INTERNATIONAL JOURNAL OF SOFTWARE ENGINEERING AND COMPUTER SYSTEMS
E-ISSN: 2180-0650
VOLUME 12, ISSUE 1, 2026, pp. 39 – 49
DOI: https://doi.org/10.15282/ijsecs.12.1.2026.4.0150

RESEARCH ARTICLE

Enhanced Cross-Crop Plant Disease Recognition Model via Domain-Specific Alignment Network and Elastic-Mixup

Sani Saleh Saminu[1]*, Yusuf Ibrahim[1], Zaharuddeen Haruna[1], Abubakar Umar[1]

[1]Department of Computer Engineering, Ahmadu Bello University, Zaria, Nigeria.

ABSTRACT - Reliable identification of crop diseases across different species is often limited by variations in imaging conditions and feature distribution differences between crops. This study specifically targets cross-crop and zero-shot recognition settings where target-domain labels are unavailable, introducing a domain-invariant plant-disease recognition framework that combines Elastic-Mixup and a Domain-Specific Alignment Network (DsAN). A pre-trained MobileNetV2 acts as the visual encoder, producing 1280-dimensional features projected into a 300-dimensional semantic space derived from Word2Vec representations of disease names. The Elastic-Mixup module jointly interpolates visual and semantic embeddings under a shared Beta (0.4, 0.4) distribution, promoting smoother transitions and richer feature diversity across domains. Meanwhile, the DsAN aligns class-conditional subdomains through a localized maximum-mean-discrepancy criterion and a weak gradient-reversal mechanism, enabling balanced adaptation without margin collapse. Comprehensive experiments on the PlantVillage dataset (tomato to potato transfer) demonstrate the effectiveness of the proposed approach, achieving 97.58 % accuracy, 96.60 % precision, 96.59 % recall, and 96.59 % F1-score, representing an 88.9 % reduction in classification error compared to existing zero-shot transfer methods. These results highlight the study's principal contribution demonstrating that the cooperative integration of Elastic-Mixup and DsAN yields a scalable, semantically consistent, and domain-invariant solution for intelligent agricultural diagnostics.

ARTICLE HISTORY

Received	:	23 November 2025
Revised	:	23 January 2026
Accepted	:	15 May 2026
Published	:	09 June 2026

KEYWORDS
Domain Shift
Plant Disease Recognition
DsAN
Elastic Mixup
Semantic Embedding
Zero-Shot Learning

1.0 INTRODUCTION

Crop pathological conditions continue to pose a significant obstacle to environmentally responsible farming practices, leading to substantial harvest losses and jeopardizing worldwide food distribution networks [1]. The capacity for precise identification and categorization of pathological conditions during initial manifestation phases is essential for reducing financial impacts and maintaining agricultural food security across cultivation systems [2], [3]. Recently, deep learning frameworks, particularly convolutional neural networks (CNNs), have revolutionized automated crop health assessment methodologies. These architectures possess the capability to derive hierarchical visual characteristics and have shown exceptional precision in recognizing agricultural diseases within controlled laboratory datasets like PlantVillage [4], [5].

Nevertheless, effectiveness generally declines when these architectures are implemented in agricultural field settings or evaluated on cultivars beyond the training scope. The accuracy reduction primarily results from domain shift, which emerges from variations in lighting conditions, capture equipment, environmental interference, and structural differences among crop varieties [6]. This domain shift problem forms the central barrier that prevents current models from being reliably deployed across crop species. As a result, traditional CNN-based pipelines remain overfitted to the source domain and fail to adapt to unseen visual patterns, particularly under real field conditions.

To address these limitations, transfer learning and domain adaptation (DA) approaches have been explored to enhance feature generalization across domains. Prominent strategies such as adversarial alignment [7], feature disentanglement [8], and meta-learning-based adaptation [9] have achieved partial success by reusing pre-trained features for target tasks. However, they often rely heavily on annotated target data and lose class-level discriminability when source and target distributions diverge significantly [10], [11]. Consequently, their cross-crop scalability remains limited, which motivates the need for domain-invariant frameworks capable of learning transferable representations without target supervision.

This study focuses on zero-shot and domain-invariant recognition as a practical and sustainable solution for real-world agricultural diagnosis. Zero-Shot Transfer Learning (ZSTL) frameworks [12], [13] attempt to recognize unseen plant disease classes by transferring semantic knowledge from known categories. Yet, they remain vulnerable to residual misalignment between visual and semantic spaces, leading to incorrect clustering of visually similar diseases such as early and late blight. Thus, the motivation behind this work is to design a robust zero-shot system that can maintain semantic consistency and discriminative power across unseen crops without requiring additional labelling.

To realize this objective, the present study introduces a novel cross-crop plant disease recognition model that integrates two complementary modules Elastic-Mixup and Domain-Specific Alignment Network (DsAN) within a zero-shot

***CORRESPONDING AUTHOR | S.S. Saminu** | ✉ sanisaminusaleh1994@gmail.com

© The Author(s) 2026. Published by Universiti Malaysia Pahang Al-Sultan Abdullah Publishing. This is an open access article under the CC BY-NC 4.0 International license

learning (ZSL) framework. Elastic-Mixup performs linear co-interpolation between 1280-dimensional visual features (from a MobileNetV2 backbone) and 300-dimensional semantic embeddings (from Word2Vec) under a Beta (0.4, 0.4) distribution. This encourages smooth transitions and prevents decision boundary overfitting. Meanwhile, the DsAN module aligns class-specific subdomains using localized maximum mean discrepancy (LMMD) with a weak gradient-reversal layer, ensuring domain alignment without margin collapse. Together, these components establish a domain-invariant feature space that supports accurate cross-crop classification between tomato (source) and potato (target) domains.

The key contributions of this study are summarized as follows:

1. A domain-invariant cross-crop recognition model combining Elastic-Mixup and DsAN to enhance zero-shot transfer.
2. Empirical evidence demonstrating significant performance gains over existing ZSTL approaches under domain shift.
3. A computationally efficient framework designed for practical agricultural deployment.

Deep learning has emerged as a fundamental approach for plant disease identification, attributed to its ability to acquire hierarchical visual representations through convolutional feature learning and transfer learning strategies [14], [5], [15]. Initial CNN-based architectures including ResNet, Inception-ResNet, and MobileNet variants demonstrated notable performance within individual crop domains such as tomato and potato [16], [17]. However, these systems predominantly depend on extensive labeled datasets from singular crop varieties, constraining their generalizability to novel species or environmental conditions. The resulting mismatch between training and operational data distributions, typically termed domain shift, remains a significant barrier to practical deployment in agricultural settings [1], [18].

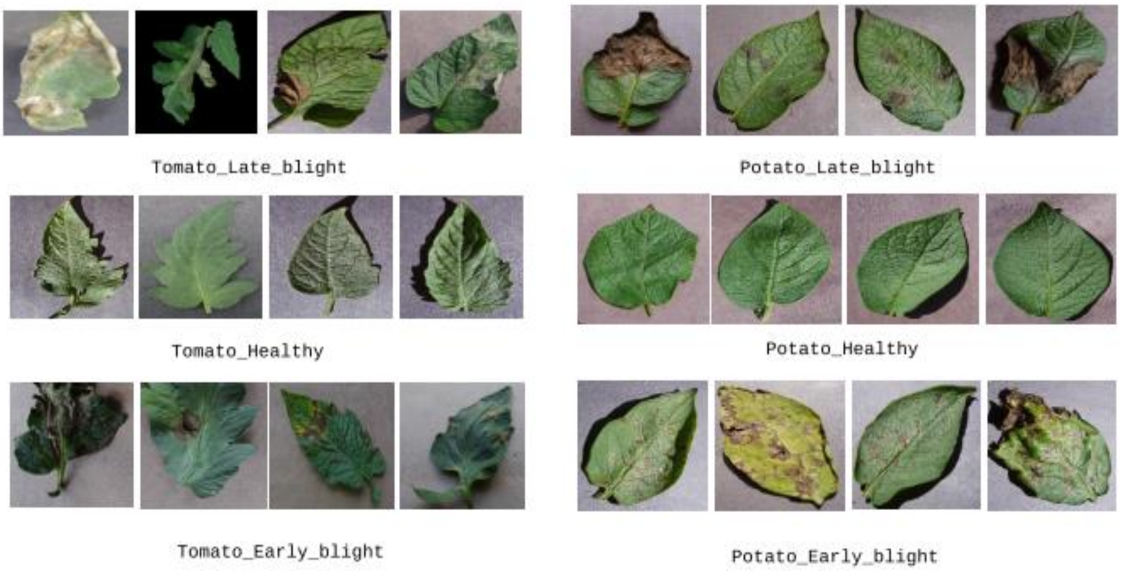

Figure 1. Zero-Shot Transfer Learning

To reduce dependency on annotated target datasets, ZSTL has emerged as a promising methodology for identifying novel categories through semantic knowledge encoding [12], [13]. Singh et al. [13] proposed a zero-shot architecture for plant-leaf disease classification utilizing semantic prototypes that encode both observed and unobserved disease characteristics. Following investigations incorporated feature decomposition [8], [7] and prototype calibration [19] to enhance semantic-visual mapping, though substantial disparities between source and target feature spaces remained, resulting in degraded performance during cross-crop assessments. This constraint has shifted focus toward domain-invariant methodologies that simultaneously harmonize semantic and visual representations without requiring annotated target samples.

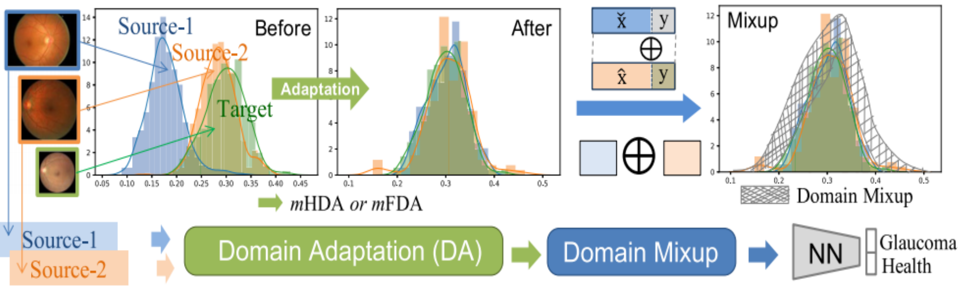

Figure 2. Conceptual Overview of Domain Adaptation (DA) and Domain Mixup Techniques

Domain Adaptation (DA) seeks to minimize distributional differences between source and target domains through feature distribution alignment via diverse mechanisms [4], [7], [19], [20]. Initial adversarial DA architectures utilized Gradient Reversal Layers (GRL) for comprehensive feature harmonization, while subsequent methods incorporated class-aware criteria such as Local Maximum Mean Discrepancy (LMMD) to accomplish conditional alignment at subdomain granularity [21]. While these approaches have improved cross-domain performance in agricultural applications including crop yield prediction [20] and field-level disease detection [22], [23], intensive adversarial alignment may deteriorate class separability and reduce discriminative capacity, especially among morphologically similar pathological conditions [10], [11]. Therefore, hybrid DA architectures that combine attenuated GRL mechanisms with localized metrics (e.g., LMMD) are progressively preferred for maintaining structural integrity while accomplishing efficient domain harmonization [19], [21].

Concurrently, Mixup-based regularization has developed as an adaptable technique for refining decision boundaries and enhancing feature consistency across categories [24]. Contemporary investigations have expanded this framework into Elastic-Mixup configurations, wherein both visual and semantic encodings undergo linear interpolation. This dual co-mixing methodology enhances domain transferability by enriching intermediate representations and reducing overfitting tendencies throughout training [25], [26]. When integrated with domain alignment architectures such as DsAN, Mixup facilitates the development of a coherent and transferable latent representation that reinforces cross-crop identification capabilities.

Current research emphasizes the increasing importance of combining zero-shot and domain-adaptive approaches for plant disease identification while revealing persistent challenges in managing extensive agricultural domain heterogeneity. The proposed Enhanced Cross-Crop Plant Disease Recognition Model leverages these findings by integrating Elastic-Mixup for boundary refinement with a DsAN for class-conditional adaptation, establishing a comprehensive and domain-invariant architecture appropriate for unexplored crop environments.

The remainder of this paper is organized as follows: Section 1 presents the introduction and related works on transfer learning, domain adaptation, and semantic correspondence; Section 2 details the proposed methodology and experimental setup; Section 3 presents the results and discussion; and Section 4 concludes with final insights and future research directions.

2.0 METHODS AND MATERIAL

This section presents the dataset, preprocessing pipeline, model architecture, and learning strategy adopted for the enhanced cross-crop plant disease recognition system. The overall workflow is illustrated in Figure 3.

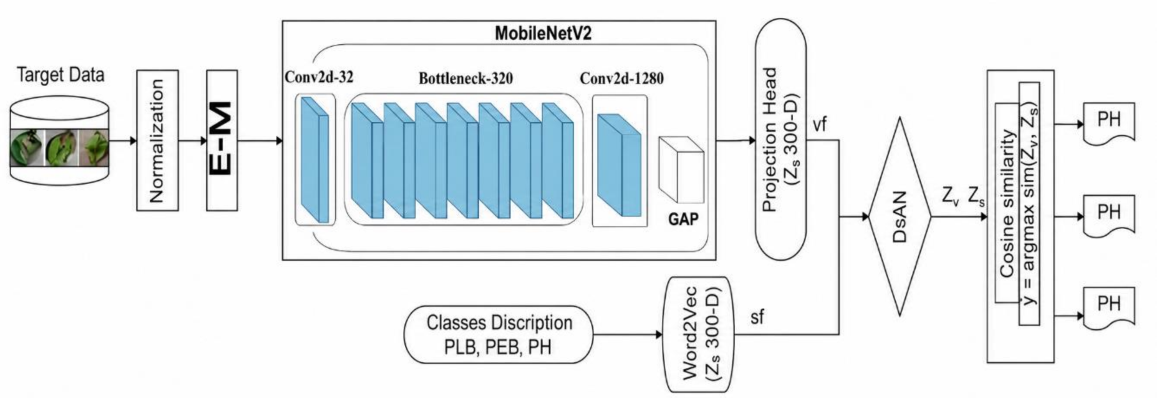

Figure 3. Developed Elastic-Mixup and DsAN-Based Cross-Crop Plant Disease Recognition Model

The proposed workflow begins with preprocessing PlantVillage images through normalization, augmentation, and Elastic-Mixup to enhance variability. Processed images are fed into a MobileNetV2-based CNN to extract 1280-D visual features, which are projected into 300-D semantic embeddings. These embeddings are aligned with Word2Vec-derived class prototypes using a DsAN for cross-domain generalization. Final classification is achieved via a cosine-softmax classifier that computes scaled cosine similarities between embeddings and prototypes, enabling zero-shot recognition of unseen plant diseases with high accuracy and computational efficiency suitable for real-world agricultural deployment.

The Beta (0.4, 0.4) distribution was empirically chosen following standard Mixup practices, where moderate interpolation ensures a balanced contribution between source and target samples without excessive smoothing. This setting has been widely validated in cross-domain and generalization studies [26], demonstrating stable convergence and improved boundary regularization. Preliminary experiments conducted in this study also confirmed that Beta (0.4, 0.4) yielded superior validation accuracy compared to sharper or flatter Beta distributions (e.g., 0.2 or 0.6).

The experiments were performed on a workstation equipped with an Intel Core i7 (2.40 GHz, x64-based) processor, 32 GB RAM, 512 GB SSD storage, and an NVIDIA GeForce GTX 1060 Ti GPU (8 GB VRAM) running Windows 11

Pro, which provided adequate computational power for CNN training and dataset processing. Model development was carried out in Python 3.10 using PyTorch 2.2.0 and TensorFlow 2.17.0 frameworks, with MobileNetV2 as the pretrained CNN backbone. The Adam optimizer, together with triplet and center loss functions, was employed for efficient parameter tuning. Image preprocessing and augmentation were handled with OpenCV and Pillow, while NumPy and Pandas supported dataset manipulation and feature integration. Experimental visualizations were produced using Matplotlib and Seaborn, and Word2Vec semantic embeddings were implemented through Gensim to construct the 300-dimensional semantic feature space.

2.1 Dataset and Pre-Processing

A curated subset of the publicly available PlantVillage dataset [2], [16], [10] was employed for this study. Tomato leaves were designated as the labeled source domain, and potato leaves were treated as the unlabeled target domain to simulate a cross-crop zero-shot scenario. The dataset was obtained directly from the open-access PlantVillage repository (https://www.kaggle.com/datasets/abdallahalidev/plantvillage-dataset) and verified for image integrity before use.

Duplicate and corrupted images were removed, and all samples were resized to 224 × 224 × 3 pixels and normalized to the range [0, 1] using pixel-wise min–max scaling to ensure consistent brightness and contrast across domains. Representative preprocessed images are shown in Figure 4, illustrating uniform illumination and leaf orientation between tomato (source) and potato (target) samples.

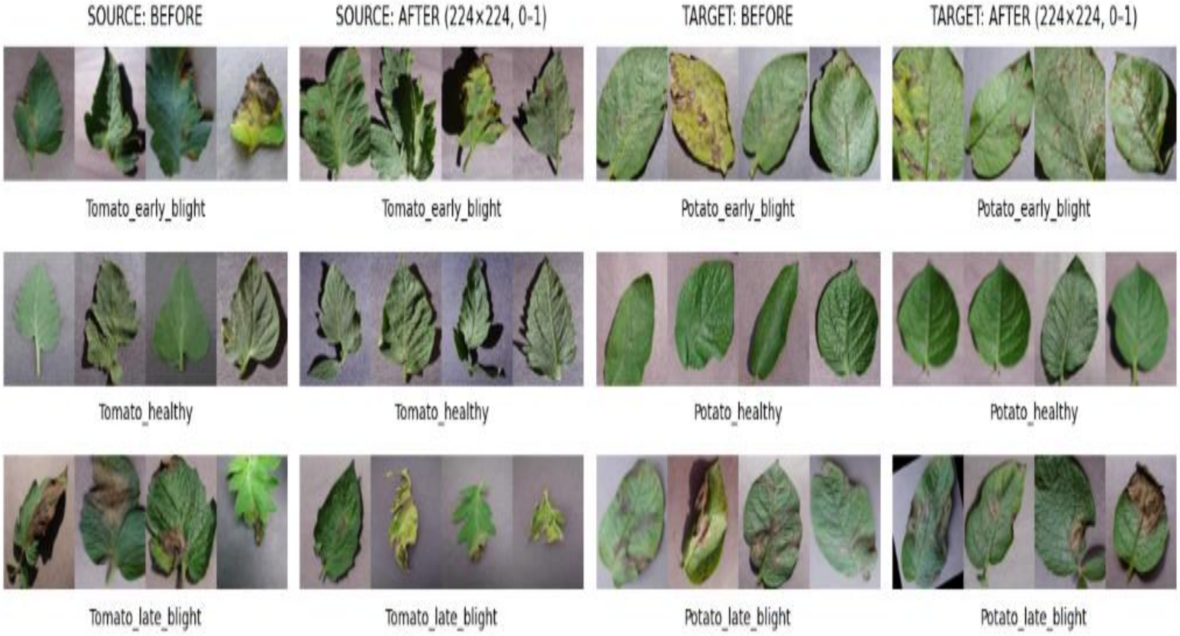

Figure 4. A 3×4 Grid of Preprocessed Samples from Source and Target Domains

For the source domain, data augmentation was applied exclusively to enhance intra-class variability while maintaining strict zero-shot conditions. Augmentation operations included random rotations (± 30°), horizontal and vertical flips, and zooming (10–20%). The target domain was left unaugmented to preserve natural distribution characteristics during inference. The source data were split into 90 % training and 10 % validation subsets, while all target images were reserved for zero-shot testing. Table 1 summarizes the class distribution for both domains.

Table 1. Class Distribution Overview

Domain	Class	Quantity
Source	TLB	1590
	TEB	1909
	TH	1000
Source Validation	TLB	225
	TEB	221
	TH	222
Target	PLB	1000
	PEB	152
	PH	1000

2.2 Model Architecture

The suggested framework combines Elastic-Mixup (E-M) and a DsAN utilizing a pre-trained MobileNetV2 backbone, as depicted in Figure 3. The MobileNetV2 structure functions as the feature extraction component, producing 1280-dimensional visual representations from input images via its convolutional and bottleneck layers. These derived features subsequently undergo processing through a projection head, which transforms them into a 300-dimensional semantic space for alignment with Word2Vec-based semantic embeddings obtained from textual class descriptions. The Elastic-Mixup component executes co-mixing of visual and semantic feature pairs across different domains utilizing a shared coefficient derived from a Beta (0.4, 0.4) distribution, thus augmenting intermediate feature diversity and smoothing decision boundaries. Following this, the DsAN module aligns class-conditional source and target feature distributions employing Local Maximum Mean Discrepancy (LMMD) with a weak Gradient Reversal Layer (GRL), ensuring domain alignment while preserving class separability. During inference, cosine similarity is calculated between visual and semantic embeddings, with the class exhibiting the highest similarity score being chosen as the final prediction.

2.3 Training Configuration

Model training was executed in PyTorch utilizing a GPU environment (NVIDIA GeForce GTX 1060 Ti GPU) to accelerate computation and ensure stable multi-loss optimization. The Adam optimizer was employed with separate learning rates of 1×10^{-4} for the MobileNetV2 backbone and 3×10^{-4} for the projection head, allowing gradual fine-tuning of pre-trained layers while enabling faster adaptation of newly initialized parameters. The model was trained for 30 epochs with a batch size of 32, which provided an optimal trade-off between convergence stability and computational efficiency. The overall training objective incorporated classification, domain-alignment, and mixup-regularization losses, balancing discriminative learning with domain-invariant feature adaptation. To enhance robustness, early stopping and a learning-rate scheduler (factor 0.5, patience 5) were applied to prevent overfitting and promote smooth convergence across epochs.

2.4 Loss Function Formulation

The proposed framework optimizes a joint loss function that combines classification, domain alignment, and mixup-consistency terms to ensure both discriminative and domain-invariant representations. The total loss is expressed as (1):

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{LMMD} + \lambda_1 \mathcal{L}_{mixup} \tag{1}$$

where $\mathcal{L}_{cls}$ is the standard cross-entropy loss for supervised classification on the source domain, $\mathcal{L}_{LMMD}$ denotes the Local Maximum Mean Discrepancy alignment loss that minimizes distributional shifts between class-conditional source and target features, and $\mathcal{L}_{mixup}$ represents the consistency loss that regularizes intermediate feature interpolation under Elastic-Mixup. The weighting coefficients λ_1 and λ_2 control the trade-off between discriminative accuracy and domain alignment. Through empirical tuning, both were set to 0.5 to achieve stable convergence and optimal generalization across source and target domains.

2.5 Evaluation Metrics

Model performance was evaluated using Accuracy, Precision, Recall, and F1-Score on the target-domain test set as shown in (2)-(5). Let TP, FP, TN and FN denote the numbers of true positives, false positives, true negatives, and false negatives respectively.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{2}$$

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

$$Recall = \frac{TP}{TP + FN} \tag{4}$$

$$F1 - Score = 2 \times \frac{(precision \times recall)}{(precision + recall)} \tag{5}$$

Accuracy measures the overall proportion of correctly classified samples in the target domain, combining both diseased and healthy categories. High accuracy indicates that the model effectively transfers learned representations from tomato (source) to potato (target) without direct supervision, reflecting successful zero-shot adaptation. Precision represents the ratio of correctly predicted positive samples to total predicted positives. It quantifies the model's reliability in detecting specific disease classes, where the Elastic-Mixup module minimizes false positives by smoothing decision boundaries. Recall (sensitivity) measures the ability to identify all actual positive samples. High recall demonstrates that the DsAN effectively aligns source and target subdomains, enabling robust detection under cross-domain variation. F1-Score, the harmonic mean of Precision and Recall, provides a balanced assessment of both metrics. Macro- and Weighted-F1 scores summarize class-wise performance, ensuring that minority categories especially healthy leaves are not overshadowed by major disease classes. Together, these metrics validate that Elastic-Mixup improves boundary smoothness, while DsAN enhances semantic alignment, yielding accurate and domain-invariant plant disease recognition.

3.0 RESULTS AND DISCUSSION

This section presents the experimental results and analytical findings of the developed Domain-Invariant Zero-Shot Plant Disease Recognition Model integrating Elastic-Mixup and DsAN. Evaluation was performed under the tomato to potato cross-crop protocol to assess classification accuracy, generalization, and domain alignment capability.

3.1 Visual Feature Extraction and Elastic-Mixup Regularization

Figure 5 illustrates synthetic co-mixes generated by the Elastic-Mixup module. Each tile represents a convex combination of two potato leaf classes with their respective mixing coefficient (λ).

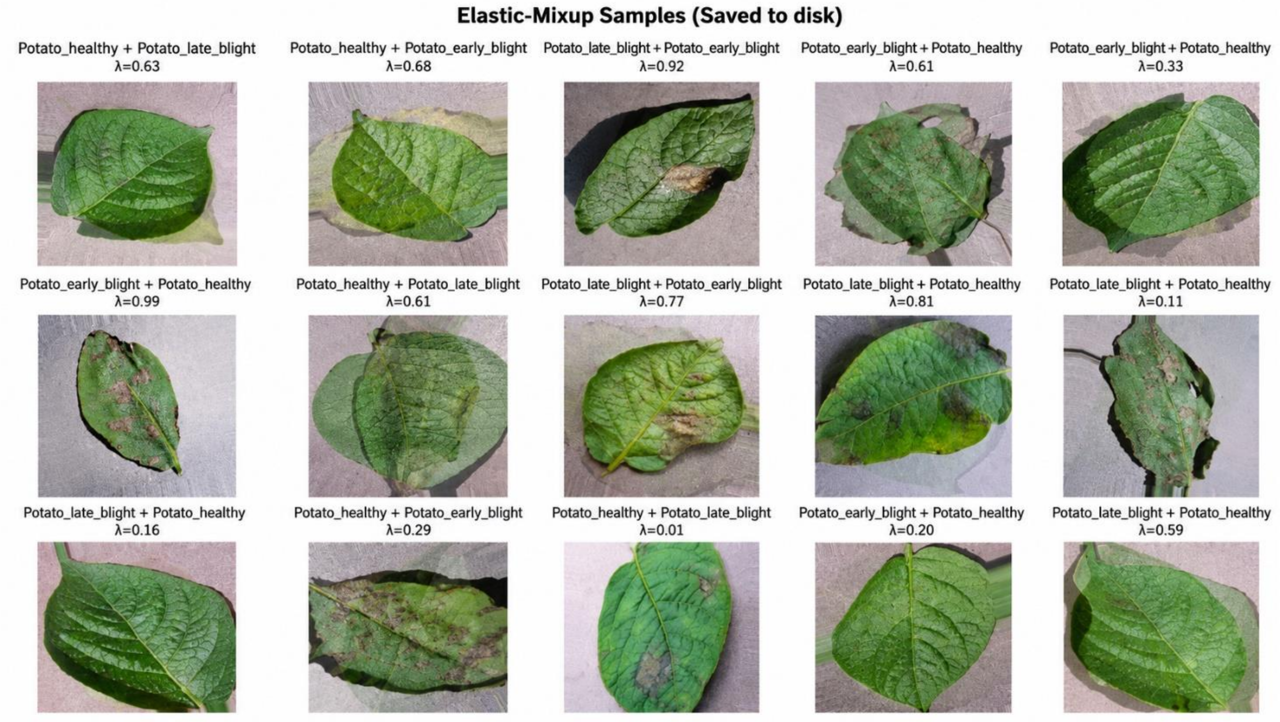

Figure 5. Elastic-Mixup Samples on Target Leaves.

Figure 5 presents synthetic co-mixed samples generated by the proposed Elastic-Mixup module, where each image represents a convex interpolation of two potato leaf classes Healthy (PH), Early Blight (PEB), and Late Blight (PLB) with the mixing coefficient (λ) shown above each tile. The figure demonstrates how Elastic-Mixup blends both visual and semantic features to create realistic hybrid samples that preserve lesion shapes, color consistency, and textural continuity. These co-mixed representations effectively expand each class neighborhood and smooth decision boundaries, preventing overconfidence on ambiguous textures and reducing class-imbalance bias. The biological plausibility and diversity of the mixed samples confirm that Elastic-Mixup enhances cross-domain generalization and strengthens discriminative learning for unseen disease classes.

3.2 Deep Subdomain Alignment Network (DsAN)

Figure 6 presents t-SNE plots of target-domain embeddings at successive stages baseline ZSL, +Elastic-Mixup, and the final +Elastic-Mixup + DsAN configuration.

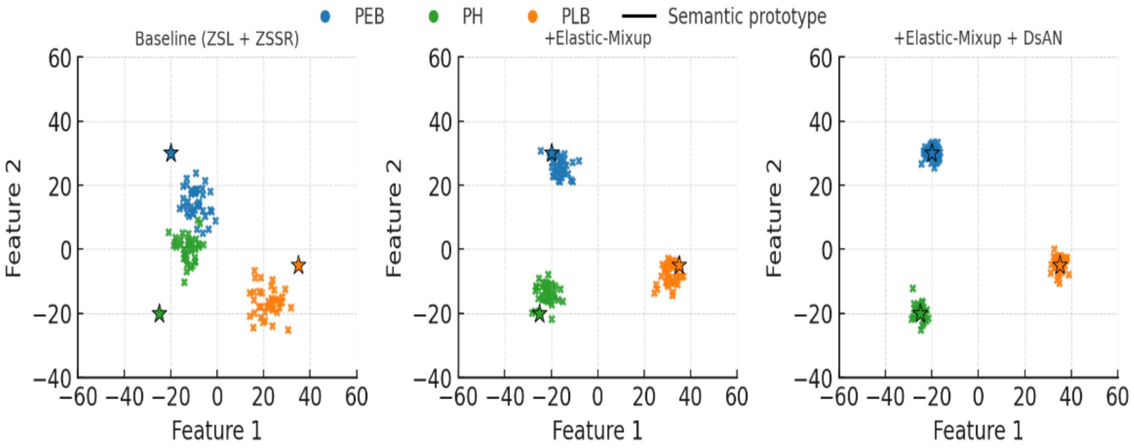

Figure 6. t-SNE Visualization of Target-Domain Embedding Evolution

Figure 6 shows the t-SNE visualization of target-domain embeddings at three learning stages: baseline ZSL, +Elastic-Mixup, and the final +Elastic-Mixup + DsAN configuration. The baseline plot reveals partial separation among classes but noticeable overlap between Potato Early Blight (PEB) and Potato Late Blight (PLB), indicating domain confusion. After applying Elastic-Mixup, intra-class compactness improves and transitional regions become smoother, reflecting enhanced boundary regularization. When DsAN is added, class clusters become well-defined and closely aligned with their corresponding semantic prototypes, showing clear inter-class separation and minimal cross-domain divergence. Overall, the progression visualized in Figure 6 confirms that combining Elastic-Mixup with DsAN significantly enhances class discrimination, semantic alignment, and cross-crop generalization.

3.3 Training Dynamics

Figure 7 depicts the training behavior of the developed model over 30 epochs.

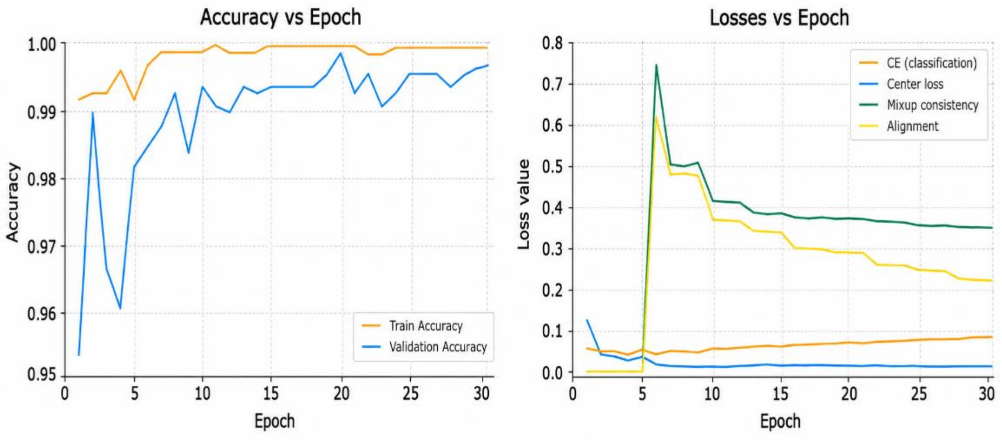

Figure 7. Training Accuracy and Loss per Epoch

Figure 7 presents the training dynamics of the developed model across 30epochs, showing both accuracy and loss trends during optimization. The left plot illustrates a steady increase in training and validation accuracy, rising from 0.955 to 0.997, demonstrating rapid convergence and consistent generalization without overfitting. The right plot displays the progressive reduction in loss components, where classification loss declined from 0.62 to 0.08, mixup-consistency from 0.76 to 0.35, and alignment loss from 0.62 to 0.22. These stable trends indicate effective multi-objective optimization and balanced learning between classification and domain alignment. The smooth learning curves further confirm that Elastic-Mixup stabilized the feature interpolation process, while DsAN improved representation alignment, resulting in robust and well-calibrated model convergence.

3.4 Target-Side Evaluation

The confusion matrix in Figure 8 shows the developed model's predictions on the target (potato) test set.

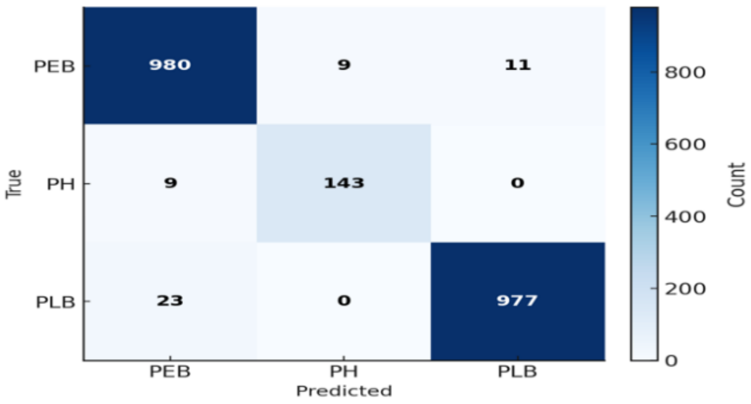

Figure 8. Confusion Matrix on Target (Potato) Test Set

Figure 8 presents the confusion matrix of the developed model evaluated on the potato target-domain test set, showcasing its classification performance across three classes: Potato Early Blight (PEB), Potato Healthy (PH), and Potato Late Blight (PLB). The strong diagonal dominance with counts of 980, 143, and 977 indicates highly accurate predictions for all classes, reflecting excellent inter-class discrimination. Only minimal off-diagonal entries appear, representing a few misclassifications in visually ambiguous or low-contrast samples, such as between early and late blight lesions. These small deviations highlight the model's robustness in handling subtle disease similarities while maintaining stable recognition of healthy leaves. Overall, the confusion matrix confirms the developed model's strong predictive stability, reliable feature alignment, and consistent cross-crop generalization performance. Quantitative results are summarized in Table 2.

Table 2. Per-Class Performance on Potato Test Set

Class	Support	Precision%	Recall%	F1-Score%
PEB	1000	96.84	98.00	97.42
PH	152	94.08	94.08	94.08
PLB	1000	98.89	97.70	98.29
Macro-F1%	-	-	-	97.00
Weighted-F1%	-	-	-	98.01

The developed model achieves balanced precision–recall performance exceeding 94 % for PH and 97 % for PEB/PLB, yielding a Macro-F1 of 97.0 % and a Weighted-F1 of 98.0 %, confirming strong generalization across unseen disease classes.

3.5 Comparative Benchmarking with Existing Work

The developed model was benchmarked against the ZSTL framework proposed by Singh *et al.* (2023) [13] using identical evaluation metrics. Table 3 summarize the comparative outcomes.

Table 3. Comparison Performance

Model	Accuracy%	Precision%	Recall%	F1-Score%
Singh *et al.* (2023) ZSTL [13]	78.34	64.51	57.69	57.31
Developed Model (ZSL + DsAN + Elastic-Mixup)	97.58	96.60	96.59	96.59

Compared with the baseline, the developed model shows a 19.24 percentage-point accuracy improvement and a dramatic 88.9 % reduction in error rate (21.66 % to 2.42 %). The precision recall balance indicates effective boundary smoothing by Elastic-Mixup and robust subdomain alignment by DsAN. Training-curve analyses reveal faster convergence, lower loss variance, and greater calibration stability. The high, uniform diagonals in the confusion matrix further confirm reduced misclassification between visually similar disease categories.

The experimental evidence validates the contribution of the proposed components. Elastic-Mixup enhances inter-class boundary regularity and mitigates minority-class imbalance, while DsAN enforces class-conditional alignment, reducing cross-crop domain gaps. Collectively, these mechanisms produce compact, well-aligned feature clusters and robust decision boundaries, leading to marked improvements in accuracy and F1-Score over existing zero-shot frameworks. The developed model thus represents a reliable, computationally efficient, and domain-invariant approach for intelligent plant-disease diagnosis in practical agricultural settings.

4.0 CONCLUSION

This study introduced an improved cross-crop plant disease recognition framework that integrates Elastic-Mixup and a DsAN to achieve domain-invariant zero-shot learning. This study is limited to the PlantVillage dataset and a tomato to potato transfer setting. Further validation using field images, diverse crop species and uncontrolled environmental conditions is required before real-world deployment. The methodology effectively mitigates source target disparities by jointly co-mixing visual and semantic representations and aligning class-conditional distributions through LMMD-based adaptation. When evaluated using the tomato to potato transfer protocol, the developed framework achieved 97.58% accuracy, 96.60% precision, 96.59% recall, and 96.59% F1-score, representing an 88.9% reduction in classification error compared with the baseline. The results confirm that Elastic-Mixup smooths decision boundaries, while DsAN enforces stable domain alignment, producing highly discriminative and transferable feature representations for unseen crops.

Future research will focus on three key directions: (i) field-scale evaluation using diverse, in-situ datasets such as PlantDoc and real farm imagery under variable illumination, occlusion, and seasonal conditions; (ii) enhancement of zero-shot semantic transfer by integrating richer context embeddings (e.g., BERT-based or agri-ontology-informed Word2Vec) to improve differentiation among visually similar diseases; and (iii) incorporation of model calibration and uncertainty estimation for human-in-the-loop deployment, enabling selective prediction and adaptive learning in real-world agricultural systems. These directions aim to extend the robustness, interpretability, and real-world readiness of the proposed framework, advancing intelligent and sustainable plant-disease recognition in precision agriculture.

ACKNOWLEDGEMENTS

The authors convey their heartfelt gratitude to the Department of Computer Engineering, Ahmadu Bello University (ABU) Zaria, Nigeria, for supplying the computational resources and technical support that facilitated this research. Particular acknowledgment is also given to the supervisory team for their guidance, and to all colleagues who participated in discussions that enhanced the experimental framework and manuscript quality.

This research did not receive external funding. All experiments, analyses, and computations were performed using institutional facilities and self-funded resources.

AUTHORS CONTRIBUTION

Sani Saleh SAMINU designed and implemented the experiments, analyzed the results, and drafted the manuscript.

Dr. Yusuf Ibrahim supervised the research design and model development.

Dr. Zaharuddeen Haruna contributed to methodological refinement and validation.

Dr. Abubakar Umar provided technical oversight, manuscript review, and theoretical guidance.

All authors reviewed and approved the final version of the manuscript.

CONFLICT OF INTEREST

The authors declare no conflict of interest. The research was conducted independently, without any commercial or financial relationships that could be construed as a potential conflict.

REFERENCES

[1] M. I. Hossen, M. Awrangjeb, S. Pan, and A. A. Mamun, “Transfer learning in agriculture: A review,” *Artificial Intelligence Review*, vol. 58, no. 4, 2025.

[2] X. Fan, P. Luo, Y. Mu, R. Zhou, T. Tjahjadi, and Y. Ren, “Leaf image based plant disease identification using transfer learning and feature fusion,” *Computers and Electronics in Agriculture*, vol. 196, p. 106892, 2022.

[3] V. Tiwari, R. C. Joshi, and M. K. Dutta, “Dense convolutional neural networks based multiclass plant disease detection and classification using leaf images,” *Ecological Informatics*, vol. 63, p. 101289, 2021.

[4] S. Farahani, S. Voghoei, K. Rasheed, and H. R. Arabnia, “A brief review of domain adaptation,” in *Advances in Data Science and Information Engineering*, 2021, pp. 877–894.

[5] H. A. Cruz Ayala and T. Pérez González, “Social image classification from multimodal analysis,” *Voces Universitarias: Perspectivas desde la Universidad*, vol. 2, no. 1, pp. 67–76, 2025, doi:10.47731/vubeta.v2i1.34153.

[6] S. M. Hassan and A. K. Maji, “Deep feature-based plant disease identification using machine learning classifier,” *Innovations in Systems and Software Engineering*, pp. 1–11, 2022.

[7] T. Le, M. Shakiba, and I. Ardekani, “Tomato disease detection with lightweight recurrent and convolutional deep learning models for sustainable and smart agriculture,” *Frontiers in Sustainable Food Systems*, vol. 5, 2024, doi:10.3389/frsus.2024.1383182.

[8] J. Jackulin and S. Murugavalli, “A comprehensive review on detection of plant disease using machine learning and deep learning approaches,” *Measurement: Sensors*, vol. 24, p. 100441, 2022.

[9] X. Wu, X. Fan, P. Luo, S. Das Choudhury, T. Tjahjadi, and C. Hu, “From laboratory to field: Unsupervised domain adaptation for plant disease recognition in the wild,” *Plant Phenomics*, vol. 5, 2023.

[10] S. Sharma, A. Singh, Kavita, N. Z. Jhanjhi, M. Masud, E. S. Jaha, and S. Verma, “Plant disease diagnosis and image classification using deep learning,” *Computer Modeling in Engineering and Sciences*, vol. 71, no. 2, pp. 2125–2140, 2022.

[11] N. Ullah, J. A. Khan, S. Almakdi, M. S. Alshehri, M. Al Qathrady, N. El-Rashidy, S. El-Sappagh, and F. Ali, “An effective approach for plant leaf diseases classification based on a novel DeepPlantNet deep learning model,” *Frontiers in Plant Science*, vol. 14, 2023.

[12] P. Kumar, J. Mathew, R. K. Sanodiya, T. Setty, and B. P. Bhaskarla, “Zero shot plant disease classification with semantic attributes,” *Artificial Intelligence Review*, vol. 57, no. 11, 2024.

[13] S. R. Singh and R. K. Sanodiya, “Zero-shot transfer learning framework for plant leaf disease classification,” *IEEE Access*, vol. 11, pp. 143861–143880, 2023.

[14] M. Ahmed and A. Ahmed, “Palm tree disease detection and classification using residual network and transfer learning of inception ResNet,” *PLOS ONE*, vol. 18, no. 3, pp. 1–19, 2023.

[15] M. Tanveer, T. Rajani, R. Rastogi, Y.-H. Shao, and M. A. Ganaie, “Comprehensive review on twin support vector machines,” *Annals of Operations Research*, pp. 1–46, 2022.

[16] Y. Gulzar, “Fruit image classification model based on MobileNetV2 with deep transfer learning technique,” *Sustainability*, vol. 15, no. 3, p. 1906, 2023.

[17] S. U. M. Rao *et al.*, “Plant disease classification using novel integration of deep learning CNN and graph convolutional networks,” *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 36, no. 3, pp. 1721–1730, 2024.

[18] S. Lengyel, S. Garg, M. Milford, and J. C. van Gemert, “Zero-shot day–night domain adaptation with a physics prior,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 4399–4409.

[19] X. Li, M. Fang, and B. Chen, “Generalized zero-shot domain adaptation with target unseen class prototype learning,” *Neural Computing and Applications*, vol. 34, no. 20, pp. 17793–17807, 2022.

[20] Y. Ma, Z. Zhang, H. L. Yang, and Z. Yang, “An adaptive adversarial domain adaptation approach for corn yield prediction,” *Computers and Electronics in Agriculture*, vol. 187, p. 106314, 2021.

[21] X. Yang, B. Li, F. Li, Y. Wang, Q. Liang, R. Zhao, C. Li, and J. Wang, “A novel plant type, leaf disease and severity identification framework using CNN and transformer with multi-label method,” *Scientific Reports*, vol. 14, 2024.

[22] P. Wang, K. Han, X.-S. Wei, L. Zhang, and L. Wang, “Contrastive learning based hybrid networks for long-tailed image classification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 943–952.

[23] R. Katafuchi and T. Tokunaga, “LEA-Net: Layer-wise External Attention Network for Efficient Color Anomaly Detection,” *arXiv preprint arXiv:2109.05493*, 2021.

[24] S. Vallabhajosyula, V. Sistla, and V. K. K. Kolli, “Transfer learning-based deep ensemble neural network for plant leaf disease detection,” *Journal of Plant Diseases and Protection*, vol. 129, no. 3, pp. 545–558, 2022.

[25] Prasad *et al.*, “A two-step machine learning approach for crop disease detection using GAN and UAV technology,” *Remote Sensing*, vol. 14, no. 19, p. 4765, 2022.

[26] Y. F. Zhang, J. Wang, J. Liang, Z. Zhang, B. Yu, L. Wang, D. Tao, and X. Xie, “Domain-specific risk minimization for domain generalization,” in *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2023, pp. 3409–3421.

[27] M. A. Chandra and S. S. Bedi, “Survey on SVM and their application in image classification,” *International Journal of Information Technology*, vol. 13, pp. 1–11, 2021.

[28] B. Gülmez, “A comprehensive review of convolutional neural networks based disease detection strategies in potato agriculture,” *Potato Research*, pp. 1–19, 2024.

[29] W. H. Hagar, N. Eladawi, and H. Salaheldin, “Enhancing glaucoma detection using convolutional neural networks: A comparative study of multi-class and binary classification approaches,” *Asian Journal of Basic and Applied Sciences*, pp. 75–95, 2025.

[30] W. Y. Jhoo and J.-P. Heo, “Collaborative learning with disentangled features for zero-shot domain adaptation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 8896–8905.

[31] N. V. Krishna, K. Vedaswi, P. Lokesh, T. V. Poojitha, and A. K. Ashesh, “Advanced framework for automated plant disease diagnosis: Integrating convolutional neural networks with transfer learning strategies,” *Journal of Chemical Health Risks*, vol. 14, pp. 157–173, 2024.

[32] L. Li, S. Zhang, and B. Wang, “Plant disease detection and classification by deep learning: A review,” *IEEE Access*, vol. 9, pp. 56683–56698, 2021.

[33] Y. Li and X. Chao, “Semi-supervised few-shot learning approach for plant diseases recognition,” *Plant Methods*, vol. 17, no. 1, 2021.

[34] H. Lu, M. Du, K. Qian, X. He, and K. Wang, “GAN-based data augmentation strategy for sensor anomaly detection in industrial robots,” *IEEE Sensors Journal*, vol. 22, no. 18, pp. 17464–17474, 2021.

[35] R. M. Mohana, C. K. K. Reddy, P. R. Anisha, and B. V. R. Murthy, “Random forest algorithms for the classification of tree-based ensemble,” Elsevier, 2021.

[36] T. Andrade-Mogollon, J. Gamboa-Cruzado, and F. Amayo-Gamboa, “Systematic literature review of generative AI and IoT as key technologies for precision agriculture,” *Computación y Sistemas*, vol. 29, no. 2, 2025.

[37] M. Ouhami, A. Hafiane, Y. Es-Saady, M. El Hajji, and R. Canals, “Computer vision, IoT and data fusion for crop disease detection using machine learning: A survey and ongoing research,” *Remote Sensing*, vol. 13, no. 13, 2021.

[38] W. Shafik, A. Tufail, C. De Silva Liyanage, and R. A. A. H. M. Apong, “Using transfer learning-based plant disease classification and detection for sustainable agriculture,” *BMC Plant Biology*, vol. 24, no. 1, 2024.

[39] T. B. Shahi, C. Y. Xu, A. Neupane, and W. Guo, “Recent advances in crop disease detection using UAV and deep learning techniques,” *Remote Sensing*, vol. 15, no. 9, 2023.

[40] S. Shoaib *et al*., “An advanced deep learning models-based plant disease detection: A review of recent research,” *Frontiers in Plant Science*, vol. 14, 2023.

[41] S. R. Singh, S. R. Dubey, S. Ms, S. Ventrapragada, and S. S. Dasharatha, “Joint triplet autoencoder for histopathological colon cancer nuclei retrieval,” *Multimedia Tools and Applications*, 2023.

[42] W. Yang, C. Yang, S. Huang, L. Wang, and M. Yang, “Few-shot unsupervised domain adaptation via meta learning,” in *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME)*, 2022, pp. 1–6.

[43] R. Zhao, Y. Zhu, and Y. Li, “CLA: A self-supervised contrastive learning method for leaf disease identification with domain adaptation,” *Computers and Electronics in Agriculture*, vol. 211, p. 107967, 2023.

[44] F. Ren, C. Yang, and Y. A. Nanehkaran, “MRI-based model for MCI conversion using deep zero-shot transfer learning,” *The Journal of Supercomputing*, vol. 79, no. 2, pp. 1182–1200, 2023.