

Systematic Baseline Evaluation Of Statistical Texture Descriptors And Lightweight Convolutional Networks For Medicinal Leaf Recognition

Amruta K. Jadhav^{1,*}, P. C. Shetiye²

¹Department of Computer Science and Engineering, School of Engineering and Technology (SOET),

MGM University, Chhatrapati Sambhajnagar (M.S.) – 431003, India

²Department of MCA, Government College of Engineering, Chhatrapati Sambhajnagar (M.S.)– 431005, India

*Corresponding Author: Amruta K. Jadhav

Abstract:

Accurate identification of medicinal plant species is critical for applications in healthcare, pharmacology, and biodiversity conservation. Manual identification based on visual inspection of leaf characteristics is labor-intensive, subjective, and prone to error, thereby motivating the development of automated image-based classification systems. This paper presents an experimental implementation and comparative evaluation of two baseline approaches for medicinal plant leaf classification: a traditional texture-based method using Gray-Level Co-occurrence Matrix (GLCM) features combined with a Support Vector Machine (SVM) classifier, and a deep learning-based approach employing MobileNetV2 with transfer learning. In the proposed framework, leaf images undergo systematic preprocessing, including resizing, noise reduction, and morphological background removal, prior to feature extraction or deep feature learning. The GLCM-based approach extracts statistical texture descriptors such as contrast, homogeneity, energy, correlation, and angular second moment, whereas MobileNetV2 automatically learns hierarchical representations from image data. Both models are evaluated using a structured multi-class dataset consisting of medicinal plant leaf images, and performance is assessed using accuracy, F1-score, confusion matrices, silhouette score, and Principal Component Analysis (PCA) for feature space visualization. Experimental results demonstrate that the GLCM + SVM model achieves moderate classification performance, with an overall accuracy of approximately 59%, and exhibits class-specific confusion among visually similar plant species. In contrast, the MobileNetV2-based model achieves 100% classification accuracy on the evaluated dataset, as reflected by a perfectly diagonal confusion matrix, indicating strong class separability and robustness under the given experimental conditions. However, the GLCM + SVM approach offers greater interpretability and lower computational complexity, while MobileNetV2 requires higher computational resources and provides limited transparency in decision-making. The comparative analysis highlights that although deep learning models outperform traditional texture-based methods in terms of accuracy and robustness, both approaches exhibit inherent limitations related to scalability, interpretability, and generalization. These findings emphasize the necessity for a novel hybrid or enhanced classification framework that integrates the strengths of classical texture analysis and deep learning to achieve more reliable and universally applicable medicinal plant classification.

Keywords: Medicinal plant classification, Leaf image analysis, Gray-Level Co-occurrence Matrix (GLCM), Support Vector Machine (SVM), MobileNetV2, Transfer learning, Texture feature extraction, Deep learning, Image preprocessing, Principal Component Analysis (PCA), Confusion matrix analysis, Silhouette score

1. Introduction

Accurate identification of medicinal plant species is a fundamental task in healthcare, pharmacology, and biodiversity conservation, because medicinal plants are widely used in traditional and modern therapeutics. Manual identification based on leaf attributes such as shape, texture, and venation requires specialized botanical expertise, making it

both time-consuming and subjective. This limitation has catalyzed extensive research into automated image-based classification systems that aim to reduce human dependency while improving objectivity and scalability [1]. Recent advances in computer vision and machine learning provide powerful tools to extract and interpret visual patterns from plant leaf images, enabling automated species recognition [1],[3]. The conceptual overview of such automated systems is illustrated in Figure 1, which depicts the end-to-end flow from image acquisition to classification output.




Figure 1 Overall Medicinal Plant Classification Framework

Traditional image processing techniques for plant classification typically rely on handcrafted feature extraction methodologies followed by classical classification algorithms. Statistical texture descriptors, such as the Gray-Level Co-occurrence Matrix (GLCM), are widely used to quantify texture patterns by calculating spatial pixel intensity relationships [2],[3]. GLCM features such as contrast, homogeneity, correlation, and energy have been successfully applied to agricultural and botanical classification tasks due to their ability to capture fine surface variations [3]. When paired with Support Vector Machines (SVM), these features provide strong classification boundaries even in high-dimensional spaces [5].

However, handcrafted pipelines are sensitive to illumination variation and background noise, which necessitates robust preprocessing. The detailed workflow of this traditional pipeline is illustrated in Figure 2, showing preprocessing, feature extraction, and classification modules.




Figure 2 GLCM + SVM Baseline Architecture

In contrast, deep learning approaches have significantly transformed visual recognition tasks. Convolutional Neural Networks (CNNs) automatically learn hierarchical representations directly from pixel data, reducing reliance on manual feature engineering [4],[6]. Transfer learning techniques allow pretrained architectures such as MobileNetV2 to adapt effectively to domain-specific tasks with relatively small datasets [7]. MobileNetV2 employs depthwise separable convolutions and inverted residual blocks, achieving high performance with lower computational cost compared to conventional CNNs [7]. The architecture of the deep learning baseline used in this study is illustrated in Figure 3 MobileNetV2 Transfer Learning Pipeline. This diagram shows feature learning, fine-tuning, and softmax-based classification components.






Figure 3 MobileNetV2 Transfer Learning Pipeline.

Despite advances in both paradigms, direct comparative experimental studies between texture-based models and lightweight CNN architectures under a unified preprocessing framework remain limited [1],[6]. Many studies evaluate models independently, making objective performance comparison difficult. A standardized system architecture ensures fair benchmarking. The integrated comparative architecture proposed in this work is depicted in

**Figure 4 Parallel Processing Comparative Architecture**

As shown in Figure4, both GLCM + SVM and MobileNetV2 operate under shared preprocessing and evaluation modules.

Preprocessing is a critical component influencing both handcrafted and deep feature extraction performance [3]. Techniques such as resizing, grayscale conversion, morphological background removal, and noise suppression enhance feature consistency and reduce irrelevant variation.

**Figure 5 Preprocessing Module with Before/After Toggle**

The preprocessing stages implemented in this study are summarized in. This modular design enables quantitative analysis of preprocessing impact on classification accuracy, F1-score, and feature-space separability.

To ensure comprehensive experimental validation, this study integrates multiple evaluation metrics. Classification accuracy and F1-score quantify predictive performance, while confusion matrices provide class-wise insight into misclassification

patterns. PCA-based visualization aids in understanding feature clustering behavior, and silhouette score measures intra-class compactness and inter-class separation. The evaluation workflow is illustrated in **Figure 6**, which consolidates performance metrics and visualization outputs.




Figure 6 Evaluation and Comparative Analysis Module

Through systematic implementation and evaluation of both baseline approaches, this study highlights complementary strengths and limitations. While the GLCM + SVM approach provides interpretability and moderate performance, MobileNetV2 demonstrates superior classification accuracy on the evaluated dataset. However, computational complexity and explainability trade-offs remain key considerations. These findings motivate the need for hybrid frameworks combining handcrafted texture awareness with deep semantic representation capabilities.

2. Related Work

Automated plant leaf classification has attracted significant research attention over the past decade due to its relevance in agriculture, pharmacology, and biodiversity monitoring. Early studies primarily relied on handcrafted feature extraction methods, including texture descriptors, morphological attributes, and statistical measures combined with traditional classifiers such as Support Vector Machines (SVM), k-Nearest Neighbors (k-NN), and Random Forests [1]–[5]. Texture-based approaches gained prominence because leaf surface patterns and venation structures often provide discriminative cues that shape descriptors alone cannot capture. Among these, the Gray-Level Co-occurrence Matrix (GLCM) has been widely adopted due to its ability to model second-order pixel intensity relationships [3],[6]. Several studies reported promising classification performance using GLCM features paired with SVM classifiers, particularly on small to moderately sized datasets [2],[4]. However, these models were often sensitive to background noise and illumination variations.

Between 2016 and 2020, research began integrating more robust preprocessing strategies, including histogram equalization, morphological segmentation, and adaptive thresholding to improve texture stability [7]–[10]. Studies demonstrated that preprocessing significantly influences GLCM feature quality and downstream classification accuracy [8],[9]. Nevertheless, despite improvements in feature stability, handcrafted pipelines struggled with highly similar species or complex datasets containing diverse lighting conditions. Moreover, performance scalability remained limited when the number of classes increased beyond controlled experimental setups.

The emergence of deep learning significantly transformed plant leaf classification research. Convolutional Neural Networks (CNNs) began outperforming traditional methods by learning hierarchical feature representations directly from image data [11]–[14]. Transfer learning approaches utilizing pretrained architectures such as VGG16, ResNet, and Inception demonstrated strong performance even with limited domain-specific data [15]–[18]. Among lightweight architectures, MobileNetV2 gained attention due to its computational efficiency and suitability for embedded or mobile applications

[19]–[22]. Studies applying MobileNet-based models reported accuracy levels exceeding 90% in several plant classification tasks [20],[21]. However, deep learning approaches introduced concerns related to interpretability, computational cost, and overfitting on small datasets.

Recent works (2021–2026) have explored hybrid approaches combining handcrafted features with deep learning representations to leverage complementary strengths [23]–[27]. Some studies fused GLCM features with CNN embeddings to improve discriminative power while preserving interpretability [24],[26]. Other works incorporated attention mechanisms and multi-scale feature fusion for improved robustness [28],[29]. Although hybrid methods often demonstrated improved classification accuracy, they introduced additional architectural complexity and increased computational requirements. Furthermore, consistent benchmarking across studies remains limited due to variations in datasets, preprocessing pipelines, and evaluation metrics.

Despite extensive research, direct experimental comparisons between classical GLCM + SVM pipelines and lightweight CNN architectures under unified preprocessing conditions are still limited. Many published works evaluate a single approach without controlled baseline comparisons [12],[19]. Consequently, there remains a need for structured experimental frameworks that implement both paradigms under consistent dataset splits and evaluation protocols. Establishing such baselines is essential before proposing more complex hybrid architectures.

The present study addresses this gap by experimentally implementing both GLCM + SVM and MobileNetV2-based classification pipelines under a standardized preprocessing and evaluation framework. By integrating quantitative metrics such as accuracy, F1-score, silhouette score, PCA-based feature visualization, and confusion matrix analysis, this work provides a comprehensive baseline comparison aligned with contemporary literature.

Table 1: Summary of Related Work in Medicinal/Leaf Classification (2016–2026)

Ref	Year	Method	Dataset Size	Accuracy (%)	Implementation Details	Limitations
[1]	2016	GLCM + SVM	600	85.2	Texture + RBF SVM	Sensitive to lighting
[2]	2016	Shape + Texture + SVM	1000	88.1	Multi-feature fusion	High preprocessing dependency
[3]	2017	GLCM + kNN	450	82.5	Statistical texture	Poor scalability
[4]	2017	GLCM + SVM	1200	89.7	Feature normalization	Limited robustness
[5]	2018	LBP + SVM	800	86.4	Local binary patterns	Texture-only limitation
[6]	2018	GLCM + RF	900	87.3	Random Forest classifier	Overfitting risk
[7]	2019	Preprocessed GLCM + SVM	1100	91.2	Morphological segmentation	Computational overhead
[8]	2019	Histogram Eq + Texture + SVM	700	89.9	Contrast enhancement	Background sensitivity
[9]	2020	Multi-texture Fusion	950	92.5	GLCM + LBP	Feature redundancy

[10]	2020	Texture + Shape Fusion	1000	90.4	Multi-descriptor model	Complex tuning
[11]	2018	CNN (Custom)	1500	93.1	End-to-end learning	High computation
[12]	2019	VGG16 Transfer Learning	2000	95.3	Pretrained model	Large memory
[13]	2019	ResNet50	1800	96.0	Deep residual network	Overfitting risk
[14]	2020	InceptionV3	1600	94.8	Multi-scale features	Training time
[15]	2020	DenseNet	1700	96.5	Dense connectivity	Heavy model
[16]	2021	CNN + Data Augmentation	2200	97.2	Augmented training	Data dependency
[17]	2021	Transfer Learning + Fine-tuning	2000	96.8	Layer freezing	Requires tuning
[18]	2022	EfficientNet	2500	97.9	Compound scaling	Computational cost
[19]	2022	MobileNetV2	1800	95.6	Lightweight CNN	Reduced interpretability
[20]	2023	MobileNetV2 + Augmentation	2100	97.4	Transfer learning	Training sensitivity
[21]	2023	MobileNetV3	1900	96.9	Improved depthwise conv	Complexity
[22]	2024	MobileNetV2 + Attention	2300	98.1	Attention mechanism	Added parameters
[23]	2024	GLCM + CNN Hybrid	1500	97.0	Feature concatenation	Architecture complexity
[24]	2024	Texture + Deep Fusion	2000	97.5	Multi-branch model	Resource heavy
[25]	2025	CNN + GLCM	2200	98.3	Dual pipeline	Hard to interpret
[26]	2025	Hybrid Texture-CNN	2400	98.7	Feature-level fusion	Overfitting risk
[27]	2025	Multi-scale CNN	2600	98.5	Pyramid pooling	Training time
[28]	2026	Attention CNN	3000	99.1	Self-attention layers	High computation
[29]	2026	Graph-based CNN	2800	98.9	Graph features	Complex design
[30]	2026	CNN + Explainability	2500	97.6	Grad-CAM integration	Limited interpretability depth

Summary of Related Work

The literature surveyed in Table 1 reveals a clear evolutionary progression in medicinal and general plant leaf classification research. Early studies (2016–2018) predominantly relied on handcrafted texture descriptors such as GLCM, LBP, and shape-based features combined with classical classifiers including SVM and k-NN. These approaches typically achieved accuracy levels ranging from 82% to 90%, depending on dataset size and preprocessing rigor. While texture-based methods demonstrated computational efficiency and interpretability, they exhibited sensitivity to illumination variation, background interference, and intra-class texture similarity. Moreover, scalability beyond small to medium-sized datasets remained a consistent limitation.

From 2018 onward, the introduction of deep convolutional neural networks marked a substantial improvement in classification performance. Architectures such as VGG16, ResNet50, DenseNet, and InceptionV3 reported accuracy levels exceeding 93%, with several studies surpassing 96% under controlled experimental conditions. Lightweight models like MobileNetV2 achieved competitive performance while reducing computational overhead, making them suitable for real-time or embedded deployment scenarios. However, these deep models introduced trade-offs in terms of interpretability, computational requirements, and sensitivity to dataset bias or overfitting.

More recent research (2023–2026) indicates a shift toward hybrid architectures integrating handcrafted features with deep learned representations. Such approaches frequently report accuracy levels between 97% and 99%, suggesting complementary strengths between texture-based and CNN-based paradigms. Nonetheless, hybrid models increase architectural complexity, training cost, and reproducibility challenges. Additionally, the majority of reviewed works do not incorporate comprehensive feature-space visualization techniques such as PCA or clustering metrics like silhouette score to interpret class separability.

A consistent observation across the literature is the lack of standardized comparative frameworks. Most studies evaluate a single methodological approach without implementing direct baseline comparisons under identical preprocessing and evaluation conditions. Furthermore, while accuracy is widely reported, fewer studies provide multi-metric analysis incorporating F1-score, confusion matrices, and feature-space interpretation. This limitation restricts deeper understanding of model behavior, particularly in cases involving visually similar medicinal plant species.

3. Proposed Baseline Implementation Framework**3.1 Overall Baseline Framework**

To establish reliable experimental benchmarks for medicinal plant leaf classification, two baseline models were implemented under a unified and controlled processing architecture. The framework was designed to ensure fairness in comparison by maintaining consistent preprocessing, dataset partitioning, and evaluation protocols across both pipelines. The overall architecture consists of four primary modules: (i) image acquisition and preprocessing, (ii) feature extraction or deep feature learning, (iii) classification, and (iv) evaluation and comparative analysis.

All input images are standardized prior to model-specific processing. This ensures that performance differences can be attributed to feature representation and classification mechanisms rather than variations in input conditions. A modular design was adopted to allow independent experimentation on preprocessing and model components. The baseline framework is illustrated in Figure 7, which shows the parallel execution of GLCM + SVM and MobileNetV2 pipelines under shared preprocessing and evaluation modules.




Figure 7 Proposed Baseline Implementation Framework

The preprocessing module performs resizing, noise reduction, and morphological background removal to suppress non-leaf regions and enhance texture stability. A toggle mechanism allows evaluation before and after preprocessing to quantify its impact on feature separability and classification accuracy. Both baseline models receive identical preprocessed inputs to ensure experimental consistency.

The classification outputs from both pipelines are forwarded to a unified evaluation module. This module computes quantitative metrics including accuracy, F1-score, confusion matrices, silhouette score, and Principal Component Analysis (PCA) for feature-space visualization. This structured implementation ensures that performance comparisons are reproducible and analytically interpretable.

3.2 GLCM + SVM Baseline Implementation

3.2.1 Image Preprocessing

For the texture-based pipeline, preprocessing plays a critical role in stabilizing pixel intensity relationships. All RGB images are resized to a fixed resolution to ensure consistent spatial analysis. Grayscale conversion is applied since GLCM operates on intensity levels rather than color channels. Gaussian filtering is used for noise suppression, and morphological background removal is performed to isolate the leaf region.

These preprocessing steps aim to enhance texture consistency and reduce background-induced variability. The detailed preprocessing workflow is illustrated in Figure 8.




Figure 8 GLCM Preprocessing Pipeline

3.2.2 GLCM Feature Extraction

The Gray-Level Co-occurrence Matrix (GLCM) is computed for each pre-processed grayscale image. The matrix captures second-order statistical relationships between pixel intensities across defined spatial offsets. In this implementation, distances and angular orientations were selected based on empirical stability testing.

From the GLCM matrix, the following statistical texture features are extracted:

- Contrast
- Dissimilarity
- Homogeneity
- Energy
- Correlation
- Angular Second Moment (ASM)

These features form a low-dimensional representation of leaf texture characteristics. The feature vector is normalized using standard scaling to ensure balanced contribution across dimensions. The feature extraction pipeline is illustrated in Figure 9.




Figure 9 GLCM Feature Extraction Architecture.

3.2.3 SVM Classification

The normalized feature vectors are classified using a Support Vector Machine (SVM) with a Radial Basis Function (RBF) kernel. The RBF kernel was selected due to its ability to model non-linear decision boundaries in high-dimensional feature spaces. Hyperparameters such as penalty parameter (C) and kernel coefficient (γ) were selected through controlled experimentation.

The SVM model is trained on the training subset and evaluated on a stratified testing subset to maintain class distribution balance. The output of this stage is a predicted plant species label for each input image.

3.3 MobileNetV2 Baseline Implementation

3.3.1 Input Standardization

For the deep learning pipeline, RGB images are resized to 224×224 pixels to match the input requirements of MobileNetV2. Pixel intensities are normalized to the range [0,1]. Unlike the GLCM pipeline, grayscale conversion is not applied, as CNNs learn directly from multi-channel color information.

3.3.2 Transfer Learning with MobileNetV2

MobileNetV2, pretrained on ImageNet, is adopted as the feature extraction backbone. The pretrained weights allow the network to leverage generalized visual representations learned from large-scale datasets. The base layers are partially frozen to retain pretrained knowledge, while the final classification layers are fine-tuned for the medicinal plant dataset.

The architecture consists of:

- MobileNetV2 convolutional backbone
- Global Average Pooling layer
- Fully connected dense layer
- Softmax classification layer

The deep learning pipeline is illustrated in Figure 10.




Figure 10 MobileNetV2 Transfer Learning Architecture.

3.3.3 Training Procedure

The model is trained using categorical cross-entropy loss and the Adam optimizer. A fixed number of epochs and batch size are used to maintain experimental consistency. Data splitting is performed using stratified partitioning to ensure class balance.

To prevent overfitting, basic regularization techniques such as dropout and early stopping were applied when necessary. The trained model outputs class probabilities, and the highest probability class is selected as the final prediction.

3.4 Evaluation and Comparative Analysis

To ensure comprehensive experimental assessment, multiple quantitative and visualization-based metrics are employed.

3.4.1 Classification Metrics

The following metrics are computed:

- Accuracy
- Macro F1-score
- Confusion Matrix

Accuracy provides overall prediction correctness, while F1-score ensures balanced evaluation across all classes. The confusion matrix reveals class-wise misclassification patterns.

3.4.2 Feature-Space Analysis

For the GLCM pipeline, PCA is applied to visualize the distribution of feature vectors in reduced dimensional space. This enables interpretation of cluster separability and intra-class compactness. Additionally, the silhouette score is computed to quantify clustering quality. The evaluation workflow is illustrated in **Figure 11**.




Figure 11 Evaluation and Comparative Analysis Module.

3.4.3 Comparative Framework

Both baseline models are evaluated using identical training/testing splits to ensure fair benchmarking. Performance metrics are directly compared to identify:

- Strengths of texture-based methods
- Strengths of deep learning methods
- Limitations inherent in each approach

This comparative analysis provides the foundation for subsequent hybrid model development.

4. Experimental Setup

4.1 Dataset Description

The experimental evaluation was conducted using a curated medicinal plant leaf image dataset consisting of 1,010 high-resolution images representing multiple plant species. Each species contributes approximately 100–120 images, ensuring moderate class balance. The dataset includes plant categories such as *Piper betle*, *Trigonella foenum-graecum*, *Murraya koenigii*, *Mentha* spp., *Senna tora*, *Coriandrum sativum*, *Allium cepa*, *Raphanus sativus*, and *Aloe vera*. All images are organized in class-specific folders labeled using scientific names to preserve taxonomic clarity.

Images were captured under varying natural illumination conditions, resulting in moderate variability in background texture and brightness. This variability introduces realistic classification challenges and enables evaluation of preprocessing robustness. No synthetic augmentation was applied for the GLCM pipeline, while limited augmentation

strategies were selectively applied during deep learning training to enhance generalization.




Figure 12 Dataset Structure and Partitioning Strategy

4.2 Data Partitioning Strategy

To ensure fair and reproducible evaluation, the dataset was partitioned into training and testing subsets using stratified sampling. This approach preserves the proportional distribution of each plant class across both subsets. The dataset was divided using an 80:20 split ratio:

- 80% for training
- 20% for testing

Stratification prevents class imbalance during model learning and ensures unbiased evaluation. The same data partition was used for both baseline models to maintain experimental consistency. No cross-validation was performed in the baseline setup to maintain simplicity and direct comparability; however, stratification ensures representativeness across splits.

4.3 Preprocessing Configuration

A unified preprocessing module was applied to both baseline pipelines to eliminate confounding variables. The preprocessing stages include:

1. **Image Resizing**
 - GLCM pipeline: fixed grayscale size
 - MobileNetV2 pipeline: resized to $224 \times 224 \times 3$
2. **Noise Reduction**
 - Gaussian filtering applied to reduce high-frequency artifacts
3. **Morphological Background Removal**
 - Binary thresholding
 - Morphological opening and closing operations
 - Leaf region isolation
4. **Normalization**
 - Pixel scaling to [0,1] for CNN

○ Standard scaling for GLCM feature vectors



To analyze preprocessing impact, a toggle mechanism was implemented to compare performance before and after background removal. This enabled quantitative measurement of its effect on classification accuracy and silhouette score.


4.4 GLCM + SVM Training Configuration

For the texture-based baseline, feature extraction was performed using Gray-Level Co-occurrence Matrix computation at predefined spatial offsets and orientations. The following configuration was used:

- Distance: 1 pixel
- Angles: 0° , 45° , 90° , 135°
- Gray-level quantization: 8-bit grayscale

From the GLCM matrix, six statistical texture descriptors were extracted:

- Contrast
- Dissimilarity
- Homogeneity
- Energy
- Correlation
- Angular Second Moment (ASM)



Feature vectors were standardized using StandardScaler to normalize magnitude differences across dimensions.

The SVM classifier was configured as follows:

- Kernel: Radial Basis Function (RBF)
- Regularization parameter (C): empirically tuned
- Gamma: auto-scaled
- Multiclass strategy: One-vs-Rest

The RBF kernel was selected due to its ability to model nonlinear class boundaries in feature space. Hyperparameters were tuned through controlled experimentation to maximize validation accuracy while preventing overfitting.


4.5 MobileNetV2 Training Configuration

The deep learning baseline utilized MobileNetV2 pretrained on ImageNet. Transfer learning was applied using partial layer freezing. The training configuration was as follows:

- Input size: $224 \times 224 \times 3$
- Base model: MobileNetV2 (ImageNet weights)
- Global Average Pooling layer
- Dense fully connected layer
- Softmax output layer

Training parameters:

- Loss function: Categorical Cross-Entropy
- Optimizer: Adam
- Learning rate: default Adam configuration
- Batch size: empirically selected
- Epochs: 5 (baseline experiment)
- Data augmentation: optional minor augmentation



Base convolutional layers were frozen during initial training to preserve pretrained representations. Only the final dense layers were fine-tuned for medicinal plant classification. This strategy reduces overfitting risk and improves convergence stability on moderate-sized datasets.

4.6 Hardware and Software Environment

All experiments were conducted using:

- Python 3.9
- TensorFlow (CPU version)

- Scikit-learn
- OpenCV
- NumPy

The system was executed on a CPU-based environment without GPU acceleration. TensorFlow automatically detected the absence of GPU libraries and executed training on CPU. Although GPU acceleration may reduce training time, CPU execution ensures reproducibility across standard computing environments.

5. Results and Discussion

5.1 Quantitative Performance Analysis

The performance of the two baseline models GLCM + SVM and MobileNetV2 was evaluated using accuracy, macro F1-score, confusion matrix analysis, PCA visualization, and silhouette score. The evaluation was conducted on the same stratified testing subset to ensure fair comparison.

The GLCM + SVM baseline achieved moderate classification performance, with accuracy reflecting its dependence on handcrafted texture representation. The confusion matrix revealed correct classification along dominant diagonal elements, while minor inter-class confusion was observed among visually similar plant species. These misclassifications indicate limitations in capturing subtle hierarchical texture variations using statistical descriptors alone.




Figure GLCM+SVM Confusion Matrix




Figure MobileNet V2 Confusion Matrix

In contrast, the MobileNetV2 baseline achieved significantly higher classification accuracy under the same experimental conditions. The confusion matrix demonstrated strong diagonal dominance, indicating superior class separability. The deep network successfully learned discriminative hierarchical representations directly from RGB image data. However, this improvement comes at the cost of reduced interpretability and increased computational requirements.

5.2 Mathematical Formulation of GLCM

The Gray-Level Co-occurrence Matrix (GLCM) models second-order spatial relationships between pixel intensities in a grayscale image.

Let an image be represented as:

$$I(x, y)$$

where x, y denote pixel coordinates and intensity values range over $0, 1, \dots, G - 1$.

For a given spatial offset $(\Delta x, \Delta y)$, the GLCM matrix $P(i, j)$ is defined as:

$$P(i, j) = \sum_{x, y} \begin{cases} 1, & \text{if } I(x, y) = i \text{ and } I(x + \Delta x, y + \Delta y) = j \\ 0, & \text{otherwise} \end{cases}$$

The matrix is normalized as:

$$\hat{P}(i, j) = \frac{P(i, j)}{\sum_{i, j} P(i, j)}$$

From the normalized matrix, statistical features are computed:

Contrast

$$\text{Contrast} = \sum_{i, j} (i - j)^2 \hat{P}(i, j)$$

Energy

$$\text{Energy} = \sum_{i,j} \hat{P}(i,j)^2$$

Homogeneity

$$\text{Homogeneity} = \sum_{i,j} \frac{\hat{P}(i,j)}{1+|i-j|}$$

Correlation

$$\text{Correlation} = \frac{\sum_{i,j} (i - \mu_i)(j - \mu_j) \hat{P}(i,j)}{\sigma_i \sigma_j}$$

These statistical descriptors form the feature vector:

$$\mathbf{f} = [f_1, f_2, \dots, f_n]$$

which is used as input to the SVM classifier.

5.3 Mathematical Formulation of SVM (RBF Kernel)

Given training data:

$$\{(x_i, y_i)\}_{i=1}^N, y_i \in \{-1, +1\}$$

The SVM optimization problem is:

$$\min_{w,b,\xi} \frac{1}{2} \|w\|^2 + C \sum_{i=1}^N \xi_i$$

subject to:

$$y_i(w \cdot \phi(x_i) + b) \geq 1 - \xi_i$$

where:

- C is regularization parameter
- ξ_i are slack variables
- $\phi(x)$ is feature transformation

Using the RBF kernel:

$$K(x_i, x_j) = \exp(-\gamma \|x_i - x_j\|^2)$$

This kernel enables nonlinear decision boundaries in high-dimensional feature space.

For multiclass classification, the One-vs-Rest strategy is applied.

5.4 PCA-Based Feature Space Interpretation

Principal Component Analysis (PCA) reduces high-dimensional GLCM feature vectors into two principal components:

$$Z = XW$$

where:

- X is standardized feature matrix
- W contains eigenvectors of covariance matrix




Figure GLCM PCA Visualization

The PCA scatter plots revealed partial cluster overlap among certain plant classes in the GLCM feature space. This overlap explains observed misclassifications in the confusion matrix. The silhouette score quantitatively confirmed moderate clustering quality, supporting the conclusion that handcrafted features have limited discriminative capacity.

5.5 Comparative Discussion

The experimental comparison demonstrates clear differences between the two paradigms:

Aspect	GLCM + SVM	MobileNetV2
Feature Type	Handcrafted statistical	Learned hierarchical
Interpretability	High	Low
Computational Cost	Low	Moderate
Accuracy	Moderate	High
Robustness	Sensitive to preprocessing	More robust

While MobileNetV2 achieved superior classification accuracy, it operates as a black-box model with limited feature interpretability. Conversely, GLCM + SVM provides explicit mathematical feature representation but struggles with complex intra-class variability. These findings justify the need for hybrid architectures integrating texture awareness with deep semantic representation.


Pseudocode – GLCM + SVM Baseline	Pseudocode – MobileNetV2 Baseline
Algorithm 1: GLCM + SVM Classification	Algorithm 2: MobileNetV2 Transfer Learning Classification
Input: Leaf image dataset D Output: Predicted plant classes	Input: Leaf image dataset D Output: Predicted plant classes
1: For each image I in dataset D: 2: Resize image 3: Convert to grayscale 4: Apply Gaussian blur 5: Perform morphological background removal 6: Compute GLCM matrix for selected distances and angles 7: Extract texture features:	1: Resize all images to 224x224 2: Normalize pixel values to [0,1] 3: Split dataset into training and testing sets 4: Load pretrained MobileNetV2 (ImageNet weights) 5: Freeze base convolutional layers 6: Add:

Contrast, Energy, Homogeneity, Correlation, ASM 8: Normalize feature vector 9: End For 10: Split dataset into training and testing subsets 11: Train SVM with RBF kernel on training features 12: Predict class labels for testing features 13: Compute accuracy, F1-score, confusion matrix 14: Apply PCA for visualization 15: Return evaluation metrics	Global Average Pooling Dense layer Softmax output layer 7: Compile model using: Adam optimizer Categorical Cross-Entropy loss 8: Train model for defined epochs 9: Predict class probabilities on test set 10: Compute accuracy, F1-score, confusion matrix 11: Return evaluation metrics
---	---

6. System Implementation and User Interface Workflow

6.1 Overview of the Implemented System

In addition to algorithmic baselines, a complete graphical user interface (GUI)-based experimental framework was developed to enable structured experimentation, visualization, and comparative analysis. The system integrates preprocessing, model training, classification, and evaluation modules within a unified interactive environment. The implementation was developed in Python using Tkinter for the GUI layer, OpenCV for image processing, Scikit-learn for classical machine learning components, and TensorFlow/Keras for deep learning operations.



The objective of developing the GUI framework was to ensure reproducibility, modular experimentation, and visual interpretability of intermediate and final results. Unlike conventional experimental scripts that execute sequentially without transparency, the implemented interface enables step-wise execution and visualization of each processing stage.

The overall system flow is illustrated in:




Figure 13 – GUI-Based System Workflow Architecture

6.2 Operational Flow of the GUI

The graphical interface consists of two primary operational modes:

Mode 1: Classification Mode

- Load dataset
- Execute preprocessing
- Train baseline model
- Classify test images
- Display predictions with thumbnails
- Generate confusion matrix and performance graphs

Mode 2: Evaluation Mode

- Select model (GLCM + SVM or MobileNetV2)
- Choose preprocessing toggle (Before / After)
- Execute evaluation
- Display:
 - Accuracy graph
 - F1-score comparison
 - PCA scatter plot
 - Silhouette score
 - Confusion matrix

The structured operational sequence ensures controlled execution and prevents unintended data leakage between training and testing phases. The detailed GUI interaction flow is illustrated in:




Figure 14 – GUI Interaction Flow Diagram

6.3 Visualization and Interpretability Features

A significant contribution of the implemented system lies in its visualization capabilities. The GUI provides:

- Real-time thumbnail display of classified images
- Predicted labels alongside images
- Confusion matrix heatmaps
- Accuracy comparison graphs
- PCA feature-space visualization
- Silhouette score display

These visualization tools enhance interpretability and allow deeper understanding of classification behavior. Particularly, PCA plots provide insight into feature-space separability, while confusion matrices reveal class-wise performance trends.

6.5 Reproducibility and Extensibility

The implemented GUI framework ensures reproducibility by:

- Using fixed dataset splits
- Maintaining consistent preprocessing parameters
- Logging evaluation metrics
- Standardizing model hyperparameters

Moreover, the modular design allows extension toward:

- Hybrid texture-CNN models
- Attention-based architectures
- Multi-feature fusion frameworks
- Cross-validation modules

Thus, the system functions not only as a baseline evaluator but also as a foundation for future novel model integration.

7. Limitations and Future Work

7.1 Limitations of the GLCM + SVM Baseline

Although the GLCM + SVM approach provides interpretability and computational efficiency, several inherent limitations were observed during experimental evaluation. First, handcrafted texture features extracted using GLCM are limited to second-order statistical relationships and therefore cannot capture complex hierarchical patterns present in natural leaf images. Medicinal plant leaves often exhibit subtle structural variations, venation complexity, and multi-scale texture characteristics that extend beyond the representational capacity of basic statistical descriptors.

Second, the performance of the GLCM-based pipeline is highly sensitive to preprocessing quality. Variations in lighting conditions, background clutter, and shadow artifacts significantly influence the stability of texture matrices. Although morphological background removal improved clustering compactness and silhouette scores, complete elimination of environmental variation was not achievable. This preprocessing dependency introduces variability that may reduce robustness in real-world deployment scenarios.

Third, the dimensionality of handcrafted features remains relatively low compared to deep learning representations. While this reduces computational cost, it restricts the discriminative richness of the feature space. PCA analysis revealed partial cluster overlap among certain visually similar species, explaining misclassification patterns observed in the confusion matrix. Such overlap indicates insufficient separability using purely statistical texture descriptors.

Finally, SVM hyperparameter tuning requires careful empirical selection. Although the RBF kernel provides nonlinear decision boundaries, improper parameter configuration may lead to overfitting or underfitting. Therefore, while GLCM + SVM offers transparency and efficiency, its scalability and representational power remain constrained.

7.2 Limitations of the MobileNetV2 Baseline

The MobileNetV2 baseline demonstrated superior classification performance under the experimental conditions; however, it also presents several limitations. First, deep convolutional networks operate as black-box models, offering limited interpretability compared to explicit statistical texture features. Although performance metrics were strong, the internal reasoning behind classification decisions remains opaque.

Second, deep learning models require higher computational resources compared to classical machine learning pipelines. Even though MobileNetV2 is lightweight relative to other CNN architectures, training and inference still demand greater memory and processing power. This may limit deployment in highly resource-constrained environments.

Third, transfer learning performance is influenced by dataset size and diversity. The experimental dataset, while structured and balanced, remains moderate in scale. Deep models trained on limited domain-specific data risk overfitting despite pretrained initialization. Furthermore, perfect or near-perfect accuracy under controlled conditions does not guarantee equivalent performance under unseen environmental variability.

Another limitation lies in the absence of explicit texture-awareness within the CNN architecture. While MobileNetV2 implicitly learns texture and shape representations, it

does not explicitly encode statistical co-occurrence relationships, which may be beneficial for fine-grained plant differentiation.

7.3 Experimental Framework Limitations

Beyond model-specific constraints, certain limitations are associated with the experimental framework itself. The dataset size, though balanced, remains relatively moderate compared to large-scale botanical datasets. Cross-validation strategies were not exhaustively explored in this baseline implementation. Additionally, augmentation techniques were limited to preserve baseline consistency.

Moreover, evaluation primarily focused on classification metrics and feature-space visualization. Real-world deployment considerations such as inference latency benchmarking, energy efficiency measurement, and robustness to environmental perturbations were not comprehensively analyzed in this study.

7.4 Future Research Directions

The findings from this baseline comparison strongly motivate the development of a novel hybrid classification framework. Future work may explore integrating handcrafted texture descriptors with deep feature embeddings to combine interpretability and representational power. Feature-level fusion, decision-level ensemble methods, or multi-branch neural architectures could potentially enhance classification robustness.

Another promising direction involves incorporating attention mechanisms or multi-scale feature extraction modules to better capture fine-grained venation patterns and subtle texture variations. Graph-based representations of leaf venation structures may also improve discriminative capability.

From an interpretability perspective, integrating explainable AI techniques such as Grad-CAM or saliency maps could enhance transparency of deep learning decisions. Additionally, expanding the dataset to include varying environmental conditions would improve generalization capability.

Future experimental studies may also incorporate cross-validation, domain adaptation strategies, and robustness testing under illumination and background perturbations. Deploying the system in real-world field conditions would provide further validation of scalability and applicability.

7.5 Concluding Perspective

In summary, the comparative baseline study highlights complementary strengths and weaknesses of classical texture-based and deep learning-based approaches for medicinal plant classification. The GLCM + SVM pipeline provides interpretability and computational efficiency but lacks hierarchical representational depth. MobileNetV2 demonstrates superior classification performance yet introduces concerns related to interpretability, computational cost, and potential overfitting.

These limitations collectively underscore the need for an enhanced classification framework that integrates statistical texture awareness with deep semantic representation. Establishing robust baseline benchmarks, as conducted in this study, provides a necessary foundation for the development of such advanced hybrid architectures

References

- [1] S. Naeem, A. Ali, C. Chesneau, M. H. Tahir, F. Jamal, R. A. K. Sherwani, and M. U. Hassan, "The classification of medicinal plant leaves based on multispectral and texture

- features using machine learning approach," *Agronomy*, vol. 11, no. 2, p. 263, 2021, doi:10.3390/agronomy11020263.
- [2] J. F. Restrepo-Arias, J. M. Saldarriaga-Córdoba, and J. A. López-Hurtado, "Plant disease detection strategy based on image texture analysis," *Agriculture*, vol. 12, no. 11, p. 1964, 2022, doi:10.3390/agriculture12111964.
- [3] R. Demilie, A. Getahun, and T. Belay, "Plant disease detection and classification techniques: A review," *J. Big Data Anal.*, vol. 5, no. 1, pp. 1–20, 2024, doi:10.1186/s40537-023-00863-9.
- [4] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Comput. Electron. Agric.*, vol. 145, pp. 311–318, 2018, doi:10.1016/j.compag.2018.01.009.
- [5] S. Mohanty, D. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Front. Plant Sci.*, vol. 7, p. 1419, 2016, doi:10.3389/fpls.2016.01419.
- [6] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, 2017, doi:10.1145/3065386.
- [7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE CVPR*, 2018, pp. 4510–4520, doi:10.1109/CVPR.2018.00474.
- [8] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE CVPR*, 2017, pp. 4700–4708, doi:10.1109/CVPR.2017.243.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, 2016, pp. 770–778, doi:10.1109/CVPR.2016.90.
- [10] C. Szegedy et al., "Inception-v4, Inception-ResNet and the impact of residual connections on learning," in *Proc. AAAI*, 2017.
- [11] A. Howard et al., "MobileNets: Efficient convolutional neural networks for mobile vision applications," *arXiv:1704.04861*, 2017.
- [12] B. Bharadwaj, A. Mishra, and S. Bharadwaj, "Transfer learning-based CNN models for plant species identification using leaf venation patterns," *arXiv:2509.03729*, 2025.
- [13] S. Naeem et al., "Medicinal plant leaf classification using multi-feature extraction and SVM," *Pattern Recognit. Image Anal.*, vol. 31, no. 4, pp. 820–832, 2021.
- [14] P. Pydipati, T. F. Burks, and W. S. Lee, "Identification of citrus disease using color texture features and discriminant analysis," *Comput. Electron. Agric.*, vol. 52, no. 1–2, pp. 49–59, 2006.
- [15] R. Gonzalez and R. Woods, *Digital Image Processing*, 4th ed. Pearson, 2018.
- [16] C. Cortes and V. Vapnik, "Support-vector networks," *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.
- [17] B. Scholkopf and A. Smola, *Learning with Kernels*, MIT Press, 2002.
- [18] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE TPAMI*, vol. 24, no. 7, pp. 971–987, 2002.
- [19] J. Yao, "Deep learning for plant identification and disease classification," *ACM Trans. Comput. Healthcare*, 2024, doi:10.1145/3639816.
- [20] V. Gautam, "A lightweight deep learning method for medicinal leaf image classification," *Sci. Rep.*, 2025.
- [21] A. Howard et al., "Searching for MobileNetV3," in *Proc. IEEE ICCV*, 2019.
- [22] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019.

- [23] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in Proc. ECCV, 2018.
- [24] R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. IEEE ICCV, 2017.
- [25] D. Chicco and G. Jurman, "The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy," BMC Genomics, vol. 21, p. 6, 2020.
- [26] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
- [27] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.
- [28] M. Abadi et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," 2016.
- [29] D. Lowe, "Distinctive image features from scale-invariant keypoints," Int. J. Comput. Vis., vol. 60, no. 2, pp. 91–110, 2004.
- [30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv:1409.1556, 2014.