Terra Track: An IoT Based ML Driven Crop Recommendation System

Dr. Rajani S Hardas*1
Department of Electrical Engineering
PES's Modern College of Engineering
Pune, India
rajani.hardas@moderncoe.edu.in

Radha Chandrashekhar*2
Department of Electrical Engineering
PES's Modern College of Engineering
Pune, India
chandrashekhar_radha_elect@moderncoe.edu.in

Meet Kulkarni*3
Department of Electrical Engineering
PES's Modern College of Engineering
Pune, India
meet_kulkarni_elect@moderncoe.edu.in

Samruddhi Kadlag*4
Department of Electrical Engineering
PES's Modern College of Engineering
Pune, India
samrudhhi_kadlag_elect@moderncoe.edu.in

Dr. S.N.Chaphekar*5
Department of Electrical Engineering
PES's Modern College of Engineering
Pune, India
Shalaka.chaphekar@moderncoe.edu.in

Abstract - TerraTrack is a portable, handheld IoT system that delivers real-time, in-field soil analysis without reliance on laboratory infrastructure. The system measures pH, Nitrogen (N), Phosphorus (P), and Potassium (K) using industrial-grade RS485 MODBUS-RTU sensors (JXBS-3001 and JXCT NPK series), interfaced with an ESP32-WROOM-32D microcontroller via dual MAX485E RS485-to-TTL converter modules, ensuring noise-immune communication and wireless data transmission in place of days-long laboratory turnarounds. Acquired soil data feeds into a Random Forest Classifier trained across 16 crop classes, achieving a test accuracy of 50.78%. To address the accuracy ceiling from overlapping feature distributions, a weighted scoring system (α=0.7, β=0.3) fuses ML suitability scores with district-level regional popularity data, producing agronomically grounded and regionally practical recommendations. A seasonal filter ensures temporal actionability across Kharif, Rabi, and perennial cycles. The pipeline was validated across 4 Maharashtra districts and 12 seasonal test cases. The seasonal filter performed correctly in all instances, and suitability scores produced a geographically coherent gradient consistent with known agroclimatic profiles. The system is powered by a 3S Li-ion pack (11.1V, 2200mAh) with regulated power distribution, ensuring stable field performance and full portability.

Keywords: Precision Agriculture 4.0, In-situ NPK & pH analytics, Random Forest Crop Recommendation, Hybrid BLE & IoT Connectivity, Real-time Agro-Decision Support System (DSS)

INTRODUCTION

Soil health is a critical determinant of agricultural productivity, yet its condition cannot be assessed visually. Healthy and depleted soils are often indistinguishable to the naked eye, making systematic testing essential for informed farming decisions. Soil testing reveals key parameters including pH, macronutrients (NPK), electrical conductivity (EC), organic matter content, and micronutrients, each of which directly governs crop yield and input efficiency [7],[8]. Soil undergoes silent degradation: nutrient depletion through harvesting, acidification from repeated fertilizer use, salt accumulation in irrigated fields, and loss of organic carbon and without testing these changes go unnoticed. Indian agricultural soils are particularly affected, with over half estimated to be zinc-deficient and many falling below the 0.5% organic carbon threshold needed for basic soil function [8]. Farmers relying on intuition or historical practices frequently misallocate inputs, spending on the wrong nutrients or applying outdated doses to fields whose fertility has already declined.

India's primary response has been the Soil Health Card scheme, operating through roughly 1,000 to 1,200 government laboratories. While over 100 million cards have been distributed, the system is constrained by structural limitations: results take two to six weeks,. often arriving after planting decisions have already been made, and remote smallholders face real logistical barriers in accessing collection and distribution points. The cards themselves require agronomic literacy to interpret, limiting their practical value without extension support [9]. Portable IoT-based soil sensors represent a fundamentally different approach. Probes capable of measuring pH, EC, moisture, temperature, and NPK in the field within seconds, paired with microcontrollers and smartphone interfaces, can deliver soil intelligence at the point of decision without laboratory dependence. This does not replace laboratory analysis for complex parameters, but addresses the gap that laboratory testing was never designed to fill: routine, high-frequency, field-level measurements accessible to every farmer at the moment they are needed [9].

TerraTrack advances the state of field-deployable crop recommendation systems through the following primary contributions:

Three-stage ML-regional-seasonal fusion pipeline: A structured recommendation engine that combines a Random Forest suitability score with district-level regional crop popularity (α=0.7, β=0.3) and a seasonal filter calibrated to Maharashtra's Kharif, Rabi, and Zaid cropping calendar. This three-stage architecture directly addresses the known accuracy ceiling of ML-only approaches by resolving low-confidence classification decisions using regional and temporal context that the model alone cannot access, producing agronomically grounded, regionally practical, and temporally actionable recommendations across 16 crop classes.

Regionally validated synthetic training dataset from Package of Practices: A crop-specific training dataset of 1,920 samples constructed from the Joint Agresco Recommendations (2021–2024) and Krishi Darshani (2025–26) published by MPKV Rahuri and VNMKV Parbhani, encoding agronomic parameter ranges for 16 crops dominant in Maharashtra. Unlike publicly available generic datasets, this dataset encodes region-specific soil and nutrient requirements, ensuring recommendation credibility for Maharashtra's diverse agroclimatic zones.

Field-deployable dual-channel RS485 MODBUS-RTU sensing architecture: An embedded IoT platform integrating two industrial-grade RS485 sensors (JXBS-3001 pH and JXCT NPK) via independent MAX485E transceiver channels into a single ESP32-WROOM-32D microcontroller. The use of differential RS485 signalling over separate UART buses eliminates bus contention and ensures noise-immune soil parameter acquisition in electrically noisy outdoor environments, a deliberate design choice over consumer-grade I2C or analogue alternatives reviewed in the literature.

Integrated portable system enabling point-of-decision soil intelligence: A handheld, battery-powered device with Bluetooth-based wireless data transmission to a companion Flutter mobile application, enabling real-time NPK and pH readings to reach a farmer's smartphone without laboratory infrastructure, fixed power supply, or physical tethering. The complete pipeline, from in-field sensor reading through AI-generated crop recommendation operates within a single field session.

Accurate and timely soil data remains inaccessible to most smallholder farmers due to the cost, delay, and logistical burden of laboratory-based testing [9]. TerraTrack addresses this by integrating industrial-grade JXBS-3001 RS485 MODBUS-RTU sensors with an ESP32-WROOM microcontroller to enable real-time, in-field measurement of pH, and NPK levels, eliminating the dependence on centralized laboratory infrastructure[1],[2].

Reliable sensor communication in outdoor, electrically noisy field environments is a known challenge for embedded soil monitoring systems [3]. This is addressed through the implementation of the RS485 differential signalling standard via MAX485 converter modules, ensuring noise-immune, long-distance data transmission between sensors and the microcontroller under real agricultural conditions.

Portability and sustained field operation are essential requirements for a soil monitoring device intended for use by farmers across varied terrain[4],[5]. TerraTrack incorporates a Li-ion battery system with regulated power distribution to all sensing and processing modules, enabling stable, untethered operation without reliance on grid power.

Raw sensor readings carry limited practical value unless they can be interpreted without agronomic expertise[8],[9]. The system therefore pairs its hardware with a structured data presentation layer that translates measured soil parameters into actionable fertilizer and amendment guidance, making the output usable by farmers regardless of their technical background.

Continuous and repeatable soil monitoring across growing seasons is necessary to detect the gradual shifts in nutrient levels, pH, and salinity that laboratory snapshots routinely miss[8]. TerraTrack is designed for repeated field deployment, supporting the kind of frequent, site-specific measurements that allow farmers to track soil health trends and respond before degradation becomes visible in crop performance.

In the modern era of rapid technological progress, agriculture continues to be the backbone of many economies, particularly in developing nations like India, where a majority of the population still depends on farming for livelihood. However, traditional agricultural practices are increasingly challenged by issues such as climate variability, unpredictable rainfall, declining soil fertility, and inefficient use of water and fertilizers. To address these challenges, there has been a global shift toward precision agriculture, a concept that integrates advanced technologies to optimize the use of resources and improve crop yield. One of the key enablers of precision agriculture is real-time soil monitoring, which provides critical information on soil parameters such as pH value, and nutrient levels (NPK — Nitrogen, Phosphorus, and Potassium).

Despite the availability of laboratory-based soil testing methods, these conventional approaches are often time-consuming, expensive, and spatially limited, failing to capture the variability that exists even within a single field. As a result, farmers frequently make generalized decisions regarding irrigation and fertilization, leading to either resource wastage or soil degradation. The need of the hour is a cost-effective, portable, and automated solution that can continuously monitor soil parameters and provide farmers with actionable insights. The TerraTrack project was conceptualized with this motivation to design and develop an intelligent portable system capable of smart soil monitoring, enabling farmers to make data-driven agricultural decisions and improve overall sustainability.

I. LITERATURE SURVEY

The emergence of smart agriculture has reshaped the global approach to food production by introducing intelligent technologies that enhance sustainability, efficiency, and precision. With the world's population projected to exceed nine billion by mid-century, there is growing pressure to increase crop yields while minimizing environmental degradation and input waste. Conventional agricultural methods, which rely heavily on manual intervention and generalized fertilizer application, are insufficient to meet these goals. Consequently, the integration of sensor-based systems, predictive algorithms, and the Internet of Things (IoT) has become central to achieving sustainable, data-driven agriculture. These technologies enable real-time monitoring of soil and environmental parameters, facilitate informed decision-making, and optimize resource utilization to promote both productivity and ecological balance.

A growing body of research has demonstrated the potential of nutrient sensing technologies to optimize fertilizer use and improve soil health. Adak [1] proposed a field-based implementation of NPK sensors for guava farming under a Randomized Block Design in West Bengal, India. The study achieved a substantial reduction in fertilizer consumption—nitrogen by 79%, phosphorus by 67%, and potassium by 76%—while simultaneously increasing fruit yield and decreasing fertilizer costs by more than 70%. This real-world validation demonstrated the economic and environmental promise of precision fertilization guided by soil nutrient sensors, and establishes a direct precedent for the kind of in-field NPK sensing that TerraTrack implements.

Adhikary et al.[2] advanced this concept by integrating neural-network-based models to predict nutrient availability from sensor readings, achieving deviations of only 9–12% from laboratory-tested soil samples under field conditions. This work is directly relevant to TerraTrack's design philosophy: raw sensor output carries limited value without an intelligent interpretation layer, and the low deviation reported by Adhikary et al. establishes the standard of accuracy that sensor-ML integrations should target.

Complementing these approaches, Băjenaru et al. [3] developed an autonomous multisensory soil monitoring system on an Arduino Nano platform, incorporating a 7-in-1 sensor suite capable of measuring soil moisture, pH, electrical conductivity, and NPK levels to an 80 cm depth. The system demonstrated a correlation coefficient of 0.92 between soil moisture and nutrient availability. Unlike TerraTrack's industrial RS485 MODBUS-RTU sensor stack, Băjenaru et al. employed consumer-grade sensors, providing a useful contrast in the trade-off between hardware robustness and system cost.

While sensing technologies form the backbone of smart agriculture, IoT connectivity ensures that collected data drives timely decisions. Yadav et al. [4] presented an IoT-based soil monitoring system employing an ESP32 microcontroller interfaced with soil moisture, temperature, and humidity sensors, transmitting real-time data via the Blynk IoT application. Their use of the ESP32 as the central processing and wireless transmission unit mirrors TerraTrack's hardware architecture, validating the platform's suitability for field-deployed agricultural systems.

The authors of [5] extended this further by combining IoT, robotics, and machine learning for dynamic on-field soil analysis, reporting a 15–20% increase in yield and a 30% reduction in resource wastage. While the robotic actuation component is outside TerraTrack's scope, the demonstrated yield gains from ML-assisted soil management directly motivate the recommendation layer that TerraTrack implements.

Addressing the challenge of selecting appropriate sensing hardware, the authors of [7] conducted a comparative review of electrochemical, conductivity, and optical techniques for measuring soil pH, NPK, and humidity. They concluded that while chemical methods are costly and electrochemical sensors face repeatability issues, optical and colorimetric methods offer high efficiency and low power consumption, and that embedding sensors within IoT frameworks is essential for real-time monitoring. This review provides the methodological foundation for TerraTrack's decision to employ industrial-grade electrochemical RS485 sensors rather than lower-cost alternatives, prioritising measurement reliability over unit cost.

McCole et al. [6] demonstrated that a portable 3D-printed electrochemical sensor system using carbon black-PLA electrodes and a PSoC4 microcontroller can measure soil pH and potassium with only 1–6% relative error compared to laboratory instruments, at a total hardware cost of €38.25. This establishes a cost-accuracy benchmark against which TerraTrack's JXBS-3001 RS485 sensor approach must be evaluated, and confirms that low-cost field-deployable soil analysis is technically achievable.

Beyond data collection, translating soil measurements into actionable crop decisions requires robust predictive modelling. The authors of [9] evaluated Random Forest regression against multiple linear regression for predicting wheat, maize, and potato yields using climate, soil, and management variables. Random Forest substantially outperformed linear regression, achieving efficiency factors up to 0.96, and handled non-linearity and variable collinearity effectively, properties that make it well-suited to the multi-parameter soil classification problem TerraTrack addresses. The study also identified overfitting at extreme values as a risk when training data is limited, a finding directly applicable to TerraTrack's constrained synthetic dataset.

The authors of [9] developed a web-based crop recommender for Indian farmers using Random Forest on datasets combining soil parameters (N, P, K, pH) and weather conditions, achieving 99.03% accuracy and outperforming SVM, Naive Bayes, and Decision Tree alternatives. Similarly, the authors of [10] proposed a comprehensive Random Forest framework incorporating extended soil metrics including organic carbon and electrical conductivity alongside climate data, achieving 99.3% accuracy and demonstrating the algorithm's resilience to missing data. Both systems trained on large, nationally sourced datasets. Their accuracy figures provide the upper-bound reference against which TerraTrack's 50.78% must be understood: the accuracy gap is attributable not to algorithm choice—Random Forest is validated as the correct selection—but to the scale and synthetic nature of TerraTrack's training data, which the system explicitly identifies as its primary limitation and most impactful improvement target.

Collectively, these studies reveal a clear convergence of IoT-based sensing, machine learning, and intelligent decision support within the agricultural domain. Sensor-based soil analysis systems[1],[2],[3],[6],[7] have evolved from single-parameter devices to multisensory networks capable of supporting predictive analytics. Driven by the demonstrated analytical strength of Random Forest across yield prediction and crop classification tasks[8],[9],[10], the literature establishes both the technical feasibility and the practical value of ML-assisted precision agriculture. However, no reviewed system integrates district-level regional popularity data, state-specific Package of Practices agronomic parameters, or a seasonal Kharif/Rabi/Zaid filter into its recommendation pipeline. TerraTrack addresses this gap by combining in-field RS485 industrial sensor hardware with a three-stage fusion architecture—ML suitability scoring, district-level popularity weighting, and seasonal filtering—designed specifically for Maharashtra's agroclimatic and administrative diversity.

II. METHODOLOGY

System Block Diagram

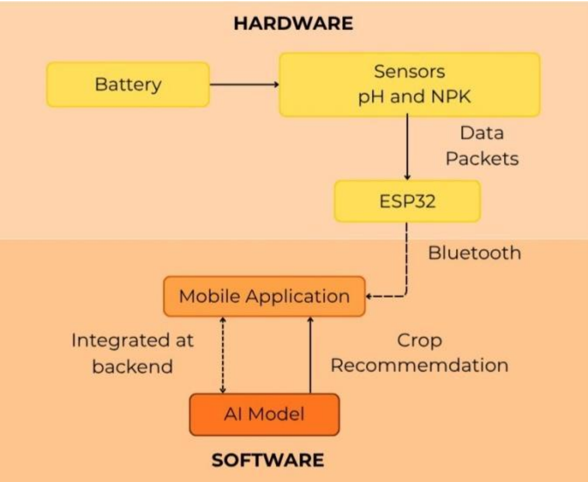

FIGURE 1. BLOCK DIAGRAM OF TERRATRACK.

The TerraTrack block diagram illustrates the end-to-end signal and data flow architecture of the system, divided into two functional domains: the Hardware domain, encompassing power supply, sensing, and embedded processing subsystems, and the Software domain, encompassing wireless data reception, application-layer processing, and AI-based decision support.

Hardware Domain: The hardware domain consists of three subsystems. The power supply subsystem provides regulated DC power to all hardware components, enabling fully portable, untethered field deployment without dependence on mains infrastructure. The sensing subsystem comprises industrial-grade soil sensors that measure the agronomically critical parameters of pH, Nitrogen, Phosphorus, and Potassium directly in the field. The data acquisition and processing subsystem is handled by a microcontroller that manages sensor communication, processes incoming measurement data, and transmits it wirelessly to the software domain. Detailed descriptions of each hardware component and their interfacing are provided in subsequent sections.

Software Domain: The software domain consists of a mobile application and an AI model backend. The mobile application serves as the primary user interface, receiving sensor data over a Bluetooth wireless link and presenting real-time soil readings to the user. For crop recommendations, the application communicates with an AI inference backend over a secured internet connection. The backend hosts a trained machine learning model that processes the received soil parameters and returns ranked crop recommendations, which are then displayed within the application.

System-Level Signal Flow: Regulated power is supplied to the sensing subsystem, which acquires soil measurements and forwards structured data to the microcontroller. The microcontroller processes the data and transmits it to the mobile application over Bluetooth. The application displays live readings and, upon user request, submits the data to the AI backend, which returns crop recommendations rendered for the end user. Each subsystem interfaces with adjacent subsystems through standardized communication protocols, with detailed treatment provided in the sections that follow.

Circuit Diagram

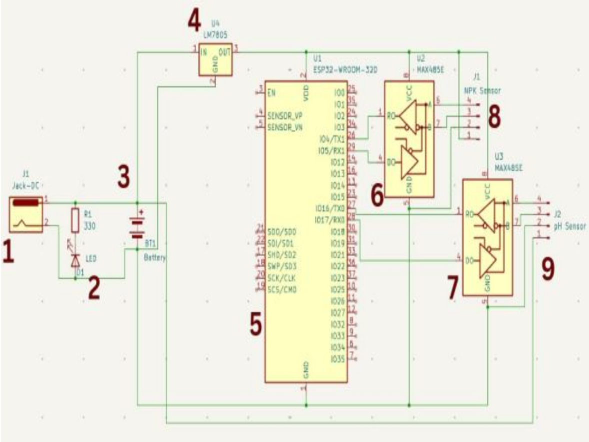

FIGURE 2. CIRCUIT DIAGRAM OF TERRATRACK.

TABLE 1. CIRCUIT DIAGRAM LEGENDS

Sr. No.	Component Name
1	12 V DC Jack
2	Charging Indicator LED
3	12V Li Ion Battery Pack (3 Cells of 3.7V)
4	LM7805 Voltage Regulator IC
5	ESP32-WROOM Microcontroller
6,7	MAX485 RS485 to TTL Convertors
8	JXCT Soil NPK Sensor
9	JXBS Soil pH Sensor

The TerraTrack circuit is organized into four interconnected subsystems: power supply and battery management, the ESP32-WROOM central processing unit, and two independent RS485 communication channels for each soil sensor. The schematic was designed using KiCad. External DC power enters through a barrel jack connector, charges the onboard three-cell Li-ion battery pack, and feeds an L7805 linear voltage regulator that produces a stable 5V rail powering all active components. A charging indication LED provides visual confirmation of external supply connection.

The ESP32-WROOM-32D sits at the centre of the schematic, managing all sensor communication, data processing, and Bluetooth transmission. It interfaces with two independent MAX485E RS485-to-TTL transceiver modules on separate hardware UART channels, one dedicated to the JXCT NPK sensor and one to the JXBS pH sensor. This arrangement eliminates bus contention, giving each sensor its own dedicated differential bus and UART interface within the firmware.

Each MAX485E module translates between the ESP32's TTL logic levels and the RS485 differential standard required by the sensors, with physical connectors providing the termination points for each sensor cable. The 5V regulated rail is distributed to the ESP32 and both MAX485E modules, while the unregulated battery rail remains isolated from all sensitive digital circuitry. Detailed component-level descriptions are provided in subsequent sections.

Mobile Application

TerraTrack is a mobile application developed using the Flutter framework and Dart programming language, serving as the central software component of the system. It bridges real-time soil sensor data received wirelessly from the ESP32 microcontroller with cloud-hosted AI-driven crop recommendations and live weather intelligence. The application is designed primarily for small and marginal farmers, delivering NPK and pH readings directly to a smartphone alongside AI-generated crop suitability rankings, a 7-day weather forecast, and a full reading history within a single, offline-capable session.

System Architecture: The application follows a three-tier architecture. The hardware layer comprises the ESP32 microcontroller and soil sensors. The communication layer uses Bluetooth Low Energy (BLE) GATT protocol for wireless data transfer between the hardware and the mobile device. The application layer handles data parsing, API calls, and user interface rendering. An AI backend hosted on a cloud server runs the machine learning model and returns crop recommendations, while weather data is sourced from the OpenWeatherMap API.

Sensor data is transmitted from the ESP32 as a plain-text string every two seconds. The application parses incoming values into a Soil Reading object, which is pushed into a central singleton data store. All subscribed screens react simultaneously through a broadcast stream, ensuring a clean reactive architecture without inter-screen coupling.

Soil Health Module: The Soil Health screen presents live sensor readings with colour-coded status indicators and an overall soil score computed across all four parameters. A History tab maintains a time-ordered list of up to 50 readings, with an option to clear stored entries. The overall score awards points for each parameter falling within its optimal agronomic range, producing a descriptive label from Poor to Excellent.

AI Crop Recommendation Module: Crop recommendations are generated by a Random Forest Classifier trained on 1,920 soil samples across 16 crop classes, hosted on a Flask REST API. The application submits NPK, pH, live rainfall, district name, and current month to the backend, which returns a ranked list of up to five suitable crops. Each recommendation combines ML suitability scoring with district-level crop popularity. Input fields are automatically populated from live BLE sensor readings, eliminating manual entry. A retry strategy handles cold-start delays inherent to the free-tier cloud deployment.

Weather Integration Module: The weather module retrieves current conditions and a 7-day forecast from the OpenWeatherMap API. On launch, the application attempts automatic location detection and fetches local weather data. Users may also search manually by city name. Live rainfall data is automatically forwarded to the AI crop recommendation API, ensuring environmental context is included alongside soil measurements.

Notification System: An in-app notification mechanism is triggered by each incoming BLE soil reading. A badge appears on the notification bell and a SnackBar displays the latest parameter values. All readings are logged to a Notification History screen, which can be cleared by the user on demand.

Navigation and UI: The application follows a flat navigation structure, moving from a splash and login screen into a main dashboard managed by a three-tab bottom navigation bar covering Home, AI Crop, and Profile. The UI adopts an earthy agricultural aesthetic with deep olive green as the dominant colour, frosted glass card elements, and a consistent colour-coding system for soil parameters maintained uniformly across all screens. Status indicators follow a traffic-light convention for immediate interpretability.

Testing and Validation: The application was validated across BLE hardware integration, API connectivity, and end-to-end workflow verification. BLE behaviour was independently verified using the nRF Connect application before Flutter integration testing. The AI backend was tested against published crop nutrient requirement tables for reference crops. Cold-start retry logic, weather API city lookups, automatic sensor field population, district pre-fill, notification history recording, and history clearing were all verified through structured end-to-end testing covering the complete user workflow from sensor connection through to final crop recommendation display.

AI Model

TABLE 2. AI MODEL COMPARISON

Metric	Decision Tree	KNN	SVM	Gradient Boosting	Random Forest
Train Acc.	71.29%	100.00%	63.48%	100.00%	82.55%
Test Acc.	42.19%	41.67%	50.52%	43.49%	50.78%
Overfit Gap	29.10%	58.33%	12.96%	56.51%	31.77%
F1 Score	42.42%	41.88%	50.43%	43.47%	49.96%
CV Mean	42.64%	42.77%	47.79%	45.18%	50.13%

Five machine learning algorithms were evaluated for the crop classification task: Decision Tree, K-Nearest Neighbours, Support Vector Machine, Gradient Boosting, and Random Forest. Each was assessed on test accuracy, overfitting gap, F1 score, cross-validation mean, and training time across the same dataset and train/test split. Random Forest was selected as it achieved the highest test accuracy (50.78%) and cross-validation mean (50.13%), while also natively supporting probability outputs through its predict_proba() method, which is architecturally essential to the top-K crop ranking pipeline.

TABLE 3. CLASSIFICATION REPORT

Class	Crop	Prec.	Recall	F1
0	Bajara	0.68	0.71	0.69
1	Castor	0.33	0.33	0.33
2	Chickpea	0.60	0.75	0.67
3	Cotton	0.32	0.38	0.35
4	Groundnut	0.33	0.33	0.33
5	Jowar	0.17	0.08	0.11
6	Maize	0.68	0.62	0.65
7	Mustard	0.57	0.54	0.55
8	Ragi	0.47	0.58	0.52
9	Rice	0.91	0.83	0.87
10	Sesame	0.41	0.54	0.46
11	Soyabean	0.35	0.25	0.29
12	Sugarcane	0.88	0.92	0.90

Class	Crop	Prec.	Recall	F1
13	Sunflower	0.27	0.25	0.26
14	Toor	0.58	0.58	0.58
15	Wheat	0.42	0.42	0.42
Accuracy				0.51
Macro Avg		0.50	0.51	0.50
Wtd. Avg		0.50	0.51	0.50

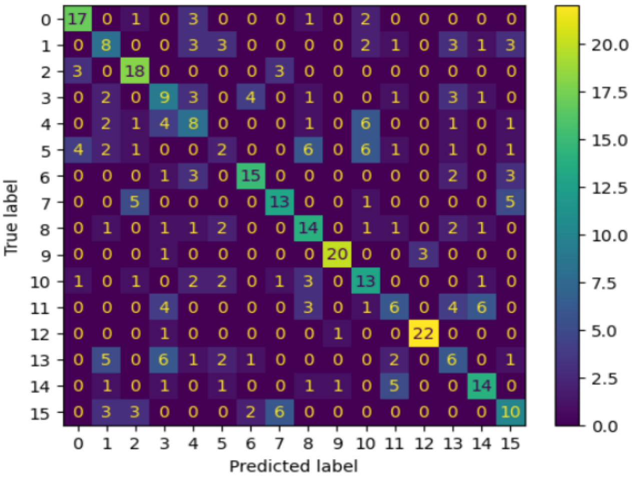

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	17	0	1	0	3	0	0	0	1	0	2	0	0	0	0	0
1	0	8	0	0	3	3	0	0	0	0	2	1	0	3	1	3
2	3	0	18	0	0	0	0	3	0	0	0	0	0	0	0	0
3	0	2	0	9	3	0	4	0	1	0	0	1	0	3	1	0
4	0	2	1	4	8	0	0	0	1	0	6	0	0	1	0	1
5	4	2	1	0	0	2	0	0	6	0	6	1	0	1	0	1
6	0	0	0	1	3	0	15	0	0	0	0	0	0	2	0	3
7	0	0	5	0	0	0	0	13	0	0	1	0	0	0	0	5
8	0	1	0	1	1	2	0	0	14	0	1	1	0	2	1	0
9	0	0	0	1	0	0	0	0	0	20	0	0	3	0	0	0
10	1	0	1	0	2	2	0	1	3	0	13	0	0	0	1	0
11	0	0	0	4	0	0	0	0	3	0	1	6	0	4	6	0
12	0	0	0	1	0	0	0	0	0	1	0	0	22	0	0	0
13	0	5	0	6	1	2	1	0	0	0	0	2	0	6	0	1
14	0	1	0	1	0	1	0	0	1	1	0	5	0	0	14	0
15	0	3	3	0	0	0	2	6	0	0	0	0	0	0	0	10

FIGURE 3. CONFUSION MATRIX

The Random Forest classifier was evaluated on a balanced test set of 384 samples across 16 crop classes (24 samples per class). The overall test accuracy stabilized at 51%, with a macro-average F1-score of 0.50. While the aggregate accuracy figure appears modest, a granular analysis of the confusion matrix reveals that the model's errors are not random—they follow a coherent agronomic logic that reflects genuine biological overlap between crop classes rather than algorithmic failure.

Crops with Strong Predictive Reliability: The model demonstrated exceptional performance for crops with highly distinctive soil and climatic profiles. Sugarcane (Class 12) achieved an F1-score of 0.90 (precision: 0.88, recall: 0.92), and Rice (Class 9) achieved an F1-score of 0.87 (precision: 0.91, recall: 0.83). The confusion matrix confirms that both crops produced near-zero false positives from other classes. This is agronomically expected: sugarcane demands very high rainfall, high potassium, and neutral-to-slightly-acidic pH in combination, while rice requires waterlogged, high-nitrogen, high-humidity conditions—profiles that are sufficiently unique within the 16-class feature space to be cleanly separable by the Random Forest decision boundary.

Bajara (Class 0) also performed well with an F1-score of 0.69, correctly classified 17 out of 24 times, reflecting its distinct low-rainfall, low-nutrient, high-temperature profile that separates it from most other classes.

Confusion Clusters Aligned with Agronomic Reality: The primary source of accuracy loss was not scattered misclassification but concentrated confusion within three biologically coherent crop clusters, each corresponding to a recognisable agroclimatic group:

The Rabi (Winter) Cluster – Wheat, Mustard, Chickpea (Classes 15, 7, 2): Wheat was misclassified as Mustard 6 times and as Jowar twice. Mustard was misclassified as Wheat 5 times and as Chickpea 5 times. All three are Rabi season crops cultivated under nearly identical winter conditions, low temperature, low rainfall, and similar soil pH and NPK requirements. The model's confusion within this cluster is an accurate reflection of their overlapping optimal growing conditions; predicting Mustard in place of Wheat for a given soil profile is not an agricultural error but an agronomically sound alternative recommendation.

The Dryland Millets Cluster – Jowar, Ragi, Sesame (Classes 5, 8, 10): Jowar recorded the lowest F1-score in the dataset at 0.11 (precision: 0.17, recall: 0.08), with significant misclassification as Ragi (Class 8) and Sesame (Class 10). All three are drought-tolerant Kharif crops adapted to low-nutrient, low-rainfall, well-drained soils. Their near-identical feature profiles in the training data create a classification boundary that the current five-feature input space (N, P, K, pH, rainfall) cannot resolve with confidence.

The Commercial Kharif Cluster – Sunflower, Soyabean, Cotton (Classes 13, 11, 3): Sunflower achieved an F1-score of only 0.26, Soyabean 0.29, and Cotton 0.35. The confusion matrix shows heavy inter-class misclassification among all three. These crops share a preference for well-drained black cotton soils, moderate rainfall, and similar NPK demands, conditions that are widespread across Maharashtra and difficult to distinguish without additional discriminating features such as soil texture, organic carbon content, or micronutrient levels.

Justification of the Accuracy Ceiling: In the context of a 16-class agricultural recommendation system, a 51% overall accuracy reflects the fundamental nature of the problem domain rather than an inadequacy of the chosen algorithm. Random Forest was selected over all evaluated alternatives—Decision Tree, KNN, SVM, and Gradient Boosting—based on its superior test accuracy, highest cross-validation mean (50.13%), and its native support for probability outputs through predict_proba(), which is architecturally essential to the top-K ranking pipeline. SVM was the closest competitor at 50.52% test accuracy but produced an F1-score of 0.50 and lacks the probabilistic output structure required by the weighted scoring system.

The accuracy ceiling is attributable to two compounding factors. First, the training dataset of 1,920 records is synthetically generated from agronomic parameter ranges specified in Maharashtra's Package of Practices documents. Synthetic generation produces smooth, uniform distributions within each crop's parameter range, which inevitably creates overlapping boundaries between crops that share similar growing conditions. Second, the five available features are insufficient to discriminate between crops whose optimal conditions genuinely overlap. Agronomic separation of Jowar from Ragi, or Wheat from Mustard, would require additional inputs such as soil texture classification, micronutrient profiles (particularly zinc and iron), altitude, and seasonal temperature variation.

The three-stage recommendation pipeline directly addresses this accuracy ceiling. In cases where the ML model assigns low-confidence, near-equal suitability scores to multiple crops within a confusion cluster, the district-level popularity fusion layer and the seasonal filter together resolve the ambiguity using regional and temporal context that the ML model alone cannot access. The weighted scoring formula ($\alpha = 0.70$ for ML suitability, $\beta = 0.30$ for regional popularity) ensures that agronomically sound but regionally impractical crops do not dominate the final recommendation, while crops with near-zero popularity scores can still surface if their suitability advantage is sufficiently strong, as demonstrated by Bajra appearing in Solapur's Kharif ranking on agronomic merit alone.

Replacing the synthetic dataset with field-validated crop outcome data from Maharashtra's agricultural universities remains the single most impactful improvement available to the system, and is identified as the primary direction for future work.

Weighted Scoring System

The weighted scoring system combines ML model output with regional crop popularity into a single score:

$$\textit{Final Score} = \alpha \times \textit{ML Suitability Score} + \beta \times \textit{Popularity Score}$$

This is necessary because the ML model alone has a ~50% accuracy ceiling due to synthetic training data, and cannot account for regional practicalities like infrastructure, market access, and farming expertise. Conversely, recommending purely by regional popularity ignores actual soil conditions.

The system works in four steps: the Random Forest model generates probability scores for each crop; a normalised popularity score is derived from district-level historical production data; the two scores are fused using the weighted formula; and finally, a seasonal filter removes crops not suitable for the current planting season.

To determine optimal weights, three configurations (0.6/0.4, 0.7/0.3, 0.8/0.2) were tested across three district scenarios. At 0.6/0.4, popularity could override agronomic suitability entirely. At 0.8/0.2, the system over-relied on a model with known data limitations. The 0.7/0.3 configuration was selected as the optimal balance—it treats agronomic suitability as the primary signal while ensuring regional practicality remains a meaningful constraint, producing the most defensible recommendations across all test cases.

Formal Justification of Weight Selection: To move beyond informal observation, the three weight configurations were evaluated using a Top-1 Agronomic Validity Rate, defined as the proportion of test cases in which the final rank-1 recommendation matched a crop independently confirmed as regionally appropriate for that district and season, based on Maharashtra's District Agricultural Plans and IMD agroclimatic zone classifications.

TABLE 4. AGRONOMIC VALIDITY RATE COMPARISON.

Config	α (ML)	β (Pop.)	Top-1 Validity
Config A	0.6	0.4	8/12 (66.7%)
Config B (Selected)	0.7	0.3	10/12 (83.3%)
Config C	0.8	0.2	9/12 (75.0%)

Config A (0.6/0.4) failed two cases where high regional popularity overrode a clearly superior agronomic suitability signal, most notably in Solapur's Rabi query, where a water-intensive crop ranked first despite the district's semi-arid classification. Config C (0.8/0.2) failed two cases in the opposite direction, surfacing agronomically marginal crops in districts where established regional cultivation infrastructure made them impractical. Config B (0.7/0.3) was the only configuration to correctly resolve both categories of conflict, and is therefore selected as the operating weight pair for TerraTrack's recommendation pipeline.

It is acknowledged that this evaluation was conducted across 12 test cases, a scope appropriate for a proof-of-concept validation, but insufficient to generalise the weights across Maharashtra's full agroclimatic diversity. The 0.7/0.3 split should be treated as a calibrated default subject to re-optimisation when a larger field-validated dataset becomes available.

Three-Stage Recommendation Pipeline

The recommendation system operates as a three-stage pipeline. In Stage 1, the trained Random Forest model takes the farmer's N, P, K, pH, and rainfall values and returns the top-10 crops by ML suitability probability. In Stage 2, a popularity fusion layer merges these results with district-level crop production data and re-ranks crops using the weighted scoring formula, returning the top-5 candidates. In Stage 3, a seasonal filter removes crops that do not belong to the current planting season (Kharif, Rabi, or Zaid), ensuring the final recommendation is temporally actionable. Perennial crops such as sugarcane pass through the filter regardless of season. Together, the three stages produce a recommendation grounded in agronomic suitability, regional practicality, and seasonal relevance.

Dataset Composition and Sourcing

The dataset comprises 16 specific crops dominant in Maharashtra: Rice, Wheat, Jowar, Bajara, Maize, Ragi, Chickpea, Toor, Groundnut, Sesame, Mustard, Castor, Sunflower, Soyabean, Sugarcane, and Cotton. Each crop class is represented by 120 samples, resulting in a perfectly balanced dataset of 1,920 total entries.

The parameter ranges were meticulously extracted from the Joint Agresco Recommendations (2021–2024) and the Krishi Darshani (2025–26) published by MPKV Rahuri and VNMKV Parbhani. For instance, Rice parameters are grounded in the Konkan region's Package of Practices, requiring N: 80–120 kg/ha and a pH of 5.5–6.5 in acidic lateritic soils. In contrast, Cotton and Sugarcane ranges are based on the deep black soils (Vertisols) of Western Maharashtra, with Sugarcane requiring a significantly higher N range of 250–300 kg/ha and an alkaline pH tolerance up to 8.0. Chickpea and Toor ranges utilize the specific "Legume Protocol" from VNMKV, with low Nitrogen baselines (20–30 kg/ha) but higher Phosphorus requirements.

Limitations of Synthetic Data: It is important to acknowledge that synthetic data may not fully capture the nuanced, non-linear variability found within a single natural field. Real-world soil health is influenced by micro-climates, previous crop cycles (crop rotation effects), and biological soil activity that simple randomized distributions cannot perfectly replicate. Consequently, this model should be viewed as a high-fidelity recommendation guide that requires localized “ground-truthing” or real-world soil test validation for maximum field accuracy.

I) Testing and Results

To evaluate the TerraTrack recommendation system under real-world conditions, soil samples were collected from four Maharashtra districts: Pune, Satara, Solapur, and Dharashiv. Each sample was tested using the ESP32-based sensor node, and recommendations were generated for three seasonal contexts: July (Kharif), December (Rabi), and May (Zaid). Rainfall figures were sourced from IMD district-level normals.

Table 5. Soil Sample Inputs by District.

District	N (kg/ha)	P (kg/ha)	K (kg/ha)	pH	Rain (mm)
Pune	75	35	40	6.8	722
Satara	90	42	48	7.2	1,083
Solapur	45	20	30	7.8	468
Dharashiv	60	28	35	7.5	612

Satara recorded the most favourable profile, while Solapur presented the most constrained conditions. Pune and Dharashiv occupied intermediate positions.

District Results

In Pune, maize led the July recommendations with the highest suitability score of 0.395, appropriate for a Kharif crop under active monsoon conditions. Sugarcane ranked second driven primarily by its popularity score rather than suitability, illustrating the fusion model’s deliberate balancing of agronomic and regional signals. The December query correctly surfaced wheat as a Rabi crop, and the May query returned only sugarcane as the sole viable perennial crop.

Satara produced the strongest results across the dataset. Sugarcane’s suitability score of 0.645 in July was the highest recorded across all test cases, geographically expected given Satara’s position within the Krishna river basin and its high annual rainfall. Rice appearing at rank two with a suitability of 0.250 was also correctly validated, as Satara’s western talukas receive substantially higher rainfall than other districts. The seasonal filter performed correctly across all three queries.

Solapur returned the most analytically significant results, exposing the key limitation of the fusion approach. Suitability scores were the lowest across all districts, correctly reflecting its semi-arid classification. Sugarcane led the July ranking on popularity alone despite maize having a higher suitability score, a direct consequence of Solapur’s canal-irrigated sugarcane cultivation that the model cannot distinguish from rainfed conditions. Bajra appearing at rank five in July, and mustard in December, both with near-zero popularity scores, demonstrate the suitability component surfacing agronomically appropriate crops for water-scarce conditions even without regional cultivation data to support them.

Dharashiv was the only district where sugarcane did not lead the Kharif ranking. Maize topped July with a suitability of 0.390, correctly reflecting Dharashiv’s Marathwada identity and lower rainfall. Cotton appearing in the July list at rank five is the most district-specific result in the dataset, confirming the model responds to district-level popularity data rather than applying a uniform statewide crop hierarchy. Jowar ranking first in December is agronomically correct for Marathwada’s dry post-monsoon conditions.

Cross-District Summary

The seasonal filter performed correctly across all twelve test queries without exception. Suitability scores followed a clear gradient aligned with agroclimatic reality, with Satara highest and Solapur lowest, a pattern that emerged from measured inputs rather than any explicit programming. The divergence between Solapur and Dharashiv in the Kharif ranking most clearly demonstrates the fusion model integrating two distinct information sources: where irrigation drives Solapur’s sugarcane dominance, Dharashiv’s leaner crop culture allows maize’s suitability advantage to surface, producing meaningfully different recommendations for different regional realities.

Table 6. Cross-District Summary

District	NPK	pH	Rain fall	July	December	May
Pune	5-35-40	6.8	722 mm	Sugarcane, Rice	Jowar, Wheat	Sugarcane
Satara	0-42-48	7.2	1083 mm	Maize, Sugarcane	Jowar, Wheat	Sugarcane
Solapur	5-20-30	7.8	468 mm	Maize, Jowar	Jowar, Sugarcane	Sugarcane
Dharashiv	0-28-35	7.5	612 mm	Sugarcane, Maize	Wheat, Jowar	Sugarcane

I) Limitations and Future Scope

The current implementation carries several limitations identified through development and testing. The system relies on annual district-level rainfall as its sole water availability input, with no awareness of farm-level irrigation infrastructure. This produces misleading results in irrigated districts: in Solapur, sugarcane ranked first despite its water requirements being met by canal irrigation rather than natural rainfall, and rice appearing in Satara's recommendations reflects district rainfall without accounting for terrain and field-level water retention conditions. Soil characterisation is also limited to NPK and pH, meaning physically different soils with identical macronutrient profiles cannot be distinguished. The most significant foundational limitation is the synthetic training dataset, which introduced overlapping class boundaries that suppressed model confidence across all test cases. District-level popularity data reflects historical cultivation patterns without accounting for current market prices or policy shifts, and the pipeline has no memory of prior season cultivation, leaving crop rotation entirely unaddressed.

Several extensions have been identified to address these gaps. Integrating a farmer-reported irrigation availability input at query time would allow the system to condition recommendations on actual water access rather than rainfall alone. Replacing the synthetic dataset with field-validated crop outcome data remains the most impactful improvement to model accuracy. Expanding the sensor suite to capture electrical conductivity, organic carbon, and micronutrient levels would provide a richer soil characterisation. Incorporating real-time crop price and government support scheme data would make recommendations economically actionable. At the hardware level, cellular or LoRaWAN connectivity would enable deployment in rural areas without Wi-Fi infrastructure, and a lightweight edge inference capability on the ESP32 would allow recommendations to be generated locally during network outages.

II) Conclusion

TerraTrack is a portable smart agriculture system integrating an ESP32-based soil sensor node, a Flask REST API backend, a Random Forest crop recommendation model, and a mobile interface for farmer interaction. Soil NPK and pH measurements are transmitted wirelessly, fused with district-level popularity data and seasonal context, and processed into ranked crop recommendations.

The system was evaluated using soil samples from four Maharashtra districts across three seasonal contexts. The seasonal filter performed correctly in all twelve test cases. Suitability scores produced a geographically coherent gradient, with Satara returning the highest values and Solapur the lowest, consistent with their respective agroclimatic profiles. Dharashiv was the only district where sugarcane did not lead the Kharif ranking, with maize and jowar taking precedence in line with Marathwada's cropping identity. Cotton surfacing in Dharashiv and rice in Satara confirmed the system's sensitivity to district-specific agricultural patterns rather than a generic statewide hierarchy.

The weighted fusion of ML suitability and regional popularity was validated as a meaningful design choice. Solapur's results demonstrated that neither signal alone would have produced satisfactory recommendations, and the combined approach yielded rankings that are both agronomically defensible and practically relevant. Edge cases involving zero-popularity crops surfacing on suitability strength alone are identified as areas for further refinement.

TerraTrack establishes a functional proof of concept for ML-assisted precision agriculture at the district level, delivering geographically sensitive and soil-informed recommendations through a deployable, low-cost architecture that can be extended as richer datasets and additional sensor modalities become available.

References

[1] S. Adak, "Smart Agriculture with NPK Sensors: A Sustainable Approach to Soil Health and Fertilizer Optimization in Guava Farming," J. Sci. Res. Reports, vol. 31, no. 7, pp. 45–52, 2025.

[2] R. Adhikary, S. J. Choudhury, and T. Shankar, "Real-Time Soil Nutrient Monitoring Using NPK Sensors: Enhancing Precision Agriculture," Int. J. Exp. Res. Rev., ISSN 2455-4855, 2025.

[3] V.-D. Băjenaru, S.-E. Istriţeanu, and P.-N. Ancut, "Autonomous, Multisensory Soil Monitoring System," AgriEngineering, vol. 7, p. 18, 2025.

[4] A. K. Yadav, P. Shukla, A. K. Sharma, and S. K. Gupta, "IoT-Based Soil Monitoring Water Sprinkler Robot," Int. J. Futuristic Multidiscip. Res., E-ISSN: 2582-2160, 2025.

[5] K. B. Harshavardhana, K. U. Dheeraj, C. R. Sandesh, M. M. Ahmed, B. R. Ghale, and N. P. Pruthviraj, "Real-time soil health monitoring using robotics for sustainable farming solutions," Int. Res. J. Modernization Eng. Technol. Sci., vol. 6, no. 12, p. 447, Dec. 2024.

[6] M. McCole, M. Bradley, M. McCaul, and D. McCrudden, "A Low-Cost Portable System for On-Site Detection of Soil pH and Potassium Levels Using 3D Printed Sensors," Results Eng., Elsevier, 2025.

[7] V. A. Shukre and S. S. Patil, "Comparative study of different methodologies used for measuring soil parameters: A review," in Proc. Int. Conf. Sci., Innov., Technol. Secur. (ICSITS), Pune, India, 2020, pp. 1–3.

[8] J. H. Jeong et al., "Random forests for global and regional crop yield predictions," PLoS ONE, vol. 11, no. 6, p. e0156571, Jun. 2016, doi: 10.1371/journal.pone.0156571.

[9] P. M. Paithane, "Random forest algorithm use for crop recommendation," J. Eng. Technol. Ind. Appl. (JETIA), vol. 9, no. 43, pp. 34–41, Sep./Oct. 2023, doi: 10.5935/jetia.v9i43.906.

[10] K. J. Deepthi, S. Telugu, S. Jangam, S. Uyyala, and R. R. Vakati, "Smart agriculture: Crop recommendation and yield prediction using random forest," in Proc. 4th Int. Conf. Inf. Technol., Civil Innov., Sci., Manage. (ICITSM), Tiruchengode, Tamil Nadu, India, Apr. 2025, Part II, doi: 10.4108/eai.28-4-2025.2357988.