

Artificial Intelligence-Based Intelligent Fertigation Recommendation Systems for Precision Agriculture A Structured Review with an Applied Case Study on Western Maharashtra Soil–Crop–Fertilizer Data

SHRADDHA S. TAYADE^{1*}, DR. YOGESH V. CHIMATE²

^{1, 2} M. Tech Scholar, AI/ML in Agriculture, D. Y. Patil Agriculture and Technical University, Talsande, Kolhapur, Maharashtra, India

*Corresponding author: Shraddha Shivshankar Tayade

Abstract- Artificial-intelligence-enabled fertilizer and fertigation recommendation has progressed from single-task crop or fertilizer classification toward integrated decision-support architectures. However, predictive classification, agronomic dose calculation, irrigation scheduling, explainability, real-time sensing, and farmer-facing delivery are often studied separately. This paper presents a structured narrative review of these strands and an applied machine-learning case study using the publicly available Crop and Fertilizer Dataset for Western Maharashtra. The case study contains 4,513 records spanning five districts, 16 crops, and 19 fertilizer classes. Seven classifiers were compared using an 80:20 stratified train-test split with five-fold stratified cross-validation on the training set. XGBoost achieved 97.34% test accuracy and $94.21\% \pm 1.15\%$ cross-validated accuracy, while Random Forest achieved 93.58% and $90.89\% \pm 0.80\%$, respectively. TreeSHAP analysis of Random Forest identified crop identity, potassium, and nitrogen as the leading predictors of the historical fertilizer class. These results are interpreted as a computational baseline rather than proof of agronomic optimality because the target label represents recorded fertilizer choices. The review also incorporates evidence on IoT/edge-cloud sensing, multilingual agricultural advisory, and federated learning. It concludes that RF/SHAP, IoT sensing, and multilingual interfaces are established capabilities; a more defensible research direction is an auditable pipeline that separates fertilizer identity, nutrient dose, and application timing, connects explainable prediction to sequential scheduling, and independently benchmarks outputs against authoritative agronomic guidance. The proposed Intelligent Fertigation Recommendation System (IFRS) is therefore presented as a research framework requiring multi-season and field validation before claims of yield, water, nutrient-use-efficiency, or adoption benefits.

Index Terms: fertigation; precision agriculture; machine learning; explainable AI; SHAP; Q-learning

I. INTRODUCTION

Nutrient management is a central component of sustainable crop production because fertilizer decisions influence crop productivity, production cost, nutrient-use efficiency, and environmental losses. Conventional recommendations can be insufficient when soil fertility, crop type, climate, irrigation conditions, and management practices vary within and between farms. Precision agriculture seeks to reduce this mismatch by adapting inputs to measured or inferred field conditions.

Fertigation—the application of soluble nutrients through irrigation—creates a particularly demanding decision-support problem. A useful system must reconcile soil nutrient status, crop identity and growth stage, irrigation events, weather, water availability, fertilizer formulation, and timing. This differs fundamentally from a single-shot classification task because an agronomically meaningful recommendation may require both a dose and a schedule.

Recent studies demonstrate progress in several related directions. Interpretable crop recommendation has used Random Forest, SHAP, and LIME [1]; explainable-AI research has formalized post-hoc feature attribution for agricultural decision support [2]; IoT systems have connected soil sensing with machine-learning recommendations [4]; and federated-learning research has addressed privacy and non-IID heterogeneity in distributed agricultural data [5]. Importantly, an independent 2025 Scientific Reports study applied TabNet and SHAP to the same

Western Maharashtra dataset used here and reported 95.24% fertilizer-class accuracy and 96.21% crop-class accuracy [3]. Therefore, RF/SHAP on this dataset cannot be claimed as a novel contribution by itself.

Multilingual agricultural advisory is also becoming an operational capability rather than a purely academic feature. Government initiatives such as Kisan e-Mitra and Bharat Vistar, together with Farmer.CHAT, demonstrate that language-accessible agricultural information is already being deployed at scale [6]–[8]. Consequently, the present work does not treat multilingual output as its headline novelty. Instead, it proposes a narrower role for language generation: translating an evidence-linked fertigation recommendation and its explanation into the farmer's preferred language.

The objectives of this paper are therefore:

- (i) to critically synthesize AI-based fertilizer/fertigation recommendation research across ML, XAI, IoT/edge-cloud, multilingual advisory, and privacy-preserving learning;
- (ii) to assess dataset, validation, and reproducibility limitations;
- (iii) to provide an applied comparative baseline on the Western Maharashtra dataset; and
- (iv) to formulate a cautious IFRS architecture that clearly separates demonstrated components from proposed components.

II. REVIEW METHODOLOGY

A. Nature of the Review

This paper follows a structured-search-inspired approach rather than a formally registered systematic review. The search was designed to identify representative and directly relevant primary studies and authoritative sources, but no pre-registered protocol, independent dual-reviewer screening, interrater agreement calculation, or PRISMA identification flow was generated. Accordingly, the manuscript does not claim exhaustive systematic coverage.

B. Search Strategy

Searches were organized into five thematic clusters: (1) machine learning for soil-nutrient, fertilizer, and crop recommendation; (2) explainable AI, particularly

SHAP and LIME, for agricultural decision support; (3) IoT, edge, and cloud architectures for smart irrigation and fertigation; (4) multilingual farmer advisory systems; and (5) federated learning and agricultural data governance. Searches emphasized 2020–2026 literature, while foundational algorithm papers were retained where required. Publisher platforms, scholarly indexes, government portals, and original dataset records were prioritized.

C. Inclusion and Exclusion Criteria

Included sources addressed agricultural nutrient/fertilizer recommendation, XAI on agricultural tabular data, smart irrigation/fertigation sensing architectures, multilingual agricultural advisory, or privacy-preserving learning. Studies limited to crop disease image classification, remote-sensing-only yield prediction, or generic XAI/federated-learning applications without a direct agricultural decision-support connection were excluded from the main synthesis.

D. Verification and Evidence Discipline

Reference details were checked against original publisher pages, primary conference/journal records, official government sources, or the original dataset record. Secondary web summaries were not used as the bibliographic basis of the final reference list. Claims are explicitly bounded when a source is a preprint, a web resource, or an applied dataset rather than a peer-reviewed experimental paper. Reported performance values are attributed to their original studies or, for the present case study, to the analysis described in this manuscript.

E. Review Limitation

Because the literature search was targeted rather than a database-registered systematic review, the synthesis should be read as a structured narrative review. A future extension can add Scopus/Web of Science queries, screening counts, quality-assessment scores, and a formal PRISMA flow before making exhaustive claims.

III. AI-BASED SOIL-NUTRIENT AND FERTILIZER RECOMMENDATION

Agricultural recommendation datasets commonly contain nitrogen (N), phosphorus (P), potassium (K),

pH, rainfall, temperature, humidity, crop identity, soil descriptors, and a recommendation label. Supervised models can learn associations between these features and historical crop or fertilizer classes. Random Forest is a strong tabular baseline because it captures nonlinear interactions and is comparatively robust [10], while gradient-boosted trees such as XGBoost provide a high-performance alternative for structured prediction [11].

A critical distinction is required between fertilizer-class prediction and agronomically valid dosing. A classifier can reproduce a historical fertilizer label without demonstrating that the historical choice was optimal. A deployable fertigation system should therefore distinguish at least three outputs: fertilizer identity, nutrient dose, and application timing. The first can be formulated as classification; the latter two require agronomic rules, optimization, or sequential decision-making.

The Western Maharashtra dataset provides a useful regional benchmark because it contains district, soil-color, N, P, K, pH, rainfall, temperature, crop, and fertilizer fields [9]. Nevertheless, the dataset should not be treated as a direct ground-truth source for agronomic optimality. Its fertilizer field is a recorded recommendation and may reproduce the assumptions or inefficiencies present in historical practice.

IV. EXPLAINABLE AI FOR AGRICULTURAL DECISION SUPPORT

SHAP (SHapley Additive exPlanations) provides a game-theoretic framework for attributing model output to input features [12]. TreeSHAP extends this idea efficiently to tree-based models and supports consistent local and global interpretation [13]. In agricultural decision support, the practical value is not simply producing a ranking of features, but attaching a traceable explanation to a specific recommendation. Bouni et al. demonstrated interpretable crop recommendation using Random Forest, SHAP, and LIME [1], while Turgut et al. presented an edge-oriented explainable crop-recommendation framework using LIME, SHAP, ELI5, and counterfactual explanations [14]. Pai et al. reviewed XAI applications in agriculture and emphasized the importance of matching explanation methods to the

data and model type [2]. The 2025 TabNet study on the Western Maharashtra dataset further confirms that SHAP-based interpretability is already being applied to fertilizer recommendation in this regional data setting [3].

The present study therefore treats SHAP as a transparency mechanism rather than a novelty claim. It also distinguishes association from causation. If crop identity is the strongest predictor of a historical fertilizer class, the correct interpretation is that crop identity strongly predicts the recorded class in the dataset—not that crop identity alone causes the agronomic fertilizer requirement.

V. IOT, EDGE, AND CLOUD ARCHITECTURES

IoT systems can provide the sensing layer required for near-real-time decision support. Sunandini et al. described an IoT-based architecture using NPK, temperature, pH, humidity, and rainfall sensing with wireless transmission and machine-learning-based recommendation [4]. For fertigation, edge processing can perform unit checking, missing-data detection, plausibility checks, outlier filtering, and temporary buffering before cloud inference.

A safe architecture should not allow an incomplete or anomalous sensor packet to automatically trigger fertilizer application. Model inference should be accompanied by timestamps, sensor-quality flags, calibration status, and out-of-distribution checks. Cloud services can maintain model versions, explanations, recommendation histories, and audit trails, while edge components provide resilience during intermittent connectivity.

The literature reports potential benefits from AI-enabled irrigation and agricultural sensing, but system-level gains are strongly dependent on crop, location, hardware, management, and trial design. Consequently, this manuscript does not transfer any water-saving or yield percentage reported elsewhere to the proposed IFRS.

VI. MULTILINGUAL FARMER ADVISORY

Kisan e-Mitra is an AI-powered voice chatbot used by the Government of India for farmer queries and has been reported to support 11 Indian languages [6]. Bharat Vistar was launched in 2026 as an AI-based digital platform intended to make farming information available in local languages through a simple access channel [8]. Farmer.CHAT, developed by Digital Green and Gooye.AI, combines agricultural knowledge resources with multilingual conversational access [7].

These developments materially narrow the novelty space for a general multilingual farmer chatbot. A more defensible IFRS advisory layer is therefore task-specific: it should explain a fertilizer or fertigation recommendation, state the principal evidence, communicate uncertainty or missing-data warnings, and provide a concise agronomic rationale in the farmer's selected language.

VII. FEDERATED LEARNING AND DATA GOVERNANCE

Agricultural datasets are distributed across farms, cooperatives, laboratories, extension offices, and public repositories. Federated learning can support collaborative model development without centrally collecting raw data, but non-IID data remain a major challenge because soil, climate, crop choice, and management practices vary across regions [5].

For IFRS, federated learning is therefore treated as a future privacy-preserving extension rather than a

demonstrated component. A credible evaluation would require multiple independently owned datasets, explicit partitioning protocols, communication constraints, privacy assumptions, and performance analysis under realistic non-IID conditions.

VIII. COMPARATIVE LITERATURE SYNTHESIS AND RESEARCH GAP

The reviewed evidence supports four conclusions. First, RF/SHAP interpretability is established rather than novel by itself [1], [2], [3]. Second, IoT sensing and cloud-connected agricultural decision-support architectures are established design patterns [4]. Third, multilingual agricultural advisory is an active government and industry capability [6]–[8]. Fourth, sequential fertigation scheduling is conceptually different from single-shot fertilizer classification and requires its own state representation, action space, reward function, and temporal validation.

Accordingly, the most defensible research gap is not that AI-based fertilizer recommendation has not been studied. Instead, the gap is the limited evidence for a reproducible end-to-end fertigation decision pipeline in which fertilizer identity, nutrient dose, and application timing are explicitly separated; prediction is explainable; sequential scheduling is independently evaluated; recommendations are auditable against authoritative agronomic guidance; and the complete pipeline is eventually tested in multi-season field trials. The exact combination proposed here should be treated as a research hypothesis rather than a proven first-of-its-kind contribution.

Table I. Comparative synthesis of the most directly relevant sources

Study / Source	Year	Focus	Method	Key evidence	Gap relative to IFRS
Bouni et al. [1]	2024	Interpretable crop recommendation	RF + SHAP + LIME	Interpretable crop-class prediction	Crop recommendation; not a complete fertigation schedule
Pai et al. [2]	2025	XAI in agriculture	Review of SHAP/LIME and related methods	XAI selection depends on	Not fertigation-specific; no field scheduler

				model/data type	
Venkateswara & Padmanaban [3]	2025	Fertilizer and crop recommendation	TabNet + SHAP	95.24% fertilizer; 96.21% crop accuracy	Same Western Maharashtra dataset; does not validate sequential fertigation
Sunandini et al. [4]	2023	IoT soil/fertilizer monitoring	Sensors + WSN + ML	Sensor-to-cloud agricultural pipeline	No independently validated dose/schedule engine
Dembani et al. [5]	2025	Federated learning in agriculture	Privacy/FL review	Non-IID and privacy challenges	Fertigation-specific FL remains unvalidated
Kisan e-Mitra [6]	2025	Government farmer advisory	Voice AI chatbot	11-language support reported	General scheme/advisory scope, not fertigation-specific
Farmer.CHAT [7]	2023	Multilingual agricultural advisory	Knowledge-grounded conversational AI	Multilingual farmer access	General advisory; not a dose/schedule validation system
Bharat Vistar [8]	2026	National AI agricultural platform	AI + agricultural information systems	Local-language farming information	National platform; IFRS must remain domain-specific
Turgut et al. [14]	2024	Edge/XAI crop recommendation	Edge ML + LIME/SHAP/counterfactuals	Local/global explanations	Crop recommendation; not validated fertigation scheduling

IX. APPLIED CASE STUDY: WESTERN MAHARASHTRA DATASET

A. Dataset Description

The applied analysis uses the publicly available Crop and Fertilizer Dataset for Western Maharashtra [9]. The dataset is district-wise and contains fields for district name, soil color, nitrogen, phosphorus, potassium, pH, rainfall, temperature, crop, and fertilizer. The analysis used 4,513 records spanning five districts—Kolhapur, Solapur, Satara, Sangli, and Pune—16 crops, and 19 fertilizer classes. The dataset

is suitable for regional benchmarking, but the fertilizer field represents a recorded recommendation rather than an independently measured agronomic optimum. An independent 2025 study subsequently used the same dataset for TabNet-based fertilizer and crop recommendation with SHAP [3]. This establishes that the dataset is not a novel contribution of the present work. The value of the present case study is instead the transparent comparison of multiple classical models and the interpretation of the resulting baseline for a proposed fertigation architecture.




Fig. 2. Distribution of records across the 16 crops and 19 fertilizer classes.




Fig. 3. Correlation matrix of the six continuous soil and environmental parameters.




Fig. 4. District-wise distribution of nitrogen, phosphorus, and potassium across the five districts.

B. Preprocessing and Experimental Setup

Categorical variables (District_Name, Soil_color, and Crop) were label-encoded, while Fertilizer was encoded separately as the 19-class target. The dataset was divided using an 80:20 stratified split, producing

3,610 training records and 903 test records. Five-fold stratified cross-validation was applied to the training partition. Standardization was used for Logistic Regression, KNN, and SVM, while tree-based models were trained on unscaled inputs.

The compared models were Logistic Regression, K-Nearest Neighbors (K=7), Decision Tree, Random Forest (300 trees), Gradient Boosting, XGBoost, and Support Vector Machine with an RBF kernel. Performance was summarized using test accuracy, weighted precision, weighted recall, weighted F1-score, and five-fold cross-validation accuracy.

C. Model Comparison

The results show a clear advantage for tree ensembles on this 19-class classification task. XGBoost produced the highest test accuracy, while Random Forest provided a strong second baseline with a mature explanation mechanism. These values should be independently reproduced from the analysis code before journal submission.




Table II. Performance comparison for 19-class fertilizer prediction

Model	Test Accuracy	Precision (wtd.)	Recall (wtd.)	F1-score (wtd.)	5-Fold CV Accuracy
XGBoost	97.34%	0.9742	0.9734	0.9735	94.21% ± 1.15%
Random Forest	93.58%	0.9379	0.9358	0.9359	90.89% ± 0.80%
Decision Tree	79.73%	0.8012	0.7973	0.7934	78.75% ± 2.61%
Gradient Boosting	72.20%	0.7246	0.7220	0.7205	70.14% ± 5.36%
K-Nearest Neighbors	67.44%	0.6839	0.6744	0.6760	63.66% ± 1.81%
SVM (RBF)	44.63%	0.4218	0.4463	0.4169	45.46% ± 0.97%
Logistic Regression	41.09%	0.3936	0.4109	0.3841	41.55% ± 1.32%

XGBoost achieved the highest test accuracy (97.34%) and five-fold cross-validated accuracy (94.21% ± 1.15%). Random Forest achieved 93.58% test accuracy and 90.89% ± 0.80% cross-validated accuracy. The results should be interpreted as performance on the supplied classification task, not as evidence that 97.34% of fertilizer decisions are agronomically correct. Because the target represents

Fig. 5. Reported test-set accuracy across the seven compared algorithms.




Fig. 6. Confusion matrix for the reported XGBoost model.

historical recommendation classes, the classifier may learn the decision patterns embedded in the source data rather than an independently validated optimum.

D. Model Selection and SHAP Analysis

Random Forest is retained as the primary explanatory baseline because TreeSHAP provides a mature mechanism for local and global tree-model

explanations [13]. XGBoost remains the strongest predictive benchmark in the supplied experiment. A deployment study should compare both models using calibration, class-wise performance, explanation stability, inference latency, out-of-distribution detection, and agronomic validity rather than accuracy alone.




Fig. 7. Random Forest feature importance reported in the applied analysis.




Fig. 8. Global mean absolute SHAP importance for the Random Forest fertilizer classifier.

TreeSHAP explanations computed for a representative sample of 300 test records ranked crop identity, potassium, and nitrogen as the leading predictors, followed by phosphorus, rainfall, and temperature, with pH, soil color, and district contributing less. The ranking is plausible as a description of the historical label structure, but it should not be interpreted causally. In particular, crop identity may dominate because fertilizer classes in the dataset are strongly conditioned on crop choice.

X. PROPOSED INTELLIGENT FERTIGATION RECOMMENDATION SYSTEM (IFRS)

The proposed IFRS is organized into four functional layers plus an independent reference/validation layer. The design intentionally separates demonstrated components from future components. The Random Forest/SHAP classification baseline is the empirically analyzed component; IoT sensing, sequential scheduling, federated learning, and farmer-facing deployment remain design-stage or future-validation elements.

Figure 1. Proposed Four-Layer IFRS Architecture




Fig. 1. Proposed four-layer IFRS architecture with an independent reference/validation layer.

A. Layer 1—IoT Sensing

Field sensors can capture soil N, P, K, pH, moisture, and temperature, together with weather information where available. Every measurement should carry a timestamp, unit, calibration status, quality flag, and missing-data indicator. This layer is proposed for future deployment and is not field-validated in the present study.

B. Layer 2—Edge Processing

Edge processing can perform unit normalization, plausibility checks, outlier detection, temporary storage, and communication management. The edge layer should not independently authorize fertilizer application; it should prepare trustworthy observations for the decision layer and provide safe fallback behavior during connectivity or sensor failures.

C. Layer 3—Explainable ML and Sequential Scheduling

The demonstrated component is the supervised fertilizer-class prediction module. Random Forest is retained for the explainability demonstration, while XGBoost is the strongest predictive benchmark in the present experiment. SHAP explanations can be attached to each prediction [12], [13].

A separate Q-learning scheduler is proposed for fertigation timing. Q-learning is a model-free reinforcement-learning method for sequential decision problems [15]. In IFRS, the state could include crop stage, soil moisture, N/P/K status, recent irrigation, weather forecast, cumulative nutrient application, and recent model recommendations. Actions could represent dose/timing choices. The reward should penalize nutrient surplus, water stress, excessive water use, cost, and deviation from agronomic constraints while rewarding acceptable crop-state trajectories. This scheduler is proposed and requires simulation and multi-season field validation.

D. Layer 4—Farmer Advisory

The advisory layer should convert a validated recommendation and its explanation into concise English, Hindi, and Marathi messages. A useful message should report the recommended action, principal evidence, confidence or uncertainty, and a safety note when sensor data are missing or outside the model's training domain. Its contribution should remain fertigation-specific explanation rather than general conversational agriculture.

E. Reference and Validation Layer

The reference layer provides an independent agronomic sanity check. The Government of India's Soil Health Card system provides soil-test and fertilizer-recommendation functions [16]. Tamil Nadu

Agricultural University's fertilizer-management guidance emphasizes using soil analysis and crop requirements when determining fertilizer use [17]. These sources should be treated as external references and constraints rather than as hidden labels used to train the classifier.

The architecture should therefore produce a traceable decision record: input data → model prediction → explanation → agronomic reference check → dose/timing decision → farmer-facing message. This audit trail is a more defensible systems contribution than presenting a high-accuracy classifier as a complete fertigation solution.

XI. DISCUSSION

The central lesson from the literature and applied analysis is that high classification accuracy is necessary but insufficient for fertigation decision support. A model can reproduce a historical fertilizer class accurately while still reproducing historical over- or under-application. The 97.34% XGBoost and 93.58% Random Forest values should therefore be interpreted as predictive performance for the stated classification task, not as agronomic decision accuracy.

The same caution applies to explainability. SHAP identifies how a model uses features; it does not establish that a feature causes crop response. In the present dataset, crop identity is the strongest predictor partly because the target label is crop-dependent. An expert validation layer is therefore essential.

The revised research gap is consequently narrower and more defensible than a claim of a completely new AI fertilizer recommender. The strongest contribution is the proposed integration of explainable classification, explicit separation of fertilizer identity from dose and timing, sequential scheduling, and independent agronomic benchmarking. The literature review does not establish that this exact combination is unprecedented; it establishes that these functions are commonly reported separately and that end-to-end field validation remains necessary.

Evaluation should also move beyond accuracy. Recommended computational metrics include macro-

F1, class-wise recall, calibration error, confusion patterns among similar fertilizer classes, explanation stability, and out-of-distribution performance. Agronomic metrics should include dose error, nutrient-use efficiency, water productivity, input cost, yield, and environmental indicators. Farmer-facing evaluation should measure comprehension and correct action rather than linguistic fluency alone.

XII. LIMITATIONS

- No controlled field trial or multi-season agronomic experiment has been conducted for the proposed IFRS.
- The applied ML results are based on historical/reported records and should be independently reproduced from the dataset and analysis code before publication.
- The fertilizer label is not an independently verified agronomic optimum; high predictive accuracy can therefore reproduce historical practice rather than optimal practice.
- Q-learning scheduling, IoT deployment, and federated learning are proposed/design-stage components and are not demonstrated by the reported classifier experiment.
- The literature search was targeted rather than a formal PRISMA-compliant systematic search; exhaustive novelty claims are therefore inappropriate.
- Multilingual advisory is not a strong standalone novelty claim because government and industry platforms already provide multilingual agricultural assistance.
- The manuscript does not establish causal relationships between SHAP feature importance and agronomic outcomes.

XIII. FUTURE RESEARCH

- Release a reproducible code archive containing preprocessing, fixed random seeds, model configurations, class distributions, confusion matrices, and SHAP settings.
- Evaluate models using repeated or nested cross-validation, class-wise metrics, calibration, and out-of-distribution detection.

- Construct a dose-level dataset containing crop stage, soil-test results, irrigation volume, fertilizer formulation, application timing, and observed crop response.
- Develop and validate the Q-learning scheduler using real multi-season fertigation trajectories and compare it with farmer practice, fixed schedules, and agronomic recommendations.
- Conduct controlled field trials across representative Western Maharashtra districts and report yield, nutrient-use efficiency, water productivity, cost, and environmental indicators.
- Validate the reference layer using current authoritative agronomic recommendations and Soil Health Card information.
- Evaluate multilingual recommendations with farmers and extension personnel for comprehension, trust, and correct action.
- Test federated learning across genuinely independent farm or institutional datasets with realistic non-IID distributions.

XIV. CONCLUSION

This paper presents a structured review and applied case study of AI-enabled fertilizer and fertigation recommendation for precision agriculture. The literature shows that RF/SHAP interpretability, IoT sensing, and multilingual advisory are established capabilities. The Western Maharashtra case study provides a transparent classical-ML baseline: XGBoost achieved 97.34% test accuracy and Random Forest 93.58%, while SHAP analysis identified crop identity, potassium, and nitrogen as the leading model predictors.

These results are useful as a computational benchmark but do not demonstrate agronomic optimality, nutrient savings, water savings, yield improvement, or farmer adoption. The revised research gap is therefore a systems and validation problem: future fertigation platforms should connect explainable prediction with explicit dose and timing logic, sequential scheduling, independent agronomic validation, reproducible evaluation, and field trials. The proposed IFRS formalizes this direction while clearly separating validated ML evidence from design-stage IoT, Q-learning, and federated-learning components.

REFERENCES

- [1] M. Bouni, B. Hssina, K. Douzi, and S. Douzi, "Interpretable Machine Learning Techniques for an Advanced Crop Recommendation Model," *Journal of Electrical and Computer Engineering*, vol. 2024, Art. no. 7405217, 2024. [Crossref](#)
- [2] D. G. Pai, M. Balachandra, and R. Kamath, "Explainable AI in agriculture: review of applications, methodologies, and future directions," *Engineering Research Express*, vol. 7, no. 3, Art. no. 032202, 2025. [Crossref](#)
- [3] S. M. Venkateswara and J. Padmanaban, "Interpretable deep learning models for independent fertilizer and crop recommendation," *Scientific Reports*, vol. 15, Art. no. 41721, 2025. [Crossref](#)
- [4] M. Sunandini, K. H. Sree, R. Deepiga, and A. Gokulapriya, "Smart soil fertilizer monitoring and crop recommendation system by using IoT and machine learning technology," *International Journal of Engineering Research & Technology*, vol. 12, no. 3, 2023. [Crossref](#)
- [5] R. Dembani, I. Karvelas, N. A. Akbar, S. Rizou, D. Tegolo, and S. Fountas, "Agricultural data privacy and federated learning: A review of challenges and opportunities," *Computers and Electronics in Agriculture*, vol. 232, Art. no. 110048, 2025. [Crossref](#)
- [6] Press Information Bureau, Government of India, Ministry of Agriculture & Farmers Welfare, "KISAN E-MITRA and IoT enabled systems to improve crop productivity," Apr. 1, 2025. [Online]. Available: <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2117392>
- [7] Digital Green and Gooley.AI, "Farmer.CHAT: An AI assistant for farmers," 2023. [Online]. Available: <https://www.help.gooley.ai/farmerchat>
- [8] Press Information Bureau, Government of India, Ministry of Agriculture & Farmers Welfare, "Record increase in food grain production, agricultural credit reaches ₹28.69 lakh crore, AI-based 'Bharat Vistar' platform launched for farmers," Mar. 12, 2026. [Online]. Available: <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2238885>
- [9] S. S. Gholap, "Crop and Fertilizer Dataset for Western Maharashtra," *Kaggle*, 2023. [Online]. Available: <https://www.kaggle.com/datasets/sanchitagholap/crop-and-fertilizer-dataset-for-westernmaharashtra>
- [10] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, 2001. [Crossref](#)
- [11] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794. [Crossref](#)
- [12] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30*, 2017, pp. 4765–4774.
- [13] S. M. Lundberg, G. Erion, H. Chen, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, "From local explanations to global understanding with explainable AI for trees," *Nature Machine Intelligence*, vol. 2, no. 1, pp. 56–67, 2020.
- [14] Ö. Turgut, İ. Kök, and S. Özdemir, "AgroXAI: Explainable AI-Driven Crop Recommendation System for Agriculture 4.0," in *Proc. 2024 IEEE International Conference on Big Data (BigData)*, 2024, pp. 7208–7217. [Crossref](#)
- [15] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3, pp. 279–292, 1992. [Crossref](#)
- [16] National Informatics Centre, Department of Agriculture & Farmers Welfare, Government of India, "Soil Health Card Portal," 2026. [Online]. Available: <https://soilhealth.dac.gov.in/>
- [17] Tamil Nadu Agricultural University, "Good Management Practices for Fertilizer Use," *Agritech Portal*. [Online]. Available: https://www.agritech.tnau.ac.in/gap_gmp_glp/gmp_fertilizer%20use.html