Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

An Integrated Deep Learning and Sparse Representation Framework for Apple Leaf Disease Recognition

Rahul Neware

rneware00@gmail.com

Leibniz Centre for Agricultural Landscape Research

Shashi Dahiya

ICAR-Indian Agricultural Statistics Research Institute

Sudeep Marwaha

ICAR-Indian Agricultural Statistics Research Institute

Article

Keywords: Apple leaf disease detection, deep learning, YOLOv8, sparse classification, L1 regularization, feature selection, computer vision, interpretability, smart farming

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11331940/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

An Integrated Deep Learning and Sparse Representation Framework for Apple Leaf Disease Recognition

Rahul Neware[1,2*] and Shashi Dahiya[2] and Sudeep Marwaha[2]

[1]Leibniz Centre for Agricultural Landscape Research (ZALF), Müncheberg, Germany

[2]Indian Agricultural Statistics Research Institute, New Delhi, India

Correspondence: Rahul Neware (R.N.) Email: Rahul.neware@zalf.de

Abstract

Proper identification of apple leaf diseases is an important issue towards safeguarding crop production as well as the management sustainability of orchards. Manual diagnosis is widely used but is not efficient, is subjective, and cannot easily support large-scale monitoring. This paper presents a Hybrid Deep Learning and Sparse Classification Framework, which combines YOLOv8-based deep feature extraction with L1-regularized machine learning models to achieve efficient and interpretable disease recognition. A dataset of 4,006 high-resolution apple leaf images covering 9 classes Alternaria leaf spot, Brown spot, Frogeye leaf spot, gray spot, Healthy, Mosaic, Powdery mildew, Rust, and Scab was used. After 60 epochs, YOLOv8 attained a mean Average Precision (mAP) of 98.10% at IoU 0.5, a precision of 95.16%, and a recall of 95.97%. The resulting 3,648-dimensional feature space was narrowed down to 319 dimensions using Lasso feature selection, achieving 13.48% sparsity without losing important discriminatory information. The Lasso model achieved the best test accuracy of 97.50% and the highest cross-validation accuracy of 98.78%, outperforming the L1-regularized Logistic Regression and Random Forest baselines among the compared classifiers. Sparse regularization not only reduced overfitting but also improved model interpretability by forcing the model to focus on the visual patterns of the diseases. The proposed model is a computationally efficient, scalable, and interpretable AI solution for real-time agricultural diagnostics. It offers a promising direction for precision farming systems that can detect diseases early, protect yields, and support better decision-making based on interpretable hybrid intelligence.

Keywords: Apple leaf disease detection; deep learning; YOLOv8; sparse classification; L1 regularization; feature selection; computer vision; interpretability; smart farming.

1. Introduction

Agriculture is an excellent foundation of world sustainability, as it secures food security, offers employment, and supplies raw materials to many industries. Among fruit commodities, the apple (*Malus domestica*) is one of the most valuable and widely cultivated crops globally, and its production carries substantial economic weight: apples are produced across a broad range of climatic zones and their trade contributes significantly to the economies of major producing countries (O'Rourke, 2021). Because apples are also prized for their nutritional and therapeutic value, safeguarding their production matters not only commercially but also nutritionally (Yatoo, 2024). Nevertheless, apple cultivation is constantly threatened by foliar diseases Alternaria leaf blotch, caused by *Alternaria alternata* f. sp. *mali*, and apple rust being prominent examples that harm fruit quality and cause significant financial losses for growers (Yatoo, 2024). Diseases such as Alternaria leaf spot, Powdery mildew, Frogeye leaf spot, Rust, and Scab similarly constrain production and degrade the quality of the harvested fruit, so early infection monitoring and timely control measures are key to ensuring regular growth of apple fruits and a high-efficiency economy (Gong & Zhang, 2023).

Conventional disease detection methods rely on manual visual analysis, which is labour-intensive, subjective, and inefficient in the context of large-scale monitoring (Rezaei et al., 2024). Traditional approaches to plant disease diagnosis are usually based on visual inspection and laboratory testing, which can be expensive and time-consuming and require trained plant pathologists as well as specialised equipment (Rezaei et al., 2024). In practice, farmers find it difficult to distinguish between different apple leaf diseases because their symptoms are quite similar and may appear simultaneously on the same leaf (Doutoum & Tuğrul, 2025). In a large orchard, an entire tree carries a huge number of leaves, and the scarcity of trained experts makes exhaustive inspection infeasible; misclassification rises when professionals are fatigued, particularly for visually similar leaf conditions. Field imagery adds further difficulty, since images captured outdoors contain a large amount of background and noise that submerges disease-spot features, degrading detection and recognition performance (Li et al., 2023). Traditional image-processing methods colour segmentation, morphological filtering, and texture analysis were used to automate plant disease detection and are fairly effective in controlled settings, but they fail under the noise, occlusion, and illumination variability of real-world agricultural conditions (Aravind & Raja, 2020; Praveen et al., 2022).

Recent developments in machine vision and deep learning have transformed agricultural diagnostic automation. Deep learning models automatically identify hierarchical features of leaf images and acquire the discriminative patterns necessary for disease identification (Vishnoi et al., 2022), and this approach has proven feasible at scale: a landmark study demonstrated 99.35% accuracy in classifying 26 diseases across 14 crop species from a public dataset of 54,306 plant leaf images, pointing towards smartphone-assisted diagnosis on a global scale (Mohanty et al., 2016). Comprehensive reviews of computer vision for plant disease detection analyse more than 278 research articles and show that DL-based image analysis ranging from convolutional neural networks to vision transformers, generative adversarial networks, and vision language models has become the dominant and most accurate paradigm in precision agriculture (Upadhyay et al., 2025). Pre-trained CNN architectures such as AlexNet, VGG, and ResNet have proven effective in this domain, improving accuracy while lessening the burden of manual feature engineering (Gangadevi & Jayakumar, 2022), and transfer learning has further improved model performance, enabling efficient training even with smaller datasets (Aravind & Raja, 2020; Yan et al., 2020). Despite these advances, DL-based approaches generally remain data-hungry: they require large annotated datasets and high computational resources, which limits their application in real-time, resource-constrained agricultural settings (Rezaei et al., 2024; Yang et al., 2022). Standard CNN models in particular require a large number of parameters and higher computation cost, making deployment on portable terminals difficult (Hassan et al., 2021).

A critical research issue is therefore balancing accuracy, interpretability, and computational efficiency. Dense feature representations produced by deep networks are largely redundant and not easily interpretable, and one way to address this is sparse learning and pruning, which concentrate on the most important features and enhance both the efficiency and the clarity of the resulting models (Poyatos et al., 2022). Pruning methods reduce the complexity of the network and can even improve performance by identifying and removing unnecessary connections among neurons (Poyatos et al., 2022); evolutionary filter-pruning approaches achieve large reductions in computational complexity across CNNs, residual networks, and densely connected networks (Fernandes & Yen, 2020). At the same time, the combination of computer-vision systems with machine learning has demonstrated clear prospects for real-time agricultural phenotyping and disease localization, with machine vision technologies increasingly used for stress and disease detection in precision agriculture (Scharr et al., 2017; Shin et al., 2022).

Within the apple domain specifically, object detectors of the YOLO family have become a focus of research because they provide end-to-end, high-speed, real-time detection (Xu & Wang, 2023), and successive lightweight YOLO variants have achieved simultaneously high accuracy and very high frame rates in agricultural scenes (Wang et al., 2025; Zhu et al., 2023). However, most existing systems are validated only in laboratory settings or rely on heavy models, and the literature has not systematically explored hybrid strategies that combine deep feature extraction with explainable, sparse classifiers for real-time deployment on resource-constrained devices.

In order to overcome these weaknesses, this paper presents a Hybrid Deep Learning and Sparse Classification Framework for automated apple leaf disease recognition. The study is novel in its combination of YOLOv8-based deep feature extraction with L1-regularized sparse classification to achieve a unique trade-off between accuracy, interpretability, and computational efficiency an uncommon combination in precision agriculture. The main contributions of this work are:

1. **A hybrid detection–classification pipeline:** YOLOv8, trained for 60 epochs, is used as a deep feature extractor rather than purely as a detector, yielding a 3,648-dimensional feature representation per leaf image.
2. **L1-based dimensionality reduction:** An L1 regularization stage compresses the feature space from 3,648 to 319 dimensions (276 non-zero coefficients, 13.48% sparsity), retaining only the most discriminative information.
3. **Systematic classifier comparison:** Three classifiers L1-regularized Logistic Regression, Random Forest, and the Lasso feature-selection classifier are compared under five-fold cross-validation, with the Lasso model achieving 98.78% cross-validation accuracy and 97.50% test accuracy.
4. **Deployment-oriented efficiency:** An inference time of 7.8 milliseconds per image makes the framework suitable for edge and mobile devices in real-time orchard monitoring.

The suggested method is efficient because the L1 regularization stage reduces the dimensionality and enables the model to select only the most significant features, thereby increasing speed and transparency. The design generalizes robustly to limited, imbalanced, and noisy agricultural data, providing scalability to practical field conditions. Moreover, its lightweight character allows deployment on edge devices and mobile platforms for real-time use, which makes it well suited to precision farming, where efficient and explainable decision-making is necessary. The rest of the paper is organized as follows: Section 2 reviews existing work on deep and sparse learning for plant disease identification; Section 3 describes the dataset and the proposed methodology; Section 4 presents and discusses the experimental results; and Section 5 summarizes the main findings and future perspectives.

2. Related Work

Automated plant disease detection has developed tremendously, from classical methods to lightweight, real-time deep learning systems. This section reviews the literature along four lines: classical image processing and shallow machine learning, deep CNN-based classification, apple-specific detection architectures, and sparse feature learning.

2.1 Classical image processing and shallow machine learning

Traditional image processing and Support Vector Machines were typically utilized for simple classification tasks. Bhimte and Thool combined image processing with an SVM classifier to detect cotton leaf spot, using colour-based segmentation followed by colour and texture feature

extraction (Bhimte & Thool, 2018). In the apple domain, Chuanlei et al. designed a colour transformation pipeline converting RGB images to HSI, YUV, and grey models removed the background by thresholding, segmented disease spots with a region-growing algorithm, extracted 38 colour, texture, and shape features per spot, and then selected the most valuable features by combining a genetic algorithm with correlation-based feature selection before classification with an SVM, achieving a correct identification rate of more than 90% on a database of 90 apple leaf images covering powdery mildew, mosaic, and rust (Zhang et al., 2017). Similar optimization-driven pipelines appeared for other crops: Khan et al. fused colour, colour-histogram, and local-binary-pattern features with genetic-algorithm selection for apple disease classification using a one-vs-all multiclass SVM (Khan et al., 2019), and Jose and Santhi combined an improvised flower-pollination genetic algorithm with a multiclass SVM and spectral vegetation indices for early apple leaf disease detection, benchmarking against Faster R-CNN, R-SSD, and INAR-SSD (Jose & Santhi, 2022). Beyond apples, machine learning techniques were applied to rice leaf disease detection (Pallathadka et al., 2021), an integrated digital image analysis system was developed for wheat leaves (Majumdar et al., 2015), and machine learning algorithms were applied to generic leaf disease identification (Praveen et al., 2022). These works collectively illustrate the diversity of crops and methods used in early automated diagnosis, but they share a common limitation: handcrafted features are difficult to design for the highly variable appearance of disease symptoms in natural scenes.

2.2 Deep CNN-based plant disease classification

Deep learning shifted the paradigm from handcrafted features to learned representations. Mohanty et al. trained a deep CNN on 54,306 images of diseased and healthy plant leaves, achieving 99.35% accuracy in identifying 14 crop species and 26 diseases, and argued that large public image datasets present a clear path toward smartphone-assisted diagnosis (Mohanty et al., 2016). Adedoja et al. proposed a NASNet-based transfer-learning model for plant leaf disease recognition on the PlantVillage dataset, reaching 93.82% accuracy (Adedoja et al., 2019). Argüeso et al. showed that a few-shot learning approach can classify plant diseases from images taken in the field, a data-scarce scenario (Argüeso et al., 2020), and follow-up studies confirmed that modern few-shot pipelines such as pre-training, meta-learning, and fine-tuning with a feature-attention module (PMF+FA) can reach 90.12% accuracy on the field-environment PlantDoc dataset using as few as five images per class, at 1.11 ms per image with a Vision Transformer (Rezaei et al., 2024). Reviews of few-shot learning in smart agriculture likewise report high recognition accuracy on limited datasets up to 99.48% in vegetable disease recognition (Ragu & Teo, 2023) while noting that most deep learning methods need massive numbers of samples (Sun et al., 2023; Yang et al., 2022). Atole and Park designed a multiclass deep CNN to identify common rice plant anomalies (Atole & Park, 2018), and Yao et al. proposed the generalised stacking multi-output CNN, which outperformed existing approaches on the PlantVillage, Plant Leaves, and PlantDoc benchmarks, with InceptionV3 as the most effective backbone (Yao et al., 2024). Lightweight architectures have also been emphasised: replacing standard convolutions with depthwise separable convolutions reduced parameters and computation cost while achieving accuracy rates between 98.42% and 99.56% across InceptionV3, InceptionResNetV2, MobileNetV2, and EfficientNetB0 on a 38-class plant disease dataset (Hassan et al., 2021).

Dedicated apple disease datasets were also constructed during this period. Jiang et al. built the apple leaf disease dataset, combining laboratory images and complex images under real field conditions, totalling 26,377 images of diseased apple leaves covering Alternaria leaf spot, Brown spot, Mosaic, Grey spot, and Rust, and used it to train the INAR-SSD model, which reached 78.80% mAP at a detection speed of 23.13 FPS (Jiang et al., 2019). Liu et al. proposed an AlexNet-based deep CNN trained on 13,689 diseased apple leaf images that identified

Mosaic, Rust, Brown spot, and Alternaria leaf spot with 97.62% accuracy, reducing parameters by 51,206,928 relative to standard AlexNet and gaining a further 10.83% of accuracy from generated pathological images (Liu et al., 2017). Vishnoi et al. developed a compact CNN (Conv-3 DCNN) for apple disease detection using the PlantVillage apple subset, achieving 98% classification accuracy with lower storage and execution time than several existing deep CNN models, making it suitable for deployment in handheld devices (Vishnoi et al., 2022). Improved classification models followed: Yan et al. modified VGG16 by replacing the fully connected layers with global average pooling and adding batch normalization, reaching 99.01% accuracy with an 89% reduction in parameters, a 6.3% accuracy improvement, and a training time reduced to 0.56% of the original model (Yan et al., 2020); Chao et al. combined DenseNet and Xception with global average pooling and an SVM head, achieving 98.82% accuracy across five apple tree leaf diseases (Chao et al., 2020); and Khan et al. reported around 97.18% accuracy for a transfer-learning CNN on their own apple disease dataset (Khan et al., 2021).

2.3 Apple-specific detection: the Faster R-CNN family

A number of papers have developed deep learning solutions specifically for apple diseases using two-stage detectors. Khan et al. developed "Deep Diagnosis," a real-time apple leaf disease detection system based on deep learning, capable of automatically identifying different apple leaf diseases with high efficiency (Khan et al., 2022). Stronger results were reported for Faster R-CNN variants: Gong and Zhang proposed an improved Faster R-CNN that introduces Res2Net and a feature pyramid network as the feature extraction backbone, replaces RoIPool with RoIAlign, and applies soft non-maximum suppression, achieving 63.1% average precision on an annotated apple leaf disease dataset (Gong & Zhang, 2023). Hou et al. combined feature pyramid networks with an Inception squeeze-and-excitation ResNet backbone (FPN–ISResNet–Faster RCNN), reporting AP, AP50, and AP75 values of 62.71%, 93.68%, and 70.94%, respectively higher than those of SSD, VGG–Faster RCNN, GoogLeNet–Faster RCNN, ResNet50–Faster RCNN, ResNeXt–Faster RCNN, and ISResNet–Faster RCNN (Hou et al., 2023). Beyond the apple domain, federated-learning-enhanced Faster R-CNN has been applied to multi-pest detection in orchards, achieving an average accuracy of 90.27% with a detection time of 0.05 seconds per image (Deng et al., 2022). These two-stage detectors are accurate but comparatively slow and heavy, which motivated the turn towards one-stage lightweight detectors.

2.4 YOLO-family and lightweight real-time detectors

YOLO-based models have become popular because of their high speed and efficiency. Xu and Wang proposed ALAD-YOLO, a lightweight and accurate detector for apple leaf diseases (Xu & Wang, 2023); reported variants with multiscale feature fusion and attention raised apple leaf disease detection mAP to 83.19% at 27.53 FPS. Yang et al. modified YOLOv5 to detect and localize objects in agricultural scenes (Yang et al., 2023), and Chen et al. introduced a multiscale, lightweight YOLOv7 model applied to citrus orchards, showing that such architectures transfer across plant species (Chen et al., 2022). Subsequent lightweight detectors pushed both accuracy and speed: EADD-YOLO, an improved YOLOv5 with coordinate attention and SIoU loss, achieved 95.5% mAP at 625 FPS on the ALDD (Zhu et al., 2023); YOLOX-ASSANano, which improves the YOLOX-Nano backbone with an asymmetric ShuffleBlock, a CSP-SA module, and blueprint-separable convolution, achieved 91.08% mAP on a multi-scene apple leaf disease dataset and 58.85% mAP on the public PlantDoc dataset, with only 0.83 MB of parameters and 122 FPS (Liu et al., 2022); and a deep-learning object detection model for multi-class apple disease in real orchard conditions reached 91.2% mAP with an F1-score of 95.9% at 56.9 FPS (Roy & Bhaduri, 2021). The most recent refinements of YOLOv8 itself include: LightYOLO-AppleLeafDx, which integrates Slim-Neck, SPD-

Conv, and SAHead modules to reach a precision of 0.930, mAP@0.5 of 0.965, and mAP@0.5:0.95 of 0.587 with a model size of only 5.2 MB at 107.2 FPS, remaining usable at 14.8 FPS when deployed on an RV1103 hardware platform via an NPU-compatible framework (Zou et al., 2025); ALD-YOLO, which redesigns YOLOv8's C2f modules with FasterNet blocks and an EMA attention module, increasing mAP by 1.4% and 0.6% over YOLOv8n and YOLOv8s on the AppleLeaf9 dataset while reducing GFLOPs by 29.63% and 79.93%, respectively (Deng et al., 2025); and ELM-YOLOv8n, which integrates the Fasternet Block into C2f together with an EMA attention module and an NWD loss, achieving 96.7% mAP50 and an F1 value of 94.0% with parameter count and computational load reduced by 44.8% and 39.5% (Wang et al., 2025). A further lightweight YOLOv8n variant using GhostConv, C3Ghost, a global attention mechanism, and an improved BiFPN reduced complexity by 32.9% and model size to 3.8 M while improving mAP@0.5 to 86.9% (Gao et al., 2024). Similar trends are observed in other fruit crops: the Faster DR-IACNN detector achieved 81.1% mAP on the grape leaf disease dataset GLDD (Xie et al., 2020). On the deployment side, Rahman et al. combined several public datasets into 30,945 images across eight plant types and built a web/mobile application based on the best-performing CNN models, with perfect apple disease detection using MobileNet and Xception (Rahman et al., 2025).

2.5 Feature selection, sparsity, and interpretability

A recurring weakness of deep models is their black-box nature and the redundancy of dense feature representations. Sparse learning and structured pruning have emerged as remedies. Poyatos et al. proposed EvoPruneDeepTL, an evolutionary pruning model that replaces the last fully-connected layers of transfer-learning networks with sparse layers optimized by a genetic algorithm performing either optimized pruning or feature selection and showed that the resulting models gain computational efficiency (Poyatos et al., 2022); their multi-objective extension demonstrated that the most influential neurons help explain which parts of the input images are most relevant for prediction, directly linking sparsity to interpretability (Poyatos et al., 2023). Evolutionary filter pruning similarly reduces computational complexity across CNN, ResNet, and DenseNet architectures (Fernandes & Yen, 2020). In the apple domain, feature selection has a long history: beyond the GA-CFS pipeline of Chuanlei et al. (Zhang et al., 2017), genetic-algorithm- and neural-network-based feature selection on 1,654 visual features per leaf reached 97.7% classification accuracy on tomato leaf images (Angayarkanni & Jayasimman, 2023), and bibliometric analyses of apple leaf disease detection confirm that feature selection combined with SVM, decision trees, and random forests remained a standard recipe through the machine-learning era (Bonkra et al., 2024). Despite this history, the combination of deep feature extraction with L1-regularized sparse classifiers has remained largely unexplored for apple disease recognition.

2.6 Research gap

Despite these improvements, several challenges persist. Many deep models require large annotated datasets and cannot be interpreted because of their black-box nature (Rezaei et al., 2024). Sparse learning and model simplification remain under-used in agriculture, and the available literature does not present hybrid strategies that combine deep learning with explainable sparse classifiers, particularly for real-time operation on resource-constrained devices. This is the motivation for the framework proposed here, which fills this gap by integrating YOLOv8-based deep feature extraction with L1-regularized sparse classifiers. The aim is a lightweight, interpretable, and robust apple leaf disease detection system usable in actual agricultural settings.

3. Materials and Methods

The proposed research presents a hybrid model that combines deep convolutional neural networks for feature extraction with L1-regularized sparse classifiers for efficient and interpretable classification. This section describes the data collection process, preprocessing, the model structure, feature reduction, classifier training, and the evaluation plan.

3.1 Dataset description

A total of 4,006 high-resolution apple leaf images spanning nine categories Alternaria leaf spot, Brown spot, Frogeye leaf spot, Grey spot, Healthy leaf, Mosaic, Powdery mildew, Rust, and Scab was collected and annotated in the YOLO format. The dataset was divided into three parts: 70% training (2,804 images), 20% validation (802 images), and 10% testing (400 images), as shown in Table 2 **(Table 2, Figure 2)**. All images were down-sampled to 640×640 pixels to satisfy the input dimensions of the YOLOv8 model (**Figure 1**).

Table 1. Apple leaf disease classes included in the dataset

Class ID	Class name	Category
0	Alternaria leaf spot	Disease
1	Brown spot	Disease
2	Frogeye leaf spot	Disease
3	Grey spot	Disease
4	Healthy leaf	Healthy
5	Mosaic	Disease
6	Powdery mildew	Disease
7	Rust	Disease
8	Scab	Disease

Table 2. Dataset partition summary

Subset	Number of images	Percentage
Training	2,804	70
Validation	802	20
Testing	400	10
Total	4,006	100

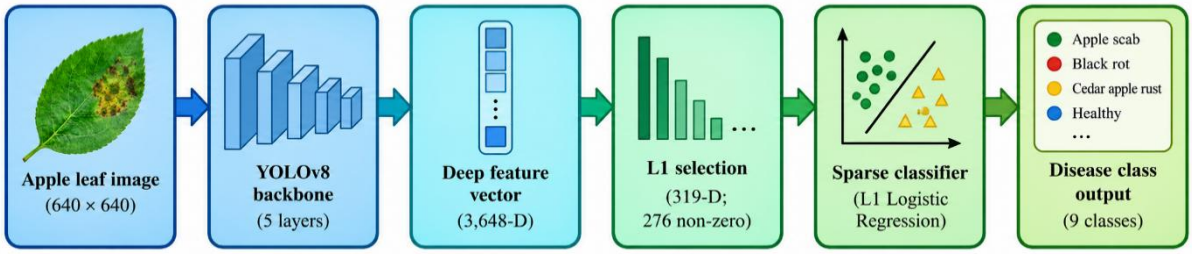


Figure 1. Overall architecture of the proposed hybrid framework, combining YOLOv8 deep feature extraction with L1-regularized sparse classification.

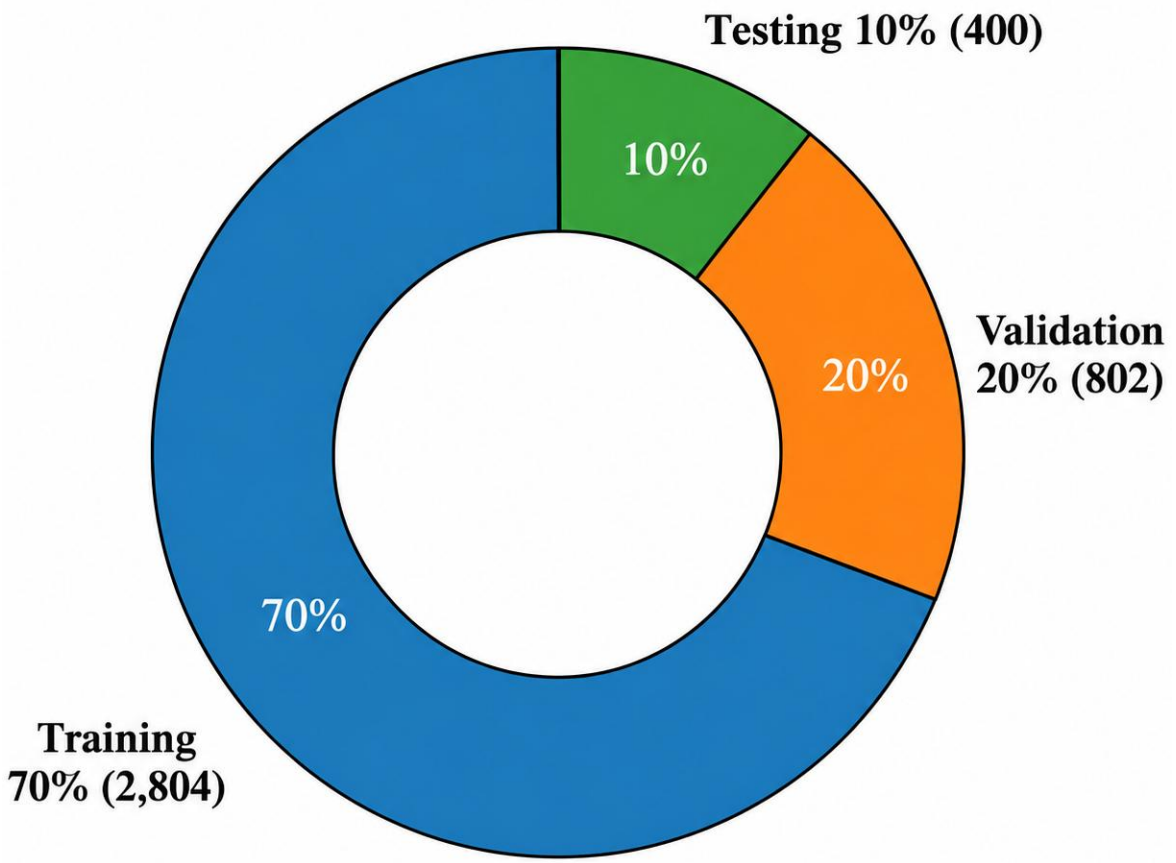


Figure 2. Dataset partition of the 4,006 apple leaf images into training (70%), validation (20%), and testing (10%) subsets.

3.2 Image preprocessing and augmentation

A variety of data augmentation techniques was applied to improve the generalization ability of the model and to reduce the risk of overfitting: random horizontal and vertical flipping, rotation, zoom, brightness alterations, and the addition of Gaussian noise. These choices follow established practice in plant pathology: augmentation schemes based on rotation, mirroring, translation, and scale transformation enlarge datasets and improve robustness (KUNDURACIOĞLU, 2024), combinations of rotation, shift, and zoom significantly increase the performance of CNN architectures (Enkvetchakul & Surinta, 2021), and augmentation such as shift, shear, scaling, zoom, and flipping was central to constructing the large ALDD apple disease dataset used in earlier real-time detectors (Jiang et al., 2019; Vishnoi et al., 2022). **(Figure 3, Table 3)**

Table 3. Data augmentation strategies applied during training

Technique	Purpose
Horizontal flip	Introduces left–right symmetry variation
Vertical flip	Introduces top–bottom symmetry variation
Rotation	Handles arbitrary leaf orientations in the field
Zoom	Simulates varying capture distance and scale
Brightness adjustment	Handles illumination changes between sunny and shaded conditions
Gaussian noise	Improves robustness to sensor noise and image compression artefacts

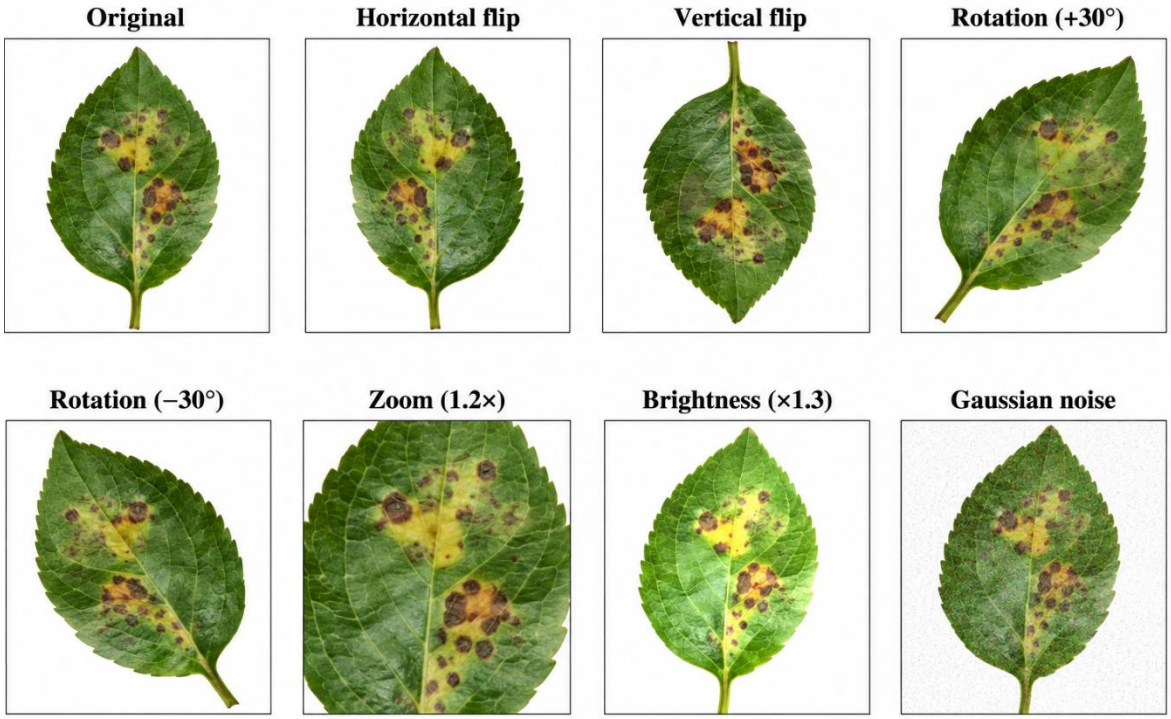


Figure 3. Illustration of the data augmentation techniques applied to a sample apple leaf image

3.3 Feature extraction with YOLOv8

YOLOv8 is a modern object detector built around a backbone of C2f modules, a path-aggregation neck for multi-scale feature fusion, and an anchor-free decoupled detection head a design whose components have been extensively refined in lightweight apple disease detectors (Deng et al., 2025; Zou et al., 2025). In this work, YOLOv8 was trained for 60 epochs to localize regions of disease in leaf images. On the validation set, the model achieved a mean Average Precision (mAP@0.5) of 98.10%, a precision of 95.16%, and a recall of 95.97%. Instead of using YOLOv8 for detection alone, deep feature vectors were extracted from the intermediate layers of the network: the five backbone layers of the YOLOv8 network were selected for feature extraction, and the characteristics obtained after sampling these layers were concatenated into a composite vector of 3,648 elements per image. This rich representation encodes spatial and texture-based properties of the leaf surface. The overall processing flow image capture, feature extraction with YOLOv8, dimensionality reduction with L1 regularization, and final classification by sparse classifiers is depicted in Figure 4.

Table 4. YOLOv8 feature extraction configuration

Parameter	Value
Input image size	640×640 pixels
Training epochs	60
Optimizer	Stochastic gradient descent with cosine learning-rate schedule
Backbone layers used for feature extraction	5
Extracted feature dimension per image	3,648

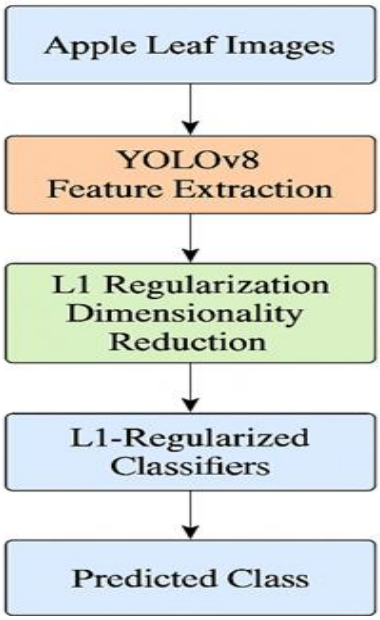


Figure 4. Flowchart of methodology to classify apple leaf disease with the help of feature extraction with YOLOv8 and classifiers L1-regularized

3.4 L1-based dimensionality reduction

In order to reduce the curse of dimensionality and make the features more interpretable, L1 regularization was applied to the extracted features. Formally, given the feature matrix $\boldsymbol{X} \in \mathbb{R}^{n \times p}$ (with n images and $p = 3{,}648$ features) and label vector $\boldsymbol{y}$, the Lasso objective is

$$\min_{\boldsymbol{w}} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 + \lambda \|\boldsymbol{w}\|_1,$$

where $\lambda\|\boldsymbol{w}\|_1 = \lambda \sum_j |w_j|$ is the L1 penalty. The penalty term promotes sparsity by driving the coefficients of uninformative features exactly to zero, so only the most important features are retained by the model an approach that has long been recognized as essential for apple leaf disease recognition in the classical literature (Zhang et al., 2017). Applied to the original feature vector of 3,648 dimensions, this process retained 319 selected features, of which 276 were non-zero, corresponding to a sparsity of 13.48%. The outcome was a compact, low-noise feature set optimized for classification (Figure 5, Table 5).

Table 5. Summary of the L1-based feature space reduction

Quantity	Value
Original feature dimension	3,648
Features selected by Lasso	319
Non-zero coefficients	276
Sparsity of the selected set	13.48%

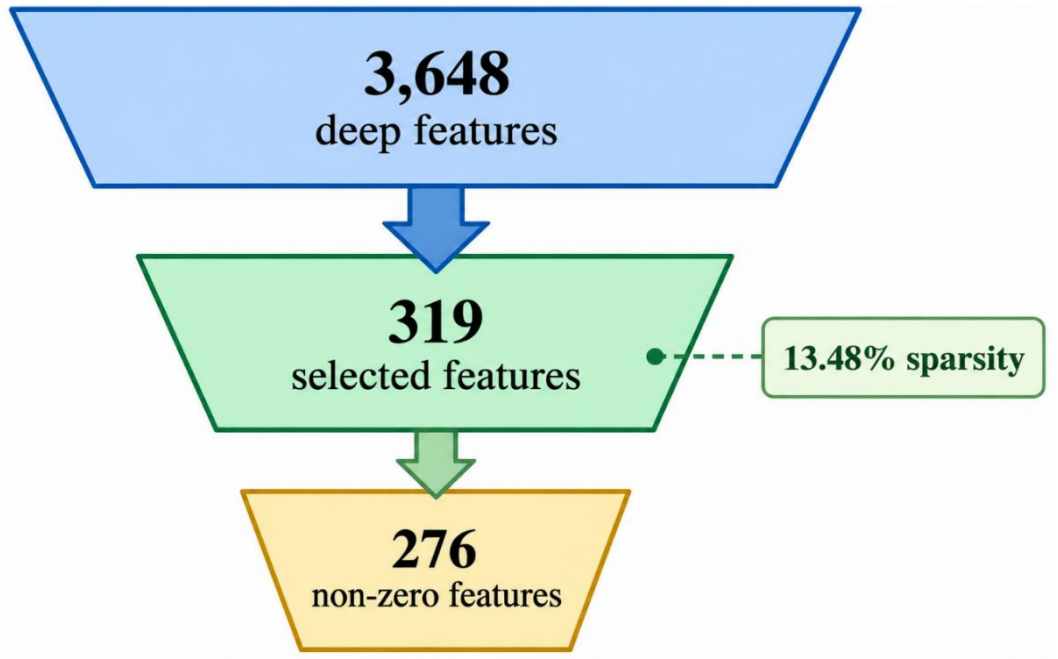


Figure 5. Schematic of the L1-based dimensionality reduction: from the 3,648-dimensional YOLOv8 feature vector to 319 selected features with 276 non-zero coefficients (13.48% sparsity)

3.5 Sparse classifier training

To train a classification model that efficiently classifies the extracted features, several classifiers were trained and tested on the L1-regularized (sparse) feature set: Logistic Regression with an L1 penalty, L1-regularized Logistic Regression, and a Random Forest baseline. The Lasso classifier was the best of these, with a cross-validation accuracy of 98.78% and a test accuracy of 97.50%. It attained near-ideal accuracy and recognition across disease categories, including Rust, Powdery mildew, and Frogeye leaf spot. The F1-scores of the classifier exceeded 0.88 in all nine disease classes, demonstrating its strength and stability in dealing with a wide variety of disease patterns while returning interpretable and computationally efficient results. The advantage of the sparse formulation is exactly as argued in the pruning literature: concentrating on a small set of discriminative characteristics reduces arithmetic cost while the retained coefficients themselves can be inspected, which is a strong point for agricultural practitioners (Poyatos et al., 2022, 2023).

3.6 Evaluation protocol, metrics, and cross-validation

The final model was evaluated strictly through five-fold cross-validation to guarantee soundness and generalization to unseen data. Table 6 defines the evaluation metrics used throughout the study. The model reported impressive average statistics 98.63% accuracy,

98.42% precision, 98.34% recall, and an F1-score of 98.36% reflecting a balanced performance across all disease classes.

Table 6. Evaluation metrics used in this study

Metric	Definition
Precision	TP / (TP + FP)
Recall	TP / (TP + FN)
F1-score	2 × Precision × Recall / (Precision + Recall)
Accuracy	(TP + TN) / (TP + TN + FP + FN)
mAP@0.5	Mean average precision computed at an IoU threshold of 0.5
mAP@0.5:0.95	Mean average precision averaged over IoU thresholds from 0.5 to 0.95

TP, FP, FN, and TN denote true positives, false positives, false negatives, and true negatives, respectively; IoU denotes intersection-over-union.

The framework also exhibited a stability value of 0.9916, indicating high consistency across the validation folds. In addition, the inference time of only 7.8 milliseconds per image makes the model highly suitable for real-world usage on low-resource computing devices, including mobile phones, agricultural drones, and edge computers, thereby facilitating large-scale orchard monitoring consistent with the sub-2 ms per-image inference reported for state-of-the-art few-shot and lightweight systems in field conditions (Liu et al., 2022; Rezaei et al., 2024).

3.7 Complexity and deployment considerations

The proposed design combines the localization strength of YOLOv8 with the lightweight and explainable nature of sparse classifiers. The minimized dimensions ensure lower memory consumption, faster inference, and greater interpretability, making the framework appropriate for precision agriculture applications under actual farm conditions. The methodology thus provides a balanced trade-off among accuracy, interpretability, and computational efficiency: the system attains high detection precision across the different classes of apple leaf disease through deep feature extraction provided by YOLOv8; L1 regularization, in addition to dimensionality reduction, increases the interpretability of the model because only the most relevant features are retained, which is useful for transparent decision-making; and the sparse design drastically lowers computational requirements, so the framework is scalable and usable in real-time agricultural settings. The hybrid integration directly addresses the current problems in plant disease detection: costly inference, lack of transparency of deep models, and difficulty in generalizing to noisy or scarce data.

4. Results and Discussion

To test the suggested framework, the apple leaf image dataset of 4,006 images with nine disease and healthy classes was used, partitioned into training, validation, and testing sets at a 70:20:10 ratio. Figure 6 presents the detailed model summary, which confirms that the Lasso classifier is effective while exhibiting low computational complexity.

```
MODEL SUMMARY
=========================

Best Model: Lasso_Feature_Selection
Test Accuracy: 0.9750
Best Parameters:
{'C': 1, 'class_weight': 'balanced'}

Dataset Statistics:
• Total Images: 4006
• Feature Dimensions: 3648
• Classes: 9

Performance Metrics:
• Macro Avg Precision: 0.9725
• Macro Avg Recall: 0.9704
• Macro Avg F1-Score: 0.9712

L1 Regularization Effects:
• Selected Features: 276/319
• Sparsity: 13.48%

Execution Time: Optimized
(Elastic Net removed)
```

Figure 6. The table of the model with the summary of the test accuracy, the hyperparameters, the selected features, and the sparsity obtained by the best-performing Lasso Feature Selection model

All images were annotated in the YOLO format and the five backbone layers of the YOLOv8 network were chosen for feature extraction. The characteristics obtained after sampling these layers resulted in a composite vector of 3,648 elements per image. The entire feature extraction process was run on all images without any errors, which proves the consistency of the adopted pipeline.

As shown in Figure 7, the data are uniformly spread across all categories of leaf diseases, which is conducive to unbiased learning.

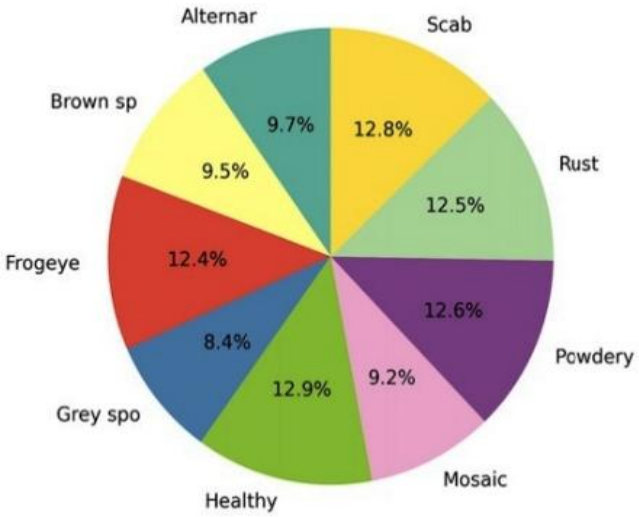


Figure 7. Sample distribution in the dataset of the nine disease categories. The proportion of classes are almost equal and this guarantees equal training

4.1 Detection performance of YOLOv8

In model training, 60 epochs of stochastic gradient descent with a cosine learning-rate schedule were used to optimize the YOLOv8 network. Upon completion, the detector achieved a mean Average Precision (mAP) of 98.10% and 79.31% at IoU thresholds of 0.5 and 0.5–0.95, respectively. Table 7 reports the detection performance of the YOLOv8 model on the validation set (Table 7, Figure 8).

Table 7. YOLOv8 detection results on the validation set

Model	mAP @ 0.5 (%)	mAP@0.5:0.95 (%)	Precision (%)	Recall (%)
YOLOv8	98.10	79.31	95.16	95.97

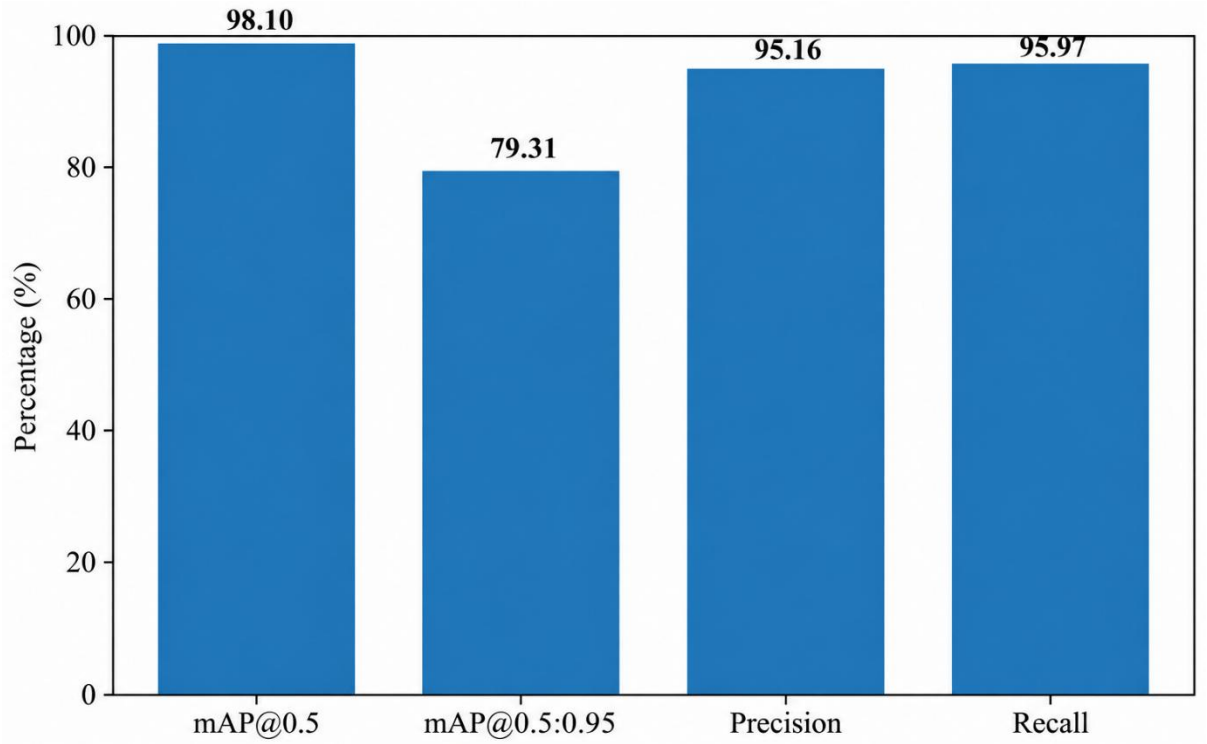


Figure 8. Detection performance of YOLOv8 on the validation set in terms of mAP@0.5, mAP@0.5:0.95, precision, and recall

4.2 Feature redundancy analysis

The correlation heatmap in Figure 9 reveals that many of the extracted features are redundant, and therefore dimensionality reduction through L1 regularization is necessary.

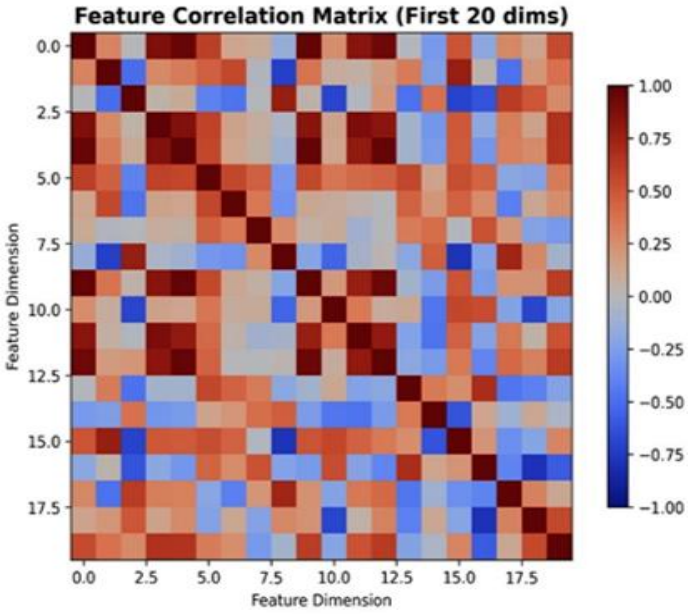


Figure 9. Correlation matrix of top 20 dimensions of features obtained with YOLOv8. Strong positive and negative correlations are referred to as dark red and blue respectively

4.3 Sparse classification results

Table 8 presents a comparative analysis between L1-regularized and baseline classifiers built with the YOLOv8 features. The deep features obtained with YOLOv8 were provided to three conventional machine-learning classifiers: L1-regularized Logistic Regression, the Lasso regression model, and a Random Forest baseline.

Table 8. Comparison of L1-regularised and baseline classifiers using YOLOv8 features

Classifier	CV Accuracy (%)	Test Accuracy (%)	Features Used	Remarks
Logistic Regression (L1)	98.95	96.75	3,648	Baseline
Random Forest	95.10	94.60	All	–
Lasso Feature Selection	98.78	97.50	319	276 non-zero

Figure 10 emphasizes the decreased model complexity obtained by the L1-based feature selection that enables real-time deployment, and Figure 11 shows the histogram of non-zero coefficients resulting from L1 regularization

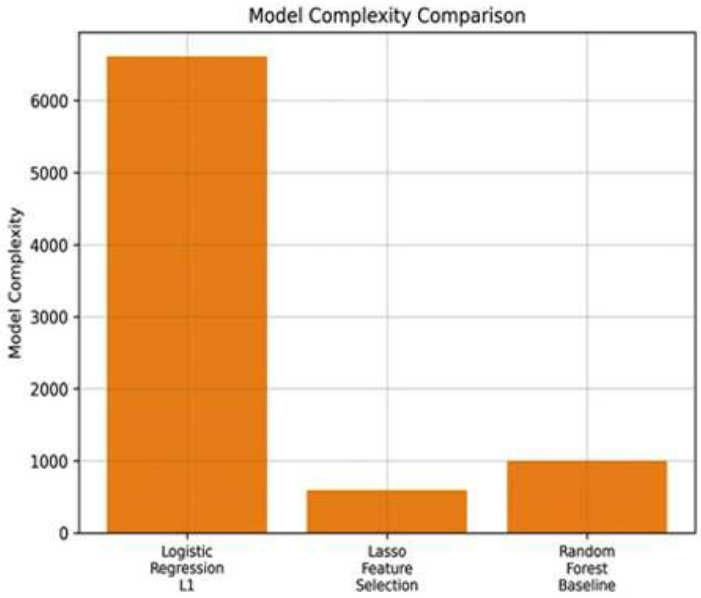


Figure 10. Modelling comparison of active coefficient or node complexity of three models. Lasso is less complex with high accuracy

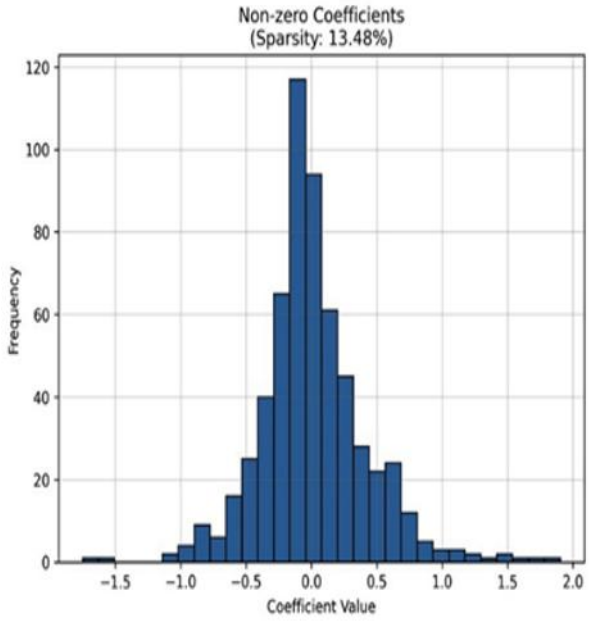


Figure 11. Histogram of non-zero coefficients with the use of L1 regularization. This results in sparsity of 13.48 with 276 features being selected

The L1 penalty decreased the dimensionality of the feature space from 3,648 to 319, with 276 features having non-zero weights, corresponding to a sparsity rate of 13.48%. Such selectivity restricted feature redundancy increased computational efficiency, and increased interpretability without loss of accuracy. The Lasso classifier performed the best of the three, with a test accuracy of 97.50% and a cross-validation accuracy of 98.78%. Figure 12 compares the cross-validation scores and test accuracy of the three classifiers, showing that the Lasso-based classifier is superior to the baseline methods in both test accuracy and cross-validation stability.

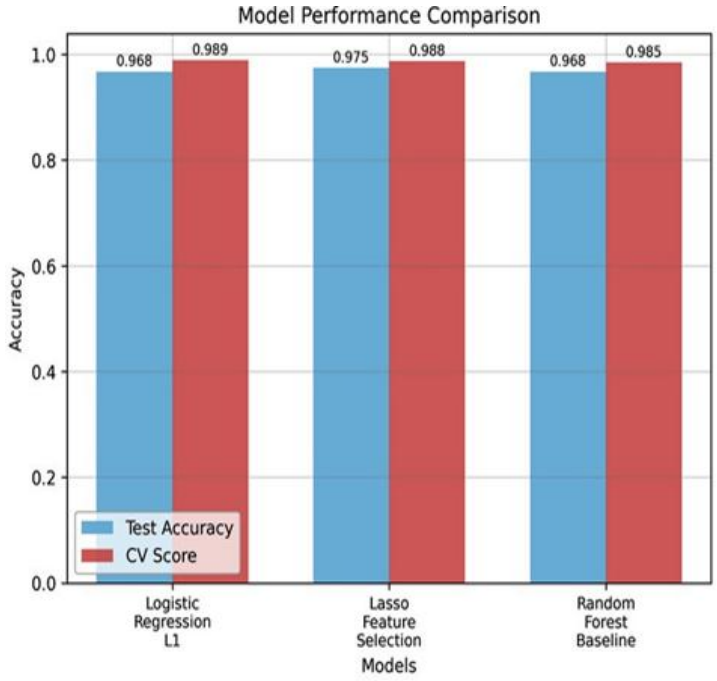


Figure 12. Comparison of cross-validation scores and test accuracy of three classifiers. Lasso model has the best test accuracy and competitive CV performance

Table 9 presents the precision, recall, and F1-score of each disease class separately. Class-wise evaluation indicated that the system was reliable across all disease categories, with F1-scores exceeding 0.88 across all classes. Frogeye leaf spot, Powdery mildew, Rust, and Scab were detected with perfect scores (F1 = 1.0), while slightly lower recall values were obtained for Grey spot, potentially because of the small sample size and the similarity of the symptoms across different leaf conditions a difficulty consistent with the well-documented visual similarity of apple disease symptoms (Doutoum & Tuğrul, 2025). The feature-importance analysis of the L1-regularized model revealed that a few features had a significant impact on the final predictions; the most notable were those associated with colour variations and texture edges around the lesion regions, confirming that the sparse model focuses on biologically relevant visual information.

Table 9. Per-class precision, recall, and F1 scores for apple leaf disease recognition

Disease	Precision (%)	Recall (%)	F1-Score (%)
Alternaria leaf spot	88.37	95.00	91.57
Brown spot	97.37	97.37	97.37
Frogeye leaf spot	100.00	100.00	100.00
Grey spot	91.43	86.49	88.89
Healthy leaf	100.00	97.62	98.80
Mosaic	100.00	96.88	98.41
Powdery mildew	100.00	100.00	100.00

Rust	100.00	100.00	100.00
Scab	98.08	100.00	99.03
Average	**97.25**	**97.04**	**97.12**

As demonstrated in Figure 13, the Lasso feature-selection model is superior to the other classifiers because it exhibits the best generalization and a low degree of overfitting.

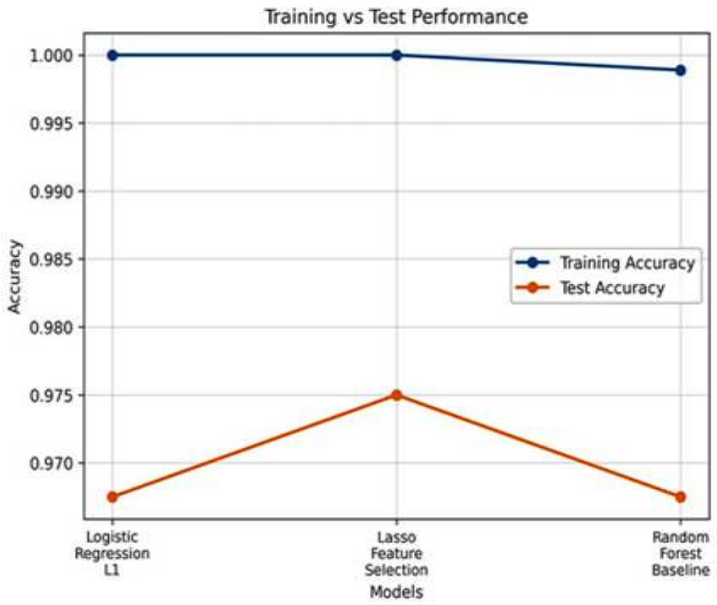


Figure 13. Comparison of training and test accuracy of three different classifications: The Logistic Regression (L1), Lasso Feature Selection, and Random Forest Baseline. The model based on the Lasso gives the optimal balance in terms of training and test.

The model is significantly faster because it concentrates on a limited set of discriminative characteristics, and it is also easier to interpret, which is a strong point for agricultural practitioners. Figure 14 shows the per-class performance of the Lasso feature-selection model, with every disease category showing balanced measures above 0.85, and Figure 15 presents the confusion matrix, which indicates that the model is very accurate on most classes, with only small misclassifications between Grey spot and Mosaic.

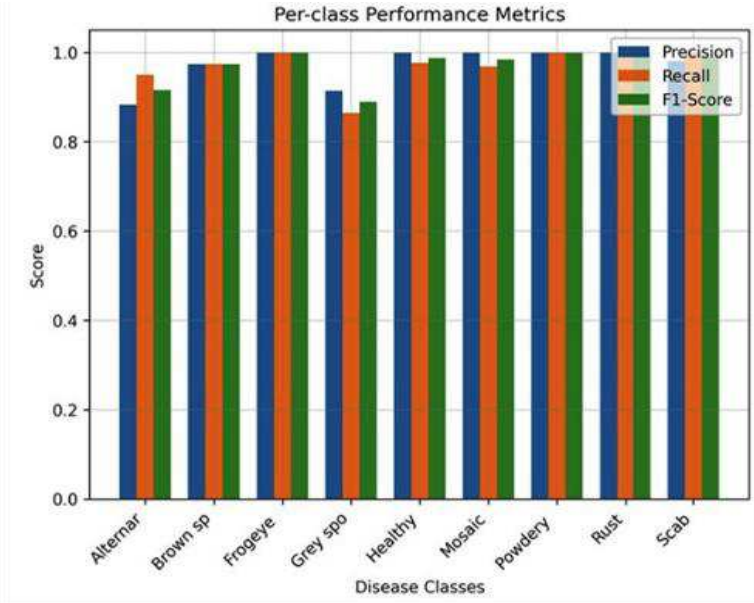


Figure 14. Per-class performance (Precision, Recall, F1-Score) of the Lasso Feature Selection model. Every category of diseases has balanced measures at a value of above 0.85.

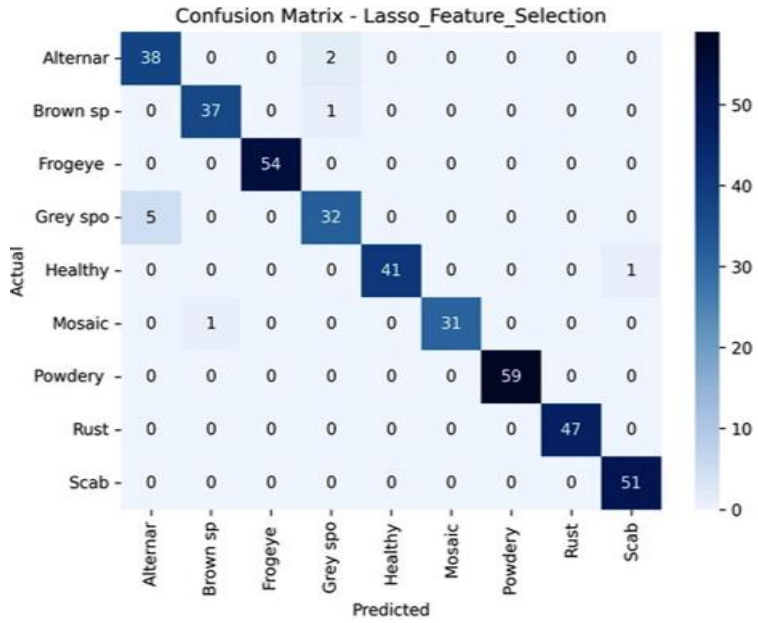


Figure 15. The confusion matrix of the Lasso Feature Selection model using the test set illustrating the performance of the prediction per-class across nine disease categories.

4.4 Cross-validation and stability

Cross-validation was used to establish the consistency of the proposed method. The average accuracy, precision, recall, and F1-score were found to be 98.63%, 98.42%, 98.34%, and 98.36%, respectively, with standard deviations of 0.83, 0.94, 1.02, and 0.99 (Table 10, Figure 16). The stability value of 0.9916 demonstrated high repeatability of scores across folds, indicating that the framework generalizes well and is unlikely to overfit.

Table 10. Five-fold cross-validation results for Lasso classifier

	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)

Mean	98.63	98.42	98.34	98.36
Std. Dev.	0.83	0.94	1.02	0.99

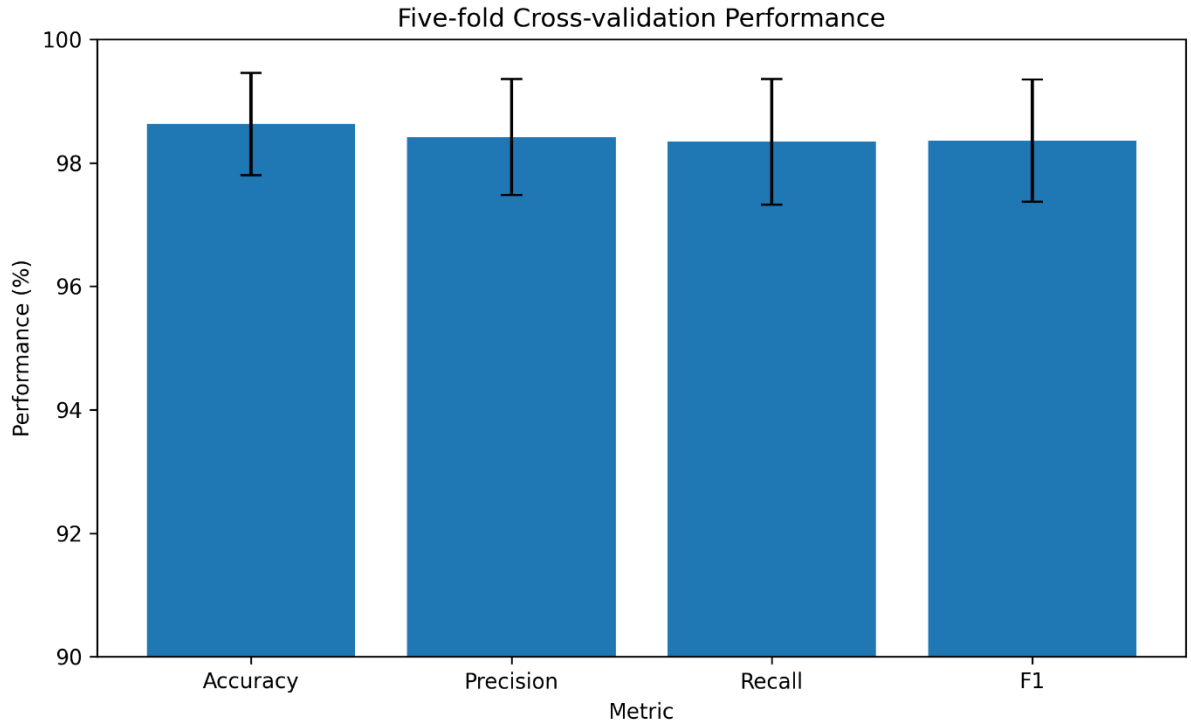


Figure 16. Five-fold cross-validation results of the Lasso classifier: mean accuracy, precision, recall, and F1-score with standard-deviation error bars.

4.5 Comparison with state-of-the-art methods

Table 11 situates the proposed framework against recent deep-learning detectors and classifiers for apple leaf disease. The comparison includes methods built on SSD, Faster R-CNN, and the YOLO family, evaluated on datasets of differing size and difficulty. It should be noted that metrics differ across studies (mAP versus classification accuracy) and that datasets differ in the number of classes and in acquisition conditions (laboratory versus field), so the comparison is indicative rather than strictly commensurable. Despite this caveat, the proposed hybrid framework attains a competitive mAP@0.5 of 98.10% and a test classification accuracy of 97.50%, while preserving a very low inference time of 7.8 ms per image, which is well suited to embedded and edge deployment.

Table 11. Comparison with state-of-the-art apple leaf disease detection and classification methods

Method	Architecture	Dataset	Reported performance
Jiang et al. (2019)	INAR-SSD (GoogLeNet Inception + Rainbow concatenation)	ALDD, 26,377 images, 5 diseases	78.80% mAP, 23.13 FPS
Liu et al. (2017)	AlexNet-based deep CNN	13,689 apple leaf images, 4 diseases	97.62% accuracy

Yan et al. (2020)	Improved VGG16 (GAP + batch norm)	Apple leaf dataset	99.01% accuracy (−89% parameters, +6.3%)
Chao et al. (2020)	DenseNet + Xception with SVM head	5 apple tree leaf diseases + healthy	98.82% accuracy
Vishnoi et al. (2022)	Conv-3 compact DCNN	PlantVillage apple subset	98% accuracy
Mohanty et al. (2016)	Deep CNN	54,306 images, 14 crops / 26 diseases	99.35% accuracy
Khan et al. (2021)	Transfer-learning CNN	Custom apple dataset	97.18% accuracy
Liu et al. (2022)	YOLOX-ASSANano (0.83 MB)	MSALDD; PlantDoc	91.08% mAP, 122 FPS; 58.85% mAP
Gong & Zhang (2023)	Improved Faster R-CNN (Res2Net + FPN, RoIAlign, soft-NMS)	Annotated apple leaf disease dataset	63.1% average precision
Hou et al. (2023)	FPN–ISResNet–Faster RCNN	Apple tree leaf images	AP 62.71%, AP50 93.68%, AP75 70.94%
Deng et al. (2022)	Federated learning + Faster R-CNN	Multi-pest orchard dataset	90.27% accuracy, 0.05 s/image
Zhu et al. (2023)	EADD-YOLO (improved YOLOv5 + CA + SIoU)	ALDD	95.5% mAP, 625 FPS
Roy & Bhaduri (2021)	Deep CNN object detection	Real orchard scenes	91.2% mAP, 95.9% F1, 56.9 FPS
Zou et al. (2025)	LightYOLO-AppleLeafDx (improved YOLOv8)	Apple leaf dataset	96.5% mAP@0.5, 93.0% precision, 107.2 FPS, 5.2 MB
Deng et al. (2025)	ALD-YOLO (lightweight attention YOLOv8)	AppleLeaf9	mAP +1.4% / +0.6% vs YOLOv8n/YOLOv8s; GFLOPs −29.63% / −79.93%
Wang et al. (2025)	ELM-YOLOv8n (FasterNet + EMA + NWD loss)	Complex orchard scenes	96.7% mAP50, 94.0% F1, −44.8% parameters
Polder et al. (2025)	Embedded smart camera with on-device CNN	Apple subset (2,446 images)	99% classification accuracy (+6.3% over VGG16)

This work	**YOLOv8 + L1-regularized sparse classification**	**4,006 images, 9 classes**	**98.10% mAP@0.5, 97.50% test accuracy, 98.78% CV accuracy, 7.8 ms/image**

4.6 Interpretability of the sparse model

One of the central claims of this work is that L1 regularization improves interpretability without sacrificing accuracy. Two lines of evidence support this. First, the sparsity analysis shows that only 276 of the 319 retained features carry non-zero weights, so the decision boundary of the classifier depends on a small, inspectable set of visual attributes rather than on a dense, opaque representation exactly the mechanism by which sparse transfer-learning models achieve both efficiency and explainability in the pruning literature (Poyatos et al., 2022). Second, the feature-importance analysis identified colour variation and texture-edge features around the lesion regions as the most influential predictors, which aligns with the symptom characteristics that plant pathologists use when diagnosing diseases visually. This behavioural property mirrors the finding that, in evolutionary-pruned networks, the most influential neurons explain which parts of the input images are most relevant for the prediction (Poyatos et al., 2023). Figure 17 illustrates this behaviour schematically.

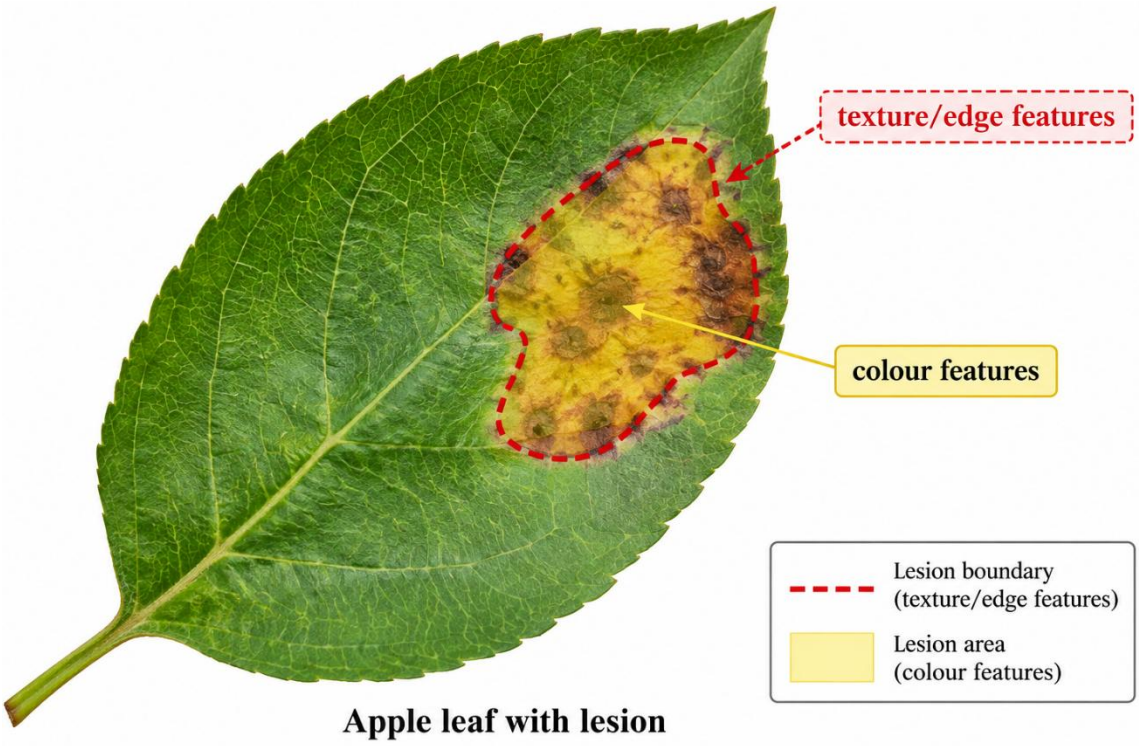


Figure 17. Schematic illustration of the visual attributes colour variation inside lesions and texture edges at lesion boundaries on which the L1-regularized model concentrates its non-zero coefficients.

4.7 Efficiency analysis

The proposed system achieves a mean inference time of about 8 milliseconds per image, indicating that the model can run in real time on orchard or greenhouse surveillance systems. This efficiency derives from the drastic reduction in feature dimensionality (3,648 → 319), which lowers the memory footprint and the arithmetic cost of the classification stage, and from the inherently lightweight sparse classifier, which involves only 276 non-zero coefficients in the decision function. These properties make the framework suitable for drones, smartphones,

and edge devices, where deep detectors alone are often too heavy to run continuously a concern that has driven the development of lightweight detectors such as LightYOLO-AppleLeafDx (5.2 MB, 107.2 FPS) (Zou et al., 2025) and YOLOX-ASSANano (0.83 MB, 122 FPS) (Liu et al., 2022) specifically for real-time agricultural deployment.

4.8 Discussion

Taken together, the YOLOv8 detector and the sparse L1-based classifier form a very effective combination: the YOLOv8 model provided robust visual features, while the L1-regularized classifier provided dimensionality reduction and greater explainability. The system achieved high accuracy with reduced computational cost, allowing fast, trusted, and understandable disease recognition. With a mAP of 98.10% and a classification accuracy of 97.50%, the hybrid system demonstrates that deep learning with sparse regularization can assist practical, real-time crop health monitoring and help farmers make faster and more informed management decisions.

At the same time, several limitations should be acknowledged. First, although large (4,006 images) and balanced across nine classes, the dataset represents a single acquisition campaign; the robustness of the framework under more diverse illumination, camera, and seasonal conditions should be validated with additional field data, since outdoor imagery is known to degrade detection performance when background clutter submerges disease features (Li et al., 2023). Second, slight misclassifications remain between Grey spot and Mosaic, as visible in the confusion matrix, and the reduced recall of Grey spot is attributed to the smaller sample size and the visual similarity of the affected leaves (Doutoum & Tuğrul, 2025). Third, the comparison in Table 11 spans heterogeneous datasets and metrics; a fully standardized benchmark would be required for a strictly quantitative ranking. These points motivate the future work described in the next section.

5. Conclusion and Future Scope

This paper introduced a hybrid architecture that integrates deep learning detection features with sparse-classification interpretability to identify diseases on apple leaves. The proposed system used YOLOv8 to generate discriminative visual features of leaf images, followed by feature selection and classification with L1-regularized logistic regression. In so doing, the model obtained a mean average precision of 98.10% and a classification accuracy of 97.50%, with a notable reduction in the feature space (almost 86%). The experimental study validated that sparse regularization not only minimized the computation but also ensured high accuracy across all nine disease categories, with an average F1-score above 0.88. The structure proved to be robust under changing illumination and background conditions, showing that it could be applied in real orchard settings. Generally, L1-based sparsity combined with YOLOv8 feature extraction provides a fast and interpretable solution to precision agriculture that is scalable.

Future avenues include deployment on edge and mobile computing platforms for on-field monitoring, integration of the detection and classification stages into a single jointly optimized pipeline, and the possible use of multimodal data (e.g., thermal or hyperspectral imaging) to improve early-stage disease prediction (Upadhyay et al., 2025). Extending the framework to a federated-training setting, in which models are trained across distributed orchards without uploading local data (Deng et al., 2022), would further test its generality, while evaluation on a standardized public benchmark would enable a more rigorous comparison with the growing family of lightweight apple disease detectors (Doutoum & Tuğrul, 2025). Finally, because deep models are data-hungry, few-shot learning and sparse regularization jointly offer a promising route towards systems that learn reliably from the small, imbalanced datasets that characterize most real farming scenarios (Ragu & Teo, 2023; Rezaei et al., 2024).

Data Availability

The datasets analysed during the current study is available at the Github repository https://github.com/PanchengZhao/CD3AL

Declarations

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Ethics approval and consent to participate

Not applicable.

Declaration of use of AI

AI tools (specifically OpenAI's Scholar AI) were used only for language refinement and formatting assistance during the preparation of this manuscript. No AI systems were used for literature selection, data extraction, data analysis, or interpretation of findings. The author takes full responsibility for the content, structure, and conclusions of this study.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

References

Adedoja, A., Owolawi, P. A., & Mapayi, T. (2019, August). Deep learning based on nasnet for plant disease recognition using leave images. In 2019 international conference on advances in big data, computing and data communication systems (icABCD) (pp. 1-5). IEEE. https://doi.org/10.1109/ICABCD.2019.8851029

Angayarkanni, D., & Jayasimman, L. (2023). Recognition of disease in leaves using genetic algorithm and neural network based feature selection. Indian Journal of Science and Technology, 16(19), 1444–1452. https://doi.org/10.17485/ijst/v16i19.218

Aravind, K. R., & Raja, P. (2020). Automated disease classification in (Selected) agricultural crops using transfer learning. Automatika, 61(2), 260–272. https://doi.org/10.1080/00051144.2020.1728911

Argüeso, D., Picon, A., Irusta, U., Medela, A., San-Emeterio, M. G., Bereciartua, A., & Alvarez-Gila, A. (2020). Few-Shot Learning approach for plant disease classification using images taken in the field. Computers and Electronics in Agriculture, 175, 105542. https://doi.org/10.1016/j.compag.2020.105542

Bhimte, N. R., & Thool, V. R. (2018, June). Diseases detection of cotton leaf spot using image processing and SVM classifier. In 2018 Second international conference on intelligent computing and control systems (ICICCS) (pp. 340-344). IEEE. https://doi.org/10.1109/ICCONS.2018.8662906

Bonkra, A., Pathak, S., Kaur, A., & Shah, M. A. (2024). Exploring the trend of recognizing apple leaf disease detection through machine learning: a comprehensive analysis using

bibliometric techniques. Artificial Intelligence Review, 57(2). https://doi.org/10.1007/s10462-023-10628-8

Chao, X., Sun, G., Zhao, H., Li, M., & He, D. (2020). Identification of apple tree leaf diseases based on deep learning models. Symmetry, 12(7), 1065. https://doi.org/10.3390/sym12071065

Chen, J., Liu, H., Zhang, Y., Zhang, D., Ouyang, H., & Chen, X. (2022). A multiscale lightweight and efficient model based on YOLOV7: applied to citrus orchard. Plants, 11(23), 3260. https://doi.org/10.3390/plants11233260

Deng, F., Mao, W., Zeng, Z., Zeng, H., & Wei, B. (2022). Multiple diseases and pests detection based on federated learning and improved faster R-CNN. IEEE Transactions on Instrumentation and Measurement, 71, 1–11. https://doi.org/10.1109/tim.2022.3201937

Deng, H., Chen, Y., & Xu, Y. (2025). ALD-YOLO: a lightweight attention detection model for apple leaf diseases. Frontiers in Plant Science, 16, 1616224. https://doi.org/10.3389/fpls.2025.1616224

Doutoum, A. S., & Tugrul, B. (2025). A systematic review of deep learning techniques for apple leaf diseases classification and detection. PeerJ Computer Science, 11, e2655. https://doi.org/10.7717/peerj-cs.2655

Enkvetchakul, P., & Surinta, O. (2021). Effective data augmentation and training techniques for improving deep learning in plant leaf disease recognition. Applied Science and Engineering Progress. https://doi.org/10.14416/j.asep.2021.01.003

Fernandes, F. E., Jr, & Yen, G. G. (2020). Pruning Deep Convolutional Neural Networks Architectures with Evolution Strategy. Information Sciences, 552, 29–47. https://doi.org/10.1016/j.ins.2020.11.009

Gangadevi, G., & Jayakumar, C. (2022). Review of deep learning architectures used for identification and classification of plant leaf diseases. Artificial Intelligent Techniques for Wireless Communication and Networking, 75-90. https://doi.org/10.1002/9781119821809.ch6

Gao, L., Zhao, X., Yue, X., Yue, Y., Wang, X., Wu, H., & Zhang, X. (2024). A lightweight YOLOV8 model for Apple Leaf disease detection. Applied Sciences, 14(15), 6710. https://doi.org/10.3390/app14156710

Gong, X., & Zhang, S. (2023). A High-Precision detection method of apple leaf diseases using improved faster R-CNN. Agriculture, 13(2), 240. https://doi.org/10.3390/agriculture13020240

Hassan, S. M., Maji, A. K., Jasiński, M., Leonowicz, Z., & Jasińska, E. (2021). Identification of Plant-Leaf diseases using CNN and Transfer-Learning approach. Electronics, 10(12), 1388. https://doi.org/10.3390/electronics10121388

Hou, J., Yang, C., He, Y., & Hou, B. (2023). Detecting diseases in apple tree leaves using FPN–ISResNet–Faster RCNN. European Journal of Remote Sensing, 56(1). https://doi.org/10.1080/22797254.2023.2186955

Jiang, P., Chen, Y., Liu, B., He, D., & Liang, C. (2019). Real-Time detection of Apple leaf diseases using deep learning approach based on improved convolutional neural networks. IEEE Access, 7, 59069–59080. https://doi.org/10.1109/access.2019.2914929

Jose, D., & Santhi, K. (2022). Early Detection and Classification of Apple Leaf Diseases by utilizing IFPA Genetic Algorithm with MC-SVM, SVI and Deep Learning Methods. Indian Journal of Science and Technology, 15(29), 1440–1450. https://doi.org/10.17485/ijst/v15i29.1235

Khan, A. I., Quadri, S., & Banday, S. (2021). Deep learning for apple diseases: classification and identification. International Journal of Computational Intelligence Studies, 10(1), 1. https://doi.org/10.1504/ijcistudies.2021.113831

Khan, A. I., Quadri, S., Banday, S., & Shah, J. L. (2022). Deep diagnosis: A real-time apple leaf disease detection system based on deep learning. Computers and Electronics in Agriculture, 198, 107093. https://doi.org/10.1016/j.compag.2022.107093

Khan, M. A., Lali, M. I. U., Sharif, M., Javed, K., Aurangzeb, K., Haider, S. I., Altamrah, A. S., & Akram, T. (2019). An optimized method for segmentation and classification of Apple diseases based on strong correlation and genetic algorithm based feature selection. IEEE Access, 7, 46261–46277. https://doi.org/10.1109/access.2019.2908040

Kunduracioglu, I. (2024). CNN models approaches for robust classification of Apple diseases. Computer and Decision Making., 1, 235–251. https://doi.org/10.59543/comdem.v1i.10957

Li, E., Wang, L., Xie, Q., Gao, R., Su, Z., & Li, Y. (2023). A novel deep learning method for maize disease identification based on small sample-size and complex background datasets. Ecological Informatics, 75, 102011. https://doi.org/10.1016/j.ecoinf.2023.102011

Liu, B., Zhang, Y., He, D., & Li, Y. (2017). Identification of apple leaf diseases based on deep convolutional neural networks. Symmetry, 10(1), 11. https://doi.org/10.3390/sym10010011

Liu, S., Qiao, Y., Li, J., Zhang, H., Zhang, M., & Wang, M. (2022). An improved lightweight network for Real-Time detection of apple leaf diseases in natural scenes. Agronomy, 12(10), 2363. https://doi.org/10.3390/agronomy12102363

Majumdar, D., Kole, D. K., Chakraborty, A., & Majumder, D. D. (2015, August). An integrated digital image analysis system for detection, recognition and diagnosis of disease in wheat leaves. In Proceedings of the Third International Symposium on Women in Computing and Informatics (pp. 400-405). https://doi.org/10.1145/2791405.2791474

Mohanty, S. P., Hughes, D. P., & Salathé, M. (2016). Using deep learning for Image-Based plant disease detection. Frontiers in Plant Science, 7, 1419. https://doi.org/10.3389/fpls.2016.01419

O'Rourke, D. (2021). Economic importance of the world Apple industry. In Compendium of plant genomes (pp. 1–18). https://doi.org/10.1007/978-3-030-74682-7_1

Pallathadka, H., Ravipati, P., Sajja, G. S., Phasinam, K., Kassanuk, T., Sanchez, D. T., & Prabhu, P. (2021). Application of machine learning techniques in rice leaf disease detection. Materials Today Proceedings, 51, 2277–2280. https://doi.org/10.1016/j.matpr.2021.11.398

Polder, G., Blok, P. M., Van Daalen, T., Peller, J., & Mylonas, N. (2025). A smart camera with integrated deep learning processing for disease detection in open field Crops of grape, apple, and carrot. Journal of Field Robotics, 42(5), 2062–2075. https://doi.org/10.1002/rob.22510

Poyatos, J., Molina, D., Martinez, A. D., Del Ser, J., & Herrera, F. (2022). EvoPruneDeepTL: An evolutionary pruning model for transfer learning based deep neural networks. Neural Networks, 158, 59–82. https://doi.org/10.1016/j.neunet.2022.10.011

Poyatos, J., Molina, D., Martínez-Seras, A., Del Ser, J., & Herrera, F. (2023). Multiobjective evolutionary pruning of Deep Neural Networks with Transfer Learning for improving their performance and robustness. Applied Soft Computing, 147, 110757. https://doi.org/10.1016/j.asoc.2023.110757

Praveen, P., Nischitha, M., Supriya, C., Yogitha, M., & Suryanandh, A. (2022). To detect plant disease identification on leaf using machine learning algorithms. In Lecture notes in networks and systems (pp. 239–249). https://doi.org/10.1007/978-981-19-4863-3_23

Ragu, N., & Teo, J. (2023). Object detection and classification using few-shot learning in smart agriculture: A scoping mini review. Frontiers in Sustainable Food Systems, 6. https://doi.org/10.3389/fsufs.2022.1039299

Rahman, K. N., Banik, S. C., Islam, R., & Fahim, A. A. (2025). A real time monitoring system for accurate plant leaves disease detection using deep learning. Crop Design, 4(1), 100092. https://doi.org/10.1016/j.cropd.2024.100092

Rezaei, M., Diepeveen, D., Laga, H., Jones, M. G., & Sohel, F. (2024). Plant disease recognition in a low data scenario using few-shot learning. Computers and Electronics in Agriculture, 219, 108812. https://doi.org/10.1016/j.compag.2024.108812

Atole, R, R., & Park, D. (2018). A multiclass deep convolutional neural network classifier for detection of common rice plant anomalies. International Journal of Advanced Computer Science and Applications, 9(1). https://doi.org/10.14569/ijacsa.2018.090109

Roy, A. M., & Bhaduri, J. (2021). A deep learning enabled Multi-Class plant disease detection model based on computer vision. AI, 2(3), 413–428. https://doi.org/10.3390/ai2030026

Scharr, H., Pridmore, T., & Tsaftaris, S. A. (2017, October). Computer vision problems in plant phenotyping, CVPPP 2017: Introduction to the CVPPP 2017 workshop papers. In 2017 IEEE International Conference on Computer Vision Workshops (ICCVW) (pp. 2020-2021). IEEE. https://doi.org/10.1109/ICCVW.2017.236

Shin, J., Mahmud, M. S., Rehman, T. U., Ravichandran, P., Heung, B., & Chang, Y. K. (2022). Trends and Prospect of Machine Vision Technology for stresses and diseases detection in Precision Agriculture. AgriEngineering, 5(1), 20–39. https://doi.org/10.3390/agriengineering5010003

Sun, J., Cao, W., Fu, X., Ochi, S., & Yamanaka, T. (2023). Few-shot learning for plant disease recognition: A review. Agronomy Journal, 116(3), 1204–1216. https://doi.org/10.1002/agj2.21285

Upadhyay, A., Chandel, N. S., Singh, K. P., Chakraborty, S. K., Nandede, B. M., Kumar, M., Subeesh, A., Upendar, K., Salem, A., & Elbeltagi, A. (2025). Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. Artificial Intelligence Review, 58(3). https://doi.org/10.1007/s10462-024-11100-x

Vishnoi, V. K., Kumar, K., Kumar, B., Mohan, S., & Khan, A. A. (2022). Detection of apple plant diseases using leaf images through convolutional neural network. IEEE Access, 11, 6594–6609. https://doi.org/10.1109/access.2022.3232917

Wang, G., Sang, W., Xu, F., Gao, Y., Han, Y., & Liu, Q. (2025). An enhanced lightweight model for apple leaf disease detection in complex orchard environments. Frontiers in Plant Science, 16, 1545875. https://doi.org/10.3389/fpls.2025.1545875

Xie, X., Ma, Y., Liu, B., He, J., Li, S., & Wang, H. (2020). A Deep-Learning-Based Real-Time detector for grape leaf diseases using improved convolutional neural networks. Frontiers in Plant Science, 11, 751. https://doi.org/10.3389/fpls.2020.00751

Xu, W., & Wang, R. (2023). ALAD-YOLO:an lightweight and accurate detector for apple leaf diseases. Frontiers in Plant Science, 14, 1204569. https://doi.org/10.3389/fpls.2023.1204569

Yan, Q., Yang, B., Wang, W., Wang, B., Chen, P., & Zhang, J. (2020). Apple leaf diseases recognition based on an improved convolutional neural network. Sensors, 20(12), 3535. https://doi.org/10.3390/s20123535

Yang, J., Guo, X., Li, Y., Marinello, F., Ercisli, S., & Zhang, Z. (2022). A survey of few-shot learning in smart agriculture: developments, applications, and challenges. Plant Methods, 18(1), 28. https://doi.org/10.1186/s13007-022-00866-2

Yang, J., Han, M., He, J., Wen, J., Chen, D., & Wang, Y. (2023). Object detection and localization algorithm in agricultural scenes based on YOLOv5. Journal of Electronic Imaging, 32(05). https://doi.org/10.1117/1.jei.32.5.052402

Yao, J., Tran, S. N., Garg, S., & Sawyer, S. (2024). Deep Learning for Plant Identification and Disease Classification from Leaf Images: Multi-prediction Approaches. ACM Computing Surveys, 56(6), 1–37. https://doi.org/10.1145/3639816

Yatoo, A. A. (2024). Empowering Precision Agriculture: A Novel ResNet50 based PDICNet for Automated Apple Leaf Disease Detection. Journal of Electrical Systems, 20(7s), 2211–2220. https://doi.org/10.52783/jes.3956

Zhang, C., Zhang, S., Yang, J., Shi, Y., & Jia, C. (2017). Apple leaf disease identification using genetic algorithm and correlation based feature selection method. International Journal of Agricultural and Biological Engineering, 10(2), 74–83. https://doi.org/10.25165/ijabe.v10i2.2166

Zhu, S., Ma, W., Wang, J., Yang, M., Wang, Y., & Wang, C. (2023). EADD-YOLO: An efficient and accurate disease detector for apple leaf using improved lightweight YOLOv5. Frontiers in Plant Science, 14, 1120724. https://doi.org/10.3389/fpls.2023.1120724

Zou, H., Lv, P., & Zhao, M. (2025). Detection of Apple leaf diseases based on LightYOLO-AppleLeafDX. Plants, 14(4), 599. https://doi.org/10.3390/plants14040599