

Few-Shot Learning for Early Detection of Coconut Whitefly Infestations using Low-Label Leaf Image Data in Sri Lanka

Vithiyasahar Vigneswaran

Faculty of Computing

Sri Lanka Institute of Information

Technology, Northern Uni

Jaffna, Sri Lanka.

vithiyasahar.v@slit.lk

Vinayagamoorthy Balalojanan

Faculty of Computing

Sri Lanka Institute of Information

Technology, Northern Uni

Jaffna, Sri Lanka.

balalojanan.v@slit.lk

Theevika Sukirthan

Faculty of Computing

Sri Lanka Institute of Information

Technology, Northern Uni

Jaffna, Sri Lanka.

theevika.a@slit.lk

Niroji Thayalan

Faculty of Computing

Sri Lanka Institute of Information

Technology, Northern Uni

Jaffna, Sri Lanka.

niroji.t@slit.lk

Kanagasabai Thiruthanigesan

Faculty of Computing

Sri Lanka Institute of Information

Technology, Northern Uni

Jaffna, Sri Lanka.

thiruthanigesan.k@slit.lk

Prasanna Sumathipala

Faculty of Computing

Sri Lanka Institute of Information

Technology

Malabe, Sri Lanka.

prasanna.s@slit.lk

Siyamalan Manivannan

Department of Computer Science

University of Jaffna, Sri Lanka.

siyam@univ.jfn.ac.lk

Abstract—Coconut whitefly (*Aleurodicus dispersus*) has emerged as a devastating pest threatening Sri Lanka's vital coconut industry, causing severe yield losses and economic hardship for over 1.5 million smallholder farmers. Early detection is critical for effective intervention, but conventional field scouting is labor-intensive and often too late to prevent widespread infestation. While deep learning offers promise for automated pest detection, its deployment is hindered by the scarcity of labeled training data. To address this challenge, we propose a novel SimCLR → FixMatch framework that combines self-supervised contrastive learning with semi-supervised fine-tuning for accurate detection of coconut whitefly infestations under extreme label scarcity. Our method leverages a new dataset, CocoInfest-SL, comprising 1,600 real-world leaf images from major coconut-growing regions in Sri Lanka, with only 120 labeled samples (30 per class), 7.5% of the full dataset. The model first learns robust feature representations via SimCLR pretraining on both labeled and unlabeled data, then fine-tunes using FixMatch with confidence-based pseudo-labeling. Experimental results demonstrate that our approach achieves 91.8% accuracy and a 0.91 F1-score, outperforming state-of-the-art few-shot and semi-supervised baselines while using minimal annotations. This work presents a scalable, low-cost solution for early pest monitoring in resource-constrained agricultural environments.

Index Terms—*Few-Shot Learning, Semi-Supervised Learning, Self-Supervised Learning, Low-Label Learning, plant disease recognition.*

I. INTRODUCTION

Coconut cultivation is a cornerstone of Sri Lanka's agriculture, economy, and rural livelihoods, with the island's tropical climate and diverse conditions for production. Despite its importance, the sector faces mounting challenges due to the increasing threat of pests like the Rhinoceros Beetle, Red Palm Weevil, and Black-Headed Caterpillar, and Whitefly. Coconut whitefly (*Aleurodicus dispersus* Russell or (*Aleurodicus dispersus*) [1], a rapidly spreading invasive pest, has emerged as a critical threat to Sri Lanka's coconut industry one of the backbone sectors of the rural economy, supporting over 1.5 million livelihoods. Infestations manifest through leaf chlorosis, sooty mold formation, and honeydew excretion, leading to reduced photosynthesis and, in severe cases, premature nut drop and canopy dieback. Field studies report yield losses of up to 60% in heavily infested plantations, posing a significant risk to food security and export revenue [2].

Traditional detection methods rely on manual scouting and visual inspection by agricultural officers, which are labor-intensive, time-consuming, and prone to subjectivity. By the time infestations are identified, they are often widespread, limiting the effectiveness of control measures. Moreover, the lack of scalable monitoring tools hinders early intervention and region-wide surveillance key requirements for sustainable pest management.

Recent advances in drone-based remote sensing and deep

learning offer transformative potential for precision agriculture. Unmanned Aerial Vehicles (UAVs) enable rapid, high-resolution imaging of large plantations, capturing subtle canopy-level changes associated with early infestation [3]. However, a major bottleneck remains: the scarcity of labeled training data required to train robust AI models. Manual annotation of infestation levels (Low, Medium, High) is costly and impractical at scale. To address this challenge, we propose a novel low-label learning framework that combines self-supervised pretraining with semi-supervised fine-tuning for accurate detection of coconut whitefly infestations using minimal labeled data. SimCLR-based [4] self-supervised Contrastive Learning pre-training to learn robust leaf-level features from unlabeled images, FixMatch-based semi-supervised learning to leverage a small labeled subset and a large pool of unlabeled data, Predictive infestation mapping to visualize spatial spread and support decision-making.

We introduce *CocoInfest-SL*, a new four-class dataset of coconut leaf images labeled with infestation severity (Healthy, Mild, Moderate, Severe), collected in 2023–2025 from key coconut-growing regions. To address this challenge, we propose a novel SimCLR → FixMatch framework that combines self-supervised contrastive pretraining with semi-supervised fine-tuning, enabling accurate detection of coconut whitefly infestations using only 30 labeled images per class. Our method achieves 91.8% accuracy and a 0.91 F1-score on the *CocoInfest-SL* dataset, significantly outperforming few-shot and semi-supervised baselines, and offering a scalable, low-label solution for early pest monitoring in resource-constrained agricultural settings.

This work presents a scalable, data-efficient solution for early pest detection in resource-constrained agricultural settings, with direct applicability to smallholder farms across tropical regions. Particularly in the early stages of infestation, where labeled data is scarce, our method enables accurate detection using minimal annotations, offering a scalable solution for real-time pest monitoring in data-limited environments.

II. RELATED WORK

Recent advancements in self-supervised learning have shown great promise in reducing the reliance on large annotated datasets for visual recognition tasks. Methods such as contrastive learning and its variants leverage unlabeled data by learning representations through instance discrimination or feature consistency across augmentations. For example, Dense Contrastive Learning has been applied to remote sensing images using Swin Transformers to capture both global and local features, improving object detection performance without manual labels. Similarly, contrastive self-supervised learning has served as a strong baseline in unsupervised hashing, demonstrating high retrieval accuracy on benchmark datasets by effectively preserving semantic similarity in the learned hash codes [5].

These self-supervised strategies are being increasingly integrated into real-world applications where labeled data is scarce or privacy-sensitive. In medical imaging, the FixMatch

framework has been adapted for Alzheimer’s disease classification, combining pseudo-labeling with consistency regularization to achieve high performance with minimal labeled data [6]. In industrial scenarios, self-supervised methods have been combined with federated learning for object detection in electric power operation sites, ensuring data privacy and efficient model training across distributed clients. Additionally, contrastive learning and adversarial training have been used to detect foreign objects on high-speed railway tracks, enhancing model robustness and detection accuracy [7] [8]. These works collectively highlight the effectiveness of self-supervised learning in bridging the gap between limited supervision and high model performance across diverse domains.

self-supervised and semi-supervised learning have significantly reduced the dependency on large labeled datasets in various domains. In remote sensing, image-text representation learning based on the CLIP model has enabled open-vocabulary object detection by leveraging pretrained vision-language models, allowing detection of diverse object categories with limited semantic information and addressing issues of class imbalance and missing detections. Similarly, dense contrastive learning has been applied to remote sensing images using Swin Transformers as the backbone, where global and local feature representations are enhanced through multi-crop augmentation and parallel projection heads, improving detection performance in a self-supervised manner [9].

In medical imaging, self-supervised pretraining has proven effective in reducing the need for labeled data, with studies showing that models pretrained on unlabeled chest X-rays can achieve performance comparable to fully supervised models using only 60% of the labeled data [10]. This approach leverages the NT-Xent loss to learn meaningful representations before fine-tuning on downstream tasks. Furthermore, the FixMatch framework has been successfully adapted for Alzheimer’s disease classification, demonstrating that combining pseudo-labeling with consistency regularization can yield results on par with supervised baselines under limited labeling. In industrial applications, federated self-supervised learning has been used for object detection in electric power operation sites, where a YOLO-based model trained with only 200 labeled images achieved an mAP50 of 0.83, closing the gap with fully supervised training. [11]–[15]

The reviewed studies demonstrate that self-supervised and semi-supervised learning methods, particularly those based on contrastive learning and consistency regularization, can significantly reduce the reliance on large labeled datasets while maintaining strong performance across diverse domains such as remote sensing, medical diagnosis, and industrial monitoring. These approaches leverage unlabeled data to learn robust feature representations or generate reliable pseudo-labels, effectively bridging the gap between limited supervision and high model accuracy. Building upon these advancements, this work proposes a two-stage framework that integrates contrastive pretraining with semi-supervised fine-tuning for accurate detection of coconut whitefly infestations under extreme label scarcity.

III. METHODOLOGY

Address the challenge of limited labeled data in real-world agricultural settings for the early detection of coconut For whitefly infestations, we propose a two-stage semi-supervised learning framework, as shown in Figure 1, utilizing SimCLR-Enhanced FixMatch for Low-Label infestation detection. First, we pre-train a ResNet-18 backbone using *SimCLR* [16], a contrastive self-supervised learning method, on the entire dataset (labeled and unlabeled images). This stage learns robust visual representations of coconut leaves by maximizing agreement between differently augmented views of the same image, enabling the model to capture subtle patterns associated with whitefly infestations, such as chlorosis, sooty mold, and honeydew residue. In the second stage, we fine-tune




Fig. 1. Proposed system architecture for Semi-supervised and few-shot learning methods on the CocoInfest-SL

the pretrained network using *FixMatch* [12], a state-of-the-art semi-supervised learning algorithm. FixMatch leverages both the small labeled subset and the large unlabeled pool by generating high-confidence pseudo-labels on weakly augmented unlabeled images and enforcing consistency under strong augmentations.

A. Our dataset

The developed dataset consists of 1,600 real-world coconut leaf images collected from major coconut-growing regions in Sri Lanka, including the North Western, Western, and Southern provinces. Images were captured using smartphones during field visits in 2023–2025 and labeled by the CVAAT annotation tool into four categories: Healthy, Mild, Moderate, and Severe infestation, based on visible symptoms such as chlorosis, honeydew, and sooty mold. To reflect real-world conditions where labeled data is limited, only 120 images (30 per class) were used as labeled samples, while the remaining 1,480 were left unlabeled. From the 120 labeled images, we performed a stratified train-test split to ensure class balance in both training and testing sets. Specifically, 90 images (22–23 per class, ~75%) were used for training, 30 images (7–8 per class, ~25%) were reserved for testing. All images were resized to 224×224 pixels and normalized for model

training. This dataset, named CocoInfest-SL as shown in Table I, supports low-label learning and is designed to enable early detection of coconut whitefly infestations in resource-constrained environments, as shown in Figure 2.




Fig. 2. Collected sample of CocoInfest-SL Healthy (A), Mild (B), Moderate (C), Severe (D)

TABLE I
DISTRIBUTION OF LABELED AND UNLABELED IMAGES ACROSS INFESTATION CLASSES

Class	Total Images	Labeled Images	Unlabeled Images
Healthy	400	30	370
Mild	400	30	370
Moderate	400	30	370
Severe	400	30	370
Total	1,600	120	1,480

B. Implementation Details

All models were trained using Colab. **SimCLR Pretraining:** We trained for 400 epochs with a batch size of 256 using the LARS optimizer and a learning rate of 0.3. Data augmentations included random cropping, color jittering, and Gaussian blur. The NT-Xent loss was used with a temperature of 0.5.

FixMatch Fine-tuning: We trained for 500 epochs with a batch size of 64, sampling 8 labeled and 56 unlabeled images per batch (1:7 ratio). We used SGD with momentum 0.9, an initial learning rate of 0.03, and a cosine annealing learning rate scheduler. Pseudo-labels were generated using a confidence threshold of 0.95, and strong RandAugment was applied to the unlabeled data. The unsupervised loss weight was ramped up as:

$$\lambda(t) = \min \left(1.0, 0.1 \times \exp \left(\frac{t}{50} \right) \right) \quad (1)$$

where t is the current training epoch in equation 3.

The ResNet-18 backbone was initialized from scratch (i.e., no ImageNet pretraining) and all input images were resized to 224×224 pixels.

C. Architecture Description

The proposed SimCLR-Enhanced FixMatch framework (Algorithm 1) follows a two-stage, label-efficient pipeline designed for accurate detection of coconut whitefly infestations under extreme label scarcity. In the first stage, a ResNet-18 backbone is pretrained using SimCLR (Simple Framework for Contrastive Learning of Representations) on the entire dataset (labeled and unlabeled images) in a self-supervised manner. Each image is augmented twice using random cropping, color jittering, and Gaussian blur to generate two different views. The encoder extracts features from both views, which are then passed through a small MLP projection head. The model is trained to minimize the NT-Xent (Normalized Temperature-scaled Cross-Entropy) loss, encouraging similar embeddings for the two views of the same image and dissimilar embeddings for different images. This pretraining step enables the network to learn robust, semantically meaningful representations of coconut leaf textures, discoloration patterns, and early infestation signs without using any labels.

In the second stage, the pretrained ResNet-18 (without the projection head) is used as the backbone for FixMatch, a consistency-based semi-supervised learning algorithm. During training, labeled images are used to compute the supervised cross-entropy loss, while unlabeled images are processed as follows: a weakly augmented version is passed through the model to generate a pseudo-label (the class with the highest prediction confidence). If the confidence exceeds a threshold $T=0.95$, this pseudo-label supervises the model's prediction on a strongly augmented version of the same image.

IV. RESULTS AND DISCUSSION

The evaluation of various semi-supervised and few-shot learning methods on the CocoInfest-SL dataset highlights the challenges of accurate plant infestation classification with limited labeled data. Accurate detection of coconut whitefly infestations under real-world agricultural constraints, particularly limited labeled data, is a significant challenge. To address this, we propose a two-stage semi-supervised learning framework, SimCLR-Enhanced FixMatch, designed to leverage a small labeled dataset (120 images) and a large pool of unlabeled images (1,480) effectively. As detailed in Algorithm 1, our method first pretrains a feature encoder using SimCLR on both labeled and unlabeled data, learning robust representations through contrastive augmentation. This pretrained encoder is then fine-tuned using FixMatch, which generates pseudo-labels for unlabeled data based on high-confidence predictions and enforces consistency under strong augmentations.

By combining self-supervised pretraining with consistency regularization, our approach maximizes the utility of scarce labeled examples while exploiting the structure of abundant unlabeled data. The following results demonstrate that

Algorithm 1: SimCLR-Enhanced FixMatch for Few-Shot Coconut Whitefly Detection

Input: Small labeled dataset $\mathcal{D}_l$, large unlabeled dataset $\mathcal{D}_u$, batch size N , temperature τ , augmentation functions $\mathcal{T}$, threshold θ

Phase 1: SimCLR Pretraining

```

for sampled minibatch  $\{x_k\}_{k=1}^N$  from  $\mathcal{D}_l \cup \mathcal{D}_u$  do
  Draw two augmentation functions  $t, t' \sim \mathcal{T}$ 
  for  $k = 1$  to  $N$  do
     $\tilde{x}_{2k-1} = t(x_k)$ ; // First augmentation
     $h_{2k-1} = f(\tilde{x}_{2k-1})$ ; // Encoder representation
     $z_{2k-1} = g(h_{2k-1})$ ; // Projection
     $\tilde{x}_{2k} = t'(x_k)$ ; // Second augmentation
     $h_{2k} = f(\tilde{x}_{2k})$ ; // Encoder representation
     $z_{2k} = g(h_{2k})$ ; // Projection
  for  $i = 1$  to  $2N$  do
    for  $j = 1$  to  $2N$  do
       $s_{i,j} = \frac{z_i \cdot z_j}{\|z_i\| \|z_j\|}$ ; // Pairwise similarity
    Define  $\ell(i, j) = -\log \frac{\exp(s_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(s_{i,k}/\tau)}$ 
   $\mathcal{L}_{\text{SimCLR}} = \frac{1}{2N} \sum_{k=1}^N [\ell(2k-1, 2k) + \ell(2k, 2k-1)]$ 
  Update  $f$  and  $g$  to minimize  $\mathcal{L}_{\text{SimCLR}}$ 

```

Discard g , keep pretrained encoder f

Phase 2: FixMatch Semi-Supervised Training

```

for minibatch from  $\mathcal{D}_l$  and  $\mathcal{D}_u$  do
  Apply weak augmentation  $a_w$  and strong augmentation  $a_s$ 
   $p = \text{softmax}(f(a_w(x_u)))$ ; // Prediction on unlabeled data
  if  $\max(p) > \theta$  then
     $\hat{y} = \arg \max(p)$ ; // Pseudo-label
     $\mathcal{L}_u = \ell(f(a_s(x_u)), \hat{y})$ ; // Consistency loss
   $\mathcal{L}_s = \ell(f(a_s(x_l)), y_l)$ ; // Supervised loss
   $\mathcal{L} = \mathcal{L}_s + \lambda_u \mathcal{L}_u$ 
  Update  $f$  to minimize  $\mathcal{L}$ 

```

Output: Trained model for *Healthy/Infested* classification with severity levels

this strategy significantly improves classification performance across all evaluation metrics compared to existing few-shot and semi-supervised baselines. Table II presents a comparative evaluation of different methods. The fully supervised model trained on full labels serves as the upper bound (95.2% accuracy), while ResNet-18 trained on few labels achieves only 76.4%, highlighting the challenge of label scarcity. Prototypical Networks (5-shot) perform slightly worse (72.1%), indicating limitations in generalization under extreme data constraints. SimCLR with linear probing improves performance to

TABLE II
COCOINFEST-SL TEST-SET RESULTS UNDER LIMITED ANNOTATIONS

Method	Acc (%)	Prec	Rec	F1
Supervised (Full Labels)	95.2	0.95	0.95	0.95
ResNet-18 (Labeled 7.5% per class)	76.4	0.75	0.76	0.74
Prototypical Networks (5-shot)	72.1	0.70	0.71	0.69
SimCLR → Linear	84.6	0.83	0.84	0.83
FixMatch (w/o pretraining)	87.3	0.86	0.87	0.86
Ours (SimCLR → FixMatch)	91.8	0.91	0.92	0.91

84.6%, demonstrating the value of self-supervised pretraining. FixMatch without pretraining reaches 87.3%, showing the effectiveness of consistency training. Notably, our SimCLR → FixMatch approach achieves 91.8% accuracy and 0.91 F1-score, closing the gap with fully supervised performance by nearly 3.4 percentage points compared to FixMatch alone.

The proposed SimCLR → FixMatch framework is built upon a two-stage learning objective, where self-supervised pretraining is followed by semi-supervised fine-tuning. The final stage of training is governed by a combined loss function that effectively leverages both labeled and unlabeled data:

$$\mathcal{L} = \mathcal{L}_{\text{sup}}(\theta) + \lambda(t) \cdot \mathcal{L}_{\text{unsup}}(\theta) \quad (2)$$

where $\mathcal{L}_{\text{sup}}(\theta)$ is the supervised cross-entropy loss on the labeled data, and $\mathcal{L}_{\text{unsup}}(\theta)$ is the consistency loss computed on unlabeled data using pseudo-labels based on equation 2. The time-varying coefficient $\lambda(t)$ controls the contribution of the unsupervised loss during training. In this work, we define it as:

$$\lambda(t) = \min(1.0, 0.1 \times \exp(t/50)) \quad (3)$$

where t is the current training epoch. This exponential equation 3 ramp-up schedule ensures that the model initially focuses on learning from the small labeled set, minimizing the risk of error propagation from low-confidence pseudo-labels. As training progresses, $\lambda(t)$ increases, gradually incorporating high-confidence predictions on unlabeled data into the optimization process. This dynamic weighting balances supervised and unsupervised learning, enabling stable convergence and robust generalization in low-label regimes.

To visualize the impact of this dynamic weighting strategy, Figure 3 shows the supervised loss decreases rapidly as the model learns from the limited labeled data, reflecting fast convergence on known infestation patterns. The raw unsupervised loss decreases more gradually, indicating steady improvement in feature learning from unlabeled images via contrastive and consistency training. The weighted unsupervised loss ($\lambda(t) \cdot \mathcal{L}_{\text{unsup}}$) starts near zero and increases over time as $\lambda(t)$ ramps up, gradually incorporating high-confidence pseudo-labels into training. The total loss combines both labeled and unlabeled signals, optimizing for accurate detection while leveraging abundant unlabeled data for robust feature representation. The Figure 4 weighting factor $\lambda(t)$ follows a gradual ramp-up schedule defined as $\lambda(t) = \min(1.0, 0.1 \times \exp(t/50))$, where t is the training epoch. This strategy ensures stable and effective




Fig. 3. Training dynamics of loss components in the SimCLR → FixMatch framework




Fig. 4. Scaling Factor Evolution

semi-supervised learning by controlling the contribution of the unsupervised loss over time. In the early stages (epochs 0–50), $\lambda(t)$ remains small, prioritizing learning from the limited labeled data and preventing the model from being misled by low-confidence pseudo-labels. During mid-training (epochs 50–100), $\lambda(t)$ increases exponentially, gradually incorporating high-confidence pseudo-labels from the unlabeled data and encouraging feature consistency under strong augmentations. In the late stages (100+ epochs), $\lambda(t)$ approaches its maximum value, achieving a balanced optimization between supervised and unsupervised signals. This dynamic weighting enables robust feature extraction and improves generalization, particularly for distinguishing subtle visual patterns in early-stage coconut whitefly infestations.

A. Performance Analysis

The proposed SimCLR → FixMatch framework achieves 91.8% accuracy in detecting coconut whitefly infestations using only 30 labeled images per class, demonstrating high data efficiency and strong generalization. The model rapidly learns diagnostic features such as chlorosis, honeydew, and sooty mold from limited labeled data, then refines its representations using a large pool of unlabeled images. This synergistic learning strategy enables accurate early detection in real-world agricultural settings, such as smallholder coconut farms in Sri

Lanka, where labeled data is scarce and timely intervention is critical.




Fig. 5. Detection accuracy of the SimCLR → FixMatch framework across training epochs on the CocoInfest-SL test set

V. CONCLUSION

In this study, we have presented a highly effective and data-efficient framework, SimCLR → FixMatch, for the early detection of coconut whitefly (*Aleurodicus dispersus*) infestations in Sri Lanka, where labeled data is extremely scarce and timely intervention is critical for preserving livelihoods. To address the challenge of limited annotations, we introduced *CocoInfest-SL*, a new dataset of 1,600 real-world coconut leaf images labeled with four severity levels (Healthy, Mild, Moderate, Severe), with only 120 images manually annotated (30 per class). Our two-stage approach first leverages SimCLR-based self-supervised learning to pretrain a robust feature encoder on both labeled and unlabeled data, capturing subtle visual patterns such as chlorosis, honeydew, and sooty mold. In the second stage, FixMatch fine-tunes the encoder using a consistency-based semi-supervised strategy, generating high-confidence pseudo-labels on weakly augmented views and enforcing prediction consistency under strong augmentations.

The total loss function, defined as $\mathcal{L} = \mathcal{L}_{\text{sup}}(\theta) + \lambda(t) \cdot \mathcal{L}_{\text{unsup}}(\theta)$, dynamically balances supervised and unsupervised learning through a time-ramped weighting factor $\lambda(t)$, preventing error propagation in early training stages. Our method achieves a classification accuracy of 91.8% and an F1-score of 0.91, significantly outperforming baseline approaches such as ResNet-18 with few labels (76.4%), Prototypical Networks (72.1%), and FixMatch without pretraining (87.3%). This performance demonstrates that self-supervised pretraining substantially enhances the model's ability to generalize in low-label regimes.

The proposed framework is particularly suited for deployment in resource-constrained agricultural environments, where expert labeling is impractical. By reducing the need for large annotated datasets, our approach enables scalable, real-time monitoring of coconut plantations using smartphone or drone-captured imagery. Future work will focus on extending this model to multispectral drone data for plantation-wide infestation mapping and integrating it into a mobile application

for farmer-facing diagnostics. This work paves the way for intelligent, low-cost pest management systems across tropical agriculture.

REFERENCES

- [1] K. Elango, K. P. Elango, N. A. kumar, N. Kumar, V. Dhanasekar, and V. Dhanasekar, "Ovicidal and insecticidal effects of microbial pathogens against invasive coconut rugose spiralling whitefly (rsw), aleurodicus rugioperculatus martin and impact on their chrysopid predator mallada boninensis (okamoto)," *Egyptian Journal of Biological Pest Control*, 2022.
- [2] Parliament of Sri Lanka, "Agricultural challenges in sri lanka," tech. rep., Parliament of Sri Lanka, December 2023. Accessed: 2024-04-05.
- [3] N. Thayalan, J. Pratheepan, and S. Thayalan, "Exploring the advances in computer applications with the internet of things (iot) for optimized agricultural farm management," *International Journal of Scientific and Research Publications (IJSRP)*, vol. 12, no. 6, 2022. ISSN: 2250-3153.
- [4] S. Kong, "Self-supervised image classification using convolutional neural network," in *2023 IEEE 3rd International Conference on Power, Electronics and Computer Applications (ICPECA)*, pp. 1283–1290, 2023.
- [5] H.-F. Yang, "Contrastive self-supervised learning as a strong baseline for unsupervised hashing," in *2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP)*, pp. 1–5, 2022.
- [6] C.-C. Wang, S. Xu, J. Fu, Y. Liu, and B. Wang, "Kd-fixmatch: Knowledge distillation siamese neural networks," in *2023 IEEE International Conference on Image Processing (ICIP)*, pp. 341–345, 2023.
- [7] Y. Fan, P. Dai, H. Song, and H. Li, "Detection of high-speed railway track bed foreign object based on contrastive learning and adversarial training," in *2025 IEEE 5th International Conference on Power, Electronics and Computer Applications (ICPECA)*, pp. 73–77, 2025.
- [8] R. Muralitharan, N. Inuvaisubramaniam, U. Jayasinghe, R. G. Ragel, A. Rajeevan, and G. Niladhari, "Plant leaves classification system," in *2023 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, pp. 1–6, 2023.
- [9] S. Manivannan, R. Wang, and E. Trucco, "Hierarchical mix-pooling and its applications to biomedical image classification," in *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*, pp. 541–544, 2016.
- [10] S. H. P. Abeyagunasekera, Y. Perera, K. Chamara, U. Kaushalya, P. Sumathipala, and O. Senaweera, "Lisa : Enhance the explainability of medical images unifying current xai techniques," in *2022 IEEE 7th International conference for Convergence in Technology (I2CT)*, pp. 1–9, 2022.
- [11] K. Thiruthanigesan, R. D. Nawarathna, and R. G. Ragel, "Advancing suspicious object detection in unconstrained environments using computer vision," *Applied Intelligence*, vol. 55, p. 871, 2025. Accepted: 23 June 2025, Published: 15 July 2025.
- [12] Y. Zhong, F. Wang, C. Wang, and B. Han, "Pixelfixmatch: A semi-supervised image segmentation method based on fixmatch with pixel attention," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, 2024.
- [13] K. Thiruthanigesan, R. D. Nawarathna, and R. G. Ragel, "Enhanced yolov4 for facilitating public safety management amidst protests and riots," in *2023 IEEE 17th International Conference on Industrial and Information Systems (ICIIS)*, pp. 1–6, 2023.
- [14] K. Thiruthanigesan, R. D. Nawarathna, and R. G. Ragel, "Suspicious object detection in environments with tear gas," in *2021 IEEE 16th International Conference on Industrial and Information Systems (ICIIS)*, pp. 185–190, 2021.
- [15] K. Thiruthanigesan, R. D. Nawarathna, and R. G. Ragel, "A systematic approach for object detection in unconstrained environments," in *2023 IEEE 17th International Conference on Industrial and Information Systems (ICIIS)*, pp. 215–220, 2023.
- [16] H. Sheng, J. Zhang, L. Wan, Z. Huang, P. Wu, and Q. Wang, "Incremental learning for object detection based on self-supervised learning," in *2024 IEEE International Conference on Progress in Informatics and Computing (PIC)*, pp. 249–253, 2024.