

Explainable AI for Agricultural Climate-Risk Prediction: A Systematic Review of Progress and Strategic Roadmap

Jacintha Menezes

`jacintha.menezes@majancollege.edu.om`

Majan University College

Rishabh Kant Jha

Lakehead University

Jia Bernadette D Souza

Ontario Tech University

Systematic Review

Keywords: Explainable AI (XAI), Agriculture, Crop Yield Prediction, Soil Salinity Stress, SHAP, LIME, PRISMA

Posted Date: September 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10686074/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

The global agricultural sector faces unprecedented threats from climatic variability, including drought, soil salinity, and yield instability. While Artificial Intelligence (AI) and Deep Learning (DL) have achieved high predictive accuracy in managing these risks, their "black-box" nature creates a transparency deficit that hinders adoption by agronomists and farmers. This study presents a PRISMA-guided systematic literature review of shortlisted papers published between 2021 and 2026, capturing the period where Explainable AI (XAI) moved to the forefront of agricultural research.

The review categorizes findings across five primary domains: crop yield, salinity stress, drought, plant disease, and climate-risk forecasting. Key findings indicate that while models achieve exceptional benchmarks such as scores of 0.975 in yield prediction and 99.63% accuracy in disease detection their explanations are often post-hoc and regionally bounded. Technical evaluations highlight SHAP's axiomatic superiority for global consistency, contrasted with LIME's local instability and the failure of marginal distribution methods (ALE/PDP) under the multicollinearity of climate data. The study identifies critical research gaps, including a Global North bias, a lack of uncertainty communication, and the "silo effect" of modeling risks in isolation. Finally, the review provides a strategic roadmap for the field, advocating for "Explainability-by-Design," unified federated-XAI for privacy-aware diagnostics, and the standardization of uncertainty-aware probabilistic forecasting (e.g., DeepQuantreg) to ensure that AI-driven insights are robust, trustworthy, and operationally ready for field-level deployment.

1. Introduction

1.1 Background and Justification

The global agricultural sector is under persistent threat from climatic variability, with drought, water scarcity, and soil salinity presenting the most formidable challenges to productivity and food security. Regional and country-level investigations into drought forecasting emphasize that while climatic fluctuations can be monitored via machine learning algorithms, the efficacy of these methods is highly contingent upon localized climate and data conditions (Bouregaa, 2025; Ozupek et al., 2025; Tanriverdi & Batmaz, 2025a). Soil salinity presents a specialized climate-related risk, particularly within delta agriculture and irrigated systems where degraded soil productivity directly compromises global yields (Mohammadifar et al., 2022; Sahoo et al., 2025). Consequently, developing reliable predictive systems for these risks is paramount for safeguarding agricultural stability.

Artificial Intelligence (AI) has emerged as a disruptive paradigm in this domain, providing the computational capacity to process vast multi-dimensional agricultural datasets encompassing soil metrics, remote-sensing indicators, and climatic variables at speeds unattainable by traditional statistical techniques (Cheema et al., 2026; Hosseinpour et al., 2025). Recent advances in machine learning (ML) and deep learning (DL) have refined crop yield forecasting by integrating soil moisture models and environmental variables into sophisticated prediction pipelines (Rajbongshi et al., 2026).

Despite these gains in predictive accuracy, the increasing complexity of DL and ML architectures has precipitated a "black box" problem. These models often provide negligible insight into their internal decision-making logic, creating a significant barrier to adoption among stakeholders. Agronomists, farmers, and other "less expert" users who lack specialized technical knowledge frequently cannot analyze the model's working process leading to a lack of trust in automated results (Ozupek et al., 2025). This transparency deficit is a recognized obstacle in the management of plant stress and crop diseases (Alsanoosy & Malik, 2025; Shawon et al., 2026). Therefore, Explainable AI (XAI) is essential for establishing the transparency required to translate high-dimensional predictive power into actionable trustworthy agricultural results.

1.2 Research Objectives

The primary objectives of this systematic literature review are:

- To examine recent studies that apply machine learning and deep learning models to climate-risk prediction in agriculture.
- To identify interpretable and Explainable AI (XAI) techniques utilized in agricultural prediction models including SHAP, LIME, feature importance, partial dependence, attention mechanisms, and transparent model structures.
- To review XAI applications across five specific agricultural climate-risk domains: crop yield prediction, salinity stress prediction, drought and water-stress forecasting, plant disease prediction, and climate-risk forecasting.
- To evaluate the technical strengths and limitations of interpretable ML models in supporting agricultural decision-making.
- To identify critical research gaps concerning explainability, model transparency, generalizability, data quality, and real-world deployment in climate-smart agriculture.

1.3 Research Questions

1. What machine learning and deep learning models are most commonly used for climate-risk prediction in agriculture?
2. What explainable AI and interpretable machine learning techniques are used to interpret agricultural prediction models?
3. How are interpretable ML models applied in crop yield prediction, salinity stress prediction, drought forecasting, plant disease prediction, and climate-risk forecasting?
4. What are the major strengths and limitations of current XAI approaches in agricultural climate-risk prediction?
5. What methodological gaps remain in the development of explainable, trustworthy, and decision-support-oriented agricultural prediction systems?

1.4 Scope and Boundaries

The temporal scope of this review is strictly defined between 2021 and 2026, capturing the period during which XAI has moved to the forefront of agricultural AI research. The review is bounded by five primary agricultural domains: crop yield prediction; salinity stress and soil salinity prediction; drought, water stress, and irrigation forecasting; plant or crop disease prediction; and agricultural climate-risk forecasting. This review specifically prioritizes the intersection of predictive performance with model interpretability to address the operational needs of the agricultural sector.

2. Methodology

2.1 Systematic Review Framework and Database Selection

This study adheres to the PRISMA 2020 (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework to ensure methodological rigor, reproducibility, and transparency. Three primary databases were selected to ensure a balanced, interdisciplinary perspective:

- SpringerLink: Selected for its robust coverage of agricultural modeling, environmental science, and "sustainability research," alongside its deep repository of applied XAI studies.
- IEEE Xplore: Chosen to capture technically oriented research focusing on DL, computer vision, sensor-based agriculture, and precision farming technologies.
- Taylor & Francis Online: Utilized for its emphasis on interdisciplinary studies relating to "environmental risk," "crop modeling," and data-driven agricultural decision-making.

2.2. Search Strategy and Boolean Logic

The search strategy was constructed using a four-block conceptual keyword approach:

- Explainable and Interpretable AI: Focusing on XAI, transparent AI, and specific techniques like SHAP and LIME.
- Agriculture and Crop Systems: Targetting farming, crops, and agroecosystems.
- Climate-Risk Prediction Domains: Identifying yield forecasting, drought, salinity, and disease.
- Machine Learning and Deep Learning Models: Specifying algorithms like Random Forest, XGBoost, CNN, and LSTM.

Master Search String: (("explainable artificial intelligence" OR XAI OR "interpretable machine learning" OR "interpretable AI" OR "explainable machine learning" OR "transparent AI" OR SHAP OR LIME OR "feature importance") AND (agriculture OR farming OR crops OR "crop production" OR agroecosystem*) AND ("crop yield prediction" OR "yield forecasting" OR drought OR "drought forecasting" OR "drought prediction" OR salinity OR "salinity stress" OR "soil salinity" OR disease OR "plant disease prediction" OR "climate risk" OR "climate-risk assessment" OR "climate risk assessment" OR "climate forecasting") AND ("machine learning" OR "deep learning" OR forecasting OR prediction OR "random forest" OR XGBoost OR "support vector machine" OR LSTM OR CNN))

2.3 Eligibility Criteria

The criteria for study inclusion and exclusion are detailed in the table below.

Table 1: Eligibility Criteria for Study Selection

Inclusion Criteria	Exclusion Criteria
Temporal Scope: Published between 2021–2026. Older studies were considered only if foundational, though ultimately none were included.	Unrelated Topic: Not related to agriculture or crop systems; studies focused on general AI without agricultural prediction.
Topical Relevance: Must address XAI or interpretable ML within the five defined agricultural climate-risk domains.	No Interpretability: Studies utilizing ML/DL but lacking any component of interpretability or explainability.
Methodological Requirement: Use of ML/DL approaches (e.g., Random Forest, CNN, SHAP, LIME, attention mechanisms).	Inappropriate Source: Editorials, commentaries, news articles, blog posts, or other non-academic sources.
Study Type: Peer-reviewed journal articles, conference proceedings, or applied AI review papers.	Out of Scope: Studies published before 2021; studies focused on remote sensing or GIS without ML/DL prediction.
Language: English-language publications only.	Technical/Structural Removal: Duplicate records or studies with missing abstracts/insufficient title relevance.

2.4 PRISMA Flow Diagram

Fig 1. Provides the PRISMA flow followed for systematic literature review. The screening and study selection process followed the four standardized PRISMA stages:

1. Identification of studies via databases: Initial records were retrieved from SpringerLink, IEEE Xplore, and Taylor & Francis Online.
2. Duplicate Removal: Records were normalized by title and removed to ensure a unique dataset for review.
3. Title and Abstract Screening: Records were filtered based on the absence of agricultural relevance, missing DL/ML components, or lack of explainability.

Final Inclusion: The shortlisted papers evaluated thematic relevance, methodological alignment, and suitability for qualitative synthesis. Selected papers were then grouped into the five review categories: crop yield prediction, salinity stress prediction, drought and water-stress forecasting, plant disease prediction and irrigation/climate-risk forecasting.

Based on the data visualized in Fig 2, the distribution of the 33 shortlisted papers demonstrates a significant upward trajectory within the 2021–2026 research window. While the methodological search scope (as defined in Section 1.4) began in 2021, the analysis confirms that this specific year yielded 0

results. The literature only began to appear in 2022, highlighting a distinct recency bias and the exceptionally rapid emergence of explainability as a prerequisite in agricultural modeling. This growth is driven by the technical transition described in Section 3.1, where research has shifted from traditional tree-based ensembles toward complex deep learning architectures (e.g., CNNs and LSTMs) that necessitate Explainable AI (XAI) to overcome the "black-box" problem. In 2025 and 2026 the volume of published research on agricultural XAI reaches its peak. This surge is directly aligned with the "Scope and Boundaries" of the review, representing the historical moment when XAI moved from a niche technical concern to the forefront of agricultural AI research. This concentration of literature reflects the industry's pivot toward establishing the transparency required to translate high-dimensional predictive power into actionable, trustworthy results for agronomists and stakeholders.

3. Synthesis of Agricultural Outcomes and Environmental Drivers

This section synthesizes current advancements in predictive performance across diverse agricultural domains, identifying the primary environmental drivers that govern climate-related risks.

The thematic distribution of the 33 shortlisted studies, as illustrated in Fig 3, reveals a concentrated research focus on high-impact agricultural domains where predictive transparency is most critical for operational decision-making.

The analysis shows that crop yield prediction and plant disease prediction are the most saturated themes, each accounting for 24.2% (n=8) of the total literature. These domains are followed by salinity stress prediction at 18.2% (n=6), highlighting growing concern over soil degradation in vulnerable regions like the Nile Delta. When specific drought and water-stress studies are consolidated, they represent 15.2% (n=5) of the research focus. The remaining studies are distributed among irrigation and climate-risk forecasting (9.1%, n=3), general XAI reviews (6.1%, n=2), and crop recommendation (3.0%, n=1). This grouped distribution aligns with the five primary climate-risk domains defined in the review's objectives (Section 1.2), providing a balanced roadmap of current research priorities in the field.

3.1 Crop Yield Prediction Performance and Drivers

Recent literature indicates a definitive shift from tree-based ensemble models (e.g., LightGBM, Random Forest) toward deep learning and hybrid architectures. This transition is driven by the inherent limitations of tree-based models in modeling non-Euclidean spatial dependencies and long-term temporal sequences critical to the crop lifecycle. Deep learning frameworks, such as CropNet-XAI and multi-head architectures offer superior capacity for capturing these complex spatial-temporal dynamics. Furthermore, recent advancements in ensemble and deep neural network (DNN) designs have demonstrated an exceptional ability to estimate complex non-linear relationships, establishing new benchmarks for predictive fidelity.

Table 2: Summary of Model Performance and Yield Drivers

Model Type	Primary Environmental Drivers Identified	Performance Metrics
Random Forest (RF) and Deep Neural Networks (DNN)	Temperature, humidity, and soil quality	R ² :: 0.975; RMSE: 0.029 (Kaur et al., 2026)
CropNet-XAI (1D-CNN)	Fertilizer use, rainfall, and irrigation status	93.46% test accuracy (Al Mahmud Joy et al., 2026)
MHCNN-LSTM-MHA	Precipitation, solar radiation, vapor pressure, soil pH, and clay content	RMSE: 3.75 bushels/acre; R ² : 0.905 (Cheema et al., 2026)
Hybrid LSTM-XGB	Multi-seasonal satellite archives and historical climate indicators	MSE: 0.0491 (Osibo, 2023; cited in Rajbongshi et al., 2026)
LightGBM-SHAP	Historical rainfall, temperature, humidity, and soil nutrients (NPK)	Effective modeling of nonlinear interactions (Dangi & Rathore, 2025)
Hybrid Ensemble (LightGBM/CatBoost)	Rainfall, fertilizer application, and irrigation	High accuracy and well-calibrated (>91%) (Gupta et al., 2026)

Table 2 analysis across these models consistently identifies temperature, humidity, and soil quality as the most influential predictors often accounting for more than half of the total explanatory power.

3.2 Salinity Stress Mapping and Hydrological Drivers

Studies conducted in the Nile Delta, Minab plain and the Karkheh River basin demonstrate that soil salinity is governed by a multifaceted interaction of topographic, chemical, and hydrological factors (Sahoo et al., 2025; Mohammadifar et al., 2022; Amini et al., 2026). Stacking ensembles and 1D-CNN-LSTM hybrids and Gradient Boosting Regressors have proven highly effective at projecting salinity intensification, both spatial and temporal scales, particularly under varying climate and discharge scenarios.

- Core Environmental Drivers:
 - Hydrological & Climatic: Precipitation (Sahoo et al., 2025), evaporation rates, and vertical distance to water channels (Mohammadifar et al., 2022), river discharge (Amini et al., 2026).
 - Hydro-Chemical: Sodium and sulfate dominant contributors to total dissolved solids (Amini et al., 2026).
 - Soil & Topographic: Clay content, sand content, cation exchange capacity, elevation, and soil pH (Mohammadifar et al., 2022; Sahoo et al., 2025; Amini et al., 2026).
- Predictive Success: In the Nile Delta, hybrid stacking designs successfully identified north-eastern and eastern zones as high-risk areas for future salinity intensification by 2030 (Sahoo et al., 2025). Complementing these spatial mappings, longitudinal analyses using 50-year monthly archives

demonstrate that interpretable ML can reveal the process-relevant hydrological controls of salinity over time (Amini et al., 2026).

3.3 Drought Forecasting and Physiological Indicators

Drought forecasting increasingly relies on identifying physiological stress markers and atmospheric demand. High-performance models like SAM-L and CatBoost integrate meteorological indices with vegetation response data to delineate stress levels (Ozupek et al., 2025; Alsanoosy & Malik, 2025).

Top Drought Predictors:

- Temporal Persistence: Lagged indices (e.g., SPEI-12) representing historical drought memory (Ozupek et al., 2025).
- Physiological State: Chlorophyll content, soil moisture levels, rooting conditions, soil-water retention, and land-surface properties (Alsanoosy & Malik, 2025; Tanriverdi & Batmaz, 2025), leaf water potential and shoot-growth (Costanza et al., 2026).
- Root-Zone Profiling Indicators: NDVI and solar radiation, identified as critical drivers for depth-specific soil moisture dynamics (Alahmad et al., 2026).
- Atmospheric Demand: Temperature and wind speed (Ozupek et al., 2025).

3.4 Crop and Plant Disease Prediction

Vision-based architectures have revolutionized the identification of pathologies by isolating biological markers such as symptomatic lesions, chlorotic streaks, and necrotic patches. Recent research emphasizes a shift toward architectures that combine high diagnostic precision with visual transparency to support real-world agricultural adoption (Potharaju et al., 2026; Peyal et al., 2024).

- Diagnostic Performance: CNN-SVM hybrids have achieved 98.34% accuracy for corn leaf diseases while utilizing only 0.95 million parameters, making them highly efficient for edge-based or mobile deployment (Peyal et al., 2024). Similarly, the ViT-B/16 model with transfer learning provides superior identification of spatially distributed symptoms by capturing long-range spatial relationships across the leaf surface (Potharaju et al., 2026).
- Architectural Robustness: Ensemble CNN frameworks (integrating VGG16, ResNet101V2, and InceptionV3) improve classification reliability across varying lighting conditions, leaf textures, and infection stages (Oad et al., 2024).
- Visual and Temporal Explainability: Tools like Grad-CAM and GradCAM++ generate class-specific heatmaps to highlight pathological regions, while the PalmRachis-BiLSTM-Attn framework provides "where–when" explanations by modeling disease progression along biologically meaningful anatomical pathways (Alghamdi et al., 2026; Potharaju et al., 2026).

- **Pathology Identification:** Explainable frameworks have successfully isolated specific diagnostic symptoms for Angular Leaf Spot, Bean Rust, and Maize Streak Virus achieving 96% accuracy. Their use of SHAP ensured that predictions were grounded in chlorotic streak patterns rather than image artifacts (Kalyango & Ntanda, 2023).

3.5 Actionable Decision Support and Recommendation Outcomes

Modern XAI frameworks translate predictive data into tangible agronomic advice moving beyond simple classification toward operational decision support:

- **Irrigation Optimization:** Depth-specific water application advice based on root-zone profiling where XAI identifies NDVI as a key indicator for shallow soil moisture and solar radiation for deeper root-zone dynamics (Alahmad et al., 2026).
- **Nutrient Management:** XAI-enabled systems provide fertilizer adjustment recommendations (e.g., nitrogen, phosphorus, and potassium) tailored to specific soil nutrient profiles (Kumar & Kumar, 2025).
- **Infrastructure Priorities:** Prioritization of land-reclamation and drainage investments in salinity-prone zones (Sahoo et al., 2025).

4. Technical Evaluation of XAI Techniques

This section provides a rigorous critique of the XAI methodologies employed in agricultural research, evaluating their technical stability, computational costs, and interpretability.

4.1 SHAP and LIME

The deployment of post-hoc local explanations in agronomic modeling is dominated by SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations). While both frameworks seek to decompose model predictions into individual feature contributions, they diverge significantly in their mathematical rigor and reliability for decision-support systems.

Mathematical Foundations and Feature Attribution SHAP is rigorously grounded in Cooperative Game Theory, utilizing Shapley values to provide a unique solution for additive feature attribution. SHAP satisfies four fundamental axioms Efficiency, Symmetry, Dummy, and Additivity, which ensure that the total model output is fairly distributed among features. This axiomatic foundation provides a guarantee of global consistency that is absent in other post-hoc methods making it the preferred choice for complex salinity models where feature interactions are dense (Kaur et al., 2026; Sahoo et al., 2025). Conversely, LIME operates by training local surrogate models typically weighted linear approximations around specific instances. Since LIME lacks a theoretical guarantee of consistency, its explanations are purely local and do not necessarily align with the model's global logic.

The primary bottleneck for SHAP is its computational complexity, particularly when calculating the exact Shapley values for high-dimensional deep learning ensembles or multi-layered remote sensing data. As noted in recent salinity and yield modeling (Sahoo et al., 2025), the exponential nature of feature permutations requires significant overhead often necessitating the use of KernelSHAP or TreeSHAP approximations which can still be intensive when processing massive spatio-temporal datasets.

LIME’s technical limitations center on its inherent instability. The linear surrogate is highly dependent on the choice of the proximity kernel and the sampling distribution used to generate the local neighborhood. Small perturbations in these hyperparameters can lead to fundamentally different linear approximations for the same data instance. This sensitivity often results in inconsistent results across similar agricultural instances, undermining the trust required for field-level intervention (Al Mahmud Joy et al., 2026) .

Table 3: Technical Trade-offs: SHAP vs. LIME

Feature	SHAP	LIME	Supporting Sources
Consistency	High. It is rooted in cooperative game theory and satisfies mathematical properties like consistency, local accuracy, and additivity. It ensures feature values are fairly distributed across all possible feature coalitions.	Lower. It approximates model behavior locally using a simple linear surrogate on perturbed samples. This makes it subject to the Rashomon Effect, where minor changes in data or models can lead to significantly different or unstable explanations.	(Patil & Kinger, 2025; Sahoo et al., 2025; Al Mahmud Joy et al., 2026)
Computational Speed	Lower. It is identified as computationally expensive because it must calculate the marginal contribution of each feature across all possible combinations. This results in higher overhead for high-dimensional agricultural datasets.	Higher. It is computationally simpler and faster. It avoids complex combinations by training a basic interpretable model only on a localized set of samples near a specific instance.	(Patil & Kinger, 2025; Sahoo et al., 2025 ; Al Mahmud Joy et al., 2026)
Robustness	High. It offers a holistic view by providing both local and global interpretability. It is cited as the most effective algorithm for capturing complex interactions in tasks like soil salinity prediction.	Moderate. While excellent for human-centered, instance-level explanations, it lacks the global robustness of SHAP and is weaker for deriving broad conclusions across an entire dataset.	(Sahoo et al., 2025; Patil & Kinger, 2025; Al Mahmud Joy et al., 2026)

4.2 Global Feature Attribution and Response Patterns (ALE/PDP)

Agro-meteorological modeling frequently utilizes Feature Importance rankings and Partial Dependence Plots (PDP) to visualize the directional response of the model to specific inputs. These methods attempt

to isolate the marginal effect of a single variable by averaging the model's predictions over the distribution of all other features.

Despite their limitations, these methods serve as effective "first-level explanations" for identifying high-stakes agronomic signals across different risk domains.

- **Plant-Stress Dynamics:** Feature importance analysis identifies chlorophyll content as the most influential indicator of photosynthetic health, followed by soil moisture as a water-availability marker (Alsanoosy & Malik, 2025).
- **Nutrient Management:** The ranking of nitrogen, phosphorus, and potassium (NPK) allows black-box recommendations to be translated into actionable advice regarding soil fertility and nutrient adjustments (Kumar & Kumar, 2025).
- **Root-Zone Dynamics:** Recent studies reveal that NDVI serves as a dominant indicator for shallow soil moisture, while solar radiation drivers signal a need for scheduling adjustments to counter deeper subsoil depletion.
- **Temporal Memory:** These rankings distinguish between immediate climatic fluctuations and long-term trends by highlighting the influence of lagged drought-index values (e.g., $SPEI(t-1)$) in driving model outputs (Ozupek et al., 2025).

A critical technical failure occurs when ALE and PDP are applied to highly coupled data environments. Agricultural systems are defined by intense multicollinearity among meteo-hydrological variables and coupled drivers. Since marginal distribution methods like PDP assume feature independence during the averaging process, they inadvertently create "mathematical artifacts" by calculating model outputs for feature combinations that are physically impossible in an agroecosystem, such as extreme atmospheric demand in the absence of corresponding energy inputs. This masks the non-linear interactions and feedback loops specifically those between soil moisture and atmospheric demand that are essential for forecasting multivariate climate disasters (Alahmad et al., 2026).

Technical Failures in Coupled Data Environments

- **Marginalization Over Impossible States:** PDPs calculate averages across feature combinations that cannot coexist in nature leading to logically flawed response curves.
- **Dilution of Attribution:** Multicollinearity causes the importance of primary coupled drivers to be split among redundant variables, obscuring the true physical cause of plant stress.
- **Absence of Structural Context:** Unlike design-based approaches such as PalmRachis, these marginal methods cannot represent the biological or physical "anatomical" progression of a climate-risk event.
- **Insufficient Nonlinearity Capture:** These methods fail to visualize the cascading interactions critical to root-zone dynamics and climate-risk forecasting.

4.3 Visual Explainability: Grad-CAM and Attention Mechanisms

Visual interpretability in agricultural deep learning distinguishes between spatial localization (identifying where a symptom is) and global contextual reasoning (understanding the relationship between scattered symptoms) to bridge the gap between model accuracy and clinical diagnosis (Alghamdi et al., 2026).

- **CNN and Grad-CAM (Spatial Localization):** Grad-CAM enhances transparency in CNN-based models by utilizing gradient information from the final convolutional layers to produce post-hoc heatmaps (Peyal et al., 2024). This method is highly effective for identifying localized symptoms, such as lesions in CNN-SVM hybrids (Peyal et al., 2024). However, as a post-hoc method, it primarily reveals where the model looked but does not always explain the underlying biological reasoning process (Alghamdi et al., 2026).
- **Vision Transformers (ViT) and Attention (Contextual Reasoning):** Unlike the local filters of CNNs, ViTs divide images into a sequence of patches and utilize self-attention mechanisms to capture long-range dependencies and global contextual patterns across the entire leaf surface (Potharaju et al., 2026). While this allows for global patch-level reasoning, the sources caution that attention does not always equal explanation as attention weights may reflect statistical artifacts or data bias rather than causal biological symptoms (Alghamdi et al., 2026).
- **Hybrid Synergies:** Advanced frameworks combine MobileNetV2 and ViT architectures to integrate efficient local symptom detection with broader visual context which is critical for identifying spatially distributed or subtle discoloration patterns (Özüpak et al., 2025).
- **Spatial-Temporal "Where-When" Explanations:** The PalmRachis-BiLSTM-Attn framework represents a technical peak by modeling disease as a spatial-temporal progression guided by the anatomical ordering of the rachis. It links spatial localization (Grad-CAM) with temporal-attention timelines to explain both where symptoms appear on the leaf and when specific anatomical segments (rachis-aligned vascular pathways) drive the disease progression (Alghamdi et al., 2026).

4.4 Uncertainty-Aware XAI and Confidence Estimation

Advanced models now integrate uncertainty quantification to move beyond deterministic 'Black Box' predictions. The application of DeepQuantreg to establish 95% confidence limits (Mohammadifar et al., 2022) and Monte Carlo Dropout to generate 95% prediction intervals (Tanriverdi & Batmaz, 2025) is technically superior to standard point estimates. By distinguishing between high-confidence risks and areas of high uncertainty, these approaches provide a more rigorous basis for high-stakes decisions in environmental management such as soil reclamation and regional drought adaptation (Mohammadifar et al., 2022).

4.5 Technical Synthesis of Model Constraints

The current evaluation of XAI methodologies highlights a profound Operational Disconnect between post-hoc experimental visualizations and the requirements for deployable field systems (Alghamdi et al., 2026).. Most frameworks treat interpretability as a secondary validation layer an "after-the-fact" diagnostic rather than an integrated component of the model architecture (Alghamdi et al., 2026). While

post-hoc tools like SHAP and LIME provide descriptive utility, they remain structurally disconnected from the underlying biological processes (Patil & Kinger, 2025). This is contrasted by emerging design-integrated approaches such as the PalmRachis framework which utilize attention mechanisms to capture "anatomical meaning" and the temporal progression of stress representing a more technically robust path forward (Alghamdi et al., 2026).

Technical Summary Methodological stagnation in the field is evidenced by a persistent prioritization of predictive accuracy metrics (e.g., R^2 and RMSE) over the technical stability and agronomic validity of the explanations provided(Cheema et al., 2026). While current models achieve high statistical precision, the reliance on unstable local surrogates or flawed marginal distribution methods for meteo-hydrological variables limits their utility as trustworthy decision-support tools (Alghamdi et al., 2026). Without transitioning toward design-integrated XAI and addressing the instabilities in post-hoc feature attribution, these systems will remain "lab-bound" diagnostic exercises rather than robust frameworks for managing agricultural climate risk.

4.6 Comparative Technical Summary Table

Table 4: Comparative Technical Summary of XAI Techniques

XAI Technique	Primary Agricultural Application	Technical Strength	Critical Limitation
SHAP	Yield & Salinity Mapping	Additive feature attribution; global consistency.	Computationally expensive for large ensembles.
LIME	Crop Recommendations	Intuitive local surrogates for farmers.	Instability due to sampling sensitivity.
ALE/PDP	Drought Forecasting	Visualizes directional response of variables.	Fails under high multicollinearity of climate data.
Grad-CAM	Disease Detection	Precise spatial localization of symptoms.	Post-hoc; sensitive to final layer architecture.
Attention	Disease Progression	Captures global context and temporal dynamics.	Focus does not strictly equal causal reasoning.
DeepQuantreg	Salinity Risk Mapping	Epistemic uncertainty/confidence estimation.	Requires dense, high-quality ground-truth data.

5. Synthesis of Research Gaps in Agricultural XAI

5.1. Critical Imbalance: Predictive Accuracy vs. Explanation Quality

A pervasive methodological stagnation in current agricultural climate-risk research is the prioritization of predictive performance metrics over the agronomic validity and actionability of Explainable AI (XAI) outputs. While modern deep learning and ensemble frameworks achieve unprecedented statistical precision reporting high R^2 , RMSE, and F1-scores they frequently relegate interpretability to a secondary validation role.

In models like CropNet-XAI (Al Mahmud Joy et al., 2026) and the soybean yield ensemble (Patil & Kinger, 2025), SHAP and LIME are treated as "after-the-fact" diagnostics rather than core components that shape the model's recommendation logic or field-level advice.

Frameworks like MHCNN-LSTM-MHA achieve exceptional precision ($R^2 = 0.905$), yet the resulting explanations remain analytical. They identify influential predictors like precipitation and clay content but fail to convert these attributions into operational prescriptions such as irrigation scheduling or planting-window adjustments.

A significant gap remains in stakeholder-centered evaluation. Current literature consistently fails to involve agronomists or farmers in the review of XAI outputs (Patil & Kinger, 2025; Hosseinpour et al., 2025). Without this human-in-the-loop verification, the explanations provided by XAI remain academic exercises rather than operational insights. This stakeholder gap is further intensified by a pronounced Global North bias and a lack of consideration for low-literacy farmers in developing regions (Rajbongshi et al., 2026).

5.2. The Post-hoc Limitation: Interpretation vs. Core Design

There is a fundamental divide between models that apply interpretability as an "after-the-fact" external layer and those that embed it directly into the model's architecture. Most current systems rely on post-hoc methods to describe what influenced a prediction rather than influencing the models internal reasoning logic.

Hybrid frameworks utilizing LightGBM, XGBoost, and CatBoost (Gupta et al., 2026) provide transparency through feature-attribution but lack an inherent understanding of the biological processes they model. These "black-box ensembles" are descriptively useful but structurally disconnected from agronomic reality.

Emerging approaches like the PalmRachis-BiLSTM-Attn framework (Alghamdi et al., 2026) represent a significant shift by reorganizing image features according to the anatomical structure of the date-palm rachis and using attention mechanisms. This model achieves "anatomical meaning" which provides

"where- when" explanations that align with the biological progression of disease rather than just statistical correlation.

Despite these architectural advances, even design-integrated models fail to transition into "prescriptive recommendation systems". As identified in the analysis of the SAM-L framework(Alsanoosy & Malik, 2025) models can identify stress drivers but they lack the decision logic required to issue confident, actionable management instructions

5.3. Fragmented Intelligence: Isolated Risk Domains vs. Compound Realities

The current research landscape is hampered by a "silo effect" wherein climate risks such as yield loss, drought, salinity, and disease are modeled as independent variables. This ignores the systemic feedback loops of real-world agroecosystems. Current XAI provides, at best, "univariate explanations for multivariate disasters".

A farmer does not manage risks in "academic categories" yet modeling remains fragmented. Isolated salinity mapping in the Nile Delta (Sahoo et al., 2025) or drought forecasting in Türkiye (Ozupek et al., 2025) fail to map the cascading feedback loops inherent in climate stress. For instance, water scarcity leads to the use of marginal irrigation which intensifies root-zone salinity; this in turn reduces a crop's disease resistance. Current XAI cannot yet reveal these interactions between soil moisture, evaporation, and nutrient uptake across risk domains. A shift toward "systems-oriented XAI" that treats the agroecosystem as an integrated complex whole is necessary.

5.4 Limited Model Generalizability Across Crops, Regions, and Climatic Zones

XAI models in agriculture face a severe generalizability barrier because their interpretability logic is often grounded in localized datasets shaped by specific soil, climate, and management interactions. Most existing frameworks are regionally bounded, such as soybean models in India, salinity mapping in the Nile Delta, and wheat-growth modeling in Iran. These models often rely on drivers like specific coastal seawater intrusion or surface soil moisture that do not transfer reliably to different geographic or industrial contexts.

Furthermore, explainable models show significant temporal and visual sensitivity; drought and disease explanations often depend on specific crop-growth stages, lighting conditions, or anatomical structures like the palm rachis which limits their universal application. A critical technical failure occurs when a model maintains predictive accuracy in a new region while its explanation stability fails leading to agronomically invalid feature rankings.

5.5. Deficiency in Uncertainty Communication

There is a critical deficiency in communicating confidence levels within high-stakes agricultural decision-making. Most models report "black-box" point predictions, which users often perceive as absolute truths rather than stochastic estimates.

The DeepQuantreg salinity study (Mohammadifar et al., 2022) is the rare exception that proves this rule of deficiency. By providing 95% confidence limits, it allows decision-makers to distinguish between high-risk/low-certainty and high-risk/high-certainty scenarios. Without this distinction, XAI outputs can lead to risky resource allocation or incorrect chemical applications, as seen in the critique of recent yield and disease detection studies (Peyal et al., 2024; Cheema et al., 2026). In an agricultural context an explanation without a confidence interval is fundamentally incomplete and potentially hazardous.

6. Future Research Directions for XAI in Agriculture

The current state of agricultural artificial intelligence has reached a critical juncture where high predictive accuracy no longer suffices for real-world efficacy. To achieve true climate-risk resilience, there must be a systematic transition from experimental "black-box" architectures toward operational, decision-support-oriented systems. This roadmap identifies and addresses the prevailing research gaps specifically the historical prioritization of predictive metrics over interpretability and the fragmentation of risks into isolated research silos. By establishing a framework for transparency and domain integration, we aim to provide the actionable intelligence necessary for field-level agronomic management under increasing climatic variability.

6.1 Integrated Explainability-by-Design Models

Future research should move beyond applying explainable artificial intelligence only after prediction and instead integrate explainability into the initial design of agricultural climate-risk models. Many reviewed studies use SHAP, LIME, Grad-CAM, feature importance, or attention mechanisms after model training to explain outputs produced by Random Forest, XGBoost, CNN, LSTM, Vision Transformer, or hybrid deep learning models. While this improves transparency, it often positions XAI as a post-hoc visualization layer rather than as a core component of the model architecture or decision process (Al Mahmud Joy et al., 2026; Özüpak et al., 2025; Patil & Kinger, 2025). Future models should therefore be designed to generate interpretable outputs from the beginning linking prediction, explanation, uncertainty, and recommendation in a single workflow. This would allow agricultural AI systems to explain not only what is likely to happen but also why it is happening and what action may be appropriate. Future research must also transition toward unified explainable-federated frameworks to establish a new standard for privacy-aware XAI in agriculture. By integrating Federated Learning (FL), these architectures enable decentralized diagnostics where multiple stakeholders collaboratively train shared global models without the need to transmit sensitive raw data thereby preserving data sovereignty (Shawon et al., 2026).

6.2 Domain-Specific Agricultural XAI Frameworks

XAI to achieve legitimacy in the field, its outputs must align with biological and physical laws rather than merely highlighting statistical correlations. Model reasoning must be validated against the physiological realities of crop growth such as interpreting nitrogen and chlorophyll within a nutrient-use efficiency context (Alsanoosy & Malik, 2025). For instance, in salinity modeling, feature selection should reflect soil chemical properties; recent findings by Sahoo et al. (2025) used Boruta feature selection to identify soil pH as a dominant predictor providing an agronomically sound basis for validating model reasoning against soil-texture and chemical dynamics. Future XAI frameworks should therefore incorporate crop physiology, soil–water dynamics, phenological stages, salinity processes, drought response, and disease-development pathways into the explanation process (Alsanoosy & Malik, 2025; Cheema et al., 2026; Hosseinpour et al., 2025; Sahoo et al., 2025). Explanations failing to reflect established plant-pathological pathways or hydrological principles must be deemed technically insufficient, regardless of the model's top-level predictive accuracy.

6.3 Stakeholder-Centered Explanation Design

The gap between technical XAI outputs (e.g., Grad-CAM heatmaps) and stakeholder action remains a significant barrier to adoption (Kumar & Kumar, 2025). Future systems must translate abstract feature-attribution scores into actionable insights that reflect specific interventions for farmers and extension officers. Future systems should translate model explanations into clear practical statements such as which factor is driving risk, whether the risk is water, salinity, disease, or nutrient related, and what management response may be considered. This is particularly important for crop recommendation, irrigation scheduling, plant-stress prediction, disease detection, and drought-risk forecasting, where explanation quality directly affects user trust and adoption (Kumar & Kumar, 2025). Visual tools must distinguish between image artifacts and genuine pathological markers (Alghamdi et al., 2026). Grad-CAM heatmaps should be overlaid with high-resolution imagery to assist farmers in early lesion identification (Peyal et al., 2024).

6.4 Holistic Multi-Risk Climate Modeling

Current research silos often ignore the interconnected nature of climate-related stresses. Integrated frameworks are required to model compound risks, where one stress factor intensifies another in a cascading fashion.

Table 5
Comparison of Isolated vs. Integrated Compound-Risk Approaches

Isolated Risk Domain	Integrated Compound-Risk Approach
Focuses exclusively on rainfall deficit or heat stress (Ozupek et al., 2025).	Models the cascading effect where drought necessitates reliance on marginal irrigation, intensifying root-zone salinity (Sahoo et al., 2025).
Explains yield loss through soil moisture variables only (Hosseinpour et al., 2025).	Explains how heat-induced evaporative demand compounds with nutrient deficiencies to accelerate yield decline (Al Mahmud Joy et al., 2026).
Treats disease detection as an isolated visual task.	Explains disease susceptibility as a function of environmental vigor and water-stress history (Alsanoosy & Malik, 2025).

Future research should develop integrated climate-risk models that combine yield prediction, salinity stress, drought forecasting, disease detection, irrigation demand, and crop recommendation within a unified framework. Future studies should therefore move toward multi-task learning, graph-based agroecosystem models, spatial–temporal risk modeling, and compound-risk forecasting systems that explain how multiple climate risks interact simultaneously.

6.5 Uncertainty-Aware Probabilistic Forecasting

Standardizing the communication of uncertainty is essential for high-stakes agricultural decisions. Following the DeepQuantreg benchmark for salinity mapping (Mohammadifar et al., 2022), future work in yield and drought forecasting (Al Mahmud Joy et al., 2026; Ozupek et al., 2025) must move beyond point estimates toward probabilistic outputs. Models must specify reliability using standardized benchmarks such as the 95% confidence limits to ensure that users can distinguish between certain risks and speculative predictions. Without this distinction, XAI outputs can lead to "hazardous" decisions such as incorrect chemical applications or risky resource allocation. Future systems should communicate these technical intervals in stakeholder-friendly categories like "high confidence" or "field verification required".

6.6 Cross-Region and Cross-Crop Transferability

Generalizability remains a major hurdle, as models for soybean (Patil & King, 2025) and wheat (Hosseinpour et al., 2025) are often bounded by their original geographic training zones. Future XAI models must adhere to a new standard of robustness that mandates cross-region and cross-season validation. A model is only considered mature if its explanations remain stable and agronomically valid when transferred across diverse cultivars, management histories, and soil textures. Transfer learning and domain adaptation must be utilized to ensure that intelligence generated in one climatic zone is applicable and interpretable in another.

6.7 Real-Time XAI-Enabled Decision-Support Systems

The field should transition from retrospective and laboratory-based modeling toward real-time XAI-enabled agricultural decision-support systems. Many current studies demonstrate model performance

using historical datasets, remote-sensing archives, field observations, or benchmark image datasets but fewer studies show how XAI is operationally integrated into real farming or policy workflows. Future systems should connect explainable models with real-time sensor networks, IoT devices, satellite platforms, weather forecasts, GIS dashboards, mobile applications, and extension-service platforms. Such systems could provide live guidance for irrigation scheduling, drought alerts, salinity-risk monitoring, disease diagnosis, crop suitability assessment, and yield-risk management (Alahmad et al., 2026; Alsanoosy & Malik, 2025; Kumar & Kumar, 2025; Özüpak et al., 2025).

6.8 Agronomic Validation of AI Explanations

The final phase of this roadmap mandates that AI "explanations" undergo rigorous validation against expert knowledge. The type of validation must match the explanation's complexity; for instance, simple spatial localization via Grad-CAM requires different verification than the temporal modeling of anatomical progression seen in frameworks like PalmRachis-BiLSTM-Attn (Alghamdi et al., 2026). Visual disease diagnostics must be verified by plant pathologists to ensure the model targets pathological markers rather than noise (Peyal et al., 2024) while salinity drivers must be cross-referenced with known hydrological and soil-chemical processes (Amini et al., 2026; Mohammadifar et al., 2022). This validation is necessary because computational explanations may appear plausible while still being unstable, biased, or weakly connected to real agronomic mechanisms.

7. Conclusion

This systematic review confirms that while AI achieves exceptional accuracy in agricultural climate-risk prediction, a "transparency deficit" remains a primary barrier to real-world adoption. Key findings reveal that current XAI frameworks often suffer from technical instability and a "silo effect" that treats interconnected risks like drought and salinity as isolated variables. To bridge the gap between technical precision and stakeholder trust, future research must transition toward "Explainability-by-Design" and unified federated-XAI frameworks that ensure data sovereignty while maintaining high performance. Furthermore, standardizing uncertainty-aware forecasting via the DeepQuantreg benchmark is essential to prevent "hazardous" field-level decisions based on deterministic point estimates. Ultimately, for AI to become a legitimate tool in safeguarding global food security, its explanations must undergo rigorous agronomic validation to ensure they reflect genuine physiological and hydrological realities.

Declarations

Data Availability: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare that there are no conflicts of interest regarding the publication of this manuscript.

Funding Statement: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contribution: Jacintha Menezes and Rishah Kant Jha conceived the study, developed the master search string, and conducted the primary screening and data extraction. Jacintha Menezes and Jia Bernadette D Souza performed the independent secondary screening, analyzed the technical benchmarks of XAI techniques, and prepared the figures. Jacintha Menezes supervised the research, resolved disagreements during the eligibility assessment, and reviewed the final manuscript for intellectual content. All authors have read and agreed to the published version of the manuscript

Ethics Statement: Not applicable.

Competing Interests: The author declares that there are no competing interests that could have influenced the research, analysis, or publication of this manuscript.

References

1. Al Mahmud Joy, A., Islam, Md. F., Islam, Md. S., Ahmed, S., Morshed, M. N., & Ahmed, H. (2026). CropNet-XAI: An Explainable 1D-CNN Framework for Transparent Crop Yield Prediction. *2026 5th International Conference on Electrical, Computer & Telecommunication Engineering (ICECTE)*, 1–6. <https://doi.org/10.1109/ICECTE69292.2026.11429350>
2. Alahmad, T., Neményi, M., & Nyéki, A. (2026). Explainable XGBoost-based models of root-zone soil moisture profiling using coupled Sentinel-2 and IoT data in loam and silt loam soil. *Discover Applied Sciences*, 8(6), 672. <https://doi.org/10.1007/s42452-026-08673-3>
3. Alghamdi, M., Abadleh, A., Mnasri, S., Alrashidi, M., & Sulaiman Alkhazi, I. (2026). PalmRachis-BiLSTM-Attn: An Anatomically Guided Explainable Deep Learning Framework for Spatial–Temporal Progression Modeling of Date-Palm Leaf Diseases. *IEEE Access*, 14, 33636–33650. <https://doi.org/10.1109/ACCESS.2026.3669421>
4. Alsanoosy, T., & Malik, J. A. (2025). Predicting plant stress using SAM-L: novel self-adaptive-meta learner with XAI based on soil moisture and chlorophyll analysis. *Scientific Reports*, 15(1), 42304. <https://doi.org/10.1038/s41598-025-26184-w>
5. Amini, H., Shakeri, R., Ghaderi, N., Fakheri, F., Morovati, K., Zahraie, B., & Ahmadian, R. (2026). Interpretable machine learning for river salinity dynamics in arid basins. *Scientific Reports*, 16(1), 18310. <https://doi.org/10.1038/s41598-026-49042-9>
6. Bouregaa, T. (2025). Comparative evaluation of machine learning models for regional agricultural drought prediction in Algeria using SHAP analysis. *Natural Hazards*, 121(20), 24275–24346. <https://doi.org/10.1007/s11069-025-07719-w>
7. Cheema, I. A., Hanif, M. K., Khokhar, A. M. A., Rehman, M. A. U., & Zaheer, R. (2026a). Enhancing crop yield prediction accuracy with a novel interpretable deep learning model: MHCNN-LSTM-MHA. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-53616-y>

8. Cheema, I. A., Hanif, M. K., Khokhar, A. M. A., Rehman, M. A. U., & Zaheer, R. (2026b). Enhancing crop yield prediction accuracy with a novel interpretable deep learning model: MHCNN-LSTM-MHA. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-53616-y>
9. Costanza, L., Pedrero Salcedo, F., Marti Martinez, C. M., Lorente, B., Maldera, F., Lopriore, G., & Garofalo, S. Pietro. (2026). Use of explainable artificial intelligence with high-resolution satellite imagery to assess water status and vine shoot growth in grapevine (*Vitis vinifera* L.). *Modeling Earth Systems and Environment*, 12(1), 25. <https://doi.org/10.1007/s40808-025-02656-w>
10. Gupta, M., Jain, S., & Kumar, M. (2026). A Crop-Wise Hybrid Ensemble Learning Framework for Accurate and Explainable Crop Yield Prediction. *2026 International Conference on Computing, Sciences and Communications (ICCSC)*, 1–6. <https://doi.org/10.1109/ICCSC67078.2026.11468595>
11. Hosseinpour, S., Pirdashti, H., Hosseinpour, D., Mousavi, H., & Mohammadpour, S. (2025). Machine and deep learning-based wheat yield prediction: the critical role of soil moisture and remote sensing data. *Modeling Earth Systems and Environment*, 11(4), 275. <https://doi.org/10.1007/s40808-025-02456-2>
12. Kalyango, M. F., & Ntanda, K. M. (2023). Interpretable Deep Learning for Diagnosis of Maize Streak Disease. *2023 First International Conference on the Advancements of Artificial Intelligence in African Context (AAIAC)*, 1–6. <https://doi.org/10.1109/AAIAC60008.2023.10465315>
13. Kaur, G., Ashanand, Singh, R., Kamal, P., Dua, R., & Kukreti, K. (2026). Explainable Artificial Intelligence for Crop Yield Prediction Using Agro-Climatic Data. *2026 2nd International Conference on Big Data and Machine Learning, ICBTML 2026*. <https://doi.org/10.1109/ICBDML68582.2026.11544417>
14. Kumar, S., & Kumar, M. (2025). Developing an XAI-Based Crop Recommendation Framework Using Soil Nutrient Profiles and Historical Crop Yields. *IEEE Transactions on Consumer Electronics*, 71(2), 6950–6959. <https://doi.org/10.1109/TCE.2025.3569736>
15. Mohammadifar, A., Gholami, H., & Golzari, S. (2022). Assessment of the uncertainty and interpretability of deep learning models for mapping soil salinity using DeepQuantreg and game theory. *Scientific Reports*, 12(1), 15167. <https://doi.org/10.1038/s41598-022-19357-4>
16. Oad, A., Abbas, S. S., Zafar, A., Akram, B. A., Dong, F., Talpur, M. S. H., & Uddin, M. (2024). Plant Leaf Disease Detection Using Ensemble Learning and Explainable AI. *IEEE Access*, 12, 156038–156049. <https://doi.org/10.1109/ACCESS.2024.3484574>
17. Özüpak, Y., Alpsalaz, F., Aslan, E., & Uzel, H. (2025). Hybrid deep learning model for maize leaf disease classification with explainable AI. *New Zealand Journal of Crop and Horticultural Science*, 53(5), 2942–2964. <https://doi.org/10.1080/01140671.2025.2519570>
18. Ozupek, E., Teke, A., Celik, N., & Kavzoglu, T. (2025). Explainable artificial intelligence to explore the intrinsic characteristics of climatic parameters governing meteorological drought forecasting: opening the black box. *Stochastic Environmental Research and Risk Assessment*, 39(8), 3201–3222. <https://doi.org/10.1007/s00477-025-03007-y>

19. Patil, J., & Kinger, S. (2025). Explainable Ensemble Learning Framework for Soybean Yield Prediction: A Data-Driven Approach for Agricultural Resilience in India. *2025 9th International Conference on Computing, Communication, Control and Automation (ICCCBEA)*, 01–09. <https://doi.org/10.1109/ICCUBEA65967.2025.11283744>
20. Peyal, H. I., Mondal, Md. N. I., & Miraz, S. (2024). An Efficient Explainable AI Method Combining CNN and SVM for Corn Leaf Disease Detection and Visualization. *2024 27th International Conference on Computer and Information Technology (ICCIT)*, 2958–2962. <https://doi.org/10.1109/ICCIT64611.2024.11021795>
21. Potharaju, S., Singh, A., Singh, D., Tambe, S. N., Kantipudi, P. M., & Kiranmai, B. (2026). An explainable vision transformer model with transfer learning for accurate bean leaf disease classification. *Scientific Reports*, 16(1), 10402. <https://doi.org/10.1038/s41598-026-41723-9>
22. Rajbongshi, A., Johora, F. T., Hossain, A., Sarker, Md. S., Rahman, M. H., Rahman, M. W., Alotaibi, F. T., & Moni, M. A. (2026). Leveraging explainable AI for sustainable agriculture: a comprehensive review of recent advances. *Artificial Intelligence Review*, 59(3), 105. <https://doi.org/10.1007/s10462-025-11459-5>
23. Sahoo, S., Singha, C., & Govind, A. (2025). Enhanced soil salinity index prediction using hybrid stacking ensemble machine learning with explainable artificial intelligence (XAI) technique: a case study of the Nile Delta, Egypt. *Discover Applied Sciences*, 8(1), 36. <https://doi.org/10.1007/s42452-025-08038-2>
24. Shawon, S. M., Niha, F. L., Sinan, O. H., Abed, M. J., Al Qwaid, M., Sarker, M. T., & Zubair, H. T. (2026). Integrating Federated Learning and Explainable AI for Plant Disease Detection: A Systematic Review of Recent Advances in Precision Agriculture. *International Journal of Computational Intelligence Systems*, 19(1), 138. <https://doi.org/10.1007/s44196-026-01219-w>
25. Tanriverdi, I., & Batmaz, İ. (2025). AI-driven U.S. drought prediction using machine learning and deep learning. *Climate Dynamics*, 63(6), 249. <https://doi.org/10.1007/s00382-025-07720-w>

Figures




Figure 1

PRISMA flow diagram




Figure 2

Shortlisted papers by year of publication




Figure 3

Distribution of shortlisted papers by theme