

Original Research

Received 15 June 2026

Revised 30 July 2026

Accepted 12 Aug 2026

Natural Resources for Human Health 2026, Vol. 6 (16s): 557–570

KEYWORDS:

Deep Learning, Transfer Learning, ResNet-18, Medicinal Plants, Leaf Classification, Flower Recognition, Computer Vision.

Transfer Learning-Enabled Residual Networks for Fine-Grained Identification of Indian Medicinal Flora Across Leaves, Flowers, and Whole Plants

Balaji Laxmi Pati¹, Prabira Kumar Sethy², P Bharat Siva Varma³, Vallam Reddy Bhargavi Reddy⁴, Santi Kumari Behera⁵ & Aziz Nanthaamornphong^{*6}

¹ Department of Computer Science and Engineering, Sambalpur University Institute of Information Technology, Jyoti Vihar, Burla, Odisha 768019

² Department of Electronics and Communication Engineering, Sambalpur University Institute of Information Technology, Jyoti Vihar, Burla, Odisha 768019

³ Department of Computer Science and Engineering, SRKR Engineering College(A), Bhimavaram, West Godavari District, Andhra Pradesh- 534204

⁴ Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Andhra Pradesh, 522302, India

⁵ Department of Computer Science and Engineering, VSSUT Burla

⁶ College of Computing, Prince of Songkla University, Phuket campus, Thailand

Corresponding Author: Aziz Nanthaamornphong

Email: aziz.n@phuket.psu.ac.th

ABSTRACT: Accurate taxonomic identification of medicinal flora is vital for safeguarding indigenous biodiversity, standardizing Ayurvedic and herbal pharmacopoeia, and intercepting commercial botanical adulteration. Conventional manual identification by taxonomists is unscalable and prone to subjective error due to subtle intra-class morphometric variations and high inter-class phenotypic similarities. In this paper, a multi-organ botanical classification framework based on deep transfer learning is developed and evaluated across three distinct modalities: whole plants, plant leaves, and flowers. A fine-tuned ResNet-18 architecture with tailored dense projection heads and differential layer-wise learning rates was trained on three comprehensive Indian medicinal flora benchmarks comprising 40 whole-plant classes, 80 leaf classes (6,900 images), and 28 flower classes (6,213 images). An online stochastic augmentation pipeline—incorporating planar rotations ($\pm 20^\circ$), horizontal mirroring, and spatial translations—was deployed to prevent overfitting on background noise and preserve geometric invariance. Trained with the Adam optimizer over 8 epochs using stratified 70/15/15 partitions, the proposed pipeline achieved generalization test accuracies of 98.10% on whole plants, 94.79% on leaves, and 99.79% on flowers. These empirical outcomes confirm that deep residual representations generalise across distinct botanical organs, providing a computationally lightweight foundation for embedded and field-deployable botanical diagnostic tools.

1. INTRODUCTION

India possesses one of the world's richest medicinal floras, anchored historically in traditional therapeutic systems such as Ayurveda, Siddha, and Unani. An estimated 80% of developing populations rely on plant-derived raw extracts for primary

healthcare needs. However, the commercial herbal pipeline faces a major bottleneck: reliable species-level authentication at harvest and processing stages. Inadvertent misidentification or deliberate adulteration with morphologically analogous, inactive, or toxic weeds undermines pharmacological efficacy and jeopardizes consumer safety.

Historically, botanical taxonomy depends on manual inspection of key reproductive and vegetative organs, including leaf venation patterns, phyllotaxy, floral symmetry, petal color distributions, and overall plant habit. This paradigm requires specialized taxonomic training, is labor-intensive, and degrades when specimens present non-standard physical damage, drying artifacts, or developmental variations.

Recent breakthroughs in Convolutional Neural Networks (CNNs) have automated visual pattern recognition across agricultural and plant pathology domains. While several models target isolated organs—predominantly foliage due to its year-round availability—few automated systems address the complete morphological spectrum spanning leaves, flowers, and whole plants under an integrated experimental pipeline. Plant leaves present significant identification challenges owing to planar contours, overlapping vein structures, and subtle geometric variances. Flowers display sharp, discriminative chromatic and structural petal landmarks. Whole-plant imagery introduces real-world background clutters, stem geometries, and complex spatial growth habits.

To address these diverse morphological profiles, this paper presents a unified deep transfer learning system utilizing an adapted ResNet-18 architecture. The contributions of this work are threefold:

1. **Multi-Organ Benchmark Assessment:** We evaluate a unified deep residual learning pipeline across three distinct Indian medicinal benchmarks: a 40-class whole-plant dataset, an 80-class leaf dataset, and a 28-class flower dataset.
2. **Transfer Learning Pipeline:** We implement a layer-graph surgical adaptation of ResNet-18 that incorporates differential learning rate factors for fast head adaptation alongside on-the-fly affine augmentation to mitigate background overfitting.
3. **Rigorous Experimental Validation:** We report high-confidence classification accuracies (98.10% on plants, 94.79% on leaves, and 99.79% on flowers), analyzing how morphological variability directly affects deep feature discriminability across vegetative and reproductive organs.

2. LITERATURE REVIEW

Automated plant identification has transitioned from engineered hand-crafted visual descriptors to end-to-end deep feature extractors.

A. Hand-Crafted Feature Extractors

Early computer vision approaches relied on manual geometric parameterization and texture descriptors. Researchers widely used Hu's seven invariant moments, Fourier shape descriptors, Gray-Level Co-occurrence Matrices (GLCM), Scale-Invariant Feature Transform (SIFT), and Speeded-Up Robust Features (SURF) to describe leaf borders and venation networks. Classifiers such as Support Vector Machines (SVM), Random Forests (RF), and k -Nearest Neighbors (k -NN) were subsequently trained on these tabular descriptors. Although computationally lean, these pipeline-dependent systems were fragile; they failed under variable illumination, partial occlusions, and background field noise, and could not scale across diverse botanical genera.

B. Deep Convolutional Architectures and Transfer Learning

The advent of deep CNNs eliminated manual feature engineering by learning hierarchical spatial abstractions directly from raw pixel arrays. Rather than training models from scratch on scarce botanical datasets, transfer learning leverages networks pretrained on broad databases (e.g., ImageNet), fine-tuning their generic edge, color, and texture filters for specific botanical tasks.

Architectures such as VGG-16, VGG-19, Inception-v3, DenseNet, and ResNet have been adapted for agricultural tasks and disease recognition. In residual networks (He et al.), identity skip connections mitigate the vanishing gradient problem, allowing deeper layers to learn residual mappings:

$$\mathcal{F}(x) = \mathcal{H}(x) - x \quad (1)$$

This structural stability allows deep networks to capture delicate venation traits and subtle petal variations without numerical divergence or optimization stalling.

C. Organ-Specific Classification Studies

Most existing literature focuses on single-organ botanical identification:

- **Foliar Systems:** Leaves are available year-round, making them the most frequently studied plant organ. However, foliar identification suffers from high inter-class visual similarity—multiple distinct species share lanceolate, ovate, or elliptic green margins with similar secondary venation.
- **Floral Systems:** Floral structures provide dense chromatic and structural variation, yielding high classification precision; however, floral blooming periods are seasonal, limiting year-round identification.
- **Whole-Plant Imagery:** Whole-plant images provide structural habit and stem architecture, but introduce field clutter, variable canopy boundaries, and soil backgrounds.

Comparative, multi-dataset studies that evaluate the same underlying deep architecture across leaves, flowers, and whole plants under a standardized training paradigm remain sparse. This study bridges that gap by systematically applying an adapted ResNet-18 transfer learning framework across all three botanical modalities.

3. ABOUT THE DATASETS

This research utilizes three separate image datasets representing Indian medicinal flora across three distinct anatomical profiles. The dataset distribution and class profile of leaves, plants and flowers is illustrated in Table 1. The dataset of plant leaves [11], plants [11] and flowers [12] are available in Mendeley repository.

Table 1. Dataset Distribution and Class Profiles for Leaves, Whole Plants, and Flowers		
Modality	Number of Classes	Total Image Volume
Medicinal Leaves	80 species	6,900 images
Whole Plants	40 species	5,953 images
Medicinal Flowers	28 species	6,213 images

A. Indian Medicinal Leaves Dataset (80 Classes)

The leaf corpus comprises 6,900 images distributed across 80 distinct medicinal plant classes, exhibiting natural collection imbalances that range from a minimum of 8 samples per class (Lemongrass, Pepper) to a maximum of 177 samples (Tulsi), with a mean of 86.25 samples and a median of 79.5. This structural profile presents significant intra-family morphological overlap, challenging the deep neural network to isolate fine-grained micro-textural epidermal traits and subtle tertiary venation patterns across high-similarity taxa. The complete class-wise sample distribution of the Indian Medicinal Leaves dataset is detailed in Table 2. Some image samples of Indian medicinal leaves are illustrated in Figure 1.




Figure 1. Samples of Indian Medicinal Plant Leaves

Table 2. Dataset distributions of Indian Medicinal Leaves

#	Plant Leaf	Samples	#	Plant Leaf	Samples	#	Plant Leaf	Samples	#	Plant Leaf	Samples
1	Aloe vera	118	21	Coriender	115	41	Kohlrabi	73	61	Pomegranate	75
2	Amla	67	22	Curry	168	42	Lantana	76	62	Pumpkin	92
3	Amruthaballi	91	23	Doddpathre	142	43	Lemon	123	63	Raddish	40
4	Arali	89	24	Drumstick	56	44	Lemongrass	8	64	Rose	106
5	Astma_weed	82	25	Ekka	81	45	Malabar Nut	51	65	Sampige	61
6	Badipala	76	26	Eucalyptus	80	46	Malabar Spinach	79	66	Sapota	44
7	Balloon_Vine	61	27	Ganigale	75	47	Mango	103	67	Seethaashoka	47
8	Bamboo	118	28	Ganike	63	48	Marigold	93	68	Seethapala	114
9	Beans	97	29	Gasagase	79	49	Mint	135	69	Spinach1	67
10	Betel	114	30	Ginger	82	50	Neem	132	70	Tamarind	176
11	Bhrami	104	31	Globe Amarnath	81	51	Nelavembu	90	71	Taro	69
12	Bringaraja	73	32	Guava	128	52	Nerale	62	72	Tecoma	69

13	Caricature	76	33	Henna	80	53	Nooni	72	73	Thumbe	74
14	Castor	129	34	Hibiscus	118	54	Onion	92	74	Tomato	62
15	Catharanthus	134	35	Honge	113	55	Padri	73	75	Tulsi	177
16	Chakte	68	36	Insulin	89	56	Palak(Spinach)	149	76	Turmeric	39
17	Chilly	69	37	Jackfruit	110	57	Papaya	135	77	ashoka	81
18	Citron lime (herelikai)	99	38	Jasmine	49	58	Parijatha	66	78	camphor	66
19	Coffee	83	39	Kambajala	59	59	Pea	47	79	kamakasturi	67
20	Common rue(naagdalli)	67	40	Kasambruga	48	60	Pepper	8	80	kepala	76

B. Indian Medicinal Whole-Plant Dataset (40 Classes)

The plant dataset comprises 40 distinct species with a cumulative sum of 5,953 images. The distribution features a consistent baseline of 146 samples across 30 of the 40 classes, ranging from a minimum of 115 samples (*Ganike*) to a maximum of 187 samples (*Jasmine*). The images capture multi-stem growth habits, structural foliage arrangement, and complex backgrounds. The samples of Indian Medicinal whole plants are illustrated in Figure 2.




Figure 2. Samples of Indian Medicinal whole plants

Table 3. Dataset distributions of Indian Medicinal whole plants.

#	Medicinal Plant	Samples	#	Medicinal Plant	Samples
1	Aloevera	164	21	Hibiscus	165
2	Amla	146	22	Honge	146
3	Amruta_Balli	146	23	Insulin	146
4	Arali	146	24	Jasmine	187
5	Ashoka	146	25	Lemon	146
6	Ashwagandha	146	26	Lemon_grass	146
7	Avacado	146	27	Mango	146
8	Bamboo	146	28	Mint	153
9	Basale	146	29	Nagadali	152
10	Betel	151	30	Neem	146
11	Betel_Nut	146	31	Nithyapushpa	146
12	Brahmi	146	32	Nooni	146
13	Castor	160	33	Pappaya	146
14	Curry_Leaf	146	34	Pepper	146
15	Doddapatre	146	35	Pomegranate	146
16	Ekka	146	36	Raktachandini	146
17	Ganike	115	37	Rose	168
18	Gauva	146	38	Sapota	146
19	Geranium	146	39	Tulasi	146
20	Henna	150	40	Wood_sorel	146

C. Indian Medicinal Flower Dataset (28 Classes)

The flower corpus consists of 6,213 images spanning 28 botanical flower categories. Class sizes range from 117 images (*Spider Lily*) to 358 images (*Jatropha*), with a mean of approximately 221.89 images per class. The high visual divergence in corolla symmetry, stamen arrangement, and petal pigmentation provides strong discriminative markers for visual recognition. The samples of Indian Medicinal flower are illustrated in Figure 3.




Figure 3. Samples of Indian Medicinal flower

Table 4. Dataset distributions of Indian Medicinal Flowers

#	Medicinal Flower	Samples	#	Medicinal Flower	Samples
1	Adathoda	128	15	Marigold	337
2	Banana	163	16	Nagapoovu	141
3	Bush Clock Vine	138	17	Nityakalyani	228
4	Champaka	228	18	Pinwheel	232
5	Chitrak	315	19	Rose	275
6	Common Lanthana	227	20	Shankupushpam	241
7	Crown Flower	126	21	Spider Lily	117
8	Datura	231	22	Sunflower	331
9	Four O' Clock Flower	210	23	Thechi	300
10	Hibiscus	209	24	Thumba	179
11	Honey Suckle	242	25	Touch Me Not	201
12	Indian Mallow	304	26	Tridax procumbens	237
13	Jatropha	358	27	Wild potato vine	181
14	Malabar Melastome	243	28	Yellow Daisy	194

4. METHODOLOGY

The investigation adopts the Framework of the Proposed Botanical Classification Methodology, executed across six phases. The frame work is illustrated in Figure 4.

A. Phase 1: Ingestion and Dataset Invariance Profiling

Raw imagery is ingested from hierarchical local storage using an automated image datastore configured with recursive subfolder directory parsing:

$$\text{Labels} \leftarrow \text{DirectoryNames}(\text{Path}) \quad (2)$$

Class boundaries and label distributions are validated programmatically to configure the target output dimensionality K for each respective dataset ($K \in \{80,40,28\}$).

B. Phase 2: Dual-Stage Stratified Data Partitioning

To prevent sampling bias, each dataset is split using a dual-stage randomized stratified procedure:

1. An initial 70% partition is assigned to the training set ($\mathcal{D}_{\text{train}}$), leaving a 30% intermediate pool.
2. The intermediate pool is divided equally (0.50 factor), yielding 15% for validation monitoring ($\mathcal{D}_{\text{val}}$) and 15% for holdout test evaluation ($\mathcal{D}_{\text{test}}$).

This guarantees identical class-frequency ratios across all three partitions:

$$|\mathcal{D}_{\text{train}}| : |\mathcal{D}_{\text{val}}| : |\mathcal{D}_{\text{test}}| = 0.70 : 0.15 : 0.15 \quad (3)$$

C. Phase 3: Pre-processing and Stochastic Augmentation

To match the receptive field of the backbone architecture, all inputs are standardized to $224 \times 224 \times 3$. Any single-channel grayscale images are expanded to 3-channel RGB:

$$\mathbf{I}_{\text{RGB}} = [\mathbf{I}_{\text{gray}}, \mathbf{I}_{\text{gray}}, \mathbf{I}_{\text{gray}}] \quad (4)$$

To minimize background bias and teach the network rotation-invariant leaf and petal representations, an online stochastic data augmenter applies affine transformations to $\mathcal{D}_{\text{train}}$:

- Planar Rotations: $\theta \sim \mathcal{U}(-20^\circ, +20^\circ)$
- Horizontal Shift: $\Delta x \sim \mathcal{U}(-10 \text{ px}, +10 \text{ px})$
- Vertical Shift: $\Delta y \sim \mathcal{U}(-10 \text{ px}, +10 \text{ px})$
- Horizontal Mirroring: Random reflections across the vertical axis with probability $p = 0.5$

The validation ($\mathcal{D}_{\text{val}}$) and holdout test ($\mathcal{D}_{\text{test}}$) sets undergo only deterministic spatial resizing and channel normalization, preserving the integrity of the evaluation benchmarks.




Figure 4. Methodological framework of the proposed transfer learning pipeline for multi-class Indian medicinal plant classification using an adapted ResNet-18 architecture.

E. Phase 5: Supervised Fine-Tuning

Network parameters are optimized using the Adam optimizer, which computes individual adaptive learning rates for different parameters from first and second moments of the gradients:

$$\theta_{t+1} = \theta_t - \frac{\alpha}{\sqrt{\hat{v}_t + \epsilon}} \hat{m}_t \quad (5)$$

- Initial Learning Rate (α): 1×10^{-4}
- Mini-Batch Size: 32 images
- Epoch Budget: 8 complete passes
- Data Shuffling: Per-epoch mini-batch reshuffling
- Validation Frequency: Evaluated every 20 iterations against $\mathcal{D}_{\text{val}}$
- Hardware Acceleration: Automatic assignment to CUDA-enabled GPU hardware

F. Phase 6: Evaluation and Model Serialization

The converged network models are saved as .mat binaries. Final generalization is evaluated strictly on the unseen holdout test set ($\mathcal{D}_{\text{test}}$) to compute overall test classification accuracy:

$$\text{Accuracy} = \frac{\sum_{i=1}^{|\mathcal{D}_{\text{test}}|} \mathbb{I}(\hat{y}_i = y_i)}{|\mathcal{D}_{\text{test}}|} \times 100\% \quad (6)$$

A complete $K \times K$ confusion matrix is computed for each modality to evaluate class-by-class precision and identify structural misclassifications.

3. RESULTS AND DISCUSSION

A. Comparative Multi-Organ Performance Analysis

The proposed deep transfer learning framework converged successfully across all three botanical datasets.




Figure 5. Comparative classification performance of the fine-tuned ResNet-18 model across the three botanical modalities: Medicinal Leaves (94.79%), Whole Plants (98.10%), and Medicinal Flowers (99.79%) on the unseen holdout test partitions

The comparative classification accuracy of the proposed transfer learning architecture across medicinal leaves, whole plants, and flowers is illustrated in Fig. 5. In addition, detailed class-wise performance is characterized by the confusion matrices evaluated on the unseen test datasets, as shown in Fig. 6 (leaves), Fig. 7 (whole plants), and Fig. 8 (flowers).




Figure 6. Confusion matrix for leaves




Figure 7. Confusion matrix for Whole Plants




Figure 8. Confusion matrix for Flowers

B. Biological Morphometry and Performance Discrepancies

The variation in classification performance across modalities underscores how biological complexity directly shapes deep feature discriminability. The flower modality achieved near-perfect accuracy (99.79%), misclassifying only two samples in the holdout partition due to highly discriminative floral landmarks—such as vivid chromatic palettes, actinomorphic versus zygomorphic corolla symmetries, and distinct pistil/stamen configurations—which formed well-separated latent clusters. Whole-plant recognition similarly attained high performance (98.10%) despite cluttered field environments, as the residual backbone successfully learned multi-scale abstractions spanning canopy profiles, stem branching geometry, and macroscopic foliage arrangements. Conversely, foliar classification proved the most challenging (94.79%), a performance delta driven by a large 80-class decision space, widespread visual homogeneity across uniformly green palettes that forced reliance on subtle venation and serration cues, and severe sample imbalances in tail classes like Lemongrass and Pepper that induced boundary confusion with visually analogous species such as Bamboo.

C. Hyperparameter and Convergence Dynamics

Setting the learning rate factors to 10 for the ‘new_fc’ layer facilitated rapid loss convergence across the 8 epochs, as demonstrated by the sharp decline in validation loss during Epochs 1 and 2 while the new classification head adapted to upstream ImageNet filter representations. Concurrently, the conservative base learning rate 1×10^{-4} protected the foundational residual blocks from catastrophic forgetting, while the mini-batch size of 32 delivered stable gradient updates that balanced computational throughput with robust generalization performance.

6. CONCLUSION

This study developed and evaluated an end-to-end multi-organ deep transfer learning framework for identifying Indian medicinal flora using an adapted ResNet-18 architecture, demonstrating across leaves (80 classes), whole plants (40 classes), and flowers (28 classes) how anatomical differences directly affect deep feature representations. The framework achieved generalization test accuracies of 99.79% on flowers, 98.10% on whole plants, and 94.79% on plant leaves, confirming that while floral organs present the most discriminative visual markers, lightweight residual networks can also reliably classify complex

foliar contours and cluttered whole-plant imagery. Building on these findings, future research will explore multi-modal fusion architectures that concurrently process leaves, flowers, and whole plants via joint latent embeddings, incorporate explainable AI through Gradient-weighted Class Activation Mapping (Grad-CAM) to verify taxonomically grounded decision boundaries, and apply post-training quantization to enable real-time, low-power edge deployment on handheld devices for field botanists and Ayurvedic manufacturers.

AUTHOR CONTRIBUTIONS

Balaji Laxmi Pati: Conceptualization, Data Curation.

Prabira Kumar Sethy: Methodology, Analysis.

P Bharat Siva Varma: Validation, Formal Analysis.

Vallam Reddy Bhargavi Reddy: Visualization, Result Analysis

Santi Kumari Behera: Conceptualization, Draft Writing

Aziz Nanthaamornphong: Finding Analysis, Super Vision and Editing

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

ETHICS STATEMENT

Ethical approval was not required for this study.

DATA AVAILABILITY STATEMENT

The datasets are available in "Indian Medicinal Leaves Image Datasets", Mendeley Data, V3, doi: 10.17632/748f8jkphb.3 and "IMFI: Indian Medicinal Flower Image Dataset", Mendeley Data, V1, doi: 10.17632/kk8h65th9n.1.

REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778.
- [2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2009, pp. 248–255.
- [3] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Represent. (ICLR)*, San Diego, CA, USA, 2015, pp. 1–15.
- [4] S. H. Lee, C. S. Chan, P. Wilkin, and P. Remagnino, "Deep-plant: Plant identification with convolutional neural networks," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, 2015, pp. 452–456.
- [5] J. Wäldchen and P. Mäder, "Plant species identification using computer vision techniques: A systematic literature review," *Arch. Comput. Methods Eng.*, vol. 25, no. 2, pp. 507–543, 2018.
- [6] A. Camargo and J. S. Smith, "An image-processing based algorithm to automatically identify plant disease visual symptoms," *Biosyst. Eng.*, vol. 102, no. 1, pp. 9–21, 2009.
- [7] M. Dyrmann, H. Karstoft, and H. S. Midtby, "Plant species classification using deep convolutional neural networks," *Biosyst. Eng.*, vol. 151, pp. 72–80, 2016.
- [8] S. Sladojevic, M. Arsenovic, A. Anderla, D. Culibrk, and D. Stefanovic, "Deep neural networks based recognition of plant diseases by leaf image classification," *Comput. Intell. Neurosci.*, vol. 2016, Art. no. 3289801, 2016.
- [9] K. P. Ferentinos, "Deep learning models for plant disease detection and diagnosis," *Comput. Electron. Agric.*, vol. 145, pp. 311–318, 2018.

- [10] C. R. Sun, Y. Liu, and M. Wang, "Computer-aided identification of medicinal plants using botanical features and deep residual representations," *IEEE Access*, vol. 9, pp. 78210–78221, 2021.
- [11] B. R. Pushpa; Rani, Shobha (2023), "Indian Medicinal Leaves Image Datasets", Mendeley Data, V3, doi: 10.17632/748f8jkphb.3
- [12] B, Arjun ; Nair, Bipin; S, Abhishek ; NM, Abhinav; Madhavan, Varun (2024), "TMFI: Indian Medicinal Flower Image Dataset", Mendeley Data, V1, doi: 10.17632/kk8h65th9n.1