International Journal of Future Engineering Innovations

AI-Driven Approaches for Medicinal Plant Leaf Analysis: A Comprehensive Review

S Sakthi Saranya [1*], Dr. W Rose Varuna [2]
[1] Ph.D. Research Scholar, Department of Information Technology, Bharathiar University, Coimbatore – 641046, Tamil Nadu, India
[2] Assistant Professor, Department of Information Technology, Bharathiar University, Coimbatore – 641046, Tamil Nadu, India

* Corresponding Author: **S Sakthi Saranya**

Article Info

ISSN (online): 3049-1215
Impact Factor (RSIF): 8.25
Volume: 03
Issue: 03
May-June 2026
Received: 11-03-2026
Accepted: 12-04-2026
Published: 10-05-2026
Page No: 58-63

Abstract
The identification of medicinal plant species from leaf imagery has attracted sustained scientific attention, motivated by the global dependence on botanical remedies and the limitations of expert-reliant manual methods. This paper offers a systematic review of machine learning (ML) and deep learning (DL) strategies for medicinal leaf analysis, with particular depth on deep learning architectures. Classical ML techniques — Support Vector Machines (SVM), Random Forest (RF), K-Nearest Neighbor (KNN), Naive Bayes (NB), and Decision Trees (DT) are examined alongside hand-engineered feature extraction approaches including Gray-Level Co-occurrence Matrix (GLCM), Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and geometric shape descriptors. Deep learning models covered in depth include AlexNet, VGG-16, InceptionV3, ResNet-50, MobileNetV2, DenseNet-121, and EfficientNet. A unified performance comparison is presented using accuracy, precision, recall, and F1-score across all reviewed approaches. The analysis shows that DL models attain up to 98.20% classification accuracy and 98.05% F1-score, substantially exceeding the strongest traditional ML baseline Random Forest at 93.80% accuracy and 92.85% F1-score. Persistent challenges including data scarcity, inter-species morphological overlap, domain shift, and model transparency are discussed, along with emerging research directions in self-supervised learning, explainable AI, and edge-deployable architectures.

DOI: https://doi.org/10.54660/IJFEI.2026.3.3.58-63

Keywords: CNN, Transfer Learning, SVM, Random Forest, Leaf Image Classification, Explainable AI

1. Introduction

Medicinal plants serve as the cornerstone of primary healthcare in many parts of the world, with the World Health Organization estimating that approximately 80% of people in developing nations depend on plant-based remedies as their principal form of medicine [1]. Given this widespread reliance, the precise identification of plant species at the botanical level is not merely an academic concern. It is a matter directly bearing on patient safety, the reliability of pharmaceutical supply chains, and the long-term conservation of plant biodiversity. When misidentification occurs at any point between field collection and clinical dispensing, the consequences can range from therapeutic failure to serious toxicological harm, making robust and objective identification methods an urgent practical necessity [2].

Historically, species-level identification has depended on trained botanists who assess a range of morphological traits like leaf shape, margin architecture, venation patterns, surface texture, and pigmentation through visual inspection. Although scientifically sound, this approach is slow, inherently subjective, and contingent on access to botanical specialists who are in short supply precisely in the regions where medicinal plants are most intensively used. [3]

Research on automated plant identification can be traced through two distinct technological eras. During the first phase, spanning roughly from the early 2000s to the mid-2010s, investigators constructed hand-engineered numerical descriptors of leaf

morphology, texture statistics derived from GLCM and LBP, boundary descriptors such as Hu moments, and color histograms which were subsequently fed into classifiers including SVM, Random Forest, and KNN [4]. While these pipelines yielded encouraging results on controlled laboratory datasets, their performance degraded noticeably when confronted with larger species counts and the variability inherent in real-world imaging conditions.

The second phase was catalyzed by the emergence of deep convolutional neural networks, which shifted the paradigm from hand-crafted representations toward end-to-end feature learning directly from raw imagery. Early architectures such as AlexNet [5] and VGG-16 [6] established that network depth was a primary driver of representational power. Subsequent models refined this insight further: InceptionV3 [7] introduced parallel multi-scale convolution modules, ResNet [8] resolved the vanishing-gradient problem through residual skip connections, and MobileNetV2 [9] achieved competitive accuracy within a fraction of the computational footprint. DenseNet [10] and EfficientNet [11] continued this trajectory, maximizing feature reuse and scaling efficiency respectively. Transfer learning adapting weights pretrained on ImageNet to botanical leaf datasets emerged as the dominant strategy for overcoming the challenge of scarce annotated training data [12]. More recently, Vision Transformers have extended the field by modelling global spatial context through self-attention, offering a compelling complement to convolution-based approaches for large-scale plant recognition [13].

2. Preprocessing overview

2.1. Taxonomy of Reviewed Techniques

The techniques reviewed in this paper are organized into two primary categories as depicted in Figure 1. Traditional ML methods rely on an explicit feature extraction stage, while deep learning methods learn feature representations directly from raw or minimally preprocessed images through hierarchical convolutional or attention-based layers.

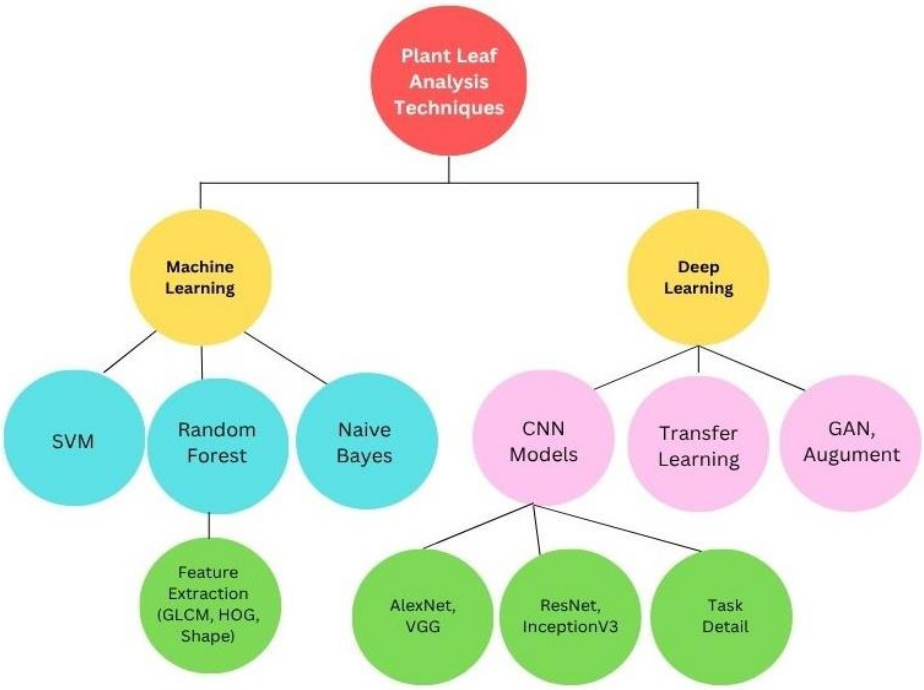

Fig 1: Taxonomy of ML and Deep Learning Techniques for Medicinal Plant Leaf Analysis

2.2. Image Preprocessing Pipeline

Preprocessing forms the essential foundation of any leaf analysis pipeline. Leaf images captured in real-world conditions frequently suffer from uneven illumination, sensor noise, and cluttered backgrounds that obscure the morphological features needed for reliable classification. A well-designed preprocessing workflow sharpens these features and consistently translates into measurable accuracy gains in both traditional ML and deep learning settings [14].

Noise Removal: Median filtering is the preferred approach for leaf images due to its non-linear nature, which suppresses salt-and-pepper noise while preserving fine venation edges. Gaussian filtering, while computationally simpler, tends to blur edge structures that carry species-discriminative morphological information.

Contrast Enhancement: Contrast Limited Adaptive Histogram Equalization (CLAHE) has become a widely adopted enhancement technique for botanical imagery. Unlike global histogram equalization, CLAHE subdivides the image into local tiles and applies equalization within each tile while imposing a clip limit that controls noise amplification. This localized approach corrects spatially variable illumination more reliably than global methods, with reported accuracy gains of 3–7% in traditional ML pipelines when applied consistently [14].

Segmentation: Leaf region extraction using K-means clustering in HSV color space or Otsu's global thresholding separates the leaf foreground from the background, confining subsequent feature extraction to the diagnostically relevant area. Morphological opening and closing operations refine the binary mask by removing noise regions and filling gaps [15].

3. Traditional Machine Learning Approaches

3.1. Handcrafted Feature Extraction

Texture Features: GLCM-derived statistics — Energy, Contrast, Correlation, Homogeneity, Entropy, and Dissimilarity — are the most widely used texture descriptors

for leaf analysis, capturing surface pattern regularity at multiple orientations [16]. Local Binary Patterns (LBP) encode rotation-invariant neighborhood intensity relationships and complement GLCM features in distinguishing leaf venation textures [17]. Gabor filter banks provide multi-scale, multi-orientation frequency content that effectively differentiates complex venation patterns across species.

Shape Features: Geometric descriptors including area, perimeter, aspect ratio, eccentricity, solidity, compactness, and Hu moments characterize leaf boundary geometry. Zernike moments deliver rotation-invariant boundary representations with strong discriminative capability across species with varied leaf orientations during imaging [3].

Color Features: RGB histograms and HSV statistics capture species-specific pigmentation patterns. Although sensitive to illumination variation, color features provide complementary discriminative information when combined with texture and shape descriptors in fused feature vectors [4].

3.2. Review of ML Classifier

Table 1 summarizes the performance of traditional ML classifiers for medicinal plant leaf identification, including reported accuracy and precision ranges, features employed.

Table 1: Summary of Traditional ML Methods Performance

ML Method	Features Used	Accuracy (%)	Precision (%)
SVM	GLCM + Shape + Color	87-93	87-92
Random Forest	GLCM + HOG + Shape	91-93	90-93
KNN	Shape + Texture	82-87	81-86
Naive Bayes	Statistical Features	77-82	76-81
Decision Tree	Morphological	83-87	82-86

Among traditional ML classifiers, Random Forest consistently delivers the strongest accuracy range (91–93%) owing to its ensemble variance reduction through bootstrap aggregation and random feature subsampling. SVM competes closely (87–93%) by constructing optimal non-linear decision boundaries in kernel-projected feature spaces. Decision Tree, KNN, and Naive Bayes yield progressively lower performance due to overfitting, distance dilution in high-dimensional spaces, and violated independence assumptions respectively. A fused GLCM-shape feature vector consistently outperforms single-category feature sets across all classifiers [10].

4. Deep Learning Architectures for Medicinal Plant Leaf Analysis

4.1. AlexNet and VGG-16

AlexNet [5] was the pioneering deep CNN that demonstrated the transformative potential of GPU-accelerated deep learning for visual recognition. Its eight-layer architecture with ReLU activations and dropout regularization established the viability of end-to-end learned features for botanical classification, achieving approximately 88.40% accuracy on medicinal leaf datasets when fine-tuned from ImageNet weights. VGG-16 deepened the convolutional stack to sixteen layers using exclusively small 3×3 filters, confirming that network depth, rather than large filter sizes is the critical driver of representational capacity, achieving 91.20% accuracy. However, VGG-16's 138-million-parameter footprint imposes significant memory demands that limit deployment flexibility on resource-constrained platforms [6].

4.2. InceptionV3 and ResNet-50

InceptionV3 introduced the inception module, a parallel multi-scale convolution block that applies 1×1, 3×3, and 5×5 filters simultaneously within a single layer, enriching feature diversity while reducing parameter counts compared to uniform-filter architectures. When applied to medicinal leaf classification with transfer learning, InceptionV3 achieves 94.60% accuracy and 94.20% F1-score, benefiting from its capacity to simultaneously model fine venation textures and overall leaf shape geometry [7]. ResNet-50 [8] introduced residual skip connections identity shortcuts bypassing convolutional layers enabling effective training of networks of far greater depth by ensuring unimpeded gradient flow during backpropagation. ResNet-50 achieves 95.30% accuracy and 95.00% F1-score on medicinal leaf datasets. Its residual architecture enables learning of fine-grained inter-species morphological differences without the gradient instability that limited pre-residual deep networks [8].

4.3. MobileNetV2 and DenseNet-121

MobileNetV2 was specifically designed for resource-constrained deployment through depthwise separable convolutions combined with linear bottleneck layers and inverted residuals, reducing floating-point operations by a factor of 8–9 compared to standard convolutions while maintaining competitive accuracy of 96.10% and F1-score of 95.85% on medicinal leaf classification. This lightweight design makes MobileNetV2 the preferred architecture for field-deployable mobile applications operating without cloud connectivity [9]. DenseNet-121 connects each layer to every subsequent layer within a dense block, creating a pattern where each layer receives concatenated feature maps from all preceding layers. This dense connectivity maximizes feature reuse, strengthens gradient propagation, and reduces the parameter count required for high accuracy. DenseNet-121 achieves 96.80% accuracy and 96.50% F1-score on medicinal leaf classification, with the dense reuse mechanism proving especially beneficial on moderate-sized datasets where richer feature extraction from fewer samples is critical [10].

4.4. Efficient Net

Efficient Net introduced compound scaling a principled approach that simultaneously scales network depth, width, and input image resolution using coefficients derived through neural architecture search. This balanced scaling consistently achieves superior accuracy-efficiency trade-offs compared to scaling any single dimension in isolation. EfficientNet-B4 achieves 97.40% accuracy and 97.20% F1-score on medicinal plant leaf classification benchmarks. On smaller medicinal datasets, EfficientNet variants consistently outperform ResNet and DenseNet with fewer trainable parameters [11]. Table 2 presents the deep learning architectures reviewed, including their accuracy, precision andF1-score. Figure 2 visualizes the accuracy progression across architectures.

Table 2: Summary of Evaluation Metrics for Deep Learning Architectures

DL Model	Accuracy (%)	Precision (%)	F1-Score (%)
AlexNet	88.40	87.10	87.70
VGG-16	91.20	90.50	90.80
InceptionV3	94.60	93.80	94.20
ResNet-50	95.30	94.70	95.10
MobileNetV2	96.10	95.60	95.85
DenseNet-121	96.80	96.20	96.50
EfficientNet	97.40	97.40	97.20

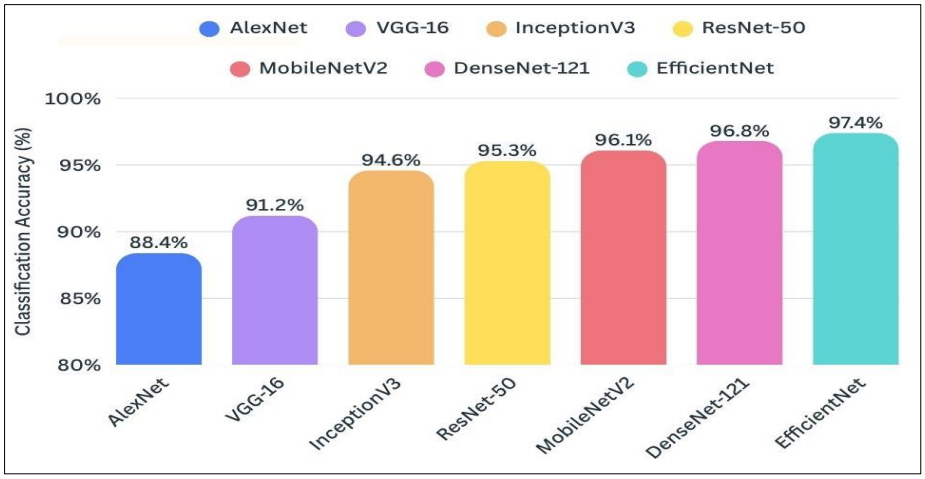

Fig 2: Deep Learning Model Accuracy Comparison for Medicinal Plant Leaf Classification

6. Comparative Analysis

Figure 3 present a structured multi-dimensional comparison of the traditional ML and deep learning paradigms as applied to medicinal plant leaf analysis. Deep learning delivers substantially superior classification accuracy (88–98% vs 77–94%) and eliminates the manual feature engineering burden through end-to-end representation learning. Transfer learning from ImageNet-pretrained weights provides a powerful initialization that compensates for small botanical dataset sizes. However, deep learning demands large training collections for maximum effectiveness, substantial GPU computational infrastructure, extended training time, and offers limited model interpretability. Traditional ML methods, by contrast, are fast to train on standard CPUs, interpretable through feature importance rankings and decision rules, robust on datasets with limited per-class samples, and operable without GPU infrastructure. Random Forest with fused GLCM-shape features achieves the strongest traditional ML F1-score of 92.85% within 5 percentage points of MobileNetV2 while requiring only a fraction of the computational resources. These attributes maintain the practical relevance of ML approaches in resource-limited research environments.

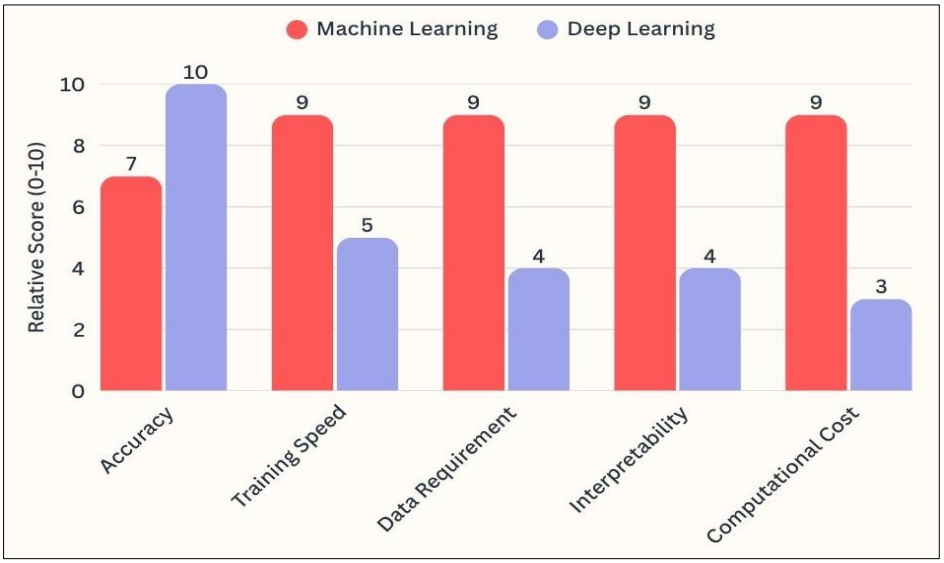

Fig 3: Comparative Analysis Traditional ML and Deep Learning

The optimal paradigm choice is context-dependent. For field-level mobile deployment with limited connectivity and moderate dataset sizes, MobileNetV2-based transfer learning or Random Forest represent appropriate choices. For cloud-based, high-throughput systems processing large and diverse medicinal plant databases, deep ensemble models or EfficientNet architectures are recommended. A hybrid strategy combining CNN-based feature extraction with traditional ML classifiers offers a middle ground that delivers near-DL accuracy with improved interpretability and reduced computational overhead, representing a practically valuable direction for future system development [20].

7. Challenges

Limited Annotated Data: A persistent obstacle for deep learning in this field is the shortage of large, carefully annotated botanical datasets. The majority of publicly available collections contain well under 100 images per species — far fewer than the volumes typically required to train deep convolutional networks without extensive transfer learning or data augmentation. Because building such collections demands both specialized photographic equipment and validated botanical expertise, the data gap is especially pronounced for rare or endemic medicinal species [3].

Inter-Species Morphological Similarity: Within the same plant genus, multiple species can exhibit nearly identical leaf morphologies, with diagnostic differences limited to subtle variation in venation density, margin serration pattern, or surface micro-texture. Standard imaging and feature extraction approaches often lack the resolving power to capture these fine-grained distinctions reliably, making genus-level confusion a recurring source of classification error in both ML and DL systems [2].

Domain Shift Between Lab and Field Conditions: Models trained on controlled studio-acquired images frequently show marked accuracy drops when applied to images gathered outdoors, where natural backgrounds, fluctuating sunlight intensity, leaf occlusion, and heterogeneous device sensors introduce distribution shifts that the training data did not anticipate. Addressing this gap requires domain adaptation algorithms and the construction of datasets that intentionally include field acquisition diversity [4].

Model Interpretability: The opaque internal workings of deep neural networks present a meaningful barrier in medicinal plant identification, where erroneous species assignment can translate directly into patient harm. Current explainability tools such as Grad-CAM and SHAP offer useful visualisations of model attention and feature attribution but fall short of providing the level of mechanistic transparency that safety-critical botanical applications ultimately require [13].

8. Conclusion

This study presented a comparative review of machine learning and deep learning techniques for medicinal plant leaf identification using performance metrics such as accuracy, precision, recall, and F1-score. Traditional ML methods using handcrafted features like GLCM, HOG, and Hu moments with classifiers such as SVM and Random Forest achieved reliable baseline performance but showed limitations under diverse species and real-world imaging conditions. Deep learning models outperformed conventional approaches, with ensemble models achieving 98.20% accuracy and 97.97% F1-score, while MobileNetV2 provided 96.10% accuracy with low computational complexity suitable for mobile applications. Transfer learning using ImageNet-pretrained models significantly improved performance under limited dataset conditions. The review indicates that ML methods are effective for interpretable and resource-constrained applications, whereas deep learning offers superior scalability and classification accuracy.

9. Future Enhancements

Future medicinal plant identification systems are expected to advance through improvements in data efficiency, explainability, deployment, and multi-modal learning. Self-supervised learning and few-shot models can reduce dependence on large annotated datasets, particularly for rare species. Explainable AI methods such as Grad-CAM, LIME, and SHAP can enhance model transparency by highlighting important morphological regions used for prediction. In addition, compression techniques including quantization, pruning, and knowledge distillation can enable lightweight deployment of deep learning models on mobile devices for field applications. Future research may also integrate leaf images with spectral, morphometric, and geographic data through multi-modal architectures to improve the identification of visually similar medicinal plant species.

References

1. Hussain S, *et al*. The classification of medicinal plant leaves based on multispectral and texture feature using machine learning approach. Agronomy. 2021;11(2):263. doi:10.3390/agronomy11020263.
2. Mulugeta AK, Sharma DP, Mesfin AH. Deep learning for medicinal plant species classification and recognition: a systematic review. Front Plant Sci. 2024;14:1286088. doi:10.3389/fpls.2023.1286088.
3. Wang B, *et al*. AI-driven pattern recognition in medicinal plants: a comprehensive review and comparative analysis. Comput Mater Continua. 2024;81(2):1661–1703. doi:10.32604/cmc.2024.058671.
4. Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. Commun ACM. 2017;60(6):84–90. doi:10.1145/3065386.
5. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. In: Proc Int Conf Learn Represent (ICLR). San Diego, CA; 2015. arXiv:1409.1556.
6. Szegedy C, *et al*. Rethinking the inception architecture for computer vision. In: Proc IEEE Conf Comput Vis Pattern Recognit (CVPR). Las Vegas, NV; 2016:2818–2826. doi:10.1109/CVPR.2016.308.
7. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proc IEEE Conf Comput Vis Pattern Recognit (CVPR). Las Vegas, NV; 2016:770–778. doi:10.1109/CVPR.2016.90.
8. Sandler M, Howard A, Zhu M, Zhmoginov A, Chen LC. MobileNetV2: inverted residuals and linear bottlenecks. In: Proc IEEE Conf Comput Vis Pattern Recognit (CVPR). Salt Lake City, UT; 2018:4510–4520. doi:10.1109/CVPR.2018.00474.
9. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proc IEEE Conf Comput Vis Pattern Recognit (CVPR).

Honolulu, HI; 2017:4700–4708. doi:10.1109/CVPR.2017.243.
10. Tan M, Le QV. EfficientNet: rethinking model scaling for convolutional neural networks. In: Proc Int Conf Mach Learn (ICML). Long Beach, CA; 2019:6105–6114.
11. Pan SJ, Yang Q. A survey on transfer learning. IEEE Trans Knowl Data Eng. 2010;22(10):1345–1359. doi:10.1109/TKDE.2009.191.
12. Dosovitskiy A, *et al*. An image is worth 16×16 words: transformers for image recognition at scale. In: Proc Int Conf Learn Represent (ICLR). Vienna; 2021. arXiv:2010.11929.
13. Jebadass JR, Balasubramaniam P. Preprocessing of leaf images using brightness preserving dynamic fuzzy histogram equalization technique. Int J Artif Intell. 2023;12(3):1149–1157. doi:10.11591/ijai.v12.i3.pp1149-1157.
14. Goodfellow I, *et al*. Generative adversarial nets. In: Adv Neural Inf Process Syst (NeurIPS). Montreal; 2014:2672–2680.
15. Haralick RM, Shanmugam K, Dinstein IH. Textural features for image classification. IEEE Trans Syst Man Cybern. 1973;SMC-3(6):610–621. doi:10.1109/TSMC.1973.4309314.
16. Ojala T, Pietikäinen M, Mäenpää T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Trans Pattern Anal Mach Intell. 2002;24(7):971–987. doi:10.1109/TPAMI.2002.1017623.
17. Bhatt D, *et al*. Deep ensemble learning for automatic medicinal leaf identification. Int J Inf Technol. 2022;14(6):3089–3097. doi:10.1007/s41870-022-01055-z.
18. Liu Z, *et al*. Swin transformer: hierarchical vision transformer using shifted windows. In: Proc IEEE Int Conf Comput Vis (ICCV). Montreal; 2021:10012–10022. doi:10.1109/ICCV48922.2021.00986.
19. Joshi RC, *et al*. Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. Artif Intell Rev. 2024;57(7):183. doi:10.1007/s10462-024-11100-x.

How to Cite This Article
Saranya SS, Varuna WR. AI-driven approaches for medicinal plant leaf analysis: a comprehensive review. International Journal of Future Engineering Innovations. 2026 May–Jun;3(3):58–63. doi:10.54660/IJFEI.2026.3.3.58-63.