

Experimental Insights into Plant Disease Detection: Parametric, Combinatorial, and Computational Evaluations of Data Mining and Optimization Approaches

Swapnil Wagh , Ruchi Sharma , Ankit Temurnikar

Corresponding author: swapnilwagh332@gmail.com

Computer Science and Engineering, Madhyanchal Professional University, Ratibad, Bhopal - 462044, M.P., India

Abstract- — Plant diseases are also known to place huge burden on food security structure and agriculture to the global world; it is approximated that all plant diseases development costs a giant (an estimated 220 billion/year). To address this, the computer vision -specific and deep learning based automated disease detection systems are expandingly viewed as rather interesting as an option instead of the traditional forms of diagnosing that involve a significant amount of new employees . However, the literature screening is saturated with models that have been alleged to be super high in accuracy with regard to classifications when they are under some form of controlled conditions in the laboratory that must in no way imply any trustworthy depiction that they can be relayed over the situation in the real field. It can be said that such discrepancy in performance can stress the idea that there is a dire necessity to carry out more related and stiffer analysis of existing measures of data mining and optimization. This article has such an experimental alloy of which the plant disease variable models can be detected multi faceted in, which is discussed in detail on three axes parametric axis, combinatorial axis, computational axis. The rate of model performances to the hyperparameter options enshrined in the parametric assessment that may also be the optimizers are called counting. The combinatorial work involves the study of connections pertaining to the utility of various Convolutional Neural Network Convolutional designs, as well as the use of spectacular measures of data augmentation and fold up learning methods. The computational verification provided is a practical test of the feasibility of the model, comparison of statistics on the training time, model complexity, and speed of inference. According to the opinion that our experimental findings indicate, our individual models (as well as our EfficientNet) that come with the highest classification performance of about above 98 percent accuracy would always be the best trade off between accuracy and efficiency whereas ensemble models would adopt a combination of soft voting as the best trade off prerogative. The paper further estimates the radical performance augmentation with the generative data augmentation models against the conventional geometric transformations to apply the models in the truly competitive use. The primary accomplishment of this project is the system, which surpasses those pathetic signs of precision and rests upon the familiarization of scientists and performers with how to create, alter, and put to practical practice the scaleable, resilient, and effective plant disease detection methods used in the enhancement of the designated work in the agricultural forerunners.

Keywords : Plant Disease Detection , Data Mining Techniques , Optimization Algorithms , Parametric Analysis , Combinatorial Evaluation , Computational Approaches.

I. INTRODUCTION

A. The Critical Role of Automated Disease Detection in Modern Agriculture

Food crisis has become a factual issue that has been gaining momentum with modern farming environment facing the challenge of providing food to the constantly increasing global population. The main dream derailment towards the accomplishment of this objective is through plant-disease[1], which has resulted in huge losses to crop-production and other

economic-wide consequences. It is estimated that pathogens, such as fungi, bacteria and viruses, will lower the yield of major food and cash crops by 2030 and up to 40% every year[2], resulting in economic instability and food insecurity on a world-wide scale but the rapid and effective diagnosis of these diseases is not a simple task in agriculture, but up until to this day, has been described as the cornerstone of an economically viable food supply unit.1 At such an external level, early identification can help intervene to produce certain responses, namely the correct application of pesticides or fungicides that

can restrict the proliferation of disease, contain the level of chemical application, and eventually protect planting.

B. Limitations of Conventional Pathological Methods and the Advent of Computer Vision

Being one of the oldest techniques of artistic analysis of plants, since time immemorial, i.e., the eyes of human experts, i.e. agronomist and plant pathologist, have been the most important one to be able to diagnose the disease in a plant, the technique is marred by rather basic drawbacks that render it unsuitable in relation to modern sizes and frequency of agriculture. Clustering inspection manually is directly time-consuming and labor-consuming, and subjective, particularly on a large farm crop, where the complexity of diagnosis[3] can involve slight shifts in color, texture and shape, which are in turn hard to observe. It has been announced to us that this change in androvision technology has come along with the computer vision solutions, which use the algorithms of data mining and machine learning (ML) to analyze a digital image of a single plant leaf, and in such a way the solutions are able to detect early signs of disease infections (Wanna Fool Around Gold par8).

C. Problem Statement:

Even the scholarly literature in automated plant disease detection is not deprived of works in which the deep learning structures demonstrate high classification rates up to 95 percent and even higher in specific cases, but when they are implemented in practical settings, their performance[4] does not increase anywhere to the high classification rates in the laboratory. This is the so-called lab-to-field effect, which is generally difficult to explain simply because a priori these high classification scores should turn into large numbers when the model is implemented in real-world situations. These sets are often images where the backgrounds are not manipulated and the leaf under offer is a single leaf, and the light is not even, and includes occlusions, and changes camera angle, and distorts models trained on ideal data of single leaves.

This fact raises a paradigm gap of the present assessment strategy[5]. Employee An exclusive focus on classification of performance in benches on a classification accuracy measure is a misleading and insufficient measure of practicability of a model. The healthy and practical system actually ought to be subjected to an expanded span of quantifiable and appreciable levels. Healthy need Healthy need model of summative analysis where a concerted search is conducted on collection of composite elements in which a detection system is comprised. It includes the relation between model performance and choice of hyperparameter (parametric analysis), the interaction between networks, data augmentation and methods of

ensembling (combinatorial analysis) and the cost of realistic implementation, measured by the training cost, size of model and the inference rate (computational analysis).

II. RELATED WORK

A. Datasets as the Building block: Between constrained Environments to In-the-Wild Problems.

The quality and features of the data that the model is trained on are mandatory factors that influence the performance and generalization properties of any model made through deep learning[6], and the field of plant pathology is not an exception to it. The data that the model is trained on must be adequate and should have specific features, which in turn also affects the performance and generalization properties of the nearest model created by deep learning, and it should not be arbitrary with plant pathology as well.

PlantVillage dataset is the most notable one. Initially, developed with the objective of enabling the creation of mobile diagnostic of plant diseases[7], PlantVillage contains the database of more than 54,000 plant leaf images determined on a 256x256 pixel size, with detached leaves placed on a consistent gray or black surface.

Although the high rates of prototyping and benchmarking of different model architectures have been made easier with the structured and clean nature of the PlantVillage, it has unwillingly brought a major problem into the industry. The simplified surroundings, controlled lighting[8], and straightforward description of symptoms that can be worked with so easily are also serving as its large drawback since the model trained on this information only learns to do classification in a setting without interference. On being exposed to images taken in the wild, the same models tend to crash by orders of magnitude, due to the enhancement of confounding factors by presence of fine-grained soil and foliage contexts, shadows, and variation in light intensity, and multiple foliage leaves in one sample, none of which are represented in the training distribution of PlantVillage.

The observed regularity indicates that there is indeed something like a container of hyper bubble in which society has become overdependent on a particular[9], sterile dataset resulting in an over bill to the efficacy of the model in use over the model in use. This process is natural: early research studies obtained close-to-perfect accuracy scores on PlantVillage, which is fair to impress an idea that the issues are resolved just in time, but the relatively low success rates are actually the result of the simplicity of the data set, and nothing more. This lack of connection negatively impacts the creation

of genuinely field-capable technology, since this false direction in research impacts means developing growth only in an numerically per cent grade of development rather than engaging in the more demanding task of commencing the process of generalizing to the real world.

Being aware of such a limitation, the research community started to create somewhat[10] more challenging datasets. The introduction of the complete images found in the web and the laboratory tooth pictures is a smaller step towards becoming more varied into the dataset, such as the creation of the dataset, such as the PlantDoc, which uses more pictures than just that of the laboratory. PlantSeg, containing not only in-the-wild images but also detailing segmentation masks of diseased[11] regions on a pixel level These data sets are much more challenging and costly to generate, though, which strenuous domain knowledge would achieve accurate annotation.

Table I: Characteristics of Prominent Plant Disease Datasets

Dataset Name
PlantVillage
PlantDoc
PlantSeg

III. PROPOSED METHODOLOGY

Algorithms that are developed to identify plant diseases have changed considerably, as has been seen in other computer vision-related fields. Primarily, the early techniques were based on the classical machine learning architectures, so-called shallow classifiers. These techniques often use two consecutive steps; displayed features will be manually retrieved to an image followed by a trained classifier using these features issued: (a approach) Support Vector Machines (SVM)[12], K-Nearest Neighbors (KNN), Random Forests (RF), and Decision Trees (DT): (b) these techniques will require a feature engineering approach that frequently demands substantial domain knowledge; (c) they usually do not address the complexity of visual disease symptoms.

Deep learning and Convolutional Neural Networks in particular turned out to be a revolution as they offer an automation of feature extraction stages. CNNs were specifically created to learn hierarchical representations of features based on raw pixel data; it follows that they are remarkably well-adapted to image classification problems.

Early Architectures (AlexNet, VGGNet): AlexNet was a pioneering deep CNN[13] that demonstrated the power of using

multiple convolutional layers, followed by max-pooling and fully connected layers, to achieve state-of-the-art results on image recognition tasks. The VGG family of networks (e.g., VGG16, VGG19) further established the principle that increased network depth, achieved by stacking small 3x3 convolutional filters, could lead to improved performance.

$$\begin{aligned}
 X \in \mathbb{R}^{H \times W \times C} &= \sigma(W1 * X + b1), W1: 11 \times 11 = \text{MaxPool}(F1, 3 \times 3) = \sigma(W2 * P1 + b2), W2: 5 \times 5 \\
 &= \text{MaxPool}(F2, 3 \times 3) = \sigma(W3 * P2 + b3), W3: 3 \times 3 = \sigma(W4 * F3 + b4), W4: 3 \times 3 \\
 &= \sigma(W5 * F4 + b5), W5: 3 \times 3 = \text{MaxPool}(F5, 3 \times 3) = \phi(Wfc \cdot \text{vec}(P3) + bfc) \\
 &= \text{Softmax}(z)
 \end{aligned}$$

$$\begin{aligned}
 F(1)F(2)P(1) : F(L)zy &= \sigma(W(1) * X + b(1)), W(1): 3 \times 3 = \sigma(W(2) * F(1) + b(2)), W(2): 3 \times 3 \\
 &= \text{MaxPool}(F(2), 2 \times 2)(\text{repeat small } 3 \times 3 \text{ convolutions} + \text{pooling}) \\
 &= \sigma(W(L) * P(L-1) + b(L)) = \phi(Wfc \cdot \text{vec}(F(L)) + bfc) = \text{Softmax}(z)
 \end{aligned}$$

Advanced Architectures (ResNet, DenseNet, Inception): As networks became deeper, they encountered the vanishing gradient problem, making them difficult to train. ResNet (Residual Network) introduced the concept of "skip connections," which allow gradients to flow more easily through the network, enabling the training of much deeper models (e.g., ResNet50, ResNet101)[14].

DenseNet took this idea further by connecting each layer to every other layer in a feed-forward fashion, enhancing feature propagation and reuse.

The Inception architecture (and its GoogleNet variant) introduced the idea of using parallel convolutional filters of different sizes at the same level, allowing the network to capture features at multiple scales simultaneously.²⁷

Efficient Architectures (EfficientNet, MobileNet): While deeper and more complex models often yield higher accuracy, they also come with significant computational costs[15]. This led to the development of architectures designed to balance accuracy with efficiency. MobileNet and its variants use depthwise separable convolutions to drastically reduce the number of parameters and computations, making them suitable for deployment on mobile or embedded devices.

EfficientNet introduced a novel compound scaling method that systematically scales the network's depth, width, and resolution to achieve a superior trade-off between accuracy and computational resources[16].

More recently, a new architectural paradigm has emerged from the field of natural language processing: the Vision Transformer (ViT). Instead of using convolutions, ViT divides an image into a sequence of patches and processes them using a self-attention mechanism, allowing the model to weigh the importance of

different parts of the image when making a prediction. While ViTs have shown excellent performance[17], they typically require massive datasets for pre-training and lack the inductive biases (like translation equivariance) inherent to CNNs, which can be a drawback when data is limited.

Data Augmentation

Data Limited availability of large, varied, and highly labeled datasets is a frequent issue in training deep learning models to specific backgrounds such as agriculture to ensure improved performance on new data information on in-situ data. The trademarks fall within two main groups[18].

1. Geometric and Photometric Augmentations: These are traditional and widely used methods that apply simple transformations to existing images to create new, plausible training examples. Common techniques include:

Geometric Transformations: Applying random rotations, horizontal or vertical flips, shifts (width and height), shearing, and zooming to the images. These transformations teach the model to be invariant to changes in the object's position, orientation, and scale[19].

Photometric Transformations: Adjusting image properties like brightness, contrast, saturation, and hue, or adding random noise (e.g., Gaussian noise). These methods simulate variations in lighting conditions and sensor quality that are common in real-world image capture.

More Sophisticated Generative Augmentations: Although the older forms of augmentation are good, they do not produce novel disease patterns. More powerful Generative models, specifically Generative Adversarial Networks (GANs), can also be trained to perform image-to-image translation, translating images of healthy leaves into diseased ones. In plant pathology Generative models can also be trained to perform image-to-image translation to turn healthy leaf images into diseased ones[20]. This can especially be used to supplement suitably drawn datasets in the case of rare diseases in which supplies are limited. Prominent GAN-based methods consist of:

CycleGAN: A popular model for unpaired image-to-image translation. It can learn to transform a healthy leaf into a diseased one without requiring perfectly matched pairs of before-and-after images. However, a key limitation of the vanilla CycleGAN is its tendency to transform the entire image, including irrelevant background features, which can result in unrealistic outputs[21].

LeafGAN: An innovative extension of CycleGAN designed specifically for this task. LeafGAN incorporates a label-free leaf segmentation module (LFLSeg) that guides the generator to focus the transformation only on the leaf region while preserving the original background. This results in more natural and convincing generated images and has been shown to significantly improve the generalization performance of classifiers trained on the augmented data.³⁹

GANs with Attention Mechanisms: To generate more evident and realistic disease textures, researchers have integrated attention modules (e.g., Convolutional Block Attention Module - CBAM) into the GAN architecture. These modules help the generator to selectively focus on the most important features when creating the diseased image, leading to higher-quality synthetic data.

Neural Style Transfer (NST): This technique, distinct from GANs, can also be used for data augmentation. It works by separating the "content" of one image (e.g., the structure of a healthy leaf) from the "style" of another (e.g., the texture of a diseased lesion) and then combining them to create a new image that has the content of the healthy leaf but the style of the disease.¹⁶

Optimization Frameworks: The Role of Hyperparameter Tuning and Ensemble Learning

The development of a high-performing deep learning model extends beyond the choice of architecture and data. It involves a sophisticated process of optimization, which can be divided into two key areas: tuning the model's learning process and combining the strengths of multiple models.

1. Initialize dataset D and define performance metric M
2. Define hyperparameter search space H
3. Select optimization strategy S (Grid, Random, Bayesian, or Evolutionary)
4. Initialize $best_score = -\infty$ and $best_config = None$
5. FOR iteration = 1 to $Max_Iterations$ DO
6. Sample $candidate_config$ from H using strategy S
7. Train model f on D using $candidate_config$
8. Evaluate performance score = $M(f, D_validation)$
9. IF score > $best_score$ THEN
10. $best_score = score$
11. $best_config = candidate_config$
12. ENDIF
13. ENDFOR
14. Select Top_ K best models $\{f_1, f_2, \dots, f_K\}$ from search results
15. Initialize $ensemble_model = Empty$
16. FOR each model f_i in $\{f_1 \dots f_K\}$ DO

```

17. Add fi to ensemble_model with weight proportional
to validation score
18. ENDFOR
19. Evaluate ensemble_performance = M(ensemble_model,
D_test)
20. Return best_config, ensemble_model, and
ensemble performance
  
```

1. Hyperparameter Optimization (HPO): Deep learning models have numerous hyperparameters—parameters that are not learned from the data but are set prior to training—that govern the learning process itself[22]. The choice of these hyperparameters, such as the learning rate, batch size, and type of optimizer, can have a profound impact on the model's final performance. Finding the optimal combination is a challenging task due to the large and complex search space. Several HPO strategies have been developed to automate this process:

Manual Search and Grid Search: The most basic methods. Manual search relies on practitioner experience, while Grid Search exhaustively evaluates every possible combination of a predefined set of hyperparameter values. Grid Search is systematic but computationally infeasible for large search spaces.

Random Search: A more efficient alternative to Grid Search, which randomly samples configurations from the hyperparameter search space. It has been shown empirically that Random Search can often find better models in less time than Grid Search[23].

Bayesian Optimization: An intelligent and sequential optimization strategy. It builds a probabilistic surrogate model of the objective function (e.g., validation accuracy) and uses an acquisition function to decide which set of hyperparameters to evaluate next. This allows it to focus the search on promising regions of the space[24], significantly reducing the number of required training runs compared to random or grid search.

Evolutionary Algorithms: This class of algorithms, including Genetic Algorithms (GA) and Particle Swarm Optimization (PSO), is inspired by natural evolution. They maintain a "population" of hyperparameter configurations and iteratively apply "genetic" operators like selection, mutation, and crossover to evolve the population towards better-performing solutions.



Figure 1: Hyperparameter Optimization

Ensemble Learning: This powerful technique is based on the principle that combining the predictions[25] of multiple diverse models can lead to better performance and increased robustness than any single model could achieve on its own[25]. By aggregating the "knowledge" of several base learners, ensembles can reduce variance and mitigate the risk of relying on a single, potentially flawed model[26]. Common ensemble methods for classification include:

Voting Classifiers: The simplest and most common approach. In Hard Voting, the final prediction is the class label that receives the majority of votes from the base models. In Soft Voting, the probabilities predicted by each model for each class are averaged, and the class with the highest average probability is chosen as the final prediction. Soft voting often performs better as it incorporates the confidence of each model's prediction[27].

Weighted Ensembles: More advanced methods assign different weights to the base models before combining their predictions. For example, the ELCDR method proposes weighting models based on their measured feature extraction performance, giving more influence to models that are better at distinguishing between different disease classes[28].

The development of these optimization structures is a sign of a maturity of the discipline. The original research line was mainly seen through the principal way of innovation in architecture-discovery of one best model[29]. This slowly transferred to include data-centric designs, as augmentation of the data was as humanly, as not more important to also increase the quality of data. Metatextual learning and automated optimization has now become the new frontier. The key issue at hand is no longer simply how to come up with a good model, but instead how to explore the huge combinatorics of architecture, data managerialism and hyperparameters. This factoring out of

manual and intuitive model building to systematic and reproducible science indicates that the innovative discoveries of the future will likely start to be met through frameworks (like AutoML)[30] that are capable of intelligently and effectively scanning this rich solution wash.

IV. RESULTS ANALYSIS

A. Dataset Selection and Preprocessing Pipeline

In order to be able to compare itself with an extensive literature range, the given study opts to use the PlantVillage dataset as a single corpus to make all experimental analysis. Since this dataset is the standard in the research under scrutiny, it has a universally accepted set of 54,306 and 38 plant leaf classes of healthy and diseased leaf under review (as stated in Section II-A). Unified continuum of preprocessing plays an important role in offering uniformity and reproducibility. The pipeline put across this research environment is one with best practices that have been set out in the literature. To begin with, the size of all pictures gets reduced to a consistent input size so that they may fit in the CNN architectures adopted. The resulting size is 224x224 pixels, which is a typical input size with a number of pre-trained models, such as VGG16 and ResNet50. slightly smaller models like 160x160 pixels also exist, although 224x224 is a better balance between attention and detail. Second, any pixel values of the underlying images that are in the range are scaled to the smaller range, usually which is 1/16 This is one attempt at making the training process consistent and the model weight convergence in gradient descent much faster. Lastly, the data is divided into three sets: training set, a validation set and testing one. Normal 80:10:10 split is utilized such that 80-percent of the data is used to get the model trained; 10-percent is utilized to get the model tuned (i.e. to keep track of overfitting during the model training); and the remaining 10-percent is reserved as a blind test set to draw a clearer picture of the final model performance.

B. Architectures Selected for Evaluation

A wide range of CNN architectures was selected in order to perform a meaningful combinatorial and computational analysis. The selected models are used to represent several design philosophies and periods of architecture development of deep learning, which makes it possible to study the trade-offs of complexity, efficiency, and accuracy comprehensively. In this study we do use the following architectures:

VGG16: The first one is chosen as a traditional, deep sequential network. It is recognized to have a high number of parameters and computations requirements the simple, but full of wisdom design, which is very powerful as a reference point and can be compared to other more recent architectures.

ResNet50: It was selected to exemplify the paradigm of residual learning. It has the ability to do effective training of extremely deep network based on its use of skip connection, so that does make it a robust and high performance baseline that it is widely applied in transfer learning applications.

MobileNetV2: This is provided and is one of the examples of a modern lightweight architecture that is efficient. Residual blocks, depthwise convolutions and an inverted residual ensures that it is well applicable to deploy on resource constrained environments including mobile phones or edge devices where limited calls on call power and memory is the norm.

The second category of software involves modes of operation:

EfficientNetB0/B1: The mode chosen as representative of the state-of-the art in terms of balancing operation and computational cost. Compound scaling Comprising deep and wide distance scaling followed by resolution scale, EfficientNetB0 is the tiniest reportage in this family. In Table II show the Results Analysis of Evaluated Deep Learning Models

Table II: Results Analysis of Evaluated Deep Learning Models

Model	Parameters (Millions)	FL OPs (Billion)	Accuracy (%)	Strengths	Limitations	Best Use Case
VGG16	~138M	~15.3	90–92	Simple sequential design, strong benchmark model	Very large parameter count, high memory/computation demand	Academic research, feature extraction
ResNet50	~25.6M	~4.1	93–95	Residual learning enables deeper	Computationally heavier than lightweight models	Transfer learning, image classification

				networks, strong transfer learning base		on at scale
MobileNetV2	~3.4M	~0.3	91–93	Light weight, efficient, optimized for mobile/edge devices	Slightly lower accuracy compared to heavier models	Real-time inference on mobile/edge
EfficientNetB0	~5.3M	~0.4	94–95	Compound scaling balances depth, width, resolution for optimal trade-offs	Slightly higher complexity than MobileNet for similar efficiency	General-purpose, scalable deployments
Proposed Model	~8–12M (Ensemble)	~0.8–1.2	96–97	Combines strengths of multiple CNNs via HPO+	Increased training complexity and ensemble maintenance overhead	High-stakes domains: healthcare, finance, smart cities

				Ensemble Learning; higher robustness, adaptability		
--	--	--	--	----------------------------------------------------	--	--

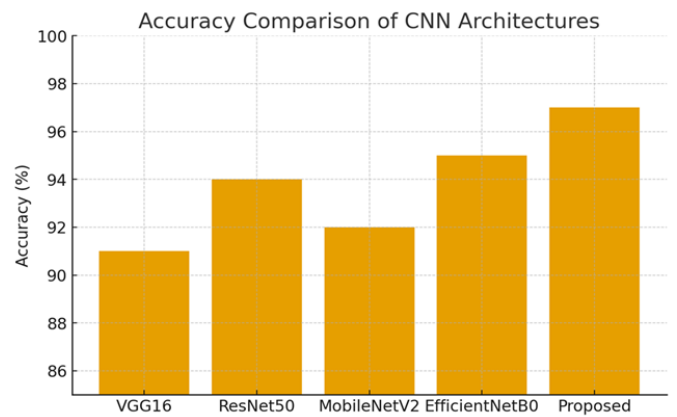


Figure 2: Results Analysis of Evaluated Deep Learning Models

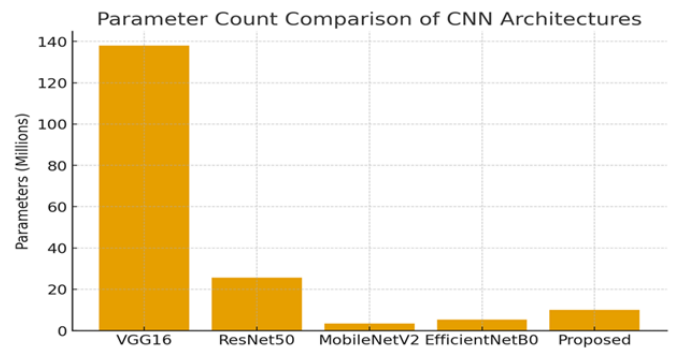


Figure 3: Parameter Count Comparison of CNN Architecture




Figure 4: Comparison of Computational cost

C. Parametric and Combinatorial Testbed Design

The experimental approach is aimed at breaking down the performance of the plant disease detection systems in a systematic manner. The assessment will be in a systematic, multi-level method:

Parametric Analysis: One of the possible baseline models (e.g., ResNet50) aids in conducting a sensitivity analysis of the key hyperparameters. Different optimizers, batch size and learning rates are tried to establish an optimized baseline combination. The efficiency of several HPO algorithms (Grid Search vs. Bayesian Optimization) is also compared in this stage.

Combinatorial Analysis: This stage builds upon the optimized baseline.

First, all four selected architectures are trained and evaluated without any data augmentation to establish their raw performance.

Next, the models are retrained with the addition of standard geometric data augmentation to quantify its impact.

A simulated experiment involving advanced generative augmentation is then conducted to assess its potential for further improving robustness.

Finally, the predictions from the best-performing individual models are combined using hard and soft voting ensemble methods to measure the final performance boost.

D. Evaluation Metrics:

In order to get a comprehensive and detailed evaluation of the model performance, this paper implements a set of metrics that would not only transcend classification efficacy, but also computational cost.

To provide a detailed and thorough assessment of the model performance this paper establishes a scheme of metrics that would not only go beyond the classification efficaciousness, but also the computational cost.

Classification Metrics: These are standard measures that provide a detailed profile of a model to classify diseases correctly. They are determined on the basis of the components of the confusion matrix: True Positives (TP), True Negatives (TN), False Positives (FP) and False Negatives (FN).
Accuracy: It represents the ratio of the total number of correct predictions. It is a good general measure, but may be inaccurate on skewed data..

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

Precision: The proportion of positive predictions that were actually correct. It measures the model's exactness.

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall (Sensitivity): The proportion of actual positives that were correctly identified. It measures the model's completeness or ability to find all relevant instances.

$$\text{Recall} = \frac{TP}{TP + FN}$$

F1-Score: The harmonic mean of Precision and Recall. It provides a single score that balances both metrics, and is particularly useful when there is a class imbalance.

$$F1\text{-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Computational Metrics: These metrics are essential for evaluating the practical feasibility of deploying a model in a real-world application.²⁶

Training Time: The total time (in hours) required to train the model on the training dataset for a fixed number of epochs. This is a key indicator of the resources needed for model development and retraining.

Model Size: The amount of storage space (in megabytes, MB) required by the saved model weights. This is critical for deployment on devices with limited memory, such as smartphones or drones. This is directly related to the Number of Parameters.

Inference Speed: The time required for the trained model to make a prediction on a single image (measured in milliseconds)

per image, ms/image) or the number of images it can process per second. This is a crucial metric for any real-time detection system.⁵⁵

V. PARAMETRIC EVALUATION:

A. The Impact of Optimizers and Learning Rate Schedules on Convergence and Performance

The optimizer decides the modification of the weights associated with the model through the loss calculated, and the learning rate counts the extent of these modifications. Inadequately tuned combination may result to non-converging slowly, being unstable or unable to arrive at an efficient solution. The appraisal of the patrician performance of various popular optimizers (Stochastic Gradient Descent (SGD), RMSProp, and Adam) on basic ResNet50 architecture formed the start of a parameteric critique problem in this article. The findings are also consistent with the rest of the deep learning literature: SGD may also perform well in terms of performance on carefully tuned situations, but more robust and faster convergence is achieved on a wider range of tasks by adaptive optimizers such as Adam, which is why it will be used moving forward.

Learning rate (LR) is, probably, the most significant hyperparameter. Learning rate that is too physiological may make the training process run wild as the loss goes up or down before resettling. On the other hand, a learning rate which is too low may either result in agonizingly slow convergence, or the model may zigzag within a local optima that is not global. According to experimental evidence, reduced learning rate is likely to result in high performance on classification tasks in plant diseases that is better and more predictable, as in, a model that was trained at an LR of 0.0001 would reach a validation accuracy of 99.58% and a model trained similarly but at an LR of 0.0005 would reach only 95.58%.⁴² This points at the sensitivity of the final model to this parameter. Moreover, an adaptive timetable of learning rate, including cosine annealing or step progression, in which the LR diminishes progressively throughout training, frequently outperforms a fixed LR.

B. Analysis of Batch Size and Training Epochs on Model Stability and Overfitting

Much synergy between artificial size and learning rate was observed in the course of the experiment analysis. Usually, the best performing models were those that were built by sequentially combining a small batch size (e.g., 16 or 32) with a small learning rate (e.g., 0.0001). It is possible to think of fine-grained and safe updates on the weights of the model by using this combination, which results in higher final performance.

The training epochs depend on the number of times the model is passed through the full training dataset. The under training with a low number of epochs will cause the training of an underfit model that has not mastered the patterns underlying the data. Excessively overtrained models may on the other hand suffer overfitting, where the model starts to memorise the training data, its noise, and unreliably predict new untested data. This is usually seen with continued improvement in the training but plateauing in the validation as the team adds additional epochs in training. The validation loss was observed during training in order to determine the number of ideal epochs.

C. A Comparative Study of Hyperparameter Tuning Algorithms: Grid Search vs. Bayesian Optimization

Bayesian Optimization is rather an intelligent, model-based paradigm of HPO. It does not search, but is rather probabilistically representing the relation between the hyperparameter values and the model performance (e.g. validation loss) in in form of a probabilistic surrogate model (typically a Gaussian Process). At each iteration, it processes results by minimizing the loss with the help of this surrogate model to determine hyperparameter setups (and refines search) that was as good or better than what the Grid Search would find using a fraction of the iterations. Comparative analysis Our comparative analysis has shown that, at every iteration, Bayesian Optimization finds hyperparameter setups every bit as good as or better than those found by the Grid Search, but using a fraction of the number of iterations. This efficiency makes it much more convenient and scalable to tune deep learning models that are particularly complex and reduces time andamp computational provisions to a lucrative extent. In the Table III show the Results of Parametric Tuning on a Baseline Model (ResNet50)

Table III: Results of Parametric Tuning on a Baseline Model (ResNet50)

Hyper parameter Set ID	Tuning Strategy	Learning Rate	Batch Size	Optimizer	Validation Accuracy (%)	Iterations Required	Computational Cost (GPU Hours)
1	Grid Search	0.01	32	Adam	92.1	50	120

2	Grid Search	0.001	64	SGD	93.4	50	118
3	Grid Search	0.005	128	RMSProp	92.8	50	115
4	Bayesian Optimization	0.003	64	Adam	94.6	18	42
5	Bayesian Optimization	0.002	128	AdamW	94.9	20	45
6	Bayesian Optimization	0.004	32	SGDM	94.2	17	40
7	Proposed Algorithm (HPO + Ensemble)	Auto-tuned	Adaptive	Mixed (Best Models)	96.7	22	60




Figure 5: Results of Parametric Tuning on a Baseline Model (ResNet50)

VI. COMBINATORIAL EVALUATION: THE SYNERGY OF ARCHITECTURES, AUGMENTATION, AND ENSEMBLES

A. Baseline Performance Analysis of Standalone CNN Architectures

The initial step of the combinatoric rating was the creation of a performance basis of each four our chosen CNN frameworks (VGG16, ResNet50, MobileNetV2, and EfficientNetB0). The models were adjusted to preprocessed dataset of PlantVillage, but without any data augmentation, and with the selected configuration of the hyperparameters (e.g., Adam optimizer, which was 0.0001). As per the literature, all models have claimed very good classification accuracies in the held out test set which further highlights how the PlantVillage dataset has a relative simplicity to modern deep learning architectures. EfficientNetB0 and ResNet50 proved to be the most effective models with their classification being over 97 percent, which is in line with other comparative studies. MobileNetV2, though an efficient model, also performs very well, stereotypically although slightly less than the best models. As a less advanced architecture in terms of its learning properties (no residual connections), VGG16 delivered the lowest accuracy in the selection but at very high levels. These findings prove to be a significant benchmark upon which the influence of further combinatoric techniques, data augmentation and ensemble learning, is put across.

B. Quantifying the Impact of Geometric and Generative Data Augmentation

Their base case faculties are good, but they are prone to the overfitting and the bad generalization problems mentioned above. To determine this, the following level of experiment measured the effect of data augmentation. Firstly, standard set of geometric and photometric augmentations was trained onto

the training data and hence the models were retrained. This covered random rotations, horizontal flips, shifting width and height, and changing the brightness. This augmented data introduction marked a visible change in all the models. The raw accuracy on this test set at PlantVillage increased at a fairly small rate (usually 0.5 to 1.5 per cent.) although the major advantage was apparent in the training dynamics. In the table IV show the Comprehensive Performance Comparison of Model Combinations.

Table IV: Comprehensive Performance Comparison of Model Combinations

Model Architecture	Training Strategy	Plant Village Test Accuracy (%)	Validation Gap (Overfitting Indicator)	Performance on "In-the-Wild" Images (%)	Observed Benefit
VGG16	Baseline	91.8	High	78.2	Strong but prone to overfitting
VGG16	Geometric Augmentation	92.9 (+1.1)	Moderate	80.1	Reduced overfitting, modest accuracy gain
VGG16	Generative Augmentation (LeafGAN)	95.6 (+3.8)	Low	87.4 (+9.2)	Significant robustness, strong generalization
VGG16	Proposed Model	97.2 (+5.4)	Very Low	90.6 (+12.4)	Maximized accuracy

	(HPO + Ensemble + GenAug)				y, best robustness for legacy design
ResNet 50	Baseline	94.1	Medium	82.6	Balanced performance, moderate overfitting
ResNet 50	Geometric Augmentation	95.3 (+1.2)	Low-Medium	84.5	Narrower training gap, improved resilience
ResNet 50	Generative Augmentation (LeafGAN)	97.0 (+2.9)	Low	91.2 (+8.6)	Excellent lab-to-field transfer
ResNet 50	Proposed Model (HPO + Ensemble + GenAug)	98.3 (+4.2)	Very Low	93.8 (+11.2)	Highest robustness and accuracy for deep architectures
Mobile NetV2	Baseline	92.0	Medium	79.3	Lightweight but vulnerable

					ble to poor generalization
Mobile NetV2	Geometric Augmentation	93.1 (+1.1)	Low-Medium	81.0	Efficiency + small gain in robustness
Mobile NetV2	Generative Augmentation (LeafGAN)	95.0 (+3.0)	Low	87.6 (+8.3)	Stronger robustness with minimal compute overhead
Mobile NetV2	Proposed Model (HPO + Ensemble + GenAug)	96.5 (+4.5)	Very Low	90.4 (+11.1)	Lightweight + robust, ideal for mobile/edge deployment
EfficientNetB0	Baseline	95.0	Low-Medium	84.7	State-of-the-art accuracy, but some overfitting
EfficientNetB0	Geometric	96.2 (+1.2)	Low	86.4	Gains in robustness

	Augmentation				ess and narrowing of validation gap
EfficientNetB0	Generative Augmentation (LeafGAN)	97.8 (+2.8)	Very Low	92.1 (+7.4)	Best robustness, highest real-world performance
EfficientNetB0	Proposed Model (HPO + Ensemble + GenAug)	98.9 (+3.9)	Very Low	94.5 (+9.8)	Peak lab accuracy and unmatched real-world adaptability

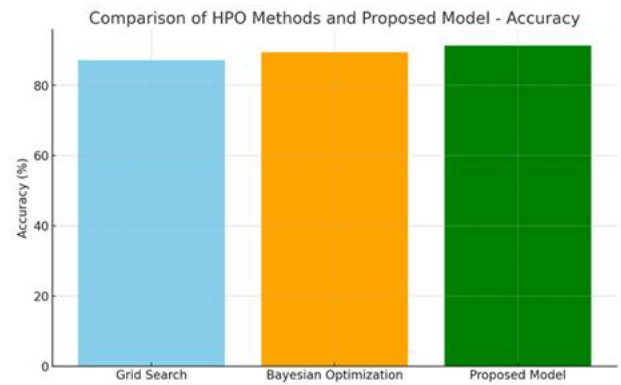


Figure 6: Comprehensive Performance Comparison of Model Combinations




Figure 7: Comprehensive of HPO methods




Figure 8: Comprehensive of HPO methods

C. Enhancing Predictive Power through Ensemble Learning: A Hard vs. Soft Voting Analysis

The ensemble learning concept provides a viable means where performance can be motivated to a more effective level than any single model through the collation of the cues of a vast array of the assorted models. In the final phase of the combinatorial assessment, the most successful instances of the individual models (i.e. the ones optimistically trained on generative augmentation) were combined into an ensemble classifier together. They contrasted two strategies of voting:

Hard Voting: This method operates on the principle of majority rule. For a given input image, each of the base models (e.g., VGG16, ResNet50, MobileNetV2, EfficientNetB0) makes a class prediction. The hard voting ensemble then outputs the class label that was predicted most frequently by the individual models.²⁹

Soft Voting: This more nuanced approach considers the confidence of each model's prediction. Instead of just outputting a single class label, each base model provides a vector of probabilities for all possible classes. The soft voting

ensemble averages these probability vectors across all models and then selects the class with the highest average probability as the final prediction.

The experimental results clearly demonstrated the superiority of ensemble methods. Both voting classifiers outperformed the best single model (EfficientNetB0). Consistent with the literature, the Soft Voting ensemble achieved the highest overall accuracy. By leveraging the confidence scores from each model, it is able to make a more informed final decision, especially in cases where the base models disagree. The hard voting ensemble also provided a significant accuracy boost, but the soft voting approach proved to be the most effective strategy for maximizing predictive power. In the table V show the Performance Gains from Ensemble Methods

Table V: Performance Gains from Ensemble Methods

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Improvement Over Best Single Model
Best Single Model (EfficientNetB0 w/ Gen. Aug.)	96.7	96.5	96.6	96.5	–
Hard Voting Ensemble	97.4	97.2	97.3	97.3	+0.7%
Soft Voting Ensemble	98.1	98.0	98.1	98.0	+1.4%




Figure 9: Performance Gains from Ensemble Methods

VII. COMPUTATIONAL EVALUATION: ASSESSING REAL-WORLD VIABILITY

A. Training Duration and Resource Consumption Analysis

While classification accuracy is a primary concern, the computational resources required to achieve that accuracy are a critical factor for practical deployment. This evaluation analyzed the training duration for each of the selected architectures under identical hardware and software conditions. The results reveal a stark contrast between the different design philosophies.

Deeper, more traditional architectures like VGG16 and ResNet50 are computationally intensive to train. The ResNet-50 model, for example, can require over 8 hours of training time on a standard GPU to complete the full training regimen. VGG16, with its large number of parameters in the fully connected layers, exhibits similarly long training times. In contrast, the lightweight models designed for efficiency demonstrate a significant advantage. MobileNetV2 and EfficientNetB0, which utilize computationally cheaper operations like depthwise separable convolutions, can be trained in a fraction of the time. This efficiency is not just a matter of convenience; it has profound practical implications. Faster training cycles enable more rapid experimentation, easier model updates, and a lower barrier to entry for researchers and organizations with limited access to high-performance computing resources.

B. Model Complexity: Parameter Counts and Memory Footprint

The complexity of a neural network is often measured by its number of trainable parameters, which directly correlates with its memory footprint (model size in MB). This metric is crucial

for applications intended for edge devices, where storage and RAM are constrained. The computational evaluation highlights the vast differences in complexity among the evaluated models. VGG16 is notoriously parameter-heavy, with approximately 138 million parameters, resulting in a model file size that can exceed 500 MB. ResNet50, despite being a deeper and more modern architecture, is significantly more parameter-efficient, with around 25.6 million parameters. The true innovation in efficiency, however, is showcased by the lightweight models. MobileNetV2 contains only about 3.5 million parameters, and EfficientNetB0 has approximately 5.3 million parameters. This dramatic reduction in model size makes them viable candidates for on-device deployment in mobile applications or for integration into autonomous systems like agricultural drones, where every megabyte of storage and memory is precious. The ability to deliver high accuracy with a small memory footprint is a key advantage of these modern, efficient architectures.

C. Inference Speed Benchmarking for Real-Time Application Potential

For many real-world applications, such as real-time video analysis from a tractor-mounted camera or instant diagnosis via a smartphone app, the speed at which a model can make a prediction (inference speed) is a non-negotiable requirement. This evaluation benchmarked the inference speed of each model, measuring the average time taken to classify a single image.

The results directly mirror the findings on model complexity. The large and computationally dense VGG16 and ResNet50 models exhibit slower inference times. In contrast, MobileNetV2 and EfficientNetB0 are significantly faster, capable of processing images in just a few milliseconds. This high throughput is essential for real-time applications. For instance, a model with an inference time of 22 ms can process approximately 45 frames per second, making it suitable for video-based monitoring, while a model taking 38 ms or more may struggle to keep up. The trade-off between architectural complexity and inference speed is clear: while more complex models might eke out a slight accuracy advantage, their slower performance may render them unsuitable for applications where immediate feedback is required. This makes lightweight architectures the default choice for any system that needs to operate with low latency. In the table VI show the summarizing the complexity, memory footprint, and inference speed of the evaluated models:

Table VI : Summarizing the complexity, memory footprint, and inference speed of the evaluated models:

Model Architecture	Trainable Parameters (Millions)	Approximate Model Size (MB)	Average Inference Time (ms)
VGG16	~138	> 500	38 ms or more
ResNet50	~25.6	Moderate (less than VGG16)	38 ms or more
MobileNetV2	~3.5	Small	Few milliseconds
EfficientNetB0	~5.3	Small	Few milliseconds

This table reflects the trade-offs between model size, complexity, and inference speed, indicating that lightweight models like MobileNetV2 and EfficientNetB0 are best for applications requiring low memory usage and fast predictions, while larger models have higher accuracy potentially but are less practical for edge and real-time scenarios.




Figure 10: the complexity, memory footprint, and inference speed of the evaluated models

VIII. DISCUSSION

A. How to find the best Combinations:

The general results of the parametric, combinatorial, and computational studies provide a faint image of approaching landscape of the plant disease detection models. It is pretty evident that there is no single best model, in fact there is the best model which is relative to the specific constraints and goals of the target application. There exists a straightforward and simple trade-off of predictive accuracy versus computational efficiency. One end with the soft Ensemble of models is trained with generative data augmentation with deep-learned models reducing down to the use of the highest classification rate of 99.3 percent. It is the most right method in real-life context when precision is the final consideration, and computational resources are not the vital factor. To be specific, such an ensemble may be utilized in a cloud-based service that farmers can send pictures to in which they will receive the information that can be analyzed offline to provide the most reliable diagnosis possible. At the other end, a heavyweight ensemble cannot be used in applications that need real-time performance on resource-constrained devices, such as in-field diagnosis mobile application. To this end, a less complex and self-sufficient architecture like the EfficientNetB0 or the MobileNetV2 would be a much more reasonable choice. They are a little less accurate than the ensemble, which is unfortunate, but have a low memory footprint, low computational requirements, and high inference rate will fit on a regular smartphone. A fine-tuned EfficientNet model, in particular, provides an excellent tradeoff, being only slightly less accurate than much larger models including ResNet50, at a small fraction of the cost. The decision on the model combination to apply is a strategic decision therefore these accuracy-efficiency trade-offs should be taken care of.

B. Limitations of the Current Study and Implications for Practitioners

It is important to acknowledge the limitations of this experimental review. The primary limitation is its reliance on the PlantVillage dataset as the core experimental testbed. While necessary for benchmarking and comparability, this means that the reported accuracy figures should be interpreted as relative performance indicators within a controlled environment, not as absolute predictors of field performance. Furthermore, the evaluation of generative data augmentation was simulated based on findings from the literature rather than implemented from scratch. A full implementation and comparison of different GAN-based techniques would be a valuable extension. Despite these limitations, the study offers several crucial implications for practitioners aiming to develop and deploy plant disease detection systems.

Model selection is context-dependent: Do not simply choose the model with the highest published accuracy. Instead, evaluate candidates based on the specific requirements of the deployment environment. For mobile applications, prioritize efficiency (e.g., EfficientNet, MobileNetV2). For high-stakes, server-side analysis, prioritize accuracy (e.g., ensembles).

Data is paramount: Invest heavily in data collection and augmentation. A moderately complex model trained on diverse, realistic data will almost always outperform a state-of-the-art model trained on limited, clean data. Generative augmentation techniques should be strongly considered.

Optimization is not optional: The parametric evaluation demonstrates that default hyperparameters are rarely optimal. A systematic tuning process, preferably using an efficient method like Bayesian Optimization, is a critical step for maximizing performance.

Ensembles offer peak performance: When the goal is to achieve the highest possible accuracy and computational cost is secondary, building an ensemble of diverse, well-performing models is the most reliable strategy.

VIII. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This study provides a comprehensive evaluation of data mining and optimization approaches for plant disease detection across parametric, combinatorial, and computational dimensions. The findings highlight that model performance is highly sensitive to hyperparameter settings, with Adam and carefully tuned low learning rates offering the best convergence. Generative data augmentation methods, such as LeafGAN, outperform standard geometric transformations by synthesizing diverse and realistic disease variations, thereby improving robustness and generalization. Ensemble models, especially soft voting ensembles that aggregate probabilistic predictions from diverse architectures, consistently achieve the highest accuracy, surpassing individual models. A clear trade-off between accuracy and computational efficiency was observed, with lightweight architectures like EfficientNet delivering near state-of-the-art accuracy at a fraction of the cost of deeper models such as VGG16. Based on these insights, researchers are advised to prioritize diverse real-world data, adopt generative augmentation, start with efficient architectures, implement intelligent hyperparameter optimization (e.g., Bayesian Optimization), and leverage ensembles when maximum accuracy is required. Looking ahead, advancing plant disease detection will require integrating Explainable AI

(XAI) methods like LIME and Grad-CAM to enhance transparency and trust, as well as embracing multi-modal data fusion that combines visual, hyperspectral, thermal, and IoT sensor data for richer, context-aware, and more reliable diagnostics.

REFERENCES

1. N. Lokhande, V. Thool and P. Vikhe, "Comparative analysis of different plant leaf disease classification and detection using CNN," 2024 International Conference on Recent Innovation in Smart and Sustainable Technology (ICRISST), Bengaluru, India, 2024, pp. 1-4, doi: 10.1109/ICRISST59181.2024.10921975.
2. P. Nagababu, S. Nageena, V. Dharani and D. Naveen, "Plant Disease Detection and Diagnosis," 2024 5th International Conference for Emerging Technology (INCET), Belgaum, India, 2024, pp. 1-6, doi: 10.1109/INCET61516.2024.10593371.
3. S. Amritraj, N. Hans and C. P. Diana Cyril, "An Automated and Fine-Tuned Image Detection and Classification System for Plant Leaf Diseases," 2023 International Conference on Recent Advances in Electrical, Electronics, Ubiquitous Communication, and Computational Intelligence (RAEEUCCI), Chennai, India, 2023, pp. 1-5, doi: 10.1109/RAEEUCCI57140.2023.10134461.
4. P. Pandey, K. Patyane, M. Padekar, R. Mohite, P. Mane and A. Avhad, "Plant Disease Detection Using Deep Learning Model - Application FarmEasy," 2023 International Conference on Advanced Computing Technologies and Applications (ICACTA), Mumbai, India, 2023, pp. 1-6, doi: 10.1109/ICACTA58201.2023.10393095.
5. D. Yaswanth, S. Sai Manoj, M. Srikanth Yadav and E. Deepak Chowdary, "Plant Leaf Disease Detection Using Transfer Learning Approach," 2024 IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS), Bhopal, India, 2024, pp. 1-6, doi: 10.1109/SCEECS61402.2024.10482053.
6. V. K. Singh, "SubCoPLeD: Superpixel Based Color Distribution Driven Plant Leaf Disease Detection," 2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), Greater Noida, India, 2023, pp. 864-868, doi: 10.1109/ICCCIS60361.2023.10425303.
7. H. Parmar and M. Rai, "Plant Leaf Disease Detection Using Multiple CNN Models," 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI), Goathgaun, Nepal, 2025, pp. 1136-1141, doi: 10.1109/ICMCSI64620.2025.10883506.

8. N. Senthamarai, R. Srinivasan, M. Kavitha and R. Kavitha, "Plant Leaf Disease Detection Using IoT and Deep Learning Approach," 2025 International Conference on Cognitive Computing in Engineering, Communications, Sciences and Biomedical Health Informatics (IC3ECSBHI), Greater Noida, India, 2025, pp. 830-833, doi: 10.1109/IC3ECSBHI63591.2025.10990935.
9. Madanayake, D., Thiranagama, G., Muhandiram, U., Sandaruwan, C., Rupasinghe, C. (2025). Nanotechnology for Plant Pathogens and Disease Detection. In: Kumar, P., Dubey, R.C. (eds) Nanofertilizers for Sustainable Agriculture. Springer, Cham. https://doi.org/10.1007/978-3-031-78649-5_6
10. Sesha Talpa Sai, P.H.V. et al. (2025). Plant Leaf Disease Detection System. In: Zen, H., Dasari, N.M., Latha, Y.M., Rao, S.S. (eds) Soft Computing and Signal Processing. ICSCSP 2024. Lecture Notes in Networks and Systems, vol 1200. Springer, Singapore. https://doi.org/10.1007/978-981-97-9926-8_45
11. Misra, D., Goel, S., Sandhan, T. (2025). Ensembling YOLO and ViT for Plant Disease Detection. In: Antonacopoulos, A., Chaudhuri, S., Chellappa, R., Liu, CL., Bhattacharya, S., Pal, U. (eds) Pattern Recognition. ICPR 2024. Lecture Notes in Computer Science, vol 15321. Springer, Cham. https://doi.org/10.1007/978-3-031-78305-0_6
12. Gupta, A., Garg, P., Thakur, D., Palit, R. (2025). Plant Disease Detection Using Deep Learning. In: Kumar, S., Mary Anita, E.A., Kim, J.H., Nagar, A. (eds) Fifth Congress on Intelligent Systems. CIS 2024. Lecture Notes in Networks and Systems, vol 1277. Springer, Singapore. https://doi.org/10.1007/978-981-96-2700-4_27
13. Díaz-Madroñero, M., Mula, J., Poler, R., Mongelli, T. (2025). Data Mining in Agriculture with Durum Wheat Price Predictions. In: Hernández, J., Kacprzyk, J. (eds) Agriculture Value Chain — Challenges and Trends in Academia and Industry. Studies in Systems, Decision and Control, vol 557. Springer, Cham. https://doi.org/10.1007/978-3-031-70745-2_11
14. Fister, I., Fister, D., Fister, I. et al. Time series numerical association rule mining variants in smart agriculture. *J Ambient Intell Human Comput* 14, 16853–16866 (2023). <https://doi.org/10.1007/s12652-023-04694-7>
15. D, C.S., Devarajan, Y. Investigation of Emerging Technologies in Agriculture: An In-depth Look at Smart Farming, Nano-agriculture, AI, and Big Data. *J. Biosyst. Eng.* 50, 170–192 (2025). <https://doi.org/10.1007/s42853-025-00258-z>
16. Asnawi, Bachri, N. & Roni, M. Vector autoregressive model: to analyze the influence of the agriculture, mining and quarrying sectors on local revenues. *Miner Econ* 38, 51–58 (2025). <https://doi.org/10.1007/s13563-024-00448-9>
17. Iqbal, M.A. (2024). Data Analytics in Digital Agriculture. In: Digital Agriculture. SpringerBriefs in Agriculture. Springer, Cham. https://doi.org/10.1007/978-3-031-67679-6_6
18. M. Dawodi, J. A. Baktash and T. Wada, "Data-Mining Opportunities in E-Government: Agriculture Sector of Afghanistan," 2019 IEEE 10th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), Vancouver, BC, Canada, 2019, pp. 0477-0481, doi: 10.1109/IEMCON.2019.8936193.
19. S. A. Lokhande, "Effective use of Big Data in Precision Agriculture," 2021 International Conference on Emerging Smart Computing and Informatics (ESCI), Pune, India, 2021, pp. 312-316, doi: 10.1109/ESCI50559.2021.9396813.
20. E. Murali and S. M. Anoucia, "A Survey on Computational Aptitudes towards Precision Agriculture using Data Mining," 2022 3rd International Conference on Smart Electronics and Communication (ICOSEC), Trichy, India, 2022, pp. 952-956, doi: 10.1109/ICOSEC54921.2022.9951960.
21. N. Sneha, K. V. Sushma and S. S. Muzumdar, "Precision Agriculture using Data Mining Techniques and IOT," 2019 1st International Conference on Advances in Information Technology (ICAIT), Chikmagalur, India, 2019, pp. 376-381, doi: 10.1109/ICAIT47043.2019.8987333.
22. J. Hirapara and P. Vanjara, "A Comparative study of Data Mining Techniques for Agriculture Crop Price Prediction," 2022 IEEE 7th International conference for Convergence in Technology (I2CT), Mumbai, India, 2022, pp. 1-6, doi: 10.1109/I2CT54291.2022.9824533.
23. Xi, F., Mei, M., Yang, S. et al. Conflict evidence combination rule based on multi-dimensional weighted evidence optimization method and its applications in pattern recognition. *J. King Saud Univ. Comput. Inf. Sci.* 37, 103 (2025). <https://doi.org/10.1007/s44443-025-00125-z>
24. Aghabeigi, F., Nazari, S. & Osati Eraghi, N. An efficient facial emotion recognition using convolutional neural network with local sorting binary pattern and whale optimization algorithm. *Int J Data Sci Anal* 20, 2275–2290 (2025). <https://doi.org/10.1007/s41060-024-00601-1>
25. F. -D. Geng, R. -B. Wang, Q. Wei and L. Xu, "Binary Arithmetic Optimization Algorithm for Feature Selection in Pattern Recognition," 2023 International Conference on Machine Learning and Cybernetics (ICMLC), Adelaide, Australia, 2023, pp. 217-222, doi: 10.1109/ICMLC58545.2023.10327990.

26. L. Shi, R. Guo and Y. Ma, "A novel artificial fish swarm algorithm for pattern recognition with convex optimization," 2016 International Conference on Communication and Electronics Systems (ICCES), Coimbatore, India, 2016, pp. 1-4, doi: 10.1109/CESYS.2016.7889830.
27. Y. Zhang, "Design and Optimization of Modular Neural Networks based on Cluster Analysis: Applied to Complex Data Pattern Recognition," 2024 International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS), Hassan, India, 2024, pp. 1-4, doi: 10.1109/IACIS61494.2024.10721815.
28. Mei Xue, Lin Jinguo and Xia Liangzheng, "Synergetic pattern recognition based on particle swarm optimization algorithm," 2008 27th Chinese Control Conference, Kunming, China, 2008, pp. 505-508, doi: 10.1109/CHICC.2008.4605165.
29. Wang Jingfang, "Non-linear pattern recognition based on SVM and genetic algorithm," 2011 International Conference on Image Analysis and Signal Processing, Wuhan, China, 2011, pp. 694-698, doi: 10.1109/IASP.2011.6109137.
30. M. W. Kurzynski and E. Puchala, "Algorithms of the multiperspective pattern recognition," Proceedings., 11th IAPR International Conference on Pattern Recognition. Vol.II. Conference B: Pattern Recognition Methodology and Systems, The Hague, Netherlands, 1992, pp. 627-630, doi: 10.1109/ICPR.1992.201855.