

Medicinal Leaf Recognition

Bharath L N, Preethi B P, Sushma C Y, Yashwanth A N

Rajeev Institute of Technology, Hassan, 573201, Karnataka

Approved by AICTE, New Delhi and Affiliated to

Visvesvaraya Technological University, Belagavi, 590018, Karnataka

Email: 4ra22ci006@rithassan.ac.in, 4ra22ci039@rithassan.ac.in, 4ra22ci056@rithassan.ac.in,

4ra22ci062@rithassan.ac.in

Abstract — Identifying medicinal plants from leaf images is a task with deep roots in traditional healthcare, yet one that remains technically demanding due to the enormous variation in leaf morphology, color, and surface texture across species. This work introduces a lightweight deep learning system built around the MobileNetV2 architecture to automatically recognize and classify medicinal leaves from photographs. Because MobileNetV2 was designed with mobile and edge deployment in mind, it strikes a practical balance between classification accuracy and computational demand. A curated dataset covering a diverse set of medicinal plant species was assembled and subjected to a systematic preprocessing pipeline that included grayscale conversion, Gaussian noise suppression, Otsu thresholding, and contour-based background removal, followed by augmentation to expand training variety. Feature extraction relied on Sobel edge filters and contour analysis, and the resulting descriptors were fed to the fine-tuned MobileNetV2 classifier. The system reached a validation accuracy of 96 percent across 30 plant classes, confirming that mobile-friendly neural networks can meet the accuracy bar needed for practical ethnopharmacological applications. A Streamlit-based web interface was also built so that users can upload a leaf image and receive its common name, scientific name, medicinal properties, and known allergens in real time.

Keywords — *Medicinal Plant Recognition, MobileNetV2, Deep Learning, Transfer Learning, Image Preprocessing, Leaf Classification, Ethnopharmacology, Convolutional Neural Network.*

I. Introduction

Plants have served as humanity's oldest pharmacy. Long before synthetic compounds entered clinical practice, communities across Asia, Africa, and the Americas relied on leaves, roots, and bark to manage disease.

Today, roughly 80 percent of the global population still depends on plant-based remedies as a primary or complementary therapy [1]. Despite this enduring relevance, the accurate identification of medicinal species remains difficult for non-specialists. A single genus can contain dozens of morphologically similar species, some therapeutically valuable and others potentially toxic, making visual misidentification a genuine public-health concern.

Traditional identification methods depend on trained botanists who interpret a combination of macroscopic features — leaf shape, venation pattern, surface texture, and color — alongside microscopic and chemical analyses. This expert knowledge is concentrated in a shrinking pool of practitioners and is rarely accessible in rural or under-resourced settings where herbal medicine is most prevalent. Automated image-based identification offers a scalable alternative: a user photographs a leaf with a smartphone, and a trained model returns a species prediction within seconds.

Convolutional Neural Networks (CNNs) have transformed plant identification over the past decade by learning hierarchical visual features directly from raw pixel data rather than relying on handcrafted descriptors [2]. Among CNN families, MobileNetV2 stands out for its efficient use of inverted residual blocks and linear bottleneck layers, which sharply reduce parameter count without sacrificing representational capacity. This makes it particularly well suited to scenarios where inference must run on a mobile device or a low-power server.

The present paper describes a complete pipeline — from data collection and preprocessing through model training and interactive deployment — for recognizing 30 medicinal leaf species. The remainder of the paper is organized as follows: Section II surveys related work; Section III describes the dataset; Section IV details the

methodology; Section V presents results; Section VI discusses limitations; and Section VII concludes with directions for future work.

II. Related Works

The problem of automated plant identification from leaf images has attracted sustained research attention over the past fifteen years, and the literature reveals a clear

The availability of large-scale image datasets and GPU computing subsequently enabled CNN-based approaches. Valdez et al. [4] applied transfer learning with various backbone architectures to a ten-class medicinal plant dataset and demonstrated that pretrained features generalize well even to botanical imagery, provided the training set is sufficiently diverse.

Patel and Chaware [5] investigated model compression through quantization of MobileNet classifiers, showing that 8-bit integer inference introduces negligible accuracy degradation while enabling deployment on constrained hardware — a finding directly relevant to field-use scenarios.

Sharab et al. [6] benchmarked a VGG-16-based system on a broader collection of medicinal plants and highlighted the persistent challenge of inaccurate predictions when leaf conditions deviate from training examples (occlusion, partial damage, unusual lighting).

Collectively, these works motivate the choice of MobileNetV2 as the backbone in the current system: it inherits the accuracy gains of deep transfer learning while remaining deployable on commodity hardware, and it has not been extensively evaluated against the specific 30-species medicinal dataset used here.

III. Dataset

The dataset used in this study was obtained from a publicly available Kaggle repository of augmented medicinal leaf images [7]. It encompasses photographs of 30 distinct plant species that are widely recognized in Ayurvedic and traditional Indian medicine, including Neem (*Azadirachta indica*), Tulsi (*Ocimum sanctum*), Hibiscus (*Hibiscus rosa-sinensis*), Guava (*Psidium guajava*), Lemon (*Citrus limon*), and Mango (*Mangifera indica*), among others.

trajectory from handcrafted features toward end-to-end deep learning.

Early systems relied on geometric and statistical descriptors. Wu et al. [3] combined Principal Component Analysis (PCA) for dimensionality reduction with a Probabilistic Neural Network (PNN) classifier, achieving respectable accuracy on a controlled dataset but struggling with low-quality photographs and significant intra-class variation.

Images were collected under varied field conditions — different illumination angles, backgrounds, and growth stages — to reflect the diversity a mobile application would encounter in real-world use. Each class contains approximately 50 samples before augmentation. The dataset was partitioned into 80 percent for training and 20 percent for validation, with stratified sampling to maintain class balance across both splits.

Table I: Dataset Summary

Property	Value
Total Classes	30
Approx. Images per Class	~50 (pre-augmentation)
Training Split	80%
Validation Split	20%
Input Resolution	224 × 224 px
Color Space	RGB → Grayscale (preprocessing)

IV. Methodology

The proposed pipeline consists of four sequential stages: data preprocessing, feature extraction, model training and classification, and result generation. Figure 1 presents the overall architecture.

Figure 1: System Architecture

Dataset (Leaf Images)

**Pre-Processing**Rotation
Grayscale
Gaussian Filter

Feature Extraction (Sobel + Contours)



MobileNetV2 Classification



Result (Name + Medicinal Properties)

A. Data Preprocessing

Raw leaf images undergo a five-step preprocessing sequence before entering the feature extraction stage.

Rotation augmentation generates additional training samples at multiple orientations, alleviating the scarcity of labeled examples and reducing the model's sensitivity to how a leaf is held when photographed.

Grayscale conversion maps the three-channel RGB image to a single intensity channel using the luminosity formula: $Y = 0.30R + 0.59G + 0.11B$. This discards color information that varies with lighting and photography equipment while retaining the structural details — veins, margins, and overall silhouette — that distinguish species.

Gaussian filtering suppresses high-frequency noise by convolving the grayscale image with a 3×3 kernel whose coefficients follow a discrete Gaussian distribution (values 1, 2, 1; 2, 4, 2; 1, 2, 1, normalized by $1/16$). The smoothed result better supports subsequent edge detection.

Otsu thresholding converts the smoothed grayscale image to binary by automatically selecting the intensity value that minimizes intra-class variance between foreground (leaf) and background pixels. This produces a clean binary mask without requiring manual threshold selection.

Background subtraction uses the binary mask derived from the detected leaf contour to isolate the leaf from its background, setting non-leaf pixels to white. The result is a standardized, background-free leaf image ready for feature extraction.

B. Feature Extraction

Two complementary edge-detection techniques are applied to each preprocessed image. Sobel filtering computes the image gradient in the horizontal and vertical directions and combines them to highlight sharp intensity transitions — primarily leaf margins and primary veins. Contour analysis then traces the outer boundary of the thresholded leaf, providing a compact representation of its overall shape that is invariant to small texture variations. Together, these two representations supply the geometric and structural information that differentiates similar-looking species.

The feature-enriched images are resized to 224×224 pixels and normalized to the $[0, 1]$ range before being passed to the classifier.

C. Classification with MobileNetV2

MobileNetV2 [8] is a convolutional architecture designed around two innovations: *inverted residuals*, which expand the feature representation in a low-dimensional bottleneck before applying depthwise-separable convolutions and then project back down, and *linear bottlenecks*, which omit the ReLU nonlinearity at the projection step to prevent information loss in low-dimensional manifolds.

The network contains two types of building blocks. A stride-1 residual block processes features at constant spatial resolution and adds a skip connection between input and output. A stride-2 downsampling block halves spatial dimensions, enabling the network to build progressively more abstract representations. Within each block, the three convolutional layers are: (1) a 1×1 Conv + ReLU6 expansion, (2) a 3×3 depthwise Conv + ReLU6, and (3) a 1×1 linear projection without activation.

Transfer learning is applied by loading ImageNet-pretrained weights into the MobileNetV2 base and freezing them, then attaching a global average pooling layer, a 1024-unit dense layer with ReLU activation, and a softmax output layer with 30 units corresponding to the plant classes. Only the appended layers are trained initially, allowing the pretrained feature extractor to be

reused without overfitting on the relatively small medicinal leaf dataset.

The model is compiled with the Adam optimizer (learning rate 0.0001), categorical cross-entropy loss, and accuracy as the evaluation metric. Training runs for up to 15 epochs with early stopping based on validation loss.

Table II: MobileNetV2 Configuration

Component	Details
Base Model	MobileNetV2 (ImageNet weights)
Base Layers	Frozen during initial training
Pooling	Global Average Pooling
Dense Layer	1024 units, ReLU
Output Layer	30 units, Softmax
Optimizer	Adam (lr = 0.0001)
Loss Function	Categorical Cross-Entropy
Batch Size	32

D. Deployment Interface

A Streamlit-based web application wraps the trained model, offering a login system, an upload page where users submit leaf photographs, and a results page that displays the predicted species name alongside its common name, scientific name, documented medicinal properties, and known allergens. The interface is designed to be accessible to non-technical users, including farmers, community health workers, and amateur botanists.

V. Results and Analysis

The trained MobileNetV2 model was evaluated on the held-out validation set. The primary metric is top-1 classification accuracy. Figure 2 shows the training and validation accuracy curves over 15 epochs; both curves converge smoothly, with validation accuracy plateauing near 96 percent after approximately epoch 8, which suggests that the model has learned generalizable

representations rather than memorizing training examples.

Figure 2: Training & Validation Accuracy (schematic) — Training Acc -- Validation Acc 0 15 Epochs Accuracy

The loss curves (Figure 3) follow the expected pattern: training loss descends steeply in early epochs and levels off, while validation loss tracks closely with only minor divergence, indicating that regularization through augmentation and the small learning rate was effective in limiting overfitting.

Figure 3: Training & Validation Loss (schematic) — Training Loss -- Validation Loss 0 15 Epochs Loss

The confusion matrix (Figure 4) reveals that the vast majority of plant species are classified with near-perfect precision and recall. Minor confusion occurs between a small number of morphologically similar species — particularly those sharing elliptic leaf shapes and smooth margins — but these misclassifications are rare and do not materially affect the overall accuracy figure.

Table III: Classification Performance Summary

Metric	Value
Overall Validation Accuracy	96%
Number of Classes	30
Model Parameters (MobileNetV2)	~3.4 M
Training Epochs	15
Inference Time (per image)	< 200 ms (CPU)

These results compare favorably with prior work on similar tasks. The VGG-16 system described by Sharrab et al. [6] reported lower accuracy on a comparable 30-class problem, partly because the deeper and wider architecture is prone to overfitting on limited training data. MobileNetV2's lighter footprint, combined with transfer learning from ImageNet, appears to provide stronger regularization in the low-data regime typical of specialized botanical datasets.

VI. Limitations

Several constraints deserve acknowledgment. First, the dataset covers only 30 species, a small fraction of the estimated 70,000 medicinal plant species documented globally. Scaling the system to hundreds of classes would require substantially more labeled data and may necessitate architectural changes.

Second, the current preprocessing pipeline converts images to grayscale, discarding color information. While color introduces lighting-related variability, it is also a genuine diagnostic feature for certain species pairs, and future work could explore strategies that retain color selectively or use it as an auxiliary input channel.

Third, the model was trained on controlled or near-controlled photographs. Its accuracy on heavily occluded, partially damaged, or disease-affected leaves — which are common in field settings — has not been systematically evaluated.

Finally, the 96 percent figure reflects top-1 accuracy on a balanced validation set. In imbalanced real-world deployment, per-class precision and recall may differ, and a misclassification that leads a user to consume an incorrect plant could have health consequences. Uncertainty quantification — for instance, presenting the top-three predictions with associated confidence scores — would make the system safer in practice.

VII. Conclusion and Future Work

This paper presented a medicinal leaf recognition system that couples a systematic image preprocessing pipeline with the MobileNetV2 deep learning architecture and a transfer learning strategy. The system achieved 96 percent classification accuracy across 30 plant species and was packaged into an interactive Streamlit application capable of returning species identity and medicinal metadata from a single leaf photograph.

The core contribution is demonstrating that a mobile-friendly architecture, when combined with careful preprocessing and data augmentation, can meet the accuracy threshold needed for practical ethnopharmacological tools — without requiring the computational infrastructure associated with heavier networks such as VGG-16 or ResNet-50.

Several directions merit future exploration. Expanding the dataset to cover a broader taxonomic range, incorporating hyperspectral or multispectral imaging for richer feature representation, developing a native mobile application for offline inference, and integrating geographic location data to constrain the prior distribution over species are all promising avenues. In parallel, addressing the safety concerns raised in Section VI — through confidence thresholding, human-in-the-loop verification, and rigorous field testing — will be essential before deployment in healthcare-adjacent contexts.

References

- [1] World Health Organization, "WHO Traditional Medicine Strategy 2014–2023," WHO Press, Geneva, 2013.
- [2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [3] 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," in *Proc. IEEE Int. Symp. Signal Process. Inf. Technol.*, 2007, pp. 11–16.
- [4] D. B. Valdez, C. J. G. Aliac, and L. S. Feliscuzo, "Medicinal plant classification using convolutional neural network and transfer learning," in *Proc. IEEE Int. Conf. Artificial Intelligence Eng. Technol. (IICAET)*, Kota Kinabalu, Malaysia, Nov. 2022.
- [5] R. Patel and A. Chaware, "Quantizing MobileNet models for classification problem," in *Proc. 8th Int. Conf. Computing Sustainable Global Development (INDIACom)*, 2021, pp. 348–351.
- [6] Y. Sharrah, D. Al-Fraihat, and M. Tarawneh, "Medicinal plants recognition using deep learning," presented at the Int. Conf. Inf. Technol., Valencia, Spain, Jul. 2023.
- [7] Kaggle, "Medicinal Plant Dataset – Augmented," [Online]. Available: <https://www.kaggle.com/datasets/medicinal-plant-dataset-augmented>. [Accessed: 2024].
- [8] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and

linear bottlenecks," in *Proc. IEEE/CVF Conf. Computer Vision Pattern Recognition (CVPR)*, 2018, pp. 4510–4520.

[9] J. R. Janani and A. Gopal, "Identification of selected medicinal plant leaves using image features and ANN," in *Proc. IEEE Int. Conf. Advances in Electronic Systems (ISAES)*, 2018, pp. 238–242.

[10] R. Rajini and B. Veena, "Study on identification and classification of medical plants," *Int. J. Advances Sci. Eng. Technol.*, 2019.