

An Efficient Hybrid Deep Learning Model For Skin Cancer Detection And Comparative Analysis

Rashmi Sinha¹, Dr. Anil Kumar Sagar²

¹Sharda University, Greater Noida, India

²Professor, Sharda University, Greater Noida, India

Corresponding Author: Dr. Anil Kumar Sagar²

Abstract

Skin cancer remains one of the most prevalent and life-threatening forms of cancer globally, with early detection being critical for successful treatment outcomes. The rapid advancement of deep learning (DL) in medical image analysis has opened new avenues for automated, accurate, and non-invasive diagnostic tools. This study proposes an efficient hybrid deep learning model for skin cancer detection, leveraging the complementary strengths of a pre-trained Convolutional Neural Network (CNN) for feature extraction and a custom classifier for precise classification. A comprehensive comparative analysis is conducted against state-of-the-art models including VGG16, ResNet50, InceptionV3, and a standalone CNN. The models were trained and evaluated on the publicly available HAM10000 dataset, which comprises over 10,000 dermatoscopic images across seven diagnostic categories. The proposed hybrid model, which integrates a fine-tuned EfficientNetB0 backbone with a multi-head attention mechanism and optimized dense layers, achieved superior performance, with an overall accuracy of 93.8%, precision of 93.2%, recall of 93.5%, and an F1-score of 93.3%. These results significantly outperform the benchmark models, demonstrating the efficacy of the hybrid architecture. The findings underscore the potential of hybrid deep learning approaches in developing robust, automated skin cancer screening systems that can assist clinicians in making faster and more accurate diagnoses, ultimately improving patient outcomes.

Keywords: Skin Cancer Detection; Deep Learning; Hybrid Model; Convolutional Neural Network; Medical Image Analysis; HAM10000; EfficientNet; Transfer Learning

1. Introduction

Skin cancer is a global health concern, with its incidence rising steadily over the past decades. According to the World Cancer Research Fund, skin cancer accounts for a significant proportion of all new cancer cases worldwide, with melanoma, the deadliest form, being particularly aggressive if not detected early (World Cancer Research Fund, n.d.). Early and accurate diagnosis is paramount, as it dramatically increases the survival rate and reduces the morbidity associated with invasive treatments. Traditional diagnosis relies on visual inspection by dermatologists followed by histopathological examination of biopsies. This process is subjective, time-consuming, and heavily dependent on the clinician's expertise, leading to potential delays and diagnostic variability (Tschandl et al., 2018).

The advent of artificial intelligence (AI), particularly deep learning (DL), has revolutionized medical image analysis, offering powerful tools to augment clinical decision-making (Helaly et al., 2024). DL models, especially Convolutional Neural Networks (CNNs), have demonstrated remarkable capabilities in recognizing complex patterns in medical images, achieving expert-level performance in tasks such as breast cancer imaging, neurological disorder analysis, and knee osteoporosis diagnosis (Luo et al., 2024; Hussain & Nazir, 2024; Sarhan et al., 2024). In dermatology, deep learning has shown immense promise in classifying skin lesions from dermatoscopic images, with models like CNNs achieving accuracy comparable to board-certified dermatologists (Tschandl et al., 2018; Maduranga & Nandasena, 2022).

Despite these advances, challenges remain. The performance of single CNN architectures can be limited by dataset size, class imbalance, and the subtle visual differences between benign and malignant lesions (Alwakid et al., 2022). Hybrid models, which combine the strengths of multiple architectures or integrate attention mechanisms, have emerged as a promising solution to these limitations. By leveraging pre-trained networks for robust feature extraction and incorporating additional layers for refined classification, hybrid models can capture both global and local features more effectively (Reyad et al., 2024; Devarakonda et al., 2024). Furthermore, the integration of explainability techniques can enhance trust and adoption in clinical settings (Mehta et al., 2024).

This study aims to develop and validate an efficient hybrid deep learning model for skin cancer detection. The primary objectives are: (1) to design a hybrid model that combines a pre-trained CNN (EfficientNetB0) with a multi-head attention mechanism and a custom classifier; (2) to evaluate its performance on the HAM10000 dataset; and (3) to conduct a comparative analysis against standard pre-trained models (VGG16, ResNet50, InceptionV3) and a baseline CNN. This research contributes a robust, high-performance model that can serve as a reliable tool for automated skin cancer screening, thereby supporting the broader integration of AI in healthcare systems (Kumari et al., 2023; Badawy et al., 2023).

2. Materials and Methods

2.1. Dataset


This study utilized the HAM10000 ("Human Against Machine with 10000 training images") dataset, a large, publicly available collection of multi-source dermatoscopic images of common pigmented skin lesions (Tschandl et al., 2018). The dataset comprises 10,015 dermatoscopic images across seven diagnostic categories:

- akiec: Actinic keratoses and intraepithelial carcinoma
- bcc: Basal cell carcinoma
- bkl: Benign keratosis-like lesions
- df: Dermatofibroma
- mel: Melanoma
- nv: Melanocytic nevi
- vasc: Vascular lesions

The dataset exhibits a significant class imbalance, with 'nv' (melanocytic nevi) being the most prevalent and 'df' (dermatofibroma) and 'vasc' (vascular lesions) being the least. The images were pre-processed by resizing to 224x224 pixels, normalizing pixel values to the range [0,1], and applying data augmentation (rotation, width/height shift, horizontal flip, zoom) to mitigate overfitting and address class imbalance. The dataset was split into training (70%, 7,010 images), validation (15%, 1,502 images), and testing (15%, 1,503 images) sets, stratified by class to preserve the class distribution.

2.2. Proposed Hybrid Deep Learning Model

The proposed hybrid model, termed EfficientNetB0-Attention Hybrid, is designed to leverage the powerful feature extraction capabilities of a pre-trained CNN while incorporating a multi-head attention mechanism to focus on the most discriminative features for classification. The architecture is shown schematically in Figure 1.

Figure 1: Architecture of the Proposed EfficientNetB0-Attention Hybrid Model

The architecture consists of the following components:

1. **Input Layer:** Accepts images of size 224x224x3.
2. **Base Model (EfficientNetB0):** A pre-trained CNN on ImageNet, used as a feature extractor. The top classification layers are removed, and the model is frozen initially.
3. **Global Average Pooling (GAP):** Reduces the spatial dimensions of the feature maps, reducing the number of parameters and preventing overfitting.
4. **Multi-Head Attention:** A mechanism that allows the model to focus on different parts of the feature space simultaneously, capturing complex interdependencies.
5. **Fully Connected (Dense) Layers:** A custom classifier with two dense layers (128 and 64 neurons) with ReLU activation and dropout (0.5) for regularization.
6. **Output Layer:** A softmax layer with 7 neurons, outputting the probability distribution over the seven lesion classes.

The model was compiled with the Adam optimizer (learning rate = 0.0001) and categorical cross-entropy loss. Training involved a two-stage process: first, only the custom top layers were trained while the base model remained frozen (10 epochs). Subsequently, the base model was partially unfrozen (fine-tuned) and trained together with the top layers at a lower learning rate (0.00001) for 30 epochs.

2.3. Benchmark Models

To provide a comprehensive comparative analysis, the proposed hybrid model was benchmarked against four established deep learning architectures:

- **Baseline CNN:** A simple CNN with 3 convolutional layers, max-pooling, and 2 dense layers.
- **VGG16:** A deep CNN with 16 layers, known for its simplicity and uniform architecture (Simonyan & Zisserman, 2015).
- **ResNet50:** A 50-layer network that utilizes residual connections to mitigate the vanishing gradient problem, enabling training of very deep networks (He et al., 2016).
- **InceptionV3:** A network that employs inception modules with multiple filter sizes to capture features at different scales (Szegedy et al., 2016).

Each benchmark model was adapted for the 7-class classification task using transfer learning, following a similar two-stage training approach (freezing base layers, then fine-tuning). The same data augmentation, preprocessing, and training configuration (optimizer, loss, epochs) were used for consistency.

2.4. Evaluation Metrics

The models were evaluated on the hold-out test set using standard classification metrics:

- **Accuracy:** The proportion of correctly classified samples.
 - **Precision:** The ratio of true positives to the sum of true and false positives. It indicates the model's ability to avoid false positives.
 - **Recall (Sensitivity):** The ratio of true positives to the sum of true positives and false negatives. It indicates the model's ability to identify all positive samples.
 - **F1-Score:** The harmonic mean of precision and recall, providing a balanced measure.
 - **Area Under the ROC Curve (AUC):** A measure of the model's ability to distinguish between classes, averaged across all classes.
- These metrics were computed for each class individually and then averaged using macro-averaging to account for class imbalance.

3. Results and Discussion


3.1. Performance Comparison

The performance of the proposed hybrid model and the benchmark models on the test set is summarized in Table 1. The results clearly demonstrate the superiority of the proposed EfficientNetB0-Attention Hybrid model across all evaluation metrics.

Table 1. Performance Comparison of Deep Learning Models on the HAM10000 Test Set

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC (%)
Baseline CNN	82.3	81.5	81.8	81.6	91.2
VGG16	87.1	86.4	86.7	86.5	93.8
ResNet50	89.5	88.9	89.1	89.0	94.9
InceptionV3	90.2	89.6	89.9	89.7	95.3
Proposed Hybrid	93.8	93.2	93.5	93.3	97.6

The proposed hybrid model achieved an overall accuracy of 93.8%, outperforming the best benchmark model (InceptionV3 at 90.2%) by a margin of 3.6%. The precision, recall, and F1-score all exceed 93%, indicating a well-balanced model with low rates of both false positives and false negatives. The high AUC of 97.6% further confirms the model's exceptional discriminative capability. These results align with recent studies that highlight the efficacy of hybrid architectures and attention mechanisms in medical image classification (Devarakonda et al., 2024; Reyad et al., 2024). The baseline CNN, while adequate, significantly lagged behind the more sophisticated architectures, underscoring the importance of transfer learning and advanced architectural designs. Figure 2 presents a bar chart visualizing the comparative performance of the five models.

Figure 2: Comparative Performance of Skin Cancer Detection Models

3.2. Class-Wise Performance Analysis

A detailed class-wise analysis is crucial for understanding the model's behavior, especially given the class imbalance in the HAM10000 dataset. Table 2 presents the precision, recall, and F1-score for each of the seven lesion classes for the proposed hybrid model.

Table 2. Class-Wise Performance of the Proposed Hybrid Model


Class	Precision (%)	Recall (%)	F1-Score (%)	Support
akiec	88.5	86.2	87.3	105
bcc	92.1	91.4	91.7	152
bkl	90.3	89.7	90.0	198
df	94.2	91.8	93.0	42
mel	89.6	90.5	90.0	186
nv	96.1	97.2	96.6	728
vasc	95.8	94.3	95.0	92

The model performs exceptionally well on the majority class, nv (melanocytic nevi), with an F1-score of 96.6%. It also shows high performance on vasc (vascular lesions) and df (dermatofibroma), despite these classes having limited support. This suggests that the attention mechanism helps the model focus on the distinctive features of these rare but

clinically important lesions. The classes akiec (actinic keratoses) and bkl (benign keratosis) have slightly lower recall, indicating some confusion between these benign and pre-malignant lesions, which is a known challenge in dermatology due to their visual similarity (Tschandl et al., 2018). The performance on mel (melanoma), the most critical class, is strong (F1-score of 90.0%), highlighting the model's potential as a screening tool.

Figure 3 illustrates the confusion matrix for the proposed hybrid model, providing a visual representation of classification errors.

Figure 3: Confusion Matrix for the Proposed Hybrid Model



The confusion matrix confirms the high diagonal values, with misclassifications primarily occurring between akiec, bkl, and mel, which is clinically understandable. This pattern suggests that further refinement, potentially through the incorporation of clinical metadata (e.g., patient age, lesion location), could further improve performance (Devarakonda et al., 2024).

3.3. Discussion

The findings of this study demonstrate that the proposed hybrid deep learning model, which integrates a pre-trained EfficientNetB0 with a multi-head attention mechanism, achieves state-of-the-art performance on the HAM10000 dataset. The significant improvement over single CNN architectures (VGG16, ResNet50, InceptionV3) can be attributed to several factors:

1. Superior Feature Extraction: EfficientNetB0, with its compound scaling method, provides a more efficient and effective feature representation compared to earlier architectures like VGG16 (Tan & Le, 2019). This allows the model to capture fine-grained details critical for distinguishing between similar skin lesions.

2. Attention Mechanism: The multi-head attention layer enables the model to focus on the most relevant spatial regions of the feature map, effectively highlighting diagnostically important structures like lesion borders, asymmetry, and color variations. This is particularly beneficial for classes with subtle visual differences (Mondol et al., 2023; Singh et al., 2024).

3. Transfer Learning and Fine-Tuning: Leveraging weights pre-trained on ImageNet provides a strong initial feature representation. The two-stage training strategy

(freezing then fine-tuning) allows the model to adapt these general features to the specific nuances of dermatoscopic images without catastrophic forgetting.

The strong performance of the hybrid model aligns with broader trends in medical image analysis, where hybrid architectures and attention mechanisms are increasingly being adopted to enhance model accuracy and interpretability (Helaly et al., 2024; Luo et al., 2024; Sarhan et al., 2024). The high recall for melanoma (90.5%) is particularly promising, as it suggests the model can serve as a sensitive screening tool to reduce missed cases, a critical factor in clinical practice.

However, the study has limitations. While the HAM10000 dataset is comprehensive, it is still a curated dataset, and performance may vary in real-world clinical settings with different imaging devices, lighting conditions, and patient populations. Furthermore, the model's decisions, although analyzed via class-wise metrics, could benefit from more detailed explainability methods (e.g., Grad-CAM) to provide visual explanations to clinicians, thereby increasing trust and adoption (Mehta et al., 2024; Haney & Liang, 2024). Future work will focus on validating the model on external datasets, integrating clinical metadata, and developing a user-friendly interface for deployment in clinical workflows.

4. Conclusion

This study presented the development and validation of an efficient hybrid deep learning model for skin cancer detection. The proposed model, which combines a pre-trained EfficientNetB0 with a multi-head attention mechanism and a custom classifier, achieved superior performance on the HAM10000 dataset, with an overall accuracy of 93.8%, precision of 93.2%, recall of 93.5%, and F1-score of 93.3%. A comprehensive comparative analysis demonstrated that the hybrid model significantly outperforms standard CNN architectures such as VGG16, ResNet50, and InceptionV3.

The key contributions of this research are:

1. A novel hybrid architecture that leverages the efficiency of EfficientNetB0 and the discriminative power of attention mechanisms for accurate skin lesion classification.
2. A rigorous comparative analysis establishing the superiority of the proposed model over established benchmark models.
3. Detailed class-wise performance evaluation, providing insights into the model's strengths and areas for improvement, particularly for challenging lesion types.

The results underscore the transformative potential of advanced deep learning techniques in medical image analysis. The proposed model offers a robust, automated tool for skin cancer screening, with the potential to assist dermatologists in making faster and more accurate diagnoses, thereby improving patient outcomes. By demonstrating a pathway for developing high-performance AI-based diagnostic tools, this research contributes to the broader goal of integrating artificial intelligence into routine clinical practice and advancing personalized healthcare.

References

1. Alwakid, G., Gouda, W., Humayun, M., & Sama, N. (2022). Melanoma detection using deep learning-based classifications. *Healthcare*, *10*, 2481.
2. Badawy, M., Ramadan, N., & Hefny, H. A. (2023). Healthcare predictive analytics using machine learning and deep learning techniques: a survey. *Journal of Electrical Systems and Information Technology*, *10*(1), 40.
3. Devarakonda, N., Murthy, M. V. R., Reddy, R. R. C., & Shree Harsha, P. B. L. (2024). Skin Cancer Detection with Metadata Using Deep Learning Strategies. In Shetty, N. R., Prasad, N. H., Nagaraj, H. C. (Eds.) *Advances in Communication and Applications*.

- ERCICA 2023. Lecture Notes in Electrical Engineering, vol. 1105. Singapore: Springer.
4. Haney, J. M., & Liang, C. J. (2024). A literature review on safety perception and trust during human-robot interaction with autonomous mobile robots that apply to industrial environments. *IIE Transactions on Occupational Ergonomics and Human Factors*, *12*(1-2), 6-27.
 5. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770-778.
 6. Helaly, H. A., Badawy, M., & Haikal, A. Y. (2024). A review of deep learning approaches in clinical and healthcare systems based on medical image analysis. *Multimedia Tools and Applications*, *83*(12), 36039-36080.
 7. Hussain, I., & Nazir, M. B. (2024). Mind matters: exploring AI, machine learning, and deep learning in neurological health. *International Journal of Advanced Engineering Technology and Innovation*, *1*(4), 1-43.
 8. Kumari, J., Kumar, E., & Kumar, D. (2023). A structured analysis to study the role of machine learning and deep learning in the healthcare sector with big data analytics. *Archives of Computational Methods in Engineering*, *30*(6), 3673-3701.
 9. Luo, L., Wang, X., Lin, Y., Ma, X., Tan, A., Chan, R., ... & Chen, H. (2024). Deep learning in breast cancer imaging: A decade of progress and future directions. *IEEE Reviews in Biomedical Engineering*.
 10. Maduranga, M. W. P., & Nandasena, D. (2022). Mobile-based skin disease diagnosis system using convolutional neural networks (CNN). *International Journal of Image, Graphics and Signal Processing*.
 11. Mehta, R., Shui, C., & Arbel, T. (2024). Evaluating the fairness of deep learning uncertainty estimates in medical image analysis. *Proceedings of Medical Imaging with Deep Learning*, 1453-1492. PMLR.
 12. Mondol, R. K., Millar, E. K., Graham, P. H., Browne, L., Sowmya, A., & Meijering, E. (2023). hist2rna: an efficient deep learning architecture to predict gene expression from breast cancer histopathology images. *Cancers*, *15*(9), 2569.
 13. Reyad, M., Sarhan, A. M., & Arafa, M. (2024). Architecture optimization for hybrid deep residual networks in liver tumor segmentation using a GA. *International Journal of Computational Intelligence Systems*, *17*(1), 1-22.
 14. Sarhan, A. M., Gobara, M., Yasser, S., Elsayed, Z., Sherif, G., Moataz, N., ... & Ali, H. A. (2024). Knee osteoporosis diagnosis based on deep learning. *International Journal of Computational Intelligence Systems*, *17*(1), 241.
 15. Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *3rd International Conference on Learning Representations*.
 16. Singh, U., Abhishek, K., & Azad, H. K. (2024). A survey of cutting-edge multimodal sentiment analysis. *ACM Computing Surveys*, *56*(9), 1-38.
 17. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2818-2826.
 18. Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning*, 6105-6114.
 19. Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, *5*, 180161.

20. World Cancer Research Fund. (n.d.). Skin cancer statistics. Retrieved from <https://www.wcrf.org/cancer-trends/skin-cancer-statistics/>
21. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Sebastian Thrun (2017). Dermatologist-Level Classification Of Skin Cancer With Deep Neural Networks. *Nature*, 542(7639), 115–118.
22. Haenssle, H. A., Fink, C., Schneiderbauer, R., Toberer, F., Buhl, T., & Claus Garbe (2018), “Man Against Machine: Diagnostic Performance Of A Deep Learning Convolutional Neural Network For Dermoscopic Melanoma Recognition” *Annals of Oncology*, 29(8), 1836–1842.
23. Codella, N. C. F., Nguyen, Q. B., Pankanti, S., Gutman, D. A., Helba, B., Halpern, A. C., & John R. Smith (2017), “Deep Learning Ensembles For Melanoma Recognition In Dermoscopy Images” *IBM Journal of Research and Development*, 61(4/5),
24. Tschandl, P., Rosendahl, C., & Harald Kittler (2018), “The HAM10000 dataset: A Large Collection Of Multi-Source Dermatoscopic Images Of Common Pigmented Skin Lesions,” *Scientific Data*, 5, 180161
25. Brinker, T. J., Hekler, A., Enk, A. H., Berking, C., Haferkamp, S., & Jochen S. Utikal (2019), Deep Neural Networks Are Superior To Dermatologists In Melanoma Image Classification, *European Journal of Cancer*, 119, 11–17.