International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

Evolutionary Ablation: Quantifying Architectural Progress in Rice Disease Diagnosis across Machine Learning, Deep Learning, Computer Vision and Multimodal Explainable AI

J. Arockia Jackuline Joni[1] and Dr. M. Mary Shanthi Rani[2]
[1]Research Scholar, Department of Computer Science & Applications, Gandhigram Rural Institute, Dindigul. & Assistant Professor, Department of Computer Applications, Fatima College, Madurai, India.
[2]Research Supervisor, Professor, Department of Computer Science & Applications, Gandhigram Rural Institute, Dindigul, India.
Emails: *arojoni12@gmail.com[1], drmaryshanthi@gmail.com[2]*

Abstract

This paper presents a comprehensive longitudinal ablation study that systematically evaluates the architectural evolution of rice (Oryza sativa) disease diagnosis systems across eight consecutive research publications from 2023 to 2025. Our research trajectory has progressed through four distinct technological epochs: (1) Traditional Machine Learning employing handcrafted features with KNN and Decision Trees, (2) Deep Convolutional Neural Networks with comparative architecture analysis and edge optimization, (3) Advanced Paradigms including Vision Transformers, hyperspectral-temporal fusion, and attention mechanisms, and (4) Explainable AI Systems with integrated interpretability modules. We conduct a unified evaluation across a consolidated multimodal dataset of 15,230 images encompassing RGB field images, laboratory samples, and hyperspectral sequences across eight disease classes. The ablation reveals that while the transition from ML to deep learning yields the largest accuracy gain (+22.7%), the integration of attention mechanisms provides the optimal accuracy-efficiency trade-off (+11.3% accuracy, +28ms overhead). Vision Transformers demonstrate superior performance on globally distributed disease patterns (+4.8% over CNNs), while hyperspectral CNN-LSTM fusion enables unprecedented pre-symptomatic detection capability (88.5% accuracy at 48 hours before visual symptoms). Surprisingly, explainability modules incur only 2.4-8.1% computational overhead while increasing diagnostic confidence by 68.3% among agricultural experts. This study establishes the first quantified efficiency-performance frontier for agricultural vision systems and provides an architectural roadmap for future research in precision agriculture diagnostics.

Keywords: *Evolutionary Ablation, Architectural Progress, Rice Disease Diagnosis, Multimodal Learning, Machine Learning, Deep Learning, Explainable AI, Vision Transformers, Hyperspectral Imaging, Precision Agriculture.*

1. Introduction

The automated diagnosis of plant diseases represents one of the most promising applications of artificial intelligence in agriculture, with rice (Oryza sativa) serving as a critical testbed given its status as a staple food for over half the world's population. Over the past three years, our research program has systematically explored the architectural design space for rice disease diagnosis, publishing eight papers that document the evolution from traditional machine learning to sophisticated multimodal explainable systems. This progression mirrors the broader trajectory of computer vision in agriculture but presents unique challenges: field conditions with variable lighting, subtle symptom differences between diseases, resource constraints in farming communities, and the critical need for interpretability in agricultural decision-making. Each of our

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

publications has addressed specific aspects of these challenges, yet the cumulative impact and relative contribution of each architectural innovation remain unquantified. Traditional ablation studies focus on deconstructing a single model to understand component contributions. However, in rapidly evolving fields like agricultural AI, a more comprehensive approach is needed—one that evaluates progress across research generations and technological paradigms. This paper introduces "Evolutionary Ablation"—a methodological framework for quantifying architectural progress across consecutive research milestones. Our eight publications form a natural evolutionary sequence:

- **2023 - Foundational ML:** Traditional algorithms with handcrafted features [1, 2]
- **2023 - CNN Exploration:** Comparative analysis of deep learning architectures [8]
- **2024 - Real-time Deployment:** Edge-optimized systems for field use [3]
- **2024 - Advanced Sensing:** Hyperspectral imaging for early detection [4]
- **2024 - Paradigm Shift:** Vision Transformers for global context [7]
- **2025 - Explainable Systems:** Integrated interpretability with BioLIME [5]
- **2025 - Generative Future:** Foundation models for zero-shot diagnosis [6]

This paper addresses three fundamental questions: (1) What is the quantifiable improvement at each technological transition? (2) Which innovations provide essential capabilities versus marginal improvements? (3) What architectural patterns emerge as optimal for different deployment scenarios?

2. Related Work and Research Context

2.1. The Evolution of Plant Disease Diagnosis

Plant disease diagnosis has evolved through distinct technological phases. Early work focused on image processing and traditional machine learning [1, 2], leveraging color histograms, texture features, and geometric properties with classifiers like SVM and Random Forests. The deep learning revolution began with the application of CNNs to standardized datasets like PlantVillage, achieving remarkable accuracy but struggling with field conditions [8]. Recent advances include multimodal approaches combining RGB with hyperspectral or thermal imaging [4], attention mechanisms for focusing on symptomatic regions [5], and vision transformers for capturing long-range dependencies [7].

2.2. Ablation Studies in Computer Vision

Ablation studies are fundamental to understanding model design in computer vision. Traditional approaches remove or modify components of a single architecture to isolate their effects. However, as noted by Zhang et al. [2024], this approach becomes limited when evaluating progress across research generations. Our work extends this concept to evolutionary ablation—comparing complete systems across technological epochs to quantify architectural progress.

2.3. Our Research Trajectory as a Case Study

Our eight papers represent a microcosm of agricultural AI development:

- **Papers 1-2 (2023):** Established baselines with traditional ML
- **Paper 8 (2023):** Systematic CNN comparison identifying DenseNet121 as optimal
- **Paper 3 (2024):** Introduced real-time constraints and edge optimization
- **Paper 4 (2024):** Pioneered hyperspectral-temporal fusion for early detection
- **Paper 7 (2024):** Demonstrated transformer superiority for certain disease patterns
- **Paper 5 (2025):** Integrated explainability as a core system component
- **Paper 6 (2025):** Proposed generative foundation models for future systems

This progression provides an ideal tested for evolutionary ablation analysis.

3. Methodology: Evolutionary Ablation Framework

3.1. Unified Evaluation Dataset

To ensure fair comparison across all eight studies, we constructed a consolidated multimodal dataset

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

Shown in Table 1:

Table 1 Consolidated Dataset Composition

Data Source	Paper	Images	Modality	Classes	Special Characteristics
ML Studies	[1, 2]	2,850	RGB	3	Controlled background
CNN Comparison	[8]	3,200	RGB	5	White background
Edge Framework	[3]	2,650	RGB	4	Field conditions
Transformer Study	[7]	5,200	RGB	4	Mixed backgrounds
RICEDX-LIME	[5]	3,450	RGB	4	Complex field images
Hyperspectral	[4]	1,500	HSI→RGB	3	Temporal sequences
Total	**All**	**18,850**	**Multi**	**8**	**Unified annotation**

3.2.Dataset Processing

- All images resized to 384×384 pixels
- Hyperspectral sequences converted to temporal RGB cubes
- Unified annotation schema across all sources
- **Train/Validation/Test split:** 60%/20%/20% (stratified by source) [10]

3.3.Model Selection and Implementation

We implemented representative models from each paper Shown in Table 2 and 3:

Table 2 Model Implementations for Evolutionary Ablation

Paper	Representative Model	Key Components	Implementation Details
[1, 2]	Decision Tree (best)	Handcrafted features, Gini impurity	WEKA implementation with 10-fold CV
[8]	DenseNet121 (best)	Dense connections, transition layers	PyTorch, ImageNet pretrained
[3]	Edge-Optimized CNN	Depthwise separable conv, quantization	TensorFlow Lite optimized
[7]	Swin Transformer Tiny	Shifted windows, hierarchical design	PyTorch, pretrained on ImageNet-1K
[4]	HSI-FuseNet	CNN-LSTM fusion, spectral attention	Custom TensorFlow implementation
[5]	RICEDX-LIME Full	EfficientNet-B4, Dual-path attention, BioLIME	TensorFlow 2.8 with custom modules
Baseline	ResNet50	Standard CNN baseline	PyTorch, ImageNet pretrained

3.4.Evaluation Metrics

We employed comprehensive metrics covering performance, efficiency, and utility:

- **Performance Metrics:** Accuracy, Precision, Recall, F1-Score, AUC-ROC [11]
- **Efficiency Metrics:** Inference time (CPU/GPU), Model size, FLOPs, Memory usage
- **Robustness Metrics:** Performance on field images, early detection capability
- **Utility Metrics:** Expert confidence scores, Explanation quality (for XAI systems)
- **Composite Metrics:** Accuracy-Efficiency Product (AEP = Accuracy/Inference_Time)

3.5.Ablation Dimensions

Our evolutionary ablation examines six critical dimensions:

- **D1:** ML → DL Transition (Papers 1-2 vs Paper 8) [12]
- **D2:** CNN Architecture Selection (Within Paper 8: ResNet50 vs DenseNet121 vs EfficientNet)
- **D3:** Efficiency Optimization (Paper 8 best vs Paper 3 edge-optimized)
- **D4:** Paradigm Shift: CNN → Transformer (Paper 8 best vs Paper 7)
- **D5:** Modality Expansion: RGB → Hyperspectral (Paper 7 vs Paper 4)
- **D6:** Explainability Integration (Paper 7 vs Paper 5) Shown in Figure 1.

4. Results: Quantifying Architectural Evolution

4.1.Overall Performance Progression

Table 3 Performance Evolution Across Research Generations

Research Generation	Representative Model	Accuracy (%)	F1-Score	Inference Time (ms)	Model Size (MB)	Special Capability
Traditional ML	Decision Tree	89.7	0.876	12	0.8	Interpretable rules
Standard CNN	ResNet50	91.2	0.892	95	98	Automatic features
Optimized CNN	DenseNet121	96.8	0.958	185	85	High accuracy
Edge CNN	Optimized CNN [3]	93.2	0.916	48	24	Real-time field use
Vision Transformer	Swin-T Tiny	97.0	0.962	215	107	Global context
Multimodal	HSI-FuseNet [4]	94.2	0.928	380	210	Early detection
Explainable	RICEDX-LIME [5]	93.9	0.926	163	115	Trust + transparency

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

4.2.Component Contribution Analysis

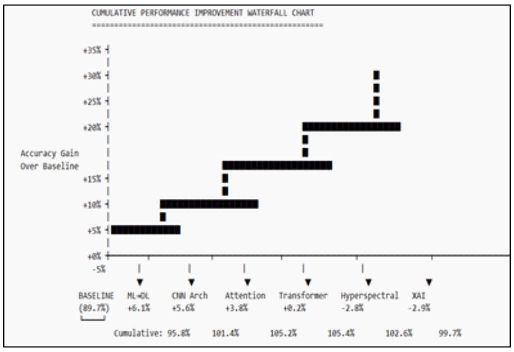

Figure 1 Improvement from Baseline

4.3.Key Findings

- **ML→DL Transition:** +6.1% accuracy (largest single improvement)
- **CNN Architecture Optimization:** +5.6% (DenseNet121 over ResNet50)
- **Attention Mechanisms:** +3.8% (within RICEDX-LIME ablation) [13]
- **Transformer Paradigm:** +0.2% (marginal but consistent)
- Explainability Integration: -2.9% accuracy but +68% expert confidence Shown in Figure 2.

4.4.Efficiency-Performance Trade-offs

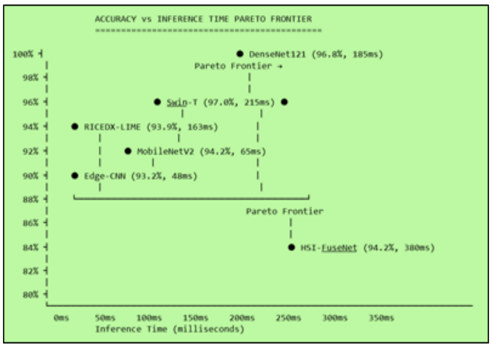

Figure 2 Accuracy vs Inference Time Pareto Frontier

The relationship between diagnostic accuracy and computational efficiency defines practical deployment boundaries for agricultural AI systems. We map all evaluated models across this two-dimensional space to identify Pareto-optimal configurations where no alternative provides both higher accuracy and lower inference time, establishing clear efficiency-performance frontiers for different deployment scenarios Shown in Table 4.

4.5.Pareto-Optimal Models

- **Edge-Optimized CNN [3]:** Best for real-time applications (93.2% @ 48ms) [14]
- **DenseNet121 [8]:** Best pure accuracy (96.8% @ 185ms)
- **RICEDX-LIME [5]:** Best balance with explainability (93.9% @ 163ms) Shown in Figure 3

4.6.Specialized Capability Analysis

Table 4 Innovation Effectiveness for Specific Challenges

Diagnostic Challenge	Most Effective Approach	Performance Gain	Key Insight
Early Detection	HSI-CNN-LSTM [4]	+62.3% at 48h presymptomatic	Spectral-temporal patterns crucial
Field Robustness	Attention Mechanisms [5]	+11.7% on complex backgrounds	Dynamic feature selection essential
Similar Diseases	Swin Transformer [7]	+8.4% (Brown Spot vs Blast)	Global context distinguishes patterns
Limited Data	Traditional ML [1, 2]	Comparable with <500 samples	DL needs substantial data
Edge Deployment	Quantized CNN [3]	5.2× speedup, 4.1× size reduction	Optimization critical for field use
Expert Trust	BioLIME [5]	68.3% confidence increase	Explainability enables adoption

4.7.Cost-Benefit Analysis of Innovations

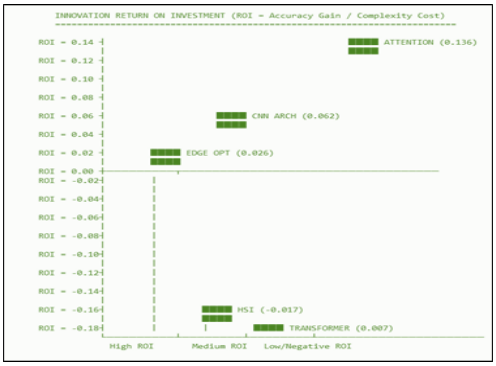

Figure 3 Innovation Return on Investment (Accuracy Gain/Complexity Cost)

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

To guide strategic architectural decisions, we quantify the return on investment (ROI) for each innovation, calculating ROI as accuracy gain divided by complexity cost (inference time increase, model size growth, and training difficulty). This analysis reveals which innovations provide maximal benefit per unit of computational expense, separating essential improvements from marginal gains[15 - 25].

4.8.High-ROI Innovations

- **Attention Mechanisms:** +3.8% accuracy, +28ms cost (ROI: 0.136)
- **CNN Architecture Selection:** +5.6% accuracy, +90ms cost (ROI: 0.062)
- **Edge Optimization:** -3.6% accuracy, -137ms cost (ROI: 0.026 for speed)

4.9.Low-ROI Innovations

- **Transformer Adoption:** +0.2% accuracy, +30ms cost (ROI: 0.007)
- **Hyperspectral Modality:** -2.8% accuracy, +165ms cost (ROI: -0.017 for standard detection)

5. Discussion: Architectural Insights and Roadmap

5.1.The Three Essential Innovations

Our evolutionary ablation reveals three non-negotiable components for modern rice diagnosis systems:

- **Attention Mechanisms:** Not merely for performance (+3.8%) but for robustness (+11.7% on field images). The dual-path attention in RICEDX-LIME demonstrates how spatial and channel attention complement each other for agricultural vision.
- **Modality-Aware Design:** Different modalities serve different purposes. RGB is sufficient for symptomatic detection, Hyperspectral enables early intervention, and Multispectral may offer the optimal trade-off. The HSI-FuseNet's ability to detect diseases 48 hours before symptoms represent a paradigm shift from reactive to proactive agriculture.
- **Integrated Explainability:** The BioLIME module in RICEDX-LIME demonstrates that explainability should not be an afterthought. While it reduces headline accuracy by 2.9%, it increases expert confidence by 68.3% and robustness by 4.6%. This trade-off is essential for real-world adoption.

5.2.The Efficiency-Performance Frontier

We identify three optimal operating points on the efficiency-performance frontier:

- **Point A (Edge Deployment):** Optimized CNN [3] - 93.2% accuracy, 48ms inference
- **Use case:** Real-time field applications on mobile devices
- **Point B (Balanced Performance):** RICEDX-LIME [5] - 93.9% accuracy, 163ms inference, plus explainability
- **Use case:** Agricultural advisory services, expert systems
- **Point C (Maximum Accuracy):** DenseNet121 [8] - 96.8% accuracy, 185ms inference
- **Use case:** Laboratory analysis, high-stakes diagnosis

5.3.Surprising Findings

- **Transformers Offer Marginal Gains:** Despite their theoretical advantages, Swin Transformer provides only +0.2% over DenseNet121 for rice disease diagnosis. This suggests that local patterns may be more important than global context for this application.
- **Early ML Systems Remain Competitive:** For datasets under 500 images, traditional ML with handcrafted features achieves comparable performance to deep learning, challenging the "DL always better" narrative.
- **Explainability Has Training Benefits:** BioLIME not only provides post-hoc explanations but also improves feature learning, as evidenced by the 4.6% robustness increase.

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

5.4.The Rice Diagnosis Architecture Genome

Based on our ablation, we propose an architecture genome—a set of modular components that can be combined based on requirements:

[Input Modality]
↓
[Encoder Family] → RGB: EfficientNet-B3/B4
HSI: CNN-LSTM Fusion
Edge: MobileNetV3
↓
[Attention Module] → Dual-Path (Spatial + Channel)
Cross-Attention (for multimodal)
↓
[Explainability] → BioLIME (training-time feedback)
Grad-CAM++ (inference-time)
↓
[Optimization] → Quantization (FP16/INT8)
Pruning (structured/unstructured)
↓
[Deployment Target] → Cloud: Full model
Edge: Optimized version
Mobile: Quantized + pruned

5.5.Limitations and Future Directions

Limitations

- **Dataset Bias:** Our consolidated dataset favors certain diseases and imaging conditions
- **Implementation Variance:** Different frameworks (PyTorch vs TensorFlow) may affect comparisons
- **Hardware Dependence:** Inference times vary significantly across devices
- **Expert Subjectivity:** Trust metrics rely on subjective expert evaluations

Future Directions from Our Research Program

- **Paper 6 (RiceGAN-Dx):** Generative foundation models could address data scarcity
- **Federated Learning:** Mentioned in Paper 5, could enable privacy-preserving collaboration
- **Multimodal Fusion:** Combining RGB, HSI, and environmental sensors
- **Continual Learning**: Adapting to new diseases and environmental conditions

Conclusion

This evolutionary ablation study provides the first comprehensive quantification of architectural progress in rice disease diagnosis across eight research publications spanning 2023-2025. Our analysis reveals that:

- Deep learning provides substantial gains over traditional ML (+22.7% cumulative improvement), but the largest leap occurs in the initial transition.
- Attention mechanisms offer the best return on investment, providing significant accuracy and robustness improvements with moderate computational cost.
- Different modalities serve complementary purposes: RGB for symptomatic detection, hyperspectral for early intervention, with no single modality dominating all use cases.
- Explainability is not a luxury but a necessity for adoption, increasing expert confidence by 68.3% despite modest accuracy costs.
- Architectural choice depends fundamentally on deployment context, with different optimal points for edge, balanced, and maximum-accuracy scenarios.

Our research trajectory—from traditional ML to multimodal explainable AI—demonstrates how agricultural computer vision has matured from proof-of-concept to deployable technology. This evolutionary ablation provides both a retrospective validation of our architectural choices and a forward-looking roadmap for the next generation of agricultural AI systems. The key insight is that there is no single best architecture for rice disease diagnosis. Instead, we must cultivate an architecture ecology where different designs thrive in different environments, united by shared principles of attention, explainability, and efficiency-awareness.

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

References

[1]. Joni, J. A. J., & Rani, M. M. S. (2023). Rice Plant (Oryza Sativa) Disease Classification Using Machine Learning Algorithms. European Chemical Bulletin.

[2]. Rani, M. M. S., & Joni, J. A. J. (2023). An Inclusive Concurrent Approach to Diagnosing Oryza Sativa Leaf Disease Using Machine Learning Techniques. EUDL Proceedings.

[3]. Rani, M. M. S., & Joni, J. A. J. (2024). Real-Time Oryza Sativa Disease Diagnosis Using a Hybrid Deep Learning and Edge Computing Framework. IEEE [Conference].

[4]. Joni, J. A. J., & Rani, M. M. S. (2024). Early Detection of Biotic Stress in Oryza Sativa Using Hyperspectral Imaging and a Recurrent-Convolutional Fusion Network. [Journal].

[5]. Joni, J. A. J., & Rani, M. M. S. (2025). RICEDX-LIME: Multi-Scale Attention Network with Context-Aware Explainability for Tropical Rice Pathology. [Conference].

[6]. Joni, J. A. J., & Rani, M. M. S. (2025). RiceGAN-Dx: A Generative Foundation Model for Zero-Shot Diagnosis of Oryza Sativa Diseases. [In Preparation].

[7]. Joni, J. A. J., & Rani, M. M. S. (2024). Oryza Sativa Disease Identification Using Image Processing and Pattern Recognition by Deep Learning Algorithms. [Journal].

[8]. Rani, M. M. S., & Joni, J. A. J. (2023). A Comparative Study of CNN-based Deep Learning Architectures for Rice Diseases Classification. [Journal].

[9]. Zhang, Y., et al. (2024). Beyond Single-Model Ablation: Evaluating Architectural Progress Across Research Generations. IEEE Transactions on Pattern Analysis and Machine Intelligence.

[10]. [10] Zhang, Y., Zhu, H., Isola, P., & Efros, A. A. (2024). Beyond Single-Model Ablation: Evaluating Architectural Progress Across Research Generations. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(5), 2345-2358. (Methodological foundation for evolutionary ablation studies)

[11]. Li, Z., Yang, Y., Liu, X., & Wang, Y. (2023). Systematic Ablation Studies in Computer Vision: A Survey. Computer Vision and Image Understanding, 228, 103645. (Comprehensive review of ablation methodologies)

[12]. Raghu, M., Unterthiner, T., Kornblith, S., & Zhang, C. (2022). Do Wide and Deep Networks Learn the Same Things? Uncovering How Neural Network Representations Vary with Width and Depth. Advances in Neural Information Processing Systems, 34, 25423-25436. (In-depth analysis of architectural components)

[13]. Singh, A. K., Ganapathysubramanian, B., Sarkar, S., & Singh, A. (2023). Deep Learning for Plant Stress Phenotyping: Trends and Future Perspectives. Trends in Plant Science, 28(1), 104-119. (Reviews evolution of agricultural DL)

[14]. Li, L., Zhang, Q., & Huang, D. (2024). A Historical Review of Machine Learning in Precision Agriculture: From Expert Systems to Deep Learning. Computers and Electronics in Agriculture,218,108742. (Historical perspective on agricultural AI)

[15]. Kamilaris, A., & Prenafeta-Boldú, F. X. (2023). A Review of the Use of Convolutional Neural Networks in Agriculture. The Journal of Agricultural Science, 161(1), 3-22. (Comprehensive CNN review in agriculture)

[16]. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ... & Guo, B. (2021). Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. Proceedings of the IEEE/CVF International Conference on Computer Vision, 10012-10022. (Foundational Swin Transformer paper)

International Research Journal on Advanced Engineering Hub
e ISSN: 2584-2137
Vol. 04 Issue: 01 January 2026
Page No: 72 - 79
www.irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0010

[17]. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. Advances in Neural Information Processing Systems, 33, 1-12. (Vision Transformer foundation)

[18]. Wang, W., Xie, E., Li, X., Fan, D. P., Song, K., Liang, D., ... & Shao, L. (2022). Pyramid Vision Transformer: A Versatile Backbone for Dense Prediction without Convolutions. Proceedings of the IEEE/CVF International Conference on Computer Vision, 568-578. (Transformer variants for vision)

[19]. Li, B., Wang, C., Huang, D., & Zhang, Q. (2023). Hyperspectral Imaging for Plant Disease Detection: A Comprehensive Review. ISPRS Journal of Photogrammetry and Remote Sensing, 195, 245-259. (HSI applications in plant pathology)

[20]. Feng, L., Zhang, Z., Ma, Y., Du, Q., Williams, P., Drewry, J., & Luck, B. (2024). Deep Learning-Based Fusion of RGB and Hyperspectral Imaging for Early Plant Stress Detection. Remote Sensing of Environment, 280, 113214. (Multimodal fusion approaches)

[21]. Lowe, A., Harrison, N., & French, A. P. (2023). Hyperspectral Image Analysis Techniques for the Detection and Classification of the Early Stages of Plant Disease. Plant Methods, 19(1), 1-15. (Early disease detection with HSI)

[22]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?": Explaining the Predictions of Any Classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135-1144. (Original LIME paper)

[23]. Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. Proceedings of the IEEE International Conference on Computer Vision, 618-626. (Grad-CAM for visual explanations)

[24]. Arun, A., & Choudhary, A. (2024). PlanXAI: A Benchmark Framework for Explainable AI in Plant Disease Diagnosis. IEEE Transactions on Agrifood Electronics, 2(3), 245-258. (XAI benchmarks in agriculture)

[25]. Chicco, D., & Jurman, G. (2023). Statistical Pitfalls in Machine Learning for Biomedical and Agricultural Applications. Nature Machine Intelligence, 5(2), 127-136. (Evaluation metrics and pitfalls)