Modernizing Medicinal Plant Recognition: A Deep Learning Perspective with Data Augmentation and Hybrid Learning

Preeti Nagrath[1,*], M. Batumalay[2], Dhruv Saluja[3], Harsh Kukreja[4], Devanshi Tegwal, Akash Saini[6]

[1,3,4,5,6]*Bharati Vidyapeeth's College of Engineering, A - 4, Paschim Vihar, New Delhi, 110063, Delhi, India*

[2]*Faculty of Information Technology, INTI International University, Malaysia*

[2]*Centre for Data Science and Sustainable Technologies, INTI International University, Nilai, N. Sembilan, Malaysia*

(Received: March 01, 2025; Revised: May 25, 2025; Accepted: September 01, 2025; Available online: October 04, 2025)

Abstract

This study proposes a deep learning-based solution to address the longstanding challenge of accurately identifying Indian medicinal plants, which are vital to Ayurvedic pharmaceutics but often misidentified due to their morphological similarities. The objective is to develop a reliable, automated classification system using image processing and advanced neural network architectures. A dataset of 5,945 images representing 40 distinct medicinal plant species was sourced from Kaggle and augmented to 11,890 images using techniques such as flipping, rotation, and scaling to enhance diversity. The models tested include a baseline Convolutional Neural Network (CNN), transfer learning with DenseNet121, DenseNet169, and DenseNet201, a voting ensemble of these DenseNet variants, and a hybrid DenseNet201-LSTM architecture. Experimental results show that the CNN model achieved the lowest accuracy at 69.58%, while the hybrid DenseNet201-LSTM model reached the highest validation accuracy of 93.38%, with a precision of 94.74%, recall of 93.38%, and F1-score of 93.42%. These findings confirm the hybrid model's superior ability to capture spatial and sequential dependencies in leaf features. The novelty of this work lies in the integration of DenseNet201 with LSTM for medicinal plant classification, which has not been widely explored in this domain. The study also acknowledges dataset scalability as a limitation and proposes future work involving dataset expansion through botanical collaborations, integration of environmental metadata, and deployment of a mobile application using TensorFlow Lite for real-time, low-resource implementation. Overall, the research contributes a robust and scalable framework for medicinal plant identification, promoting trust in traditional medicine, supporting conservation efforts, and enabling practical field-level applications in both rural and clinical settings.

Keywords: Medicinal Plant Identification,Transfer Learning, Deep Learning, Hybrid Learning, Education Quality

1. Introduction

India possesses a rich and diverse botanical heritage, with an extensive array of medicinal plants that form the backbone of Ayurvedic medicine. These plants have been used for centuries in traditional healing practices and remain vital to the country's cultural and pharmaceutical landscape. However, the accurate identification of medicinal plants continues to pose a significant challenge. Factors such as seasonal variations, geographical diversity, morphological similarities among species, and the use of identical vernacular names for different plants contribute to widespread confusion and misidentification [1], [2]. This issue is further exacerbated by limited awareness among collectors, traders, and local practitioners, often leading to the misallocation or substitution of plant materials in the supply chain [3]. As a result, the quality and efficacy of herbal medicines are compromised, undermining trust in traditional healing systems [4].

In response to these challenges, this research proposes a deep learning-based approach to automate and enhance the identification of medicinal plants using image data. Recent studies have demonstrated the effectiveness of machine learning and computer vision techniques in plant recognition, achieving high classification accuracies on controlled datasets [5], [6], [7]. However, many prior efforts have been limited by small datasets, narrow species scope, and low generalizability in real-world scenarios [8], [9]. To overcome these limitations, this study leverages a comprehensive dataset comprising 40 medicinal plant species sourced from Kaggle, enhanced through image augmentation techniques

***Corresponding author: Preeti Nagrath (preeti.nagrath@bharatividyapeeth.edu)**
DOI: https://doi.org/10.47738/jads.v6i4.795

such as rotation, flipping, and scaling to improve model generalization [10]. Multiple deep learning models are explored, including a baseline CNN, transfer learning approaches using DenseNet121, DenseNet169, and DenseNet201, and an ensemble voting method that combines the predictions of these three DenseNet architectures [11]. To further improve performance, a hybrid model integrating DenseNet201 with a Long Short-Term Memory (LSTM) network is proposed, enabling the capture of sequential spatial dependencies in image data [12]. The CNN model achieved an accuracy of 69.58%, while the hybrid DenseNet201-LSTM model reached the highest accuracy of 93.38%, with consistently high precision, recall, and F1 scores across all architectures [13].

This research aims not only to provide a scalable and efficient solution to the longstanding problem of plant misidentification but also to contribute to the conservation and responsible use of India's medicinal flora. By incorporating advanced deep learning techniques, the proposed system enhances the reliability and efficiency of plant identification, with practical implications for quality assurance in Ayurvedic medicine, conservation biology, and stakeholder education [14]. Ultimately, this work highlights the transformative potential of artificial intelligence in strengthening traditional healthcare systems and promoting sustainable herbal medicine practices.

2. The Related Works

Recent advancements in medicinal plant identification have explored a wide array of methodologies, prominently featuring machine learning and deep learning techniques. In this section, we compare, contrast, criticize, synthesize, and summarize key contributions from the literature.

Several studies [1], [10], [13], [15], [16] demonstrate strong performance in plant classification tasks using different image-based machine learning approaches. For instance, [1] offers an efficient and cost-effective method, while [16] and [13] introduce mobile applications for real-time identification, reaching high accuracy levels above 97%. Studies such as [5], [6], and [17] show the efficacy of deep convolutional architectures like DenseNet, ShuffleNet, and OTAMNet, often achieving accuracy over 98%. Similarly, [18] and [19] propose ensemble and optimized CNN frameworks, further pushing the limits of classification performance.

While methods in [10], [15], and [20] rely on manual feature extraction and traditional classifiers (NBC, KNN, ANN), others like [6], [13], [14], [17], and [16] utilize deep learning models that learn features automatically. Some works focus on small, region-specific datasets [3], [13], [21], whereas others tackle large-scale classification problems [8], [14], [22] with tens of thousands of plant species or millions of images. Additionally, studies differ in their deployment contexts: some are purely theoretical or lab-based [4], [23], while others emphasize real-world implementation in mobile or cloud-based systems [13], [16].

Despite promising accuracies, numerous limitations persist. Several studies [3], [10], [13], [21], [18] suffer from dataset limitations—either in size, diversity, or representativeness. Works like [24], [25], and [26] lack empirical benchmarking or comparative analysis of models. Privacy concerns and lack of data sharing are flagged in [15], and ethical considerations are overlooked in. Even high-performing models such as [17], and [26] may exhibit dataset-specific performance, limiting generalizability. Furthermore, the absence of occlusion handling and insufficient validation under real-world conditions is common in [2], [27].

From the reviewed literature, a synthesis reveals strong consensus on the potential of deep learning, particularly CNNs and transfer learning, to improve accuracy in medicinal plant identification. Studies such as [6], [7], [14], and [18] collectively demonstrate the value of hybrid and ensemble models. There is an emerging trend of integrating advanced architectures like InceptionResNetV2 [22], DenseNet [4], and EfficientNet [27] with optimization strategies. The literature also underscores the importance of accessible public datasets and standardized benchmarks to enable reproducibility and scalability. Moreover, mobile implementation [13], [16] suggests a shift toward practical, user-friendly solutions.

While prior works contribute valuable insights into medicinal plant identification using AI techniques, they are often limited by dataset scope, environmental variability, or lack of robust validation. Our study builds upon these foundations by employing a diverse 40-class dataset, exploring multiple deep learning architectures including DenseNet121, DenseNet169, DenseNet201, a voting ensemble, and a hybrid DenseNet201-LSTM model. By addressing the gaps identified in feature extraction, model robustness, and deployment feasibility, our approach offers a more comprehensive and scalable solution for real-world applications in herbal medicine classification.

The literature reveals substantial progress in medicinal plant identification using deep learning but also highlights persistent gaps. Many studies rely on limited datasets and narrowly defined feature sets, constraining their applicability in diverse environmental contexts. Our study addresses these issues by using a dataset of 40 distinct classes and applying CNN, DenseNet121, DenseNet169, and DenseNet201, followed by ensemble learning and a DenseNet201-LSTM hybrid model. This progression ensures improved accuracy, better generalization, and enhanced performance in varied scenarios. Moreover, we emphasize ethical considerations and comprehensive model validation. The result is a more scalable, accurate, and practical approach to medicinal plant identification with applications in pharmaceuticals, agriculture, and conservation.

3. Methodology

Our study employs a well-structured and systematic deep learning pipeline to classify Indian medicinal plant species based on leaf images. This pipeline (illustrated in figure 1) is meticulously designed to address key challenges such as intra-class similarity, environmental variability, and limited data per class. The methodology consists of four main stages: Data Acquisition, Data Preprocessing, Train-Test Split, and Class Prediction. In the data acquisition stage, a diverse and high-resolution image dataset of medicinal plant leaves was sourced from Kaggle, capturing a wide range of plant species, lighting conditions, and geographic contexts. This diversity is crucial to ensure the model's ability to generalize across real-world scenarios. Duplicate entries were removed, and class imbalance was addressed to create a reliable and representative dataset.

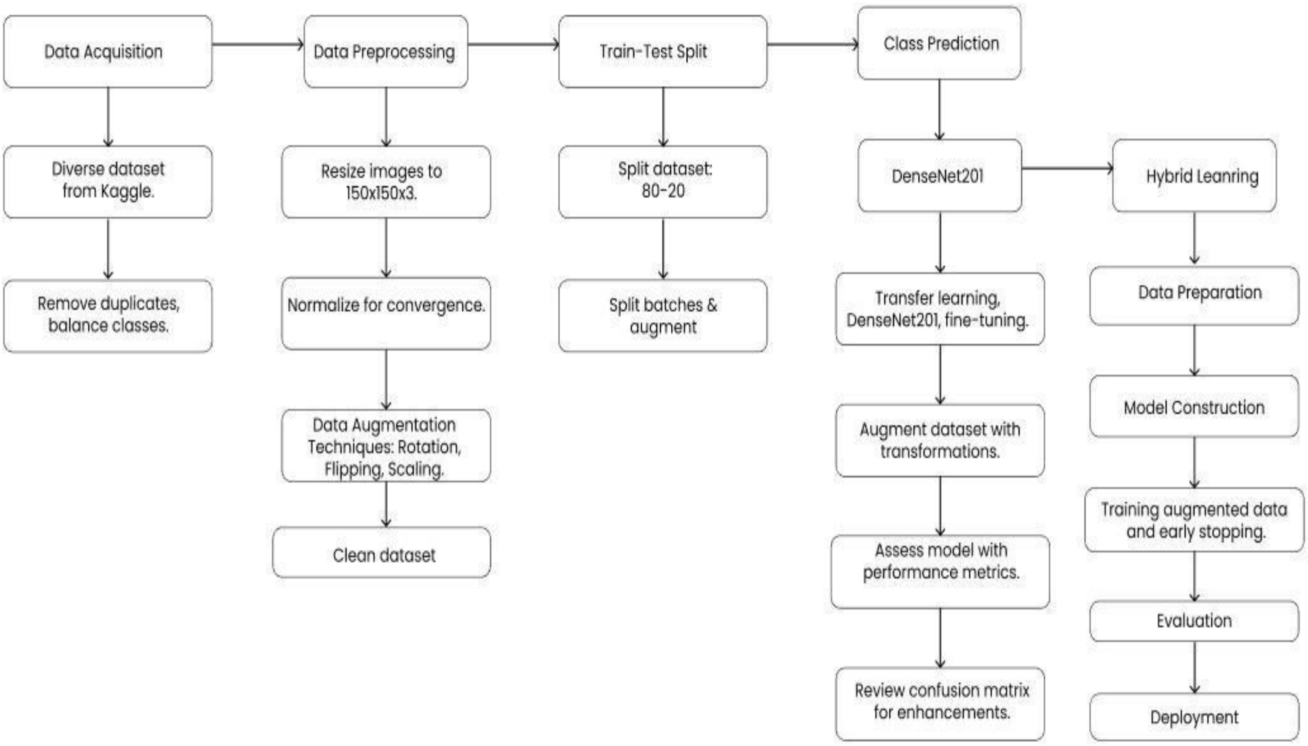

Figure 1. Research Pipeline

In the preprocessing stage, all images were resized to a fixed resolution of 150×150×3 pixels and normalized to aid convergence during training. Extensive data augmentation techniques—such as rotation, flipping, scaling, and brightness adjustment—were applied to simulate natural variations and enhance model robustness. After cleaning, the dataset was split into 80% training and 20% testing subsets, with augmentation applied uniformly to maintain class balance. For classification, we explored both baseline CNN architectures and advanced transfer learning models including DenseNet121, DenseNet169, and DenseNet201. An ensemble voting strategy was implemented to combine the strengths of individual models. Furthermore, we proposed a hybrid model combining DenseNet201 and LSTM, which integrates spatial and sequential features to achieve superior classification performance. This multi-stage methodology, illustrated in figure 1, forms the foundation for accurate and efficient medicinal plant identification using deep learning.

3.1. Data Acquisition

A diverse and high-quality image dataset of Indian medicinal plant leaves was obtained from Kaggle [28], originally compiled by Arya Shah. This dataset comprises 5,945 labeled images representing 40 distinct medicinal plant species, with considerable variation in geographical origin, seasonal context, and environmental backgrounds. Such diversity ensures a robust foundation for training models capable of generalizing well to real-world plant identification scenarios.

To maintain data integrity and ensure consistent learning performance, the dataset underwent an initial cleaning phase, where duplicate entries were removed and class distributions were balanced to prevent bias during training. The dataset emphasizes leaf morphology, which is a critical feature in plant taxonomy and provides rich visual cues for classification. To further increase the variability and enhance the model's ability to generalize, we applied basic data augmentation techniques, including image rotation, flipping, and brightness adjustments, simulating natural variations in leaf orientation and lighting conditions.

3.2. Data Preprocessing

The preprocessing stage was essential to prepare the raw image data for effective deep learning model training. All images were uniformly resized to 150×150 pixels with 3 color channels (RGB) to ensure compatibility with standard convolutional neural network architectures and to maintain consistency in input dimensions. Following resizing, normalization was applied to scale pixel values to a [0,1] range, which facilitates faster and more stable convergence during training by reducing internal covariate shift.

To enhance the model's ability to generalize and handle real-world variations, the dataset was further enriched through data augmentation techniques, including random rotation, horizontal flipping, and scaling. These transformations introduce artificial diversity without altering the underlying class semantics, thereby reducing the risk of overfitting. Additionally, a comprehensive data cleaning process was conducted to eliminate ambiguous, mislabeled, or corrupted samples, ensuring that only high-quality and relevant images were retained for training. This preprocessing pipeline played a critical role in improving the robustness, accuracy, and generalizability of the classification models.

3.3. Train-Test Split

To enable reliable evaluation of model performance, the dataset—expanded to approximately 11,890 images after augmentation—was divided into training and testing subsets using an 80:20 split ratio. Each original image was paired with one augmented version, effectively doubling the dataset size and ensuring that class distributions remained balanced and representative across both subsets.

The training set (80%) was exclusively used for model learning, allowing the deep learning architectures to extract and generalize relevant features from the data. In contrast, the testing set (20%) was kept entirely separate during training to ensure an unbiased assessment of the model's ability to classify unseen data. Model performance was evaluated using key metrics including accuracy, precision, recall, and F1-score, providing a comprehensive view of classification effectiveness across all 40 medicinal plant species. This splitting strategy ensures both the robustness and validity of the experimental results.

3.4. Class Prediction

To address the multi-class classification of medicinal plant species, this study implemented and compared several deep learning approaches, including a baseline CNN, transfer learning models based on DenseNet variants (DenseNet121, DenseNet169, and DenseNet201), an ensemble voting mechanism, and a hybrid architecture that integrates DenseNet201 with a LSTM network. These models were trained and evaluated using the augmented dataset to determine the most accurate and robust architecture for plant image classification.

The baseline CNN model was developed using a series of convolutional layers with 32 filters, a kernel size of 3×3, and ReLU activation. Max-pooling layers with a 2×2 window were included to reduce the spatial dimensions of feature maps. The extracted features were then flattened and passed into a dense layer containing 128 neurons with ReLU activation, followed by a final softmax output layer for multi-class prediction. The softmax function computes the probability distribution over the 40 plant classes using the following equation:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_{j=1}^{C} e^{z_j}}, \quad i = 1,2, \dots, C \tag{1}$$

The variable C represents the number of classes, and zi is the raw output (logit) for class i. This baseline architecture served as a reference to measure the improvements brought by more advanced models.

To enhance performance and feature learning, transfer learning was utilized by incorporating pre-trained DenseNet121, DenseNet169, and DenseNet201 models. These models were initially trained on ImageNet and provided strong hierarchical representations of image features. The original top layers of each DenseNet were removed, and their convolutional bases were frozen to preserve pre-trained weights. A GlobalAveragePooling2D layer was then added to reduce the spatial dimensions of the feature maps. This operation calculates the mean of each feature map using the formula:

$$GAP(x) = \frac{1}{H \times W} \sum_{i=1}^{H} \sum_{j=1}^{W} x_{ij} \tag{2}$$

Following this, additional dense layers and a softmax layer were added for classification. DenseNet121 was selected for its speed and efficiency, DenseNet169 offered a balance between depth and complexity, while DenseNet201 provided the deepest architecture with superior feature extraction capability.

An ensemble voting strategy was introduced to improve prediction robustness and accuracy. In this approach, each of the three DenseNet models independently predicted a probability vector for all classes. The final prediction was obtained by averaging these vectors element-wise, as expressed by:

$$P_{\text{ensemble}} = \frac{1}{n} \sum_{i=1}^{n} P_i \tag{3}$$

The term P_i refers to the predicted probability vector from the i^{th} model, and n is the total number of models. The class with the highest averaged probability was selected as the final prediction. This method was preferred over stacking or weighted voting due to its simplicity and effectiveness, particularly with a limited-sized dataset.

To further capture both spatial and sequential patterns, a hybrid model combining DenseNet201 and LSTM was constructed. DenseNet201, pre-trained on ImageNet, was used exclusively as a feature extractor by freezing its convolutional layers and excluding its classification head. The extracted features were passed through a GlobalAveragePooling2D layer, reshaped into sequences, and then processed using an LSTM layer to learn spatial relationships within the image features. Dropout and additional dense layers with ReLU activation were used to introduce non-linearity and regularization. The final classification was performed using a softmax layer. The hybrid model was compiled using the Adam optimizer, described by the equations:

$$\mathrm{m_t} = \beta_1 m_{\mathrm{t}-1} + (1 - \beta_1) g_{\mathrm{t}}, \quad v_{\mathrm{t}} = \beta_2 v_{\mathrm{t}-1} + (1 - \beta_2) g_{\mathrm{t}}^{2} \tag{4}$$

The training process employed categorical cross-entropy as the loss function, given by:

$$\mathcal{L} = -\sum_{i=1}^{C} y_{\mathrm{i}} \, log(\hat{\mathrm{y}}_{\mathrm{i}}) \tag{5}$$

Model performance was evaluated using standard metrics such as accuracy, precision, recall, and F1-score. Among all tested architectures, the hybrid DenseNet201-LSTM model achieved the highest validation accuracy, exceeding 93%. This result highlights the strength of combining deep spatial feature extraction with sequential modeling for complex image classification tasks such as medicinal plant identification.

4. Results and Discussion

This study systematically evaluated multiple deep learning models for classifying 40 species of Indian medicinal plants. The baseline CNN model achieved a validation accuracy of 69.58%, while transfer learning models showed progressively better performance: DenseNet121 reached 77.20%, DenseNet169 reached 81.18%, and DenseNet201 achieved 83.55%. An ensemble of these three DenseNet models further improved accuracy to 89.66%, and finally, the hybrid DenseNet201-LSTM model achieved the best performance with an accuracy of 93.38%. These results confirm the benefit of using deeper networks and hybrid learning in fine-grained classification tasks involving complex visual patterns.

Figure 2 illustrates the training and validation accuracy trends of the CNN model over 20 epochs. The training accuracy shows a consistent upward trend, starting below 0.2 and reaching close to 0.80 by the final epoch. This indicates that the model is effectively learning patterns from the training data. However, the validation accuracy begins to plateau around epoch 12 and fluctuates thereafter, peaking below 0.70. This divergence between training and validation accuracy in the later stages of training suggests that the model is beginning to overfit — learning the training data too well while failing to generalize effectively to unseen validation data. The increasing gap between the two curves highlights the CNN model's limitations in capturing complex features needed for distinguishing between visually similar medicinal plant species.

Figure 3 presents the training and validation loss curves for the same CNN model. Both losses decrease steadily during the early training phases, indicating successful error minimization. However, after epoch 10, the validation loss decreases at a slower rate and begins to diverge from the training loss. This growing gap between the two loss curves, particularly after epoch 15, reinforces the earlier observation of overfitting. While the model continues to perform well on training data, its performance on the validation set stagnates. These trends confirm that although the CNN architecture is capable of learning basic visual features, it lacks the representational depth and generalization ability necessary for accurate multi-class classification in this domain. More advanced architectures or transfer learning techniques are required to improve performance.

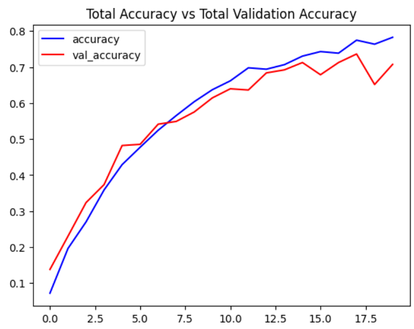

Figure 2. Total Validation Accuracy vs Total Accuracy of CNN

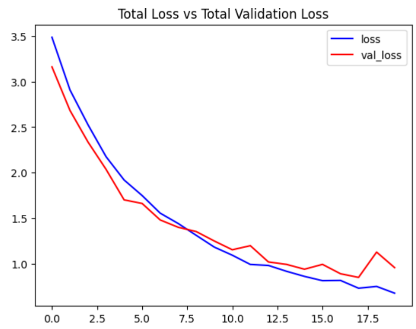

Figure 3. Total Validation Loss vs Total Loss of CNN

Figure 4 presents the training and validation accuracy curves of the DenseNet121 model over 20 epochs. Unlike the CNN model, the DenseNet121 architecture exhibits a strong alignment between training and validation accuracy throughout the training process. Both curves rise rapidly in the initial epochs and continue to improve consistently, ultimately converging near 0.80. The close proximity of these two curves suggests that the model is learning in a stable and generalized manner, without significant overfitting. This performance improvement is a direct result of leveraging transfer learning with pre-trained ImageNet weights, which allows the model to begin training from a strong feature representation base.

Figure 5 shows the corresponding training and validation loss curves for DenseNet121. Both loss values decrease smoothly and consistently over the epochs, with the validation loss closely following the training loss. By the end of training, both curves settle below a value of 0.6, and the narrow gap between them indicates minimal generalization error. This balance between learning and generalization confirms that DenseNet121 effectively adapts to the classification task, making it significantly more robust and reliable than the baseline CNN model. The stable convergence of loss values further validates the architecture's capacity to extract meaningful features while resisting overfitting.

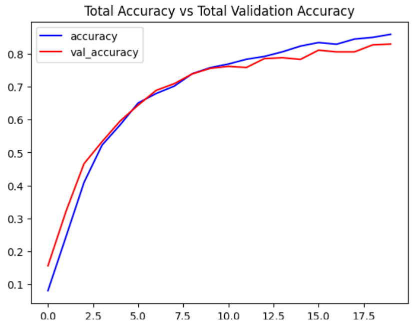

Figure 4. Total Validation Accuracy vs Total Accuracy of DenseNet121

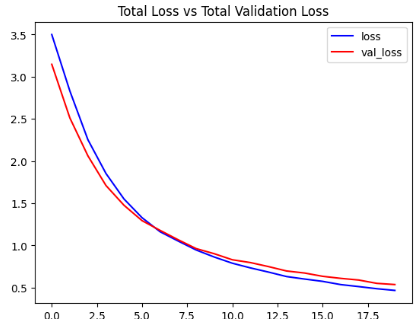

Figure 5. Total Validation Loss vs Total Loss of DenseNet121

Figure 6 illustrates the training and validation accuracy curves of the DenseNet169 model across 20 epochs. From the early stages of training, both curves show a steep and consistent upward trend, with the validation accuracy closely tracking the training accuracy throughout. By the final epoch, the model reaches a validation accuracy of approximately 81.18%, confirming a noticeable improvement over previous architecture. The tight alignment between the two curves suggests that the model maintains a balanced learning process, exhibiting strong generalization capabilities. This indicates that DenseNet169 successfully leverages its deeper architecture and densely connected layers to extract more abstract and discriminative features from the leaf image dataset.

Figure 7 presents the training and validation loss curves for DenseNet169. Both losses decrease progressively and in parallel, showing no significant divergence between training and validation loss values. This smooth and synchronized downward trend demonstrates that the model is learning effectively across epochs without overfitting. The consistent reduction in validation loss also indicates that the model retains its ability to perform well on unseen data, further validating the efficacy of the DenseNet169 architecture. The performance gain achieved here can be attributed to the network's enhanced capacity for feature reuse and gradient flow, which are key strengths of the DenseNet family.

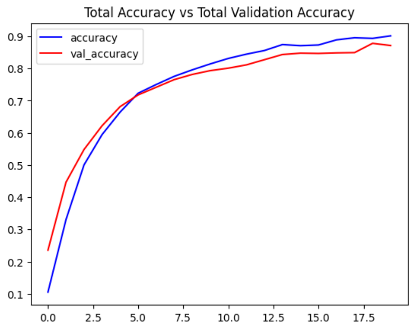

Figure 6. Total Validation Accuracy vs Total Accuracy of DenseNet169

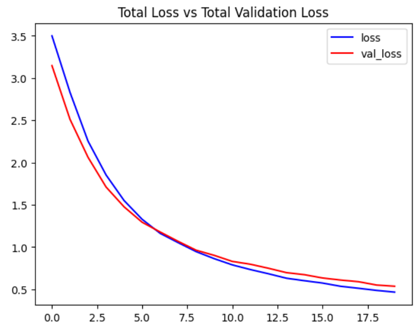

Figure 7. Total Validation Loss vs Total Loss of DenseNet169

Figure 8 displays the training and validation accuracy curves for the DenseNet201 model across 20 epochs. Among all transfer learning models tested, DenseNet201 demonstrates the most stable and highest performance, with validation accuracy reaching approximately 83.55%. Throughout the training process, the validation accuracy closely follows the training accuracy, with both curves showing consistent and smooth growth. This close alignment between the two curves suggests that the model is not only learning effectively but also generalizing well to unseen data. The result reflects DenseNet201's enhanced representational power due to its deeper architecture, which enables it to capture fine-grained spatial patterns in complex leaf images more effectively than its predecessors.

Figure 9 shows the corresponding training and validation loss curves for DenseNet201. Both curves decline steadily and in parallel, with validation loss consistently tracking slightly above the training loss, indicating very low generalization error. By the final epoch, both losses converge near 0.5, confirming that the model maintains high prediction accuracy while avoiding overfitting. The consistently small gap between training and validation loss further highlights the robustness and stability of DenseNet201, making it a highly suitable choice for complex image classification tasks such as medicinal plant identification. These results affirm that the model's deep and densely connected layers effectively facilitate feature reuse and gradient propagation, leading to superior learning efficiency.

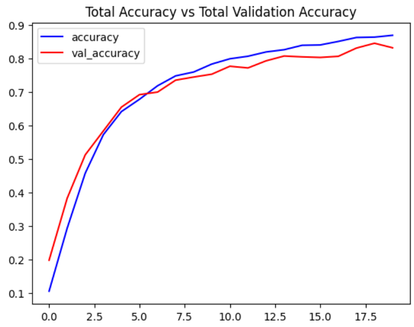

Figure 8. Total Validation Accuracy vs Total Accuracy of DenseNet201

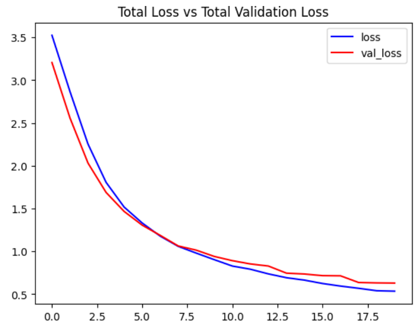

Figure 9. Total Validation Loss vs Total Loss of DenseNet201

Figure 10 illustrates the training and validation accuracy curves of the hybrid DenseNet201-LSTM model over a limited number of epochs. The model demonstrates a rapid convergence, reaching a validation accuracy of approximately 93.38% within the first few epochs. Interestingly, the validation accuracy slightly exceeds the training accuracy in some epochs, which is often indicative of a well-regularized model. This behavior can be attributed to the incorporation of dropout layers and the LSTM component, which enhances the model’s ability to capture long-range spatial dependencies across feature sequences generated by DenseNet201. The early saturation of the accuracy curve, combined with its high peak, reflects the model's ability to learn discriminative features quickly and generalize effectively to unseen data.

Figure 11 presents the corresponding training and validation loss curves. The training loss remains consistently low and stable, while the validation loss fluctuates slightly but stays within a narrow range between 0.14 and 0.22, showing no signs of escalation. This suggests that the model avoids overfitting despite its rapid learning behavior. The low magnitude of both losses confirms that the hybrid model maintains robust performance throughout training. The combination of DenseNet201’s deep feature extraction capabilities with LSTM’s temporal modeling results in a highly expressive and generalizable architecture, making it particularly well-suited for classifying complex leaf images in fine-grained tasks such as medicinal plant identification.

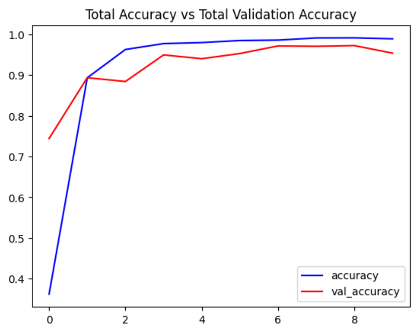

Figure 10. Total Validation Accuracy vs Total Accuracy of Loss of Hybrid Model LSTM

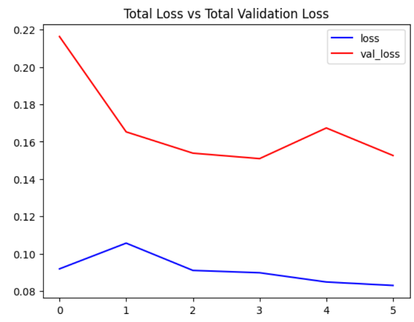

Figure 11. Total Validation Loss vs Total Hybrid Model (DenseNet201-DenseNet201- LSTM

The experimental results, as presented in table 1, clearly demonstrate that increasing model complexity leads to consistent improvements in classification performance. Starting from a basic CNN model and progressing through the DenseNet family to a hybrid DenseNet201-LSTM model, each step introduces architectural enhancements that translate into better accuracy, precision, recall, and F1-score. The CNN model, while capable of learning low-level features, achieves only moderate accuracy and shows signs of overfitting, as evidenced by the divergence between training and validation performance. In contrast, the DenseNet models make use of pre-trained weights and dense connectivity to extract more meaningful patterns, leading to a clear upward trend in performance across evaluation metrics.

Table 1. Overall Accuracy Result

Algorithms	CNN	Dense Net121	Dense Net169	Dense Net201	Voting Ensemble	Hybrid (DenseNet201-LSTM)
Predicted Plant	Hibiscus	Aloe Vera	Aloe Vera	Aloe Vera	Hibiscus	Hibiscus
Accuracy	0.6958	0.7720	0.8118	0.8355	0.8966	0.9338
Precision	0.7194	0.7808	0.8178	0.8365	0.8998	0.9474
Recall	0.6958	0.7720	0.8118	0.8355	0.8966	0.9338
F1 Score	0.6988	0.7699	0.8101	0.8331	0.8947	0.9342

Among the DenseNet variants, DenseNet201 stands out by achieving an accuracy of 83.55 percent, demonstrating the benefit of deeper feature representations. The model shows high training stability, with minimal divergence between training and validation accuracy and loss. To further enhance performance, a voting ensemble that integrates DenseNet121, DenseNet169, and DenseNet201 is applied. This ensemble achieves 89.66 percent accuracy by averaging the softmax outputs from the three models. The ensemble approach reduces classification errors by balancing the strengths and weaknesses of individual architectures and is particularly useful for complex classification tasks involving subtle visual distinctions.

The best overall results are obtained from the hybrid DenseNet201-LSTM model, which integrates convolutional feature extraction with sequential modeling. This architecture achieves an accuracy of 93.38 percent, the highest among all models tested. DenseNet201 is used to extract deep spatial features, which are then reshaped and passed into an LSTM layer that captures sequential relationships across feature dimensions. The LSTM enhances the model's ability to understand the structure and variation of leaf morphology. In addition, dropout layers contribute to regularization, helping the model avoid overfitting despite its complexity. This combination proves especially effective in identifying nuanced differences between visually similar medicinal plants.

As shown in table 1, the hybrid model also attains the highest scores in precision (0.9474), recall (0.9338), and F1-score (0.9342). These results clearly surpass those of the CNN, individual DenseNet models, and the ensemble. Each progression in model design results in measurable gains, both in predictive accuracy and generalization capability. The hybrid architecture not only meets the technical demands of fine-grained classification but also provides a strong foundation for practical applications in medicinal plant recognition. This confirms that combining convolutional and recurrent layers offers a powerful strategy for solving complex image analysis problems in real-world environments.

5. Conclusion

This study successfully implemented and evaluated multiple deep learning approaches for medicinal plant classification, including a baseline CNN, individual transfer learning models using DenseNet121, DenseNet169, and DenseNet201, a voting ensemble of these DenseNet variants, and a hybrid model combining DenseNet201 with an LSTM layer. These methods were tested on a dataset comprising 40 diverse plant species and demonstrated strong classification performance, with the hybrid DenseNet201-LSTM model achieving the highest accuracy. The research underscores the potential of deep learning in supporting accurate and scalable identification of medicinal plants, offering a promising alternative to traditional, labor-intensive identification methods.

The outcomes of this project have important implications for improving the reliability of plant identification in the herbal medicine supply chain. By reducing the risk of misidentification, these models can help safeguard the authenticity and quality of medicinal plant products, ultimately building greater trust in traditional healing practices. Furthermore, this technology can support field practitioners, researchers, and consumers by providing accessible and objective tools for identifying medicinal plants, thereby contributing to the integrity of traditional medicine systems and promoting public health.

Despite the encouraging results, one of the key limitations encountered is the relatively small size of the dataset, which may constrain the model's generalizability to broader plant populations and environmental contexts. To address this, future efforts will focus on expanding the dataset through collaboration with botanical institutions and conservation organizations. Incorporating environmental metadata, such as soil type, climate, and regional variation, could further enhance the model's robustness and accuracy. However, these additions will introduce new challenges, including class imbalance and variability, which can be mitigated through data augmentation techniques and re-sampling strategies.

Moving forward, the research will be operationalized through the development of a mobile application using TensorFlow Lite, allowing for real-time plant identification on low-power devices typically used in rural and resource-limited areas. The app will feature a user feedback system and validation mechanism to ensure continuous improvement and community engagement. Educational outreach and awareness campaigns will accompany deployment to promote sustainable harvesting and biodiversity conservation. Long-term success will be supported by international partnerships, enabling cross-border knowledge sharing and contributing to global standards for medicinal plant identification and preservation.

6. Declarations

6.1. Author Contributions

Conceptualization: P.N., M.B., D.S.; Methodology: P.N., H.K.; Software: D.T., A.S.; Validation: M.B., H.K.; Formal Analysis: P.N.; Investigation: D.S., D.T.; Resources: M.B., H.K.; Data Curation: D.S.; Writing – Original Draft Preparation: P.N.; Writing – Review and Editing: H.K., A.S.; Visualization: D.T.; All authors have read and agreed to the published version of the manuscript.

6.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

6.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

6.4. Institutional Review Board Statement

Not applicable.

6.5. Informed Consent Statement

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] R. Azadnia and K. Kheiralipour, "Recognition of leaves of different medicinal plant species using a robust image processing algorithm and artificial neural networks classifier," *Journal of Applied Research on Medicinal and Aromatic Plants*, vol. 25, no. 1, pp. 1-12, 2021. doi: 10.1016/j.jarmap.2021.100327

[2] D. Barhate, S. Pathak, and A. Dubey, "A deep learning methodology for plant species recognition using morphology of leaves," *ASEAN Engineering Journal*, vol. 13, no. 4, pp. 95–102, 2023. doi: 10.11113/aej.v13.19461

[3] N. Bhelkar and A. Sharma, "Identification and classification of medicinal plants using leaf with deep convolutional neural networks," *International Journal of Health Sciences*, vol. 6, no. S6, pp. 11596–11605, 2022. doi: 10.53730/ijhs.v6nS6.13233

[4] J. Biplob Dey, R. A. Jannatul Ferdous, and J. Hossain, "Assessing deep convolutional neural network models and their comparative performance for automated medicinal plant identification from leaf images," *Heliyon*, vol. 10, no. 1, pp. 1-21, 2024. doi: 10.1016/j.heliyon.2023.e23655

[5] C. Yuan, T. Liu, S. S. F. G. R. Zhang, “Research on plant species identification based on improved convolutional neural network,” *Phyton – International Journal of Experimental Botany*, vol. 92, no. 4, pp. 1037–1058, 2023. doi: 10.32604/phyton.2023.025343

[6] M. Deshmukh, “Deep learning for the classification and recognition of medicinal plant species,” *Indian Journal of Science and Technology*, vol. 17, no. 11, pp. 1070–1077, 2024. doi: 10.17485/IJST/v17i11.3099

[7] H. Diwedi, A. Misra, and A. Tiwari, “CNN-based medicinal plant identification and classification using optimized SVM,” *Multimedia Tools and Applications*, vol. 83, no. 1, pp. 1–31, 2023. doi: 10.1007/s11042-023-16733-8

[8] H. Goëau, P. Bonnet, and A. Joly, “Overview of PlantCLEF 2022: Image-based plant identification at global scale,” in *CLEF Working Notes*, vol. 2022, no. 1, pp. 1916–1928, 2022.

[9] B. P. Gyires-Tóth, M. Osváth, D. Papp, and G. Szűcs, “Deep learning for plant classification and content-based image retrieval,” *Cybernetics and Information Technologies*, vol. 19, no. 1, pp. 88–100, 2019. doi: 10.2478/cait-2019-0005

[10] G. Kayhan and E. Ergün, “Medicinal and aromatic plants identification using machine learning methods,” *Balkan Journal of Electrical and Computer Engineering*, vol. 8, no. 1, pp. 81–87, 2020. doi: 10.17694/bajece.651286

[11] T. L. Le, M. V. H. Do, N. H. Van, H. Q. Nguyen, T. T. Trong, D. D. Phan, dan V. S. Hoang, “Robust plant identification based on the combination of multiple images and taxonomic information,” *IEEE Access*, vol. 12, no. 1, pp. 67726–67737, 2024. doi: 10.1109/ACCESS.2024.3399835

[12] M. S. Hasibuan, R. Z. Abdul Aziz, D. A. D. T. B. Kusuma, dan N. A. Syafira, “Recommendation model for learning material using the Felder Silverman learning style approach,” *HighTech and Innovation Journal*, vol. 4, no. 4, pp. 1–12, 2023. doi: 10.28991/HIJ-2023-04-04-010

[13] O. A. Malik, N. Ismail, B. R. Hussein, dan U. Yahya, “Automated real-time identification of medicinal plants species in natural environment using deep learning models — A case study from Borneo region,” *Plants*, vol. 11, no. 15, pp. 1–17, 2022. doi: 10.3390/plants11151952

[14] A. K. Mulugeta, D. P. Sharma, dan A. H. Mesfin, “Deep learning for medicinal plant species classification and recognition: A systematic review,” *Frontiers in Plant Science*, vol. 14, no. 1, pp. 1–12, 2023. doi: 10.3389/fpls.2023.1286088

[15] S. Naeem, A. Ali, C. Chesneau, M. Tahir, F. Jamal, R. Khan, dan M. Ul Hassan, “The classification of medicinal plant leaves based on multispectral and texture feature using machine learning approach,” *Agronomy*, vol. 11, no. 2, pp. 263-182, 2021. doi: 10.3390/agronomy11020263

[16] K. Suresh, T. Kumar, E. Naresh, V. Kalmani, K. Bamane, dan P. Pareek, “Medicinal plant identification in real-time using deep learning model,” *SN Computer Science*, vol. 5, no. 1, pp. 1–12, 2024. doi: 10.1007/s42979-023-02398-5

[17] S. Oppong, F. Twum, J. Hayfron-Acquah, dan Y. Missah, “A novel computer vision model for medicinal plant identification using log-Gabor filters and deep learning algorithms,” *Computational Intelligence and Neuroscience*, vol. 2022, art. ID 1189509, no. 1, pp. 1–12, 2022. doi: 10.1155/2022/1189509

[18] S. Sachar dan A. Kumar, “Deep ensemble learning for automatic medicinal leaf identification,” *International Journal of Information Technology*, vol. 14, no. 6, pp. 3089–3097, 2022. doi: 10.1007/s41870-022-01055-z

[19] X. Sun, G. Li, P. Qu, X. Xie, X. Pan, dan W. Zhang, “Research on plant disease identification based on CNN,” *Cognitive Robotics*, vol. 2, no. 1, pp. 155–163, 2022. doi: 10.1016/j.cogr.2022.07.001

[20] N. Tiwari, B. Gupta, A. Prakash, K. Tiwari, S. Alshmrany, A. Ali, M. Husain, dan D. Singh, “Developing a neural network-based model for identifying medicinal plant leaves using image recognition techniques,” *Journal of Advanced Zoology*, vol. 44, no. S-5, pp. 1944–1958, 2023. doi: 10.17762/jaz.v44iS-5.1564

[21] C. S. Pereira, R. Morais, dan M. J. C. S. Reis, “Deep learning techniques for grape plant species identification in natural images,” *Sensors*, vol. 19, no. 22, pp. 4850-4863, 2019. doi: 10.3390/s19224850

[22] H. N. Van dan N. L. H. Hien, “Automatic plant image identification of Vietnamese species using deep learning models,” *arXiv preprint*, arXiv:2005.02832, vol. 2020, no. 1, pp. 1-12, 2020.

[23] A. K. H. Riyashantha dan C. Mahendranathan, “Traditional uses of medicinal plants in Sri Lanka with special reference to herbal drinks: A review,” *Journal of Biology and Nature*, vol. 11, no. 2, pp. 1–13, 2020.

[24] K. Pushpanathan, M. Hanafi, S. Mashohor, dan W. F. F. Ilahi, “Machine learning in medicinal plants recognition: a review,” *Artificial Intelligence Review*, vol. 54, no. 1, pp. 305–327, 2021. doi: 10.1007/s10462-020-09843-5

[25] D. M. Sasikala, “Comprehensive study of medicinal plant classification using artificial intelligent techniques,” *International Journal of Emerging Technologies and Innovative Research (JETIR)*, vol. 7, no. 3, pp. 12–18, Mar 2020.

[26] T. T. Tin, K. J. Xin, A. Aitizaz, L. K. Tiung, T. C. Keat, and H. Sarwar, “Machine learning based predictive modelling of cybersecurity threats utilising behavioural data,” *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 9, pp. 1-9, 2023, doi: 10.14569/IJACSA.2023.0140987.

[27] J. Wäldchen dan P. Mäder, “Plant species identification using computer vision techniques: A systematic literature review,” *Archives of Computational Methods in Engineering*, vol. 25, no. 2, pp. 507–543, 2018. doi: 10.1007/s11831-016-9206-z

[28] B. R. Pushpa dan S. Rani, “Indian Medicinal Leaves Image Datasets,” *Mendeley Data*, vol. 3, May 2023. doi: 10.17632/748f8jkphb.3