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*, Yusuf Ibrahim, Zaharuddeen Haruna, Abubakar Umar

Department of Computer Engineering, Ahmadu Bello University, Samaru Campus, Zaria, Kaduna State, 810107, Nigeria

Abstract – The 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, and introduces a domain-invariant plant disease recognition framework that combines elastic mixup and a domain-specific alignment network (DsAN). A pretrained 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 the domains. Meanwhile, the DsAN aligns class-conditional subdomains through a localised 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) demonstrated 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 with existing zero-shot transfer methods. These results highlight the principal contribution of this study, 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. Introduction

Crop pathological conditions continue to pose a significant obstacle to environmentally responsible farming practices, causing substantial harvest losses and jeopardising global food distribution networks [1]. Precisely identifying and categorising pathological conditions early is essential to reduce financial impacts and maintain agricultural food security across cultivation systems [2]-[3]. Recently, deep learning frameworks, particularly convolutional neural networks (CNNs), have revolutionised automated crop health assessment methodologies. These architectures can extract hierarchical visual features and have shown exceptional precision in recognising agricultural diseases in controlled laboratory datasets, such as PlantVillage [4]-[5]. Nevertheless, their effectiveness generally declines when implemented in agricultural field settings or evaluated on cultivars beyond the training scope. This accuracy reduction primarily results from domain shift, which arises from variations in lighting conditions, capture equipment, environmental interference, and structural differences among crop varieties [6]. This domain shift problem is the central barrier preventing the reliable deployment of current models across crop species. Consequently, 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, researchers have explored transfer learning and domain adaptation (DA) approaches to improve feature generalisation across domains. Prominent strategies, such as adversarial alignment [7], feature disentanglement [8], and meta-learning-based adaptation [9], have achieved partial success by reusing pretrained features for target tasks. However, they often rely heavily on annotated target data and lose class-level discriminability when the source and target distributions diverge significantly [10]-[11]. Consequently, their cross-crop scalability remains limited, motivating the need for domain-invariant frameworks that learn transferable representations without supervision.

This study focuses on zero-shot and domain-invariant recognition as practical, sustainable solutions for real-world agricultural diagnosis. Zero-Shot Transfer Learning (ZSTL) frameworks [12]-[13] attempt to recognise unseen plant disease classes by transferring semantic knowledge from known categories, as shown in Figure 1. However, 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, this study aims to design a robust zero-shot system that maintains semantic consistency and discriminative power across unseen crops without additional labelling. To achieve this objective, this 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 learning (ZSL) framework. Elastic-Mixup performs linear co-interpolation between 1280-dimensional visual features (from the 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 a localised 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 summarised as follows.

*CORRESPONDING AUTHOR | S.S. Saminu | ✉ sanisaminusaleh1994@gmail.com
© 2026 The Author(s). Published by Universiti Malaysia Pahang Al-Sultan Abdullah Press. This is an open access article under the CC BY-NC 4.0 license

i) A domain-invariant cross-crop recognition model combining Elastic-Mixup and DsAN to enhance zero-shot transfer.
ii) Empirical evidence demonstrating significant performance gains over existing ZSTL approaches in the case of domain shifts.
iii) We designed a computationally efficient framework for practical agricultural deployment.

Deep learning has emerged as a fundamental approach for plant disease identification because it can learn hierarchical visual representations through convolutional feature learning and transfer learning strategies [14]-[15]. Initial CNN-based architectures, including ResNet, Inception-ResNet, and MobileNet variants, demonstrated notable performance within individual crop domains, such as tomatoes and potatoes [16]-[17]. However, these systems largely depend on extensive labelled datasets from single crop varieties, limiting their generalisability to novel species or environmental conditions. This mismatch between training and operational data distributions, typically termed a domain shift, remains a major barrier to practical deployment in agricultural settings [1], [18].



Figure 1. Zero-shot transfer learning

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

Domain Adaptation (DA) seeks to minimise distributional differences between the source and target domains through feature distribution alignment using diverse mechanisms [4], [7], [19]-[20]. Early adversarial DA architectures used Gradient Reversal Layers (GRL) for broad feature harmonisation, while later methods incorporated class-aware criteria, such as Local Maximum Mean Discrepancy, to achieve 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 localised metrics (for example, LMMD) are progressively preferred for maintaining structural integrity while accomplishing efficient domain harmonisation [19], [21]. Concurrently, Mixup-based regularisation has been developed as an adaptable technique for refining decision boundaries and enhancing feature consistency across categories [24]. Recent work has extended this framework to elastic mixup configurations, in which both visual and semantic encodings undergo linear interpolation. This dual co-mixing methodology enhances domain transferability by enriching intermediate representations and reducing overfitting tendencies during training [25]-[26]. When integrated with domain alignment architectures such as DsAN, Mixup helps

develop a coherent, transferable latent representation that strengthens cross-crop identification capabilities, as shown in Figure 2.

Current research emphasises 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 DsAN for class-conditional adaptation, establishing a comprehensive, domain-invariant architecture suitable for unexplored crop environments. The remainder of this paper is organised 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.

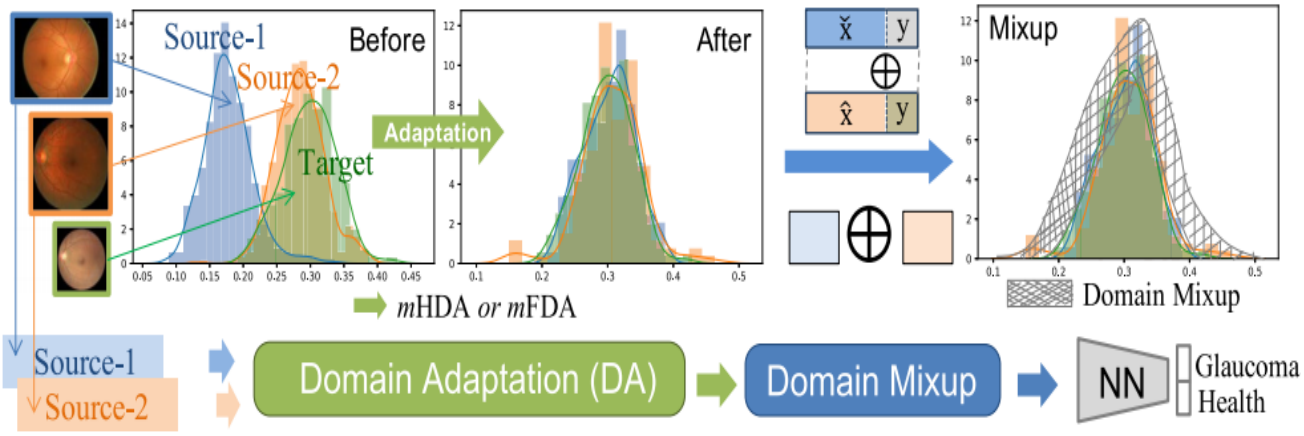


Figure 2. Conceptual overview of domain adaptation and domain mixup techniques

2. Methods and Materials

This section presents the dataset, preprocessing pipeline, model architecture, and learning strategy adopted for the enhanced cross-crop plant disease recognition system. Figure 3 shows the overall workflow. The proposed workflow begins by preprocessing PlantVillage images through normalisation, augmentation, and Elastic-Mixup to increase variability. The processed images were fed into a MobileNetV2-based CNN to extract 1280-D visual features, which were projected into 300-D semantic embeddings. We aligned these embeddings with Word2Vec-derived class prototypes using DsAN for cross-domain generalisation. The final classification uses 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. We empirically chose the Beta (0.4, 0.4) distribution 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 generalisation studies [26], demonstrating stable convergence and improved boundary regularisation.

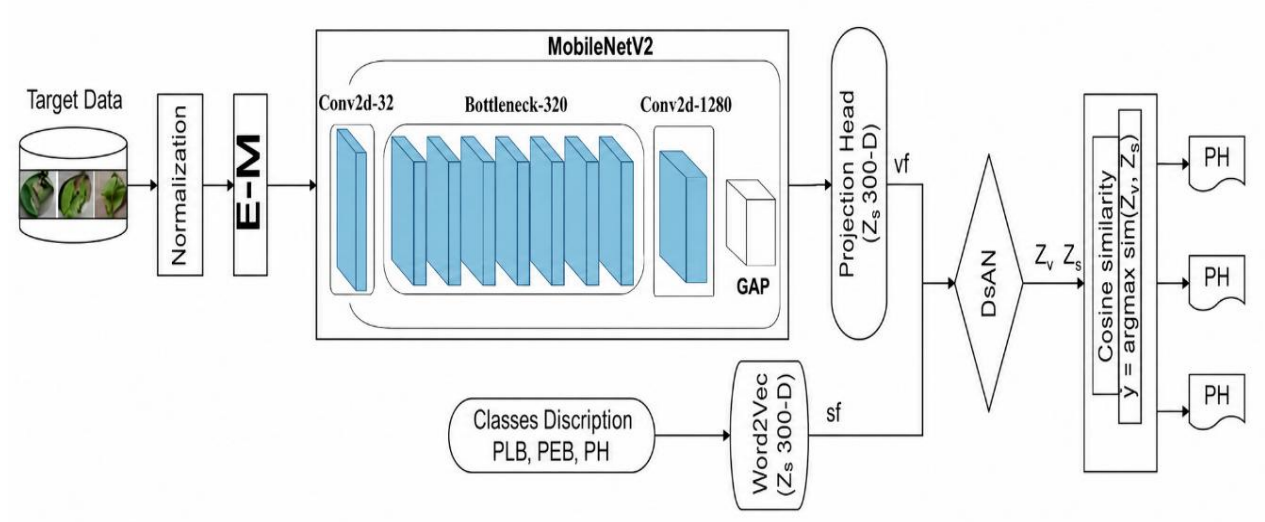


Figure 3. Developed elastic-mixup and DsAN -based cross-crop plant disease recognition model

Preliminary experiments conducted in this study also confirmed that Beta (0.4, 0.4) yielded superior validation accuracy compared with sharper or flatter Beta distributions (for example, 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 the CNN training and dataset processing. The model was developed in Python 3.10 using PyTorch 2.2.0 and TensorFlow 2.17.0, with MobileNetV2 as the pretrained CNN backbone. The Adam optimiser, together with triplet and centre loss functions, was employed for efficient parameter tuning. Image preprocessing and augmentation were performed using OpenCV and Pillow, whereas NumPy and Pandas supported dataset manipulation and feature integration. Matplotlib and Seaborn produced experimental visualisations, and Gensim implemented Word2Vec semantic embeddings to construct the 300-dimensional semantic feature space.

2.1 Dataset and Pre-Processing

This study used a curated subset of the publicly available PlantVillage dataset [2], [10], [16]. Tomato leaves were designated as the labelled source domain, and potato leaves were treated as the unlabelled target domain to simulate a cross-crop zero-shot scenario. We obtained the dataset directly from the open-access PlantVillage repository (https://www.kaggle.com/datasets/abdallahalidev/plantvillage-dataset) and verified image integrity before use. We removed duplicate and corrupted images, resized all samples to 224 × 224 × 3 pixels, and normalised them to the range [0, 1] using pixel-wise min–max scaling to ensure consistent brightness and contrast across domains. Figure 4 shows representative pre-processed images, illustrating uniform illumination and leaf orientation between the tomato (source) and potato (target) samples.

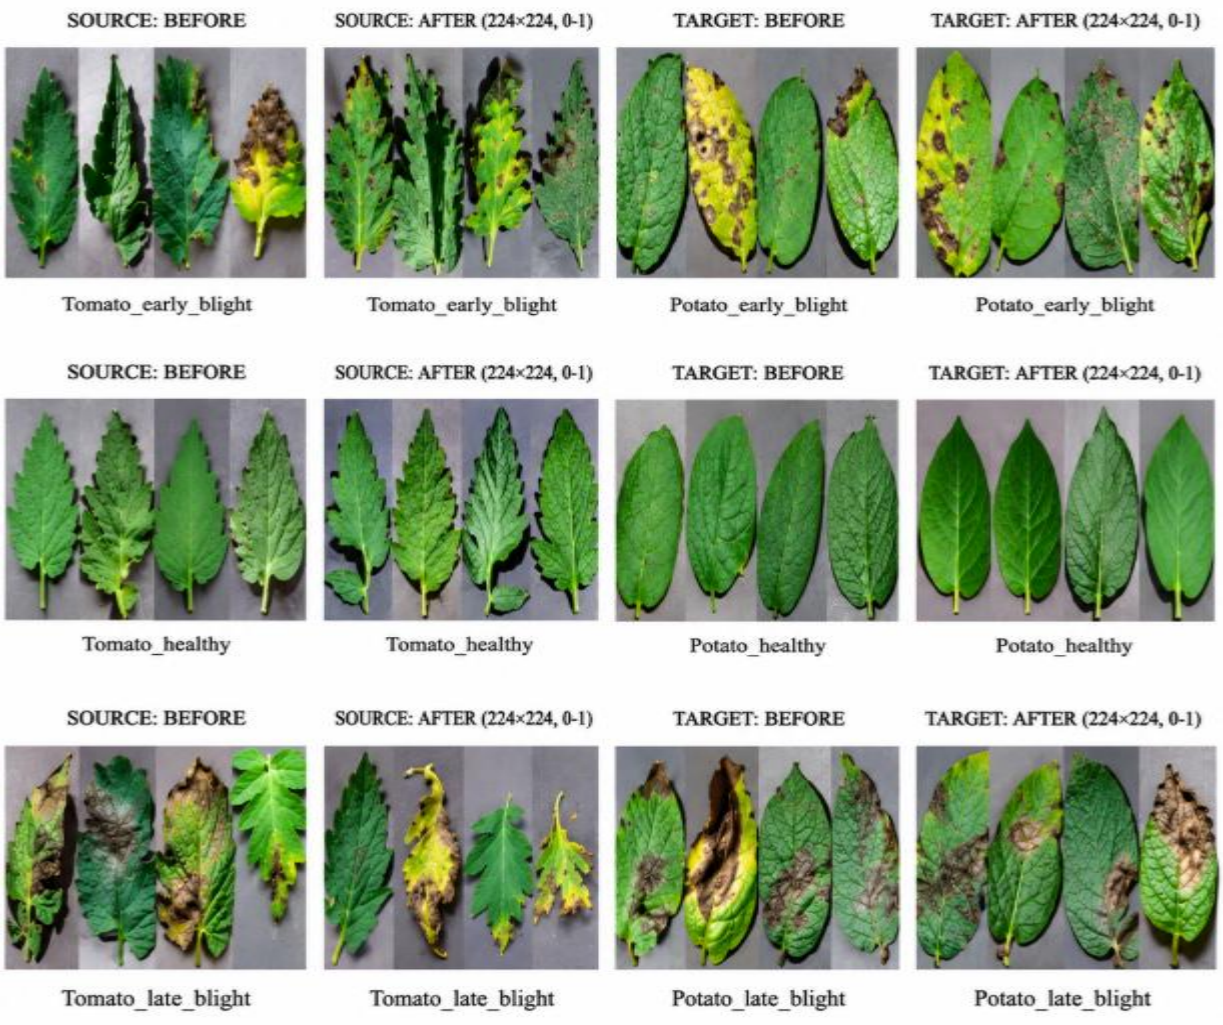


Figure 4. A 3×4 grid of pre-processed samples from source and target domains

For the source domain, we applied data augmentation exclusively to enhance intra-class variability while maintaining strict zero-shot conditions. The augmentation operations included random rotations (± 30°), horizontal and vertical flips, and zooming (10–20%). We did not augment the target domain to preserve its natural distribution during inference. We split the source data into 90% training and 10% validation subsets, and reserved all target images for zero-shot testing. Table 1 summarises the class distributions 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 DsAN using a pretrained MobileNetV2 backbone, as depicted in Figure 3. The MobileNetV2 structure serves as a feature extractor, producing 1280-dimensional visual representations from input images through its convolutional and bottleneck layers. These derived features were then processed through a projection head, which transformed them into a 300-dimensional semantic space for alignment with Word2Vec-based semantic embeddings obtained from textual class descriptions. The Elastic-Mixup component co-mixes visual and semantic feature pairs across domains using a shared coefficient drawn from a Beta (0.4, 0.4) distribution, increasing intermediate feature diversity and smoothing decision boundaries. Next, the DsAN module aligns the class-conditional source and target feature distributions using the Local Maximum Mean Discrepancy with a weak Gradient Reversal Layer (GRL), ensuring domain alignment while preserving class separability. During inference, the cosine similarity between the visual and semantic embeddings is calculated, and the class with the highest similarity score is chosen as the final prediction.

2.3 Training Configuration

Model training was performed in PyTorch on a GPU (NVIDIA GeForce GTX 1060 Ti) to accelerate computation and ensure stable multi-loss optimisation. We used the Adam optimiser 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 the pre-trained layers while enabling faster adaptation of newly initialised 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 incorporates classification, domain alignment, and mixup-regularisation losses, balancing discriminative learning with domain-invariant feature adaptation. To improve robustness, we applied early stopping and a learning rate scheduler (factor 0.5, patience 5) to prevent overfitting and promote smooth convergence across epochs.

2.4 Loss Function Formulation

The proposed framework optimises a joint loss function that combines the classification, domain alignment, and mixup-consistency terms to ensure both discriminative and domain-invariant representations. The total loss is expressed as follows: (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 minimises distributional shifts between class-conditional source and target features, and $\mathcal{L}_{mixup}$ represents the consistency loss that regularises 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, we set both to 0.5 to achieve stable convergence and optimal generalisation across the source and target domains.

2.5 Evaluation Metrics

The 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.

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

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

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

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

Accuracy measures the overall proportion of correctly classified samples in the target domain, combining both the diseased and healthy categories. High accuracy indicates that the model effectively transferred 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 the total number of predicted positives. It quantifies the model's reliability in detecting specific disease classes, where the Elastic-Mixup module minimises false positives by smoothing the decision boundaries. Recall (sensitivity) measures the ability to identify all the actual positive samples. High recall indicates that DsAN effectively aligns the source and target subdomains, enabling robust detection under cross-domain variations. F1-Score, the harmonic mean of Precision and Recall, provides a balanced assessment of both metrics. Macro- and Weighted-F1 scores summarise 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, whereas DsAN enhances semantic alignment, yielding accurate and domain-invariant plant disease recognition.

3. 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. We evaluated the model under the tomato-to-potato cross-crop protocol to assess classification accuracy, generalisation, and domain-alignment capability.

3.1 Visual Feature Extraction and Elastic-Mixup Regularisation

Figure 5 illustrates the synthetic co-mixes generated by the Elastic-Mixup module. Each tile represents a convex combination of two potato leaf classes with their respective mixing coefficients (λ). 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. Figure 5 shows how Elastic-Mixup blends visual and semantic features to create realistic hybrid samples that preserve lesion shapes, colour consistency, and textural continuity. These co-mixed representations effectively expand each class neighbourhood and smooth decision boundaries, preventing overconfidence in ambiguous textures and reducing class-imbalance bias. The biological plausibility and diversity of the mixed samples confirmed that Elastic-Mixup enhanced cross-domain generalisation and strengthened discriminative learning for unseen disease classes.

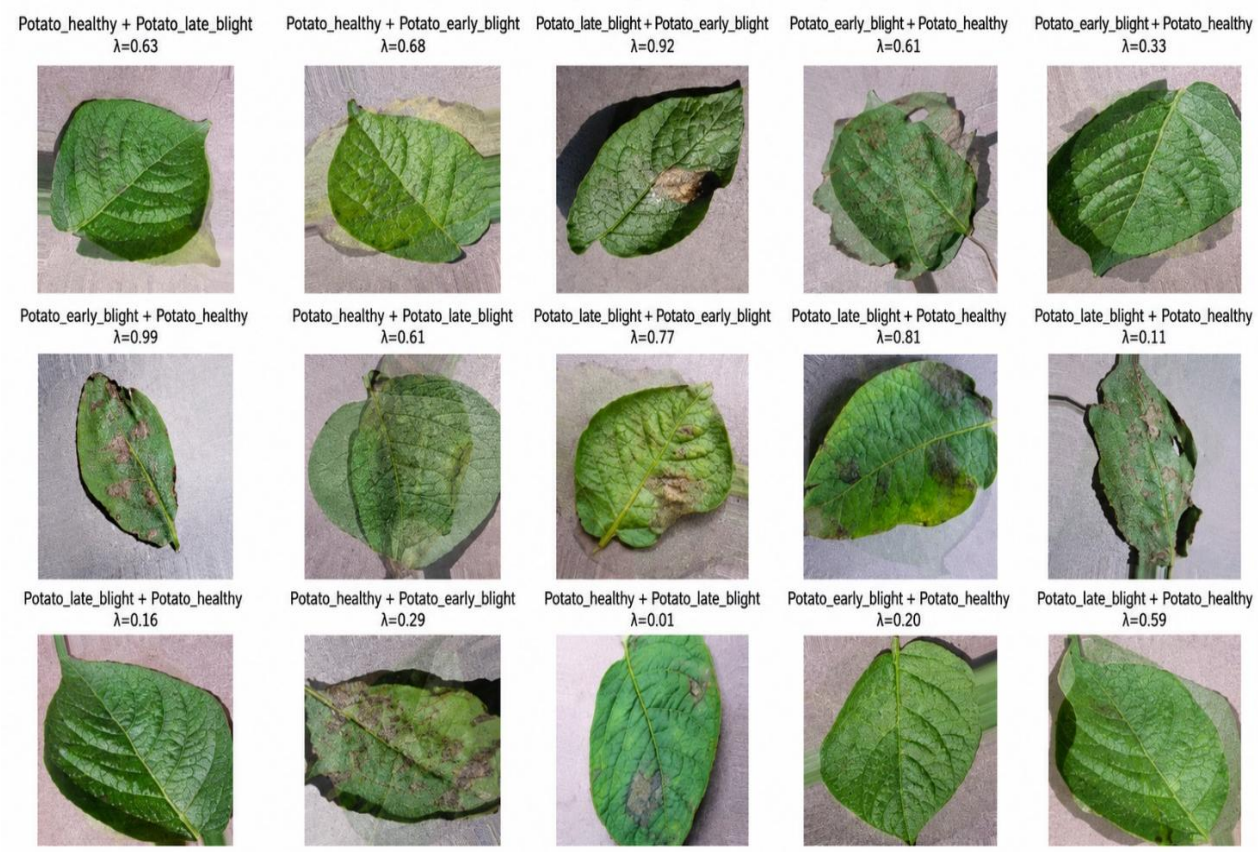


Figure 5. Elastic-mixup samples on target leaves

3.2 Deep Subdomain Alignment Network

Figure 6 presents the t-SNE plots of the target domain embeddings at successive stages: baseline ZSL, +Elastic-Mixup, and the final +Elastic-Mixup + DsAN configuration. Figure 6 shows the t-SNE visualisation of target domain embeddings at three learning stages: baseline ZSL, +Elastic-Mixup, and the final +Elastic-Mixup + DsAN configuration. The baseline plot shows partial class separation but noticeable overlap between Potato Early Blight and Potato Late Blight (PLB), indicating domain confusion. After applying Elastic-Mixup, intra-class compactness improves, and the transitional regions become smoother, reflecting enhanced boundary regularisation. With DsAN, the 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 shown in Figure 6 confirms that combining Elastic-Mixup with DsAN significantly enhances class discrimination, semantic alignment, and cross-crop generalisation.

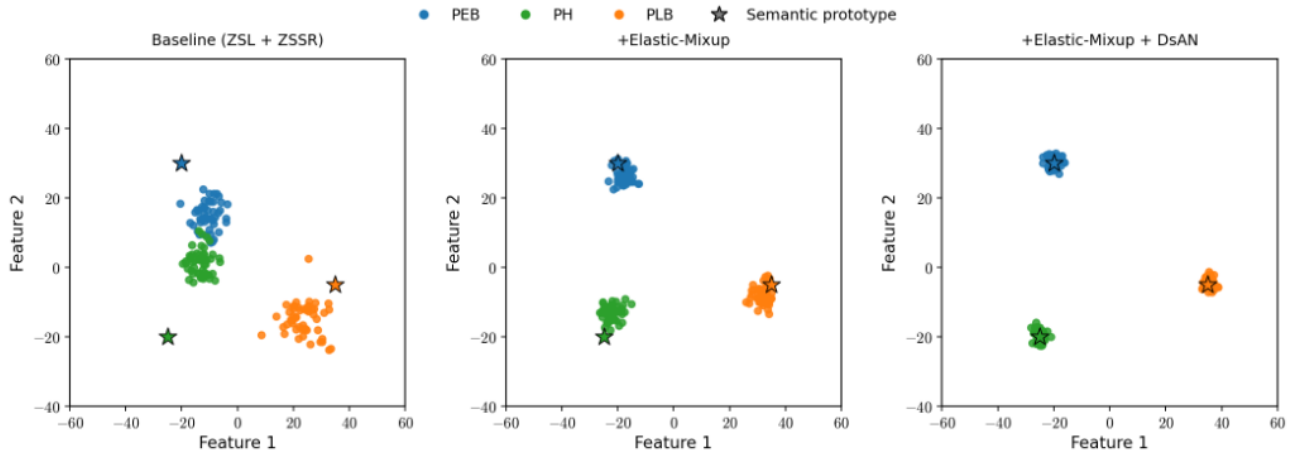


Figure 6. t-SNE visualisation of target-domain embedding evolution

3.3 Training Dynamics

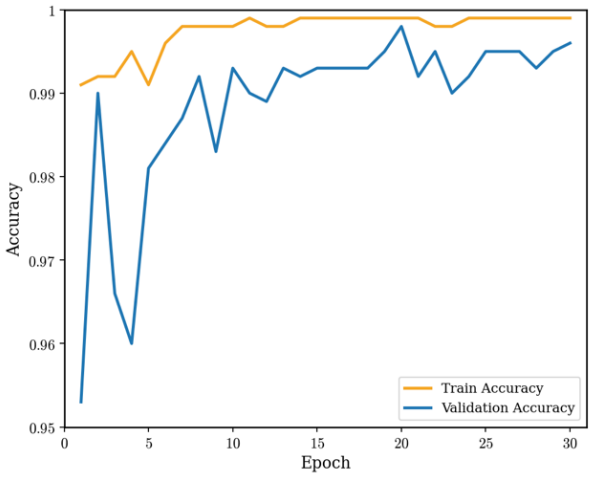
Figure 7 shows the model's training behaviour over 30 epochs. Figure 7 presents the training dynamics of the developed model across 30 epochs, showing accuracy and loss trends during optimisation. The left plot illustrates a steady increase in training and validation accuracy, rising from 0.955 to 0.997, demonstrating rapid convergence and consistent generalisation without overfitting. The right plot shows the progressive reduction in the loss components: 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 optimisation and balanced learning between the classification and domain alignment. The smooth learning curves further confirm that Elastic-Mixup stabilised feature interpolation, while DsAN improved representation alignment, resulting in robust, well-calibrated model convergence.

3.4 Target-Side Evaluation

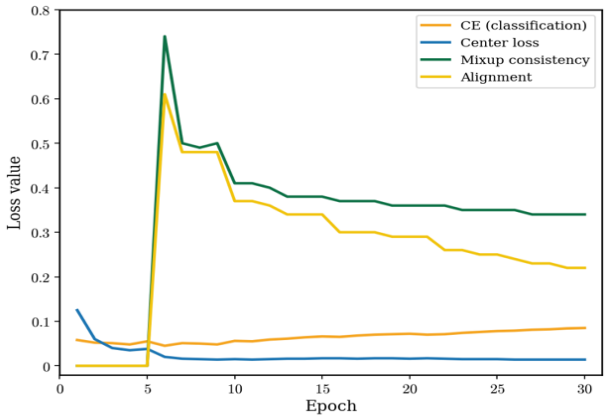
The confusion matrix in Figure 8 shows the developed model's predictions for the target (potato) test set. Figure 8 presents the confusion matrix of the developed model evaluated on the potato target-domain test set, showcasing its classification performance across three classes: Potato Early Blight, Potato Healthy, and Potato Late Blight. The strong diagonal dominance, with counts of 980, 143, and 977, indicates highly accurate predictions for all classes and reflects excellent interclass discrimination. Only minimal off-diagonal entries appear, representing a few misclassifications in visually ambiguous or low-contrast samples, such as between the 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 confirmed the model's strong predictive stability, reliable feature alignment, and consistent cross-crop generalisation performance. Table 2 summarises the quantitative results. The developed model achieved a 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 generalisation across unseen disease classes.

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



(a) Accuracy versus epoch



(b) Losses versus epoch

Figure 7. Training accuracy and loss per epoch

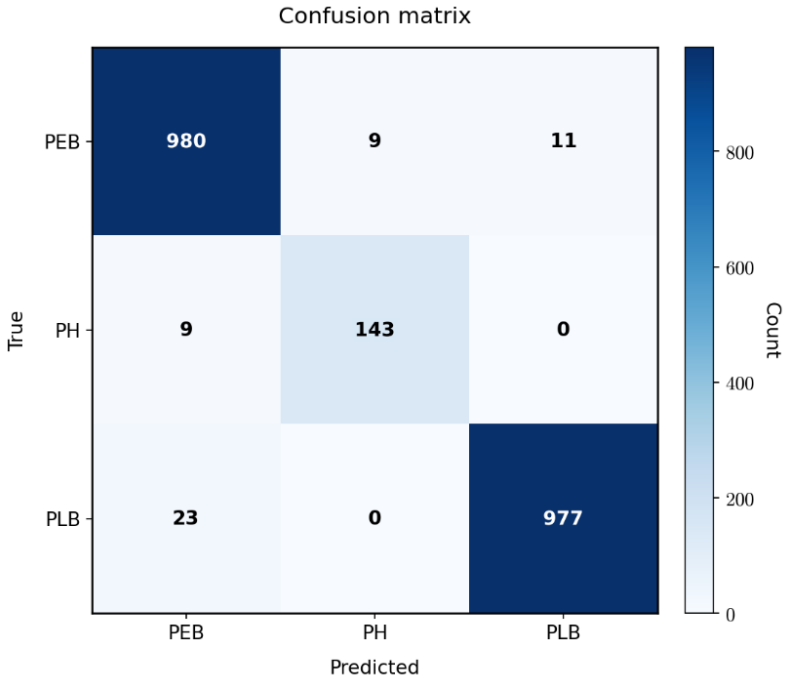


Figure 8. Confusion matrix on target (Potato) test set

3.5 Comparative Benchmarking with Existing Work

We benchmarked the developed model against the ZSTL framework proposed by Singh et al. [13] using the same evaluation metrics. Table 3 summarises the comparative outcomes. Compared with the baseline, the developed model showed a 19.24 pp accuracy improvement and a dramatic 88.9% reduction in the 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 revealed 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 contributions of the proposed components. Elastic-Mixup enhances inter-class boundary regularity and mitigates minority-class imbalance, whereas DsAN enforces class-conditional alignment and reduces 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. Thus, the developed model represents a reliable, computationally efficient, and domain-invariant approach for intelligent plant disease diagnosis in practical agricultural settings.

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

4. Conclusions

This study introduced an improved cross-crop plant disease recognition framework that integrates Elastic-Mixup and DsAN to achieve domain-invariant zero-shot learning. This study was 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. This methodology effectively mitigates source-target disparities by jointly co-mixing visual and semantic representations and aligning class-conditional distributions through LMMD-based adaptation. Under the tomato-to-potato transfer protocol, the framework achieved 97.58% accuracy, 96.60% precision, 96.59% recall, and 96.59% F1-score, an 88.9% reduction in classification error compared with the baseline. The results confirm that Elastic-Mixup smooths decision boundaries, whereas 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 (for example, 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, thereby 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. The authors also acknowledge the supervisory team for their guidance and all colleagues who participated in discussions that enhanced the experimental framework and manuscript quality.

Funding

This study was not supported by any grants from funding bodies in the public, private, or not-for-profit sectors.

Declaration of Competing Interest

The authors declare no conflict of interest.

Credit Authorship Contribution Statement

Sani Saleh SAMINU: Conventionalisation; Design; Methodology; Analysis; Writing-original draft;
Yusuf Ibrahim: Conventionalisation; Design; Methodology; Supervision; Writing-review & revise
Zaharuddeen Haruna: Methodology; Analysis; Validation; Writing-review & revise
Abubakar Umar: Investigation; Validation; Writing-review & revise

Availability of Data and Materials

Data sharing is not applicable to this article, as this study generated or analysed no new datasets.

Ethics Declarations

This study did not involve human participants or animal subjects. Therefore, ethical approval was not required.

Generative Artificial Intelligence Declarations

No generative artificial intelligence tools were used to prepare this manuscript. All intellectual content is original and is solely attributable to the authors.

References

[1] M. I. Hossen, M. Awrangjeb, S. Pan, and A. A. Mamun, "Transfer learning in agriculture: A review," *Artificial Intelligence Review*, vol. 58, Art. no. 97, 2025, https://doi.org/10.1007/s10462-024-11081-x

[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, Art. no. 106892, 2022, https://doi.org/10.1016/j.compag.2022.106892

[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, Art. no. 101289, 2021, https://doi.org/10.1016/j.ecoinf.2021.101289

[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, https://doi.org/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*, vol. 20, no. 4, pp. 789–799, 2024, https://doi.org/10.1007/s11334-022-00513-y

[7] A. 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 Sustainability*, vol. 5, Art. no. 1383182, 2024, https://doi.org/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, Art. no. 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, Art. no. 0038, 2023, https://doi.org/10.34133/plantphenomics.0038

[10] R. Sharma, A. Singh, Kavita, *et al.*, "Plant disease diagnosis and image classification using deep learning," *Computers, Materials & Continua*, vol. 71, no. 2, pp. 2125–2140, 2022, https://doi.org/10.32604/cmc.2022.020017

[11] N. Ullah, J. A. Khan, S. Almakdi, *et al.*, "An effective approach for plant leaf diseases classification based on a novel DeepPlantNet deep learning model," *Frontiers in Plant Science*, vol. 14, Art. no. 1212747, 2023, https://doi.org/10.3389/fpls.2023.1212747

[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, Art. no. 305, 2024. https://doi.org/10.1007/s10462-024-10950-9.

[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, https://doi.org/10.1109/ACCESS.2023.3343759

[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, Art. no. e0282250, 2023. https://doi.org/10.1371/journal.pone.0282250.

[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*, 2022, https://doi.org/10.1007/s10479-022-04575-w

[16] Y. Gulzar, "Fruit image classification model based on MobileNetV2 with deep transfer learning technique," *Sustainability*, vol. 15, no. 3, Art. no. 1906, 2023, https://doi.org/10.3390/su15031906

[17] S. U. M. Rao, K. Sreekala, P. S. Rao, N. Shirisha, G. Srinivas, and E. Sreedevi, "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. https://doi.org/10.11591/ijeecs.v36.i3.pp1721-1730.

[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, Art. no. 106314, 2021, https://doi.org/10.1016/j.compag.2021.106314

[21] B. Yang, M. Li, F. Li, *et al.*, "A novel plant type, leaf disease and severity identification framework using CNN and transformer with multi-label method," *Scientific Reports*, vol. 14, Art. no. 11664, 2024, https://doi.org/10.1038/s41598-024-62452-x

[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, arXiv:2109.05493, 2021. [Online]. Available: https://arxiv.org/abs/2109.05493.

[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] A. Prasad, N. Mehta, M. Horak, and W. D. Bae, "A two-step machine learning approach for crop disease detection using GAN and UAV technology," Remote Sensing, vol. 14, no. 19, Art. no. 4765, 2022. https://doi.org/10.3390/rs14194765.

[26] Y. F. Zhang, J. Wang, J. Liang, et al., "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.