

Multi-Agent Generative AI-Based Digital Twin Framework for Precision Agriculture

Mayuri Mangade¹, Pranjal Mahakal², Dr. Kavita Dhakad³

^{1,2}M.Sc. CA, School of Information and Technology, Indira University, Pune, India

³Associate Professor, School of Information Technology, Indira University, Pune, India

Abstract— Precision agriculture relies on integrating information from various sources and turning the findings into practical choices. Using public agricultural datasets, this research introduces AgroTwin, an academic prototype of a multi-agent digital-twin framework for crop recommendation, soil-moisture interpretation, irrigation advice, weather analysis, and farm decision support. The prototype integrates meteorological records, soil-moisture readings, and crop-recommendation data into a Python-based workflow. Crop labels are predicted by a Random Forest classifier, and crop, soil, irrigation, weather, decision-support, and generative-AI explanation activities are arranged by software agents. The crop classifier's test accuracy on the crop-recommendation dataset was 99.55% in the reported experimental run. With 4,117 recordings, the initial plant-vase soil dataset indicated a mean moisture value of roughly 0.4245. The results are shown in an interactive dashboard that lets users investigate potential modifications to farm conditions. These findings demonstrate that a student research prototype can incorporate several public datasets and modular software agents. They don't show real-world agronomic efficacy or water savings at the field level; those results would need external confirmation and controlled field experiments.

Keywords— Random Forest, Irrigation decision-making help, Crop Recommendation, Digital twin, Precision agriculture, Multi Agent Systems, and Generative AI.

I. INTRODUCTION

Farming is affected by various interacting elements, which include soil condition, needs of crops for nutrients, water availability, and climate. Individual tools can analyze these factors separately, but farm decisions often require their outputs to be considered together. Digital twin is a digital copy of a physical system that uses data to describe its state and explore possible conditions or actions. Agricultural digital-twin approaches have been discussed for crop monitoring, irrigation, farm operations and decision support [1], [2]. Important challenges remain, including interoperability, data availability, model validation, and the relationship between a virtual representation and the physical farm [2].

This paper presents AgroTwin, a prototype that combines public agricultural datasets, a modular agent workflow, and a dashboard.

It is intended to show how crop prediction, soil-moisture interpretation, weather summaries, and decision-oriented explanations can be presented through one software interface. The main contributions are: (1) integrating public agricultural datasets from different sources; (2) organizing crop, soil, irrigation, weather, decision-support, and explanation tasks into separate modules; (3) evaluating a Random Forest model for crop recommendation; and (4) developing a dashboard for displaying outputs and exploring hypothetical input changes.

II. RELATED WORK

Previous works have considered digital twins for representing agricultural systems and connecting data, models, and decision-making. Verdouw et al. discuss the possibility of applying digital twins in agriculture and the orchestrating difficulties that come from the integration of different system components [1]. The following review explores applications in the different stages of the agricultural life-cycle and emphasizes the difficulty of managing heterogeneous data and concurrent operations [2]. This justifies the adoption of modular systems that explicitly define data sources and decision-making.

Machine learning techniques are applied to predict crop recommendations based on nutrient levels in soil and other environmental characteristics. Algorithms like Random Forest can capture the nonlinear nature of relationships in tabular data. Yet, the performance of the algorithm depends on the dataset, labels, data split, and the way performance is evaluated. Thus, achieving high performance on a public dataset alone does not necessarily guarantee good performance on the farm or regionally.

Contrary to a demonstration which is purely based on one model, the AgroTwin framework takes crop prediction into a larger framework of the software workflow. At present, this framework is still at an academic level; it is not a deployed system that synchronizes in real time both the physical and virtual farms. The agent architecture and the dashboard are the foundations for further integration and validation.



International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 10, October 2026)

III. SYSTEM ARCHITECTURE

The agro twins architecture consists of five layers. (1) Data layer: public crops, soil moisture, and weather data is loaded using CSV files from the project's datasets folder. (2) Preparation layer: dataset-specific preparation process makes the fields ready to be used in the future and eliminates duplicated weather data if needed. (3) Prediction and analysis layer: Random Forest classifier predicts the crop labels based on nutrients and environment features, and the soil and weather routines analyze observations. (4) Multi-agent orchestration layer: crop, soil, irrigation, weather, decision support, and generative AI-related agents generate specific output. (5) Presentation layer: Streamlit dashboard offers overview, virtual digital twin, agents' status, scenario simulation, model laboratory, resources planning, advisory, analytics, and reports view.

Workflow has been designed to separate tasks rather than treating the whole system as one prediction. Crop component provides crop label output, while soil component summarizes the soil moisture sensors' values; Irrigation component uses defined soil moisture threshold value; Weather component summarizes the available weather data, Decision support component combines all this information together, while Explanation component generates readable output. Prototype has some parts that are based on publicly available historical data and rule-based thresholds. System is not connected to a real farm or to a real irrigation system.

IV. DATASETS AND PREPROCESSING

The prototype used three datasets. The crop recommendation dataset had 2200 records and eight features N, P, K, temperature, humidity, pH, rainfall, and crop name. There were no missing or duplicated records reported at the initial check. The first plant vase soil dataset had 4,117 records and 12 features including timestamp fields, five moisture features named moisture0–moisture4, and an irrigation status feature which is named irrigation in the source dataset. Another plant vase dataset had 10,289 records and the same structure. The weather dataset initially had 533,312 records, but after eliminating duplicates, the record size became 517,048 records; that is a decrease of 16,264 records.

Pre-processing steps depend on datasets. White space in crop variables is stripped, irrigation data for soil is made boolean when possible, and duplicates in the weather dataset are eliminated before analysis. To ensure reproducibility, dimension of datasets, variable names, number of missing values, duplicates, and descriptive statistics are included. Original spellings of source columns are kept in the documentation.

These public datasets operate under different circumstances and should not be considered as synchronized records from a common farm. When using Agro Twin, these datasets are unified at the application level for purposes of demonstration. This unification does not produce a synchronized sensor history of one particular field.

V. MACHINE-LEARNING METHODOLOGY

This is a multiclass supervised crop recommendation problem. Let $x = [N, P, K, \text{temp}, \text{humidity}, \text{pH}, \text{rainfall}]$ is the input feature vector and y represent the crop label. A Random Forest classifier aggregates predictions of a number of decision tree models and assigns the class supported by the aggregate vote. The existing implementation evaluates the model on a held-out test set. Test accuracy is calculated as the ratio of the number of correctly classified test examples to the total number of test examples.

Test accuracy alone does not provide a complete assessment of a classifier. Before publication, the training/test split, random seed, stratification setting, hyper parameters, class-wise precision and Recall, F1-score, confusion matrix and cross-validation results should be confirmed from the executed model output. Feature-importance values can describe how the fitted model uses its inputs, but they do not establish causal effects on crop yield.

The virtual-twin interface accepts environmental inputs that users can change and displays a model prediction or decision summary. In a what-if scenario, one or more inputs are changed and the output is recalculated. These scenarios are exploratory. In the absence of either process models or field data, they cannot be regarded as yield, water savings, or crop response predictions.

VI. EXPERIMENTAL RESULTS

The results of the existing project run suggest that the accuracy of the test using Random Forest algorithm for the crop classification is 99.55% (0.9955). Given the input values $N=90$, $P=42$, $K=43$, $\text{temp} \approx 20.88^\circ\text{C}$, $\text{humidity} \approx 82.00\%$, $\text{pH} \approx 6.50$, $\text{rainfall} \approx 202.94$, the output is rice. According to the dashboard, rice (90%) and jute (10%) are the most prominent options with the rest of the presented option being equal to 0%. Such numbers can be referred to as the probabilities if the implementation proves it to be true; otherwise, it will be model scores.

In regard to the first plant-vase soil data set, the project indicates 4,117 entries with a mean moisture value of 0.4245.

The dashboard utilizes the irrigation threshold of 0.55; according to the defined criteria, moisture value lower than the threshold is identified as “Irrigation required.” It is a prototype threshold and not an absolute agricultural one. In addition, the dataset shows that mean weather temperature is 12.63 °C and mean wind speed is 12.04. It is worth mentioning that the provided figures are just the dataset statistics, and they do not represent the actual weather condition.

DATASET CHARACTERISTICS SUMMARY, MODEL PERFORMANCE, AND IRRIGATION DECISION RESULTS

Component	Reported result	Interpretation
Crop dataset	2,200 × 8	No missing values or duplicate rows reported
Random Forest	99.55% test accuracy	Verify data split and full evaluation metrics
Plant-vase soil dataset 1	4,117 records	Mean moisture ≈ 0.4245
Weather dataset	533,312 → 517,048	16,264 duplicate records removed
Irrigation rule	Threshold = 0.55	Demonstration rule; requires crop- and soil-specific calibration

VII. DISCUSSION

From the above output, it can be clearly seen that the crop classifications, soil analysis, weather summary, and irrigation rule can be combined into one single application. However, the above accuracy may not necessarily reflect the overall accuracy of the algorithm since it depends on the data chosen. The result cannot be extrapolated to other geographical areas, seasons, types of crops, or farms without further external validation. In addition, the average soil moisture content and irrigation status based on threshold do not indicate reduction in water usage.

The key contribution of AgroTwin from an engineering standpoint comes from its agent responsibility segregation and presentation of different data perspectives via a single dashboard. In this paper, “digital twin” stands for a data-driven virtual replica that allows what-if analysis. A fully functional digital twin needs to have a well-defined physical object, proper synchronization in terms of time, consistent data flow, state estimation tuning, and validation based on physical measurements. Any explanations derived from generative-AI should be tied to model outputs and not provide agronomic recommendations.

VIII. LIMITATIONS AND FUTURE WORKS

There are several limitations to this study. First, the publicly available dataset was not generated via synchronized observations from the same farm.

Second, the classification result should be reproducible using a proper documentation of the data split, random seed, parameters for the machine learning model used, and the evaluation metrics. Third, the irrigation threshold is a demonstrative example and has not been verified for different types of soil, crops, growth phases, and climates. Fourth, there is no report of any field trial or irrigation control experiment. Fifth, the multi-agent approach has not been benchmarked through a user study.

Future work will focus on: (1) publishing a reproducible training and evaluation pipeline; (2) comparing Random Forest with suitable baseline classifiers; (3) reporting precision, recall, F1-score, confusion matrix, and cross-validation results; (4) calibrating moisture thresholds using evidence specific to crops and soils; (5) integrating real-time sensor or weather feeds with timestamps and source information; (6) evaluating agent outputs against expert-reviewed cases; and (7) conducting controlled field trials before making claims about water use or crop productivity.

IX. CONCLUSION

This paper presented AgroTwin, a multi-agent digital-twin prototype for public-facing precision agriculture crop, soil-moisture, and weather datasets. The model incorporates Random Forest crop recommendations, environmental factors summarized at the dataset level, irrigation suggestions based on thresholds, agents coordination, and a dashboard interface. In the current execution, 99.55% test accuracy was achieved for the crop prediction task, while the average moisture level was recorded at about 0.4245 for the first plant-vase dataset. Reproducible evaluation, appropriate sensor integration, agronomic calibration, and field validation are still needed before the framework can be relied on for operational farm decisions.

REFERENCES

- [1] F. Verdouw, B. Tekinerdogan, A. Beulens, S. Wolfert, “Digital twins in agriculture: A state-of-the-art “Smart Agricultural Technology, vol. 3, Art. no. 100094, 2023, doi: 10.1016/j.jatech.2022.100094
- [2] Digital Twins in Agriculture: Orchestration and Control Applications, *Journal of Agricultural and Food Chemistry* 72, no. 19, pp. 10737, 10752, 2024, doi: 10.1021/acs.jafc.4c01934.
- [3] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, pp. 5, 32, 2001, doi: 10.1023/A:1010933404324
- [4] IEEE Author Center, “Structure Your Paper,” IEEE. [Online]. Available: <https://conferences.ieeeauthorcenter.ieee.org/write-your-paper/structure-your-paper/>.
- [5] IEEE Author Center, “Authoring Tools and Templates,” IEEE. [Online]. Available: <https://conferences.ieeeauthorcenter.ieee.org/write-your-paper/authoring-tools-and-templates/>. Accessed: Oct. 15, 2026.