INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY

ISSN: 2454-132X
Impact Factor: 6.078
(Volume 12, Issue 3 - V12I3-1170)
Available online at: https://www.ijariit.com

Edge-Optimized Pre-Trained Deep Learning Models for Real-Time Detection of Red Palm Weevil and Date Palm Diseases: A Review

Umar Faruk Ibrahim
ibrahimumarfaruk2000@gmail.com
Qassim University, Saudi Arabia

ABSTRACT

Date palm (Phoenix dactylifera L.) constitutes one of the most economically and culturally significant crops across arid and semi-arid regions, yet its productivity faces existential threats from the Red Palm Weevil (Rhynchophorus ferrugineus, RPW) and a spectrum of fungal and bacterial diseases. While deep learning has demonstrated remarkable classification accuracies exceeding 97% in controlled laboratory environments, the transition from academic prototypes to deployable, real-time agricultural solutions remains critically underdeveloped. This comprehensive review systematically examines recent studies, synthesizing the current landscape of deep learning applications for RPW detection and date palm disease classification. Our analysis reveals a persistent disconnect between architectural sophistication and practical deployability. Furthermore, the literature exhibits a pronounced fragmentation between pest detection and disease classification, with few studies addressing the integrated palm health ecosystem. This review identifies critical research dimensions where the current state-of-the-art falls short. By mapping these interconnected gaps across the evaluated literature, this review establishes a structured roadmap for developing lightweight, accurate, and interpretable AI systems that bridge the gap between theoretical accuracy and operational feasibility in precision agriculture.

Keywords: *Red Palm Weevil, Date Palm Disease, Edge Computing, Deep Learning, Transfer Learning, Model Compression, Explainable AI, Precision Agriculture, Real-Time Detection.*

I. INTRODUCTION

Date palm (Phoenix dactylifera L.) is one of the most commercially and culturally significant crops in arid and semi-arid regions, providing sustenance and livelihoods to millions globally [1]. However, date palm productivity faces severe threats from the Red Palm Weevil (Rhynchophorus ferrugineus, RPW) and various fungal diseases, which collectively cause substantial economic losses across major producing regions [2]. RPW has emerged as an overwhelming pest since the 1980s, spreading from Asia to the Middle East, Mediterranean basin, Europe, and Africa, causing tree mortality rates that severely compromise agricultural output.

Traditional detection methodologies have proven fundamentally inadequate against this threat. Manual visual inspection demands expertise, also remains labor-intensive across expansive plantations, and critically fails to identify infestations during the early intervention window when treatment remains viable [3]. Conventional approaches, including pheromone trapping and acoustic monitoring, demonstrate moderate effectiveness but suffer from delayed diagnosis, high operational costs, and limited scalability for large-scale agricultural operations [4].

The advent of deep learning and computer vision have demonstrated considerable shift in agricultural diagnostics. Convolutional Neural Networks (CNNs) and transfer learning architectures have achieved classification accuracies exceeding 97% in detecting RPW infestations and date palm diseases across various modalities, including RGB imagery, thermal imaging, and acoustic sensing [5]. Studies employing InceptionResNet-V2, custom CNN architectures, and hybrid deep learning models have established strong baseline performance for automated detection systems [6]. Simultaneously, edge computing has emerged as a transformative paradigm for agricultural Internet of Things (AIoT) applications, offering low-latency response, bandwidth optimization, and localized data processing that addresses the fundamental constraints of cloud-centric architectures [7].

Despite these advances, a critical translational gap persists between laboratory performance and field deployability. The current literature reveals a systematic prioritization of classification accuracy at the expense of computational efficiency, real-time feasibility, and practical integration [3]. High-performing models based on transformer architectures or deep CNNs frequently require substantial memory footprints and computational resources that exceed the capabilities of resource-constrained edge devices commonly deployed in agricultural settings [3]. While model compression techniques such as pruning, quantization, and knowledge distillation which have demonstrated success in reducing neural network size by over 90% without significant accuracy degradation in general plant disease recognition tasks, their systematic application to date palm health monitoring remains underexplored [9].

This review addresses these interconnected deficiencies by systematically examining recent studies that represent the current state-of-the-art in RPW detection and date palm disease classification. Our synthesis identifies seven critical dimensions where the literature falls short of practical deployment requirements and establishes a structured framework for advancing edge-optimized, pre-trained deep learning solutions. The rest of this paper is organized as follows: Section II establishes the theoretical foundations of transfer learning and architecture selection for palm health monitoring. Section III presents the evaluated literature. Section IV examines the state-of-the-art baseline performances and their inherent limitations. Section V analyzes model optimization and edge deployment challenges. Section VI discusses explainable AI and interpretability requirements. Section VII presents the identified research gaps, and Section VIII outlines future directions and concludes the paper.

II. THEORETICAL FOUNDATIONS AND ARCHITECTURE LANDSCAPE

A. Transfer Learning for Agricultural Diagnostics

Recent advancements in deep learning for date palm monitoring have transitioned from custom-built architectures to leveraging high-performance pre-trained models through transfer learning. Rather than training deep neural networks from scratch, researchers leverage pre-trained models which are already initialized on large-scale datasets such as ImageNet, ResNet, etc [10]. These models have learned rich, hierarchical feature representations from diverse natural images, capturing low-level features (edges, colors) in early layers and high-level semantic concepts in deeper layers. When fine-tuned on palm-specific datasets, the pre-trained weights provide a robust initialization that accelerates convergence, improves generalization, and mitigates overfitting on limited training data [11].

The effectiveness of transfer learning in date palm applications depends critically on the similarity between the source domain (natural images) and the target domain (palm fronds, thermal patterns). While ImageNet pre-training has demonstrated broad utility, domain-specific datasets such as PlantVillage containing thousands of annotated plant disease images may offer superior feature transfer for agricultural diagnostics [12]. The fine-tuning protocol itself introduces significant design decisions, that is, whether to freeze early layers and retrain only classification heads, employ gradual unfreezing, or fine-tune all parameters with differential learning rates. These choices profoundly impact both accuracy and computational cost, yet remain inconsistently addressed in the current literature [9].

B. Architectural Evolution: From CNNs to Vision Transformers

The architectural landscape for palm health monitoring has evolved through three distinct phases. Classical CNNs, including VGG16, ResNet, and Inception variants, established the foundational capability for automated feature extraction from palm imagery [14]. These architectures employ stacked convolutional layers with increasing receptive fields, enabling hierarchical feature learning from local textures to global patterns. ResNet's introduction of skip connections addressed the vanishing gradient problem in deep networks, facilitating the training of architectures with hundreds of layers that capture increasingly abstract representations.

The second phase introduced lightweight architectures specifically designed for mobile and embedded deployment. MobileNetV2 and V3 employ depthwise separable convolutions to dramatically reduce parameter counts and floating-point operations while maintaining competitive accuracy [16]. EfficientNet introduced compound scaling that is uniformly scaling network depth, width, and resolution to achieve superior efficiency-accuracy trade-offs, with EfficientNet-Lite variants optimized for edge TensorFlow deployment [17]. These architectures represent critical enablers for field-deployable palm health monitoring, yet their systematic evaluation in agricultural contexts remains limited.

The third phase witnesses the emergence of Vision Transformers (ViT), which fundamentally depart from the convolutional inductive bias. By treating image patches as sequential tokens and applying self-attention mechanisms, ViT captures long-range dependencies across the entire image, overcoming the local spatial sensitivity inherent in CNNs [18]. Recent studies have demonstrated that hybrid models combining ViT with convolutional mixers achieve exceptional classification accuracies for diseases such as brown spot and white scale, particularly in complex orchard backgrounds where global context proves discriminative [19]. However, the quadratic computational complexity of self-attention with respect to image resolution introduces significant challenges for edge deployment, necessitating careful architectural selection or intensive optimization [20].

III. METHODOLOGY AND PAPER SELECTION

The selection of the papers prioritized studies addressing the intersection of deep learning architectures, Red Palm Weevil detection, and date palm disease classification, with specific emphasis on edge deployability, real-time performance, and multi-modal sensing. The evaluated papers encompass diverse methodological approaches which include custom CNN architectures, transfer learning with pre-trained models, metaheuristic optimization, acoustic signal processing, thermal imaging, drone-based aerial surveillance, and comprehensive review studies. Table 1 provides a systematic overview of the evaluated literature, categorizing each study by its primary contribution, core architecture, dataset characteristics, and reported accuracy.

Table 1: Systematic Overview of Evaluated Literature

Authors & Year	Primary Focus	Core Architecture	Dataset Size	Reported Accuracy	Edge Consideration
Arasi et al. (2023)	RPW detection with metaheuristic optimization	Improved ShuffleNet + XGBoost	Limited	High	None
Nobel et al. (2024)	Palm leaf disease hybrid detection	ResNet50 + DenseNet201 + ECA-Net	Moderate	98.67%	None
Al-Shalout et al. (2022/2024)	Date palm disease ML classification	Custom CNN + SVM	Kaggle-based	~91%	None
Hessane et al. (2023)	White scale stage-wise classification	Pre-trained CNNs (ResNet50)	1,091 images	94.50%	None
Namoun et al. (2024)	Infected vs. healthy leaf dataset	ResNet50, EfficientNet, MobileNetV2	Standardized	~95% (EfficientNet)	None

Boulila et al. (2023)	Acoustic RPW detection	Deep learning on spectrograms	1,106 samples	100%	MobileNetV2, server only
Martin et al. (2025)	STM32-based acoustic detection	Classical ML (SVM, KNN)	47 recordings	Moderate	STM32 acquisition only
Martin et al. (2026)	Thermal imaging RPW detection	Custom CNN	Thermal dataset	98.5%	IoT mentioned, not deployed
Abdallah et al. (2025)	UAV-based palm health detection	Voting ensemble (MobileNet, etc.)	Drone imagery	99.16%	No on-drone inference
Ahmed & Ahmed (2023)	ResNet + InceptionResNet palm disease	ResNet, InceptionResNet	2,631 images	99.62%–100%	None
Pacal & Işık (2025)	Hybrid CNN-ViT for plant disease	HybridConv Mixer + ViT	Palm dataset	99%	Future work only
Safran et al. (2024)	Lightweight species classification	DPXception (custom lightweight CNN)	~2,000 images	92.9%	Mobile-oriented

IV. STATE-OF-THE-ART BASELINE PERFORMANCES

A. Custom Architectures and Metaheuristic Optimization

The earliest strand of recent literature focuses on custom-designed architectures optimized through metaheuristic algorithms. Arasi et al. [1] proposed the IRPWD-BSADL framework, integrating the Bird Swarm Algorithm with an improved ShuffleNet architecture for RPW detection. While the study achieved competitive accuracy through XGBoost classification, it exhibits several critical limitations that typify this research strand. The custom ShuffleNet was trained from scratch without leveraging pre-trained ImageNet weights, forgoing the substantial benefits of transfer learning including faster convergence and improved generalization [29]. Furthermore, the study provides no discussion of model compression techniques, quantization strategies, or deployment on actual edge hardware, despite the inherent efficiency advantages of ShuffleNet's channel shuffling architecture. The framework addresses RPW detection in isolation, neglecting the broader ecosystem of date palm diseases that concurrently threaten plantation health [1].

Similarly, Albraikan et al. [24] employed the Gorilla Troops Optimizer (GTO) to enhance RPW detection through Gabor filtering and Mask RCNN with MobileNetV2 backbone. While the reported 99.27% accuracy is impressive, the study claims edge suitability by virtue of MobileNetV2 selection without validating this claim through actual edge deployment, latency measurements, or power consumption analysis. The GTO optimization adds substantial training overhead with no demonstrable runtime benefit, and the framework focuses exclusively on adult beetle detection, while missing the critical early-stage larval infestation that determines intervention success. The absence of model compression such as quantization, pruning, and distillation and the lack of integration with disease detection capabilities further limit practical utility.

Nobel et al. [2] presented a hybrid approach combining ResNet50, DenseNet201, and ECA-Net attention mechanisms for palm leaf disease detection and therapy enhancement. While achieving 98.67% accuracy, the ensemble architecture is prohibitively large for resource-constrained edge devices, with combined parameter counts exceeding 100 million. The study provides no inference time benchmarks, latency measurements, or memory consumption analysis, which are metrics that are essential for assessing field deployability. Furthermore, the framework addresses leaf diseases (Dubas bug, honeydew) without integrating RPW detection, fragmenting the palm health monitoring problem and requiring separate systems for comprehensive plantation management.

B. Transfer Learning with Pre-Trained Models

A significant body of literature has adopted transfer learning from large-scale datasets, yet with inconsistent methodological rigor. Reddy et al. [22] employed VGG16 for date palm white scale disease detection, achieving 96% accuracy on a controlled dataset. However, the selection of VGG16 that is comprising 138million parameters demonstrates a prioritization of accuracy over deployability that contradicts edge computing requirements. The study provides no real-time inference analysis, no model compression, and no explicit discussion of ImageNet pre-training protocols or fine-tuning strategies. The Kaggle-sourced dataset of low-resolution images lacks field variability, and the framework exists in isolation without IoT integration or practical deployment architecture.

Namoun et al. [23] contributed a standardized dataset of infected versus healthy date palm leaves, benchmarking ResNet50, EfficientNet, and MobileNetV2. EfficientNet achieved the highest accuracy at approximately 95%, demonstrating the effectiveness of compound scaling for agricultural diagnostics. However, the study remains confined to desktop evaluation, with no edge hardware deployment, no quantization or pruning analysis, and no integration of RPW detection. The class imbalance acknowledged in the dataset limits model robustness across diverse field conditions.

Hessane et al. [21] focused on stage-wise classification of white scale disease using pre-trained CNNs, with ResNet50 achieving 94.50% accuracy. The two-stage classification approach: first detecting disease presence, then determining severity stage, represents a relevant formulation for precision agriculture. Nevertheless, the limited dataset of 1,091 images across four classes with significant imbalance (110 high infestation versus 320 healthy samples) raises substantial generalization concerns. No edge deployment, model compression, or multi-modal sensing integration is explored, and the framework addresses white scale in isolation without RPW co-detection.

Ahmed and Ahmed A. [28] achieved an exceptional accuracy of about 99.62%–100% using ResNet and InceptionResNet transfer learning on a 2,631-image palm leaf dataset. While demonstrating the power of deep residual architectures and transfer learning, the study provides no analysis of inference efficiency, model size, or deployment feasibility. The reliance on desktop GPU training without edge validation illustrates the accuracy-centric bias filling the literature.

C. Acoustic and Thermal Sensing Modalities

Recognizing the limitations of visual inspection for early-stage RPW detection, researchers have explored alternative sensing modalities. Boulila et al. [25] developed a deep learning classification system for acoustic signals captured from within palm trunks, converting larval chewing sounds into spectrogram images for CNN classification. MobileNetV2 was selected for mobile suitability, yet evaluation occurred exclusively on server hardware with no actual edge device testing.

The data collection methodology is destructive, potentially damaging to trees, and impractical for large-scale continuous monitoring. The small dataset (1,106 samples expanded via augmentation to 1,719) achieved suspiciously perfect 100% accuracy, suggesting potential overfitting rather than robust generalization.

V. MODEL OPTIMIZATION AND EDGE DEPLOYMENT CHALLENGES

A. Absence of Systematic Compression Techniques

Model compression represents the critical bridge between high-accuracy architectures and edge deployability, yet its application in palm health monitoring remains irregular and unsystematic. Quantization can decrease model size by 4× and accelerate inference through integer arithmetic, with typically minimal accuracy degradation when calibrated on representative validation data [30]. Also, structured pruning can reduce parameter counts by 50–90% while maintaining network topology cooperative to hardware acceleration [31]. Another technique is knowledge distillation, which enables radical size reduction with preserved accuracy [32].

Almost all the evaluated studies systematically do not apply these techniques. Arasi et al. [1] use ShuffleNet without quantization or pruning. Nobel et al. [2] deploy heavy ensembles without compression consideration. Boulila et al. [25] and Albraikan et al. [24] employ MobileNetV2 as-is, forgoing the substantial additional compression possible through post-training quantization.

B. The Accuracy-Efficiency Paradox

The models achieving highest classification accuracy are almost always the least suitable for edge deployment. Nobel et al.'s [2] hybrid ResNet50-DenseNet201-ECA-Net ensemble, while achieving 98.67% accuracy, combines two of the heaviest pre-trained architectures available, with parameter counts exceeding 50 million per backbone. Reddy et al.'s [22] VGG16 selection (138.4 million parameters) illustrates a deployment-agnostic mindset where accuracy is pursued without computational consequence. Even studies explicitly claiming edge suitability, such as Albraikan et al. [24] with MobileNetV2, fail to validate these claims through actual hardware testing, latency measurement, or power analysis.

VI. EXPLAINABLE AI AND AGRICULTURAL PRACTITIONER TRUST

A. The Black-Box Barrier to Adoption

The black-box nature of deep learning models presents a critical socio-technical barrier to adoption among agricultural practitioners. Farmers, agronomists, and plantation managers require transparent, interpretable decision-making processes to trust automated diagnostics and justify intervention costs [13]. A model achieving 99% accuracy provides little practical value if its classification rationale remains unclear to the end-user.

B. Limited XAI Integration

Among the evaluated studies, Martin et al. [26] is among who integrates explainability mechanisms. While Grad-CAM successfully highlights biologically relevant areas, the study does not evaluate explanation quality quantitatively or assess whether farmers can interpret heatmap outputs effectively.

The remaining studies treat models as pure prediction engines without interpretability consideration. This omission is particularly problematic for edge-deployable systems where farmers interact directly with mobile interfaces: without visual explanations of why a tree is flagged as diseased, user trust remains low and adoption rates suffer [27].

C. Beyond Visualization: Actionable Explanations

Future XAI integration must extend beyond passive heatmap generation to actionable, linguistically grounded justifications. Rather than merely highlighting a diseased leaf region, the system should articulate something like: “Brown spot disease detected with 94% confidence based on…”. Then Recommended intervention should also be like: “targeted fungicide application within 72 hours.” Such explanatory richness requires structured knowledge integration and natural language generation capabilities not yet explored in the palm health literature [8].

VII. SYNTHESIS OF IDENTIFIED RESEARCH GAPS

Through systematic analysis of the evaluated studies, interconnected research gaps emerge that collectively impede the transition from laboratory accuracy to field deployability. Table 2 summarizes these gaps, their prevalence across the literature, and their criticality for practical impact.

Table 2: Synthesis of Critical Research Gaps

Gap Dimension	Criticality	Current State	Required Advancement
Edge-Optimized Architectures	Critical	No systematic edge evaluation	Mandatory edge benchmarking on Raspberry Pi, Jetson Nano, Coral TPU
Model Compression	Critical	Compression entirely absent	Systematic INT8 quantization, structured pruning, knowledge distillation
Real-Time Performance	Critical	No edge-specific latency/FPS metrics	Standardized benchmarks: ms/image, FPS, power consumption
Multi-Modal Fusion	High	Single-modality isolation	Integrated visual + thermal + acoustic edge fusion
RPW-Disease Integration	High	Fragmented pest/disease focus	Unified palm health monitoring frameworks
Explainable AI	High	XAI virtually absent	Grad-CAM, LIME, SHAP with actionable explanations
End-to-End Pipelines	High	Models in isolation	Complete sensing → inference → alert → treatment loops
Cross-Regional Validation	Medium	Single-region datasets	Multi-regional validation, domain adaptation
Transfer Learning Protocols	Medium	Inconsistent fine-tuning	Standardized layer freezing, learning rate schedules
Dataset Augmentation	Medium	Basic augmentation only	GAN synthesis, federated learning, SMOTE

The most striking finding is the universality of certain gaps: almost all the studies fail to evaluate edge deployment, omit systematic model compression, and present models without complete system integration. This is not merely a collection of individual study limitations but a systemic failure of the research community to prioritize deployability alongside accuracy.

VIII. FUTURE DIRECTIONS AND CONCLUSION

Addressing the identified gaps requires a paradigm shift from accuracy-centric model development to deployability-first system design. This review proposes a conceptual framework integrating four critical innovations:

1. Hierarchical Edge-Cloud Architecture: Deploy lightweight models (MobileNetV3, EfficientNet-Lite) on mobile devices or drone-embedded edge devices for real-time inference, with uncertain cases escalated to cloud-based heavy models (ResNet, ViT) via recurrent connectivity. This hierarchical approach balances latency requirements with accuracy demands [15].
2. Progressive Multi-Modal Fusion: Implement early-stage detection through thermal and acoustic sensors, with visual confirmation triggered upon anomaly detection. Feature-level fusion using attention mechanisms can dynamically weight modalities based on environmental conditions and confidence scores, improving robustness across diverse field scenarios.
3. Continuous Compression-Accuracy Optimization: Treat model compression not as post-hoc optimization but as a co-design requirement. Neural Architecture Search (NAS) specifically targeting palm health tasks can discover efficient backbones that natively balance accuracy and efficiency, while differentiable pruning and quantization-aware training maintain accuracy through optimization.
4. Human-in-the-Loop Explainability: Integrate XAI not merely for post-hoc visualization but as an active feedback mechanism. Farmer-validated explanations can refine model attention mechanisms, while misclassification patterns identified through explanation analysis guide targeted data collection and model retraining.

This systematic review of recent studies reveals that while deep learning has achieved remarkable classification accuracy for Red Palm Weevil detection and date palm disease classification, the research community has systematically neglected the deployability dimensions essential for practical agricultural impact. The prevailing focus on server-grade accuracy metrics, evaluated on static datasets, masks critical vulnerabilities that render most published models operationally infeasible for field deployment.

The transition to edge-optimized, pre-trained deep learning models requires fundamental methodological reform. Researchers must treat model compression, edge benchmarking, multi-modal fusion, and system integration as core design requirements rather than peripheral considerations. The proposed conceptual framework provides a concrete foundation for achieving this transition.

By bridging the gap between theoretical accuracy and operational feasibility, the research community can deliver AI systems that genuinely serve agricultural practitioners through lightweight models that run on affordable edge hardware, interpretable decisions that build farmer trust, and integrated pipelines that transform detection into timely intervention. The economic and ecological stakes demand no less than this complete methodological commitment.

REFERENCES

[1] Arasi, M. A., Almuqren, L., Issaoui, I., Almalki, N. S., Mahmud, A., & Assiri, M. (2023). Enhancing red palm weevil detection using bird swarm algorithm with deep learning model. *IEEE Access, 12*, 1542–1551.

[2] Nobel, S. N., Imran, M. A., Bina, N. Z., Kabir, M. M., Safran, M., Alfarhood, S., & Mridha, M. F. (2024). Palm leaf health management: A hybrid approach for automated disease detection and therapy enhancement. *IEEE Access, 12*, 9097–9111.

[3] Mir, T. A., Bharany, S., Gupta, R., Ghoniem, R. M., Ur Rehman, A., & Taye, B. M. (2025). Hybrid Deep Learning Model for Date Palm Disease Classification: A Fusion of HybridConv Mixer and Vision Transformer. *Food Science & Nutrition, 13*(10), e71086.

[4] Yuliansyah, H., Hartanto, R., & Soesanti, I. (2024, December). Model compression using post-training dynamic range quantization for plant disease recognition. In *IET Conference Proceedings CP908* (Vol. 2024, No. 30, pp. 281–287). Stevenage, UK: The Institution of Engineering and Technology.

[5] Kondaveeti, H. K., & Simhadri, C. G. (2025). Evaluation of deep learning models using explainable AI with qualitative and quantitative analysis for rice leaf disease detection. *Scientific Reports, 15*(1), 31850.

[6] Venkatasaichandrakanth, P., & Iyapparaja, M. (2024). GNViT: An enhanced image-based groundnut pest classification using Vision Transformer (ViT) model. *PLOS ONE, 19*(4), e0301231.

[7] Amri, S., Al-shammaa, A., & Ahmad, N. (2024). MIV-PlantNet: A hybrid deep learning model for plant species classification. *Frontiers in Plant Science, 15*, 128456.

[8] Holzinger, A., Langs, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 9*(4), e1312.

[9] Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for image-based plant disease detection. *Frontiers in Plant Science, 7*, 1419.

[10] Koubaa, A., Aldawood, A., Saeed, B., Hadid, A., Ahmed, M., Saad, A., Alkhouja, H., Ammar, A., & Alkanhal, M. (2020). Smart palm: An IoT framework for red palm weevil early detection. *Agronomy, 10*(7), 987.

[11] Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. *Sensors, 18*(8), 2674.

[12] Mao, Y., Ashry, I., Wang, B., Al-Fehaid, Y., Al-Shawaf, A., Ng, T. K., Yu, C., & Ooi, B. S. (2021). Monitoring the red palm weevil infestation using machine learning and optical sensing. In *2021 Optical Fiber Communications Conference and Exhibition (OFC)* (pp. 1–3). IEEE.

[13] Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G. Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics, 4*(37), eaay7120.

[14] Alaa, H., Waleed, K., Samir, M., Tarek, M., Sobeah, H., & Salam, M. A. (2020). An intelligent approach for detecting palm trees diseases using image processing and machine learning. *Int. J. Adv. Comput. Sci. Appl*, *11*(7), 434-441.

[15] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal, 3*(5), 637–646.

[16] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 4510–4520).

[17] Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 6105–6114). PMLR.

[18] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. (2020). An image is worth 16×16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

[19] Pacal, I., & Işık, Ş. (2025). Hybrid CNN-ViT model for plant disease classification. *Neural Computing and Applications, 37*, 1–15.

[20] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., ... & Guo, B. (2021). Swin Transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 10012–10022).

[21] Hessane, A., El Youssefi, A., Farhaoui, Y., Aghoutane, B., & Amounas, F. (2023). A machine learning based framework for a stage-wise classification of date palm white scale disease. *Big Data Mining and Analytics, 6*(2), 263–272.

[22] Reddy, D. Y., Reddy, A. C. K., Prakhyat, C. H., Vamsi, D., & Kumar, A. K. (2024). Date palm white scale disease detection using convolutional neural networks (VGG16). In *Proceedings of the 2024 International Conference on Intelligent Systems for Cybersecurity (ISCS)* (pp. 1–6). IEEE.

[23] Namoun, A., Alkhodre, A. B., Abi Sen, A. A., Alsaawy, Y., & Almoamari, H. (2024). Dataset of infected date palm leaves for palm tree disease detection and classification. *Data in Brief, 57*, 110933.

[24] Albraikan, A. A., Khalid, M., Alruwais, N., Hasanin, T., Dutta, A. K., Mohsen, H., Rizwanullah, M., & Ibrahim, S. S. (2023). Automated red palm weevil detection using gorilla troops optimizer with deep learning model. *IEEE Access, 11*, 71616–71630.

[25] Boulila, W., Alzahem, A., Koubaa, A., Benjdira, B., & Ammar, A. (2023). Early detection of red palm weevil infestations using deep learning classification of acoustic signals. *Computers and Electronics in Agriculture, 212*, 108154.

[26] Martin, B., Saranya, S., & Chowdary, P. J. (2026). Smart IoT thermal imaging approach for early identification of Red Palm Weevil (RPW) infestation on palms. *Scientific Reports, 16*(1), 5392.

[27] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennettot, A., Tabik, S., Barbado, A., ... & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82–115.

[28] Ahmed, M., & Ahmed, A. (2023). Palm tree disease detection and classification using residual network and transfer learning of inception ResNet. *PLOS ONE, 18*, e0282250.

[29] Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big Data, 3*(1), 9.

[30] Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., ... & Adam, H. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2704–2713).

[31] Han, S., Mao, H., & Dally, W. (2015). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*.

[32] Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.