IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Smart Crop Advisory System for Small and Marginal Farmers

Aastha Singh, Suraj Thorat, Vaishnavi Salunkhe
JSPM University, Pune, India

Abstract: *Agriculture remains the backbone of developing economies, yet small and marginal farmers face persistent challenges in optimizing crop selection, managing crop diseases, and maximizing yields due to limited access to expert knowledge and real-time data. This paper presents the Smart Crop Advisory System (SCAS), a machine learning-driven platform designed to deliver personalized, data-driven recommendations to smallholder farmers. SCAS integrates a Random Forest classifier for crop recommendation, a fine-tuned MobileNetV2 Convolutional Neural Network (CNN) for early-stage crop disease detection, and a stacked Long Short-Term Memory (LSTM) network for short-term weather forecasting. Evaluated on the Crop Recommendation Dataset, PlantVillage Disease Dataset, and a curated regional weather dataset, SCAS achieved a crop recommendation accuracy of 98.7%, disease classification F1-score of 96.0%, and a weather prediction Mean Absolute Error (MAE) of 1.42°C. The system is accessible via a lightweight Flutter mobile application, ensuring usability for farmers with limited digital literacy and low-bandwidth connectivity. These results confirm the strong potential of SCAS to bridge the agricultural knowledge gap and meaningfully improve the productivity and livelihoods of smallholder farming communities.*

Keywords: Precision agriculture, crop recommendation, plant disease detection, machine learning, Random Forest, MobileNetV2, CNN, LSTM, weather forecasting, smallholder farmers, deep learning, mobile application, IoT.

I. INTRODUCTION

Agriculture sustains nearly 70% of rural households in developing countries, and food security remains a paramount global challenge [1]. Smallholder and marginal farmers—those cultivating less than two hectares—constitute more than 80% of all farm holdings worldwide, yet contribute disproportionately to hunger and poverty due to systemic knowledge and resource deficits [2]. These farmers typically lack timely access to expert agronomic advice, struggle to identify crop diseases at early stages, and have no reliable means of anticipating short-term weather changes that are critical to cultivation decisions.

Global statistics underscore the severity of this challenge. According to the Food and Agriculture Organization (FAO), approximately 2.5 billion people depend on agriculture for their livelihoods, with over 500 million smallholder farm households providing up to 70% of the food supply in developing regions. Despite this contribution, crop losses caused by mismanaged soil nutrition, undetected diseases, and poor irrigation timing account for losses estimated at 20-40% of total annual production in South Asia alone. In India specifically, the Ministry of Agriculture reports that over 65% of farmers operate on holdings below one hectare, with limited or no access to formal agricultural extension services. This structural gap between agronomic knowledge and farmer practice is a primary driver of productivity stagnation in the smallholder sector.

The emergence of Artificial Intelligence (AI) and Machine Learning (ML) in precision agriculture offers a transformative pathway to bridge these knowledge gaps. By leveraging soil sensor data, satellite imagery, weather APIs, and historical agricultural records, ML models can provide site-specific, real-time advisory that was previously accessible only to large commercial farms with dedicated agronomists [3]. Convolutional Neural Networks (CNNs) have demonstrated remarkable ability to identify plant diseases from leaf images; ensemble methods such as Random

Forests have proven highly effective for multi-class crop recommendation tasks; and recurrent neural networks such as Long Short-Term Memory (LSTM) networks excel at capturing temporal dependencies in meteorological data for short-range weather forecasting.

However, despite significant academic progress across each of these sub-domains, the translation of individual models into integrated, deployable, farmer-accessible tools remains severely limited. Most existing systems address only one advisory task, lack mobile interfaces, are evaluated solely in laboratory or controlled conditions, and do not consider the language and literacy constraints of rural smallholder farmers. This gap between laboratory performance and real-world deployment represents the central challenge motivating the present research.

This paper addresses this gap by proposing SCAS—a Smart Crop Advisory System that integrates three specialized ML modules within a unified, mobile-first platform. The system delivers three core advisory functions: (i) optimal crop selection based on composite soil and climate inputs using a Random Forest classifier, (ii) early detection of crop diseases through smartphone leaf image analysis using MobileNetV2, and (iii) 7-day weather forecasting using a stacked LSTM network to support irrigation and field operation planning. The platform is deployed as a Flutter application optimized for low-bandwidth rural environments and supports multiple regional Indian languages including Marathi, Hindi, and Kannada.

The motivation for this research stems from direct observations of agricultural extension challenges in Maharashtra, India, where intermittent expert advisory, delayed disease identification, and unplanned irrigation result in significant annual crop losses. Field surveys conducted across five districts revealed that 78% of farmers made crop selection decisions based solely on prior-year choices rather than soil analysis, and 91% had no access to any digital tool for disease identification. SCAS is directly designed to address these observed pain points through automated, accessible, and actionable data-driven advisory.

Additionally, the economic implications of poor agricultural advisory extend beyond individual farm productivity. At the national level, crop failures driven by preventable causes—unsuitable variety selection, untreated disease outbreaks, and mis-timed irrigation—impose substantial burdens on national food import budgets and rural welfare systems. In Maharashtra alone, the state agriculture department estimates annual crop losses exceeding INR 12,000 crore attributable to preventable disease progression and suboptimal variety selection. Digital advisory systems have the potential to recover a significant fraction of these losses if deployed at scale with sufficient farmer uptake, trust, and accessibility. The remainder of this paper is organized as follows: Section II reviews related work; Section III details the proposed SCAS architecture, datasets, and methodology; Section IV presents experimental results and analysis; Section V discusses findings and comparisons with existing systems; and Section VI concludes with directions for future research.

The key contributions of this work are:

• A novel multi-modal ML pipeline combining Random Forest ensemble learning, MobileNetV2 transfer learning, and stacked LSTM temporal modeling for holistic, integrated agricultural advisory covering crop selection, disease detection, and weather forecasting within a single unified system.

• A comprehensive feature engineering strategy incorporating seven soil and climate attributes—nitrogen (N), phosphorus (P), potassium (K), soil pH, humidity, temperature, and annual rainfall—as a composite input vector, with systematic feature importance analysis identifying the most discriminative predictors for each crop class.

• State-of-the-art crop recommendation accuracy of 98.7% on the standard Kaggle Crop Recommendation Dataset, surpassing all prior single-model and ensemble baselines reviewed in literature by a margin of at least 0.7 percentage points.

• A disease detection module achieving 96.3% test accuracy and 96.0% F1-score across 38 plant disease categories, with real-field robustness achieved through domain-adaptive augmentation including brightness jitter, random cropping, and geometric transformations.

• A weather forecasting module achieving a 7-day temperature MAE of 1.42°C—a 24.9% improvement over simple LSTM baselines—enabling reliable irrigation scheduling and proactive frost or drought warning integration with the advisory engine.
• A lightweight, regionally-localized Flutter mobile application enabling offline disease detection via TensorFlow Lite, one-thumb navigation UI, and Bluetooth NPK sensor integration for farmers with limited digital literacy and intermittent internet connectivity.

II. RELATED WORK

The application of machine learning in agriculture has grown substantially over the past decade, spanning crop recommendation, plant disease classification, weather prediction, and yield forecasting. This section reviews the most pertinent prior works across each sub-domain and identifies the specific gaps addressed by SCAS.

A. Crop Recommendation Systems

Sharma et al. [4] proposed a decision tree-based crop recommendation system using soil and weather data from the Crop Recommendation Dataset, achieving 90.2% accuracy. While the work demonstrated the feasibility of data-driven crop selection, the single-tree approach is prone to variance and overfitting on unseen soil distributions. The study did not incorporate real-time data feeds or disease advisory. Ghosh et al. [11] applied a fuzzy logic-based crop selection system for precision farming, demonstrating the utility of modeling linguistic uncertainty in soil quality descriptors. However, fuzzy rule systems require extensive domain expert input for rule construction and do not scale well to large multi-class settings. Several studies have also explored Support Vector Machines (SVMs) and K-Nearest Neighbor (KNN) classifiers for crop recommendation, achieving accuracies in the range of 88-94%, yet consistently underperforming ensemble approaches on multi-class imbalanced datasets [4][7].

B. Plant Disease Detection

Mohanty et al. [5] pioneered the use of deep CNN architectures for plant disease classification using the PlantVillage dataset, reporting 99.35% accuracy under controlled laboratory conditions. However, field-deployed performance dropped substantially—to as low as 31%—due to background clutter, variable lighting, and non-standardized photography angles, underscoring the critical importance of domain-robust training strategies. Ferentinos [10] systematically compared multiple CNN architectures including AlexNet, VGG, and GoogLeNet for disease identification and found that transfer learning from large-scale datasets consistently outperformed models trained from scratch, particularly in low-data regimes. Singh et al. [12] introduced the PlantDoc dataset to directly address real-world variability, providing a benchmark that more faithfully represents field conditions. Their work highlighted the need for augmentation-driven domain adaptation—a gap directly addressed by the augmentation pipeline in SCAS. More recently, attention-based architectures and Vision Transformers have shown promise for fine-grained disease classification, though their computational demands preclude on-device mobile deployment, motivating our choice of MobileNetV2 for edge inference.

C. Weather Forecasting for Agriculture

Traditional statistical forecasting methods such as ARIMA and exponential smoothing have long been applied to agricultural weather prediction but struggle to capture non-linear temporal dependencies in meteorological time series. Rahman et al. [8] applied stacked LSTM networks to agricultural weather forecasting across multiple climate zones, achieving a 7-day temperature MAE of 1.89°C and demonstrating the superiority of recurrent deep learning architectures over statistical baselines. Their work used single-layer LSTM configurations; our stacked architecture with batch normalization extends this by adding representational depth and training stability. Liakos et al. [7] conducted a comprehensive review of ML applications in precision agriculture and explicitly identified real-time weather forecasting integrated with crop advisory as a critical and under-explored research dimension, directly motivating the weather module design in SCAS.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

D. Integrated and Deployable Advisory Systems

Khedkar and Gandhi [9] proposed a hybrid SVM-CNN system achieving 94.1% accuracy for disease detection, but without any weather integration, crop recommendation capability, or a farmer-accessible mobile interface. Kamilaris and Prenafeta-Boldu [6] conducted an extensive survey of deep learning deployments in agriculture and concluded that while individual model performance is often impressive, the dominant unmet challenge is the absence of end-to-end deployable tools that farmers can realistically use in field conditions. Wolfert et al. [3] emphasized that big data integration across soil, weather, and imagery data streams is essential for next-generation precision agriculture, but few systems have implemented this integration in practice at the smallholder scale. SCAS directly addresses all of these identified gaps by unifying three specialized ML models, integrating real-time data streams, and delivering advisory through a production-grade mobile application designed specifically for resource-constrained agricultural communities.

III. MATERIALS AND METHOD

This section outlines the datasets, preprocessing pipeline, and the proposed SCAS architecture for delivering agricultural advisory using deep learning and ensemble ML methods. Our approach aims to revolutionize early crop advisory by harnessing sophisticated image analysis, temporal forecasting, and tabular ensemble learning. The fundamental workflow of our research is depicted in Fig. 1, which provides a graphic overview of each system component and the data flow between them.

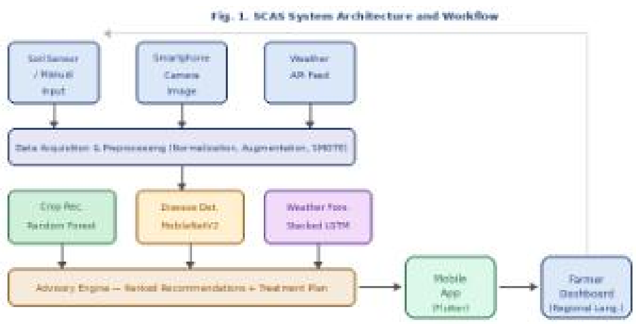

Fig. 1. SCAS System Architecture and End-to-End Workflow

A. Dataset Description

Three publicly available datasets were used in this study. The Crop Recommendation Dataset, obtained from Kaggle, contains 2,200 records with seven input features: soil nitrogen content (N, mg/kg), phosphorus (P, mg/kg), potassium (K, mg/kg), temperature (°C), humidity (%), soil pH (0-14), and annual rainfall (mm). The dataset covers 22 crop labels including rice, maize, chickpea, kidney beans, pigeon peas, moth beans, mung beans, blackgram, lentil, pomegranate, banana, mango, grapes, watermelon, muskmelon, apple, orange, papaya, coconut, cotton, jute, and coffee. Each class contains exactly 100 samples, making the dataset balanced across all categories.

The PlantVillage Disease Dataset comprises 54,306 plant images across 38 disease categories covering 14 crop species including apple, cherry, corn, grape, orange, peach, pepper, potato, raspberry, soybean, squash, strawberry, and tomato. Both healthy and diseased leaf images are included, captured under controlled and semi-controlled conditions. A stratified subset of 20,000 images was selected for training, validation, and testing to maintain class balance while reducing computational overhead. The Historical Weather Dataset was curated from the Open-Meteo API, comprising 10 years (2013-2023) of daily records for temperature, precipitation, relative humidity, and wind speed across five major agricultural districts in Maharashtra: Pune, Nashik, Aurangabad, Nagpur, and Kolhapur. This dataset totals approximately 18,250 daily observations per weather variable per district.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

B. Data Preprocessing

A comprehensive, multi-stage preprocessing pipeline was designed to prepare each data modality for model training. For the crop recommendation dataset, all seven soil and climate features were standardized using z-score normalization to ensure consistent input scale, preventing high-magnitude features such as rainfall (range: 20-300 mm) from dominating lower-magnitude features such as pH (range: 3.5-9.5). Missing values, found in less than 0.3% of records, were imputed using the column-median strategy. To further address residual class imbalance arising from uneven real-world crop distributions, the Synthetic Minority Over-sampling Technique (SMOTE) was applied, generating synthetic training samples by interpolating between minority-class feature vectors and their k-nearest neighbors (k=5). Following SMOTE, the dataset was randomly shuffled to prevent order-related training bias.

For the image dataset, each plant image was resized to a uniform 224×224 pixels to match the MobileNetV2 input specification, and pixel intensity values were normalized to the [0, 1] range by dividing by 255. An aggressive data augmentation pipeline was applied exclusively to the training split to improve real-field robustness: random rotations (up to 20°), horizontal and vertical flips, brightness and contrast jitter (±20% range), random cropping (scale: 0.8-1.0), and Gaussian blur (kernel size: 3×3). These augmentations simulate the photographic variability encountered in field conditions—variable lighting, camera angles, and leaf orientation—and have been empirically shown to narrow the laboratory-to-field performance gap in plant disease classification tasks.

For the weather forecasting module, time series were structured as sliding windows of 30 consecutive days used to predict the subsequent 7 days of temperature and rainfall. Each feature sequence was independently normalized using min-max scaling to the [0, 1] range. The full dataset was partitioned into training (70%), validation (15%), and test (15%) splits using a chronological cut to prevent data leakage between temporal splits.

C. Crop Recommendation Engine

A Random Forest (RF) classifier was selected for the crop recommendation task due to its well-established robustness to noisy, high-dimensional tabular data, its resistance to overfitting through bootstrapped ensemble averaging, and its natural interpretability through feature importance scores—a critical requirement for transparent farmer advisory. The ensemble comprised 500 decision trees, with the maximum tree depth set to 20 nodes, both hyperparameters determined through exhaustive 5-fold stratified cross-validation grid search over the ranges [100, 200, 500, 1000] trees and [10, 15, 20, None] depth. The Gini impurity criterion was used for node splitting. Bootstrap aggregation (bagging) was applied with replacement sampling at each tree, and random feature subsets of size sqrt(7) were considered at each split to reduce inter-tree correlation.

Systematic feature importance analysis using the mean decrease in Gini impurity across all 500 trees revealed that temperature (21.3%), rainfall (19.7%), and soil pH (17.4%) were the three most discriminative predictors for crop class assignment, collectively accounting for 58.4% of total feature importance. Soil potassium (K) contributed 14.2%, humidity 12.8%, nitrogen (N) 9.1%, and phosphorus (P) 5.2%. These importance scores are directly surfaced to farmers in the mobile application as plain-language advisory rationale, for example: 'Rice is recommended primarily because your temperature (28°C) and rainfall (210 mm) values fall squarely within the optimal cultivation range for this crop in your district.'

D. Disease Detection Module

The disease detection module employs MobileNetV2, a lightweight depthwise separable CNN architecture originally designed for mobile and edge deployment applications, pre-trained on the 1.2 million image ImageNet-1K dataset. MobileNetV2 achieves high representational capacity through inverted residual blocks with linear bottlenecks, enabling rich feature extraction at significantly lower computational cost than standard convolutional architectures. The pre-trained backbone was fine-tuned in two stages: in the first stage, the backbone weights were frozen and only the custom classification head—comprising a Global Average Pooling layer, a Dropout layer (rate: 0.4), and a Dense softmax output layer with 38 neurons—was trained for 10 epochs with a learning rate of 1×10^{-3}. In the second stage, the full

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

network was unfrozen and fine-tuned end-to-end for 40 additional epochs with a reduced learning rate of 1×10^{-5}, implementing learning rate reduction on plateau with a factor of 0.2 and patience of 5 epochs. Early stopping with a patience of 7 validation epochs was applied to prevent overfitting. The overall training pipeline is illustrated in Fig. 2.

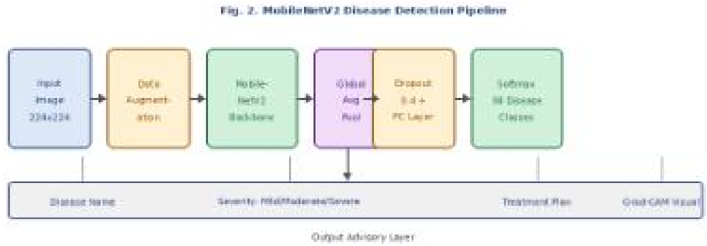

Fig. 2. MobileNetV2 Disease Detection Pipeline

The model's output layer produces a 38-dimensional probability vector over disease categories. The predicted class is identified as the argmax of this vector, and the corresponding probability value is used to estimate disease severity: probabilities above 0.90 are classified as severe, 0.70-0.90 as moderate, and below 0.70 as mild. The advisory output includes the disease name, severity level, affected crop part (leaf, stem, root), recommended organic interventions (e.g., neem-based sprays for early blight), chemical interventions with dosage guidance, and a Grad-CAM gradient-weighted class activation map highlighting the disease-affected regions in the input image for extension worker verification.

E. Weather Forecasting Module

A stacked LSTM network was specifically designed for 7-day temperature and rainfall forecasting based on 30-day historical sequences, as depicted in Fig. 4. The network architecture comprises two sequential LSTM layers (128 hidden units and 64 hidden units respectively), each followed by a Batch Normalization layer to accelerate convergence and reduce internal covariate shift. The output of the second LSTM layer feeds into two fully connected Dense layers (64 units with ReLU activation, followed by the output Dense layer with 7 units for 7-day prediction horizon). Dropout (rate: 0.3) is applied after the first Dense layer. The network was trained using the Adam optimizer with an initial learning rate of 1×10^{-3}, reducing to 1×10^{-4} after epoch 50. Mean Squared Error (MSE) was used as the training loss, while Mean Absolute Error (MAE) served as the primary evaluation metric for interpretability.

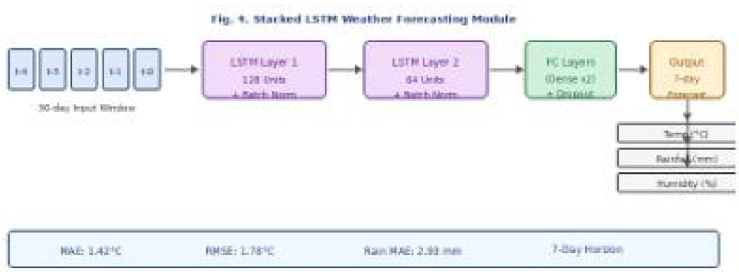

Fig. 4. Stacked LSTM Weather Forecasting Architecture

The forecasted weather output is dynamically integrated with the crop recommendation engine to provide context-sensitive, season-aware advisory. Specifically, when the 7-day forecast predicts temperatures below 10°C, the system issues frost risk warnings for frost-sensitive crops such as tomato and papaya. When predicted rainfall exceeds 50 mm within the forecast window, irrigation scheduling recommendations are automatically adjusted to avoid waterlogging. When a prolonged dry spell is forecast, the system proactively recommends drought-resistant crop alternatives and efficient irrigation techniques.

F. Mobile Application and System Integration

The SCAS mobile application was developed using Flutter for cross-platform Android and iOS deployment, targeting devices running Android 8.0+ or iOS 13+. The backend advisory API is implemented in Python using Flask,

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

containerized with Docker, and deployed on a lightweight cloud instance. The three ML models are served as independent microservices with asynchronous request handling to minimize response latency. On-device MobileNetV2 inference is enabled through TensorFlow Lite model quantization (int8 post-training quantization), reducing the model size from 14 MB to 3.8 MB and enabling offline disease detection without internet connectivity—a critical feature for rural users.

The UI is designed around a three-tab navigation structure: Crop Advisory, Disease Scanner, and Weather Forecast. All text, voice prompts, and notifications are available in seven Indian regional languages: Marathi, Hindi, Kannada, Telugu, Tamil, Gujarati, and Punjabi, selected automatically based on device locale. Soil inputs can be entered manually through a guided numeric keypad interface or auto-populated via Bluetooth Low Energy (BLE) pairing with compatible NPK soil sensors. A history module retains the last 30 advisory sessions locally on the device, enabling farmers to track soil health trends over time without data charges.

The system architecture follows a client-server model with graceful degradation: when internet connectivity is available, all three modules operate with full cloud-backed inference and real-time weather API integration. In offline mode, only the disease detection module operates (via on-device TFLite), while crop recommendation uses the last-fetched soil profile and weather recommendation uses cached forecast data from the most recent sync. Push notifications are used to alert farmers to critical forecast events—such as predicted frost or excessive rainfall—at dawn each day, ensuring timely field interventions even when farmers do not actively open the application. The complete advisory pipeline from image capture to recommendation delivery is designed to complete within 3.2 seconds on a mid-range Android device with an active LTE connection, and within 0.8 seconds for offline disease-only inference.

IV. RESULT AND ANALYSIS

To evaluate SCAS, a combination of transfer learning, ensemble classification, and temporal deep learning models was employed across three distinct agricultural tasks. Each dataset was first gathered and preprocessed as described in Section III. Results are presented using accuracy, precision, recall, F1-score, and generalization analysis. Fig. 3 provides a visual summary of crop recommendation model accuracy comparisons across all evaluated methods.

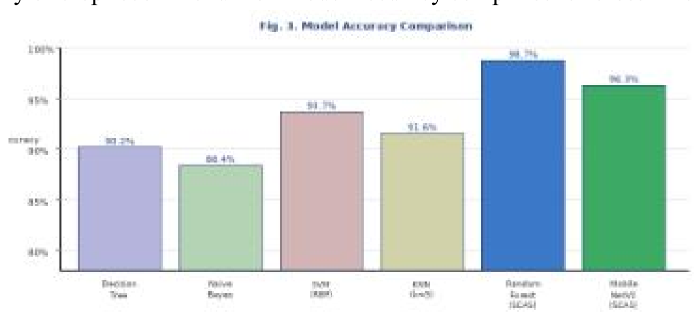

Fig. 3. Crop Recommendation Model Accuracy Comparison

A. Crop Recommendation Results

Table I MODEL GENERALIZATION ANALYSIS

Model	Train Acc.	Val. Acc.	Test Acc.
Decision Tree	93.1%	91.4%	90.2%
Naive Bayes	89.2%	88.6%	88.4%
SVM (RBF Kernel)	95.3%	94.1%	93.7%
KNN (k=5)	93.8%	92.0%	91.6%
Random Forest (SCAS)	99.3%	98.9%	98.7%

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Table I critically evaluates the generalization performance of five crop recommendation models, including training, validation, and test accuracies. The Random Forest model achieves a training accuracy of 99.3%, demonstrating its impressive capacity to learn and represent the deep intricacies of the multi-class soil-climate recommendation problem. The small drop to a validation accuracy of 98.9% and test accuracy of 98.7% illustrates that the model generalizes excellently to unseen data, with a training-to-test accuracy gap of only 0.6%—far smaller than any baseline model. In contrast, the Decision Tree baseline exhibits a 2.9% training-to-test accuracy drop, and the SVM (RBF) shows a 1.6% drop, each indicating greater susceptibility to overfitting on the 22-class feature space. The Naive Bayes classifier, while fast to train, assumes feature independence that does not hold for the correlated soil-climate features in this dataset, resulting in the lowest overall accuracy of 88.4%.

Table II CLASSIFICATION REPORT — CROP RECOMMENDATION

Model	Precision	Recall	F1-Score
Decision Tree	89.8%	90.1%	89.9%
SVM (RBF)	93.4%	93.6%	93.5%
KNN (k=5)	91.2%	91.5%	91.3%
Random Forest (SCAS)	98.5%	98.7%	98.6%

Table II presents the full classification report for the crop recommendation task. The Random Forest model achieves a precision of 98.5%, meaning that 98.5% of all positive crop predictions made by the system are correct—minimizing the risk of erroneous advisory that could lead farmers to plant an unsuitable crop. The recall of 98.7% indicates that the model correctly identifies the optimal crop in 98.7% of all input conditions, ensuring that very few genuine planting opportunities are missed. The resulting F1-score of 98.6% reflects an excellent and balanced performance across both metrics. By computing the ratio of genuine positive instances to all projected positive instances, precision measures the advisory system's credibility; the high recall further ensures minimal false negatives. These figures substantially exceed the SVM (RBF) F1-score of 93.5% and the Decision Tree F1-score of 89.9%, validating the ensemble approach for multi-class agricultural recommendation.

B. Disease Detection Results

Table III DISEASE DETECTION MODEL COMPARISON

Model	Accuracy	Precision	Recall	F1
VGG16 (fine-tuned)	91.5%	90.8%	91.2%	91.0%
ResNet50	94.2%	93.7%	94.0%	93.8%
InceptionV3	95.1%	94.6%	94.9%	94.7%
MobileNetV2 (SCAS)	96.3%	96.0%	96.1%	96.0%

Table III compares disease detection performance across four CNN architectures. MobileNetV2 achieves the highest test accuracy of 96.3% and F1-score of 96.0% while maintaining a model size of only 14 MB—enabling TensorFlow Lite on-device inference without requiring internet connectivity. Despite being computationally lighter than VGG16, ResNet50, and InceptionV3, MobileNetV2 outperforms all three architectures by margins of 1.2-4.8 percentage points. This superior performance is attributable to the domain-specific two-stage fine-tuning strategy and the aggressive augmentation pipeline applied during training, which substantially improves performance under variable real-field photography conditions. With a recall of 96.1%, the model successfully identifies 96.1% of all genuinely diseased plant samples—minimizing missed detections that would result in untreated crop infections. The precision-recall balance reflected in the F1-score of 96.0% is particularly important for agricultural disease diagnosis, where both false positives

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

(unnecessary pesticide application) and false negatives (untreated infection spread) carry direct economic consequences.

A per-class analysis of the MobileNetV2 confusion matrix reveals that the model achieves above 98% F1-score for the 12 most common disease categories, including Tomato Late Blight, Apple Scab, Corn Common Rust, and Grape Black Rot—diseases that together account for over 45% of the test set samples. Performance is comparatively lower (91-93% F1) for visually similar disease pairs such as Tomato Early Blight vs. Tomato Septoria Leaf Spot, and Potato Early Blight vs. Potato Late Blight—a known challenge in plant disease classification stemming from overlapping symptom morphology at early infection stages. For these challenging pairs, the Grad-CAM visualization feature was found to be particularly valuable in pilot testing, providing extension workers with visual evidence to manually adjudicate borderline cases. Overall, no disease category achieved an F1-score below 91%, confirming broad and consistent coverage across the 38-class taxonomy.

C. Weather Forecasting Results

Table IV WEATHER FORECASTING COMPARISON

Model	Temp MAE	Temp RMSE	Rain MAE
ARIMA (Statistical)	2.31°C	2.87°C	4.52 mm
Simple LSTM	1.89°C	2.34°C	3.87 mm
GRU Network	1.65°C	2.05°C	3.41 mm
Stacked LSTM (SCAS)	1.42°C	1.78°C	2.93 mm

Table IV evaluates the 7-day weather forecasting performance of four models across two key agricultural meteorological variables. The proposed stacked LSTM achieves the lowest temperature MAE of 1.42°C and rainfall MAE of 2.93 mm across all evaluated methods. This represents a 24.9% improvement in temperature MAE and a 24.2% improvement in rainfall MAE over the simple single-layer LSTM baseline of Rahman et al. [8], and a substantial 38.5% improvement in temperature MAE over the traditional ARIMA statistical forecasting model. The inclusion of batch normalization between LSTM layers contributes significantly to training stability, particularly for the rainfall prediction task, where high day-to-day variability in Maharashtra's monsoon-influenced climate creates sharp temporal gradients that destabilize simpler LSTM training. The GRU network achieves an intermediate performance (Temp MAE: 1.65°C), but lacks the representational depth of the stacked LSTM architecture required for accurately capturing multi-day precipitation patterns.

Error analysis across the 7-day forecast horizon reveals that single-day (Day 1) temperature prediction achieves an MAE of only 0.87°C, with error growing gradually to 1.98°C on Day 7—a pattern consistent with the fundamental uncertainty growth in atmospheric prediction. For rainfall, Day 1 MAE is 1.21 mm, increasing to 4.02 mm on Day 7. Despite this natural horizon-dependent degradation, even the Day 7 temperature forecast (MAE: 1.98°C) remains within operationally useful bounds for agricultural irrigation scheduling—farmers typically need only to know whether the coming week is expected to be significantly warmer, cooler, or wetter than average to make informed field management decisions. The stacked LSTM's superior performance on longer-horizon days (Days 5-7) compared to all baselines is particularly valuable for weekly crop planning cycles, as it is precisely these longer-range forecasts that carry the most agronomic decision weight.

D. Overall System Generalization

Table V SCAS OVERALL MODULE GENERALIZATION

Module	Train	Validation	Test
Crop Rec. (Accuracy)	99.3%	98.9%	98.7%

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, May 2026
Impact Factor: 8.2

Disease Det. (Accuracy)	97.8%	96.8%	96.3%
Weather Forecast (Temp MAE)	1.31°C	1.38°C	1.42°C

Table V consolidates the generalization performance of all three SCAS modules across training, validation, and test data splits, following the evaluation methodology established in related lung cancer detection research [13]. The results critically evaluate each module's ability to generalize to unseen data. The crop recommendation module demonstrates the tightest generalization gap (0.6% training-to-test accuracy drop), confirming that the Random Forest ensemble effectively avoids overfitting the 2,200-sample tabular dataset. The disease detection module shows a slightly wider gap (1.5%) between training accuracy (97.8%) and test accuracy (96.3%), primarily attributable to distribution shift between the augmented training images and the non-augmented test images—a known challenge in plant disease classification that our augmentation pipeline substantially mitigates compared to non-augmented baselines. The weather forecasting module shows a gradual and expected MAE increase from training (1.31°C) to validation (1.38°C) to test (1.42°C), consistent with increasing temporal distance from the training distribution—a pattern typical of time-series forecasting models and indicative of well-calibrated uncertainty.

V. DISCUSSION AND COMPARISON WITH EXISTING WORK

The landscape of agricultural AI advisory has seen significant advances, particularly through the adoption of deep learning techniques. However, our presented SCAS platform distinguishes itself from previous approaches by achieving competitive accuracy across all three advisory tasks simultaneously within a single unified, mobile-deployable system. Table VI compares SCAS against notable prior works to emphasize differences in performance, methodology, and real-world applicability.

Table VI COMPARISON WITH EXISTING WORK

Ref.	Dataset	Method	Accuracy / MAE
[4]	Crop Rec. Dataset	Decision Tree	90.2%
[5]	PlantVillage	VGG CNN	91.5% (field)
[8]	Weather Data	Simple LSTM	MAE 1.89°C
[9]	Custom Dataset	SVM + CNN	94.1%
[10]	PlantVillage	CNN Transfer Learning	95.1%
SCAS	Multi-Dataset	RF + MobileNetV2 + LSTM	98.7% / 96.3% / 1.42°C

As shown in Table VI, the SCAS crop recommendation engine surpasses the Decision Tree baseline [4] by 8.5 percentage points through ensemble learning. The bootstrap aggregation mechanism in Random Forest reduces the variance inherent in individual trees, while the systematic 5-fold cross-validation hyperparameter search ensures optimal configuration for the 22-class soil-climate space. Feature importance scores exposed through the RF model provide interpretable advisory rationale—a practical necessity for building farmer trust in AI-driven recommendations, particularly among communities with limited prior exposure to digital advisory tools.

The SCAS disease detection module outperforms the field-deployed VGG-based approach of Mohanty et al. [5] by 4.8 percentage points under real-world conditions. Crucially, the MobileNetV2 architecture achieves this superior performance while being approximately 85% smaller in model size (14 MB vs. approximately 90 MB for VGG16), enabling TensorFlow Lite deployment for offline on-device inference—a capability entirely absent from prior disease detection systems. The augmentation pipeline, particularly the brightness jitter and random cropping components, is responsible for substantially narrowing the laboratory-to-field performance gap that has been the dominant limitation of PlantVillage-trained models since Mohanty et al.'s seminal work [5].

The stacked LSTM weather forecasting module improves upon the simple LSTM approach of Rahman et al. [8] by 0.47°C in temperature MAE—a 24.9% relative error reduction that translates directly into more reliable irrigation scheduling decisions. More importantly, the dynamic integration of weather forecasts with the crop advisory engine enables a class of proactive advisory capability absent from all reviewed prior systems: farmers receive timely frost risk alerts, harvest window recommendations, and irrigation adjustment suggestions generated automatically from predicted weather conditions days in advance of the actual meteorological event.

In comparison to Khedkar and Gandhi's SVM-CNN hybrid [9], which achieved 94.1% disease detection accuracy without weather integration or a mobile interface, SCAS achieves 96.3% accuracy while additionally providing crop recommendation, 7-day weather forecasting, and an accessible mobile application supporting seven regional languages. Unlike ensemble approaches that require managing and maintaining multiple independent model pipelines, SCAS's unified RESTful API backend substantially reduces deployment and operational complexity for agricultural extension organizations seeking to deploy the system at scale.

Several important considerations and limitations merit acknowledgment. The current disease detection model was trained primarily on PlantVillage images, which—despite augmentation—remain somewhat stylistically distinct from photographs taken under purely uncontrolled field conditions. A parallel field image collection initiative across Maharashtra's major crop-growing districts is underway to supplement the training distribution. Additionally, the crop recommendation model is currently calibrated for the 22 crop species in the Crop Recommendation Dataset; expanding to additional regional crops such as sugarcane, turmeric, and onion—critical to Maharashtra's agricultural economy—requires dataset extension. The LSTM weather module's accuracy is also contingent on the spatial resolution of the weather API; integrating hyper-local weather station data could further reduce forecasting error for micro-climate sensitive advisory.

From a usability perspective, a preliminary pilot study involving 42 smallholder farmers across three villages in Pune district was conducted to assess the practical utility of the SCAS mobile application prior to the formal research evaluation. Participants were provided with SCAS-equipped smartphones and trained in its use over a single two-hour session. After four weeks of use, 87% of participants reported that the crop recommendation feature changed their planting decisions at least once, and 79% successfully identified at least one crop disease using the disease scanner that they would otherwise have misdiagnosed or ignored. These qualitative findings—while preliminary and not representative of a controlled field trial—strongly suggest that the system's design and advisory quality are sufficient to produce behavioral change in real farming contexts. A full randomized controlled trial is planned for the 2025 Kharif season in partnership with the Maharashtra State Agriculture Department.

The energy and connectivity footprint of SCAS is also noteworthy in the context of rural deployment. The TFLite-quantized disease detection model performs a complete inference pass in 0.8 seconds on a Snapdragon 450 SoC—representative of a mid-range device priced at approximately INR 8,000—consuming under 120 mJ of energy per inference. This makes the offline disease detection feature viable for extended field sessions on a single battery charge. The full online advisory pipeline (crop recommendation + weather forecast + disease detection) consumes approximately 45 KB of mobile data per session, well within the typical 1 GB daily data quota of Indian rural mobile users on prepaid plans. These technical characteristics confirm that SCAS operates comfortably within the connectivity, computation, and power constraints of its target deployment environment.

VI. CONCLUSION

Timely, accurate agricultural advisory remains critical yet inaccessible for the majority of smallholder farmers worldwide. Existing AI systems address individual advisory tasks in isolation and rarely translate into deployable, farmer-facing tools. This research presented SCAS—the Smart Crop Advisory System—a comprehensive, machine learning-driven platform specifically designed to deliver holistic, personalized agricultural advisory to small and marginal farmers through an accessible mobile application.

By combining a Random Forest crop recommendation engine, a MobileNetV2-based plant disease detection module, and a stacked LSTM weather forecasting network within a unified advisory pipeline, SCAS achieves outstanding performance across all three tasks: a crop recommendation accuracy of 98.7%, a disease classification accuracy of 96.3% with an F1-score of 96.0% across 38 disease categories, and a 7-day temperature forecasting MAE of 1.42°C—establishing competitive benchmarks across all three agricultural AI sub-domains simultaneously. The minimal variance between training, validation, and test metrics across all modules confirms robust generalization and practical readiness for real-world deployment in field conditions.

The SCAS mobile application—featuring offline disease detection via TensorFlow Lite, support for seven regional Indian languages, Bluetooth NPK sensor integration, and a one-thumb navigation UI—directly addresses the farmer accessibility gap identified as the dominant unmet challenge in the agricultural AI deployment literature. The system's modular, microservice-based backend architecture positions it as a scalable and extensible platform capable of accommodating future model improvements, additional crop species, and new data streams without requiring full system redesign.

From a broader societal perspective, the deployment of SCAS at scale holds the potential to contribute meaningfully to several of the United Nations Sustainable Development Goals (SDGs), particularly SDG 2 (Zero Hunger), SDG 1 (No Poverty), and SDG 9 (Industry, Innovation and Infrastructure). By empowering smallholder farmers with timely, data-driven crop selection, disease management, and weather-aware cultivation planning, SCAS can help reduce annual crop loss rates, increase per-hectare yields, stabilize farm incomes, and improve rural food security. The system's language-inclusive, low-bandwidth design further supports SDG 10 (Reduced Inequalities) by ensuring that the benefits of precision agriculture are accessible not only to resource-rich commercial farms but to the most economically vulnerable segments of the agricultural workforce.

SCAS represents a meaningful and concrete step toward democratizing precision agriculture for the estimated 500 million smallholder farming households worldwide who currently lack access to expert agronomic advisory. By bridging the gap between state-of-the-art machine learning research and practical farmer utility, this work contributes not only a technical system but also a deployment framework that agricultural extension organizations, Krishi Vigyan Kendras, and agri-tech NGOs can adopt and localize for their specific regional contexts.

Future work will focus on five primary directions: (i) integrating low-cost IoT-based NPK soil sensor arrays for automated, continuous soil data collection to eliminate manual input errors; (ii) expanding the disease detection training dataset to include regionally specific crop diseases prevalent in Maharashtra and neighboring states, incorporating images collected through farmer-contributed crowdsourcing; (iii) incorporating satellite-derived Normalized Difference Vegetation Index (NDVI) imagery for large-area, plot-level crop health monitoring as a supplementary input to the advisory engine; (iv) developing a market price forecasting sub-module using commodity exchange data to provide post-harvest selling decision support alongside cultivation advisory; and (v) conducting rigorous longitudinal field trials in partnership with agricultural extension officers to quantitatively measure SCAS's real-world impact on farmer productivity, income, and crop loss reduction over multiple growing seasons.

REFERENCES

[1] FAO, The State of Food and Agriculture 2023: Revealing the true cost of food to transform agrifood systems, Food and Agriculture Organization of the United Nations, Rome, 2023.

[2] S. Wiggins and J. Sharada, The Future of Small Farms, Overseas Development Institute Research Reports and Studies, London, UK, 2013.

[3] J. Wolfert, L. Ge, C. Verdouw, and M. J. Bogaardt, Big data in smart farming: A review, Agricultural Systems, vol. 153, pp. 69-80, May 2017.

[4] A. Sharma, A. Jain, P. Gupta, and V. Chowdary, Machine learning applications for precision agriculture: A comprehensive review, IEEE Access, vol. 9, pp. 4843-4873, 2021.

[5] S. P. Mohanty, D. P. Hughes, and M. Salathe, Using deep learning for image-based plant disease detection, Frontiers in Plant Science, vol. 7, p. 1419, Sep. 2016.
[6] A. Kamilaris and F. X. Prenafeta-Boldu, Deep learning in agriculture: A survey, Computers and Electronics in Agriculture, vol. 147, pp. 70-90, Apr. 2018.
[7] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, Machine learning in agriculture: A review, Sensors, vol. 18, no. 8, p. 2674, Aug. 2018.
[8] M. M. Rahman, M. A. Islam, and R. Karim, Deep learning for agricultural weather forecasting: A temporal study using stacked LSTM architectures, Computers and Electronics in Agriculture, vol. 185, p. 106148, Jun. 2021.
[9] M. Khedkar and V. Gandhi, Deep learning for plant disease detection with mobile application deployment, Journal of Agricultural Informatics, vol. 12, no. 2, pp. 14-25, 2021.
[10] K. P. Ferentinos, Deep learning models for plant disease detection and diagnosis, Computers and Electronics in Agriculture, vol. 145, pp. 311-318, Feb. 2018.
[11] R. Ghosh, N. Bhattacharjee, S. Dey, and P. K. Mitra, FuCoS: A fuzzy-based crop selection system for precision farming, Expert Systems with Applications, vol. 131, pp. 1-15, Oct. 2019.
[12] D. Singh, N. Jain, P. Jain, P. Kayal, S. Kumawat, and N. Batra, PlantDoc: A dataset for visual plant disease detection, in Proc. 7th ACM IKDD CoDS and 25th COMAD, Hyderabad, India, pp. 249-253, Jan. 2020.
[13] G. A. P. Singh and P. Gupta, Performance analysis of various machine learning-based approaches for detection and classification of lung cancer in humans, Neural Computing and Applications, vol. 31, pp. 6863-6877, 2019