E-AaSSL: Hybrid EfficientNet-DeiT Framework for Ambiguity-Aware Semi-Supervised Leaf Disease Classification

Chetan Bhatkar[1]* and Deepak D. Kshirsagar[2]

Department of Computer Science and Engineering, COEP Technological University,
Pune, Maharashtra, India
bhatkarcd24.comp@coeptech.ac.in[1]
ddk.comp@coeptech.ac.in[2]

Abstract. Deep learning has achieved high accuracy in plant disease classification under fully supervised settings; however, real-world agricultural applications are constrained by limited data and class imbalance. To address these challenges, this work proposes an enhanced ambiguity-aware semi-supervised learning (E-AaSSL) framework tailored for plant leaf disease diagnosis, aiming to substantially reduce annotation demands while maintaining strong diagnostic performance.

The proposed pipeline integrates (i) a pretrained EfficientNet-B4 warm-up phase for reliable initial pseudo-labels, (ii) ambiguity-aware pseudo-label filtering with class-balanced selection and FixMatch-style consistency regularization to suppress noise accumulation, and (iii) progressive refinement using a Vision Transformer (DeiT) through staged layer unfreezing. An ensemble of EfficientNet, a hybrid CNN–Transformer model and DeiT is further employed with an adaptive rejection mechanism to improve prediction reliability.

Experimental assessments on plant leaf disease datasets demonstrate that the framework achieves high accuracy and F1 scores under severe label scarcity, while maintaining strong selection accuracy and effective rejection of ambiguous samples. These results show that the proposed approach offers a scalable, practical solution for precision agriculture, enabling real-time crop health monitoring and minimizing economic losses in resource-constrained farming environments.

Keywords: Semi-supervised learning, plant disease classification, ambiguity rejection, knowledge distillation, precision agriculture.

1 Introduction

Crop diseases present a critical challenge to global food security because they directly lower agricultural output and degrade the standard of the crops that do survive. This is especially true for smallholder and resource constrained farming

---

* Corresponding author: bhatkarcd24.comp@coeptech.ac.in

S. Sharma et al. (eds.), *Proceedings of the International Conference on Recent Advances in Intelligent and Sustainable Technologies (RAIST 2026)*, Atlantis Highlights in Intelligent Systems 17,
https://doi.org/10.2991/978-94-6239-707-1_10

systems, where limited access to domain experts often leads to delayed diagnosis and ineffective treatment [1, 2]. Traditional diagnosis based on manual inspection by agricultural experts is time-consuming, subjective, and difficult to scale, motivating the development of automated disease recognition systems using leaf images as a key component of precision agriculture [2, 3].

Deep convolutional neural networks (CNNs) have become the dominant approach for plant disease classification, consistently outperforming handcrafted feature-based methods on benchmark datasets [4, 11]. Architectures such as Inception-v3 and VGG19, often combined with classical classifiers, can achieve accuracies exceeding 99% on curated datasets like PlantVillage. However, these fully supervised models depend heavily on large, balanced, and carefully annotated datasets, making them costly to deploy and less robust when applied to heterogeneous real-world field conditions [4, 6].

To reduce annotation costs, SSL and self-supervised strategies have attracted growing attention in agricultural image analysis [6, 8]. By leveraging unlabeled images, these methods aim to close the performance gap between fully supervised models and practical, data-scarce scenarios. For instance, confidence-guided pseudo-labeling in semi-supervised settings has been shown to improve performance over transfer-learning baselines on the PlantVillage dataset [12]. Similar gains have been reported for weed detection and cassava leaf disease grading, and recent surveys confirm that unlabeled field imagery can substantially lower annotation effort while maintaining competitive performance [6, 9].

Despite these advances, most existing semi-supervised approaches optimize overall accuracy without explicitly addressing ambiguous or low-confidence samples. In agricultural diagnosis, this limitation is critical: incorrect predictions may directly influence pesticide use, crop management, and farmer livelihoods. The ability to balance diagnostic precision against coverage therefore becomes as important as raw accuracy [7, 15].

The Ambiguity-aware Semi-Supervised Learning (AaSSL) framework was introduced to solve the problem by explicitly modeling uncertainty during pseudo-label selection. AaSSL uses per-class ambiguity rejection modules that estimate confidence thresholds using a binomial cumulative distribution, this lets optimization of selection accuracy (SA) and coverage (COV) [7]. On BRACOL coffee and banana leaf datasets, AaSSL achieves selection accuracies close to 99% while using only 5–50% labeled data, this shows effectiveness of ambiguity control under label scarcity. However, AaSSL relies on a single DeiT backbone and does not used consistency-regularized pseudo-labeling, hybrid CNN–Transformer fusion, or knowledge distillation, limiting its ability to exploit complementary local and global feature representations [7, 17].

Recent studies in agricultural vision increasingly show that hybrid and ensemble architectures provide greater robustness than single model designs [2, 17, 16]. CNN–Transformer hybrids, such as PlantHealthNet, combine the strong local feature extraction of CNNs with the global contextual modeling of Vision Transformers[5], leading to improved disease classification and severity estimation compared to standalone CNNs or Transformers [17]. Similarly, multi model deep ensembles enhance stability and accuracy across diverse crops, although most such approaches remain fully supervised and annotation-intensive [2, 14]. Motivated by these limitations, this work proposes an Enhanced Ambiguity-Aware Semi-Supervised Learning (E-AaSSL) framework that extends AaSSL [7] in three key ways. An EfficientNet-B4 teacher is used for warm-up training. Second, per-class ambiguity-aware rejection is combined with class-balanced top-K selection and FixMatch-style weak-strong consistency regularization to suppress noisy pseudo-labels and mitigate error accumulation during iterative training, thereby generating reliable pseudo-labels from limited labeled data [7, 10, 15]. Third, a hybrid CNN–Transformer model and a refined DeiT student are trained via knowledge distillation from the teacher, forming a robust ensemble suitable for resource constrained deployment [17].

Experiments on public datasets, including BRACOL coffee leaves shows improved accuracy and F1-score under limited labeled data while maintaining reliable selective predictions through explicit ambiguity rejection, making the framework a practical and reliable solution for real world plant disease diagnosis [7, 13].

2 Literature Review

2.1 Supervised deep learning for plant disease diagnosis

Early works replaced handcrafted color, texture, and shape features and classical classifiers with convolutional neural networks (CNNs), achieving substantial gains in accuracy and robustness [2, 18]. Sujatha et al. integrated Inception-v3 and VGG19 feature extractors with SVM, kNN, and other machine learning classifiers on banana, custard apple, fig, and potato leaf datasets, reporting up to 99.1% accuracy on custard apple leaves and 91.9% on banana leaves [4]. Ensemble based systems further improve robustness by aggregating multiple CNNs, but they require full annotation and significant compute, making them less suitable for smallholder or edge deployments [14].

2.2 Semi-supervised and few-shot learning in agriculture

To reduce labeling requirements, semi-supervised learning (SSL) and few-shot methods have been proposed for plant disease recognition [6, 8]. Semi-supervised few-shot strategies based on confidence-guided pseudo-label selection have shown promising performance in low-label regimes, reducing dependence on fully labeled training sets [13]. Semi-supervised segmentation driven pipelines for cassava leaf disease grading and SSL frameworks for multi class weed detection show that combining 10–30% labeled data with unlabeled images can approach the performance of fully supervised baselines, highlighting the value of SSL in agricultural imaging [9, 13, 21].

2.3 Self-supervised and representation learning approaches

Self-supervised learning (Self-SL) has been explored to pretrain encoders on large collections of unlabeled field images before fine-tuning on limited labeled data [6]. A recent systematic review concludes that self/semi-supervised pretraining often matches or slightly improves fully supervised accuracy while using a fraction of labels, especially under domain shift between lab and field conditions [6]. However, most of these methods remain focused on maximizing accuracy and give limited attention to selective prediction, confidence calibration, or explicit handling of ambiguous inputs in safety critical agronomic decisions [7, 15].

2.4 Ambiguity-aware and hybrid CNN–Transformer models

Ambiguity-aware Semi-Supervised Learning (AaSSL) directly addresses pseudo label noise and uncertain predictions in leaf disease classification [7]. Using a DeiT backbone, AaSSL estimates per class confidence thresholds via a binomial cumulative distribution model and applies ambiguity rejection both when selecting pseudo labels (ARS) and during final prediction (ARP), jointly optimizing selection accuracy (SA) and coverage (COV) [7]. On BRACOL coffee and banana leaf datasets, AaSSL achieves SA close to 99% with only 5–50% of labeled data and, at 50% labels, reaches 96.58% and 99.23% accuracy nearly identical to 100% label training [7]. In parallel, hybrid and transformer enhanced architectures such as PlantHealthNet fuse CNNs with Vision Transformers to exploit both local lesion details and global contextual cues, outperforming pure CNN or pure Transformer baselines under full supervision [5, 17]. Multi model deep ensembles similarly improve accuracy and stability across crops but remain annotation and compute intensive [2, 16].

2.5 Research Gaps

1. Supervised CNN and ensemble models achieve high classification accuracy but require fully labeled datasets and lack explicit ambiguity handling [4, 11, 14].
2. Semi-supervised, self-supervised approaches and few-shot learning methods successfully reduce label requirements; however, they mainly rely on simple confidence thresholds and optimize accuracy without explicitly controlling selection accuracy, coverage, or rejection behavior [6, 13, 15].
3. AaSSL introduces principled per class ambiguity rejection and demonstrates strong label efficiency, but it is limited to a single DeiT backbone, does not exploit FixMatch style weak strong consistency regularization, and does not integrate hybrid CNN–Transformer fusion or knowledge distillation to exploit complementary feature representations [7, 17].

This leaves a clear research gap. Few existing frameworks jointly (i) perform semi-supervised learning under severe label scarcity, (ii) enforce per-class ambiguity-aware pseudo-label selection and test-time rejection with high SA, and (iii) leverage a hybrid CNN–Transformer ensemble via knowledge distillation to maximize accuracy while remaining deployable on resource-constrained agricultural devices [7, 17, 20]. The proposed E-AaSSL framework is designed specifically to address this gap.

3 Proposed System Overview

The proposed Enhanced Ambiguity-aware Semi-Supervised Learning (E-AaSSL) framework extends AaSSL by combining an EfficientNet-B4 teacher, FixMatch-style consistency regularization, per-class ambiguity rejection, and hybrid CNN-Transformer refinement for leaf disease classification on datasets such as BRACOL and banana leaves [7, 15, 17]. Hyperparameters (learning rate, batch size, the binomial-threshold parameter α, and the knowledge-distillation parameters α_{KD} and τ) are optimized experimentally, and the modular components allow easy adaptation to other crop datasets and label budgets (e.g., 25% labeled). The pipeline has two main modules: (1) semi-supervised teacher training with ambiguity rejection and (2) hybrid refinement with ensemble evaluation, as summarized in Fig. 1, which illustrates the flow from raw images to rejected or classified outputs.

3.1 Module 1: Semi-Supervised Training with Ambiguity Rejection

1. **Dataset preparation and label budgeting:**
   Raw leaf images are resized to 224×224 and normalized using ImageNet statistics. The dataset is pre-partitioned into training, validation, and test sets, after which a fixed proportion $P_{\text{label\_ratio}}$ (e.g., 0.25) of the training data is designated as labeled and the remainder is treated as unlabeled, preserving per-class distributions and enabling label-scarce evaluation [7, 21].
2. **Augmentations and dataloaders:**
   Weak augmentations consist of random horizontal flipping and ImageNet normalization for supervised training, while strong augmentations extend this pipeline with RandAugment (e.g., num_ops = 2, magnitude = 9) to support FixMatch-style consistency checks [10, 15]. Custom dataset classes provide labeled samples (image, label) and unlabeled pairs (weak and strong views, path), and PyTorch dataloaders use a batch size of 18 with shuffling during training.
3. **Warm-up prototype model:**
   A pretrained EfficientNet-B4 backbone is fine-tuned on the labeled subset using cross-entropy loss with label smoothing $\epsilon = 0.05$ and the AdamW optimizer with a fixed learning rate (e.g., LR = 0.0044, 60 epochs). The trained network is then adopted as the initial teacher model. It generates reliable class-probability predictions that are subsequently used to estimate confidence thresholds and to select high-quality pseudo-labels during the semi-supervised refinement stage [7, 17].

4. **Per-class threshold estimation:**
   The validation set is passed through the initial teacher model to obtain prediction confidence scores. For each class, candidate thresholds are generated from these scores, and a binomial reliability test ($\alpha = 0.018$) is used to select the optimal class-specific threshold by balancing selection accuracy and coverage. The learned thresholds are then reused during semi-supervised training and applied for selective prediction at test time [7].
5. **Ambiguity-aware pseudo-label generation:**
   For each unlabeled image, the teacher predicts labels on weak and strong augmented views, and a sample is accepted as a pseudo-label if (i) the weak and strong argmax classes agree and (ii) the maximum probability exceeds λ_c for the predicted class. To avoid class imbalance and over-representation of easy samples, the number of accepted pseudo-labels is capped by selecting the top 130 high-confidence samples per class. The selected pseudo-labeled images and their predicted labels are then used for subsequent training [7, 15].
6. **Iterative teacher refinement:**
   Labeled and pseudo-labeled samples are merged into a single training set and used to iteratively fine-tune the EfficientNet-B4 teacher over several refinement rounds (e.g., 7 iterations × 10 epochs at LR = 0.00044), where pseudo-labels regenerated at each iteration are used to progressively improve training quality. The final refined teacher model provides stronger feature representations and more reliable confidence estimates under label-scarce conditions, and serves as the knowledge-distillation source for the subsequent student models [7, 21].

3.2 Module 2: Hybrid Model Refinement and Ensemble Evaluation

1. **Hybrid CNN–Transformer training with KD:**
   To leverage complementary local and global representations, a hybrid CNN–Transformer model is constructed by combining EfficientNet-B4 feature representations with DeiT feature embeddings and guiding the training process with knowledge distillation from the refined teacher, followed by a fully connected fusion head. The hybrid model is trained on the merged labeled and pseudo-labeled dataset using a combined knowledge-distillation loss defined as

$$\mathcal{L} = \alpha \cdot \tau^2 \cdot \mathrm{KL}\big(\sigma(\mathbf{y}_s/\tau) \,\|\, \sigma(\mathbf{y}_T/\tau)\big) + (1-\alpha) \cdot \mathcal{L}_{CE}, \tag{1}$$

   where KL represents the Kullback–Leibler divergence between the softened predictions of the student and teacher models, $\mathcal{L}_{CE}$ denotes the cross-entropy term, α_{KD} controls the balance between distillation and supervised learning ($\alpha_{\mathrm{KD}} = 0.25$), and τ is the temperature parameter ($\tau = 1.5$). This formulation enables the student model to benefit from both soft teacher guidance and label supervision. [17, 19].

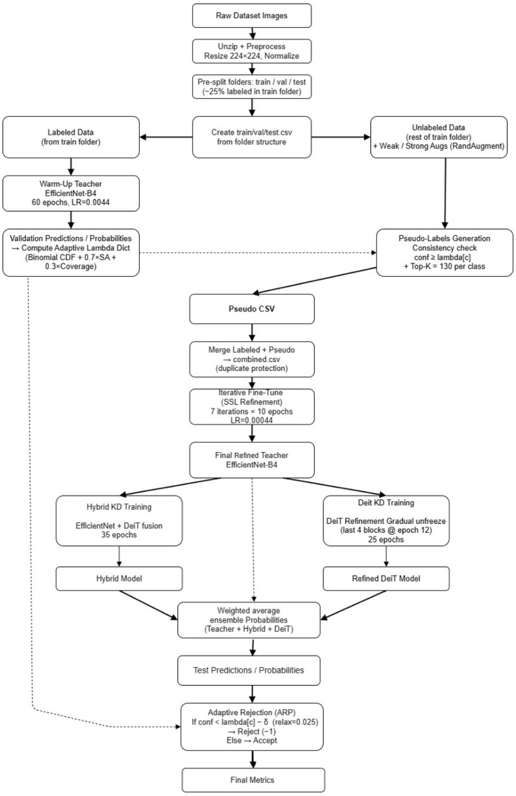

Fig. 1. Architecture of the proposed (E-AaSSL) framework for leaf disease classification.

2. **DeiT refinement with progressive unfreezing:**
   A pretrained DeiT model is further refined using the merged dataset and pseudo-labeled dataset under the guidance of knowledge distillation from the refined teacher network f_T, initially training only the classification head and gradually unfreezing deeper transformer blocks (e.g., the last four blocks after 12 epochs with a reduced learning rate) to avoid overfitting in low-label regimes. This strategy produces a compact yet context-aware transformer model suitable for efficient inference [7, 17].
3. **Ensemble inference and ambiguity-aware rejection:**
   At test time, each image is passed through three models: the refined teacher f_T, the hybrid CNN–DeiT network, and the refined DeiT model. Their outputs are combined using weighted averaging of softmax probabilities. A selective rejection strategy is then applied: if the maximum ensemble confidence for the predicted class falls below the corresponding class-specific threshold λ_c, the prediction is rejected as ambiguous; otherwise, it is accepted. A small relaxation margin (RELAX) is applied only during the final ambiguity-aware rejection stage to prevent unnecessary rejection caused by minor confidence

fluctuations in ensemble predictions around λ_c [7]. Classification performance is evaluated on the accepted predictions using precision, recall, and F1-score, while overall accuracy is computed across the entire test set. In addition, coverage and rejection accuracy are reported to measure the reliability of the selective prediction strategy. The obtained results are compared with baseline methods to highlight the effectiveness of the proposed approach under limited annotation settings [7, 2, 16].

4 Results

4.1 Dataset and Implementation Details

The BRACOL dataset comprises 2209 leaf images in five classes: Cercospora, Healthy, Miner, Phoma, and Rust [7]. The dataset was already organized into training, validation and test folders, resulting in approximately 1,539 training images and 335 images each for validation and testing. This corresponds roughly to a 70/15/15 split while maintaining class balance. To simulate a limited-annotation scenario, only about 25% of the training data was treated as labeled, while the remaining images were considered unlabeled for semi-supervised learning. All images were resized to 224×224 pixels and normalized using ImageNet statistics to ensure compatibility with pretrained models.[7, 21].

Optimized hyperparameters included learning rates of 0.0044 (warm-up phase) and 0.00044 (fine-tune phase), batch size = 18, and adaptive binomial thresholding with $\alpha = 0.018$ to select reliable pseudo-labels while balancing confidence and coverage. Knowledge distillation used $\alpha_{\text{KD}} = 0.25$ and temperature $\tau = 1.5$ for the hybrid and DeiT student models. The merged training set after pseudo-labeling comprised around 1,200–1,350 samples.

4.2 Overall Performance

The evaluation utilized an ensemble consisting of a fine-tuned EfficientNet-B4 teacher, a hybrid EfficientNet-DeiT fusion model, and a DeiT-Small student. An Ambiguity-aware rejection strategy was applied to filter uncertain predictions and improve reliability [7].

Table 1. Performance of the Proposed Ensemble Model with Ambiguity-Aware Rejection

Metric	**Value**
Accuracy	96.72%
F1-Score	97.83%
Selection Accuracy(SA)	98.03%
Precision	97.84%
Recall	97.83%
Coverage(COV)	90.75%
Reject Accuracy	16.13%

The framework achieved a Accuracy of 96.72%, with a selection accuracy of 98.03%. The model maintained strong classification performance with an F1-score of 97.83%, along with Precision 97.84% and Recall 97.83% indicating that the model maintains high discriminative power across all disease categories despite the 75% unlabeled data ratio. A Coverage of 90.75% further confirms that a large portion of samples were confidently classified, while a rejection Accuracy of 16.13% reflects the model's ability to identify uncertain predictions under the selective classification setting [7].

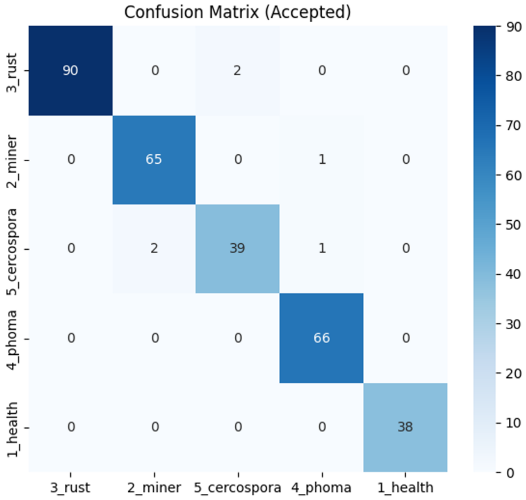

Fig. 2. Confusion Matrix of the Proposed Ensemble Model on the Non-Rejected Test Samples

4.3 Comparison with the AaSSL on BRACOL Dataset (25% Labeled Data)

Table 2. Performance comparison with baseline methods on the BRACOL dataset (25% labeled data)

Method	**Accuracy (%)**	**F1-Score (%)**	**Selection Accuracy**	**Coverage (%)**	**Rejection Accuracy**
EfficientNet-B4 (Supervised Baseline)	89.85	89.18	–	100	–
Ensemble (Without Rejection)	94.33	94.12	–	100	–
AaSSL [7]	95.98	95.51	99.08	85.88	22.65
Proposed Method	**96.72**	**97.83**	**98.03**	**90.75**	**16.13**

All methods were evaluated using the same 25% labeled training subset, and remaining data are treated as unlabeled to ensure a fair semi-supervised comparison. The supervised EfficientNet-B4 baseline achieved 89.85% accuracy and

89.18% F1-score, shows limitations of supervised learning using scarce labels. An ensemble without rejection improves accuracy to 94.33% and 94.12% F1-score , shows the benefit of model combination, but it lacks uncertainty filtering. The Aassl[7] method has achieved an accuracy of 95.98% and an F1-score of 95.51% with a selection accuracy of 99.08%, but a relatively lower coverage of 85.88% due to conservative rejection. The proposed method has achieved higher accuracy 96.72% and F1-score 97.83%, while maintaining strong selective accuracy 98.03% with improved coverage 90.75%, indicates more samples are confidently classified. The rejection accuracy of 16.13% indicates effective filtering of uncertain predictions while preserving reliable classifications.

5 Conclusion and Future Work

This work presented an enhanced ambiguity-aware semi-supervised framework that improves classification accuracy with maintaining high selection accuracy. The proposed method has achieved better accuracy , balanced F1-score and higher coverage compare to the baseline and Aassl method , demonstrating its effectiveness under limited labeled data conditions.

Future research can explore adapting the framework to large and more diverse agriculture datasets to improve generalization. Integrating lightweight models for real-time deployment on resource constrained agricultural devices and further enhancement can be achieved through improve uncertainty estimation and adaptive semi-supervised strategies. Finally, extending the framework to other crop disease monitoring applications could broaden its practical impact.

References

1. Zhao, J., Xu, L., Ma, Z., Li, J., Wang, X., Liu, Y., Du, X.: A review of plant leaf disease identification by deep learning algorithms. *Frontiers in Plant Science* **16**, 1637241 (2025)
2. Hamid, M.K., Shah, S.K., Husnain, G., Ghadi, Y.Y., Al Maaytah, S.A., Qahmash, A.: Enhancing agricultural disease detection: A multi-model deep learning novel approach. *Engineering Reports* **7**(1), e13113 (2025)
3. Jannat, M., Uddin, M.S., Hasan, M.A., Alam, M.S., Paul, A., Chowdhury, M.E.H., Haider, J.: Real-time jute leaf disease classification using an explainable lightweight CNN via a supervised and semi-supervised self-training approach. *Frontiers in Plant Science* **16**, 1647177 (2025)
4. Sujatha, R., Krishnan, S., Chatterjee, J.M., Gandomi, A.H.: Advancing plant leaf disease detection integrating machine learning and deep learning. *Scientific Reports* **15**, 11552 (2025)
5. Salman, Z., Muhammad, A., Han, D.: Plant disease classification in the wild using vision transformers and mixture of experts. *Frontiers in Plant Science* **16**, 1522985 (2025)
6. Mamun, A.A., Ahmedt-Aristizabal, D., Zhang, M., Hossen, M.I., Hayder, Z., Awrangjeb, M.: Plant disease detection using self-supervised learning: A systematic review. *IEEE Access* **12**, 171926–171943 (2024)

7. Pham, T.-C., Nguyen, T.-N., Nguyen, V.-D.: Ambiguity-aware semi-supervised learning for leaf disease classification. *Scientific Reports* **15**, 14070 (2025)
8. Zhao, L., Hao, J., Li, D., Yu, J.: Semi-supervised generative adversarial network for plant leaf disease detection. *Engineering Applications of Artificial Intelligence* **161**, 112103 (2025)
9. Velayudhan, D., Shafay, M., Khawaja, S.G., Hassan, T.: Semi-supervised segmentation-driven classification pipeline for grading cassava leaf diseases. In: *Proc. 10th Int. Conf. Engineering and Emerging Technologies (ICEET)*, Dubai, UAE (2024)
10. Zhou, S., Tian, S., Yu, L., Wu, W., Zhang, D., Peng, Z., Zhou, Z., Wang, J.: FixMatch-LS: Semi-supervised skin lesion classification with label smoothing. *Biomedical Signal Processing and Control* **84**, 104709 (2023)
11. Salka, T.D., Hanafi, M.B., Rahman, S.M.S.A.A., Zulperi, D.M., Omar, Z.: Plant leaf disease detection and classification using convolution neural networks model: A review. *Artificial Intelligence Review* **58**, 322–387 (2025)
12. Li, Y., Chao, X.: Semi-supervised few-shot learning approach for plant diseases recognition. *Plant Methods* **17**(1), 68 (2021)
13. Li, J., Chen, D., Yin, X., Li, Z.: Performance evaluation of semi-supervised learning frameworks for multi-class weed detection. *Frontiers in Plant Science* **15**, 1396568 (2024)
14. Kondaveeti, H.K., Ujini, K.G., Pavankumar, B.V.V., Tarun, B.S., Gopi, S.C.: Plant disease detection using ensemble learning. In: *Proc. 2nd Int. Conf. Computational Systems and Communication (ICCSC)* (2023)
15. Ilsever, M., Baz, I.: Consistency regularization based semi-supervised plant disease recognition. *Smart Agricultural Technology* **9**, 100613 (2024)
16. Ardah, H., Alrahhal, M., Abd-Elhafiez, W.M., Trabay, D.: Robust coffee plant disease classification using deep learning and advanced feature engineering techniques. *PeerJ Computer Science* **11**, e3386 (2025)
17. Iqbal, A.: PlantHealthNet: Transformer-enhanced hybrid models for disease diagnosis and severity estimation in agriculture. *IEEE Access* **13**, 101160–101176 (2025)
18. Mzoughi, O., Yahiaoui, I.: Deep learning-based segmentation for disease identification. *Ecological Informatics* **75**, 102000 (2023)
19. Zhao, J., Yao, L., Cheng, W., Yu, M., Shi, W., Liu, J., Jiang, Z.: Co-training semi-supervised medical image segmentation based on pseudo-label weight balancing. *Medical Physics* **52**, 3854–3876 (2025)
20. Sanyal, M., Dasgupta, S., Chaki, R., Saha, S.: Pseudo CNN: An updated CNN-based semi-supervised learning model for leaf disease detection. *Innovations in Systems and Software Engineering* **21**, 1511–1520 (2025)
21. Fan, K.-J., Liu, B.-Y., Su, W.-H., Peng, Y.: Semi-supervised deep learning framework based on modified pyramid scene parsing network for multi-label fine-grained classification and diagnosis of apple leaf diseases. *Engineering Applications of Artificial Intelligence* **151**, 110743 (2025)