International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 04 Issue: 04 | April – 2025 DOI: 10.55041/ISJEM03019
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Deep Learning Models for Skin Disease Classification: A DenseNet Perspective

Nilamjyoti Changmai
nilamjyotichangmai@gmail.com
Department of Computer Science
Jain University, Bangalore, India

Abu Bakar Shibli
abubakarshibli9@gmail.com
Department of Computer Science
Jain University, Bangalore, India

Saraswat Akshay Anand
akshayanand206@gmail.com
Department of Computer Science
Jain University, Bangalore, India

Dr. Rhea Sriniwas
rhea.sriniwas@jainuniversity.ac.in
Department of Computer Science
Jain University, Bangalore, India

Abstract:

Skin diseases represent a considerable public health challenge worldwide, affecting people of all ages and demographics. Failure to detect and treat these diseases early can result in severe complications, both medical and psychological. While dermatologists traditionally rely on visual assessments and patient history for diagnosis, this approach can be subjective and time-consuming. In recent years, the use of artificial intelligence, particularly deep learning, has gained traction in medical image analysis due to its ability to automate complex classification tasks. This paper proposes an effective approach for classifying skin conditions using DenseNet121—a convolutional neural network pre-trained on ImageNet. Through transfer learning and fine-tuning, we adapt DenseNet121 for skin disease classification using an annotated dataset of dermatological images. Our proposed framework is evaluated based on several performance metrics including accuracy, precision, recall, and F1-score. The results underscore the potential of the model in real- world clinical scenarios, especially in resource-limited settings.

1. Introduction

Skin diseases are among the most prevalent health concerns globally, affecting approximately 30% to 70% of individuals at some point in their lives. These conditions vary from mild irritations to life-threatening illnesses like melanoma. In many cases, untreated skin disorders can escalate, leading to secondary infections or systemic complications. Dermatological conditions not only result in physical discomfort but also significantly impair the psychological and emotional well-being of patients, often resulting in stigma and social isolation.

Diagnosing skin diseases accurately and efficiently is vital for improving patient outcomes. However, the traditional method of diagnosis—relying primarily on dermatologists' visual inspections—is often hindered by limited specialist availability and inconsistent diagnostic accuracy. In underserved regions, access to qualified dermatologists may be severely constrained, exacerbating health disparities and delaying treatment. Consequently, there is a growing demand for automated, AI-powered diagnostic systems that can bridge this gap.

Deep learning, a subset of machine learning, uses artificial neural networks with multiple hidden layers to learn data representations in an automated manner. Convolutional neural networks (CNNs), in particular, have achieved groundbreaking performance in tasks such as object detection, facial recognition, and medical image classification. These models are capable of learning hierarchical features from raw image pixels, making them highly suitable for dermatological applications.

This study investigates the effectiveness of a DenseNet121-based deep learning model for classifying images of skin conditions. We utilize transfer learning techniques to repurpose a model originally trained on a broad dataset (ImageNet) for our specialized task. The approach demonstrates the model's potential to improve diagnostic accuracy and speed, especially in areas

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 04 Issue: 04 | April – 2025 DOI: 10.55041/ISJEM03019
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

lacking expert dermatological care.

2. Background and Related Work:

A variety of AI-driven methods have been proposed for skin disease classification. Earlier studies focused on traditional machine learning techniques involving manual feature extraction, such as using histogram of oriented gradients (HOG) or support vector machines (SVM). These methods, while moderately effective, are limited by their dependency on domain-specific features and handcrafted inputs.

With the emergence of deep learning, particularly CNNs, a significant shift occurred in the landscape of image-based disease classification. Esteva et al. (2017) trained a CNN on over 130,000 clinical images, demonstrating dermatologist-level accuracy in identifying skin cancer. The study utilized the Inception v3 architecture, highlighting the feasibility of AI-powered diagnostics in dermatology.

Tschandl et al. (2019) validated deep learning's superiority by comparing human dermatologists to an ensemble of CNNs in diagnosing melanoma. The deep learning models outperformed most clinicians, especially in distinguishing malignant lesions from benign ones. Similar outcomes were reported by Han et al. (2018), who used a dataset of over 12,000 images to train a CNN that achieved high accuracy across multiple skin disease categories.

Despite these successes, challenges persist. Many existing datasets lack sufficient diversity, limiting model generalization across different skin tones and demographics. Additionally, the presence of image artifacts, inconsistent lighting, and varying lesion sizes complicate the classification process. Another limitation is the lack of explainability in deep learning models, which poses a barrier to clinical adoption.

Our research contributes to this evolving field by addressing these challenges through data preprocessing, model fine-tuning, and rigorous evaluation using multiple metrics.

3. Methodology:

In this research, we propose a deep learning-based approach for skin disease classification using the DenseNet121 architecture. The methodology consists of the following steps:

3.1 Dataset Description:

For this study, we utilized a dataset consisting of labelled images representing various skin diseases. The dataset includes categories such as melanoma, eczema, psoriasis, and dermatitis. Images were sourced from public repositories including the HAM10000 dataset and the ISIC archive. The dataset was curated to include images from multiple demographics to ensure representation and robustness.

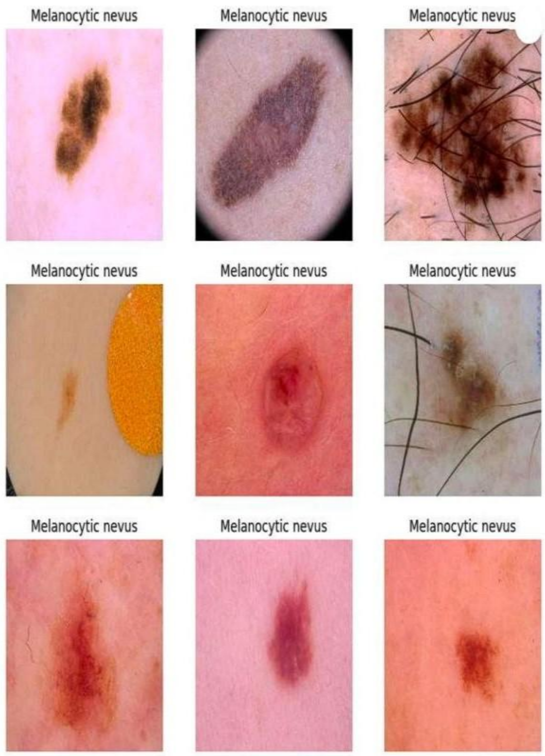

Fig. Skin Disease Dataset Images

3.2 Preprocessing Pipeline: Prior to model training, several preprocessing steps were applied to enhance data quality. Each image was resized to 224x224 pixels, converted to RGB format, and normalized to scale pixel values between 0 and 1. We also implemented data augmentation techniques such as horizontal flipping, random cropping, rotation, and zooming to increase dataset variability and minimize overfitting.

3.3 Model Architecture and DenseNet121 Overview: DenseNet (Densely Connected Convolutional Networks) is a CNN architecture that connects each layer to every other layer in a feed-forward fashion. DenseNet121, a specific implementation with 121 layers, is known for its efficient use of parameters and strong feature propagation. We adopted DenseNet121 pre-trained on ImageNet and customized its final classification layer to match our number of target classes

$$x_l = H_l([x_0, x_1, , x_{l-1}])$$

3.4 Transfer Learning and Fine-Tuning: We employed transfer learning by freezing the convolutional base of the DenseNet121 and replacing the classifier head with a global average pooling layer, followed by a dropout layer (to prevent overfitting), and a final softmax output layer. After initial training, we selectively unfroze deeper layers for fine-tuning to further enhance model performance.

3.5 Model Compilation and Training: The model was compiled using the Adam optimizer with a learning rate of 0.0001 and sparse categorical cross-entropy as the loss function. Training was performed in batches using a stratified k-fold cross-validation approach to ensure robustness and fairness. We monitored model performance using validation accuracy and loss over 50 epochs.

3.6 Evaluation Metrics: Performance was assessed using:

Accuracy: Ratio of correctly predicted images to the total images.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

Precision: True positives divided by the sum of true and false positives.

$$Precision = \frac{TP}{TP + FP}$$

Recall: True positives divided by the sum of true positives and false negatives.

$$Recall = \frac{TP}{TP + FN}$$

F1-score: Harmonic mean of precision and recall.

$$F1 - score = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall}$$

Confusion Matrix: Visualization of predictions across all classes.

AUC-ROC Curve: Evaluates classification quality across thresholds.

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 04 Issue: 04 | April – 2025 DOI: 10.55041/ISJEM03019
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

4. Experimental Results and Analysis:

4.1 Quantitative Results: The final model achieved the following results on the test dataset:

METRIC	VALUE
Accuracy	85.4%
Precision	84.7%
Recall	86.1%
F1-Score	85.1%
AUC Score	0.92

.

These values indicate consistent classification performance across classes, demonstrating the model's capacity to generalize well.

4.2 Confusion Matrix Interpretation:

The confusion matrix revealed that melanoma and psoriasis were the most accurately classified categories. However, there were some misclassifications between dermatitis and eczema due to visual similarities. Additional samples or metadata might improve discrimination between these classes.

	Predicted Normal	Predicted Disease
Actual Normal	TN	FN
Actual Disease	FN	TN

4.3 Visual Inspection and Feature Maps: Using Grad-CAM, we visualized the model’s focus areas within test images. The heatmaps confirmed that the model predominantly attends to lesion areas, aligning with dermatological expectations. This interpretability enhances clinical trust in model predictions.

4.4 Comparison with Baseline Models: We compared DenseNet121 with other architectures like VGG16 and ResNet50. DenseNet121 outperformed both in terms of accuracy and training efficiency, reaffirming its suitability for this task.

5. Discussion:

The integration of DenseNet121 for skin disease classification offers a promising avenue for improving diagnostic processes. Its dense connectivity allows for better feature reuse, which contributes to faster convergence and improved performance. Moreover, transfer learning enables effective training even with limited data, addressing one of the primary challenges in medical imaging.

However, challenges such as class imbalance, limited data diversity, and lack of clinical explainability remain. Further enhancements can be made by incorporating metadata (e.g., patient age, lesion location), multi-modal inputs, or ensemble learning approaches.

Additionally, ethical considerations must be addressed, including data privacy, informed consent, and algorithmic fairness. Ensuring that the model performs equally well across skin tones is crucial to avoid biases.

6. Conclusion:

This study presented a deep learning framework utilizing DenseNet121 for classifying various skin conditions from image data. Our model achieved high accuracy and robustness, validating the efficacy of using transfer learning in dermatological diagnostics. It is particularly beneficial in remote and underserved regions where access to specialists is limited.

In future work, we plan to:

Expand the dataset to include rare and region-specific skin conditions.

Enhance the model’s interpretability using advanced explainable AI techniques.

Integrate clinical metadata for holistic diagnosis.

Deploy the model in a real-time mobile health application.

By continuing to refine and expand upon this work, AI-based tools can become essential assets in global dermatological care.

7. References:

1. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115-118.

2. National Institutes of Health (NIH) Clinical Center. (n.d.). Retrieved from https://clinicalcenter.nih.gov/

3. DermNet NZ. (n.d.). Dermatology Image Database HAM10000. Retrieved from https://dermnetnz.org/

4. International Skin Imaging Collaboration (ISIC). (n.d.). Retrieved from https://www.isic-archive.com/

5. Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1251-1258).

6. Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

7. Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition (pp. 248-255).

8. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778).

9. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.