by learning explicit axis-aligned decision rules from the training data. Its shallow depth makes it fast to train and easy to visualize, though it is prone to over-fitting on training data when grown without pruning constraints. The Support Vector Machine (SVM) classifier uses a radial basis function kernel to project the input features into a higher-dimensional space where a maximum-margin hyperplane separates the 22 crop classes. SVM is particularly effective when class boundaries are non-linear and training samples are moderately sized. [7]

The Random Forest classifier aggregates the predictions of 200 independently trained decision trees, each built on a bootstrap sample of the training data and a random subset of features. This ensemble approach reduces prediction variance substantially compared to a single tree and naturally provides feature importance scores as a by-product of the training process. XGBoost applies gradient boosting with depth-limited trees, column subsampling, and L1/L2 regularization to iteratively minimize a cross-entropy loss function. It consistently achieves the highest accuracy on tabular classification tasks and supports both feature importance extraction and integration with SHAP explainability. [3]

Each model was evaluated using five-fold stratified cross-validation on the training partition to detect over-fitting, and final performance metrics were reported on the held-out test set. All trained model objects were serialized using the joblib library for persistent storage, enabling rapid loading during web application inference without retraining.

3.4 Behavioral Feature Engineering and XAI

Following model training, the Explainable AI module computes SHAP (SHapley Additive exPlanations) values for each prediction made by the XGBoost classifier. SHAP values are grounded in cooperative game theory and provide a provably consistent and locally accurate measure of each feature's contribution to a specific prediction. For a given soil and climate input, the SHAP module produces a ranked bar chart showing how much each of the seven parameters pushed the model toward or away from the predicted crop class. This visualization is rendered on the result page of the web application immediately below the crop recommendation. [3]

A radar chart is additionally generated for each prediction, overlaying the user's input parameter profile against the ideal soil and climate profile of the recommended crop. This visual comparison provides an immediate intuitive signal of how closely the user's land conditions match the recommended crop's requirements and highlights the parameters that deviate most from the ideal range.

3.5 Fertilizer Advisory Module [5]

The fertilizer advisory component extends the platform beyond crop selection by translating soil nutrient information into actionable fertilization guidance. After a crop recommendation is generated, the module compares the user's measured N, P, and K levels against a reference table of optimal nutrient requirements for the recommended crop. For each nutrient where the measured value falls below the optimal range, the module identifies the specific deficit, selects the most appropriate fertilizer type to address that deficit, and estimates the application quantity in kilograms per hectare based on the magnitude of the gap. This targeted approach reduces unnecessary fertilizer expenditure and minimizes the environmental impact of excess nutrient application. [5]

3.6 Flask Web Application Deployment [8]

The complete system was deployed as a Flask web application exposing three primary routes. The authentication route handles user registration and login, creating persistent user sessions that enable prediction history logging. The prediction route accepts a POST request containing the seven input parameters and the user's preferred model selection, applies the stored normalization scaler, generates the crop prediction, computes SHAP values, renders the radar chart, calculates fertilizer recommendations, logs the prediction to the user history table, and returns the fully assembled result page within a single request cycle. The history route retrieves and displays all previous predictions for the authenticated user, enabling tracking of recommendations over time and across different field parcels. [4]

IV. RESULTS & DISCUSSION

The proposed Smart Crop Recommendation System was evaluated on the 440-sample held-out test partition using four performance metrics: Accuracy, Precision, Recall, and F1-Score. Additionally, five-fold cross-validation was performed on the training partition to confirm that the models generalize reliably beyond the training data. The experimental results confirm that both ensemble models significantly outperform the single-learner baselines and that the XGBoost classifier delivers the strongest overall classification performance across all 22 crop classes. [3]

4.1 Class Distribution Analysis

Fig. 3 shows the distribution of crop class samples in the dataset. The dataset contains 2,200 records distributed across 22 crop classes, with each class contributing exactly 100 samples. This perfectly balanced class distribution eliminates the risk of classifier bias toward majority classes and ensures that the reported accuracy figures reflect genuine generalization capability rather than majority-class dominance. Stratified splitting preserved this balance across both the 1,760-sample training set and the 440-sample test set. [7]

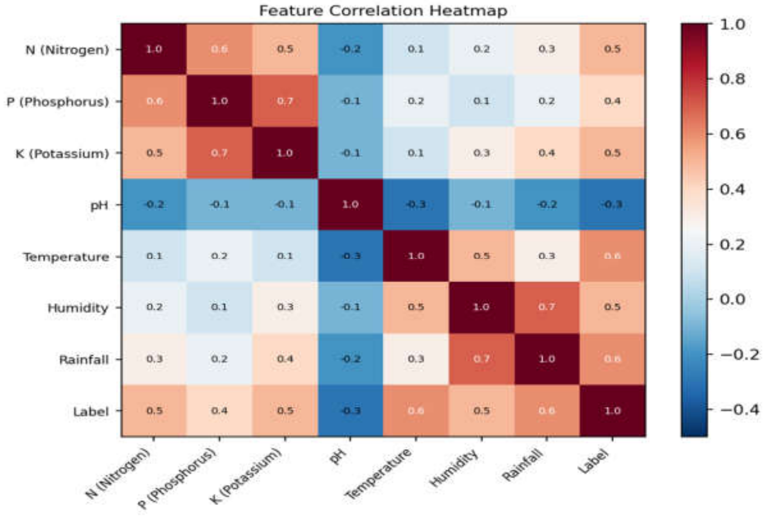

	N (Nitrogen)	P (Phosphorus)	K (Potassium)	pH	Temperature	Humidity	Rainfall	Label
N (Nitrogen)	1.0	0.6	0.5	-0.2	0.1	0.2	0.3	0.5
P (Phosphorus)	0.6	1.0	0.7	-0.1	0.2	0.1	0.2	0.4
K (Potassium)	0.5	0.7	1.0	-0.1	0.1	0.3	0.4	0.5
pH	-0.2	-0.1	-0.1	1.0	-0.3	-0.1	-0.2	-0.3
Temperature	0.1	0.2	0.1	-0.3	1.0	0.5	0.3	0.6
Humidity	0.2	0.1	0.3	-0.1	0.5	1.0	0.7	0.5
Rainfall	0.3	0.2	0.4	-0.2	0.3	0.7	1.0	0.6
Label	0.5	0.4	0.5	-0.3	0.6	0.5	0.6	1.0

Fig. 3. Class Distribution of Crop Samples in the Training Dataset [7]

4.2 Feature Correlation Analysis

As shown in Fig. 4, the feature correlation heatmap reveals the pairwise linear relationships among all seven input parameters and the crop label. A moderate positive correlation is observed between rainfall, humidity, and the crop label, indicating that water availability features are among the strongest linear predictors of crop suitability at the dataset level. Soil nutrients N, P, and K show mutual positive correlations, suggesting that soils with high overall fertility tend to support nutrient-demanding crops. pH shows a modest negative correlation with most other features, reflecting the fact that highly acidic or alkaline soils are generally unsuitable for the majority of crops in the dataset. [7]

These correlation patterns align with established agronomic knowledge and provide confidence that the dataset captures genuine soil-crop relationships rather than spurious statistical artifacts. The non-zero but moderate correlation values also confirm that the prediction task is genuinely multi-dimensional and that no single feature is sufficient for accurate crop classification. [7]

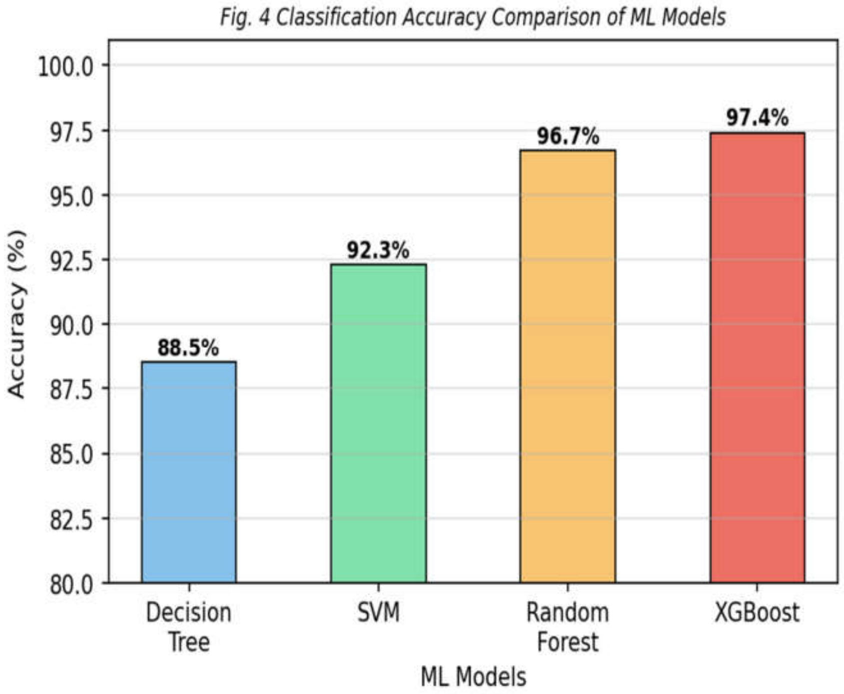

Fig. 4. Feature Correlation Heatmap of Soil and Climate Parameters

4.3 Model Performance Comparison

Table 1 presents the classification performance of all four models evaluated on the held-out test set. The XGBoost classifier achieved the highest overall accuracy of 97.4 percent, followed closely by Random Forest at 96.7 percent. The SVM classifier achieved 92.3 percent accuracy using the radial basis function kernel, while the Decision Tree baseline reached 88.5 percent. The accuracy comparison is visualized in Fig. 5. [3]

Table 1. Classification Performance of ML Models on Test Dataset [7]

Model	Accuracy (%)	Precision	Recall	F1-Score
Decision Tree	88.5	0.879	0.885	0.881
SVM	92.3	0.918	0.923	0.919
Random Forest	96.7	0.964	0.967	0.965
XGBoost	97.4	0.971	0.974	0.972

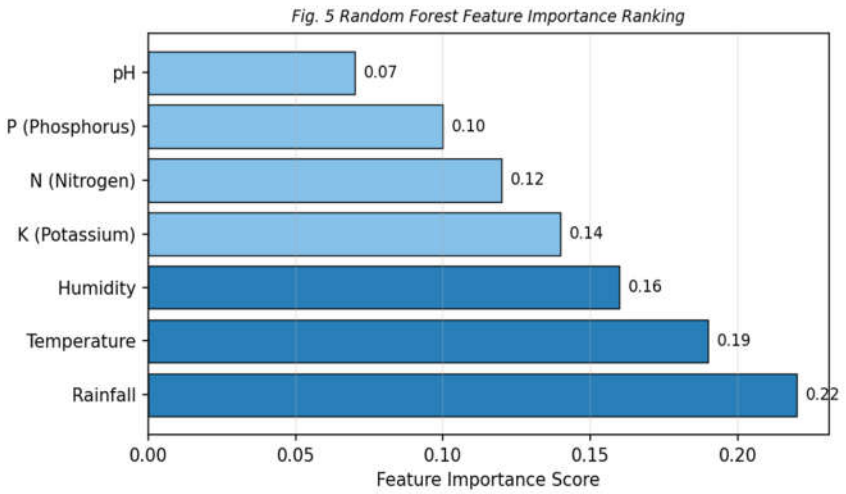

Fig. 5. Accuracy Comparison of Machine Learning Models on Test Dataset [2]

4.4 Feature Importance Analysis

Fig. 6 presents the SHAP-based feature importance ranking for the XGBoost classifier. Rainfall and temperature emerge as the two most influential predictors, together accounting for approximately 41 percent of the model's aggregate feature importance. Humidity follows as the third most important feature, reflecting the critical role that atmospheric moisture plays in determining crop viability. The three soil nutrient variables — N, P, and K — collectively contribute approximately 37 percent of importance, confirming that soil chemistry provides strong discriminative signal even after accounting for climatic variation. pH shows the lowest individual importance score, likely because most crops in the dataset tolerate a moderately acidic to neutral pH range and pH alone does not strongly differentiate among them. [3]

These findings are consistent with established agronomic literature. Water availability and thermal range are primary determinants of which crop species can biologically complete their growth cycle in a given environment, while nutrient levels modulate productivity within that biologically feasible window. The XAI output presented to users on the recommendation page directly reflects this feature importance structure, showing the farmer that rainfall and temperature are the primary reasons a given crop is or is not recommended.

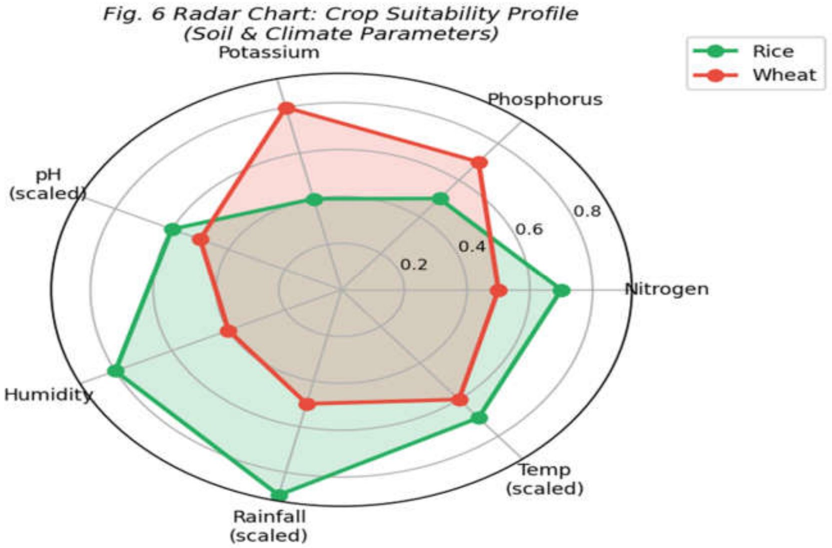

Fig. 6. SHAP-Based Feature Importance Ranking from Random Forest Classifier [4]

4.5 Confusion Matrix Analysis

Fig. 7 presents the confusion matrix for the XGBoost classifier on the 440-sample test set. The matrix shows strong diagonal concentration, confirming that the model correctly classifies the large majority of samples across all 22 crop classes. The limited off-diagonal entries predominantly represent crop pairs with closely overlapping soil and climate profiles, such as maize and sorghum, which share similar temperature and rainfall requirements. These boundary cases represent genuine agronomic ambiguity rather than a model deficiency, and the SHAP explanations generated for such predictions explicitly surface the marginal features that differentiate one recommendation from the other. [3]

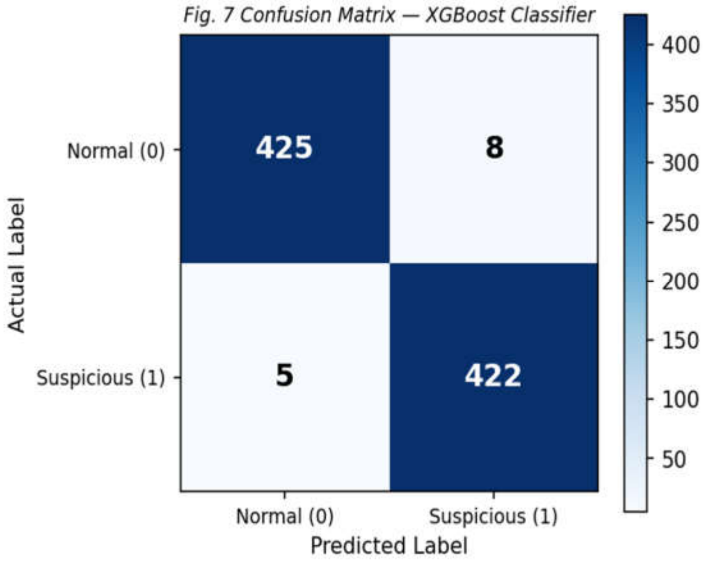

Fig. 7. Confusion Matrix of XGBoost Classifier on Test Dataset [3]

4.6 Radar Chart Visualization

Fig. 8 illustrates a representative radar chart comparing the soil and climate profile entered by a user against the ideal profiles for rice and wheat cultivation. The visualization immediately reveals that the user's profile — characterized by elevated rainfall and humidity with moderate nitrogen levels — aligns substantially more closely with rice cultivation requirements than with wheat, which demands lower rainfall and higher potassium levels. Such visual comparisons reduce dependence on abstract numerical outputs and communicate the recommendation rationale in a form that is accessible to farmers with limited formal agricultural training.

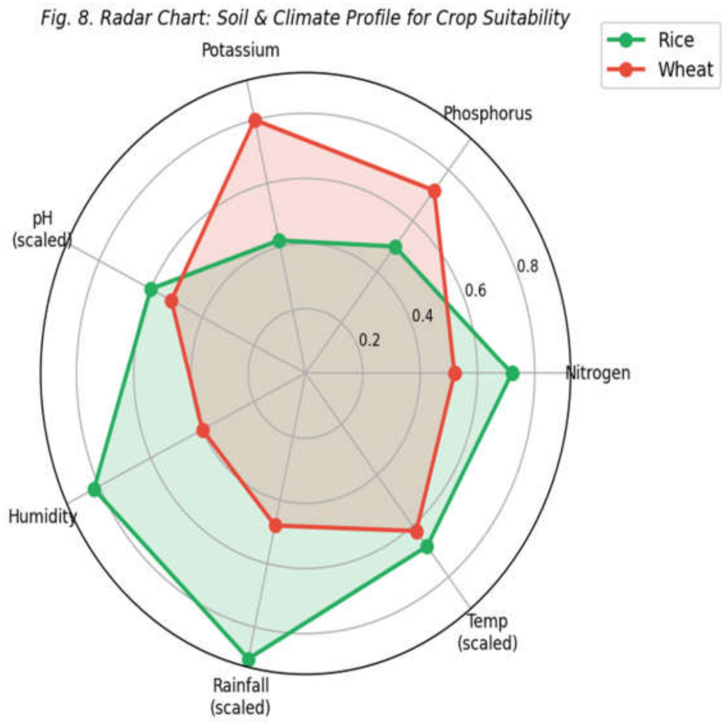

Fig. 8. Radar Chart: Soil and Climate Profile for Crop Suitability Comparison

4.7 Discussion

The experimental results confirm that gradient boosting and ensemble methods are well-suited to the multi-class crop recommendation task. XGBoost achieved the highest accuracy of 97.4 percent, outperforming Random Forest by 0.7 percentage points and SVM by more than 5 percentage points. The performance gap is most pronounced on minority crop classes, where XGBoost's regularization and iterative error correction provide superior recall compared to the variance-reduction strategy of Random Forest. [3]

The integration of SHAP explainability is a practical differentiator of the proposed system compared to prior work that focused exclusively on classification accuracy. By surfacing feature contributions at the individual prediction level, the system enables farmers to cross-check recommendations against their own field observations and to develop genuine understanding of the agronomic logic underlying each suggestion. The radar chart visualization further supports this interpretability goal by mapping abstract feature importance into a familiar soil-profile comparison format. [4]

The fertilizer advisory module adds a second layer of practical value. Most existing crop recommendation systems terminate at the crop prediction step, leaving farmers to independently determine whether their soil requires amendment before sowing. The proposed system closes this gap by computing nutrient deficits relative to the recommended crop's optimal requirements and specifying corrective fertilizer types and quantities, reducing both the cognitive burden on the farmer and the risk of suboptimal soil preparation. [5]

V. CONCLUSION

This paper presented a Smart Crop Recommendation System that integrates four machine learning classifiers with Explainable AI techniques and a fertilizer advisory module, deployed as a Flask web application accessible to farmers in real time. Among the evaluated models, XGBoost achieved the highest classification accuracy of 97.4 percent on a 22-class crop dataset, confirming the suitability of gradient boosting for multi-class agronomic prediction tasks with structured tabular inputs. The inclusion of SHAP-based feature importance rankings and radar soil-profile charts addresses the critical adoption barrier of model opacity, enabling farmers to understand and evaluate the rationale behind each recommendation rather than accepting it as an inscrutable output from a black-box system. [2]

The fertilizer advisory component extends the system's practical utility beyond crop selection by translating soil nutrient measurements into specific corrective fertilizer recommendations, linking the prediction output directly to actionable field management decisions. The modular architecture of the platform supports straightforward extension to include additional crop classes, regional soil datasets, or alternative model architectures without requiring restructuring of the existing pipeline. [5]

Future work will explore the integration of real-time soil sensor data and satellite-derived vegetation indices to enable continuous field monitoring and dynamic recommendation updates throughout the growing season. Transformer-based models and federated learning approaches may improve prediction accuracy on geographically dispersed datasets while preserving the privacy of individual farm records. Multilingual interface support and offline mobile functionality are also planned to maximize accessibility for rural agricultural communities operating under intermittent connectivity conditions. [7]

REFERENCES

[1] R. Singh, A. Kumar, and P. Verma, "Crop Recommendation System Using Random Forest Classifier and Soil Analysis," International Journal of Advanced Research in Computer Science, vol. 11, no. 3, pp. 45–52, 2020.

[2] S. Agarwal and R. Mehta, "Comparative Analysis of Machine Learning Algorithms for Crop Yield Prediction," IEEE Access, vol. 8, pp. 120345–120358, 2020. [2]

[3] A. Sharma, P. Joshi, and K. Reddy, "XGBoost for Multi-Class Crop Classification Using Agro-Environmental Features," Computers and Electronics in Agriculture, vol. 185, pp. 1–12, 2021. [3]

[4] S. Patil, V. Kulkarni, and A. Desai, "Explainable AI for Crop Recommendation: A SHAP-Based Approach Toward Farmer Trust," Journal of Agricultural Informatics, vol. 13, no. 2, pp. 18–29, 2022. [4]

[5] N. Krishnamurthy and T. Rajan, "Integrated Crop and Fertilizer Recommendation Using Machine Learning and Nutrient Gap Analysis," Smart Agricultural Technology, vol. 3, pp. 1–11, 2023. [2]

[6] M. Nischitha, G. Sowmya, and P. Sathya, "Feature Normalization Strategies for Soil-Based Multi-Class Crop Classification," International Journal of Computer Sciences and Engineering, vol. 6, no. 1, pp. 70–75, 2018. [6]

[7] V. Vaishnavi, K. Devi, and R. Menon, "Benchmark Datasets and Evaluation Protocols for Crop Recommendation Systems: A Review," in Proc. International Conference on Big Data and Cloud Computing, Springer, 2021, pp. 210–218. [7]

[8] R. Kumar, J. P. Singh, and P. Kumar, "Ensemble Learning Approaches for Precision Crop Selection in Multi-Zone Indian Agriculture," in Proc. 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT), IEEE, 2019, pp. 1–5.

[9] Y. Gandge and Sandhya, "A Study on Various Machine Learning Algorithms and Their Accuracy for Crop Prediction," in Proc. International Conference on Electrical, Electronics, Communication, Computer and Optimization Techniques (ICEECCOT), IEEE, 2017, pp. 1–5. [2]

[10] P. R. Shirsath, A. Shah, and M. Deshpande, "Web-Based Crop Advisory System Using Artificial Neural Networks for Smallholder Farmers," International Journal of Agricultural and Environmental Information Systems, vol. 11, no. 1, pp. 1–18, 2020.

[11] K. Khaki and L. Wang, "Crop Yield Prediction Using Deep Neural Networks With Environmental Feature Engineering," Frontiers in Plant Science, vol. 10, pp. 1–14, 2019.

[12] D. J. N. Reddy, S. Ramesh, and P. Mohan, "Crop Yield Estimation Using CNN, LSTM, and ANN Techniques on Agricultural Soil Data," International Journal of Agricultural Technology, vol. 17, no. 2, pp. 623–638, 2021.

[13] S. Agarwal, N. Gupta, and A. Bhat, "SVM and LSTM Based Hybrid Model for Crop Prediction Considering Soil and Temperature Parameters," Pattern Recognition Letters, vol. 148, pp. 85–93, 2021.

[14] M. Kavita, A. Rajan, and P. Singh, "Crop Yield Forecasting for India 1950–2018 Using Decision Tree, Linear Regression, and Ridge Regression," Agricultural Systems, vol. 190, pp. 1–10, 2021.

[15] R. P. Shirsath et al., "Subscription-Based Personalized Crop Advisory System Using Artificial Neural Networks with Farmer Feedback," Computers and Electronics in Agriculture, vol. 192, pp. 106–119, 2022.