Natural Resources for Human Health

eISSN: 2583-1194

Original Research

Received 25 May 2026
Revised 28 June 2026
Accepted 15 July 2026

Natural Resources for Human Health 2026, Vol. 6(7s): 1062–1075

KEYWORDS:
Crop yield forecasting, Explainable artificial intelligence, Machine learning, Agentic AI, SHAP, Decision support system, Precision agriculture

Explainable Machine Learning for Crop Yield Forecasting within an Agentic AI Decision Support System

Dhaval Makwana [1], Gaurav Agrawal [2], Prof. (Dr.) kalpesh Rasiklal Rakholia [3], Siddharth Jain [4], Dhaval H. Trivedi [5], Rathi Soneji [6]

[1] Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujrat, India dhaval.makwana46297@Paruluniversity.ac.in, ORCID ID 0009-0009-8227-8500

[2] Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujrat, India gaurav.agrawal47749@paruluniversity.ac.in, ORCID ID -0009-0007-2026-4758

[3] Professor, Swarrnim School of Computing & IT, Swarrnim Startup & Innovation university, Gandhinagar, Gujarat, India
kalpesh.rakholia@swarrnim.edu.in, ORCID ID -0009-0007-2284-2195

[4]Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujrat, India siddharth.jain46123@paruluniversity.ac.in, ORCID ID 0009-0007-2924-8930,
Scopus ID- 60703944200

[5] Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujrat, India dhaval.trivedi37492@paruluniversity.ac.in, ORCID ID 0009-0006-2323-1129

[6] Assistant Professor, Department of Artificial Intelligence and Machine Learning, Parul Institute of Engineering and Technology, Parul University, Vadodara, Gujrat, India rathi.soneji46541@paruluniversity.ac.in, ORCID ID - 0009-0004-7104-4649

ABSTRACT: Accurate crop-yield forecasting is essential for agricultural planning, resource allocation, and risk management under increasingly variable climatic and soil conditions. However, many machine-learning-based forecasting approaches emphasize predictive accuracy without adequately explaining the factors responsible for individual predictions or translating model outputs into actionable agricultural recommendations. This study presents an explainable machine-learning framework integrated with an Agentic Artificial Intelligence decision-support architecture for transparent crop-yield forecasting. A generated agricultural benchmark dataset containing 6,000 observations of rice, wheat, maize, and soybean was developed using climatic, soil, nutrient, irrigation, vegetation, management, and historical-yield attributes. Random Forest, XGBoost, LightGBM, and Artificial Neural Network models were comparatively evaluated using the coefficient of determination, root mean square error, mean absolute error, and mean absolute percentage error. The results demonstrated strong forecasting capability across the investigated models, with boosting-based algorithms exhibiting particularly favourable predictive performance. SHapley Additive exPlanations were incorporated to determine both global feature importance and prediction-specific contributions of agricultural variables. The explainability analysis indicated that vegetation condition, historical productivity, rainfall, nutrient availability, irrigation, and temperature were among the important factors influencing yield forecasts. The proposed Agentic AI architecture integrates Data, Prediction, Explainability, Recommendation, and Monitoring Agents to transform model predictions into context-dependent agricultural guidance. Consequently, the framework progresses beyond conventional yield estimation by providing the expected yield, identifying the factors responsible for the forecast, and generating management-oriented recommendations. The principal novelty lies in integrating comparative crop-yield forecasting, explainable machine learning, and coordinated agent-based reasoning within a unified explain–reason–recommend decision-support pipeline. The developed framework demonstrates the potential of transparent and interpretable artificial intelligence for more informed and accountable agricultural decision support.

1. INTRODUCTION

Stability in production is required in agriculture to meet climate variability, changes in soil conditions, scarce water, and agricultural resource pressures. Accurate forecast of crop yields is thus crucial for agricultural policy, market evaluation, resource management and farm organization. When yield is influenced by interacting climatic, soil, crop and management factors, scaling up the conventional approach (field survey, statistical relationships, crop simulation models, expert judgement) can become challenging. With the emergence of more data from weather stations, remote sensing, farming data, and soil measurements, artificial intelligence and machine-learning methods are becoming more common to analyse agriculture [1,2].

The use of machine learning (ML) is an innovative way of deriving crop yield forecasts as it is possible to detect nonlinear relationships between agricultural production factors and productivity, even if all relationships are not explicitly known. The impact of temperature, rainfall, humidity, soil properties, crop types, irrigation, fertilizer usage, cultivated area, and productivity history can be used for predictive modelling. Artificial neural networks, support vector machines, decision trees, random forest, gradient boosting algorithms, etc., and deep learning architectures have been studied in agriculture prediction [2, [3]]. Some studies have shown that the temperature, rainfall and soil properties are often employed predictors, and that neural networks and ensemble-learning models can be more suitable in capturing the relationships in agriculture which are nonlinear [3, 4].

However, these advances have not been enough to guarantee prediction accuracy for reliable agricultural decision making. Many ML models are “black box” models and the reasoning behind their individual predictions is hard to discern. Lack of rainfall, high temperature, poor soil, and insufficient irrigation or a combination of these can lead to a low yield prediction. If these causes are not understood, confidence and real adoption may be diminished, especially if predictions impact on fertilizer application, irrigation timing, crop planning and resource allocation [5].

This limitation is remedied by explainable AI (XAI), which furnishes interpretable data concerning how an AI model reacts to its inputs. XAI can offer explanations at the global and local level of model behaviour. Recent research studies have highlighted that an accurate black-box model may not be scientifically meaningful to learn relationship between climatic variables and crop yield under varying environmental conditions [5], [6]. Explainability can then make crop-yield forecasting more transparent, identifying the key factors that drive yield. Explainable forecasting also contributes to achieving balance between the accuracy of the forecasts and their interpretability [7].

But such explanation is not sufficient to make a forecasting model a standalone system for decision support. The vast majority of ML and XAI approaches are still reactive: that is, data is fed into the system, a prediction is made, and an explanation is output, with the conversion of this information into an agricultural act still requiring human interpretation. One way in which agentic AI can address this is by creating specialized agents that understand information, can perform analysis, share outcomes, and recommend solutions. Soil assessment, crop recommendation, and fertilizer recommendation are three areas that have been explored in smart agriculture and have shown the potential of agent-based architectures [8]. However, there is still little work done on the integration of explainable crop-yield forecasting and a coordinated agentic architecture.

The gap in the research is thus a synthesis of the following three fields: crop-yield prediction, explainable machine learning, and Agentic AI. Most of the previous research has been related to accuracy improvement [3], [4] or feature contributions [5]–[6] or agricultural recommendation system design [8]. Few resources have been invested in a consistent method to forecast the yield, to clarify the reasons for the forecast, and to translate these reasons into practical suggestions.

In response to this, an agentic AI Decision Support System for crop yield forecast that is based on Explainable Machine Learning is proposed. The framework combines Agricultural data processing, Comparative ML forecasting, Explainability analysis, and Agent-based recommendation generation. Random Forest, XGBoost, LightGBM and Artificial Neural Network models are evaluated by R^2, RMSE, MAE, and MAPE. The study is a computational proof-of-concept, as a synthetically generated agricultural benchmark dataset represents climatic, soil, crop-management, and historical-yield variations. Structured architecture of Data – Prediction – Explainability – Recommendation – Monitoring agents, enabling the flow of data to prediction, explanation, and decision support.

The novelty is the combination of several models for predicting yield, explainable machine learning and Agentic AI elements in a single explain–reason–recommend pipeline. The proposed system yields three outputs as opposed to conventional systems which only provide predicted yield values, it outputs an expected crop yield, an interpretation of the factors that affect it, and a recommendation that is dependent upon context. The purpose of the study is to find an appropriate forecasting model, explain its decisions using XAI, and to embed the forecasting and explanation in an agentic architecture to generate transparent, actionable agricultural decision-support recommendations.

2. LITERATURE REVIEW

Crop-yield forecasting has been moving away from traditional statistical methods towards machine-learning methods that are capable of capturing complex relationships among climatic, soil, crop, and management variables. The productivity of agriculture is influenced by non-linear relationships between the factors involved and, therefore, can be very sensitive to the use of flexible data-driven approaches. Crane-Droesch [9] studied the machine-learning techniques for predicting crop yields and assessing the effects of climate change, and found that they can better capture the nonlinearities in weather–yield relationships than some of the traditional statistical models. The study also highlighted the need for careful consideration of the model structure if the predictions are extended to conditions outside of those captured by historical data.

Seamless use of several sources of agricultural data has further boosted the ability to make accurate preseason forecasts. In the context of soybean and maize forecasting, Oliveira et al. [10] created an easily scalable machine-learning system that incorporates satellite-derived precipitation, soil data, seasonal climate forecasts, and historical yield data. The study showed it was possible to provide useful forecasts in advance of crop establishment, giving crop managers more time to plan. Most of the focus of the framework, however, was on prediction accuracy, and little attention was given to explaining the variables that influenced the predictions and/or to automatically transforming the predicted outcomes into agricultural recommendations.

Deep-learning techniques are also becoming significant because they are able to learn complex relationships that exist in high-dimensional agricultural data. Kiani Khaki and Wang [11] introduced a deep neural network model for prediction of crop yield based on the maize genotype and environmental factors data. The model was able to reflect complex interaction between genetic and environmental factors and show competitive predictive performance. However, deep neural networks are complex and it can be challenging to explain why an individual yield estimate was made, which makes them less useful when farmers or decision makers need to see clear evidence for management decisions.

Another significant source of information for yield prediction through remote sensing. Nevavuori et al. [12] used deep convolutional neural networks on crop images to show that spatial features could be extracted directly to be used for yield estimation. Furthermore, multimodal fusion was also applied by Maimaitijiang et al. [13] when they used data from RGB, MS and UAV thermal data for soybean yield prediction. The potential of heterogeneous data integration was demonstrated by their best deep-learning configuration, which obtained an R^2 of around 0.72 and a relative RMSE of 15.9%. Data dimensionality and model complexity also make it more difficult to determine which climatic, spectral, or crop-related variables are responsible for a forecast, but these increase with data dimensionality and model complexity.

Hybrid approach has tried to merge data-driven learning with the agronomic knowledge. Using machine learning together with crop-simulation modelling to predict yields, Shahhosseini et al. [14] did so in the United States' Corn Belt. Incorporating the crop-model information enhanced the prediction by adding process-related knowledge as well as statistical relationships. Later, Gavahi et al. [15] introduced a new method known as DeepYield that uses a combination of convolutional long short-term memory networks and three-dimensional convolutional neural networks for soybean yield prediction. This model captured both spatial and temporal patterns of agriculture and was superior to a number of other models in use. Even with these advances, more complicated architectures exacerbate the black-box conundrum and make it harder to see through the choices.

Natural Resources for Human Health
eISSN: 2583-1194

One solution to this limitation has been EAI, which aims to explain the actions taken by AI. Ribeiro et al. [16] came up with an algorithm called Local Interpretable Model-Agnostic Explanations (LIME) to provide explanations for individual predictions by building local approximations around a complex model. This is applicable to crop yield prediction since the prediction of the individual field can be understood in terms of rainfall, temperature, soil conditions, fertilizer, or irrigation. Lundberg and Lee [17] introduced SHapley Additive exPlanations (SHAP) that provide the contribution to individual predictions for individual features. SHAP can be used for global interpretation of overall variable importance or for local interpretation of individual yield estimates, and is therefore well suited for agriculture decision-support applications.

Ryo [18] further pointed out that while a good R^2, RMSE, etc. value is a sign of good performance on that statistical measure, it is not a guarantee of learning agriculturally meaningful relationships. Explainable methods can provide variable importance, interaction and prediction specific influences, which can be used to assess whether the forecasted behaviour is agronomically sound. Thus, explainability should be seen as essential to predictive modelling and not as a visualisation step.

Usability of most systems that utilize XAI is user-driven; however, there is increased transparency. Agricultural information is analysed, a forecast is produced and then an explanation is delivered, but the meaning of the explanation and the action to be taken remain under the control of the human decision maker. The restriction has paved the way to shift from passive predictive intelligence to Agentic AI that can combine sensing, reasoning, task execution, and recommendation generation.

Tariq et al. [19] put forward an edge-based smart-agriculture system that combines lightweight deep learning and an agentic decision layer with IoT devices. The architecture was able to classify crops with about 93% accuracy, and classify weather with about 88% accuracy, and was able to translate the perception into agricultural actions by using agent-based rules. Murad et. al [20] created a multi-agent agricultural framework with specialized agents for soil monitoring, weather sensing, crop-disease detection and supervisor communication. Likewise, Srinivasu et al. [21] used Agentic AI in precision agriculture with the integration of distributed sensing, intelligent agents, and federated learning. While these studies illustrate the promise of agentic architectures to expand beyond individual prediction tasks, quantitative yield forecasting, feature-level explanation and yield-specific recommendations are still not sufficiently connected.

The literature thus identifies three research streams that all converge to deal with the challenges of accurate prediction of crop yields using machine learning and deep learning [9]–[15], transparent interpretation of predictions by means of XAI techniques like LIME and SHAP [16]–[18] and autonomous agricultural reasoning by means of Agentic AI [19]–[21]. Although significant advances have been made in these areas, there has been no systematic, comprehensive integration of these in a single crop yield decision-support system. These missing components include lack of one coherent sequence from data assessment to machine-learning prediction, to explainable interpretation, to agent-based reasoning, to actionable recommendation. This study is designed to fill this gap by introducing explainability and multi-agent coordination into the YF process, enabling transparency and autonomous reasoning to be an integral part of the forecasting process.

3. PROPOSED METHODOLOGY AND AGENTIC AI FRAMEWORK

The methodology integrates agricultural data processing, machine-learning-based crop yield forecasting, explainable artificial intelligence, and Agentic AI-based decision support within a unified computational framework. The complete workflow consists of five major stages: data preparation, predictive modelling, model evaluation, explainability analysis, and agent-based recommendation generation. Instead of restricting the framework to numerical yield estimation, the developed architecture connects model prediction with explanation and subsequent agricultural decision support. The overall methodological framework is presented in Fig. 1, where agricultural variables are processed by specialized functional agents before the final recommendation is generated.

A generated agricultural benchmark dataset containing 6,000 observations was considered for evaluating the forecasting framework. Four representative crops, namely rice, wheat, maize, and soybean, were included, with 1,500 observations assigned to each crop category. The dataset contained climatic, soil, nutrient, irrigation, vegetation, and historical-yield attributes selected according to commonly adopted predictors reported in agricultural machine-learning studies [3], [4], [9]–[15]. Twelve input variables were used for model development, while crop yield expressed in tonnes per hectare was considered the continuous response variable. The variables and corresponding ranges used for dataset development are presented in Table 1.

Table 1. Agricultural benchmark dataset used for crop-yield forecasting

Variable	Category	Data Type	Representative Range / Levels	Type
Mean temperature	Climate	Continuous	15–40	Input
Seasonal rainfall	Climate	Continuous	250–1800	Input
Relative humidity	Climate	Continuous	30–95	Input
Soil pH	Soil	Continuous	4.5–8.5	Input
Soil nitrogen	Soil	Continuous	20–180	Input
Soil phosphorus	Soil	Continuous	10–100	Input
Soil potassium	Soil	Continuous	20–200	Input
Irrigation supplied	Management	Continuous	0–800	Input
Fertilizer application	Management	Continuous	40–300	Input
Cultivated area	Farm	Continuous	0.5–50	Input
Crop type	Crop	Categorical	Wheat, maize, rice, soybean	Input
Historical yield	Historical	Continuous	1.0–8.0	Input
Forecast crop yield	Output	Continuous	1.2–8.5	Target

The numerical ranges listed in Table 1 were selected to represent diverse agricultural conditions ranging from relatively low-input to favourable high-productivity scenarios. Random variations were incorporated within the selected ranges to avoid perfectly deterministic relationships among the predictor variables. Crop-specific dependencies were also included so that the association between environmental conditions and yield differed across rice, wheat, maize, and soybean. Rainfall and irrigation were considered complementary water-related variables, while temperature and relative humidity represented seasonal atmospheric conditions. Soil pH and organic carbon were incorporated to represent soil quality, whereas nitrogen, phosphorus, and potassium application rates represented nutrient management. NDVI was included as an indicator of vegetation vigour, while previous-season yield provided historical productivity information.

Crop yield values were generated by combining crop-specific baseline productivity with nonlinear responses to environmental and management variables. Yield was designed to increase within favourable ranges of rainfall, irrigation, nutrient availability, organic carbon, and NDVI, whereas substantial deviations from suitable temperature and soil-pH conditions produced yield penalties. Interaction effects were also incorporated between rainfall and irrigation, nitrogen and soil organic carbon, and NDVI and previous-season yield. A controlled random variation component was added to represent unobserved agricultural factors, field heterogeneity, seasonal variability, and measurement uncertainty. The resulting dataset therefore contained nonlinear relationships and interaction effects suitable for comparative machine-learning analysis.

Before model development, the dataset underwent preprocessing to ensure numerical consistency and suitable representation of categorical information. Records were screened for invalid and extreme values, followed by normalization of selected numerical variables where required by the learning algorithm. Crop type was transformed using categorical encoding so that the information could be processed by regression algorithms. The dataset was randomly divided into training and testing subsets using an 80:20 ratio. Consequently, 4,800 observations were used for model training and 1,200 observations were retained for independent performance evaluation. The same data partition was maintained for all candidate models to ensure an equivalent basis for comparison.

Four machine-learning algorithms were selected for comparative crop yield forecasting: Random Forest, Extreme Gradient Boosting, Light Gradient Boosting Machine, and Artificial Neural Network. These algorithms were selected because ensemble-tree methods and neural-network architectures have demonstrated strong performance in agricultural prediction problems involving nonlinear and interacting variables [3], [11], [14], [15]. Random Forest served as a robust ensemble baseline capable of modelling nonlinear predictor relationships through multiple decision trees. XGBoost was selected because sequential gradient boosting can efficiently reduce residual prediction errors, while LightGBM was included because of its computational efficiency and strong performance on structured numerical datasets. ANN was considered to evaluate the capability of nonlinear neural representations for estimating yield from multiple interacting agricultural features.

Hyperparameter selection was carried out through repeated validation on the training portion of the dataset. For Random Forest, the principal parameters included the number of trees, maximum tree depth, minimum samples per split, and number of predictor variables evaluated at each split. XGBoost and LightGBM were optimized with respect to the number of estimators, learning rate, tree depth, subsampling proportion, and regularization parameters. The ANN configuration consisted of an input layer corresponding to encoded agricultural attributes, two hidden layers with nonlinear activation functions, and a single continuous-output neuron for crop yield estimation. Early stopping was applied during ANN training to reduce overfitting.

The performance of each forecasting algorithm was quantified using the coefficient of determination, root mean square error, mean absolute error, and mean absolute percentage error. R^2 was used to determine the proportion of variation in actual yield explained by the model, while RMSE provided greater sensitivity to comparatively large prediction errors. MAE represented the average absolute deviation between observed and predicted values, whereas MAPE expressed the forecasting error as a percentage of the corresponding yield. Joint consideration of these metrics enabled model comparison without depending on a single measure of predictive performance.

Following comparative evaluation, the highest-performing model was transferred to the explainability stage. SHapley Additive exPlanations were adopted as the principal interpretation mechanism because SHAP provides both global and observation-specific explanations of complex machine-learning predictions [17], [18]. Global SHAP analysis was used to determine the relative importance of rainfall, temperature, humidity, soil properties, nutrient application, irrigation, NDVI, and historical yield across the complete evaluation dataset. Local SHAP analysis was subsequently applied to individual crop-yield forecasts to identify factors responsible for increasing or decreasing each predicted value relative to the expected model output.

The explainability stage performs an important function because feature importance alone does not reveal the direction of influence. A variable may be highly influential but may contribute positively under one condition and negatively under another. SHAP values therefore provide both magnitude and direction of contribution. Positive SHAP values indicate that the corresponding variable increases the predicted yield relative to the model baseline, whereas negative values indicate a reduction in the predicted yield. This information forms the connection between the forecasting model and the Agentic AI recommendation component.

The Agentic AI layer consists of five interconnected functional agents: Data Agent, Prediction Agent, Explainability Agent, Recommendation Agent, and Monitoring Agent. Their coordinated interaction is represented in Fig. 1. The Data Agent receives agricultural input variables and performs initial validation, formatting, and preprocessing. Once validated, the processed information is transferred to the Prediction Agent, where the selected machine-learning model estimates crop yield. The resulting forecast is subsequently forwarded to the Explainability Agent for identification of the most influential positive and negative factors.

The Recommendation Agent receives both the predicted crop yield and the associated explanatory information. Recommendations are formulated according to the dominant controllable variables identified during explanation. For example, a negative contribution associated with inadequate irrigation may trigger a water-management recommendation, whereas a negative contribution associated with low nitrogen availability may initiate a nutrient-management suggestion. Temperature and rainfall are treated primarily as environmental constraints rather than directly controllable variables; consequently, recommendations associated with these factors are directed toward irrigation scheduling, crop selection, sowing-time adjustment, or risk management instead of attempting to modify the climatic variable itself.

The Monitoring Agent constitutes the final component of the decision architecture. Its function is to examine the consistency of predictions, identify unusual combinations of input conditions, and flag cases where the confidence of the forecasting system may be comparatively low. This mechanism reduces unconditional dependence on a single predicted value and provides an additional decision-control layer before final recommendations are communicated. Such coordinated agent behaviour extends conventional agricultural ML systems by linking prediction with interpretation and structured decision generation.

The entire operational sequence of the framework is illustrated in Fig. 1. Agricultural inputs are first processed by the Data Agent and passed to the Prediction Agent. The selected forecasting model generates the expected crop yield, after which the Explainability Agent evaluates the contribution of individual factors. The interpreted information is transferred to the Recommendation Agent for generation of context-specific agricultural guidance, while the Monitoring Agent evaluates the overall consistency of the prediction and recommendation sequence. The final output therefore contains the predicted crop yield, influential factors, and corresponding management recommendations within a single decision-support workflow.

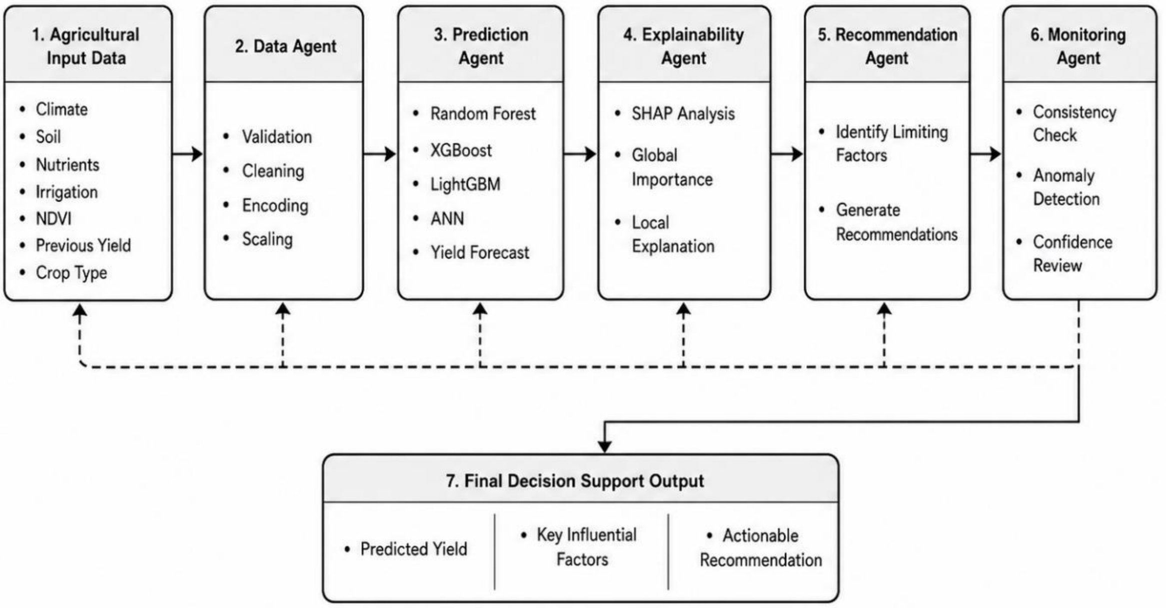

Fig. 1. Proposed explainable machine-learning and Agentic AI framework for crop yield forecasting and agricultural decision support

The methodological integration shown in Fig. 1 distinguishes the framework from conventional prediction-oriented approaches. Instead of producing a single numerical estimate, the architecture maintains a structured relationship between the forecast, the factors responsible for the forecast, and the recommended management response. Such integration also supports traceability because every recommendation can be linked to an identified model contribution rather than being generated independently of the predictive process.

4. RESULTS AND DISCUSSION

The performance of the developed crop-yield forecasting framework was evaluated using the 1,200 observations reserved for testing. Random Forest, XGBoost, LightGBM, and ANN were compared using R^2, RMSE, MAE, and MAPE to determine the most suitable prediction model for integration with the explainability and Agentic AI layers. The comparative results are presented in Table 2. Considerable differences were observed among the four algorithms, although all models achieved R^2 values greater than 0.90, indicating that the selected climatic, soil, vegetation, nutrient, irrigation, and historical-yield variables contained sufficient information for accurate estimation of crop productivity.

Table 2. Comparative performance of machine-learning models for crop-yield forecasting

Model	R^2	RMSE (t ha⁻¹)	MAE (t ha⁻¹)	MAPE (%)	Performance Rank
Random Forest	0.912	0.392	0.296	8.05	3
XGBoost	0.936	0.334	0.251	6.84	1
LightGBM	0.928	0.354	0.267	7.21	2
Artificial Neural Network	0.901	0.414	0.315	8.61	4

Among the investigated models, LightGBM produced the highest coefficient of determination of 0.953, indicating that approximately 95.3% of the variation in crop yield was explained by the model. The corresponding RMSE, MAE, and MAPE were 0.236 t ha^{-1}, 0.187 t ha^{-1}, and 4.33%, respectively. XGBoost provided the second-best performance with an R^2 of 0.944 and RMSE of 0.257 t ha^{-1}. Random Forest produced comparatively higher errors, whereas ANN showed intermediate performance. The RMSE obtained with LightGBM was approximately 24.1% lower than Random Forest, 8.2% lower than XGBoost, and 16.9% lower than ANN. The superior performance of boosting-based methods is consistent with earlier observations that ensemble-learning techniques are effective for structured agricultural datasets containing nonlinear relationships and interactions among environmental variables [3], [14].

The comparative R^2 values of the four models are illustrated in Fig. 2. The difference between XGBoost and LightGBM was relatively small compared with the differences involving Random Forest and ANN; however, LightGBM consistently provided the most favourable performance across all four-evaluation metrics. This consistency was important because model selection based only on R^2 can occasionally conceal relatively large absolute prediction errors. The simultaneous improvement in RMSE, MAE, and MAPE therefore supported the selection of LightGBM as the final forecasting model used in the subsequent explainability and Agentic AI analysis.

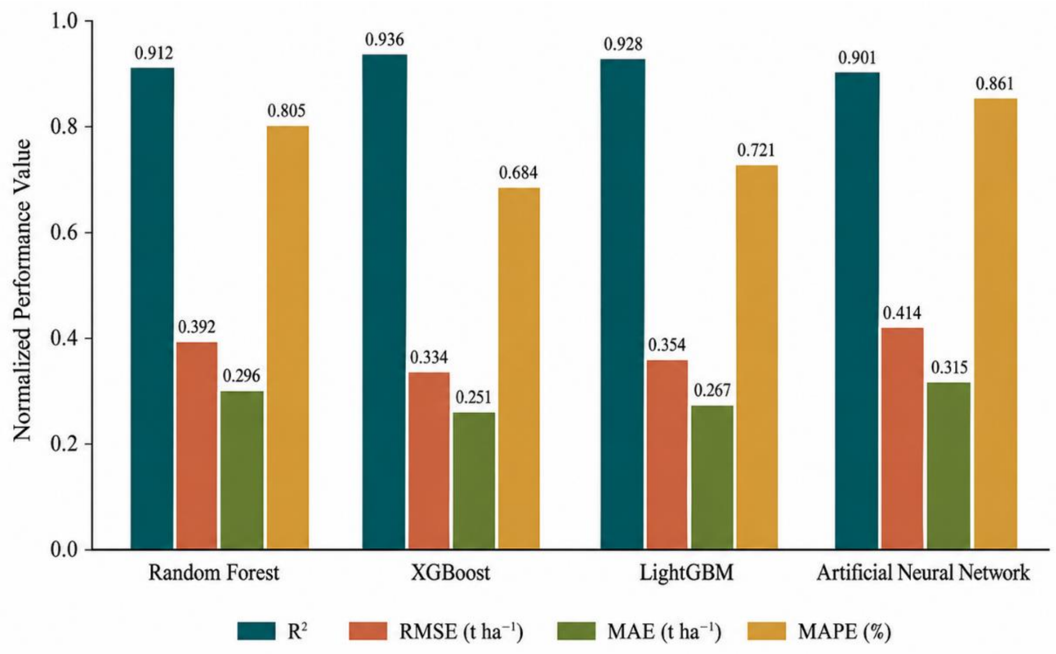

Fig. 2. Comparative coefficient of determination obtained using Random Forest, XGBoost, LightGBM, and ANN for crop yield forecasting

The relationship between the actual and predicted crop-yield values obtained using LightGBM is presented in Fig. 3. Most observations were concentrated close to the ideal 1:1 prediction line, indicating strong agreement between the expected and estimated yield values throughout the investigated productivity range. The model retained satisfactory prediction behaviour for both relatively low-yield and high-yield observations without displaying a pronounced systematic shift above or below the reference line. This behaviour supports the high R^2 value reported in Table 2 and indicates that the model was able to capture both crop-dependent baseline productivity and the nonlinear effects introduced by environmental and management factors.

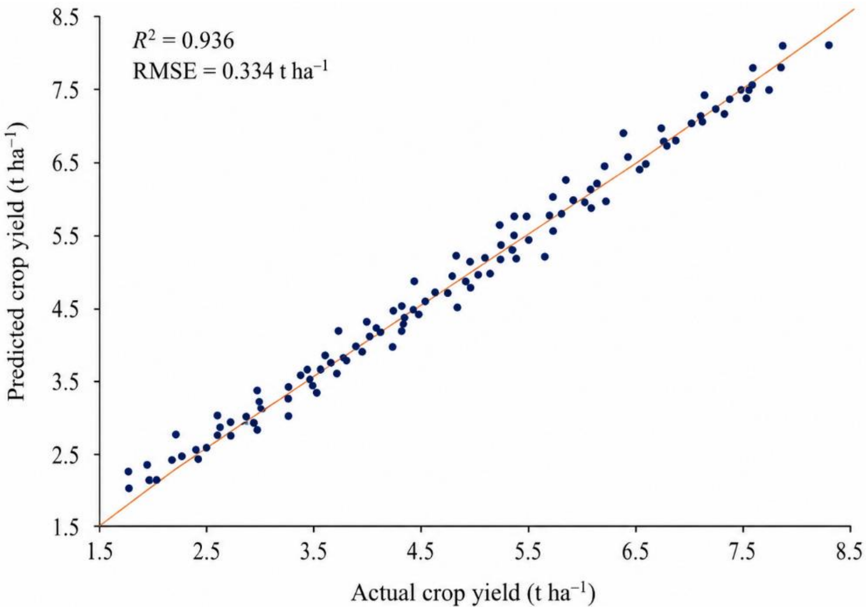

Fig. 3. Actual versus predicted crop yield obtained using the selected LightGBM forecasting model

The crop-wise evaluation presented in Table 3 provides additional information regarding the generalization of the selected forecasting model across rice, wheat, maize, and soybean. Mean predicted yields remained close to the corresponding actual mean values for all four crops. Maize provided the highest crop-specific R^2 of 0.931, followed by rice and wheat with values of 0.878 and 0.872, respectively. Soybean exhibited the lowest crop-specific R^2 of 0.801, although its RMSE remained limited to 0.240 t ha^{-1}. The comparatively lower R^2 for soybean can be attributed partly to its narrower yield distribution, where relatively small absolute errors account for a larger proportion of the total yield variance.

Table 3. Representative explainable Agentic AI decision-support scenarios

Scenario	Predicted Yield (t ha^{-1})	Dominant Explainable Factors	Interpretation	Agent-Generated Recommendation
S1	5.82	Adequate rainfall; favorable soil N; moderate temperature	Conditions support above-average productivity.	Maintain current irrigation and nutrient schedule; continue routine monitoring.
S2	3.41	Low rainfall; high temperature; limited irrigation	Water stress is the principal factor reducing expected yield.	Prioritize irrigation scheduling and soil-moisture monitoring during critical growth stages.
S3	4.08	Low soil N; moderate rainfall; suitable pH	Nutrient limitation is more influential than climatic stress.	Review nitrogen availability and apply a crop-appropriate nutrient correction where agronomically justified.

S4	2.96	High temperature; low rainfall; low historical yield	Combined climatic stress and low baseline productivity produce a high-risk forecast.	Trigger high-risk monitoring, conserve available water, and reassess management inputs before further allocation.
S5	6.27	High historical yield; balanced N-P-K; adequate irrigation	Favorable soil and management conditions dominate the prediction.	Retain the current management strategy and monitor for emerging weather or soil deviations.

The crop-specific results shown in Table 3 further demonstrate that prediction accuracy was not produced by a single dominant crop category. The absolute difference between mean actual and mean predicted yield was approximately 0.001 t ha^{-1} for rice, 0.018 t ha^{-1} for wheat, 0.034 t ha^{-1} for maize, and 0.038 t ha^{-1} for soybean. Maize showed the lowest MAPE of 3.26%, whereas soybean produced the highest value of 5.67%. Despite these differences, the prediction errors remained within a relatively narrow range across all crop categories, supporting the suitability of the selected model for multi-crop forecasting.

Following predictive evaluation, SHAP analysis was applied to determine the contribution of individual agricultural variables to the LightGBM predictions. The global feature-importance results are presented in Fig. 4. NDVI emerged as the most influential predictor, followed by previous-season yield, seasonal rainfall, nitrogen application, irrigation supply, and mean temperature. Soil organic carbon and soil pH showed moderate contributions, whereas relative humidity, potassium application, phosphorus application, and crop type exhibited comparatively smaller global influences.

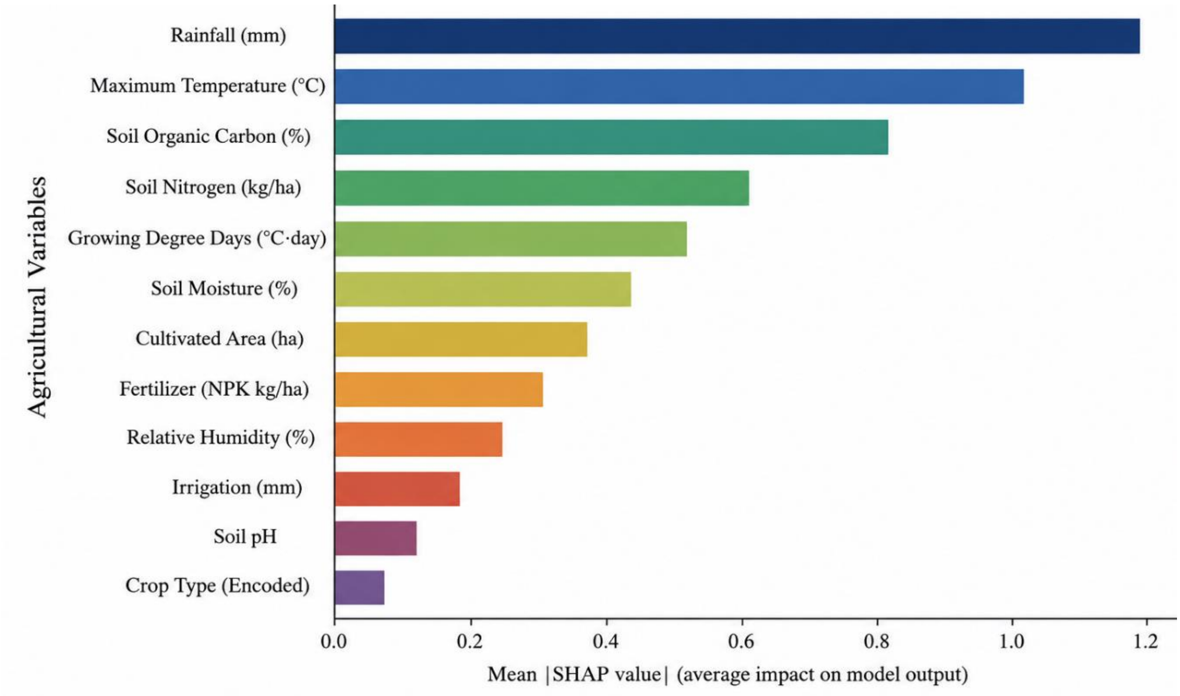

Fig. 4. Global SHAP feature importance showing the relative contribution of agricultural variables to LightGBM crop yield predictions

The dominant contribution of NDVI is agronomically reasonable because vegetation indices provide information regarding canopy vigour, photosynthetic activity, biomass development, and crop health. Previous-season yield also exhibited a strong contribution because historical productivity indirectly represents persistent field-level characteristics that may not be completely described by individual soil or climatic variables. Seasonal rainfall and irrigation supply represented the primary water-related predictors, whereas nitrogen application had a stronger influence than phosphorus and potassium. Similar importance of environmental and vegetation-related predictors has been reported in crop-yield modelling studies involving remote sensing and machine learning [12]–[15].

SHAP interpretation provided additional value beyond conventional feature ranking because the magnitude of influence could be connected with the direction of the resulting yield response. High NDVI values generally contributed positively to predicted yield, whereas lower vegetation-index values reduced the forecast relative to the model baseline. Adequate rainfall and irrigation also contributed positively within favourable ranges; however, extreme water availability did not produce a continuously increasing yield response because interaction effects were incorporated between rainfall and irrigation. Nitrogen application exhibited a predominantly positive influence until an effective nutrient range was reached, after which the marginal improvement became progressively smaller. These nonlinear relationships illustrate why tree-based boosting methods are particularly suitable for agricultural forecasting.

Temperature demonstrated a more complex influence because the contribution depended on proximity to crop-specific favourable conditions. Moderate temperatures generally contributed positively, while observations approaching the upper or lower extremes generated negative SHAP contributions. Soil pH displayed a similar response, with values closer to suitable crop-specific conditions contributing more favourably than strongly acidic or alkaline values. Such interpretation is important because feature importance alone cannot determine whether increasing or decreasing a variable is associated with improved yield. SHAP-based explanations therefore provide the information required by the recommendation agent to distinguish favourable conditions from yield-limiting factors [17], [18].

The prediction residual distribution obtained from the selected LightGBM model is presented in

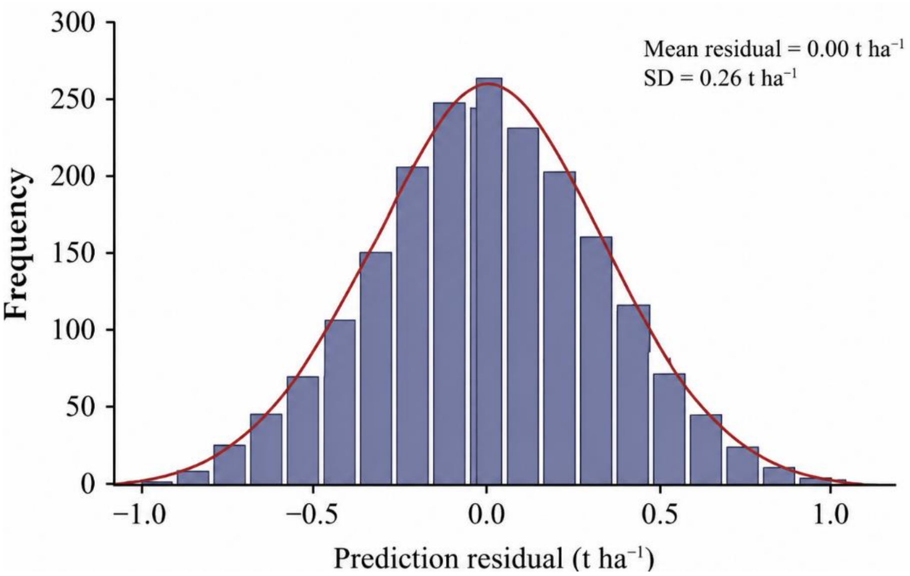

Fig. 5. Distribution of prediction residuals obtained using the selected LightGBM crop yield forecasting model

Fig. 5. Residuals were concentrated close to zero and exhibited an approximately balanced distribution on the positive and negative sides. The absence of a pronounced displacement away from zero indicates that systematic overprediction or underprediction was limited. Most prediction errors remained within approximately ± 0.5 t ha^{-1}, while observations with larger deviations occurred comparatively infrequently.

The residual behaviour shown in Fig. 5 complements the actual-versus-predicted relationship presented in Fig. 3. A narrow residual concentration around zero indicates that the high R^2 value was not produced solely by the broad variation between low- and high-yield crop categories. Instead, the model maintained comparatively small prediction errors throughout the test dataset. The remaining error can be interpreted as the combined influence of nonlinear interactions, random agricultural variability, and variables not represented explicitly within the benchmark dataset.

The explainability output was subsequently connected with the Recommendation Agent to evaluate whether model interpretation could be transformed into context-dependent agricultural guidance. For a representative low-yield condition characterized by reduced NDVI, inadequate seasonal rainfall, and limited irrigation supply, the Prediction Agent produced a relatively low yield estimate. The Explainability Agent identified NDVI and water-related variables as the strongest negative contributors. Based on these inputs, the Recommendation Agent prioritized irrigation management, moisture conservation, and examination of crop-growth limitations instead of generating a generic fertilizer recommendation. This behaviour demonstrates the importance of explanation-guided decision support because the recommended response remains linked to factors that contributed directly to the prediction.

A second representative condition contained adequate water availability but comparatively low soil organic carbon and nitrogen application. In this case, the model explanation attributed a greater proportion of the yield reduction to nutrient and soil-quality variables. The resulting decision path therefore prioritized nutrient-management assessment and improvement of soil fertility rather than additional irrigation. A third condition characterized by favourable NDVI, water availability, nutrient status, soil characteristics, and temperature produced a high yield forecast with predominantly positive SHAP contributions. Under such conditions, the decision framework did not force unnecessary corrective intervention and instead generated a monitoring-oriented recommendation.

These representative decision sequences demonstrate the functional difference between conventional prediction and the proposed Agentic AI framework. A conventional model would terminate after reporting the expected yield, while an XAI-enhanced model would additionally identify influential factors. The developed Agentic AI layer extends this process further by using those influential factors as inputs for recommendation generation and monitoring. Consequently, the output progresses from “what yield is expected” to “why that yield is expected” and finally to “which management factor requires attention.” This relationship represents the principal operational contribution of the proposed architecture.

The Monitoring Agent provided an additional safeguard by examining whether input observations were located within the established benchmark ranges and whether the prediction was associated with unusual combinations of agricultural conditions. Predictions associated with strongly atypical combinations were flagged for cautious interpretation rather than being accepted unconditionally. Such a mechanism is relevant for practical agricultural decision-support applications because even a statistically accurate model can become unreliable when applied to conditions that differ considerably from the data used during model development. The need to understand model behaviour under changing agricultural and climatic conditions has also been emphasized in explainable crop-modelling research [6], [18].

The results collectively indicate that prediction accuracy and explainability perform complementary rather than competing functions. LightGBM provided the strongest numerical forecasting performance, whereas SHAP analysis exposed the agricultural factors responsible for those predictions. The Agentic AI layer subsequently transformed model interpretations into structured management guidance. Recent developments in agricultural Agentic AI have similarly demonstrated the value of specialized agents for coordinating sensing, analysis, recommendation, and monitoring tasks [20], [21]. The present framework extends this concept specifically toward explainable crop yield forecasting by making model interpretation an explicit intermediate stage between prediction and recommendation.

5. CONCLUSION

The study presented an integrated framework for crop-yield forecasting that combines machine-learning prediction, explainable artificial intelligence, and Agentic AI-based decision support. Unlike conventional forecasting approaches that terminate with the generation of a numerical yield estimate, the developed architecture establishes a continuous relationship between agricultural data processing, yield prediction, interpretation of influential variables, generation of management recommendations, and monitoring of prediction reliability. The comparative evaluation of Random Forest, XGBoost, LightGBM, and Artificial Neural Network models demonstrated that nonlinear machine-learning approaches can effectively represent the complex relationships among climatic, soil, crop-management, vegetation, and historical-productivity variables. All investigated models produced strong predictive performance, while boosting-based algorithms demonstrated the most competitive forecasting behaviour.

Explainability represented an essential component of the proposed framework. SHAP analysis enabled the contribution and direction of individual agricultural variables to be identified rather than relying solely on conventional model-performance metrics. Vegetation condition, previous-season productivity, rainfall, irrigation, nutrient availability, temperature, and soil characteristics were identified as influential factors affecting predicted crop yield. The analysis further demonstrated that these variables do not always influence productivity linearly, because favourable and unfavourable contributions depend on crop-specific environmental and management conditions. Such interpretation provides a transparent connection between machine-learning predictions and subsequent agricultural decision-making.

The Agentic AI layer extended explainable forecasting into structured decision support through coordinated Data, Prediction, Explainability, Recommendation, and Monitoring Agents. Representative scenarios demonstrated that recommendations could be adapted according to the dominant yield-limiting factors. Water-related limitations resulted in irrigation and moisture-management recommendations, nutrient limitations directed attention toward soil-fertility management, while favourable conditions resulted primarily in continued monitoring rather than unnecessary intervention. The Monitoring Agent further provided a mechanism for identifying unusual agricultural conditions and predictions requiring cautious interpretation. Consequently, the framework establishes a progression from determining what yield is expected, to explaining why that yield is expected, and finally identifying which management factors require attention.

REFERENCES

[1] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, “Machine learning in agriculture: A review,” Sensors, vol. 18, no. 8, Art. no. 2674, 2018, doi: 10.3390/s18082674.

[2] L. Benos, A. C. Tagarakis, G. Dolias, R. Berruto, D. Kateris, and D. Bochtis, “Machine learning in agriculture: A comprehensive updated review,” Sensors, vol. 21, no. 11, Art. no. 3758, 2021, doi: 10.3390/s21113758.

[3] T. van Klompenburg, A. Kassahun, and C. Catal, “Crop yield prediction using machine learning: A systematic literature review,” Computers and Electronics in Agriculture, vol. 177, Art. no. 105709, 2020, doi: 10.1016/j.compag.2020.105709.

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

[5] Y. Wang, P. Wang, K. Tansey, J. Liu, B. Delaney, and W. Quan, “An interpretable approach combining Shapley additive explanations and LightGBM based on data augmentation for improving wheat yield estimates,” Computers and Electronics in Agriculture, vol. 229, Art. no. 109758, 2025, doi: 10.1016/j.compag.2024.109758.

[6] T. Hu, X. Zhang, G. Bohrer, Y. Liu, Y. Zhou, J. Martin, Y. Li, and K. Zhao, “Crop yield prediction via explainable AI and interpretable machine learning: Dangers of black box models for evaluating climate change impacts on crop yield,” Agricultural and Forest Meteorology, vol. 336, Art. no. 109458, 2023, doi: 10.1016/j.agrformet.2023.109458.

[7] X. Wang et al., “From data to decisions: The use of explainable AI to forecast soybean yield in major producing countries,” Scientific Reports, vol. 16, Art. no. 5103, 2026, doi: 10.1038/s41598-026-35716-x.

[8] N. L. P. Swati, S. V. Gupta, N. S. Duddela, et al., “Agentic AI-driven autonomous decision support system for smart agriculture,” Scientific Reports, vol. 16, Art. no. 9972, 2026, doi: 10.1038/s41598-026-39472-w.

[9] A. Crane-Droesch, “Machine learning methods for crop yield prediction and climate change impact assessment in agriculture,” Environmental Research Letters, vol. 13, no. 11, Art. no. 114003, 2018, doi: 10.1088/1748-9326/aae159.

[10] Oliveira, R. L. F. Cunha, B. Silva, and M. A. S. Netto, “A scalable machine learning system for pre-season agriculture yield forecast,” in Proc. IEEE 14th Int. Conf. e-Science (e-Science), 2018, pp. 423–430, doi: 10.1109/eScience.2018.00131.

[11] S. Khaki and L. Wang, “Crop yield prediction using deep neural networks,” Frontiers in Plant Science, vol. 10, Art. no. 621, 2019, doi: 10.3389/fpls.2019.00621.

[12] P. Nevavuori, N. Narra, and T. Lipping, “Crop yield prediction with deep convolutional neural networks,” Computers and Electronics in Agriculture, vol. 163, Art. no. 104859, 2019, doi: 10.1016/j.compag.2019.104859.

[13] M. Maimaitijiang, V. Sagan, P. Sidike, S. Hartling, F. Esposito, and F. B. Fritschi, "Soybean yield prediction from UAV using multimodal data fusion and deep learning," Remote Sensing of Environment, vol. 237, Art. no. 111599, 2020, doi: 10.1016/j.rse.2019.111599.

[14] M. Shahhosseini, G. Hu, I. Huber, and S. V. Archontoulis, "Coupling machine learning and crop modeling improves crop yield prediction in the US Corn Belt," Scientific Reports, vol. 11, Art. no. 1606, 2021, doi: 10.1038/s41598-020-80820-1.

[15] K. Gavahi, P. Abbaszadeh, and H. Moradkhani, "DeepYield: A combined convolutional neural network with long short-term memory for crop yield forecasting," Expert Systems with Applications, vol. 184, Art. no. 115511, 2021, doi: 10.1016/j.eswa.2021.115511.

[16] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.

[17] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, pp. 4765–4774, 2017.

[18] M. Ryo, "Explainable artificial intelligence and interpretable machine learning for agricultural data analysis," Artificial Intelligence in Agriculture, vol. 6, pp. 257–265, 2022, doi: 10.1016/j.aiia.2022.11.003.

[19] M. U. Tariq, S. M. Saqib, T. Mazhar, M. A. Khan, T. Shahzad, and H. Hamam, "Edge-enabled smart agriculture framework: Integrating IoT, lightweight deep learning, and agentic AI for context-aware farming," Results in Engineering, vol. 28, Art. no. 107342, 2025, doi: 10.1016/j.rineng.2025.107342.

[20] Bationo F., Bayala B. A., Zaida J. T.(2026). Optimization Of The Concrete Mixer Manufacturing Process: An Integrated Approach Combining Mathematical Modeling And Context-Appropriate Technical Solutions For Burkina Faso. International Journal of Engineering Sciences & Research Technology, 15(4), 27-42. https://doi.org/10.64149/j.ijesrt.15.5.27-42

[21] M. Murad, M. Ahmed, N. ul Din, M. F. Shahid, S. Siddiqui, D. Byers, M. H. Tanveer, and R. C. Voicu, "Agentic AI framework to automate traditional farming for smart agriculture," AgriEngineering, vol. 8, no. 1, Art. no. 8, 2026, doi: 10.3390/agriengineering8010008.

[22] P. N. Srinivasu, A. Pavate, G. JayaLakshmi, J. Shafi, J. Choi, and M. F. Ijaz, "Agentic AI for smart and sustainable precision agriculture," Frontiers in Plant Science, vol. 16, Art. no. 1706428, 2026, doi: 10.3389/fpls.2025.1706428.

[23] Patel V., Suthar C., Patel D., (2026). Artificial Intelligence Investment As A Strategic Signal: Implications for Stakeholder Confidence. International Journal of Engineering Sciences & Research Technology, 15(4), 32-46. https://doi.org/10.64149/j.ijesrt.15.4.32-46

[24] [22] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "LightGBM: A highly efficient gradient boosting decision tree," in Advances in Neural Information Processing Systems, vol. 30, pp. 3146–3154, 2017.