IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

Green AI for Sustainable Agriculture: An Explainable-AI (XAI) Perspective on Energy-Efficient Crop Intelligence

Aditya Singh Baliyan[1], Akhilesh Pandey[2], Payal Manwal[3], Gautam Taygi[4]
Assistant Professor, School of Computer Science and Engineering[1,4]
Assistant Professor, School of Computer Applications[2,3]
Tulas University, Dehradun, Uttarakhand, India

Abstract : *In modern agriculture precision forecast crop yields, diagnose plant disease and schedule irrigation systems rely on machine learning (ML) and deep learning (DL). Cost of training and deploying these models is weighed their accuracy. Two research threads have far developed largely in parallel. Green AI, which treats computational and energy efficiency as a first-class evaluation criterion alongside accuracy and Explainable AI(XAI) uses technology like Shapley additive explanations(SHAP), LIME and Grad-CAM to open the "black box" models so that farmers , and policymakers can act on their predictions.This paper synthesises recent literature (2017–2026) at the intersection of these threads as they apply to agriculture, covering crop-yield prediction, plant-disease detection, and soil- and irrigation-modelling studies that report either energy-efficiency measures or explainability analyses, or both. The synthesis identifies a specific and consistent gap: while explainability is now common in high-performing agricultural ML studies, almost none of them report the energy or carbon cost of training the underlying models, and the one purpose-built Green-AI agriculture framework identified in the review explicitly leaves explainability as future work rather than integrating it. Building on this synthesis, the paper proposes a unified Green-XAI methodological framework for agricultural ML that couples resource-appropriate model families – tree-based ensembles such as Random Forest, XGBoost and LightGBM, alongside a compact CNN architecture for imagery – with SHAP-based interpretability and continuous energy/carbon instrumentation across the training and inference lifecycle. The paper details the proposed dataset scope, the governing model equations, the SHAP computation, and an accuracy-versus-energy evaluation protocol, and positions the framework against a comparative synthesis of the performance figures reported by prior studies.*

Keywords: Green AI, XAI, Shap, LIME, CNN, Agriculture, Forecast, explainable, Random Forest

I. INTRODUCTION

Agriculture sits at the intersection of two pressures that are separately discussed. It feed a global population of roughly ten billion people within three decades, and it also do so while facing water scarcity, soil degradation, and an increasingly climate change [6]. Artificial intelligence (AI) and machine learning (ML) are proposed as part of the solution such as precision irrigation, yield forecasting, disease detection, and resource-efficient farm management [7]. The computational techniques that efficiency gains in the field carry a hidden efficiency cost of training and running ML and deep-learning (DL) models consumes electricity, and scale this consumption has a measurable carbon footprint.

The computational cost has become its own field of study under the banner of Green AI. [1] contrasted "Red AI" research, which chases state-of-the-art accuracy through ever-larger models regardless of computational cost, with "Green AI," which treats efficiency measured in floating-point operations, carbon emitted, or a comparable "price

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

tag"as an evaluation criterion on par with accuracy. Their position paper followed empirical work [2], who estimated that training a single large NLP model with neural architecture search could emit carbon dioxide comparable to the lifetime emissions of several cars, a finding that sharpened concern about the environmental cost of AI research generally. Subsequent work [3] surveyed the resulting body of "green-in-AI" techniques and catalogued the tools now available to measure that footprint, such as CodeCarbon and comparable carbon-tracking packages [28].

Agriculture is an explicit target of this second, "green-by-AI" category[5] recently proposed a Green AI framework for smart agriculture built around lightweight model architectures, model compression, federated learning, and edge computing, aimed specifically at reducing the carbon footprint of predictive farming systems. Their framework is notable for what it deliberately leaves out: the authors list the integration of explainable AI among the topics for future work rather than as part of the proposed system [5]. This omission is not unique to that paper; as Section 2 and Section 4 show, it reflects a broader pattern across the agricultural-AI literature that motivates the present review.

The second thread this paper addresses is explainability. As agricultural ML models have grown more accurate, they have also grown more opaque: gradient-boosted ensembles and convolutional neural networks routinely reach accuracy or R^2 figures above 0.9 on yield, soil, and disease-classification tasks (Section 4), but their internal reasoning is not directly inspectable by the agronomists, extension officers, or farmers who must act on their outputs. Explainable AI (XAI) methods address this opacity. SHapley Additive explanations (SHAP), introduced by Lundberg and Lee [4], adapt the Shapley value from cooperative game theory to attribute each prediction to the features that produced it, and has become the dominant XAI technique for tabular agricultural data, while Gradient-weighted Class Activation Mapping (Grad-CAM) plays the equivalent role for image-based disease-detection models. Recent reviews confirm that XAI adoption in agriculture has accelerated sharply since 2020 [8], [9], [10], spanning crop recommendation, yield forecasting, disease diagnosis, and precision-farming decision support.

The problem this paper addresses is the disconnect between these two threads. Green AI research in agriculture optimises for energy and carbon efficiency but, as illustrated framework [5], has not yet incorporated explainability as a design requirement. XAI research in agriculture, conversely, has matured rapidly in its ability to explain highly accurate models but – as the synthesis in Section 4 shows – almost never reports the energy, carbon, or compute cost of producing those explanations or the models that underlie them.

Based on the literature review identifies four specific research gaps:

- Most published XAIinagriculture studies report predictive performance and explanation quality but not energy consumption, training time on specified hardware, or carbon emissions, making it impossible to compare the true resource cost of competing "explainable" approaches.
- Model-family imbalance: image-based disease-detection studies default to large, pretrained CNN backbones (e.g., ResNet152, EfficientNetB7) explained via Grad-CAM, with comparatively little benchmarking of lighter-weight architectures against these backbones on an accuracy-per-joule basis.
- Explanation-cost gap: SHAP computation, particularly on large tabular datasets, is itself computationally non-trivial, yet its marginal energy cost relative to model training is essentially unreported in the agricultural-AI literature reviewed here.
- Actionability gap: SHAP and Grad-CAM outputs typically explain a prediction (e.g., which feature drove a yield estimate) but are rarely connected back to a resource-efficiency decision (e.g., reducing fertiliser or water application), leaving a gap between being 'explainable' and being 'actionable for sustainability.'

Against this background, the present paper makes three contributions. First, it synthesises recent literature at the intersection of Green AI, XAI, and agricultural ML into a structured comparison (Section 2) covering crop-yield prediction, plant-disease detection, and soil/irrigation modelling. Second, it proposes a unified Green-XAI methodological framework (Section 3) that pairs resource-appropriate model families with SHAP-based interpretability and continuous energy/carbon instrumentation across the training and inference lifecycle, directly addressing RG1–RG4. Third, it positions this proposed framework against the reviewed literature (Section 4) and sets out a concrete empirical validation plan (Section 5). The remainder of the paper is organised as follows: Section 2 reviews the related

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

literature and presents the table underlying the gap analysis above; Section 3 details the proposed methodology; Section 4 presents the comparative synthesis, the reporting-gap analysis, and the framework positioning; Section 5 outlines the planned empirical work; and Section 6 concludes.

II. LITERATURE REVIEW

This section synthesises the literature underlying the research gaps identified in Section 1. Studies through academic literature , prioritising work that published between year 2017 and 2026 that combines a machine-learning or deep-learning models used in agricultural task with an explicit explainability component (SHAP, LIME, Grad-CAM), supplemented by the foundational Green AI and XAI references that motivate the framework and by broader XAI-in-agriculture reviews.

No.	Study (Authors, Year)	Domain	Model(s)	XAI Method	Reported Performance (as stated)	Venue / DOI
1	Schwartz et al., 2020 [1]	Foundational – Green AI	Position paper (n/a)	n/a	Proposes efficiency (FLOPs / cost) as an evaluation criterion alongside accuracy	Commun. ACM 63(12):54–63; doi:10.1145/3381831
2	Strubell et al., 2019 [2]	Foundational – cost of DL	NLP Transformer models (analysis)	n/a	Estimated training-run emissions comparable to several cars' lifetime CO□	ACL 2019, 3645–3650; doi:10.18653/v1/P19-1355
3	Bolón-Canedo et al., 2024 [3]	Foundational – Green AI review	Review (n/a)	n/a	Surveys green-in-AI / green-by-AI techniques and measurement tools	Neurocomputing 599:128096; doi:10.1016/j.neucom.2024.128096
4	Lundberg & Lee, 2017 [4]	Foundational – XAI method	Model-agnostic	SHAP (introduced)	Introduces SHAP as a unified additive feature-attribution framework	NeurIPS 30, 4765–4774
5	Elbasi et al., 2025 [5]	Green AI – smart-agriculture framework	Lightweight/compressed models, federated learning, edge computing	Listed as future work only	Framework proposal for energy-efficient crop-yield and resource-management models	IEEE Access 13:204924–204953; doi:10.1109/ACCESS.2025.3637027
6	Mohan et al., 2025 [11]	Crop yield (climate factors)	Decision Tree, Random Forest, LightGBM	SHAP, LIME	R^2 up to 0.92	Frontiers in Plant Science

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

No.	Study (Authors, Year)	Domain	Model(s)	XAI Method	Reported Performance (as stated)	Venue / DOI
7	Yenkikar et al., 2025 [12]	Crop / rice yield	Hybrid RF–LSTM–XGBoost	SHAP, LIME, Counterfactual	R^2 = 0.9827 (crop), 0.9721 (rice)	MethodsX 15:103442; doi:10.1016/j.mex.2025.103442
8	Pant et al., 2025 [13]	Crop yield (comparative)	XGBoost, Decision Tree, Random Forest, Linear Regression	SHAP, LIME	Accuracy up to 99.09% (RF); R^2=0.95 (XGBoost)	2025 ICAECT
9	Mampitiya et al., 2025 [14]	Tea crop yield	CatBoost, XGBoost, LightGBM, ANN, Ridge, LASSO, ElasticNet	SHAP	R^2 = 0.904 (CatBoost, best of 7 models)	Smart Agricultural Technology
10	Kushwaha et al., 2026 [15]	Crop yield (hyper-local, India)	XGBoost Regressor	SHAP	R^2 = 0.93	2026 ICCSC
11	Kehkashan et al., 2026 [16]	Corn yield – low-carbon computing special section	Random Forest, Gradient Boosting, Ensemble Voting	SHAP	Precision/Recall/F1 ≈ 0.92; training accuracy 0.97	Scientific Reports
12	Dangi et al., 2025 [17]	Crop yield	LightGBM	SHAP	R^2 = 0.94	2025 IEEE REACS
13	Alsumaiei et al., 2026 [25]	Soil moisture / irrigation (arid)	SVM, Gaussian Process Regression, Bagged Trees Ensemble	SHAP	R^2 up to 0.987 (Bagged Trees Ensemble)	J. Irrigation and Drainage Engineering
14	Alahmad et al., 2025 [26]	Soil-moisture content	Random Forest Regression	SHAP	R^2 up to 0.97 (silt loam, 5 cm depth)	Acta AgronomicaÓváriensis
15	Gopalan et al., 2025 [18]	Corn leaf disease	ResNet152	Grad-CAM	98.34% test accuracy	BMC Plant Biology
16	Islam et al., 2025 [19]	Plant disease (15 species)	EfficientNetB7	Grad-CAM	99.90% test accuracy	2025 ICCIT

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

No.	Study (Authors, Year)	Domain	Model(s)	XAI Method	Reported Performance (as stated)	Venue / DOI
17	Joy et al., 2025 [20]	Plant disease (comparative)	InceptionV3, DenseNet121, Xception, MobileNetV2	Grad-CAM + Explainability Score	97.47–98.47% accuracy across 4 backbones	2025 ICoEIT
18	Alhammad et al., 2025 [21]	Potato leaf disease	Transfer-learning CNN	Grad-CAM	98% test accuracy	Frontiers in Artificial Intelligence
19	Al Kafi et al., 2026 [22]	Plant disease (edge computing)	Logistic Regression + MobileNetV3 (hybrid)	Grad-CAM	~3% accuracy loss vs. full DL; 77.6% faster inference	PLOS One
20	Gupta et al., 2025 [24]	Crop recommendation / health	LightGBM	SHAP	98.64% accuracy	2025 MoSICom
21	Pai et al., 2025 [8]	Review – XAI in agriculture	n/a (review)	LIME, SHAP, Grad-CAM, rule-based (surveyed)	n/a (methodological review)	Engineering Research Express
22	Rajbongshi et al., 2026 [9]	Review – XAI for sustainable agriculture	n/a (review, 2017–2024 literature)	LIME, SHAP, Grad-CAM (surveyed)	n/a (methodological review)	Artificial Intelligence Review
23	Abdul Razak et al., 2024 [10]	Scoping review – Agriculture 5.0 & XAI	n/a (84 articles, 2018–2023)	XAI methods (surveyed)	n/a (scoping review)	Emerging Science J.; doi:10.28991/ESJ-2024-08-02-024

Table 1 groups the reviewed sources into four categories: (i) foundational Green AI and XAI references; (ii) crop-yield and soil/irrigation prediction studies that use XAI; (iii) plant-disease detection studies that use XAI; and (iv) broader reviews of XAI in agriculture. For each source, the table reports the models used, the XAI technique used, the headline

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

performance metric exactly as stated in the original paper, and the publication venue with its DOI where one is publicly indexed.

Table 1. Reviewed literature at the intersection of Green AI, XAI and machine learning

Two patterns are visible across Table 1. First, SHAP is the dominant explainability technique for tabular, regression tasks (yield and soil-moisture prediction), while Grad-CAM dominates image-based classification tasks (disease detection). Second, of the fifteen application-focused XAI studies summarised in rows 6–20 of Table 1, only one – the corn-yield study of Kehkashan et al. [16], explicitly framed within a special section on sustainable, low-carbon agricultural computing – is positioned around energy efficiency at all, and even that study does not report a quantified energy, compute-time, or carbon figure alongside its accuracy and SHAP results. This asymmetry is quantified in Section 4.2 and is the central empirical gap that the proposed framework in Section 3 is designed to close.

III. METHODOLOGY

This section proposes a Green-XAI methodological framework for agricultural ML, addressing gaps RG1–RG4 identified in Section 1. The framework has not yet been empirically executed; it is presented here as a fully specified methodology – dataset scope, candidate models with their governing equations, the explainability mechanism, and the energy-instrumentation protocol – intended for the empirical validation described in Section 5.

3.1 Dataset

The framework is designed around three agricultural task families for which suitable public or field-collected data are available.

- **Crop-yield / soil-condition prediction** – tabular data combining meteorological variables (temperature, rainfall, humidity, solar radiation), soil properties (moisture, N-P-K levels, pH), and management variables (crop type, sowing date, irrigation regime), of the kind used in the crop-yield and soil-moisture studies in Table 1 [11], [12], [16], [25], [26]. Suitable sources include national agricultural-statistics datasets and open agromet / soil-sensor repositories.
- **Plant-disease detection –** RGB leaf imagery labelled by disease class, following the benchmark structure used in [19], [21], [22], [23] (e.g., PlantVillage-style corpora), supplemented where possible with field-captured images to test out-of-distribution robustness, an aspect noted as under-addressed in the reviewed disease-detection literature [20].
- **Irrigation / soil-moisture time series** – multi-depth soil-moisture and evapotranspiration sensor time series with corresponding meteorological drivers, following the structure used in [25], [26].
- For all three task families, a stratified 70/15/15 or 80/10/10 train/validation/test split is proposed, stratified by crop type or disease class where applicable, consistent with the split ratios reported in the reviewed literature (e.g., a 70/30 train/test split in [18]).

3.2 Proposed Model Framework

To keep the accuracy–energy trade-off central rather than incidental, the framework evaluates two tiers of models against each task family, both of which recur repeatedly in Table 1.

- **Tier 1** – lightweight tree-based ensembles: Random Forest (RF), XGBoost, and LightGBM for the tabular yield / soil tasks.
- **Tier 2** – compact deep architecture: a small transfer-learning CNN (e.g., a MobileNetV2- or EfficientNet-B0-class network) for the imagery task, deliberately chosen at the lighter end of the backbone spectrum represented in Table 1, where the reviewed studies mostly report on much larger networks (ResNet152, EfficientNetB7).

Each tier is paired with the SHAP explainability layer described in Section 3.4, and every training and inference run is wrapped in the energy-instrumentation layer described in Section 3.5, so that accuracy, explanation quality, and energy cost are measured on the same run rather than estimated separately after the fact.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

3.3 Model Formulations

Random Forest aggregates T decision trees h□ trained on bootstrap samples with random feature subsampling at each split. For a regression target the ensemble prediction is

$$\hat{y}(x) = (1/T)\,\Sigma\square\square\square^{T}\, h\square(x) \quad (1)$$

and for classification the prediction is the majority vote across trees.

Gradient-boosted trees (XGBoost, LightGBM) build an additive ensemble of K trees f□ sequentially, each fitted to the residual gradient of the loss from the previous stage. The regularised training objective at iteration K is

$$Obj(\theta) = \Sigma_i\, l(y_i, \hat{y}_i) + \Sigma\square\ \Omega(f\square)\ , \quad \Omega(f) = \gamma T^{f} + \tfrac{1}{2}\lambda \parallel w \parallel^{2} \quad (2)$$

where l is a differentiable loss (e.g., squared error for yield regression), T^f is the number of leaves in tree f, w the leaf weights, and γ, λ are regularisation coefficients that penalise tree complexity – the mechanism by which these methods control overfitting without the very large parameter counts associated with deep networks. LightGBM uses the same additive-boosting objective but grows trees leaf-wise with gradient-based one-side sampling, which is the source of its typically lower training-time and memory footprint relative to level-wise boosting, a property directly relevant to the energy-efficiency objective of this framework.

Convolutional neural networks for disease classification learn a hierarchy of convolutional feature maps A^k followed by a classification head producing class scores y^c. As a complementary check on SHAP (Section 3.4), the framework also computes Grad-CAM localisation maps:

$$\alpha\square^{c} = (1/Z)\,\Sigma_i\,\Sigma\square\ \ (\partial y^{c} / \partial A_{i}\square\square)\ ,L^{HradLAMc} = ReLU(\ \Sigma\square\ \alpha\square^{c}\ A\square\ ) \quad (3)$$

which produces the coarse leaf-region localisation used throughout the disease-detection literature in Table 1 [18]–[23] to indicate which part of an image drove a given classification.

3.4 Explainable AI Integration (SHAP)

Interpretability is provided primarily through SHAP [4], which is model-agnostic and applies uniformly across the Tier-1 tree ensembles and, via its DeepSHAP/GradientSHAP variants, the Tier-2 CNN. SHAP attributes a prediction f(x) to its input features by computing each feature's Shapley value – the game-theoretic solution for fairly dividing a prediction among 'players' (features):

$$\varphi_i = \Sigma\ S\subseteq F\backslash\{i\}\ \ [\ |S|!(|F|-|S|-1)!\ /\ |F|!\ ]\ \cdot\ [\ f(S\cup\{i\}) - f(S)\ ] \quad (4)$$

where F is the full feature set and f(S) is the model's expected output when only the features in subset S are known. The framework uses TreeSHAP for the Tier-1 ensembles – an efficient, polynomial-time algorithm for tree models rather than the exponential-time formulation in Eq. (4) specifically because its lower computational (and therefore energy) cost relative to KernelSHAP is itself part of the efficiency argument this paper makes. Both global explanations (mean absolute SHAP value per feature, aggregated across the test set, to identify dominant agronomic drivers such as rainfall, NDVI, or nitrogen level, consistent with [11], [16], [25]) and local explanations (per-prediction SHAP values, to justify an individual yield or disease call to an end user) are computed and logged.

3.5 Green AI Energy-Efficiency Instrumentation

Every training run, and every batch of SHAP or Grad-CAM explanation computation, is wrapped with an energy tracker (e.g., CodeCarbon [28], which samples CPU, GPU, and RAM power draw via RAPL/NVML and multiplies by run duration and local grid carbon intensity). For a run of duration T, total energy is decomposed as

$$E_total = E_CPU + E_GPU + E_RAM = \int\ \square^{T}\ [P_CPU(t) + P_GPU(t) + P_RAM(t)]\ dt \quad (5)$$

and the associated emissions estimate is

$$CO\square e = E_total \times CI_grid \quad (6)$$

where CI_grid is the regional carbon intensity of the electricity supply (kg CO□e per kWh). Reporting both E_total (a hardware-portable metric, following Schwartz et al.'s recommendation to report a hardware-independent proxy such as FLOPs alongside a "price tag" metric [1]) and CO□e (a policy-relevant but location-dependent metric) directly addresses RG1 and RG3: energy is logged separately for the model-training phase and for the explanation-generation

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

phase (TreeSHAP / Grad-CAM), so that the marginal cost of explainability itself becomes visible rather than folded silently into total training cost.

3.6 Evaluation Protocol

Three families of metric are proposed, evaluated jointly rather than sequentially.

- **Predictive performance** – R^2 and RMSE for the yield/soil regression tasks; accuracy, precision, recall and F1 for the disease-classification task,chosen to remain directly comparable to Table 1.
- **Explanation quality** – SHAP-consistency across bootstrap resamples (whether the top-k important features remain stable), and, for the imagery tier, Grad-CAM/SHAP localisation agreement, in the spirit of the Explainability Score and Jaccard-index approach used by Joy et al. [20].
- **Resource efficiency** – E_total and CO□e per training run and per 1,000 inferences (Section 3.5), plus accuracy-per-kWh and accuracy-per-explanation-second as composite efficiency indicators, enabling the accuracy-versus-energy analysis that is the central comparison this framework is designed to produce and that is largely absent from Table1.

IV. RESULTS AND DISCUSSION

4.1 Comparative Synthesis of Reported Model Performance

Figures 1 and 2 consolidate the headline performance figures reported by the studies in Table 1, grouped by task family. These are the values reported by the original authors on their own test sets – not an independent re-run or benchmark – and are reproduced here to make the current state of the art visible in one place; direct
comparison across studies should be read with the caveat that datasets, train/test splits, and evaluation protocols differ between them.

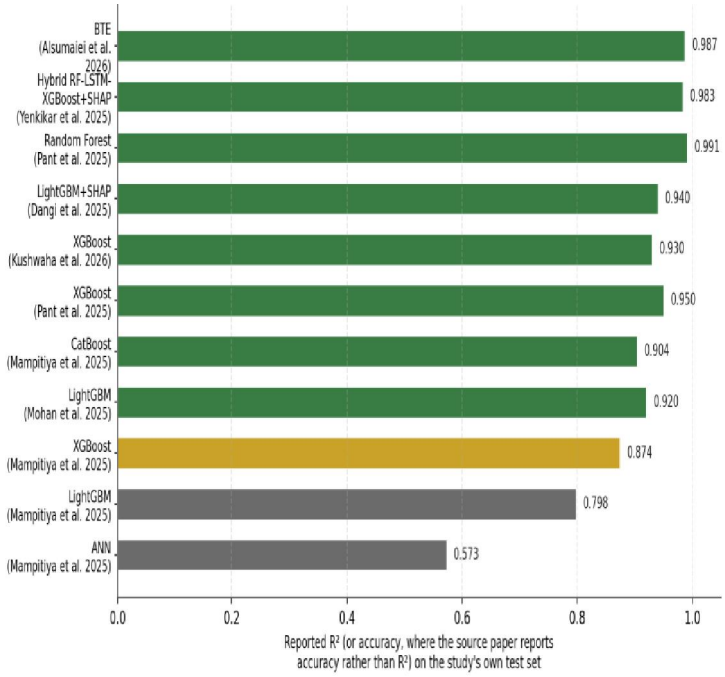


Fig 1 – Reported R^2 for crop-yield and soilmoisture models across the studies in Table 1.

For crop-yield and soil-moisture prediction (Fig. 1), tree-based ensembles are consistently the top performers: the bagged-tree ensemble of Alsumaiei et al. [25] and the SHAP-explained hybrid RF–LSTM–XGBoost model of Yenkikar et al. [12] both report R^2 above 0.98, and XGBoost- or LightGBM-based models cluster in the $R^2 = 0.87–0.95$ range across four independent studies [12], [13], [15], [17]. Simpler models trail noticeably: in the Mampitiya et al. tea-

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

yield comparison [14], CatBoost (R^2 = 0.904) clearly outperforms an unregularised ANN (R^2 = 0.573) on the same 99.90% (EfficientNetB7 [19]). Within a single controlled comparison across four backbones, Joy et al. [20] report a spread of only about one percentage point between the best and weakest backbone – suggesting that, at this level of task difficulty, backbone choice affects accuracy far less than it likely affects the resource cost of achieving that accuracy, a relationship the framework in Section 3 is designed to test directly (Hypothesis H3, Section 4.4) For plant-disease detection (Fig. 2), transfer-learning CNN backbones explained with Grad-CAM cluster at the top of the accuracy range, from 97.47% (InceptionV3 [20]) to

dataset, evaluated with the same SHAP pipeline.

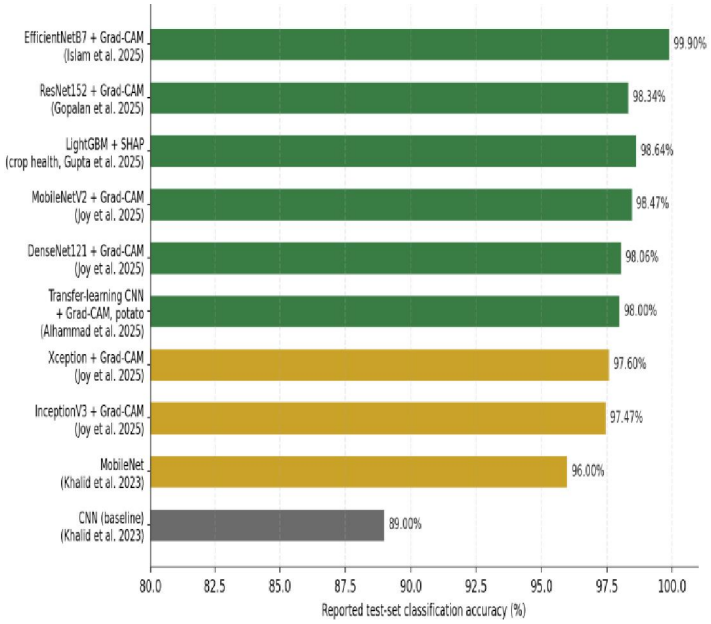
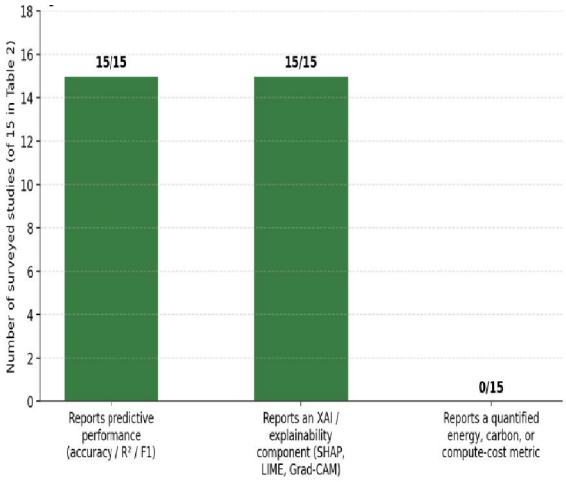


Fig 2 – Reported test-set classification accuracy of explainable deep-learning models for plant-disease detection across the studies in Table 1.

4.2 The Green-AI / XAI Reporting Gap

Figure 3 – Coverage of predictive-performance reporting, explainability reporting, and quantified energy/carbon reporting across the 15 application-focused studies in Table 1.



IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

Figure 3 quantifies the central empirical gap motivating this paper. Of the fifteen application-focused XAI-in-agriculture studies summarised in Table 1 (rows 6–20), all fifteen report a predictive-performance metric and all fifteen report an explainability component, but none report a quantified energy-, compute-time-, or carbon-related measurement alongside their results – not even Kehkashan et al.'s study [16], which is explicitly positioned within a special section on sustainable, low-carbon agricultural computing. Even the purpose-built Green-AI-for-agriculture framework of Elbasi et al. [5] does not itself incorporate an explainability layer, instead listing XAI integration as future work. Read together, these observations show that Green AI and XAI have so far been treated as separate design goals in the agricultural-ML literature reviewed here, rather than jointly optimised – precisely the gap RG1 identifies.

4.3 Positioning of the Proposed Framework

Table 2 positions the Green-XAI framework proposed in Section 3 against the general pattern observed across the reviewed literature (Table 1), rather than against any single study, since no single prior work combines all of the listed dimensions.

Dimension	Typical prior XAI-agriculture study (Table 1 pattern)	Proposed Green-XAI framework (Section 3)
Explainability method	SHAP or Grad-CAM – usually one technique, chosen by input modality	SHAP across both tiers (TreeSHAP + GradientSHAP), cross-checked against Grad-CAM for imagery
Model families compared	Usually one family or a small ad hoc set	Tiered comparison: RF / XGBoost / LightGBM vs. a compact CNN, explicitly including a lightweight tier
Energy / carbon reporting	Reported in 0 of the 15 application studies reviewed (Table 1, Fig. 3)	Reported for every training run and every explanation batch, by design (Section 3.5)
Explanation cost isolated from training cost	Not reported in any reviewed study	Measured separately via Eq. (5)–(6)
Efficiency metric	Not defined	Accuracy-per-kWh and accuracy-per-explanation-second (Section 3.6)
Validation status	Empirically executed; single-study results	Proposed and fully specified; empirical validation planned

Table 2. Qualitative comparison of the proposed framework against the pattern observed in reviewed prior work

4.4 Expected Outcomes and Hypotheses for Empirical Validation

The comparative synthesis in Sections 4.1–4.3 supports three testable hypotheses for the empirical validation described in Section 5, rather than results in themselves.

- **H1 (accuracy–efficiency trade-off):** given that tree-based ensembles already match or exceed deep-learning-level accuracy on the tabular yield/soil tasks in Fig. 1, Tier-1 models are expected to reach within a few percentage points of Tier-2 CNN accuracy on comparable tasks while requiring substantially less training energy, consistent with the generally lower FLOP cost of gradient-boosted trees relative to deep CNNs.
- **H2 (explanation overhead):**TreeSHAP's polynomial-time complexity on ensemble models is expected to add a small, near-constant energy overhead relative to model training, whereas Grad-CAM's marginal cost is expected to scale with the number of images explained rather than with model size – meaning the energy cost of 'being explainable' is expected to differ structurally between the two model tiers.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, November 2026
Impact Factor: 8.2

- **H3 (backbone diminishing returns):** consistent with the narrow accuracy spread across CNN backbones noted in Fig. 2 [20], smaller backbones (e.g., MobileNetV2-class) are expected to sacrifice little accuracy relative to the largest backbones (ResNet152,EfficientNetB7) while offering a disproportionately larger reduction in training and inference energy.

V. FUTURE WORK

The immediate next step is to implement and execute the framework specified in Section 3 against the three task families, using the datasets outlined in Section 3.1, and to test Hypotheses H1–H3 empirically, reporting accuracy, explanation-quality, and energy/carbon metrics jointly as proposed in Section 3.6. Beyond this core validation, several extensions follow naturally from the gaps identified in Section 1.

- **Federated and edge deployment** – extending the energy-instrumentation layer to federated-learning and edge-computing settings, in line with the direction flagged by Elbasi et al. [5] and the edge-focused disease-detection work of Al Kafi et al. [22], so that on-device inference energy – not just centralised training energy – is captured.
- **Explanation-faithfulness benchmarking** – moving beyond SHAP-consistency and Grad-CAM/SHAP agreement (Section 3.6) toward formal faithfulness and robustness metrics for XAI in the agricultural domain, an open problem also flagged in recent XAI-agriculture reviews [8], [9].
- **Actionable sustainability loops** – closing RG4 by connecting SHAP-identified drivers (e.g., over-fertilisation, excess irrigation) directly to farm-level resource-reduction recommendations and measuring downstream water/fertiliser savings, rather than treating explainability as an end in itself.
- **Reporting standards** – contributing toward a standard "efficiency label" for agricultural-AI publications, comparable to the "price tag" reporting Schwartz et al. proposed for AI research generally [1], so that energy and carbon costs become as routinely reported as accuracy in this literature.

VI. CONCLUSION

Artificial intelligence is increasingly central to precision agriculture, and explainability is increasingly treated as a precondition for farmer and policymaker trust in its predictions. This paper has shown, through a structured synthesis of recent literature, that the parallel push toward energy- and carbon-efficient – "Green" – AI has not yet been integrated with that explainability effort in the agricultural domain: fifteen recent XAI-in-agriculture studies report strong predictive performance and rich SHAP- or Grad-CAM-based explanations, yet none report a quantified energy or carbon measurement, and the one purpose-built Green-AI agriculture framework identified explicitly defers explainability to future work. In response, this paper has proposed a fully specified Green-XAI methodological framework – spanning dataset scope, tiered model formulations, SHAP-based interpretability, and energy/carbon instrumentation – designed to close this gap, together with the concrete hypotheses and validation plan needed to test it empirically. Carrying out that validation is the necessary next step toward agricultural AI systems that are simultaneously trustworthy and demonstrably sustainable.

REFERENCES

[1] Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. Communications of the ACM, 63(12), 54–63. https://doi.org/10.1145/3381831

[2] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL), 3645–3650. https://doi.org/10.18653/v1/P19-1355

[3] Bolón-Canedo, V., Morán-Fernández, L., Cancela, B., & Alonso-Betanzos, A. (2024). A review of green artificial intelligence: Towards a more sustainable future. Neurocomputing, 599, 128096. https://doi.org/10.1016/j.neucom.2024.128096

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

[4] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765–4774.
[5] Elbasi, E., Alzoubi, Y. I., Topcu, A. E., & Nadeem, M. (2025). Green AI for smart agriculture: Energy-efficient predictive models for crop yield and resource management. IEEE Access, 13, 204924–204953. https://doi.org/10.1109/ACCESS.2025.3637027
[6] Mana, A. A., et al. (2024). Sustainable AI-based production agriculture: Exploring AI applications and implications in agricultural practices. Smart Agricultural Technology.
[7] Gupta, S., et al. (2026). Advancing sustainable and green agriculture with AI and IoT: Trends, challenges, and future directions. Green Technologies and Sustainability.
[8] Pai, D. G., et al. (2025). Explainable AI in agriculture: Review of applications, methodologies, and future directions. Engineering Research Express.
[9] Rajbongshi, A., Johora, F. T., Hossain, A., Sarker, M. S., Rahman, M. H., Rahman, M. W., ... & Moni, M. A. (2026). Leveraging explainable AI for sustainable agriculture: a comprehensive review of recent advances. Artificial Intelligence Review, 59(3), 105.
[10] Abdul Razak, S. F., et al. (2024). Agriculture 5.0 and Explainable AI for smart agriculture: A scoping review. Emerging Science Journal, 8(2). https://doi.org/10.28991/ESJ-2024-08-02-024
[11] Mohan, R. J., Rayanoothala, P. S., & Sree, R. P. (2025). Next-gen agriculture: integrating AI and XAI for precision crop yield predictions. Frontiers in Plant Science, 15, 1451607.
[12] Yenkikar, A., Mishra, V. P., Bali, M., & Ara, T. (2025). An explainable AI-based hybrid machine learning model for interpretability and enhanced crop yield prediction. MethodsX, 15, 103442. https://doi.org/10.1016/j.mex.2025.103442
[13] Pant, H., et al. (2025). Comparative study of crop yield prediction using explainable AI and interpretable machine learning techniques. 2025 Fifth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT).
[14] Mampitiya, L., Sumanasekara, H. S., Rathnayake, N., Hoshino, Y., & Rathnayake, U. (2025). Explainable artificial intelligence to estimate the Sri Lankan (Ceylon) Tea crop yield. Smart Agricultural Technology, 11, 100999.
[15] Kushwaha, P., et al. (2026). Explainable AI for hyper-local crop yield prediction in India. 2026 International Conference on Computing, Sciences and Communications (ICCSC).
[16] Kehkashan, T., et al. (2026). Sustainable regional corn yield prediction for the United States through interpretable machine learning approach. Scientific Reports.
[17] Dangi, A., & Rathore, D. K. (2025, December). A LightGBM-SHAP Framework for Accurate and Interpretable Crop Yield Prediction. In 2025 IEEE International Conference on Recent Advances in Computing and Systems (REACS) (pp. 1-6). IEEE.
[18] Gopalan, K., Srinivasan, S., ., P. et al. Corn leaf disease diagnosis: enhancing accuracy with resnet152 and grad-cam for explainable AI. BMC Plant Biol 25, 440 (2025). https://doi.org/10.1186/s12870-025-06386-0
[19]Islam, M. N., Roy, A., & Shimu, S. A. (2025, December). Improved Plant Disease Classification by Explainable AI Techniques: A Study with Grad-CAM and EfficientNet. In 2025 28th International Conference on Computer and Information Technology (ICCIT) (pp. 2133-2138). IEEE.
[20] Joy, M. L., Jain, V., Sinha, P. K., & Gupta, N. K. (2025, July). Interpretable Deep Learning for Plant Disease Detection: A Comparative Study with Explainability Metrics. In 2025 International Conference on Engineering Innovations and Technologies (ICoEIT) (pp. 232-238). IEEE.
[21] Alhammad, S. M., Khafaga, D. S., El-Hady, W. M., Samy, F. M., & Hosny, K. M. (2025). Deep learning and explainable AI for classification of potato leaf diseases. Frontiers in Artificial Intelligence, 7, 1449329.
[22] Al Kafi, M. A., Banshal, S. K., Moni, R., Aziz, A. L., Aljuaid, M., & Bansal, R. (2026). LeafAI: Interpretable plant disease detection for edge computing. Plos one, 21(1), e0335956.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, November 2026 Impact Factor: 8.2

[23] Khalid, M. M., & Karan, O. (2024). Deep learning for plant disease detection. International Journal of Mathematics, Statistics, and Computer Science, 2, 75-84.
[24] Gupta, H., Sharma, A., Sharma, S., & Dixit, V. (2025, December). Enhancing crop prediction with explainable AI: A SHAP-based approach. In 2025 International Conference on Modeling, Simulation & Intelligent Computing (MoSICom) (pp. 7-12). IEEE.
[25] Alsumaiei, A. A., et al. (2026). Enhancing soil water prediction in arid climates using multipredictor machine-learning models and SHAP-based interpretability. Journal of Irrigation and Drainage Engineering.
[26] Alahmad, T., Neményi, M., & Nyéki, A. (2025). Enhancing Agricultural Sustainability using AI-Driven Soil Moisture Modeling: A Soil-Type and Depth Approach with SHAP Interpretability. Acta AgronomicaÓváriensis, 66(2), 5-16.
[27] Cartolano, A., Cuzzocrea, A., & Pilato, G. (2024). Analyzing and assessing explainable AI models for smart agriculture environments. Multimedia Tools and Applications, 83(12), 37225-37246.
[28] Lottick, K., Susai, S., Friedler, S. A., & Wilson, J. P. (2019). Energy usage reports: Environmental awareness as part of algorithmic accountability (basis of the CodeCarbon tool). NeurIPS Workshop on Tackling Climate Change with Machine Learning.