A Study of Explainable AI In Smart Agriculture

Lokesh Goregaonkar[1], Jay Kakade[2], Arnav Akhade[3], Shubhangi Gaikar[4]
[1,2,3]*Department of Science & Computer Science (AI/ML), MIT Arts, Commerce and Science College, Alandi(D), Pune, Maharashtra, India*
[4]*Assistant Professor, MIT Arts, Commerce and Science College, Alandi(D), Pune, Maharashtra, India*
doi.org/10.64643/ijirt.208800-459

***Abstract*—Artificial Intelligence is changing agriculture by helping farmers decide what to plant, when to water, and how to identify sick crops before they spread disease throughout a field. However, most AI systems work like "black boxes" they give answers without explaining how they arrived at them. For a farmer who relies on one harvest, a recommendation without a reason is hard to trust. This paper explores Explainable Artificial Intelligence (XAI) as the key to connecting powerful AI models with the people who need to act on their advice. We look at how techniques like SHAP, LIME, Grad-CAM, and attention mechanisms make complex AI predictions easier to understand for farmers. We also review how these methods are used in crop monitoring, disease detection, soil analysis, irrigation planning, and predicting yields. *The paper compares traditional AI with XAI*, and includes a case study on detecting diseases in tomato and apple leaves. This case study shows how a visual heat map can turn a simple label like "diseased" into a clear reason like "diseased because of the brown lesions on the leaf edges." The paper ends with a discussion on how XAI affects farmer trust and adoption, the challenges of using it in rural India, and the future of making agriculture both smart and easy to understand.**

***Index Terms*—Explainable Artificial Intelligence, XAI, Smart Agriculture, Precision Farming, SHAP, LIME, GradCAM, Crop Disease, Detection, Farmer Trust, Machine Learning**

I. INTRODUCTION

Agriculture is the foundation of India's economy and the main way of life for almost half of its people. For the past ten years, this old profession has started using a very new tool: Artificial Intelligence. Sensors check how wet the soil is, cameras look for the first sign of disease on leaves, and algorithms predict how much a field will produce before the harvest. These systems promise less uncertainty, less wasted water and fertilizer, and a better result for the farmer's work.

However, there's a hidden issue behind this promise. The deep learning models that run most modern AI systems are often called "black boxes." They take in data and give an answer, but the reasoning in between isn't visible even to the engineers who made them. When such a system tells a farmer to pull up an entire row of tomato plants or says not to water for another week, the farmer has to trust a decision they can't check. In a job where a wrong choice can mean losing an entire season's income, this lack of transparency isn't just a small problem it's a real obstacle to using AI.

Explainable Artificial Intelligence (XAI) is designed to fix this gap.
Instead of only giving a prediction, an XAI system also explains why: which part of the leaf was infected, which soil factor lowered the yield estimate, or which weather pattern caused a drought warning. This paper looks at how XAI is being used in smart farming, why it's important especially for a country with many small and poor farmers, and what still needs to be done before it becomes a common tool on farms rather than just an idea in labs.

II. LITERATURE REVIEW

Research on AI and agriculture has expanded quickly, but looking more closely at this area shows a clear change in focus: earlier studies were mostly concerned with how accurate the models were, while newer studies are more interested in whether that accuracy can be trusted [4], [5].
Traditional machine learning techniques like Random Forest, Support Vector Machines, and Decision Trees were some of the first used for tasks such as crop recommendations, soil analysis, and disease detection. This is because their decision-making process is easier

to understand. However, as deep learning models especially Convolutional Neural Networks (CNNs) for image-related tasks and Long Short-Term Memory (LSTM) networks for time-based data such as rainfall and crop yield became more popular due to their better accuracy, they also brought the problem of being hard to explain, which is the main reason for this paper.
Reviews in the field consistently show that a popular approach now is to combine high-accuracy deep learning models with layers that explain their decisions, like SHAP, LIME, or Grad-CAM [4], [5], [8].

For example, in studies about plant disease detection, once a Grad-CAM heat map [3] is shown on a leaf image, agronomists can see if the model focused on the actual disease or was distracted by other elements like soil or shadows [14], [15]. Similarly, in crop yield predictions, SHAP-based analysis [2] often highlights rainfall, vegetation health indicators such as NDVI, and soil moisture as the main factors affecting yields [11], [12], [13]. These findings match what agronomists already know, which helps build trust that the model is making logical decisions, not just finding patterns in the data by chance.

At the same time, the literature is candid about the gaps that remain. Most published systems are still trained and tested on curated laboratory datasets (such as Plant Village [9], [10]) rather than messy, real-field images with variable lighting and backgrounds, which means reported accuracies do not always transfer to real farms. Very few studies actually test whether a farmer, rather than a computer scientist, can understand the explanation being generated. This paper takes that observation as its starting point and treats farmer-level interpretability, not just model accuracy, as the true measure of success for agricultural AI.

III. UNDERSTANDING EXPLAINABLE AI

Explainable AI isn't just one algorithm; it's a group of methods meant to address three basic questions that anyone using an AI system will eventually wonder about: What did the model come up with? Why did it come up with that? And can I rely on it? [8]. Below are some of the most commonly used techniques in the agricultural field.

3.1. SHAP (SHapley Additive exPlanations)

SHAP comes from game theory and gives each input factor, like rainfall, nitrogen levels, or temperature, a number called a "credit score" that shows how much it influenced the model's prediction [2]. It's especially useful for models that predict crop yields or analyze soil data, as it clearly shows which factors had the biggest impact [11], [12], [13].

3.2. LIME (Local Interpretable Model-Agnostic Explanations)

LIME explains each prediction individually. It makes small changes to the input, like covering small parts of a leaf in an image, and checks how the prediction changes. Then it creates a simple, easy-to-understand summary of how the model behaves locally. This is helpful for a farmer who wants to know why a specific plant was marked, rather than why the whole crop was flagged.

3.3. Grad-CAM (Gradient-weighted Class Activation Mapping)

Grad-CAM is designed especially for image-focused CNN models [3]. It creates a color-coded heat map over the input image, showing precisely which pixels had the biggest impact on the diagnosis. For detecting crop diseases, this is very straightforward the part of the leaf that's affected literally shines up on the screen [14], [15].

3.4. Attention Mechanisms and Counterfactual Explanation

Attention-based models, which are often used in newer transformer architectures, show which part of the data the model paid the most attention to, like a specific time period or area. This helps in understanding different stages of crop growth [8]. Counterfactual explanations take this even further by answering questions like, "What if the soil moisture was 10% higher?
Would the yield prediction have gotten better?" This turns a problem diagnosis into clear steps you can take to improve outcomes [8].

IV. AI VERSUS XAI: A COMPARATIVE VIEW

Table 1 shows the differences between a regular AI system and one that provides explanations, especially from the perspective of a farmer [8].

Table 1: Comparison between Conventional AI and Explainable AI in Smart Agriculture

Aspect	Conventional AI (Black-Box)	Explainable AI (XAI)
Nature of the Output	Gives only a final decision, for example, "crop diseased."	Gives the decision along with the reasoning, for example, "diseased because of the lesion pattern on the leaf edge."
Transparency	The inner workings are kept secret even from the developers themselves.	Internal reasoning is shown through feature importance or visual heat maps.
Farmer Trust	Low, since advice cannot be verified against field experience.	High, since the explanation can be cross-checked against what the farmer physically observes.
Error Detection	Wrong predictions are hard to diagnose or debug.	Errors can be traced to the specific feature or region that misled the model.
Accuracy focus	Optimized purely for prediction accuracy.	Balances accuracy with interpretability, sometimes at a small accuracy trade-off.
Suitability for critical decisions	Risky for high-stakes actions like uprooting crops or large fertilizer expenditure.	Preferred for high-stakes decisions since reasoning can be audited before acting.
Adoption in rural settings	Often resisted due to "it just says so" opacity.	Better suited to build long-term trust and voluntary adoption.
Regulatory / Policy readiness	Difficult to justify in insurance, subsidy, or compliance contexts.	Easier to justify decisions to banks, insurers, and government schemes.

V. WORKFLOW OF AN XAI-ENABLED SMART AGRICULTURE SYSTEM

Figure 1 shows the overall process of turning raw data from the field into a clear answer that a farmer can use. The process starts right when the data is gathered from the field and finishes only when the farmer gets a decision that makes sense to them.

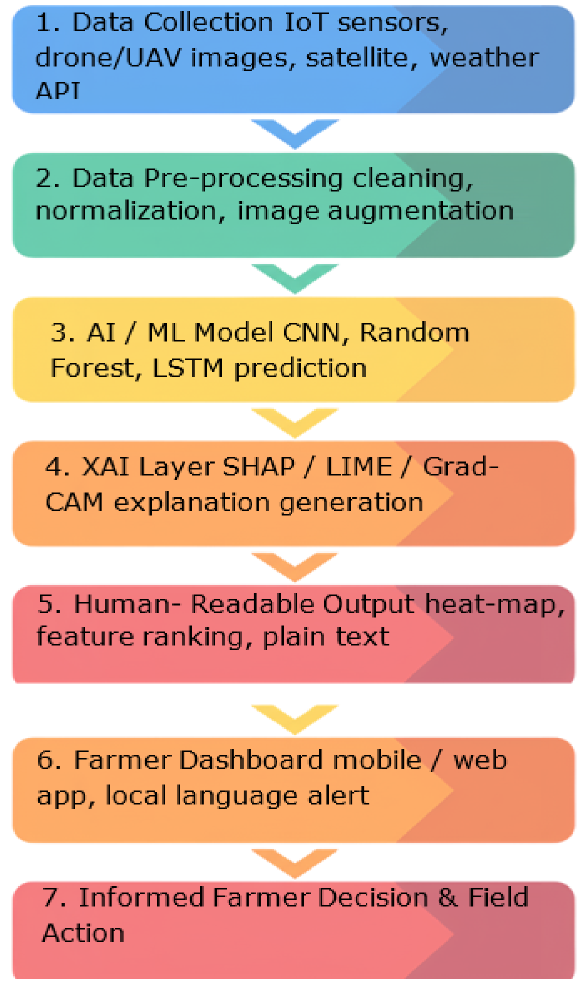

Each step has a special job.
The first step collects raw information from the field, like soil moisture levels, satellite pictures, or a photo of a leaf taken by a farmer's phone.
The second step prepares and makes this data consistent so the model can use it properly.
The third step is where the model actually makes predictions for example, a CNN might tell if a leaf is sick, or a Random Forest might estimate this season's crop output.
Importantly, the fourth step doesn't just stop at the prediction; it uses a special method to explain why the model came to that conclusion [1], [2], [3].
The fifth step turns that explanation into something easy to understand maybe a colored heat map or a short list of the main factors. Finally, the sixth step shares this explanation through a way a farmer can use it, like a mobile app in their local language, which leads to the seventh step: making a decision with confidence instead of just guessing.

VI. APPLICATIONS OF XAI IN SMART AGRICULTURE

6.1. Crop Disease and Pest Detection

This is the most advanced use of the application. CNN models combined with Grad-CAM or LIME show the exact spot on a leaf image where there is a lesion, discoloration, or damage from pests, allowing a farmer to see the problem visually instead of just getting a simple label [3], [14], [15].

6.2. Soil Health and Fertility Analysis

Using SHAP-based models on soil-sensor data helps identify exactly which nutrient (like nitrogen, phosphorus, or potassium), pH level, or moisture level is influencing a fertility recommendation. This makes it possible to avoid using too much fertilizer, which would be a waste.

6.3. Irrigation and Water Management

Explainability techniques help farmers understand if an irrigation recommendation is based on soil moisture levels, predicted rain, or increasing temperatures, so they can check if the advice matches what they observe in their fields [13].

6.4. Crop Yield Forecasting

Studies show that rainfall, plant health indicators, and growing-degree days are the main factors affecting crop predictions. This helps farmers and decisionmakers plan better by understanding the real reasons behind the predictions, instead of just relying on unclear numbers [2], [11], [12], [13].

6.5. Climate Risk and Drought Early Warning

By showing how much each factor like temperature changes, less rainfall, and plant stress affects drought, XAI-based models help tell the difference between a real drought happening and a usual temporary dry period.

VII. CASE STUDY: EXPLAINABLE CROP DISEASE DETENTION

To make the discussion clearer, this section uses a simple example based on a widely known problem: detecting disease in tomato and apple leaves with a CNN model that is made easier to understand using Grad-CAM and LIME.

7.1. Problem

A farmer sees part of the tomato plants looking sickly but isn't sure if it's because of a fungal disease, a lack of nutrients, or just too much sun. If it's a fungal infection and goes unnoticed, it can spread quickly and harm the nearby plants in just a few days.

7.2. Approach

A picture of a leaf, taken with a regular smartphone, is sent to a CNN model like a fine-tuned ResNet or EfficientNet. This model is trained to tell if the leaf is healthy or has a certain disease [9], [16]. Instead of just giving a simple result, the system also uses two methods to explain its decision. First, it uses Grad-CAM, which creates a heat map over the leaf image.

This heat map shows the exact area, like a red spot, that made the model think there was a problem [3], [16].

Then, it uses LIME, which points out the parts of the image that had the most influence on the decision. This helps check and confirm the result from Grad-CAM [1], [15].

7.3. Outcome

Instead of just getting a simple message like "Early Blight Detected," the farmer's screen shows a picture of the leaf with the infected part clearly marked. Alongside the image, there's a brief explanation like: "Diagnosis: Early Blight. Reason: A brown ring pattern was found near the leaf edge, which matches the look of a fungal infection." This small change makes a big difference. The farmer isn't just told to trust the system—they can actually see the proof, and compare it to the leaf they have. In real situations, this kind of clear, visual evidence has made farmers much more likely to follow the advice. For example, they might use a fungicide only on the damaged area instead of covering the whole field, which saves costs and helps protect the environment.

7.4. Why This Matters

This case study shows the main point of the paper in a smaller way: the benefit of XAI in agriculture isn't about making models more precise, but about making an already good model easy to use. A black-box model that's 95% accurate but a farmer doesn't trust isn't really helpful in real life compared to a model that's 90% accurate and the farmer can check how it works themselves [8].

VIII. REAL-WORLD IMPACT ON FARMERS

The advantages of XAI in agriculture go beyond just technology—they bring real results for the people who grow food.

When a suggestion comes with a clear explanation, farmers are more likely to follow it instead of ignoring it as just a machine's guess [4], [5].

Understanding the reason behind a recommendation helps prevent farmers from using too much fertilizer or water just to be safe, which saves money and protects the environment [2].

Farmers can make quicker and more confident choices when they see a visual heat map showing a diseased leaf, even if they have little formal education. This is easier to understand than a number that doesn't mean much without extra information [3].

Clear and explainable risk assessments help banks and insurance companies make better decisions, which in turn helps farmers get loans or payouts [8].

XAI systems can also help farmers who don't read much or speak technical English, as explanations can be shown through pictures, colors, or voice.

Instead of replacing a farmer's experience, XAI gives them a chance to compare machine suggestions with their own knowledge, keeping the farmer in control of the decision.

IX. CHALLENGES AND LIMITATIONS

Despite its promise, deploying XAI across Indian agriculture faces several real obstacles. Computational techniques such as SHAP can be slow, which is a problem for real-time applications like drone-based field scanning [2]. Most disease-recognition models are still trained on clean, laboratory-style datasets rather than the cluttered, poorly-lit images typical of real farms [9], [10], so accuracy and consequently the reliability of the explanation can drop noticeably in the field. Rural connectivity and the cost of AI-enabled hardware remain barriers for small and marginal farmers.

Finally, very little research has actually tested whether the explanations generated by SHAP, LIME, or Grad-CAM are genuinely understandable to a farmer rather than only to a data scientist [4], [5], which is arguably the single biggest gap standing between today's research and tomorrow's usable technology.

X. FUTURE SCOPE

Looking ahead, XAI in agriculture is expected to focus on lightweight, on-device explainability that works even on cheap smartphones without needing a constant internet connection [16]. There will also be voice-based explanations in local languages so farmers who prefer speaking over reading can understand better.

Additionally, hybrid systems that mix data from IoT sensors with drone images will allow explanations to combine information about soil, weather, and visual details all at once [13]. It will also be important for AI researchers and agricultural scientists to work closely together to ensure that what a model considers important really reflects real farming knowledge, not just random patterns in the data [11], [12].

XI. CONCLUSION

Artificial intelligence has already shown it can predict crop diseases, forecast how much crops will grow, and help with watering in a very accurate way [9], [11]. But it hasn't fully shown that farmers can trust it. This paper says that Explainable AI isn't just an extra feature for smart farming—it's a key part that's missing. It's the part that turns AI's silent answers into clear explanations that farmers can understand [8]. By looking at current methods, comparing regular AI with Explainable AI, showing a real example of disease detection, and honestly looking at what works and what doesn't in real life, this study finds that the future of AI in Indian agriculture depends not just on how accurate these systems get, but on how well they can explain their decisions to the farmers who need to act on them.

REFERENCES

[1] 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.

[2] S. Pingale, S. Shevkari, A. Jaiswal, and P. Shinde, "Gen AI-enabled personalized learning experience in higher education," International Journal of Innovative Research in Technology, vol. 12, no. 5, 2025.

[3] S. Devikar, M. Hinge, and S. S. Shevkari, "Human vs. machine: A review of AI's impact on language learning interaction or replacing human interaction," International Journal for Multidisciplinary Research, vol. 7, no. 3, 2025, doi: 10.36948/ijfmr. 2025.v07i03.49459.

[4] R. M. Abou Samra and R. R. Ali, "Explainable artificial intelligence in smart agriculture: A comprehensive review of interpretable remote sensing for sustainable decision-making," AgriEngineering, vol. 8, 2026.

[5] M. Kalyani, M. Manaswini, S. Shevkari, J. Sinkar, L. K. Tyagi, and Y. K. Sharma, "LeafNet: An intelligent system for plant disease classification with deep learning models," in Proc. 2026 13th Int. Conf. Computing for Sustainable Global Development (INDIACom), New Delhi, India, 2026, pp. 1–6, doi: 10.23919/INDIACom70271.2026.11525518.

[6] N. D. Indira, M. R., N. K. R., P. M. C., and P. S. Sankanagowda, "Smart agriculture using AI: An integrated approach for soil analysis, irrigation, crop-fertilizer recommendations, and market price analysis," in Proc. ACSCON, International Journal of Engineering Research & Technology (IJERT), 2026.

[7] S. B. Pawar, "Artificial intelligence in agriculture," in Proc. ICSITS, International Journal of Engineering Research & Technology (IJERT), 2020.

[8] A. B. Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020.

[9] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," Frontiers in Plant Science, vol. 7, Art. no. 1419, 2016.

[10] S. Shevkari, S. V. Choudhari, and A. N. Tonpe, "Analysis of implementing green accounting (GA) in India—Need of the modern era," International Journal of Scientific Development and Research, vol. 10, no. 11, 2025.

[11] "Explainable AI-based crop recommendation systems for improving yield predictions," Springer Nature, 2026.

[12] "Next-gen agriculture: Integrating AI and XAI for precision crop yield predictions," Frontiers in Plant Science, vol. 8, Art. no. 1451607, 2024.

[13] "Opening the black-box: A systematic review on explainable AI in remote sensing," arXiv:2402.13791, 2024.

[14] "A lightweight and explainable CNN model for empowering plant disease diagnosis," Scientific Reports, vol. 15, 2025.

[15] M. R. Hoque et al., "An explainable and lightweight CNN framework for robust potato leaf disease classification using Grad-CAM visualization," Applied AI Letters, 2026.

[16] "Mulberry leaf disease detection by CNN-ViT with XAI integration," *PMC*, 2025.