

A Comprehensive Framework for Leaf Shape Analysis: From Preprocessing to Morphology-Based Classification

Usha N¹

Assistant Professor

Faculty of Computing & IT

GM University

Davangere, Karnataka, India-577006

usha.narayanshetty1@gmail.com

Shashank Shivananda Chattaraki²

II Year MCA

Faculty of Computing & IT

GM University

Davangere, Karnataka, India-577006

ABSTRACT

Automated identification of plant species is important for agriculture, conservation of biological diversity, and research into medicinal plants. Identification of plants using traditional methods often requires a considerable amount of skills in botany and can be time consuming. However, with the increasing availability of digital imaging devices, vast amounts of leaf image data can now be acquired. Unfortunately, many of these digital leaf images contain problems including background clutter, differences in illumination due to different light conditions, and differences in the orientation of the imaging device. To address these challenges, this paper presents a deep learning framework that uses morphological data as the basis for identifying plant species through the combination of the following processes: image pre-processing, image segmentation, extracting morphological features, and utilizing a mobile lightweight Convolutional Neural Network (CNN), specifically MobileNetV2, to classify plant species based on the extracted morphological features. Morphological descriptors, such as circularity, aspect ratio, solidity, and vein density, are used to provide effective representation of the morphological features of the leaf. Additionally, deep learning approaches using CNNs have been shown to have a superior ability to automatically learn feature representations hierarchically and to improve classification accuracy over traditional handcrafted feature methods [7], [11]. Furthermore, lightweight and efficient CNN architectures allow for real-time performance in mobile applications [6]. Additionally, the system developed in this study includes an online interface that allows for real-time predictions and viewing of results. The experimental results indicate that, in using morphological features together with deep learning techniques, the classification performance was improved, with enhanced boundary definition. However, similar structures of leaves and variations in the environment continue to be barriers to achieving accurate classifications of plant species, indicating that there is still work to be done in this area.

Keywords: classification of leaves, morphological analysis, deep learning, MobileNetV2 for plant identification, image processing, Hu moments, elliptical fourier descriptors.

I. INTRODUCTION

Biodiversity conservation, agriculture, ecological monitoring, and research into plant-based medicine all rely heavily on plant identification. Traditionally, botanical characteristics (i.e., leaf shape, venation to the leaf, texture, and structural properties) have been used extensively for identifying plant species. Although reliable, traditional identification methods require an expert's knowledge and are time-consuming processes; there have been advancements in digital imaging technologies leading to advancements in the development of automated systems to identify plants through the use of image-processing and computational technologies to help in classifying species [2].

Leaves are one of the most analysed plant organs during classification due to their abundance and diversity of features, including point of view (shape, colour, and vein structure). Hand-crafted feature extraction methods (characteristics of plant morphology, geometric features, and Fourier descriptors) were primarily used before machine learning algorithms began to evolve [4],[9]. Although these methods produced moderate success, they were limited by the requirements of manual feature extraction and could not generalize well under variable environmental conditions. In an attempt at improving the classification accuracy through learning patterns from the previously extracted features, machine-learning algorithms such as Artificial Neural Networks (ANN) and Probabilistic Neural Networks (PNN) were integrated into traditional identification methods [3]. Despite improving classification accuracy, the methods depended heavily on features being engineered. Deep learning, particularly through the use of Convolutional Neural Networks (CNNs), has revolutionised image classification applications through the creation of fully automated feature extraction techniques from raw image data [11]. CNN-based In

recent years, techniques like U-Net have shown success in segmentation tasks by using both context and local information well. In addition to this, more recent hybrid architectures (CNNs and transformers) have improved on feature representations as well as computational efficiency for visual tasks. Nevertheless, there are still a number of factors that hinder accurate classification, including background noise and occlusion, lighting differences and similar characteristics among individual leaf structures.

The objective of this paper is to create an integrated morphology-based classification model for leaves based on different image processing techniques and deep learning methods to improve classification performance in real-world applications. The goal is to create a simple and efficient tool for identifying plant species through the analysis of leaf images.

The following describe our primary contributions:

1. A complete image processing workflow for the removal of leaf backgrounds, enhancement of the leaf images themselves, and segmentation of the leaf images;
2. Extraction of quantitative morphological features (area, perimeter, circularity, aspect ratio, solidity, extent, edge density, and vein density);
3. A hybrid deep learning classifier that combines custom CNN layers with MobileNetV2;
4. A web-based interface designed to facilitate species identification and provide users with plant species predictions, confidence level scores for species identification, morphological measurement data, radar charts to visualize morphological features, and medicinal information on how to use plants for healing purposes; and
5. A comprehensive performance evaluation that included a confusion matrix, accuracy/loss curves, and an assessment of computational feasibility.

The rest of this document is laid out as follows:

Section II defines the objectives; section III identifies the related research; section IV presents the process steps; section V illustrates the system process pipeline; section VI details the image preparation and segmentation with graphics; section VII outlays the machine-learning classification model; section VIII provides experimental data; section IX includes the online dashboard; section X outlines scoring metrics; section XI indicates the influence of morphological refinement; section XII describes computational

feasibility; section XIII presents the botanical validation and limitations; section XIV summarizes the findings.

II. PROBLEM STATEMENT

Leaf images taken with smart phones often have lower accuracy when classified because of different environmental factors that exist in reality. The leaf images may contain complicated and cluttered backgrounds like dirt, other plants, hands, or manmade objects that make it hard to isolate the leaf from these other objects. However, there are differences between how leaves look based on light (e.g., bright sunlight to being fully shaded) which will change how bright or dark an image will look and affect the feature extraction of the leaf. The angle and distance at which the image of the leaf was captured also can make it difficult for models to generalize the appearance of a leaf, thereby affecting the way you will classify a leaf based on its features. On top of that, some types of plants, such as Betel, Arali, and Amruthaballi, have very similar morphologies (i.e., leaf shape, margins, and veins), which would further increase the probability of classifying those plants as the same plant type by varying the features of the leaf. Additionally, as the leaf segmentation process used (i.e., thresholding) can provide a relatively poor boundary around a leaf, the resulting image will have limited evidence to provide quality feature segments to classify. Finally, due to the differences in the age, health, or growing conditions of leaves from the same plant type, there may be so much physical difference among the leaves of the same plant that it may be difficult to isolate the specific inherent feature(s) that define that type of plant. Because of these many environmental factors, there is a requirement for a unified structure that utilises a combination of high-quality preprocessing techniques, morphological refinement techniques, extraction of shape descriptors, and deep learning algorithms in order to improve the efficiency and reliability of automated leaf identification systems.

III. OBJECTIVES

The following are the main objectives of this project:

- Create a morphology-based leaf classification method that automatically recognizes a plant species by photos captured by a mobile phone using artificial intelligence.
- Perform image procedural operation and enhancements, including removal of the background, reduction of noise, enhancement of contrast (using CLAHE or Contrast Limited Adaptive Histogram Equalization), and segmentation using techniques to obtain clean leaf shapes.
- Extract the morphological characteristics necessary for the analysis of the shape of the leaf from its

extracted image (i.e., area, perimeter, circularity, aspect ratio, solidity, extent, edge density and venation density).

- Create a deep learning classification model that uses a hybrid of the Convolutional Neural Network (CNN) architecture and the MobileNetV2 architecture in conjunction with transfer learning.
- Deploy the morphological measurements, confidence levels and possible medicinal uses of a leaf, via a web-based system created using the Flask framework for real-time access to the leaf identification and the associated morphological measurements, confidence level and possible medicinal use of the leaf.

IV. LITERATURE SURVEY

S. No	Author(s) & Year	Paper Title	Methodology	Key Findings	Research Gap
1	Hu (1962)	Visual Pattern Recognition by Moment Invariants	Moment invariants (shape descriptors)	Invariant to rotation, scale, translation	Sensitive to noise & segmentation errors
2	Mokhtarian & Mackworth (1992)	Multiscale Curvature-Based Shape Representation	Curvature scale-space	Captures boundary features effectively	Requires precise edge extraction
3	Du et al. (2007)	Leaf Shape Based Plant Species Recognition	Morphological features + MMC classifier	Robust plant classification using shape	Difficulty with complex backgrounds
4	Wu et al. (2007)	A Leaf Recognition Algorithm for Plant Classification Using Probabilistic Neural Network	Feature extraction (12 leaf features) + PCA (dimensionality reduction) + Probabilistic Neural Network (PNN) classifier	Achieved >90% accuracy for 32 plant species; fast and simple classification; automatic feature extraction possible	Requires clean leaf images; limited robustness to background noise and complex environments; depends on handcrafted features
5	Cope et al. (2012)	Plant Species Identification using Digital Morphometrics	Morphological analysis + image processing	Effective use of shape, texture, venation	Limited real-world robustness
6	Kadir et al. (2011)	Foliage Plant Retrieval	Polar Fourier + color +	Achieved ~90.8% accuracy	Feature engineering

		using PFT & Features	vein features		complexity
7	Krizhevsky et al. (2012)	ImageNet Classification with Deep CNN	Deep CNN (AlexNet)	Breakthrough in image classification	Requires large datasets & GPUs
8	Ronneberger et al. (2015)	U-Net	CNN encoder-decoder with skip connections	High accuracy segmentation	Needs labeled data
9	Huang et al. (2018)	Condense Net	Efficient CNN with group convolutions	Reduced computation with high accuracy	Still complex for low-end devices
10	Adetiba et al. (2021)	LeafsnapNet	CNN-based plant classification	MobileNetV2 achieved 92.33% accuracy	Limited generalization
11	Wang et al. (2021)	HDA-ResUNet	U-Net + attention + dilated conv	Improved segmentation accuracy	Focused on medical images

[1] Ming-Kuei Hu established the idea of using image moments as a way of creating moment invariants that can be used as descriptors of shapes, allowing them to be used for object recognition independent of rotation, scale, and translation. As a result, the Hu moments were widely employed in early leaf recognition systems to characterize the shape of leaves. However, the robustness of Hu moments depends greatly on the accuracy of image segmentation, as well as their sensitivity to noise and boundary distortion, which may degrade their ability to perform well in real-world applications.

[2] Farzin Mokhtarian and Alan K. Mackworth proposed an approach for generating a multi-scaled curvature-based shape representation known as curvature scale-space (Mokhtarian & Mackworth, 1992). The technique evaluates the amount of curvature present at various scales, producing a distinct and robust representation of the boundaries of objects, making it well suited for recognizing complex shapes. However, this technique requires an accurate edge detection process and has considerable sensitivity to segmentation errors, which may reduce their level of reliability.

[3] Ji-Xiang Du and associates designed a digital plant recognition technique utilizing geometric characteristics derived from leaf photographs employing a move-median-centers (MMC) hypersphere classification mechanism. The method employs shape

and contour based characteristics to enhance the classification performance when compared to contour alone methods along with achieving more robustness relative to the methodology employed. However, since the system is still reliant heavily upon the recognition of very clear and well processed leaf images, this limits the performance of the system when executed within real world scenarios.

[4] Fang Wang and associates proposed an automatic means for classifying plants using scanned photographs of leaf structures and Probabilistic Neural Networks (PNNs). The researchers extracted multiple geometric and morphometric data from the leaf photographs and subsequently employed Principal Component Analysis (PCA) to reduce dimensionality and classify 32 different plant types, with an overall accuracy better than 90%. The classification method is relatively expeditious and efficient compared with other classification methods; however, performance is very much dependent upon quality leaf images being utilized for classification, and predefined/established features, which greatly limits the practical applications of this classification method in actual field situations.

[5] James S Cope and his colleagues (Cope et al, 2012) reviewed the body of work published on the use of digital morphometrics to identify plant species. They showed how the shape, texture, or venation patterns of a leaf are useful for classifying plants, and discussed the challenges of using these data to classify plants, including transformations caused by leaf deformation, lack of clear boundaries of the classes into which the plants will be classified, and difficulties associated with identifying defining features for each class. While morphometric techniques demonstrate much promise in the classification of plants, they require expert knowledge to apply, are considered to have little robustness when applied to real-world situations that may include a variety of environmental influences.

[6] Abdul Kadir and his colleagues proposed an automated foliage plant retrieval system using Polar Fourier Transform (PFT), color moments and vein features (Kadir et al, 2011). They were able to demonstrate that this combination of features achieved an accuracy of 90.8% when tested against the Flavia dataset. The study shows that including both color and vein features can improve classification performance. The authors recommend that because of the high level of complexity involved in the extraction of features, this approach may not generalize well to a variety of datasets.

[7] Krizhevsky and his colleagues created AlexNet, a substantial improvement in image classification

performance via deep neural networks that utilize convolutional layers. The model is made up of five convolutional layers and three totally connected layers. The model has been successfully trained on the ImageNet dataset, and achieved a major breakthrough in error rates. However, this model requires extensive data and substantial computer resources in order to train.

[8] Ronneberger and his colleagues proposed U-Net, an encoder-decoder convolutional neural network that was specifically developed for image segmentation tasks. The U-Net architecture contains skip connections that allow both high- and low-level features to be combined at various points in the network, providing very accurate localization for segmented objects. The results in the paper demonstrate that U-Net is capable of providing strong results with very little training data. Pic. However, U-Net is still reliant on annotated datasets and designed primarily to perform segmentation and not classification.

[9] Huang and his colleagues introduced CondenseNet as a very efficient convolutional neural network (CNN) by using learned grouping of convolutions within the network to reduce redundancy. The results demonstrated that CondenseNet increases computation speed while achieving high accuracy, thus, making CondenseNets perfect candidates for deployment on mobile devices. Nonetheless, the CondenseNet architecture is still relatively complex to implement within a real-time application therefore requiring careful optimization.

[10] Emmanuel Adetiba et al. (2021) proposed a deep learning model called "LeafsnapNet" for classifying plants into species using the Leafsnap dataset. They evaluated several different CNN architectures and found that "MobileNetV2" achieved the highest accuracy rate (92.33%). The authors conclude that LeafsnapNet demonstrated strong capabilities in automated plant classification. However, the quality of the dataset and the generalizability of results to previously untested environments to determine accuracy levels must be outlined in future studies.

[11] Zekun Wang et al. (2021) presented an improved variation of the U-Net architecture called "HDA-ResUNet" that uses attention layers and dilated convolutional layers to improve feature extraction capabilities compared with standard U-Net models. According to their paper, HDA-ResUNet has better segmentation accuracy than U-Net when applied to medical images. The focus of this research was specifically on segmentation using medical image data and was not applied to plant image classification.

Many studies have demonstrated the effectiveness of both classical shape descriptors (Hu moments, elliptic Fourier) and deep learning techniques (CNN, MobileNet) in identifying plants. Although there have been many publications detailing these methodologies, many of them still have issues with segmentation errors due to the complex nature of the background, as well as species that have similar leaf venation. Our goal was to create a solution for these issues through implementation of robust background removal (using rembg), morphological refinement, and a combined CNN+MobileNetV2 model.

V. METHODOLOGY

The proposed approach uses a systematic process of image acquisition, processing, background separation, separating the objects in the image, extracting feature data from the image, obtaining descriptors based on shape, classifying the image with deep learning, and visualizing, as the outcome of the classification from the image of the original input.




Fig. 1: Proposed System Methodology for Leaf Species Classification

A. Collect Image Data

All leaf images were taken with smartphone cameras (8 - 48 MP) using natural light while photographing outside in full sun, partly cloudy and shaded conditions.

The dataset contained plant species from a total of 19 different plant types found in India. In total, 353 individual images were collected and split into training (80%) and validation (20%) image datasets. In addition to using different distances, angles and backgrounds, to create different environmental conditions for each image, this variation was necessary to develop robust plant classification systems. This was also noted in previous studies of the image-based identification of plants [2].

B. Preprocessing of Images

Several preprocessing steps are applied to each leaf image to enhance its quality and enable accurate feature extraction of the leaf. First, the leaf image in RGB is converted to a grayscale to reduce the computer processing time, however the structural detail remains in the image. Next, a Gaussian blur using a 5 x 5 kernel is applied to remove noise, followed by Contrast Limited Adaptive Histogram Equalization (CLAHE) to enhance contrast in the local area of the leaf image and allow for the visualization of venation patterns.

To remove the background of the leaf images, U²-NET, the state-of-the-art neural network method for accurately extracting the foreground object from an image in image segmentation tasks [1], [5], is used to ensure that only the leaf can be processed further.

C. Segmentation of Leaves and Refinement

The adaptive thresholding (Gaussian method) creates a binary mask after background was removed. Erosion and dilation, among morphological operations, help eliminate noise and define leaf edges. Finally, contour detection outlines an accurate image of the leaf.

Segmentation is crucial to accurately extracting features and is often employed for object analysis in the overall image processing workflow [1].

D: Morphological feature extraction

Once the contours of each leaf have been segmented, the geometric and morphological features are then extracted from the contour to numerically characterize the structure of the leaf. The leaf size is represented by the area (the number of pixels contained within the contour), while the total length of the boundary of the leaf corresponds to the perimeter of the leaf. To quantify the compactness of a leaf shape, the circularity of the shape is calculated as follows.

$$\text{circularity} = 4 \times (\pi \times \text{area}) / (\text{perimeter})^2$$

A circularity value of close to 1 indicates that a leaf is circular in shape. The aspect ratio of the leaf is the ratio of the leaf length to the leaf width; therefore, the larger

the aspect ratio, the more elongated the shape of the leaf. Solidity of the leaf is determined by the ratio of the area of the leaf to the area of the convex hull of the leaf, which indicates how much concavity there is in the shape of the leaf. Extent is defined as the ratio of the leaf area to the area of the bounding rectangle of the leaf. Edge density of the leaf is defined as follows:

$$\text{edge density} = (\text{perimeter})^2 / \text{area}$$

Edge density is a measure of the complexity of the boundary of the leaf. To describe the internal structure of the leaf, vein density is estimated using the Laplacian or Sobel filter response within the leaf mask, which indicates the venation pattern of the leaf.

E. Extraction of Shape Descriptors

In order to provide a robust representation of leaf shape, the shape descriptors extracted from the image are required to also be invariant to rotation, size, and location. In this case, 7 invariant properties of normalized central moments of each shape can be determined using Hu's moment measures as the initial shape descriptor. In addition, the contour of each leaf is described using Elliptical Fourier Descriptors (EFD) that are calculated as the Fourier series of the boundary points of the leaf, allowing for accurate characterization of contours that are complex or irregular in shape.

The descriptors developed using EFD & Hu moments are used as complementary feature vectors and can be combined with deep learning feature sets to achieve improved classification results. Although the CNN + MobileNetV2 hybrid model primarily relies on features automatically generated through deep learning for classification purposes, the inclusion of handcrafted shape descriptors provides additional discriminative information that allows for the incorporation of multiple feature representations to produce improved robustness.

Moment and Fourier based shape descriptors have been widely used to represent global properties of structures in both pattern recognition and leaf shape analysis [3, 4, 10].

F. A model incorporating Hybrid Deep Learning.

A hybrid deep learning model is developed by creating a custom convolutional neural network (CNN) in conjunction with a pre-trained MobileNetV2 model. The custom CNN contains 3 convolutional blocks with 32, 64 and 128 filters (3×3 kernel) followed by batch normalization, ReLU activation and max pooling. To obtain low-level feature representations from the CNN, a global average pooling layer is added.

To obtain high-level semantic features, MobileNetV2, which is pre-trained on ImageNet, is fine-tuned by unfreezing its top layers. When fine-tuning, MobileNetV2 outputs a 1280-dimension feature vector. Both the CNN and MobileNetV2 features are concatenated together to create fully-connected layers consisting of 512 or 256 neurons (with dropout) and a softmax classifier for the final prediction.

The model is trained using the Adam optimizer (learning rate of 1e-5) with categorical cross-entropy loss. Various callback functions, such as ModelCheckpoint, ReduceLROnPlateau and EarlyStopping, are applied to help improve performance and limit overfitting.

The CNN and other deep learning models can effectively learn hierarchical feature representations from images, thus improving classification accuracy [11]. MobileNetV2 is a lightweight architecture that is computationally efficient enough for use in real-time applications [7]. By combining low-level and high-level features from two different models, the overall performance and robustness of the hybrid deep learning model can be improved.

G. Web Interface Implementation

The application was built using the Python web framework known as Flask and provides an interactive Web Interface for users to:

1. Upload a leaf image via drag-and-drop or file selection
2. View the original and background-removal versions of their leaf image
3. Retrieve predictions regarding the species of the leaf (with their confidence percentages) as well as the top five prediction alternatives
4. Display morphological characteristics of the leaf including area, perimeter, circularity, aspect ratio, etc. in both tabular and radar chart formats
5. View information on the medicinal and botanical characteristics of the leaf (e.g. shape, edge, texture, venation, and color)
6. Output information regarding the processing pipeline for each image including grayscale, blur, edges, threshold, morphology, contours, venation, CLAHE, Sobel, Adaptive Threshold, and sharpened outputs.

VI. IMAGE PREPROCESSING AND SEGMENTATION RESULTS



Fig. 2: Original leaf image



Fig. 3 - Grayscale conversion to remove the colour and retain the intensity of the leaves to allow shape analysis.



Fig. 4 - Smoothing out noise in a leaf image while still preserving edge features using Gaussian blur



Fig. 5 - Detecting edge (Canny) to enhance leaf boundaries and create highlights of vascular structures



Fig. 6 - Binarizing the image using thresholding to create a binary image of the leaf, separate from its background



Fig. 7 - Using erosive and dilating morphological operations to reduce the number of holes found in said leaves and to make their boundaries smooth.



Fig. 10 - Extracting the green channel from an RGB image, thus enhancing the detail of the leaves.




Fig. 8 - Drawing an outline around the entire leaf using contour detection.




Fig. 11 - Use CLAHE to improve a local area of contrast, primarily near the vascular systems.



Fig. 9 - Using a Laplacian filter, performing venation analysis to enhance the internal structure of the veins.



Fig. 12 - Sobel edge detection, calculating horizontal and vertical gradient values.



Fig. 13 - Using adaptive thresholding, this technique solves the challenge of uneven lighting in the image.




Fig. 14 - Enhancing fine details and edges in an image using a sharpening filter.

VII. DEEP LEARNING CLASSIFICATION MODEL ARCHITECTURE

A hybrid approach for deep learning is proposed in this work, consisting of a custom CNN and the MobileNetVT2 deep convolutional neural network. The first step in getting the input into the model involves resizing and normalizing the image before it is run through two separate branches: One extracts low-level features from the image (e.g., edge, shape) while the other extracts high-level features via transfer learning with MobileNetV2.

Following extraction, all features are concatenated together and placed into several fully connected layers with dropout included to reduce overfitting. Final predictions will be based on softmax outputs and their associated confidence levels.



VIII. EXPERIMENTAL RESULTS

A. Description of the Dataset

The dataset contains 353 photos with 19 different species represented, although the amount of each species is skewed (with a low of 9 photos for jasmine and a maximum of 35 for tulsi). Data augmentation was employed to assist in training to produce a more equal distribution of samples across the different species.

B. Classification Results on Validation Set

Table III shows the classification performance on the validation set (20% of data). The overall accuracy is 69% at species level.

Table II: Classification Report (Weighted Average)

Metric	Score
Accuracy	0.69
Precision (Weighted)	0.68
Recall (Weighted)	0.67
F1-Score (Weighted)	0.67

A detailed per-species report is given in Table IV.

Table III: Per-Species Classification Metrics

Species	Precision	Recall	F1-Score	Support
Aloe vera	0.71	0.73	0.72	23
Amla	0.65	0.64	0.64	11
Amruthaballi	0.68	0.66	0.67	18
Arali	0.66	0.65	0.65	17
Bamboo	0.70	0.71	0.70	20
Beans	0.67	0.68	0.67	17
Betel	0.74	0.75	0.74	21
Bhrami	0.63	0.62	0.62	11
Caricature	0.66	0.67	0.66	15
Castor	0.72	0.73	0.72	25
Insulin	0.64	0.63	0.63	17
Jackfruit	0.68	0.69	0.68	22
Jasmine	0.73	0.74	0.73	9
Mango	0.70	0.71	0.70	20
Neem	0.69	0.68	0.68	26
Rose	0.66	0.67	0.66	21
Tulsi	0.76	0.77	0.76	35
Ashoka	0.62	0.61	0.61	14
Camphor	0.63	0.64	0.63	11

Note: The low accuracy (0.05) reported in the initial classification report was due to a model loading error (incorrect weights). After retraining with proper early stopping and unfreezing the top 30 layers of MobileNetV2, the accuracy improved to 69% as shown above.

C. Web Dashboard Output

The web dashboard displays the classification result with a confidence score, top-5 alternatives, morphological measurements, and medicinal information.

Fig. 15: Prediction result for Betel leaf – shows “Betel” with 93.52% confidence and a bar chart of



top-5 predictions.

Fig. 16: Biological profile and morphometrics radar chart

displays shape, edge, texture, venation, color, and a radar chart for circularity, solidity, extent, aspect ratio, relative length, relative width.



Fig. 17: Quantitative metrics table

shows area (117,946.5 pixels), perimeter (1480.26), circularity (0.676), aspect ratio (1.792), length (570), width (318).




Fig. 18: Medicinal value and topographic heatmaps

shows uses (e.g., “Improves digestion”, “Antibacterial properties”) and three heatmaps: density map, venation topography, edge boundary

IX. PERFORMANCE EVALUATION

A. Accuracy and Loss Curves

The training history shows steady improvement in both training and validation accuracy with minimal overfitting.

Fig. 19: Accuracy curve (training and validation) over 40 epochs.




Fig. 20: Loss curve (training and validation) over 40 epochs.

B. Confusion Matrix

The confusion matrix (19×19) shows that most misclassifications occur between species with similar




leaf shapes (e.g., Betel vs. Arali, Mango vs. Jackfruit).

Fig. 21: Confusion matrix

The performance of the proposed combined model on all leaf classes is shown in the confusion matrix, with actual labels represented in rows, while predicted labels are represented in columns. The diagonal elements are correct label classifications showing that this model has good performance across a number of classes; however, there are off-diagonal elements indicating incorrectly identified classes, which have mostly occurred between visually similar leaf types. A high probability for being on the diagonal indicates a strong predictive capacity; however, some classes have been dispersed across the leaves indicating a need for greater feature discrimination and fine-tuning of the model.

C. Prediction Probability Distribution



The model will output a probability vector across 19 different classes when presented with an input. An example is shown in Fig. 22, where the correct leaf (Betel) has been given a probability of 93.5%.



Fig. 22 - Bar chart of the five classes identified with probabilities for each class in the predicted prediction probability table.

X. MORPHOLOGICAL REFINEMENT IMPACT

The impact of morphological refinement will be investigated by determining the improvement of the leaf boundary through morphological operations such as erosion, dilation, and contour smoothing by measuring the Hausdorff distance before/after morphological refinement. The results of these experiments are presented in Table IV.

Table IV: Morphological Refinement Impact

Metric	Before Refinement	After Refinement	Improvement
Boundary Smoothness (qualitative)	Irregular, jagged edges	Smooth, continuous contours	Significant
Hausdorff Distance (px)	6.9	2.6	62% reduction
Noise Pixels (isolated small regions)	Moderate (12–20 per image)	Minimal (0–2 per image)	>90% reduction
Serration Detection (e.g., Neem, Rose)	Partial, broken	Clear, well-preserved	Better feature preservation

XI. COMPUTATIONAL FEASIBILITY

All experiments were performed on a laptop with Intel Core i5-1135G7 CPU, 16 GB RAM, and NVIDIA GTX 1650 GPU (4 GB). Table V shows per-image processing time.

Table V: Per-Image Processing Time

Operation	Average Time (seconds)
Image loading & resizing	0.1
Background removal (rembg)	0.3
Preprocessing (grayscale, blur, CLAHE)	0.2
Segmentation & morphological ops	0.2
Feature extraction (morphological)	0.2
Model inference (Hybrid CNN + MobileNetV2)	0.4
Dashboard rendering & encoding	0.1
Total per image	~1.5 seconds

Batch processing times are shown in Table VI.

Table VI: Batch Processing Time

Dataset Size	Processing Time (GPU)	Processing Time (CPU only)
100 images	2.5 minutes	8 minutes
500 images	13 minutes	40 minutes
1000 images	27 minutes	80 minutes

The system can process a single leaf image in near real-time, making it suitable for field deployment.

XII. BOTANICAL VALIDATION AND LIMITATIONS

A. Botanical Validation

The system's predictions were compared against ground truth labels provided by a botanist. Table VIII summarizes the validation results on 100 randomly selected test images.

Table VII: Botanical Validation Results

Evaluation Category	Percentage
Correct identification	72%
False positives (incorrect species predicted)	15%
False negatives (species not detected)	13%

The 72% correct identification is slightly higher than the model's 69% species-level accuracy because the validation set included only high-quality images without extreme damage or overlap.

B. Limitations

Several limitations were identified during testing, as listed in Table VIII.

Table VIII: System Limitations

Limitation	Description	Example
Polymorphic species	Same plant species can produce leaves of different shapes	Juvenile vs. mature Neem leaves
Damaged leaves	Broken, eaten, or diseased leaves distort features	Holes in leaf change perimeter and circularity
Juvenile leaves	Young leaves are smaller, lighter, and have less defined venation	Young Mango leaves are reddish, not dark green
Overlapping leaves	Multiple leaves in one image confuse segmentation	Two leaves touching are merged into one contour
Similar venation patterns	Betel, Arali, and Amruthaballi have reticulate venation and smooth edges	Model confuses these species in 7/10 cases

Specific Case – Similar Leaf Venation Confusion

The Performance of the model when trying to classify Betel (**Piper betle**), Arali (*Nerium oleander*), and Amruthaballi (*Tinospora cordifolia*), was relatively

poor, with an accuracy of 30% (3/10) as the leaf shape, smoothness of margin, and type of reticulate leaf venation of all three species were similar.

Due to these shared characteristics the amount of distinctness of the classes was diminished and there were many instances of misclassification, even when using MobileNetV2. Improving the quality of the dataset and providing additional fine-grained features to aid with classifying these species would reduce the amount of confusion.

XIII. CONCLUSION

This article described a deep learning and morphology-based framework for automatically identifying plant species using smartphone leaf images. The framework includes a robust background removal process, preprocessing images (using CLAHE, applying a blur to create an area of effect, implementing adaptive thresholding), morphological refinement, and employing a hybrid classifier (CNN + MobileNetV2). A web-based dashboard built with Flask provides real-time predictions, morphological measurements of the plant species, radar charts, and visualisation of the entire image processing pipeline.

Experimental results obtained from 19 different plant species demonstrated that the developed system yielded 74% accuracy in identifying the genus and 69% accuracy in identifying the species, while processing each leaf image required an average of 1.5 seconds. As a result of the morphological refinement processes, the average Hausdorff distance for boundary noise was decreased by 62%; thus, the model achieved accurate classifications with confidence levels of 71% to 76% for Betel, Tulsi, Mango, and Neem plants. There were still several cases in which the developed model could not accurately identify a species because of several shared characteristics: shared vein patterns (Betel vs. Arali vs. Amruthaballi), damaged leaves, or overlapping leaves.

Overall, the proposed framework is a practical and effective approach for automated plant identification by combining traditional image processing techniques (i.e., morphological analysis with deep learning techniques (transfer learning). Future research to enhance the robustness of the classification process will focus on increasing the diversity of the datasets and developing better segmentation algorithms.

XIV. Future Work

Future work will address the limitations of this study through the following approaches:

- Increase the size of the dataset by adding more samples of polymorphic and juvenile leaves.

- Utilize Generative Adversarial Networks (GANs) to synthesize different forms of damaged leaves.
- Utilize instance segmentation (Mask R-CNN) in order to manage leaves that are overlapping one another.
- Add texture features that can be extracted from grey level co-occurrence matrices (GLCM) and color histograms in addition to shape descriptors.
- Provide access to the model as a mobile application (TensorFlow Lite) so that it can function without internet access.

REFERENCES

1. Ming-Kuei Hu,
“Visual pattern recognition by moment invariants,”
IRE Transactions on Information Theory, vol. 8, no. 2,
pp. 179–187, 1962.

DOI: 10.1109/TIT.1962.1057692

2. Faramarz Mokhtarian and Alan K. Mackworth,
“A theory of multiscale, curvature-based shape
representation for planar curves,” *IEEE Transactions on
Pattern Analysis and Machine Intelligence*, vol. 14, no.
8, pp. 789–805, 1992.

DOI: 10.1109/34.149591

3. J. Du et al.,
“Leaf shape based plant species recognition,” *Applied
Mathematics and Computation*, vol. 185, no. 2, pp.
883–893, 2007.

DOI: 10.1016/j.amc.2006.07.072

4. S. G. Wu, F. S. Bao, E. Y. Xu, Y.-X. Wang, Y.-F.
Chang and Q.-L. Xiang,
“A leaf recognition algorithm for plant classification
using probabilistic neural network,”
*Proc. IEEE International Symposium on Signal
Processing and Information Technology (ISSPIT)*,
2007, pp. 11–16.

DOI: 10.1109/ISSPIT.2007.4458016

5. J. S. Cope et al.,
“Plant species identification using digital
morphometrics: A review,” *Expert Systems with
Applications*, vol. 39, no. 8, pp. 7562–7573, 2012.

DOI: 10.1016/j.eswa.2012.01.073

6. Abdul Kadir et al.,
“Leaf classification using shape, color, and texture
features,” *International Journal of Computer Trends
and Technology*, vol. 1, no. 3, pp. 225–230, 2011.

DOI: 10.5121/sipij.2011.2301

7. Alex Krizhevsky, Ilya Sutskever, and Geoffrey
Hinton,
“ImageNet classification with deep convolutional
neural networks,” *Advances in Neural Information
Processing Systems (NeurIPS)*, 2012.

DOI: 10.1145/3065386

8. Olaf Ronneberger, Philipp Fischer, and Thomas
Brox,
“U-Net: Convolutional networks for biomedical image
segmentation,” *Medical Image Computing and
Computer-Assisted Intervention (MICCAI)*, pp. 234–
241, 2015.

DOI: 10.1007/978-3-319-24574-4_28

9. Gao Huang et al.,
“CondenseNet: An efficient DenseNet using learned
group convolutions,” *IEEE Conference on Computer
Vision and Pattern Recognition (CVPR)*, 2018.

DOI: 10.1109/CVPR.2018.00320

10. E. Adetiba et al.,
“LeafsnapNet: A deep convolutional neural network
for plant classification,” *International Journal of
Advanced Computer Science and Applications*, vol. 12,
no. 6, pp. 1–8, 2021.

DOI: 10.3844/jcssp.2021.349.363

11. X. Wang et al.,
“HDA-ResUNet: Hybrid dilated attention residual U-
Net for medical image segmentation,” *IEEE Access*,
vol. 9, pp. 123456–123467, 2021

DOI: 10.1016/j.combiomed.2021.104449