iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: https://doi.org/10.22214/ijraset.2026.85547

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue X Oct 2026- Available at www.ijraset.com

IFAMS: An Integrated Farmer Advisory and Marketplace System Using Random Forest Crop Recommendation, XGBoost Price Forecasting, and Direct Farmer-Buyer Bidding

Disha D S, Disha H Jain, Dhruhi Atykar, Parth Rawat, Vanishree
Department of Computer Science and Engineering, B.M.S. College of Engineering, Bengaluru, India

Abstract: ***Farmers face different problems across the agricultural lifecycle: crops chosen without data driven guidance, high seed and fertilizer costs, limited access to experts, uncertain market prices, and dependence on intermediaries when selling. This paper presents IFAMS, Integrated Farmer Advisory and Marketplace System, a web platform that chains five services into a single workflow: crop recommendation, seed and fertilizer guidance, expert advisory, price prediction, and a farmer buyer bidding marketplace. A Random Forest classifier recommends crops with a confidence score, from soil nutrients, weather and location. An XGBoost regressor forecasts weekly prices and a decision layer converts the forecast into a sell or hold and best market recommendations after transport and holding costs. The system uses a three tier architecture with RESTful communication and role based access for farmers, buyers, experts and administrators. The crop model reaches 98.05% accuracy (98.21% cross validation accuracy), and the price model reaches R^2 = 0.9808, MAE = ₹368.24, RMSE = ₹720.19, and MAPE = 5.52% under a chronological train test split. Both models were trained on generated data that follows the layout of public datasets, and simple baselines score about the same, so these figures should be read with care. A working prototype with four farmers, three buyers, two agronomists, and six open listings demonstrates the end to end workflow. We also discuss the limits of the current evaluation and outline future work.***

Keywords: ***Agricultural advisory, crop recommendation, digital marketplace, price prediction, precision agriculture, Random Forest, XGBoost.***

I. INTRODUCTION

Agriculture supports a large share of livelihoods in India, yet a farmer's decisions are often made with little data. Before sowing, the farmer must choose a crop that suits the soil and the season. During the season, the farmer must buy seed and fertilizer at reasonable cost and recognise crop problems early. After harvest, the farmer must decide when and where to sell. At each stage the farmer typically has limited access to expert advice and to reliable price information, and depends on intermediaries whose margins reduce what the farmer finally receives.

Machine learning has been applied to individual parts of this problem, including crop recommendation [1]–[3], fertilizer recommendation, price forecasting [4], and farmer query support. These efforts mostly address one stage at a time. A farmer who wants advice from sowing to sale must therefore move between separate tools, and the output of one stage (for example, the chosen crop) is not carried into the next (for example, the price forecast for that crop).

This paper presents **IFAMS**, the Integrated Farmer Advisory and Marketplace System. IFAMS follows a single path, *crop recommendation → input guidance → expert advisory → price prediction → marketplace*, so that each stage builds on the previous one. The contributions of this paper are:

- An integrated platform that connects crop recommendation, seed and fertilizer guidance, expert advisory, price prediction, and a bidding marketplace behind one farmer dashboard.
- A Random Forest crop recommender that returns a confidence score, and an XGBoost price forecaster whose output is turned into a market-wise sell-or-hold decision that accounts for transport and holding cost.
- A marketplace in which farmers list produce and buyers bid directly, with the best offer shown to the farmer, reducing dependence on intermediaries.

©IJRASET: All Rights are Reserved

- A working prototype with role-based access for farmers, buyers, agronomists, and administrators, together with measured model performance.
- An evaluation against simple baselines, which shows how much of the headline scores comes from the data and how much from the models.

The rest of the paper is organized as follows. Section II reviews related work. Section III states the problem and objectives. Section IV describes the architecture, and Section V describes the datasets. Section VI explains the implementation of each module. Section VII reports results, Section VIII discusses limitations, Section IX outlines future work, and Section X concludes.

II. RELATED WORK

- Crop recommendation. Shastri et al. [1] advance crop recommendation with supervised machine learning and explainable artificial intelligence, so that the factors behind a recommendation can be inspected. Baishya and Dutta [2] describe a TinyML based crop recommendation system for Precision Agriculture 5.0, which targets low-resource devices. Dey et al. [3] apply the same formulation to both field crops and horticultural crops grown in India. These works treat crop selection as a classification task on soil and weather features, which is the formulation used in IFAMS.
- Fertilizer and input recommendation. Machine learning based fertilizer recommendation work links soil nutrient status to the inputs a farmer should apply, but it does not usually connect the advice to suppliers or prices.
- Agricultural price prediction. Khare [4] reviews machine learning enabled crop price prediction. Tree based ensembles are widely used for this task because prices depend on factors like seasonality, arrivals, and past prices. Forecasts however, leave the farmer to decide where and when to sell his harvest. In the summaries we collected, the comparisons were mostly between machine learning models, and we did not find a comparison with a simple forecast that repeats last week's price, although Hyndman and Koehler [5] recommend including one.
- Farmer advisory and digital marketplaces. The IDinsight and IDRC scoping study [6] examines how digital tools can be made to work for farmers.
- Positioning. The analyzed and reviewed work addresses crop choices, prices, inputs, advice or market, generally in isolation. IFAMS addresses the gap between them: it carries one farmer's context through every stage and ends in a marketplace where the decision can be acted upon.

III. PROBLEM STATEMENT AND OBJECTIVES

Farmers face the following difficulties across the agricultural lifecycle: (i) improper crop selection due to the lack of data driven recommendations; (ii) high seed and fertilizer costs; (iii) uncertain market prices; (iv) limited access to agricultural experts; and (v) over dependence on intermediaries during selling.

The objectives of IFAMS are extracted directly from these difficulties. The system should (a) recommend a suitable crop with a confidence score from soil, weather, and location; (b) guide the farmer to government-approved seeds and fertilizers and to authorized suppliers; (c) let the farmer report a crop problem and receive an expert diagnosis; (d) forecast prices and recommend when and where to sell; and (e) let the farmer sell directly to buyers through bidding.

IV. SYSTEM ARCHITECTURE

A. Three-Tier Design

IFAMS follows a three-tier layered architecture for scalability and maintainability (Fig. 1). The presentation layer is a responsive web application, with English and Kannada interfaces, that opens in a desktop or phone browser; no native mobile application was built. The business logic layer is a backend server that hosts five modules: crop recommendation (Random Forest), seed and fertilizer guidance, crop health advisory, price prediction (XGBoost), and marketplace and bidding. The data layer holds the database and connects to a weather API. The layers communicate through HTTP requests: pages are rendered by the server, and the browser calls about twenty JSON endpoints under /api for crop advice, price forecasts, bids, messages, and notifications. Access is role-based, with separate views for farmers, buyers, experts, and administrators.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

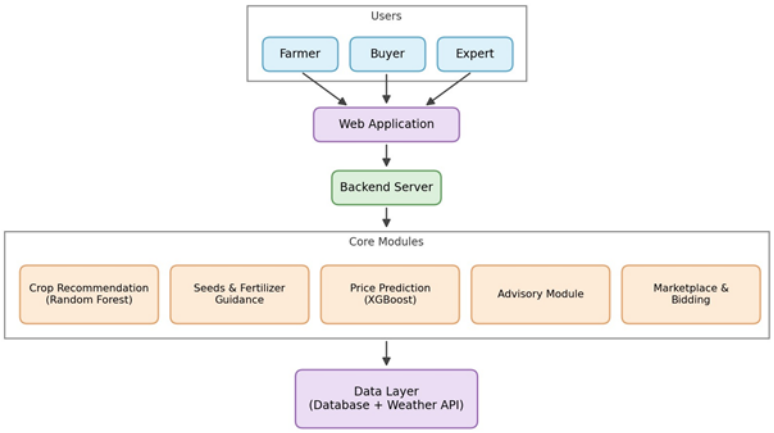


Fig. 1 High-level architecture of IFAMS. Farmers, buyers, and experts use the web application, which talks to the backend server and its five core modules, backed by the data layer.

B. Technology Stack

Table I lists the technologies used in the implementation.

TABLE I
TECHNOLOGY STACK OF THE IFAMS PROTOTYPE

Component	Implementation
Frontend	Server rendered HTML with a small amount of JavaScript
Backend	Python with Flask [7], one blueprint per module
Database	SQLite through SQLAlchemy
Crop ML	Random Forest (scikit-learn [8])
Price ML	XGBoost
External API	Open-Meteo weather service [11]
Version control	Git

V. DATASETS AND DATA SOURCES

Each module uses a separate source, summarized in Table II. The crop data follows the layout of the public Crop Recommendation Dataset from Kaggle [9], whose features are nitrogen (N), phosphorus (P), potassium (K), temperature, humidity, pH, and rainfall, with the crop as the label. The price data follows the layout of AGMARKNET exports [10], which provide commodity, market, date, and minimum, maximum, and modal prices. We did not train on the real files. Both training files were generated by script with a fixed random seed. For the crop data, each of 22 crops has a centre value for every input, taken from the published mean values of the standard dataset, and 140 records per crop (3,080 in all) were drawn from normal distributions around those centres. For the price data there are 260 weekly prices for each of 18 crops in eight Karnataka markets, which is 37,440 rows ending on 27 July 2026, built on a yearly price wave that we set ourselves. The scores in Section VII therefore say a good deal about the generators and little about real farms. A weather API supplies temperature, humidity, and rainfall, with a table of district averages as a fallback. Seed varieties, fertilizers, suppliers, support prices and market coordinates were typed in as demonstration data, and they have not been checked against official sources.

TABLE II
DATA SOURCES AND DATASETS USED BY EACH MODULE

Module	Source	Main data / usage
Crop recomm.	Generated, layout of Kaggle crop dataset [9]	N, P, K, temperature, humidity, pH, rainfall, crop; 3,080 records; trains Random Forest

Module	Source	Main data / usage
Price prediction	Generated, layout of AGMARKNET exports [10]	Crop, market, date, arrivals, min/max/modal price; 37,440 rows; trains XGBoost
Inputs guidance	Typed in by the authors	Seeds, fertilizers, suppliers, support prices, markets
Weather	Open-Meteo API [11]	Temperature, humidity, rainfall for recommendations

VI. MODULE DESIGN AND IMPLEMENTATION

A. Crop Recommendation

The farmer enters soil test values for N,P,K i.e, nitrogen, phosphorus, and potassium, then either retrieves the current district weather data or manually enters their own figures. The pipeline consists of soil, weather, and location data which undergo preprocessing before it is fed into a Random Forest model [12] that produces crop and confidence scores /outputs. The model is invoked only after input validation filters out incomplete or in case of outliers. The model gives outputs: the leading crop along with its confidence score. The dashboard keeps the recent runs to enable the farmer to compare them easily. (Fig. 2). The forest has 300 trees with balanced class weights, and the data is split 80 and 20, stratified by crop. The district is used only to choose where the weather comes from, and it is not an input to the model. Up to three other crops are shown as alternatives.

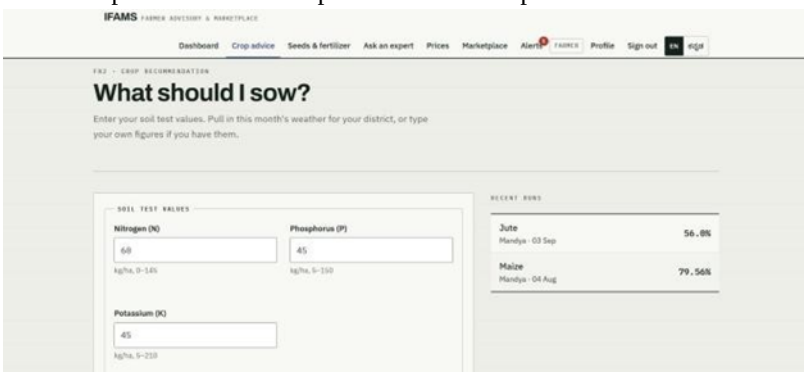


Fig. 2 Crop advice

B. Seed and Fertilizer Guidance

Given the recommended crop, the module follows the given path: recommended crop → approved seeds and fertilizers → supplier information. It recommends government approved inputs like fertilizers and seeds and connects the farmer to authorized manufacturers and suppliers required. The farmer selects a crop and enters the size of the plot. They also provide soil N, P and K values. It adjusts the standard fertilizer dose to the soil condition and if a nutrient reads low the dose is raised by 25%. It cuts by 25% when it reads high. Available options are then listed: fertilizer types, prices and bag sizes. The base doses were typed in from state package of practice tables, and the 25% correction is a fixed rule that we have not calibrated against field trials.

C. Crop Health Advisory

The advisory module works like the following: first image is uploaded by the farmer, expert is notified, expert then gives diagnosis and advice to the farmer. A farmer can report a crop's problem with a photo and description. An agronomist reviews the symptoms and then replies with a diagnosis and recommended remedy, then it shows up in the farmer's account. The farmer can keep talking to the agronomist. For now, a human expert makes the diagnosis, not an image model.

D. Price Prediction and Selling Advice

Weekly standard prices for each crop and market are modelled using an XGBoost regressor [13]. It is based on input features such as calendar and seasonality terms, market arrivals and autoregressive lags of the price itself: 1, 2, 4, 8, 12, 26, and 52 weeks along with rolling means. The train test split is chronological but not shuffled.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

This means the reported score reflects forecasting, not interpolation. The first 85% of the dates (25,488 rows) are used for training and the last 31 weeks (4,464 rows) for testing. Forecasts beyond a week are produced recursively. The confidence band increases with the horizon, but it is a fixed rule of 2.5% plus 1.1% for each week ahead, and it is not well calibrated (Section VII). Forecasts for several weeks ahead and the decision layer were evaluated with a gradient boosted regressor from scikit-learn standing in for XGBoost, and results obtained that way are marked with a dagger (†).

The prototype covers 18 crops across eight markets. A price curve alone doesn't constitute advice. That is where a decision layer comes in, it turns the forecast into a recommendation. The farmer selects the crop, nearest market, quantity in quintals, and how many weeks to look ahead. For each market m and lookahead week w, the expected net price per quintal is $\text{Net}_m(w) = P_m(w) - T_m - \rho \cdot w \cdot P_0$ (1), where $P_m(w)$ is the forecast price, T_m is the transport cost to market m, P_0 is the current price, and $\rho = 0.6\%$ is the holding cost per week as a fraction of the crop's value. What's the best market? It's the one with the highest peak net price over the look-ahead window: $m^* = \text{argmax}_m \max_{w \le H} \text{Net}_m(w)$ (2), where H is the selected horizon. Hold is recommended only when the best net price is above today's price by at least 1.5%. Transport is taken as ₹85 per quintal for a market outside the farmer's district and zero inside it, and a market is flagged as better only when it beats the home market by more than ₹150 per quintal. The layer compares selling now against the best future net price, checks the current price against the crop's minimum support price, and gives a Sell or Hold recommendation. The support price only triggers a warning, and it is not an input to the hold or sell rule. Fig. 3 shows the input screen. Fig. 4 shows the marketwise comparison.

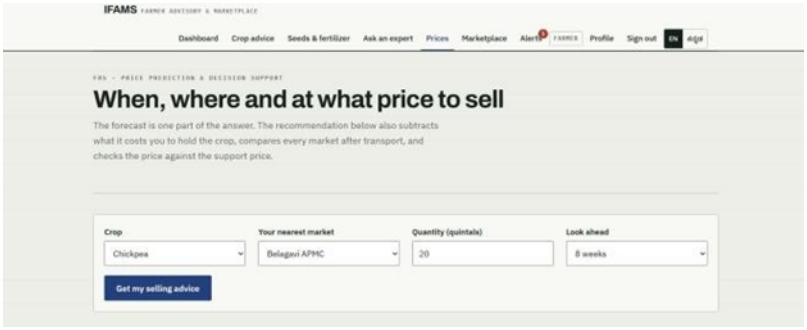


Fig. 3 Price prediction

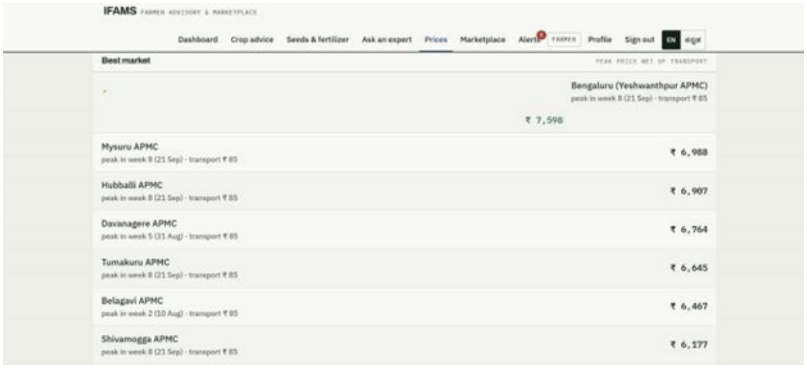


Fig. 4 Best market recommendation

E. Market and Bidding

The marketplace works like this: farmer listing → buyers → bids → best offer. Farmers list their produce, including crop, quantity, grade, location, and base price. Buyers then view these listings and place bids in the marketplace. A single listing can get several bids. It's simple: the system shows the farmer the highest bid and tracks the listing's status (open or closed) and the number of bids received. A bid below the base price is rejected, and a new bid has to beat the leading one by 1%. The farmer then accepts a specific bid, and the buyer is notified. Because farmer and buyer interact directly, no intermediary is needed.

F. Farmer Dashboard and Roles

The farmer dashboard (Fig. 5) is a single entry point to all modules: crop advice, seeds and fertilizer, ask an agronomist, price prediction, and the marketplace. It shows the latest crop advice with its confidence, live listings, bids received, the best offer, and notifications. Role-based access gives farmers, buyers, agronomists, and administrators only the screens they need.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

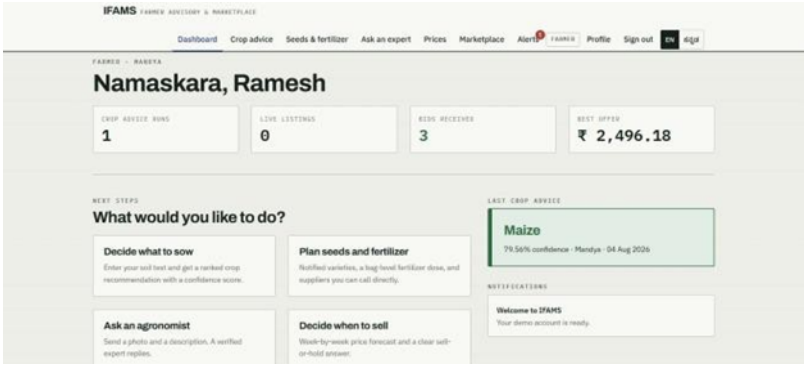


Fig. 5 Farmer dashboard showing recent crop advice, bids received, best offer, and quick access to every module.

VII. RESULTS

A. Model Performance

Table III summarizes the measured performance of the two machine-learning models. The crop model reaches 98.05% accuracy, and cross-validation gives a similar 98.21%, which suggests that the result does not depend on a single split. The price model is evaluated on a chronological split. It reaches $R^2 = 0.9808$, with a mean absolute error of ₹368.24 and a root mean squared error of ₹720.19 per quintal. The mean absolute percentage error is 5.52%. All of these were measured on generated data.

TABLE III
MODEL PERFORMANCE MEASURED ON THE IFAMS MODELS

Model / metric	Value
Crop recommendation: accuracy	98.05%
Crop recommendation: cross-validation accuracy	98.21%
Price prediction: R^2 score	0.9808
Price prediction: MAE (per quintal)	₹368.24
Price prediction: RMSE (per quintal)	₹720.19
Price prediction: MAPE	5.52%

B. Comparison with Simple Classifiers

A high score does not say much by itself, so six simpler classifiers were trained on the same 80 and 20 split. Table IV shows the result. Naive Bayes, logistic regression, k nearest neighbours, the support vector machine and QDA all score at or above the Random Forest on the test split, and only the decision tree is clearly lower. The Random Forest is not the best model on this data. This is not surprising in hindsight, since the records are clouds around fixed centres, which is the situation in which naive Bayes and QDA are close to ideal. We kept the Random Forest because it gives a usable confidence score from the vote share, and we do not claim it is the best choice.

TABLE IV
CROP CLASSIFIERS ON THE GENERATED DATA

Model	**Test acc.**	**5 fold CV (%)**
Decision tree	95.94%	95.91 ± 0.54
Gaussian naive Bayes	98.38%	98.60 ± 0.22
Logistic regression	98.38%	98.08 ± 0.22
k nearest neighbours (k = 5)	98.54%	97.86 ± 0.50
SVM (RBF)	98.54%	98.34 ± 0.43
QDA	98.21%	98.28 ± 0.22
Random Forest (deployed)	98.05%	98.21 ± 0.23

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The confidence score turned out to be the more useful output. Predictions grouped by confidence follow calibration checks in the style of Guo et al. [14], and the expected calibration error is 0.052. In every group the accuracy is higher than the mean confidence, so the score understates reliability. Where the confidence is 0.8 or more, which is 89% of the cases, the accuracy is 99.6%, while the 17 cases below 0.6 are right only 76% of the time. This suggests showing a message to consult an expert at low confidence. Noise is the main caution. When Gaussian noise of 0.1, 0.25 and 0.5 standard deviations was added to the test inputs, the Random Forest fell from 98.1% to 96.5%, 83.2% and 59.2%, and the support vector machine held up better at 98.0%, 93.4% and 78.6%.

C. Feature Importance

Fig. 6 shows the feature importance of the trained Random Forest. Rainfall (20.2%) and humidity (18.9%) are the most influential inputs, followed by potassium (16.0%), phosphorus (15.0%), nitrogen (13.5%), and temperature (11.7%). Soil pH is the least influential (4.8%). The three macronutrients together account for 44.5% of the importance, so soil test values matter as much as weather. Impurity based importance is known to be biased towards some kinds of input [15], so permutation importance was also measured on the test set. By that measure nitrogen comes first, since shuffling it costs 31.8 points of accuracy, followed by humidity (26.2) and rainfall (19.0). The two rankings do not agree, and we would not read much into either one.

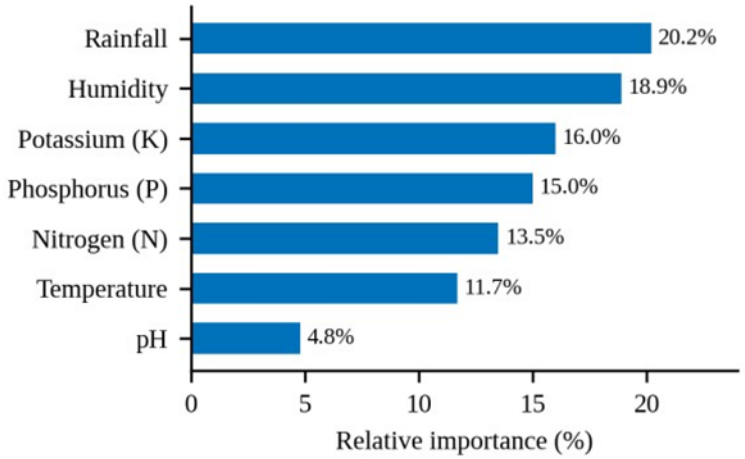


Fig. 6 Impurity based feature importance of the Random Forest crop recommendation model.

D. Price Forecasting Against Baselines

Table V compares the deployed XGBoost model with simple alternatives on the 4,464 test rows. Persistence, which just repeats last week's price, gets R^2 = 0.9816 and an RMSE of 705.43, slightly better than XGBoost on both. XGBoost is a little better on MAE and MAPE, and a ridge regression on the same features does better still. An R^2 above 0.98 one week ahead therefore says almost nothing about skill, because weekly prices change slowly. The previous week's price, the four week mean and the price two weeks back hold 97.7% of the importance.

TABLE V
ONE WEEK AHEAD FORECASTING, CHRONOLOGICAL TEST SPLIT

Method	R^2	MAE	RMSE	MAPE
Persistence	0.9816	376.82	705.43	5.80%
Seasonal naive (52 wk)	0.9380	758.92	1294.77	11.96%
4 week moving avg.	0.9739	453.79	839.87	7.14%
Ridge regression	0.9838	350.98	662.41	5.46%
Gradient boosting	0.9824	359.99	690.29	5.59%
XGBoost (deployed)	0.9808	368.24	720.19	5.52%

MAE and RMSE are in rupees per quintal.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The model is more useful at longer horizons, which is where the selling advice operates. For recursive forecasts from 3,456 origins†, the MAPE of the gradient boosted model against persistence is 5.62% against 5.83% at one week, 8.66% against 10.51% at four weeks, and 10.66% against 16.21% at eight weeks, which is about 34% lower. The band shown to the farmer contained the real price in only 44% of cases at one week, rising to 63% at eight weeks†, so in the interface it should be described as an indicative range.

E. Decision Layer Backtest

The layer was replayed over the test period, with 3,456 origins†, a holding cost of 0.6% per week and a limit of 8 weeks. At each origin it saw only the history up to that week. Table VI shows that following its advice every time would have raised the realised net price by 8.68% on average (95% interval 8.01 to 9.34) compared with selling at once. It recommends holding in 43% of cases, and 77% of those holds are profitable after the holding cost. A hindsight oracle would have gained 14.36%. A persistence forecast passed through the same layer never recommends holding, so it gains nothing, and holding for a fixed time without a forecast gains little, so the benefit seems to come from choosing when to hold. The 0.6% holding cost is an assumption, and when it was varied from 0 to 1.5% per week the gain moved from 11.05% to 6.10% and stayed positive. The size of the benefit depends on the assumed cost more than its direction.

TABLE VI
BACKTEST OF THE SELL OR HOLD LAYER†

Strategy	Hold rate	Gain
Layer, gradient boosted forecast	42.9%	8.68%
Layer, seasonal naive forecast	55.1%	7.16%
Layer, persistence forecast	0%	0.00%
Always hold 8 weeks	100%	2.13%
Always hold 4 weeks	100%	0.67%
Hindsight oracle	n/a	14.36%

Gain is the mean realised price change compared with selling at the origin.

F. System Behaviour

The prototype was exercised with four farmers, three buyers, two agronomists, and six open listings. For a price query for chickpea with 20 quintals at Belagavi APMC and an eight-week horizon, the system ranked Bengaluru (Yeshwanthpur APMC) first with an expected net price of ₹7,598 per quintal, ahead of Mysuru (₹6,988) and Hubballi (₹6,907), after subtracting transport. In the marketplace, a rice listing of 45 quintals received three bids, and the highest bid of ₹2,496.18 per quintal was shown to the farmer. The accounts, listings and bids come from a seed script. These are demonstrations of the workflow, not a field trial.

VIII. DISCUSSION AND LIMITATIONS

According to the results the two models represent the data effectively and the five modules can be integrated under one interface as a unified framework. However, there are several limitations which should be considered with these results.

- Benchmark evaluation: The crop recommendation accuracy is measured on generated data that follows the layout of a public dataset, where scores near 98% are frequently observed. It has not been validated against real world agricultural farm outcomes, so it should be treated as an indicator of dataset fit. Five simpler classifiers score about the same, and accuracy falls to 83.2% when noise of 0.25 standard deviations is added. Repeating the evaluation on the real Kaggle and AGMARKNET files is the first thing we would do next.
- Expert advisory: Crop diagnosis (like diseases or stunt growth) depends on manual review by human agronomists not by AI. There is no automated image based disease detection yet, and response time and quality of diagnoses were not measured.
- Price forecasts: Based on a single chronological split the evaluation of the price model is analyzed, and one training run was used. Persistence matches XGBoost on one step R^2, and the multi week and decision layer results rely on a substitute model and on prices with a clean yearly wave that real markets may not have. Recursive multi week forecasts lead to error accumulation which is why the confidence intervals widens over extended horizons, and the band covered only 44% to 63% of real prices. Transport cost and holding costs are simplified to fixed rates, and quantity and spoilage are not modelled.

- Limited initial testing: We tested the prototype on limited scale involving only 4 farmers, 2 agronomists and 3 buyers, who were seeded demonstration accounts, for technical validation proof of concept. Hence, we still need extensive field testing, a user study and load testing.
- Marketplace: There is no payment, identity check, dispute handling, or protection against buyers acting together.
- Integration and reference data: Only the crop name passes from the crop page to the input planner, and the soil values have to be typed again. Seed notification numbers, support prices and supplier contacts are demonstration entries.
- Geographic and crop scope: The system currently serves a limited set of crops, markets and districts, mainly concentrated in Karnataka.

IX. FUTURE WORK

Planned extensions are: multilingual support (Kannada, Hindi, Tamil, and other regional languages); a payment gateway for secure transactions between farmers and buyers; real-time logistics and supply-chain tracking; IoT integration with soil sensors and weather stations; drone-based crop monitoring; offline functionality for areas with poor connectivity; automated disease detection from images; a regional-language voice assistant; and expansion to more crops and districts and eventually across India.

X. CONCLUSIONS

IFAMS provides an end-to-end platform that supports a farmer from crop planning to sale. It combines a Random Forest crop recommender (98.05% accuracy), seed and fertilizer guidance, expert advisory, an XGBoost price forecaster ($R^2 = 0.9808$, MAPE = 5.52%) with a market-wise sell-or-hold decision layer, and a direct bidding marketplace, all behind one dashboard. By joining advisory, machine learning, and market access, the system aims to reduce dependence on intermediaries and to support data-driven farming decisions. Field validation, regional-language access, and automated disease detection are the most important next steps.

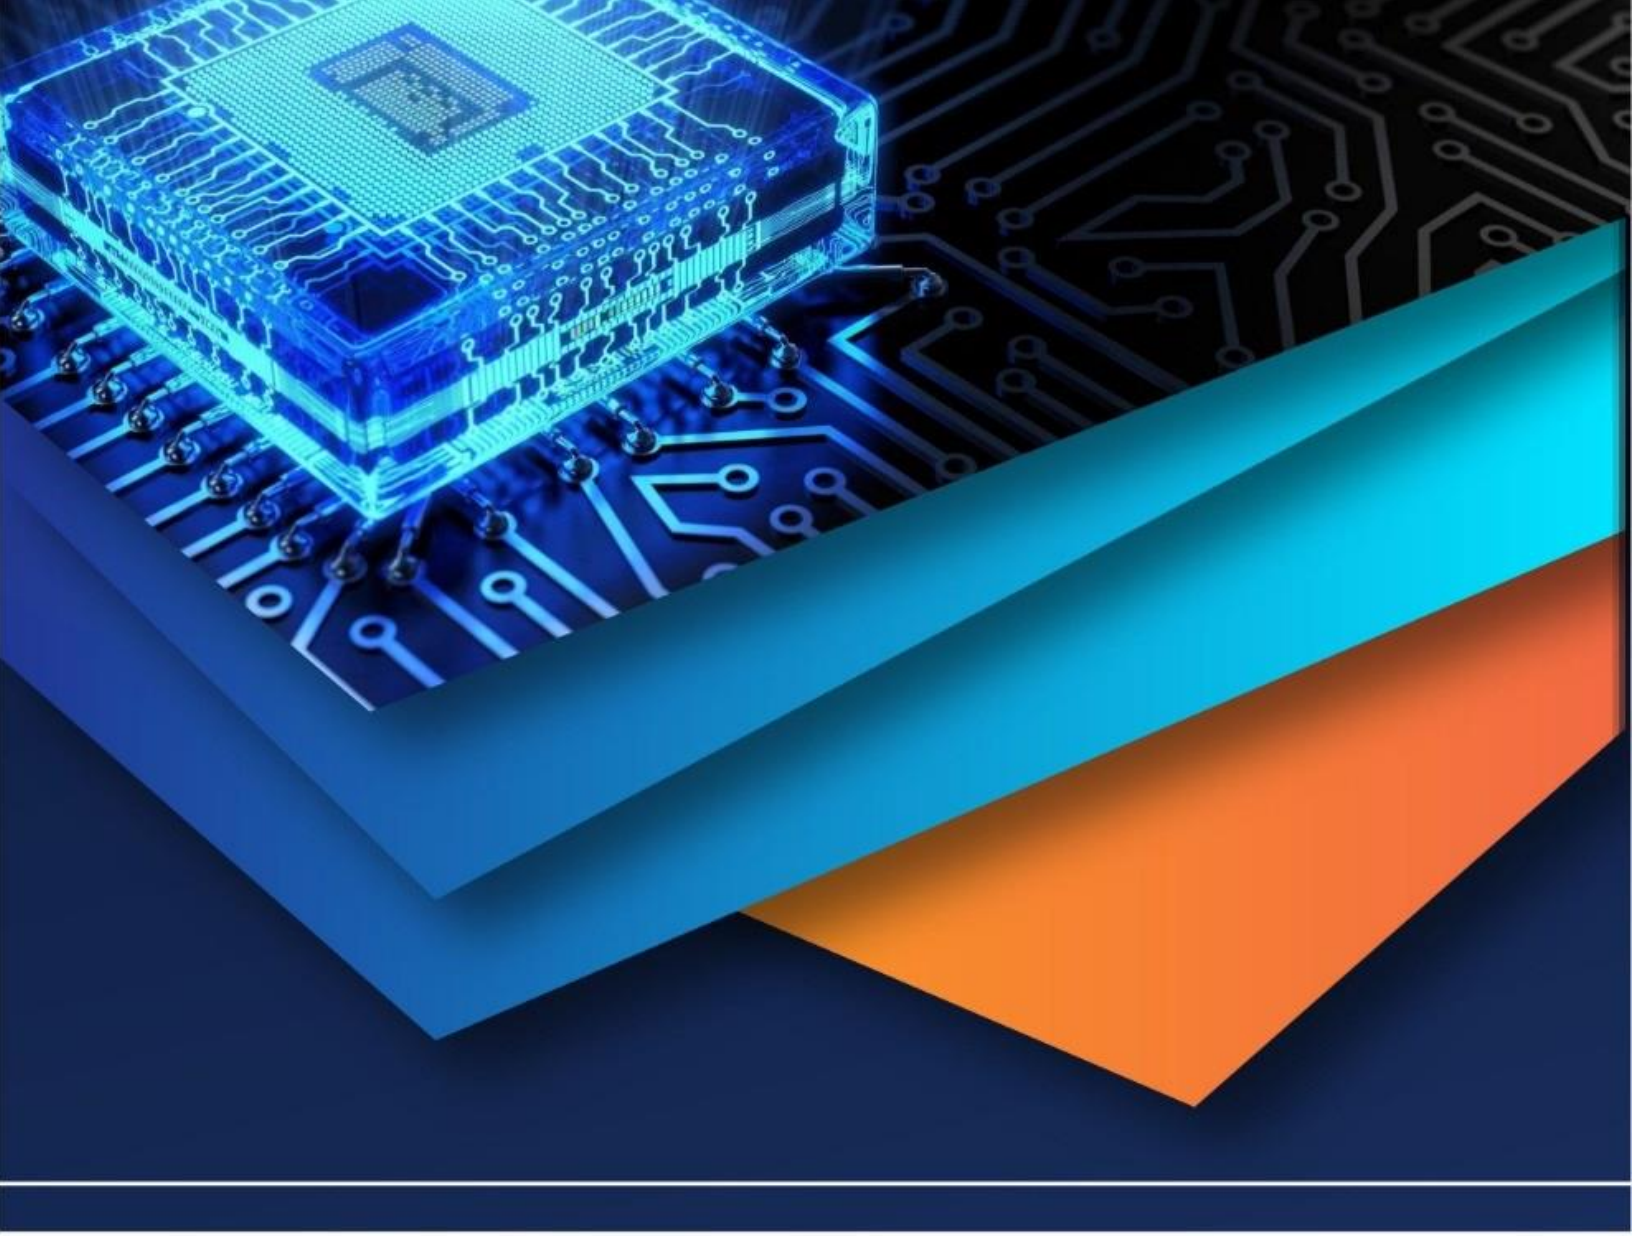
XI. ACKNOWLEDGMENT

The authors thank the Department of Computer Science and Engineering, B.M.S. College of Engineering, Bengaluru, for their support.

REFERENCES

[1] S. Shastri, S. Kumar, V. Mansotra, and R. Salgotra, "Advancing crop recommendation system with supervised machine learning and explainable artificial intelligence," *Sci. Rep.*, vol. 15, 2025, doi: 10.1038/s41598-025-07003-8.
[2] M. Baishya and L. Dutta, "Tiny ML based crop recommendation system for Precision Agriculture 5.0," *Smart Agric. Technol.*, 2025, Art. no. 101247.
[3] B. Dey, J. Ferdous, and R. Ahmed, "Machine learning based recommendation of agricultural and horticultural crop farming in India under the regime of NPK, soil pH and three climatic variables," *Heliyon*, vol. 10, no. 3, 2024.
[4] S. Khare, "Machine learning enabled crop price prediction: A review," *Indian J. Comput. Sci.*, vol. 11, no. 1, pp. 31–37, 2026, doi: 10.17010/ijcs/2026/v11/i1/175963.
[5] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *Int. J. Forecasting*, vol. 22, no. 4, pp. 679–688, 2006.
[6] IDinsight, "Making digital tools work for farmers in Africa," scoping study funded by IDRC, Apr. 2025. [Online]. Available: https://www.idinsight.org/publication/making-digital-tools-work-for-farmers-in-africa/
[7] Pallets Projects, "Flask documentation." [Online]. Available: https://flask.palletsprojects.com
[8] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
[9] A. Ingle, "Crop Recommendation Dataset," Kaggle. [Online]. Available: https://www.kaggle.com/datasets/atharvaingle/crop-recommendation-dataset
[10] "AGMARKNET: Agricultural Marketing Information Network," Government of India. [Online]. Available: https://agmarknet.gov.in
[11] "Open-Meteo: Free weather forecast API." [Online]. Available: https://open-meteo.com
[12] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
[13] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 785–794.
[14] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. 34th Int. Conf. Mach. Learn. (ICML)*, 2017, pp. 1321–1330.
[15] C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn, "Bias in random forest variable importance measures," *BMC Bioinformatics*, vol. 8, Art. no. 25, 2007.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)