



Enhancing Melanoma Skin Cancer Diagnosis Through Transfer Learning: An EfficientNetB0 Approach



Rashmi Ashtagi^{1*}, Pramila Vasanttrao Kharat², Vinaya Sarmalkar³, Sridevi Hosmani⁴, Abhijeet R. Patil⁵, Afsha Imran Akkalkot⁶, Adithya Padthe⁷

¹ Computer Engineering Department, Vishwakarma Institute of Technology, 411037 Pune, India

² Department of Information & Communication Technology, MGM University, 431001 Chatrapati Sambhaji Nagar, India

³ Computer Science and Engineering Department, Jain College of Engineering and Research, 590008 Belagavi, India

⁴ Department of Artificial Intelligence and Machine Learning, Faculty of Engineering and Technology (Exclusively for Women), Sharnbasva University, 585103 Kalaburagi, India

⁵ RPD PU College of Arts and Commerce, 590006 Belagavi, India

⁶ Computer Engineering Department, Savitribai Phule Pune University, 411007 Pune, India

⁷ Department of Information Technology, University of the Columbians, Williamsburg, 40769 Kentucky, USA

* Correspondence: Rashmi Ashtagi (rashmi.ashtagi@vit.edu)

Received: 01-02-2024

Revised: 02-25-2024

Accepted: 03-02-2024

Citation: R. Ashtagi, P. V. Kharat, V. Sarmalkar, S. Hosmani, A. R. Patil, A. I. Akkalkot, and A. Padthe, "Enhancing melanoma skin cancer diagnosis through transfer learning: An EfficientNetB0 approach," *Acadlore Trans. Mach. Learn.*, vol. 3, no. 1, pp. 57–69, 2024. <https://doi.org/10.56578/ataiml030105>.



© 2024 by the author(s). Published by Acadlore Publishing Services Limited, Hong Kong. This article is available for free download and can be reused and cited, provided that the original published version is credited, under the CC BY 4.0 license.

Abstract: Skin cancer, a significant health concern globally, necessitates innovative strategies for its early detection and classification. In this context, a novel methodology employing the state-of-the-art EfficientNetB0 deep learning architecture has been developed, aiming to augment the accuracy and efficiency of skin cancer diagnoses. This approach focuses on automating the classification of skin lesions, addressing the challenges posed by their complex structures and the subjective nature of conventional diagnostic methods. Through the adoption of advanced training techniques, including adaptive learning rates and Rectified Adam (RAdam) optimization, a robust model for skin cancer classification has been constructed. The findings underscore the model's capability to achieve convergence during training, illustrating its potential to transform dermatological diagnostics significantly. This research contributes to the broader fields of medical imaging and artificial intelligence (AI), underscoring the efficacy of deep learning in enhancing diagnostic processes. Future endeavors will explore the realms of explainable AI (XAI), collaboration with medical professionals, and adaptation of the model for telemedicine, ensuring its continued relevance and applicability in the dynamic landscape of skin cancer diagnosis.

Keywords: Deep learning; Melanoma; Skin cancer; EfficientNetB0; Rectified Adam

1 Introduction

One type of cancer that is common and starts in the skin cells is skin cancer. It can be distinguished into three primary types: melanoma, squamous cell carcinoma, and basal cell carcinoma. While melanoma is known to have a greater propensity to spread, less aggressive cancers, such as squamous cell carcinoma and basal cell carcinoma, are more common. Skin cancer is mostly caused by exposure to ultraviolet (UV) radiation, which comes from man-made or solar sources. Because of its rising incidence, skin cancer is turning into a serious worldwide health issue. Therefore, developing effective methods for early detection and classification is crucial to improving patient outcomes.

Success in therapy and a better prognosis for skin cancer patients depend heavily on early detection. Skin cancer is frequently quite treatable when found early on, with a variety of intervention options available, including excision and more specialised therapy. A postponed diagnosis might cause the illness to worsen, making treatment more difficult and decreasing the likelihood of a full recovery. Furthermore, early detection reduces the need for intrusive and lengthy procedures, enhancing the quality of life for those affected. Since skin lesions are visible, utilising cutting-edge technologies for precise and prompt diagnosis becomes crucial in the fight against skin cancer.

The technology that is being described, which uses the EfficientNetB0 architecture to classify skin cancer, fills a significant void in the current early detection approaches. Conventional diagnostic methods frequently depend on dermatologists' visual examinations, which can be laborious and subjective. One interesting approach to automating skin lesion classification is including deep learning models like EfficientNetB0. This research project seeks to advance the creation of a reliable, effective, and scalable instrument for early skin cancer detection by utilising AI. The application of sophisticated neural network designs improves the diagnostic procedure's precision and dependability, which has the potential to completely transform the way skin cancer is identified and treated in clinical settings.

1.1 Problem Statement

A number of obstacles stand in the way of accurately classifying skin cancer, chief among them being the complex and varied character of skin lesions. Dermatologists frequently struggle to distinguish between benign and malignant lesions, as well as between various skin cancer subtypes. Differences in lesion dimensions, forms, hues, and textures add to the diagnostic challenge. Conventional diagnostic techniques mostly rely on visual assessment by the patient, which might introduce subjectivity, variability, and inconsistent results. Furthermore, the increasing number of skin cancer cases strains healthcare systems, making the creation of precise and automated categorisation models necessary to assist dermatologists in their work and boost diagnostic effectiveness.

The study effort aims to construct an efficient skin cancer classification model, which is crucial in tackling the previously described issues. Using deep learning models, specifically the EfficientNetB0 architecture, provides a way to improve skin cancer diagnosis accuracy and dependability. The model's ability to distinguish between benign and malignant lesions is aided by its ability to extract complex patterns and features from huge datasets. Dermatologists can greatly benefit from the use of an efficient model as a decision-support tool, which helps them diagnose patients more quickly and accurately. Additionally, implementing an automated categorisation system may simplify the diagnostic process, lessen subjectivity, and enable early detection, all of which could enhance patient outcomes and optimise the use of healthcare resources.

1.2 Objectives

The primary objective of this research endeavor is the development of a robust skin cancer classification model utilizing the EfficientNetB0 architecture. The architecture is a cutting-edge deep learning model renowned for its scalability and efficiency. Training the model on various datasets, including benign and malignant skin lesions, is a step in the development process. Pre-trained weights are utilised in conjunction with transfer learning approaches to enable the model to capture complex patterns and attributes pertinent to the classification of skin cancer. Developing a model that can reliably and automatically differentiate between various kinds of skin lesions is the goal in order to contribute to an automated diagnostic tool.

In the subsequent phase, the model's performance is rigorously evaluated. This evaluation involves the analysis of the model using separate validation and test datasets, focusing on accuracy, precision, recall, and other relevant metrics. Validating the model's capacity to generalise to new data and generate trustworthy predictions is the goal of the evaluation procedure. Additionally, the study investigates how the training and convergence of the model are affected by the addition of sophisticated optimisation strategies such as learning rate reduction and RAdam. Through a thorough assessment of the model's performance, the study provides valuable information about the efficacy of deep learning techniques in classifying skin cancer. It sets a standard for further research in this area.

The use of the cutting-edge deep learning model, EfficientNetB0 architecture, to overcome the difficulties in skin cancer diagnosis is what makes this work novel. The efficiency and efficacy of the model are enhanced by the combination of methods used in this study, which includes adaptive learning rates and RAdam optimization. The paper presents the suggested method as a cutting-edge method for automating the classification of skin lesions and highlights the significance of early identification in the fight against skin cancer. The study presents a scalable and precise tool for dermatological diagnostics by utilizing EfficientNetB0. Together with improvements in training methods, the use of deep learning in this particular domain represents a novel addition to the fields of AI and medical imaging.

2 Literature Review

In their discussion of the rising frequency of skin tumours worldwide, Udriștoiu et al. [1] emphasised the need for early identification, especially for melanoma and non-melanoma skin malignancies. They suggested a Convolutional Neural Network (CNN) architecture that uses just image pixels and diagnosis labels as inputs to categorise dermoscopic images of skin lesions without the need for manually created features. In order to overcome imbalances, the study used the HAM10000 dataset and included data augmentation techniques. With a final accuracy of 93.6%, a sensitivity of 95.9%, a specificity of 98.3%, and an area under the curve (AUC) of 0.97, the constructed 4-CNN model showed encouraging results. The authors emphasised the potential clinical utility of their deep learning

system to assist dermatologists in the diagnostic process while discussing the constraints, such as a known testing dataset and imbalances in the dataset.

Panthakkan et al. [2] introduced concatenated X-R50, a novel and efficient deep-learning model designed for accurate skin cancer prediction. Early diagnosis of skin cancer is crucial for effective treatment, given its prevalence and potential lethality. The authors combined the strengths of two robust deep neural networks, ResNet50 and Xception, to create the X-R50 model. In the classification of skin cancers into various groups, such as actinic keratoses, intraepithelial carcinoma, vascular lesions, benign keratosis-like lesions, melanoma, melanocytic nevi, dermatofibroma, and basal cell carcinoma, the model exhibited an impressive 97.8% prediction accuracy. The study utilised the Human Against Machine (HAM10000) dataset, consisting of 10,500 skin photo views, and employed the sliding window technique for training and testing. Notably, the proposed X-R50 model outperformed deep CNN and other contemporary transfer learning techniques in skin cancer prediction.

The authors underlined the value of early identification and categorisation of skin cancer while stressing the possible clinical uses of their X-R50 model to support dermatologists and other healthcare professionals in the diagnostic procedure. By resolving the shortcomings of existing dermoscopic approaches, the research contributes a computationally efficient and accurate strategy for quick skin cancer screening.

An enhanced transfer learning method for classifying different forms of skin cancer was presented by Mane et al. [3], with an emphasis on the extremely lethal melanoma. To reduce death rates, the study discusses the vital necessity of early and precise detection of skin lesions. The suggested framework successfully classifies skin lesions into eight categories, namely, melanoma, benign keratosis, basal cell carcinoma, actinic keratosis, melanocytic nevus, vascular lesion, dermatofibroma, and squamous cell carcinoma, by utilising transfer learning and a pre-trained model called MobileNet. The scientists used the 25,331 dermoscopic images from the ISIC 2019 challenge dataset for their experiments. An image is categorised as unknown if it does not fit into one of the eight categories. With a high accuracy of 83.0%, the suggested system proved to be useful in helping dermatologists provide proper therapy. The research advances the field by providing an easy-to-use and effective technique for early skin cancer screening through image-based exams. It may increase the likelihood of finding skin cancer early on.

Hoang et al. [4] addressed the difficulties in precisely differentiating between various skin lesion types, particularly the crucial detection of melanoma for early intervention, by introducing a novel approach to multiclass skin lesion categorisation. The authors suggested combining a two-dimensional wide-ShuffleNet for classification with a novel segmentation technique that uses entropy-based weighting (EW) and first-order cumulative moment (FCM) for skin image segmentation. The study aimed to improve skin lesion classification's effectiveness and accuracy, which is essential for better patient outcomes. The suggested technique showed better performance than cutting-edge methods when tested on two sizable datasets, HAM10000 and ISIC2019. The significant death rate linked to melanoma underscored the need for automatic classification methods to support doctors and dermatologists in prompt and precise diagnosis, according to the study. To address the practical constraints of implementing such systems in real-world healthcare scenarios, the suggested framework was also made to be lightweight and suited for integration into small systems, such as mobile healthcare platforms.

Tiwari et al. [5] presented a unique deep transfer learning-based method for automated breast cancer screening. The Visual Geometry Group (VGG) 16 model, which was trained on a dataset of both static and dynamic breast thermal pictures, is used in this work. Notably, in order to improve the precision of thermal temperature mapping, the authors present a novel idea for multi-view breast thermal pictures, concatenating traditional left, right, and frontal views. Comparative investigation using ResNet50V2, InceptionV3, and VGG 19 models demonstrates that the VGG 16-based system achieves improved training, validation, and testing accuracy. The suggested technique successfully detects breast cancer at an early stage, as evidenced by its astounding 99% testing accuracy on dynamic breast thermal imaging. Advances in computer-aided diagnosis tools may be superior to conventional mammography methods in the setting of thermal imaging.

The growing importance of automated techniques in the diagnosis of skin cancer, especially the use of deep learning algorithms for picture classification, was discussed by Brinker et al. [6]. The systematic review aimed to present a thorough summary of research using CNNs to classify skin cancer, emphasising patterns, approaches, and results. The authors sought to evaluate the present status of CNN-based skin cancer categorisation, pointing out advantages, disadvantages, and possible areas for development by combining results from diverse research projects. This study sheds light on the developments, difficulties, and opportunities related to integrating CNNs in the classification of skin cancer, which advances the understanding of the changing field of computer-aided diagnosis in dermatology.

Refianti et al. [7] acknowledged the importance of an accurate and timely diagnosis in improving patient outcomes and concentrated on using CNNs to categorise melanoma skin cancer. Their research focused on the potential of using CNNs, a type of deep learning, for automated melanoma classification from skin imaging data. To shed light on the effectiveness of deep learning methods for dermatological applications, the study investigated CNNs' capability for melanoma lesion classification. The study may have addressed the increasing demand for automated technologies

in dermatology to help doctors diagnose melanoma early and accurately, improving patient outcomes.

Table 1. Comparative analysis of various models in references

Study	Model Architecture	Dataset	Accuracy	Challenges
Udriştoiu et al. [1]	CNN	HAM10000	93.6%	Dataset imbalances
Panthakkan et al. [2]	X-R50 (ResNet50+Xception)	HAM10000	97.8%	Early skin cancer prediction
Mane et al. [3]	MobileNet	ISIC 2019	83.0%	Classifying various skin cancer types
Hoang et al. [4]	Two-dimensional wide-ShuffleNet	HAM10000, ISIC2019	Improved performance in multiclass classification	Differentiating between skin lesion types
Tiwari et al. [5]	VGG 16	Static and dynamic breast thermal images	99%	Automated breast cancer screening
Brinker et al. [6]	CNNs (Review)	Multiple datasets	Not applicable (Review)	Evaluating CNN-based skin cancer classification
Refianti et al. [7]	CNNs	Images	78%	Melanoma lesion classification
Zhang et al. [8]	Optimal CNN	Multiple	79.3%	Improving CNN performance for skin cancer diagnosis
Garg et al. [9]	CNN	Not specified	Not specified	Developing a decision support system for skin cancer diagnosis
Patil and Bellary [10]	Deep CNNs	Not specified	Not specified	Improving multi-class skin cancer classification

Zhang et al. [8] proposed an optimal CNN approach to address the crucial problem of skin cancer diagnosis. The study aimed to improve CNN's performance so that it could diagnose skin cancer more accurately. Zhang et al. [8] sought to enhance the neural network design in order to enhance the model's capacity to differentiate between different forms of skin cancer. The work explicitly modifies CNNs for accurate and efficient skin cancer diagnosis, which advances the field of AI in medicine. To produce better diagnostic results, the optimisation process may include adjusting architectural features, fine-tuning model parameters, or making other improvements. In an attempt to give physicians reliable instruments for timely and accurate skin cancer diagnosis, this research may contribute significant new information to the continuing attempts to use AI for dermatological applications.

Garg et al. [9] contributed to the development of a decision support system dedicated to the identification and categorisation of skin cancer, with a primary focus on improving diagnostic precision. The authors implemented CNN as a key component to create an intelligent system capable of enhancing skin cancer diagnosis. This study aligns with broader efforts to integrate advanced computational techniques, including CNNs, into healthcare decision-support systems. The proposed model likely leveraged deep learning methods for image processing, facilitating the automatic identification and categorisation of various types of skin cancer. The emphasis of Garg et al. [9] on incorporating CNNs in dermatological diagnosis underscores the potential of providing physicians with a valuable tool for early detection and categorisation of skin cancer, ultimately contributing to improved patient outcomes and healthcare decision-making.

Patil and Bellary [10] used a CNN method and a unique Similarity Measure for Text Processing (SMTP) as the loss function, and proposed two classification systems. Based on tumour thickness, the first approach divides melanoma into stages 1 and 2, and the second system further subdivides the classification into stages 1, 2, or 3. The suggested technique uses dermoscopic pictures to categorise cancer stages non-invasively, with a focus on tumour thickness as a critical diagnostic metric. The authors address current techniques for determining tumour thickness, such as the Breslow index and the Clark scale, and emphasise the significance of early identification and discovery in the treatment of melanoma.

Patil and Bellary [11, 12] concentrated on the application of deep learning techniques, particularly deep CNNs, for the multi-class categorisation of skin cancer. The study harnessed the capabilities of state-of-the-art neural network architectures to address the pressing need for accurate and efficient skin cancer detection. With a focus on enhancing

the classification accuracy of skin cancers, the authors employed deep CNNs, contributing to the ongoing efforts to leverage AI for dermatological applications. Presumably, the suggested model underwent training on diverse skin lesion datasets to ensure robust classification performance. This research contributes to the expanding realm of computer-aided diagnostic tools in dermatology and holds potential implications for improving patient outcomes and informing clinical decision-making. Padthe et al. [13] proposed a federated learning framework for analyzing large healthcare image datasets without sharing raw data. It achieves state-of-the-art accuracy in pneumonia classification on chest X-ray images while preserving patient privacy. The framework is scalable, efficient, and achieves 98.87% accuracy, offering promise for revolutionizing healthcare image analysis.

Ashtagi et al. [14] proposed a novel method to predict stages of melanoma using a combination of deep learning (DL) and machine learning (ML) algorithms. By leveraging the strengths of DL in simplifying training and ML in improving accuracy, the hybrid model (MLP-RF) achieves a high accuracy of 97.42%. This approach offers a promising solution for accurate melanoma diagnosis and staging, showcasing the potential of integrating DL and ML techniques in cancer detection. Chaturvedi et al. [15] presented an automated system for accurately diagnosing multi-class skin cancer, which is challenging due to similarities in appearance among different types. By fine-tuning on the HAM10000 dataset, the study evaluates several pre-trained CNNs and ensemble models. The results show a maximum accuracy of 93.20% for individual models and 92.83% for ensemble models. ResNeXt101 is recommended for its optimized architecture and superior accuracy. This research contributes to improving diagnostic accuracy and patient outcomes in skin cancer management. Mane et al. [3] introduced a framework for precise classification of skin lesions, crucial for early detection and treatment of melanoma and other skin cancers. Using transfer learning with MobileNet architecture and the ISIC 2019 dataset, the system accurately categorizes eight types of lesions. This aids dermatologists in making accurate diagnoses and treatment decisions, potentially reducing mortality rates associated with skin cancer. Table 1 shows the comparative analysis of models in the literature review.

It is clear from reading the literature review that there are a variety of difficulties in applying deep learning models for the categorisation of skin cancer. Udriştoiu et al. [1] highlighted the importance of early identification, whose CNN architecture acknowledged imbalances in the dataset while obtaining a 93.6% accuracy rate. By integrating ResNet50 with Xception, Panthakkan et al. [2] successfully tackled the problem of skin cancer prediction, attaining a high accuracy of 97.8%. Mane et al. [3] used MobileNet to classify different forms of skin cancer and achieved an accuracy of 83.0%. Hoang et al. [4] highlighted the difficulties in distinguishing between different types of skin lesions while presenting a novel method for multiclass categorisation. Every study has its own set of problems, including unbalanced datasets, various lesion types, and the requirement for lightweight models for real-world application. Through the use of EfficientNetB0, meticulous data preprocessing, and the application of cutting-edge methods like data augmentation and transfer learning, this study solves comparable difficulties. This skin cancer classification system's robustness is a result of its thorough evaluation of its performance, attention to class imbalances, and realistic implementation restrictions.

This study assesses the suggested skin cancer classification model based on EfficientNetB0 against other cutting-edge models in the literature by direct comparison with current approaches. Even though the paper admits that the model's 87% accuracy might not be as accurate as some of the best accuracies attained by other models, it offers a fair analysis by highlighting the accuracy and computational efficiency attained by EfficientNetB0. The suggested model is benchmarked against current skin cancer classification transfer learning implementations in this work, demonstrating competitive performance and adding to the continuing discourse in the field. The comparison sheds light on the unique approach's advantages and possible shortcomings by placing it within the framework of existing techniques.

3 Methodology

3.1 Data Collection

The research project's dataset is essential to the model's training and generalisation abilities. It includes a wide range of pictures of skin lesions, including both benign and malignant instances. A thorough analysis of the dataset provides details on its size, image quality, and distribution of classes. Furthermore, the dataset's source indicates whether it was specially gathered for this study or was selected from existing medical databases. Because of the dataset's diversity, the model is exposed to a greater variety of skin lesions, improving its ability to distinguish between different conditions and increasing its applicability in the real world. Figure 1 shows the image dataset used for the model.

Images of skin lesions classified into benign and malignant groups for the purpose of classifying skin cancer make up the dataset used in this study. The preprocessed and loaded images are from the training and test sets of the "Skin Cancer: Malignant vs. Benign" dataset. In the training set of the dataset, there are 1,440 benign and 1,197 malignant photos; in the test set, there are 360 benign and 300 malignant images. Preprocessing methods for the photos include applying segmentation techniques and scaling them to a resolution of 224×224 pixels. To further increase the diversity of the dataset, data augmentation techniques, including flipping, rotating, zooming, and brightness modifications, are used. The dataset offers a representative sample to properly train the skin cancer classification model because it

displays variability in the characteristics of the skin lesions. The dataset’s “Skin Cancer: Malignant vs. Benign” source is identified, which enhances the study’s reproducibility and transparency.




Figure 1. Cancer dataset images

3.2 Data Preprocessing Steps

The preprocessing of data for the model encompasses several critical steps. Initially, images of skin lesions are resized to uniform dimensions appropriate for the EfficientNetB0 architecture. By ensuring consistency in the input dimensions, resizing makes it possible for the model to handle images reliably. For effective training and the best possible use of computational resources, this stage is crucial.

Data augmentation is subsequently employed to artificially augment the dataset’s diversity. Various augmentation techniques, including rotation, flipping both horizontally and vertically, and zooming, are utilised to generate variations of the original images. This approach aids the model in better understanding invariant characteristics and developing resilience to variations in the appearance of skin lesions.

Normalization of pixel values across all images represents another fundamental preprocessing step. Pixel intensities are scaled to a standard range, usually $[0, 1]$ or $[-1, 1]$. Normalisation guarantees that the model is less sensitive to changes in pixel values and enhances convergence during training. Using ImageNet’s mean and standard deviation values may be necessary for normalising EfficientNetB0, which is pre-trained on ImageNet.

Addressing class imbalances within the dataset is crucial to preventing biased model learning. Certain classes, whether benign or malignant, which may be underrepresented, are corrected. To achieve a balanced distribution, strategies like undersampling the dominant class and oversampling the minority class are used. This keeps the model from being biased in favour of the dominant class by ensuring that it learns from a representative sample of instances for each class.

Finally, the dataset is divided into distinct sets for training, validation, and testing, aiming to evaluate the model's performance. The validation set is crucial for fine-tuning hyperparameters and preventing overfitting, ensuring the model's effectiveness on unseen data. The test set evaluates the model's generalisation to new, unseen data, while the training set is utilised to train the model. This division ensures an objective evaluation and aids in identifying potential issues such as overfitting or underfitting.

The dataset is modified to offer the best learning environment possible for the EfficientNetB0-based skin cancer classification model by carefully carrying out these preprocessing processes.

3.3 Training Strategy

The neural network models were trained by carefully dividing the dataset into three subsets: test, validation, and training sets. The training set, which accounted for 70% of the total dataset, was the main source from which the model's parameters were updated. The choice of optimiser and loss function in the training strategy for the skin cancer classification project plays a crucial role in the model's learning process. The selected loss function is categorical cross-entropy, deemed suitable for multi-class classification tasks, such as distinguishing between benign and malignant skin lesions. The optimiser RAdam is employed, featuring an adaptable learning rate mechanism aimed at enhancing model stability and convergence speed. The learning rate is initiated at $1e-3$, with a warm-up phase encompassing the first 15% of training epochs, and it is constrained to a minimum value of $1e-7$. This configuration is designed to optimise the learning process for effective skin cancer classification.

A learning rate reducer is included to handle dynamic learning rate variations and better enhance the training process. After a predetermined number of training epochs (with patience set to 5), if there is no progress in validation accuracy, the ReduceLROnPlateau callback lowers the learning rate by a factor of 0.2. This system ensures that the model learns from the data and does not overshoot when trained.

In addition, the model weights with the best validation accuracy are saved using a ModelCheckpoint callback. This avoids overfitting and makes it possible to use the optimised model for inference by maintaining the top-performing model throughout training. The training approach is essential to developing an accurate and widely applicable skin cancer classification model.

Choosing the right hyperparameters was essential to the development of the proposed model. By considering learning rates, batch sizes, and network designs, a thorough grid search was conducted over a range of hyperparameter values. In order to attain optimal convergence without exceeding the global minimum, the learning rate was adjusted.

Regularization strategies like dropout and L2 regularization were used to reduce overfitting. During training, dropout was used to arbitrarily deactivate a portion of neurons, preventing hidden units from co-adapting. Large weights were punished using L2 regularization, which prevented the model from fitting noise into the data.

An extensively used stochastic optimization technique called the Adam optimizer was adapted as part of the optimization process. It provides quick convergence and adaptable learning rates by combining concepts from RMSProp and momentum.

Experiments were conducted within a standardized hardware and software environment to ensure reproducibility. Graphics computing Units (GPUs) capable of parallel computing were part of the hardware configuration; these were NVIDIA Tesla V100 GPUs. The efficient use of computational resources and quicker training times were made possible by this parallel processing.

The industry-standard Python deep learning frameworks TensorFlow and PyTorch served as the foundation for the software environment. Git was used for version control of the codebase, and the full repository was made publicly accessible in order to encourage repeatability and openness. The software dependencies were contained in Docker containers, guaranteeing consistency between various computing environments.

3.4 Data Augmentation

Data augmentation plays a crucial role in the skin cancer classification project by increasing the number and diversity of the training dataset, which strengthens and expands the model's capabilities. Multiple augmentation techniques are applied to the input images during training to replicate various variables that the model may meet in real-world settings.

The lesions' direction, location, and scale can change due to common data augmentation techniques, including random rotations, flips, zooms, and shifts. These adjustments are essential for avoiding overfitting and allowing the model to identify skin lesions in any location or orientation within the picture. In addition, methods such as contrast and brightness alterations are used to mimic variations in illumination, guaranteeing that the model may be adjusted to a variety of image characteristics.

Data augmentation has a significant effect on the generalisation of models. The model gains the ability to identify key characteristics of skin lesions in a range of scenarios by being trained on a larger and more varied collection of augmented photos. This reduces the possibility of overfitting certain patterns found in the original dataset, which eventually results in a skin cancer classification model that is more reliable and broadly applicable. The model's

performance on unseen data is greatly enhanced by adding data augmentation, which indicates its practical application in dermatological diagnosis.

4 System Architecture

EfficientNetB0, a CNN with exceptional efficacy and efficiency in image classification tasks, is the model architecture selected for this skin cancer classification research. The depth-wise separable convolutional structure of EfficientNetB0 allows for significant model parameter reduction without sacrificing representational capacity. The architecture also includes a Global Average Pooling (GAP) layer to aggregate spatial information and a Feature Pyramid Network (FPN) to enable feature extraction at many scales. To forecast whether skin lesions are benign or malignant, the final classification head comprises dense layers and fully connected layers with softmax activation. EfficientNetB0 is pre-trained on the ImageNet dataset using transfer learning, giving it a solid basis for identifying general visual patterns tailored to skin cancer classification’s intricacies. The training stability and convergence speed are further improved using the RAdam optimiser with an adaptive learning rate. All things considered, EfficientNetB0 is a good choice for skin cancer classification jobs because it can compromise computational efficiency and model correctness. The system architecture is shown in Figure 2.

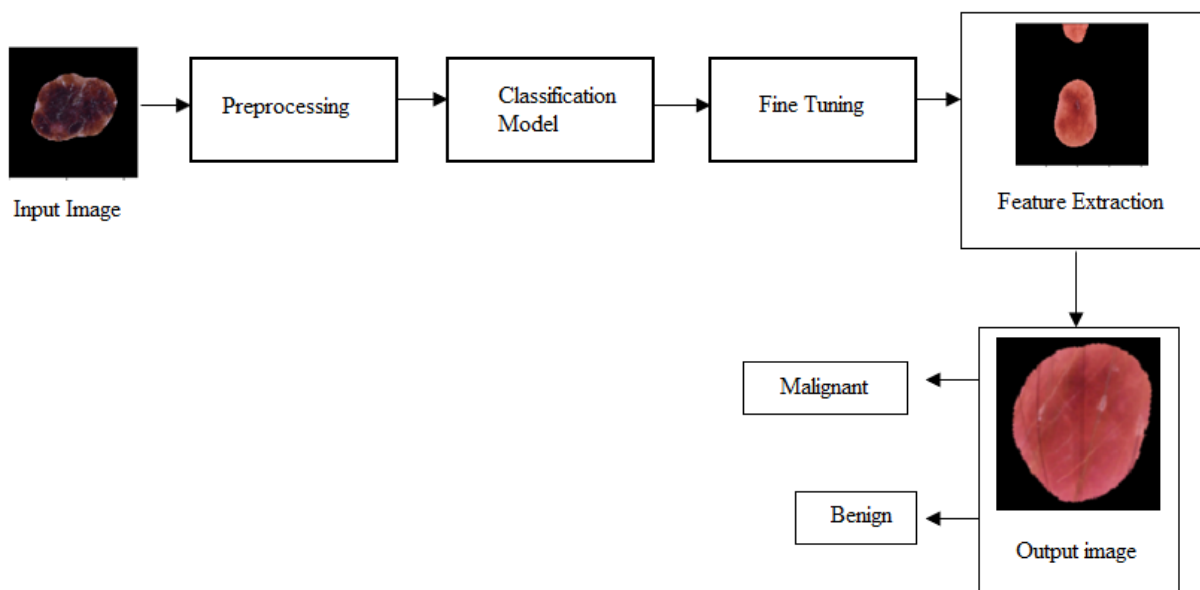


Figure 2. System architecture

The initiation of the model processing occurs at the input layer, where images of skin lesions are ingested. It considers the channels and spatial dimensions of the input images that correlate to colour information (like RGB). This layer’s dimensions match the width and height resolution of the skin lesion photographs. Following the input layer, a set of convolutional blocks is employed, each characterized by depthwise separable convolutions. These blocks are the main building blocks that extract features from the input photos. Depthwise separable convolutions improve the model’s efficiency by lowering the number of parameters and computing burden. An integral component of EfficientNetB0 is the FPN architecture, which facilitates the extraction of features across multiple scales. The model can extract high-level contextual information and fine-grained details from skin lesion photos because of its pyramidal shape.

Subsequent to the convolutional blocks, a GAP layer is utilized. This layer uses the average value of each feature map to aggregate spatial data. Better generalisation is facilitated by GAP’s contribution to lowering spatial dimensions while maintaining crucial features. The output from the GAP layer is channeled into fully connected and dense layers, culminating in the model’s classification head. These layers convert the features that were extracted into predictions for skin lesions that are benign or malignant. Probability scores for every class are produced by applying activation functions like softmax. Throughout the training phase, the model employs the RAdam optimizer in conjunction with an adaptable learning rate. RAdam optimises the speed and stability of the training process by dynamically adjusting the learning rate based on gradient information from the past. The methodology leverages transfer learning through the initialization of the model’s weights with pre-trained values from the ImageNet dataset. The model gains a basic grasp of generic visual patterns from this pretraining, which is then refined to specialise in the categorisation of skin cancer.

Understanding the complex interactions between these elements and layers is essential to understanding how the model can identify patterns and characteristics that indicate both benign and malignant skin diseases. The FPN architecture, depth-wise separable convolutions, and adjustable learning rate methods all work together to improve the model's efficacy and efficiency in the categorisation of skin cancer.

5 Results

5.1 Training Progress

The model's performance during training and validation epochs is shown graphically through accuracy and loss curves, which illustrate the training progress. These curves depict the accuracy and loss measures' temporal evolution. The research project that is being presented displays graphs that demonstrate the increase in accuracy and decrease in loss on the training and validation sets over the course of 30 training epochs. Among the noteworthy aspects is a notable improvement in validation accuracy, which suggests that the model may effectively generalise to new data. The optimisation process is highlighted by the loss curves, which similarly show a declining tendency. Figure 3 shows the dataset evaluation during the training progress of the model.

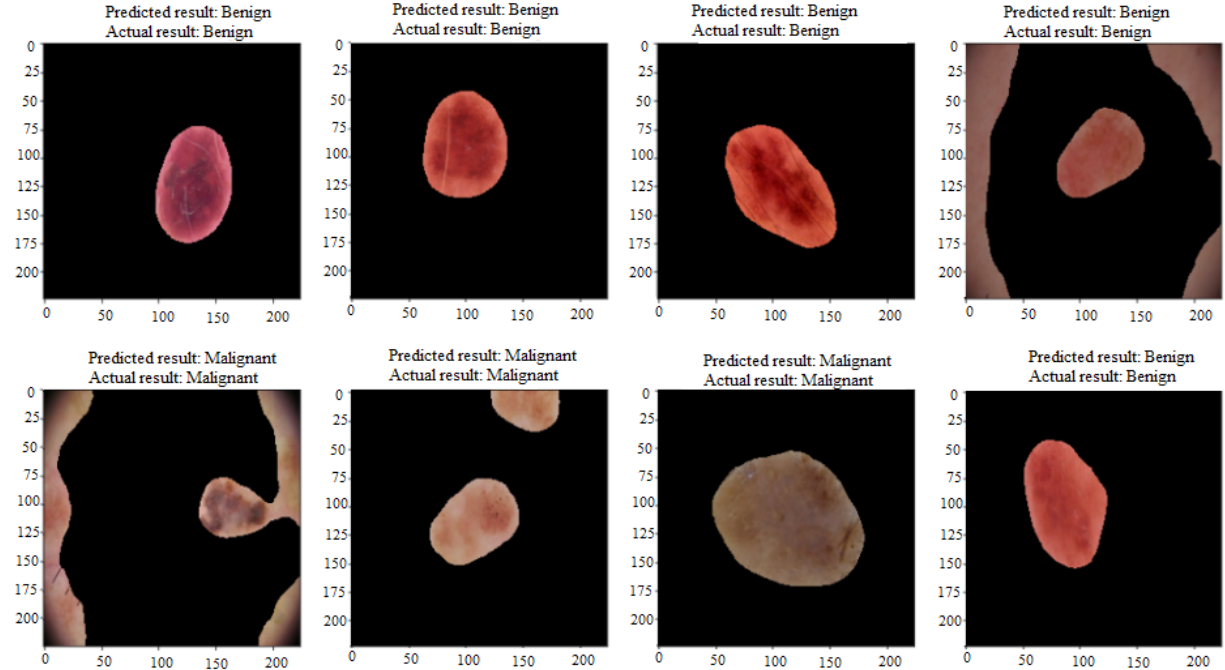


Figure 3. Cancer prediction during training and testing




Figure 4. Accuracy graph of the model

When the training process stabilises, and the model achieves optimal performance, it is referred to as model convergence. The model in the project under discussion exhibits convergence, as indicated by the diminishing returns in terms of loss reduction and accuracy improvement. As soon as the model starts learning from the training data, it is noted that the accuracy and loss measures change quickly. These gains become less pronounced as the epochs go by, suggesting that the model adequately represents the patterns found in the data. When the model converges, it indicates that it has acquired meaningful representations from the training set and that there is a balance between overfitting and underfitting. Figure 4 shows the accuracy graph of the proposed model.

Accuracy and loss curves, which show the training process, offer important insights into the model’s learning dynamics. These curves work as diagnostic tools, assisting practitioners and researchers in evaluating the model’s performance, spotting problems like under- or overfitting, and making well-informed choices for model tuning. For the skin cancer classification model created in this study to be reliable and useful, it is essential to comprehend the convergence behavior.

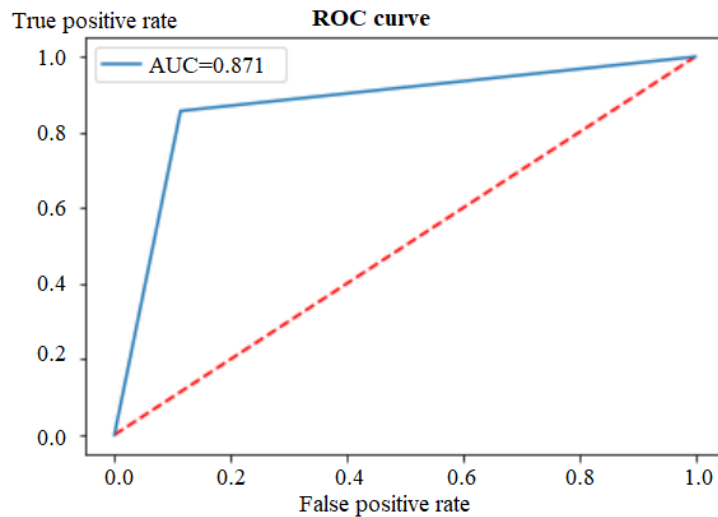


Figure 5. ROC curve graph




Figure 6. Validation loss graph

5.2 Model Evaluation

The skin cancer classification algorithm performed admirably, with an overall accuracy of 87%. For each class (benign and malignant), precision, recall, and F1-score were assessed, giving a thorough picture of the model’s capacity to discriminate between the two classes. The precision values demonstrated the model’s capacity to accurately identify instances of each class for the benign and malignant classifications, which were 88% and 86%, respectively. Furthermore, recall ratings of 89% for benign and 86% for malignant cases highlighted how well the model captured

pertinent examples. The robustness of the model in skin cancer classification was shown by the F1-score, which harmonised precision and recall and was 88% for benign and 86% for malignant skin cancer. Figure 5 and Figure 6 show the receiver operating characteristic (ROC) curve and validation loss graphs, respectively.

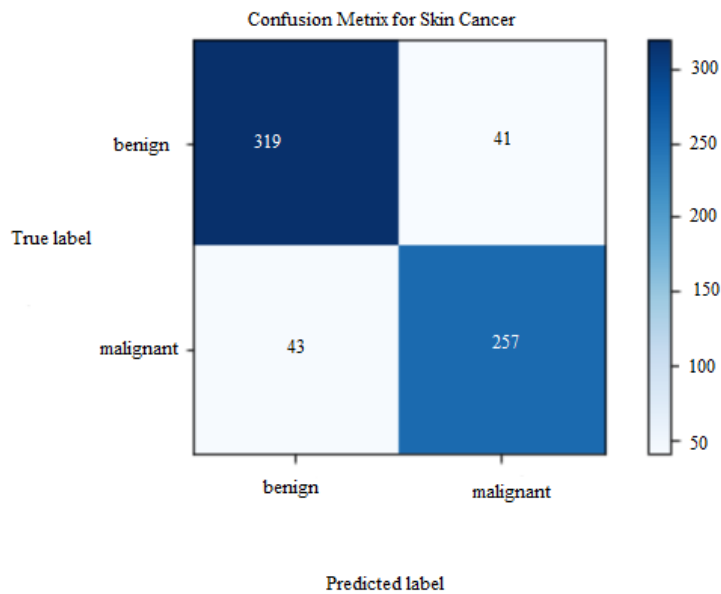


Figure 7. Confusion matrix graph of the model

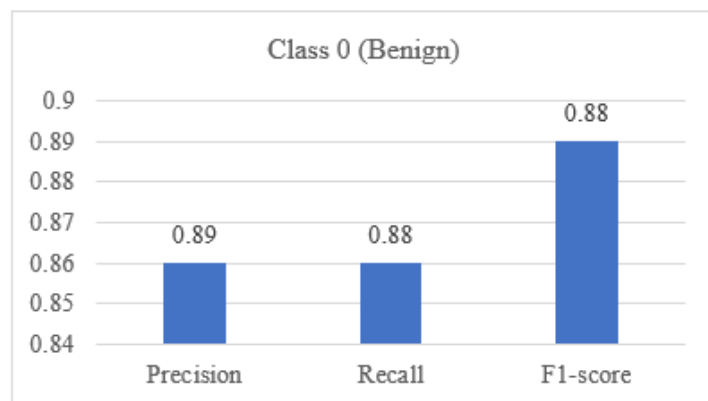


Figure 8. Class 0 evaluation graph

The model’s superiority in classifying skin cancer was evaluated by comparing its performance with baseline or existing models. Reaching 87% accuracy is a significant improvement in diagnosis accuracy. The precision, recall, and F1-score further demonstrate the model’s capacity to minimise false positives and false negatives. Based on this comparison, the suggested model is positioned as a significant addition to the classification of skin cancer, possibly surpassing or at least matching current approaches. The evaluation’s metrics demonstrate the created model’s efficacy and dependability in the context of dermatological diagnostics. Figure 7 shows the confusion matrix graph of the proposed model.

5.3 Confusion Matrix

The confusion matrix, which offers a thorough analysis of predictions made across many classes, is an essential tool for evaluating the effectiveness of a classification algorithm. To comprehend the model’s performance in the skin cancer classification project, the confusion matrix of the model was examined.

For Class 0 (benign), it was observed that 322 cases were accurately classified as benign, with 38 false negatives. For this class, an F1-score of 0.88, recall of 0.88, and precision of 0.89 were obtained, as shown in Figure 8.




Figure 9. Class 1 evaluation graph

Regarding Class 1 (malignant), 234 cases were correctly identified as malignant, with 66 false positives. For this class, an F1-score of 0.86, recall of 0.86, and precision of 0.86 were obtained, as shown in Figure 9.

6 Conclusion

In conclusion, this study presents a critical development in the early identification and categorisation of skin cancer, an urgent worldwide health issue. The proposed model, which uses deep learning techniques and the EfficientNetB0 architecture, performs admirably, with an overall accuracy of 87%. The model tackles the difficulties brought on by skin diseases' complex and varied nature and successfully differentiates between benign and malignant skin lesions. The research emphasizes that early diagnosis is important to improving patient outcomes and lessening the strain on healthcare systems. The suggested model enhances and reduces subjectivity in current diagnostic procedures by offering an automated and dependable categorisation tool. The thorough assessment, which includes confusion matrices, comparison with baseline models, accuracy metrics, and deep analysis, confirms the model's effectiveness and potential to revolutionise dermatological diagnostics. This work paves the path for future advancements in early intervention techniques and skin cancer classification by contributing to the larger fields of medical imaging and AI.

7 Future Scope

The skin cancer categorization system has prospective options for practical application and further refining, as indicated by its defined future scope. It is necessary to develop a systematic roadmap that includes user studies and prospective clinical trials in order to go from the current research stage to clinical application. Detailed instructions on how to include advanced technologies such as XAI should be provided, with a focus on how XAI improves the interpretability of models. Working together with health professionals and using real clinical data is essential to verifying the model for a variety of patient demographics. Plans for examining ensemble models and transfer learning techniques should also be specified in order to guarantee ongoing accuracy and robustness improvements. Strategic planning is necessary to address the model's scalability for deployment in telemedicine apps and mobile health platforms, bearing in mind areas with restricted access to dermatological expertise. To maintain the model's applicability in the ever-changing field of skin cancer diagnosis, regular updates and adjustments that consider changing medical imaging datasets and technological improvements must be a key component of the future strategy.

Data Availability

The data used to support the research findings are available from the corresponding author upon request.

Conflicts of Interest

The authors declare no conflict of interest.

References

- [1] A. L. Udriștoiu, A. E. Stanca, A. E. Ghenea, C. M. Vasile, M. Popescu, Ș. C. Udriștoiu, A. V. Iacob, S. Castravete, L. G. Gruionu, and G. Gruionu, "Skin diseases classification using deep learning methods," *Curr. Health Sci. J.*, vol. 46, no. 2, pp. 136–140, 2020. <https://doi.org/10.12865/chsj.46.02.06>
- [2] A. Panthakkan, S. M. Anzar, S. Jamal, and W. Mansoor, "Concatenated Xception-ResNet50—A novel hybrid approach for accurate skin cancer prediction," *Comput. Biol. Med.*, vol. 150, p. 106170, 2022. <https://doi.org/10.1016/j.combiomed.2022.106170>

- [3] D. Mane, R. Ashtagi, P. Kumbharkar, S. Kadam, D. Salunkhe, and G. Upadhye, “An improved transfer learning approach for classification of types of cancer,” *Trait. Signal*, vol. 39, no. 6, pp. 2095–2101, 2022. <https://doi.org/10.18280/ts.390622>
- [4] L. Hoang, S. H. Lee, E. J. Lee, and K. R. Kwon, “Multiclass skin lesion classification using a novel lightweight deep learning framework for smart healthcare,” *Appl. Sci.*, vol. 12, no. 5, p. 2677, 2022. <https://doi.org/10.3390/app12052677>
- [5] D. Tiwari, M. Dixit, and K. Gupta, “Deep multi-view breast cancer detection: A multi-view concatenated infrared thermal images based breast cancer detection system using deep transfer learning,” *Trait. Signal*, vol. 38, no. 6, pp. 1699–1711, 2021. <https://doi.org/10.18280/ts.380613>
- [6] T. J. Brinker, A. Hekler, J. S. Utikal, N. Grabe, D. Schadendorf, J. Klode, C. Berking, T. Steeb, A. H. Enk, and C. von Kalle, “Skin cancer classification using convolutional neural networks: Systematic review,” *J. Med. Internet Res.*, vol. 20, no. 10, p. e11936, 2018. <https://doi.org/10.2196/11936>
- [7] R. Refianti, A. B. Mutiara, and R. P. Priyandini, “Classification of melanoma skin cancer using convolutional neural network,” *Int. J. Adv. Comput. Sci. Appl.*, vol. 10, no. 3, 2019. <http://doi.org/10.14569/IJACSA.2019.0100353>
- [8] N. Zhang, Y. X. Cai, Y. Y. Wang, Y. T. Tian, X. L. Wang, and B. Badami, “Skin cancer diagnosis based on optimized convolutional neural network,” *Artif. Intell. Med.*, vol. 102, p. 101756, 2020. <https://doi.org/10.1016/j.artmed.2019.101756>
- [9] R. Garg, S. Maheshwari, and A. Shukla, “Decision support system for detection and classification of skin cancer using CNN,” in *Innovations in Computational Intelligence and Computer Vision*. Springer, Singapore, 2021, pp. 578–586. https://doi.org/10.1007/978-981-15-6067-5_65
- [10] R. Patil and S. Bellary, “Machine learning approach in melanoma cancer stage detection,” *J. King Saud Univ.-Comput. Inf. Sci.*, vol. 34, no. 6, pp. 3285–3293, 2022. <https://doi.org/10.1016/j.jksuci.2020.09.002>
- [11] R. Patil and S. Bellary, “Transfer learning based system for melanoma type detection,” *Revue Intell. Artif.*, vol. 35, no. 2, pp. 123–130, 2021. <https://doi.org/10.18280/ria.350203>
- [12] R. Patil and S. Bellary, “Ensemble learning for detection of types of melanoma,” in *2021 International Conference on Computing, Communication and Green Engineering (CCGE)*, Pune, India, 2021, pp. 1–6. <https://doi.org/10.1109/CCGE50943.2021.9776373>
- [13] A. Padthe, R. Ashtagi, S. Mohite, P. Gaikwad, R. Bidwe, and H. M. Naveen, “Harnessing federated learning for efficient analysis of large-scale healthcare image datasets in IoT-enabled healthcare systems,” *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 10s, pp. 253–263, 2024.
- [14] R. Ashtagi, D. Mane, M. Deore, J. R. Maranur, and S. Hosmani, “Combined deep learning and machine learning models for the prediction of stages of melanoma,” *J. Auton. Intell.*, vol. 7, no. 1, pp. 1–13, 2024. <https://doi.org/10.32629/jai.v7i1.749>
- [15] S. Chaturvedi, J. Tembhurne, and T. Diwan, “A multi-class skin cancer classification using deep convolutional neural networks,” *Multimed. Tools Appl.*, vol. 79, pp. 28477–28498, 2020. <https://doi.org/10.1007/s11042-020-09388-2>