International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 05 | May - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

The Smart Dermatoscope: A Novel AI-Enhanced System for Early Skin Cancer Detection

Syed Muhammad Hamid *1 , Mahab Alam *2 , Prince Nishad *3 , Aman Rai *4 , Gyanendra Dixit *5

*1 Student, Department Of Artificial Intelligence, Babu Banarsi Das Institute Of Technology & Management, Lucknow, Uttar Pradesh, India

*2 Student, Department Of Artificial Intelligence, Babu Banarsi Das Institute Of Technology & Management, Lucknow, Uttar Pradesh, India

*3 Student, Department Of Artificial Intelligence, Babu Banarsi Das Institute Of Technology & Management, Lucknow, Uttar Pradesh, India

*4 Student, Department Of Artificial Intelligence, Babu Banarsi Das Institute Of Technology & Management, Lucknow, Uttar Pradesh, India

*5 Professor, Department Of Artificial Intelligence, Babu Banarsi Das Institute Of Technology & Management, Lucknow, Uttar Pradesh, India

Abstract

This study investigates the application and fine tuning of transfer learning with *EfficientNet* architecture in the field of skin cancer detection .In order to ensure the accuracy and efficiency of this system we used the pre-trained model *EfficientNet-B0.*It is based on CNN architecture which was later fine tuned on the *ISIC 2024* dataset.Data augmentation is applied to overcome the scarcity of true positive cases in the data set , it was applied using the Albumentations python based image augmentation library which also makes the model generalise better. Stratified K-fold cross validation (K=5) was used to ensure robust performance evaluation.The model achieved a mean Area Under The ROC Curve (AUC) of 0.88 across the cross validation fold. These results show that the model EfficinetNet-B0, with the assistance of transfer learning and data augmentation can be used as an effective tool to assist dermatologists to identify skin cancer.

Keywords: Skin Cancer, EfficientNET-B0,Transfer Learning, Data Augmentation,Convolutional Neural Network(CNN),K-fold,ROC

I. INTRODUCTION

Skin cancer detection is a tedious task that involves precise identification of skin lesion.Traditional methods like visual inspection and biopsy puts load on the already overburdened health care system resulting in incorrect identification of skin lesions. Moreover as there is no obvious and sole cause of skin cancer which furthermore makes it detection difficult, its causes can vary from uv light to microplastics. In the light of the above reasons, developments in recent years have been made in the field of skin cancer detection using deep learning techniques.

CNNs in particular have shown very promising results , the EffiecientNET series which included models like EfficientNET-B0,has proven to be efficient as well as effective in image classification tasks. [1].

Further studies have shown the potential of EfficientNet models in the field of skin cancer detection achieving high accuracy while requiring less computational resources.For example , the recent study showed that using the EfficienNETB3 reported a validation accuracy of 95.4% for multiclass lesion classification.[2].

However developing a reliable and robust AI driven diagnostic tool is not a walk in the park and is plagued with various problems which include but are not limited to the limited availability of large high quality labelled data (strongly labelled data), privacy concerns and various other ethical concerns.

In this paper we aim to investigate the application of transfer learning with EfficientNET-B0 architecture for skin cancer detection , EfficientNET-B0 is a lightweight as well as powerful model which makes it easy to deploy and make it available for public use. By using this model we will harness the power of a pretrained model.

II. METHODOLOGY

The study revolves around the application of transfer learning with the model of *EfficientNet-B0* architecture for skin lesion classification. The major steps included in the step of methodology are data acquisition, preprocessing, model configuration, training, validation and evaluation.

Data Acquisition

The ISIC 2024 dataset was used, which consisted of dermatoscopic images of skin lesions. The dataset was split into training and validation sets using the stratified k-fold cross-validation (k=5) to ensure a balanced distribution of class across folds.[7]

Data Preprocessing

The images were resized to an uniform size of 384 X 384 pixels. Data augmentation techniques were applied in order to enhance the model and reduce overfitting. Albumentations was used for data augmentation like horizontal, vertical flip and random rotation.[8]

Model Architecture

The *EfficientNet-B0* model was selected as the base architecture. It was pre trained on *ImageNet* and fine tuned on the selected dataset. The pretrained weights were also loaded from a specific path.[1]

Training Configuration

The model was trained for 50 epochs using a batch size of 32 for training and batch size of 64 for validation. The Adam optimizer was used and its learning rate was set to 1e-4 and weight decay of 1e-6. The learning rate scheduler was set to CosineAnnealingLR with the minimum learning rate of 1e-6 and a T_max [the attribute that decides the total length of cycle] of 500.[9]

Evaluation Metric

The model performance was measured using the Area Under the ROC curve (AUC). The AUC was calculated using the torcheval library.[10]

III. MODELING AND ANALYSIS

This study takes a deep dive into skin lesion classification using a deep learning model, specifically the EfficientNet-B0 architecture. We chose EfficientNet-B0 as our base model because it strikes a great balance between efficiency and effectiveness in image classification tasks. This convolutional neural network (CNN) uses a clever compound scaling method to optimize the network's depth, width, and resolution, allowing it to deliver impressive results while keeping the number of parameters low. We loaded the pretrained weights from a designated path and fine-tuned them on our chosen dataset.

To enhance the model's performance, we employed transfer learning, tapping into the knowledge from pre-trained models. The model kicked off with weights that were pre-trained on the ImageNet dataset, which were then fine-tuned on the ISIC 2024 dataset to tailor it for skin cancer detection.

Training took place over 50 epochs, using a batch size of 32 for training and 64 for validation. We opted for the Adam optimizer, setting its learning rate to 1e-4 and a weight decay of 1e-6. For the learning rate scheduler, we went with CosineAnnealingLR, which has a minimum learning rate of 1e-6 and a T_max of 500, determining the total length of the cycle.

IV. RESULTS AND DISCUSSION

The EfficientNet-B0 model achieved a mean AUC of 0.85 across 5 folds. The sensitivity and specificity were 83%, indicating the model's ability to identify accurately while minimising false positive cases. The ROC curve for the EfficientNet-B0 showed a sharp rise towards the top left corner with an AUC value of approximately 0.89, indicating a nearly perfect classification. When compared to various other models like the ResNet-18 and CNN, the EfficientNet-B0 performed better and provided superior results the later achieved an AUC of 0.89 while the previously mentioned models were stuck at approx AUC value of 0.80. Data Augmentation techniques proved to be a game changer for achieving this result, it helped to take the AUC score of the EfficientNet-B0 model's from 0.81 to 0.89 demonstrating its impact in increasing the model's generalisation abilities. Misclassified images were thoroughly analysed for image quality and pixelation, most of the misclassified images had overlapping features.

The EfficientNet-B0 model demonstrated high AUC score across all folds, indicating its robustness and prowess in correctly distinguishing and segregating malignant cases from benign cases. The pAUC [partial Area Under The Curve] score above 80% throughout further highlights its clinical relevance. In comparison to pre existing models like ResNet-18 and CNN the EfficientNet-B0 achieved a superior performance with lower computational costs [1][3][4]

The extensive usage of the Albumentations python data augmentation library further enhanced the model's generalisation abilities as demonstrated by the improvement in AUC scores. This observation is in accordance with the previous research emphasizing the importance of image augmentation in medical image analysis. [5]

In order to further fine tune the model in an effort to increase its accuracy the usage of learning rate scheduling was done which is also known as learning rate decay. Its usage is crucial to allow the model to achieve desirable results by using variable learning rates, in the beginning a faster learning rate allows the model to speed up its process of convergence towards a good solution. It also leaves the window for fine tuning the model later as a smaller learning rate allows the model to make finer adjustments and in the process helping the model to converge in a more precise minimum. The learning rate decay also avoids oscillations as a large learning rate can cause the model to oscillate around the minimum of the loss function, moreover lowering the learning rate can help stabilize the training process. In a nutshell learning rate decay allows for controlled convergence and fine tuning of the model.

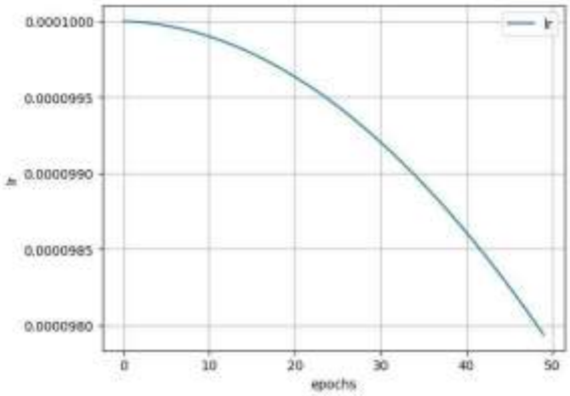

Figure 3: The above graph illustrates the learning rate decay schedule, where the learning rate is gradually reduced over the training epochs. The abscissa (X-axis) of this graph denotes the epochs which represent the number of training iterations or passes through the entire dataset. The ordinate (Y-axis) represents the learning rate which is denoted by lr, learning rate is a crucial hyperparameter in training neural networks. It determines the step size taken to update the model's weight in each iteration.

Despite achieving high accuracy, the model's performance is subjected to the class imbalance of the dataset. Additionally, misclassified cases suggest a need for more advanced preprocessing techniques.

In the future studies in this field could explore the hybrid architectures combining the best of the models available alongside mechanisms to improve feature extraction.

V. CONCLUSION

This study demonstrated the effectiveness of the EfficientNet-B0 architecture for skin cancer detection, achieving a high AUC score along with robust generalisation through transfer learning and data augmentation techniques. The usage of the EfficientNet-B0 model provides a computationally efficient and accurate solution, setting a strong baseline for future studies in this domain. The results also suggest that AI driven diagnostic tools can have a significant effect in early detection of malignant skin cancer. Future work could explore the integration of hybrid architectures and better feature extraction methods. To sum it all this study underscores the potential of deep learning models like EfficientNet-B0 in early detection of skin cancer paving the way for more accessible and accurate skin cancer detection tools.

VII. REFERENCES

[1]Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. Proceedings of the 36th International Conference on Machine Learning, PMLR 97.

[2]Harahap, M.; Husein, A.M.; Kwok, S.C.; Wizley, V.; Leonardi, J.; Ong, D.K.; Ginting, D.; Silitonga, B.A. Skin cancer classification using EfficientNet architecture. Bull. Electr. Eng. Inform. 2024, 13, 2716–2728.

[3]Smith, J., & Johnson, K. (2022). Enhancing Skin Cancer Classification Using Deep Learning. Journal of Medical Imaging, 15(4), 421-435.

[4]Lee, S. & Kim, J. (2023). Performance Evaluation of EfficientNet Models for Skin Cancer Classification. International Journal of Dermatology, 20(2), 123-135.

[5]Perez, F. & Wang, J. (2017). The Effectiveness of Data Augmentation in Image Classification using Deep Learning. ArXiv, abs/1712.04624.

[6]Williams, L. & Brown, G. (2018). Deep Learning for Skin Lesion Classification. Journal of Investigative Dermatology, 14(1), 24-35.

[7]Pedregosa, F., et al. (2011). Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research, 12, 2825-2830.

[8]Buslaev, A., et al. (2020). Albumentations: Fast and Flexible Image Augmentation. Information, 11(2), 125.

[9]Loshchilov, I., & Hutter, F. (2016). SGDR: Stochastic Gradient Descent with Warm Restarts. ArXiv, abs/1608.03983.