

DERMXNET: A HYBRID DEEP LEARNING AND GRADIENT BOOSTING APPROACH FOR EFFICIENT SKIN DISEASE DETECTION

N ANNALAKSHMI¹, S UMARANI²

¹Research Scholar, Department of Computer Science, SRM Institute of Science and Technology, Ramapuram, Chennai, India

²Professor, Department of Computer Science and Applications (BCA), SRM Institute of Science and Technology, Ramapuram, Chennai, India

E-mail: ¹an6604@srmist.edu.in, ²ravana@gmail.com

ABSTRACT

Skin diseases, including melanoma and non-melanoma, are among the most prevalent health concerns worldwide. Early and accurate detection of these conditions is critical for effective treatment and improved patient outcomes. This paper introduces DermXNet, a hybrid model combining Artificial Neural Networks (ANN) and eXtreme Gradient Boosting (XGBoost) for accurate classification of skin conditions, including melanoma, non-melanoma, and normal skin. The methodology involves systematic data pre-processing techniques such as resizing, normalization, and artifacts removal. The ANN module extracts high-dimensional features, which are classified by XGBoost, leveraging the strengths of both deep learning and gradient boosting. The proposed model was evaluated against nine existing deep learning models, including ResNet50, VGG16, and EfficientNetB0. DermXNet achieved the highest accuracy of 98.38%, surpassing other models in metrics such as precision, recall, F1-score, and AUC-ROC. Additionally, DermXNet demonstrated computational efficiency with a training time of 95 seconds and a model complexity of just 4 million parameters, making it suitable for real-world deployment. The results underscore the effectiveness of hybrid architectures in medical diagnostics. Future research can extend DermXNet to multi-class classification and integrate advanced domain-specific features to enhance its applicability.

Keywords: *Skin Disease Detection, Deep Learning, Artificial Neural Network, XGBoost, Hybrid Model, Medical Imaging, Melanoma Classification, Feature Extraction*

1. INTRODUCTION

Skin cancer is among the rapidly growing cancer incidences in the global market, which is characterized by the growth of skin cells with distorted cell structure. It mainly results from the prolonged exposure to the rays of ultraviolet, which is a component of sunrays or artificial sources like tanning equipment. Skin cancer is not a single disease and can be classified into three broad categories, including basal cell carcinoma, squamous cell carcinoma, and melanoma [1] [2]. Even though basal cell and squamous cell carcinoma are less invasive as well as curable through early diagnosis and treatment, melanoma is severe and can metastasize to other parts of the body, hence the need to expedite early diagnosis and treatment. Skin cancer is becoming more common throughout the world, which has led to significant health concerns, especially in high

radiation areas or where individuals cannot afford access to sunscreen and protective clothing. Other risk determinants include family history, type of skin, and a history of severe sun burning. Technological developments in the field of medicine enhanced dermo copier and biopsy diagnostic procedures in addition to targeted treatment, immunotherapy, and surgery. Screening such as skin examinations, advertisement on danger signs and symptoms, and teaching on issues to do with UV radiation act as the strategies that go a long way in ensuring low incidence rates and better patient outcomes [3] [4]. Skin cancer remains an active site of research on genetics and environmental precursors with an increasing number of advances to its treatment and prevention.

In the case of fighting skin cancer, several constraints go with it and make it even

harder to avoid, diagnose or cure. The first and significant problem includes low levels of health literacy regarding the necessity of timely diagnosis and skin protection from UV radiation, which results in the majority of patients receiving a diagnosis in a later, more dangerous stage [5] [6]. The skin examination is frequently times constricted, especially in resource scale regions, which hinders early detection and appropriate diagnosing. Biopsy procedures, which are otherwise efficient in this task, are also very much invasive and time consuming in nature and the possibilities of misdiagnosis as well has identified particularly in cases of atypical form. Further, some of the new drugs like immunotherapy and targeted therapy are very expensive and out of pocket reach of majority of patients. Sun protection also influenced by abnormal psychosocial factors, including incorrect use of sun protective measures like inadequate use of sunscreen, reluctance not to tan, and underestimation of risks associated with sun exposure. Limitations of the current technology call for enhanced public health awareness, cheaper health care and constant research and development of non-invasive diagnosis and curing methods [7] [8].

Deep learning models have emerged as very effective in skin cancer detection and diagnosis because of their ability to analyse medical image data. Particularly popular in the analysis of dermoscopic images are convolutional neural networks (CNN), which allows for the differentiation of skin lesions as malignant or benign with an accuracy comparable to or higher than that of dermatologists [9] [10]. These models can capture shapes and structural details that are important in melanoma and other skin cancer diagnosis, including asymmetry, border irregularity, and colour. Using deep learning in diagnosis, medical professionals can increase efficiency of the diagnosis process and reduce variability, which can lead to early diagnosis and better treatment. In addition, machine-learning algorithms integrated into the smart phone application in order to provide ordinary people with basic diagnostic tools when there is a lack of dermatological services or specialists. These models are dependent on big, clean data sets to train on, which can be a problem as these may not always be available particularly for some forms of skin cancer or on people of colour, which implies that there could be prejudice in the prognosis. Training and implementing of these models

demand computational resources, which may be expansive in many developing nations [11] [12]. In addition, such environmental factors result in overreliance on automated tools, and it is not shocking to find some users disregarding consultation from professionals in the health sector. Solving these problems means creating various datasets, enhancing model interpretability, and creating conditions in which DL tools supplement clinical knowledge rather than replace it. Figure 1 shows the various types of skin disease. The models designed to take large amounts of labelled data, which are frequently missing or inadequate in arid and remote regions. Deep learning calls for extensive computational power, a factor that makes the application of the technology expensive, thus possibly unaffordable for many a region [13] [14]. If the training datasets biased in some way, for instance, they may include certain geological conditions, it can cause wrong predictions or wrong identification of water resources. Such challenges indicate that ethical data practices, capacity development as well as a commitment towards research and development of sustainable AI technologies are important for the enhancement of the efficiency of deep learning in subsurface hydrology of arid regions [15].



Figure 1 Various Types of Skin Disease

Traditional diagnostic methods, relying on visual inspection and biopsy, are time-consuming and prone to subjective interpretation. Recent advancements in artificial intelligence (AI) and deep learning have revolutionized the field of medical imaging, offering automated, precise, and scalable solutions for disease detection. This paper introduces DermXNet, a hybrid deep learning model designed for the accurate classification of skin conditions.

DermXNet leverages the feature extraction capabilities of Artificial Neural Networks (ANN) and the classification robustness of eXtreme Gradient Boosting (XGBoost). The model employs advanced pre-processing techniques to clean and normalize the data, ensuring high-quality input for learning.

The work is concerned with binary classification of skin diseases by distinguishing malignant melanomas and non-melanomas from normal skin. While DermXNet showed high accuracy and computational efficiency of neural networks, its performance is based on the quality and diversity of the dataset. These two drawbacks will be left open for future research: the study does not approach multistage classification of other dermatoses.

The study primarily focuses on training DermXNet—a hybrid deep learning model that synergizes ANN feature extraction with XGBoost classification—to improve the diagnosis of skin disease. The major evaluation metrics include classification, accuracy, precision, recall, F1-score, AUC-ROC, and computational efficiency. The performance of the model is benchmarked with existing architectures for deep learning. Compared to the traditional deep learning models, DermXNet incorporates an ANN and an XGBoost to solve class imbalances and overfitting while achieving high values in prediction accuracy with lower computational costs. Further, the study integrates a top-end advanced pre-processing pipeline which is said to make it more efficient and suitable for real-life deployment.

1.1 Problem Statement

Skin cancer remains a pathogenically increasing worldwide malignancy that results in severe health threats from melanoma and non-melanoma variety. Early and accurate detection is key to good treatment; nevertheless, traditional diagnostic methods, including visual inspection and biopsy, insist on time, subjectivity, and specialized expertise. Despite good preeminence shown by deep learning-based models, and in particular CNNs, in reinforcement of skin disease classification, these methods still face challenges like class imbalance, overfitting, and high computational cost of learning. Moreover, these models often rely on very large labeled datasets that are hard to find and generally share some possible biases against some under-represented

classes. Comparative studies have been conducted on architectures based on CNN, including ResNet, VGG16, and EfficientNet, attaining dermatologist-level accuracy [1], [2]; however, these models might be demanding on computational capacity with limited generalizability across diverse datasets [3]. Conventional deep learning models suffer from a lack of interpretations, leading to challenges in understanding decisions by the healthcare profession. Against this backdrop, this research proposes DermXNet, a hybrid model with ANN for feature extraction and XGBoost for classification, to improve the model's accuracy, computational efficiency, and applicability in dermatology.

1.2 Research Objectives

The key objective of the study is to design, implement, and evaluate DermXNet, a hybrid deep learning model for efficient and correct classification of skin diseases. Specific objectives thus include:

- Designing and implementing an ANN-XGBoost hybrid model intrinsic of deep learning and boosting to provide robust classification of melanoma and non-melanoma skin tumors.
- Developing an optimized data pre-processing pipeline comprising image resizing, normalization, artifact removal, and augmentation that together improve the quality of inputs and boost performance.
- To evaluate the performance of DermXNet concerning state-of-the-art deep learning models, such as ResNet50, VGG16, EfficientNetB0, and Xception, some evaluation metrics used include accuracy, precision, recall, F1-score, AUC-ROC, and computing efficiency.
- To evaluate the computational efficiency of the DermXNet in relation to their potential real-time applications in mobile health platforms and resource-limited clinical settings.
- To investigate the scalability and its potential clinical deployment of DermXNet, concerning the feasibility of telemedicine, AI-assisted diagnosis, and mobile healthcare solutions.

1.3 Research hypothesis

H1: The proposed DermXNet model, which constitutes the fusion of Artificial Neural Networks (ANN) and eXtreme Gradient Boosting (XGBoost), will achieve increased classification accuracy as against conventional deep learning models as ResNet50, VGG16, and EfficientNet.

H2: The hybrid ANN-XGBoost architecture will prove to be robust against class imbalance and eventually enhance the precision, recall, and F1-score for the classification between melanoma and non-melanoma.

H3: Introduction of an optimized data pre-processing pipeline including artifact removal, denormalized images, and augmentation will yield better generalization of models on diverse datasets.

H4: DermXNet will be computationally efficient, through lower model complexity and a substantially fast inference time, and thus suitable for real-time applications in mobile and resource-constrained clinical settings.

H5: DermXNet shows promise for scalability in telemedicine and in AI-assisted diagnostics for the early detection and intervention of skin diseases.

1.4 Main Contribution of the Work

- **Hybrid Model Design:** DermXNet introduces a novel hybrid architecture combining the feature extraction power of Artificial Neural Networks (ANN) with the classification efficiency of eXtreme Gradient Boosting (XGBoost). This integration leverages the strengths of both deep learning and machine learning approaches, creating a robust framework for binary classification of skin diseases.
- **Comprehensive Pre-processing Pipeline:** A systematic data pre-processing pipeline is implemented, including image resizing, normalization, artifact removal (e.g., hair inpainting), and augmentation. These steps ensure high quality and standardized inputs, enhancing the model's ability to learn

relevant features while minimizing the impact of noise and irrelevant patterns.

- **Computational Efficiency:** DermXNet achieves computational efficiency by utilizing lightweight model architecture with minimal parameters (4 million), optimized for faster training and inference. This makes it suitable for deployment in resource-constrained environments, such as mobile or edge computing devices.
- **Extensive Comparative Analysis:** The model rigorously evaluated against nine state-of-the-art deep learning and hybrid models, including ResNet50, VGG16, and Xception. This comparative study highlights the efficacy of hybrid architectures for medical imaging tasks and provides insights into their computational and functional advantages.

The remainder of this paper organized as follows: Section 2 reviews related studies, highlighting advancements in skin disease detection using deep learning and hybrid models. It discusses existing architectures, challenges, and gaps in the field. Section 3 details the proposed DermXNet methodology, including its hybrid architecture, pre-processing pipeline, and model design. Section 4 describes the results and discussion, presenting a comprehensive evaluation of DermXNet against state-of-the-art models through performance metrics, computational efficiency, and comparative analysis. Finally, the Conclusion and Future Scope section summarizes the findings, emphasizing DermXNet's potential for clinical use and directions for future research.

2. RELATED WORKS

Machine learning or artificial intelligence aids in diagnosis and even individualized treatment plans for skin-related problems. The latest advancements in efficient manipulation of large amount of data, faster computers, and relatively cheap data storage have further expanded the role of ML in dermatology. The article provides a brief overview of the basic aspects of applying ML, including implementations, issues, and concerns involved in constructing skin cancer diagnostic and categorization systems [16]. Featured fields include disease categorization through clinical images, cancer visualization in

dermatopathology, and skin disease monetarization using smart-phones apps. Its purpose is to explain the basic concepts of ML to dermatologists and describe the necessary conditions for designing applications for the detection of skin cancer. These are the lesion tracking and full-image segmentation, and most of the surveyed techniques aim at solving these problems.

Metastatic malignant melanoma becomes incurable being one of the fatal types of skin cancer. Initial detection continues to be important as a way of increasing survival rates. Several ML techniques, including those that focus on classification and segmentation, have used in early detection [17]. The ML-based studies, their aims, constraints, and prospects described, with reference to certain works. Supervised and unsupervised classifications presented and compared to give analyses of investigation in malignancy detection. Focusing on the algorithm choice, it highlights its significance in a high diagnostic accuracy perspective, and contributes to future developments of melanoma detection and classification systems.

Melanoma occurs when there is proliferation of melanocytes, specialized cells that produce the pigment melanin and it is imperative to diagnose this kind of cancer at its early stage in order to increase the survival rate of the patients. This makes its differential diagnosis difficult especially because it can be clinically confused with benign skin lesions. A fuzzy logic-based image segmentation model integrated with deep learning for analysing skin cancer image presented. These are dermoscopy image enhancement by using pre-processing methods such as elimination of artifacts, defuzzification techniques [18]. A variant of a deep neural network used for melanoma detection using high accuracy; it is known as You Look Only Once (YOLO). The research proves that YOLO trained on ISIC datasets yields substantially higher performance levels through feature concatenation and inclusion of additional convolutional layers, outperforming the classifiers currently in existence while processing images at a significantly faster speed.

Cancer of the skin particularly melanoma is among the top causes of deaths from cancers. There has been an increase of misdiagnoses especially in large diagnostic centres hence the creation of automated diagnostic checkers. The

approach presented aims at proposing a new hybrid deep learning model between deep learning and machine learning techniques [19]. This model uses neural networks on feature extraction points and processes them using filters such as the Contourlet Transform as well as the Local Binary Pattern Histogram. The proposed model integrates both the manual and automation elements and yields 93 % accuracy and great recall values. Based on the ISIC Archive datasets, it outperforms dermatologists as well as current state-of-the-art classifiers, providing a useful tool for mitigating diagnostic mistakes.

Deep Convolution Neural Networks (DCNNs) have completely solved skin cancer detection, which the authors state encompasses a highly visual challenge. Conducting the work benefits from the EfficientNet model under transfer learning, to distinguish melanoma from non-melanoma lesions accurately [20]. To address such issues as image resolution variability or class imbalance the study expands datasets and includes metadata. The EfficientNet-B6 model with the ranger optimizer was the best performers among other models with an AUC-ROC Of 0.9681. This reveals that the architecture successfully acquires complex patterns and establishes a new high level of skin lesion classification, improving the diagnostic level in dermatology.

The developments in deep learning for skin disease diagnosis have come forth a long way, especially with CNN-based models such as ResNet, VGG16, and EfficientNet. employed NIR spectroscopy and used different machine learning algorithms aimed toward skin cancer diagnosis, obtaining high accuracies against the problem of dataset diversity and practical implementability. Similarly, [2] suggested SNC_Net, a combination of handcrafted and deep-learning-based features; however, it faced challenges in class imbalances and computational costs, limiting practical applicability in low-resource settings.

Some architectures, such as ResNet50 and EfficientNet, exhibited classifier performance at dermatologist levels (Ali et al., 2024); however, they require high computational resources and hence are less viable for real-time applications or mHealth. NBCNNs have been reported to suffer from the problem of no interpretability, resulting in a lack of acceptance in the clinical environment, as decisions are often facilitated not

just by CNNs, but also by explainable AI models in clinical settings (Zhang et al., 2024). To counter these shortcomings, different hybrid approaches that combine deep learning and machine learning have been explored. Tuncer et al. [18] proposed a lightweight CNN model for skin cancer classification, yet because of simple feature extraction capabilities, it cannot capture lesion structure complexities adequately. Further studies combining gradient boosting with CNNs [6] showed better robustness toward classification but lacked a formal pre-processing pipeline that would enable assuring a viable-quality input.

3. METHODOLOGY

The proposed methodology for skin disease detection integrates advanced data pre-processing, feature extraction, and classification within a hybrid model called DermXNet. Pre-processing includes resizing, normalization, and artifact removal to ensure clean and consistent input data. The model's architecture combines an Artificial Neural Network (ANN) for extracting high-dimensional features with eXtreme Gradient Boosting (XGBoost) for robust classification. The ANN captures intricate patterns and textures from the images, while XGBoost effectively handles non-linear relationships and class imbalances. This synergy between ANN and XGBoost ensures high performance, computational efficiency, and adaptability, making DermXNet a reliable solution for medical imaging diagnostics.

3.1 Dataset Collection

To create a strong architecture for skin disease detection, we build a large dataset that consists of melanoma and non-melanoma. It is a great support for model training and analysis. For melanoma and non-melanoma cases, we used dataset available on Kaggle. This is, amongst others, recognition as one of the best and widely used sources for sample medical image data. It is a provision of a wide base of images with lesions annotated, offering a wide variability of appearances, textures, and pigmentation, which is important in training the model to distinguish between malignant and benign lesions. We gathered images of normal skin from individuals using 5G networks to ensure that the model accurately distinguishes a diseased region from a healthy skin region and improves on its diagnostic function in actual practical situations. The manual data collection concerned obtaining images from publicly accessible medical images database and

dermatological references obtained from the Web while having ethical and data relevance to the study goals in mind. The database collectively contains 1,095 images shared amongst melanoma, non-melanoma, and normal skin images. This is immensely helpful because it helps remove biases from the model and contend with the great variety of skin problems for even more accurate diagnosis. Thus, combining open access and handpicked data sources, we intend to develop a robust and clinically relevant classifier for skin diseases. Figure 2 shows the dataset samples.




Figure 2 Dataset Samples

3.2 Data Pre-processing

3.2.1 Image resizing

The first pre-processing step is to resize the input images to fixed dimension of size 1024×1024 pixels using `cv2.resize` function. Image resizing is important for dimensionality purposes because, when training the model, images have to be of one fixed size. This done in order to exclude variability due to the arrangements of the camera setup and source of the images, and the resolutions acquired. This step also helps in easing upstream process and guarantees that the model fed with inputs of same qualities and magnitude thus helping in getting good and accurate results. The resizing process transforms an image I from its original dimensions (H, W) to a target size (H_r, W_r) . The new pixel value at position (x', y') is interpolated based on the original image:

$$I_r(x', y') = \text{Interpolate}(I(x, y)), \quad \forall x' \in [0, H_r), y' \in [0, W_r) \quad (1)$$

Where $I_r(x', y')$ is the resized image, and H_r, W_r are the target height and width.

3.2.2 Grayscale conversion

The next step is the resizing of the images then the RGB images then converted to grayscale using the `cv2.cvtColor` function with the `COLOR_RGB2GRAY` flag. Grayscale conversion makes the reduction of 3-D image into 1-D image where instead of three channels (Red, Green and Blue); there will be one intensity channel. This helps decrease computational intensity or cost greatly and yet preserves most textural and intensity information necessary for feature identification such as lesions or abnormalities. By not using any feature related with color information for classification but the intensity variation, this eliminates the need of having to extract any other structural features for classification from the pre-processing pipeline. It down samples the RGB color channels to a single tone differing in intensity. The intensity at each pixel $I_g(x, y)$ computed as a weighted sum of the RGB components:

$$I_g(x, y) = 0.2989 \cdot R(x, y) + 0.5870 \cdot G(x, y) + 0.1140 \cdot B(x, y) \quad (2)$$

Where $R(x, y)$, $G(x, y)$, $B(x, y)$ are the red, green and blue channel values of the original image at (x, y) , and $I_g(x, y)$ is the grayscale intensity at (x, y) .

3.2.3 Morphological filtering (blackhat operation)

To highlight some of the characteristic in the images of the faces the BlackHat morphological operation done on the grayscale faces images. This technique done with `cv2.morphologyEx` function, using a kernel defined with the `cv2.getStructuringElement` function, useful when we have an image with dark areas on a lighter base. While using this step in the detection of skin diseases the contrast for hair contours or any other fine structures, which may cause an interference, enhanced. In effect, such unwanted artifacts made more salient by the BlackHat operation to follow by their elimination in the subsequent stages. The BlackHat operation defined as the difference between the closing (\cdot) of the image and the original image. For an image I_g and a structuring element K :

$$I_{blackhat} = (I_g \cdot K) - I_g \quad (3)$$

Where $I_{blackhat}$ is the result of the BlackHat operation, $I_g \cdot K$ is the morphological closing: dilation followed by erosion using the kernel K , and K is typically a rectangular or elliptical kernel.

3.2.4 Thresholding

After the BlackHat operation, the binary threshold applied to the processed grayscale images using the `cv2.threshold` function. This step creates a mask that separates the image into regions of interest and non-interest, or in other words, isolates certain unwanted artefacts (for example, hair) determined in the previous step. This binary threshold makes it easier to distinguish between these artifacts and the other parts of this image, and this threshold gives a very clear pointer to every further step that should take. It is important for the purposes of generating clean and artifact-free results to clearly isolate this step. Thresholding creates a binary mask by comparing each pixel value $I(x, y)$ with a threshold T :

$$I_{threshold}(x, y) = \begin{cases} 255, & \text{if } I_{blackhat}(x, y) > T \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

Where $I_{threshold}(x, y)$ is the binary mask and T is the threshold value. Figure 3 shows the architecture of proposed model.

3.2.5 Inpainting

The last operation in the pre-processing step is the inpainting which is done using the `cv2.inpaint` function. Similarly, inpainting employs the mask produced by the thresholding process in order to determine which “repair” areas are needed. The algorithm paints over these areas with pixel values taken based on the values of the pixels in their outskirts, thus erasing all sorts of detail like hair edges. This way the original image reconstructed but the image structure integrally preserved using this approach to have clear and fundamentally correct image of the skin surface. The inpainted images are then usable for input to the final machine-learning model whereby the training sets are now clear of annoying artifacts that may complicate the feature of interest. Inpainting replaces the pixel values in the masked region $M(x, y)$ with interpolated values from the surrounding region. Using the Telea method (TELEA):

$$I_{inpainted}(x,y) = TELEA(I_r(x,y), M(x,y)) \quad (5)$$

Where $I_{inpainted}(x,y)$ is the final inpainted image, $M(x,y)$ is the binary mask generated by thresholding, and $TELEA$ represents the inpainting algorithm.

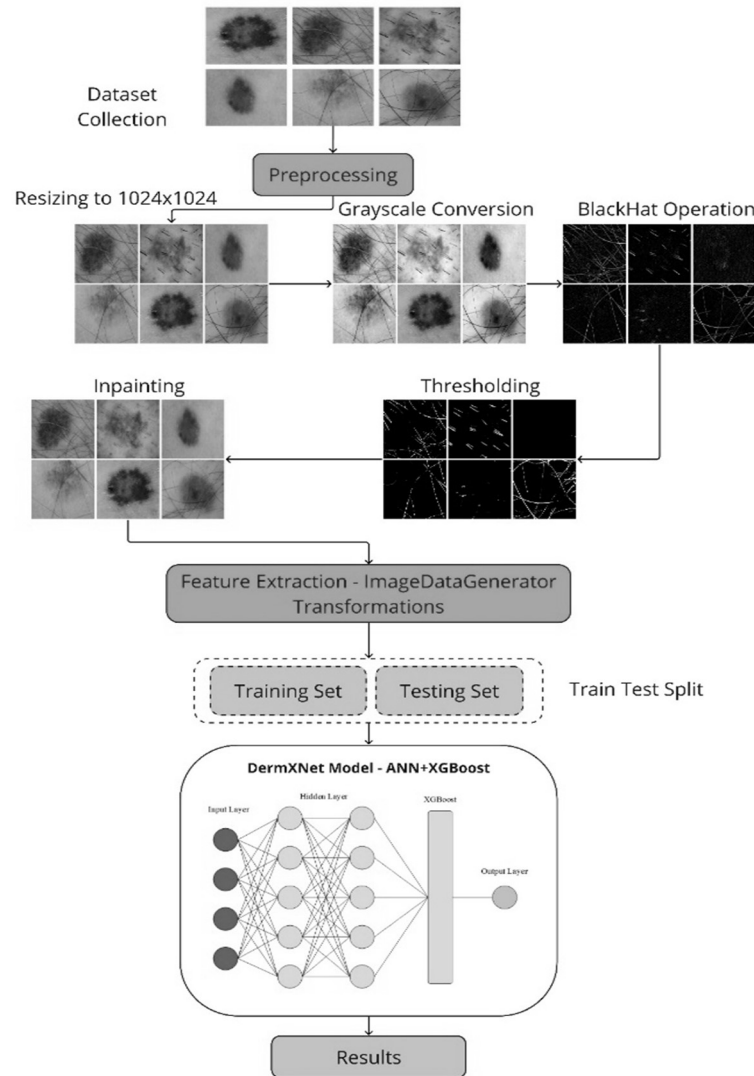


Figure 3 Architecture of Proposed Model

3.3 Extracting Features Using ImageDataGenerator

ImageDataGenerator is a useful tool provided in Keras for feature extraction as well as data augmentation in the relevant dataset of images. It intended to work before feeding image data into a model during the model training or the model building for prediction phase. Apart from this, this dynamic feature extraction increases the effectiveness of the learning procedure and makes

the resulting dataset more diverse and informative to increase the model's capability to generalize in other contexts. If used for feature extraction, ImageDataGenerator wraps images with successive transformations before passing them through the model. Such transformations include scaling of intensities by factors such as scaling the pixel values, scaling the intensity range of the acquired data and additional modifications including rotation of the data randomly, shift, flip, and zoom. For example, rescaling presupposes

dividing pixel values by a constant, for, instance, 255 in order to normalize an input for the range [0, 1] that helps to stabilize and accelerate the learning process of such neural networks. Besides improving the feature representation of the input data, it also add more diversity to the view of the same image, which can result in the reduction of overfitting. Feature-wise normalization standardizes the dataset by centering the pixel values and scaling by the standard deviation:

$$I_{normalized}(x, y, c) = \frac{I(x, y, c) - \mu_c}{\sigma_c} \quad (6)$$

Where μ_c is the mean pixel value for channel c over the entire dataset, and σ_c is the standard deviation of pixel values for channel c . Another significant feature of ImageDataGenerator involves normalizing features using image-level means and standard deviations or executing feature-wise scaling to control brightness and contrast of the images. These operations make the input data centralized and scaled to help the model capture essential patterns easily. For instance, scaling the intensity levels of the images within the dataset makes them insensitive to variations that result from image acquisition precondition such as lighting conditions and camera settings and thus more sensitive to feature preconditions such as the texture, shape, and edges of the images. The second advantage of utilizing ImageDataGenerator for feature extraction is on processing data in batches when the incoming data is large. Unlike other data-loading mechanisms, which require the data to be loaded into memory completely, it only requires the image data to be loaded at each iteration in the case of large data sets, which cannot accommodate in the memory space for processing. Due to the generator processing one batch at a time, it takes in features before passing them to the model, a procedure that enhances memory utilization and computation. Given a dataset with N images, the generator processes data in batches of size B . For the i -th batch, the set of images is:

$$Batch_i = \{I_k | k = (i - 1)B + 1, \dots, iB\} \quad (7)$$

Where I_k is the k -th image in the dataset, and B is the batch size. In addition to built-in pre-processing and augmentation functions, ImageDataGenerator allows creating a custom pre-processing pipeline. Pre-processing functions for images defined here and applying image

specific operations may include edge detection, color space conversion, and even domain specific filtering for extracting the specified features from input images. This brings the flexibility of being able to fine-tune the generator for specific application such as in disease diagnosis from images or detecting objects. ImageDataGenerator helps models extract better features and have more data, which in turn leads to better performance when the data being tested on the models for the next time. The extracted features defined to consist of summarized characteristics of the input images and augmentation makes input immune to noise, rotation and other similar factor. The processed batch of images then fed to the model:

$$\begin{aligned} & \text{Model Input} \\ &= \{I_{processed}^{(1)}, I_{processed}^{(2)}, \dots, I_{processed}^{(B)}\} \end{aligned} \quad (8)$$

Where $I_{processed}^{(k)}$ represents the preprocessed feature vector of the k -th image in the batch.

3.4 Train-Test split

In the development of the DermXNet model for skin disease detection, the dataset divided into training and testing subsets using an 80:20 split ratio. This means by far 80% of the dataset used for training the model leaving the 20% for testing the performance of the model. This split is standard in machine learning since it helps assess the model's ability of generalizing from the category of data it trained. Consequently, the training set, which comprises 80% of all the data, used to provide the model with information about the classifiers that distinguish between various skin conditions. The first part of ANN captures the abstract features, edges, and other features from images, while the second subset, XGBoost, captures decision boundaries from the feature space. This allocation ensures that model has ample data on these features to observe these patterns this reduces the likelihood of under fitting. Furthermore, the other typical strategies used for improving the quality and the variety of the training set in order to increase the model's ability to generalize.

The final 20% of data not used during the training of the system and constitutes the testing set. It used to test the calibrated model on upcoming examples by giving an impartial indication. When comparing the values of the

predicted results with the original labels, such measures as accuracy, the level of non-compliance, recall, the F - index and AUC-ROC calculated in order to assess how well the model is prepared. This is important step for assessing the validity of the model in the practical conditions. The 80:20 split gives a good equilibrium between the data to use for training and the data for testing. It does not generalize high variance by getting very little testing data or too less training data, which leads to under fitting. This method ensures that DermXNet is well trained and subjected to strict evaluation to minimize chances of making terrible mistakes that are associated with skin diseases classification.

3.5 Proposed Model - DermXNet model

DermXNet is a developing hybrid model that uses Artificial Neural Networks (ANN) and eXtreme Gradient Boosting (XGBoost) to improve differentiation and identification of skin disease. This integration combines the feature extraction capacity of deep learning in ANNs with the classification's gradient boosting efficiency of XGBoost. DermXNet tries to overcome the difficulties of diagnosing melanoma, non-melanoma and normal skin; it is a solution that is accurate, efficient and with interpretability. Specifically, DermXNet starts with an ANN module as the feature extractor and extractor. The ANN takes original or normalized image data through several layers of analysis in order to extract meaningful features the human eye is incapable of identifying plenty of details about skin tissue patterns, textures, and spatial relations necessary for the diagnosis of skin pathologies. This feature extraction mobile is important as it helps in presenting the image data in an embed size while focusing on the important information for classification. In the same way that ANN learning structure can help improve the performance of the DermXNet from a hierarchical point of view, the extracted features are confident and discriminative, capable to distinguish normal and diseased skin. In order to reduce the image features to their basic elements, the following pre-processed image taken through the ANN. Let the ANN be represented as a function f_{ANN} , parametrized by weights θ_{ANN} . The extracted feature, $F \in R^d$, are given by:

$$F = f_{ANN}(I_{preprocessed}; \theta_{ANN}) \quad (9)$$

Here F is the d - dimensional feature vector output by the ANN. Once the features

extracted, they got to the XGBoost module, which in fact behaves as a classifier. Due to the impressive efficiency of XGBoost gradient-boosting algorithm, the model demonstrates high accuracy in learning intricate decision boundaries and other forms of data scenarios. It builds a sequence of decision trees that continue to provide better classification by doing optimization over misclassified points. This component's capacity to bring structure and control to class distribution guarantees DermXNet does not plateau on more diverse data structures and on difficult cases. XGBoost, therefore, provides a systematic way of integrating with the inherently data-oriented ANN to yield high precision and recall and remains easily interpretable. The extracted features F used as input to the XGBoost classifier. Let $f_{XGBoost}$ represent the XGBoost function, parameterized by its set of decision trees τ . The probability of assigning a class label $y \in \{y_1, y_2, \dots, y_k\}$ is given by:

$$P(y|F) = f_{XGBoost}(F; \tau) \quad (10)$$

Where $\tau = \{T_1, T_2, \dots, T_n\}$ is the ensemble of n decision trees in XGBoost. In the case of DermXNet, the input images undergo resizing, normalization, and augmentation to make them invariant to issues such as size and color. Processing methods like grayscale conversion and artifact elimination (like inpainting for hair removal) pre-process the data and eliminate noise, enabling the ANN to recognize features that are helpful. The ANN then takes the pre-processed images and extracts from the features of the picture a high-dimensional feature vector of principles of the image. These features then classified by XGBoost, which uses them to come up with labels for skin conditions ranging from melanoma, non-melanoma, and normal skin condition. The final predicted class \hat{y} is the class with the highest probability:

$$\hat{y} = \arg \max_y P(y|F) \quad (11)$$

This makes having both a potential advantage for DermXNet in several ways. The hierarchical features learned by the ANN guarantees that the extracted data features are complex and comprehensive of the input data while, the modelling of nonlinear relationships and interaction by XGBoost improves on the classification. In addition, XGBoost features interpretability making it possible to explain features how they aid in such diagnoses, which is

very important in a medical diagnosis where understanding the decision-making process is vital. Altogether, these components form a system which does not oversimplify the relations between variables and which is less sensitive to noise and overfitting even if the amount of data is not very large. The model trained to minimize a composite loss function:

$$L = L_{ANN} + \lambda L_{XGBoost} \quad (12)$$

Where L_{ANN} is the loss function for training the ANN, $L_{XGBoost}$ is the loss function for XGBoost and λ is a regularization parameter that balances the contributions of the two components. Indeed, the integration of deep learning and gradient boosting used by DermXNet is a remarkable step forward in skin disease diagnosis. It gives useful possibilities to help dermatologists in making the correct diagnosis of skin diseases has high accuracy and reveals the thinking process of the specialists. This type of model not only solves the problems of the analyzation of medical images but also guarantees the focusing on the change of datasets and clinical requirements in later periods, which makes it a great asset in the healthcare industry. The overall DermXNet model can expressed as:

$$\hat{y} = f_{XGBoost}(f_{ANN}(f_{preprocess}(I); \theta_{ANN}); \tau) \quad (13)$$

3.6 Model Training Parameters

The training phases of the DermXNet model with the parameters, which are most suitable for the model and that, allow achieve successful training. The model trained for a 100 epochs; this enhances the opportunities of updating the weights and reducing the level of error. An epoch means one time when all samples of the training dataset are passed through the model and 100 epochs are just enough to make the model learn the best solutions and do not let the model memorize the data. The last used batch size of 32 that means the training data based on sample in groups of 32 processed simultaneously. This balance between the computational efficiency and gradient stability enables the model the produce a series of weight updates without overloading the memory resources, making this model particularly useful for larger datasets or when limited by hardware.

Algorithm: DermXNet

Input: Dataset $D = \{I_1, I_2, \dots, I_N\}$ (Images of skin conditions)

Labels $Y = \{y_1, y_2, \dots, y_N\}$

ANN weights θ_{ANN}

XGBoost parameters τ

Pre-processing constants T (threshold), kernel K

Output: Trained DermXNet Model (Capable of predicting skin condition for an input image I_{new})

Dataset Collection

Raw images D_{raw} from Kaggle and public sources

Split into melanoma ($D_{melanoma}$), non-melanoma ($D_{non-melanom}$), and normal skin (D_{normal})

Combined dataset $D = \{I_1, I_2, \dots, I_N\}$ with labels Y

Data Pre-processing

For each image $I \in D$

$I_r(x', y') = \text{Interpolate}(I(x, y))$ // Image Resizing

$I_g(x, y) = 0.2989 \cdot R(x, y) + 0.5870 \cdot G(x, y) + 0.1140 \cdot B(x, y)$ // Grayscale Conversion

$I_{blackhat} = (I_g \cdot K) - I_g$ // Morphological Filtering

$M(x, y) = \begin{cases} 255, & \text{if } I_{blackhat}(x, y) > T \\ 0, & \text{otherwise} \end{cases}$ // Thresholding

$I_{inpainted}(x, y) = \text{TELEA}(I_r(x, y), M(x, y))$ // Inpainting

Feature Extraction and Augmentation

$I_{augmented} = \text{Transform}(I_{normalized})$ // Apply Image Data Generator

$I_{normalized}(x, y, c) = \frac{I(x, y, c) - \mu_c}{\sigma_c}$ // Normalize pixel values

Train-Test Split

Split $D_{features}$ and Y

Model Training

$F = f_{ANN}(I; \theta_{ANN})$ // Train ANN

$P(y|F) = f_{XGBoost}(F; \tau)$ // Train XGBoost

$L = L_{ANN} + \lambda L_{XGBoost}$ // Combine ANN and XGBoost

Model Evaluation

For each $I \in D_{test}$

$F_{test} = f_{ANN}(I; \theta_{ANN})$ // Extract features

$\hat{y} = \arg \max_y P(y|F)$ // Prediction

Performance Metrics

Output

Predict class label \hat{y}

End Algorithm

The Adam optimizer known by the full name of Adaptive Moment Estimation used while optimizing the model weights. In Adam, we have adopted the merits of two standard optimization techniques viz., momentum and RMSProp while providing the learning rate that adapts itself and can converge faster. The reason is that Adam is especially useful for the deep learning tasks as it helps to work with the sparse gradients and noisy data. Specifically, the error with the predicted and actual labels computed by using the binary cross entropy loss function. Binary cross entropy is also suitable for the binary classification problem of segregating diseased and non-diseased skin tissues. Squared hinge loss is a measure of discrepancy between the probability distribution anticipated by the model and the actual binary labels, which helps the model, improve on producing probabilities that are more accurate.

3.7 Novelty of the Work

The novelty of this work lies in the design and implementation of DermXNet, a hybrid model that combines the strengths of Artificial Neural Networks (ANN) and eXtreme Gradient Boosting (XGBoost) for skin disease detection. While traditional deep learning models

excel at feature extraction, they often struggle with class imbalances and overfitting, especially in medical datasets. On the other hand, gradient-boosted models are effective classifiers but are limited in their ability to process raw image data. By merging these two approaches, DermXNet addresses these limitations, offering a unique architecture tailored for medical imaging tasks. A key advantage of the proposed model is its robust feature extraction capability using ANN. The ANN module captures complex patterns, textures, and spatial features from medical images, ensuring that the most relevant information retained. This further complemented by XGBoost, which provides robust classification by efficiently handling non-linear relationships and addressing issues such as class imbalances. This hybrid integration ensures better generalization, making DermXNet more reliable than standalone deep learning or machine learning models. Additionally, DermXNet's performance is enhanced by its comprehensive pre-processing pipeline, which includes artifact removal (e.g., hair inpainting) and augmentation, ensuring clean and standardized input data. These innovations collectively position DermXNet as a novel, practical, and highly effective solution for skin disease detection.

4. RESULTS AND DISCUSSIONS

The proposed DermXNet model has leveraged with the Python programming language in the Jupyter Notebook setup. Currently, python well known for its simplicity in coding friending flexibility in addition to having comprehensive library systems making it the most preferred programming language in machine learning and neural networks. Jupyter Notebook takes development to a new level by adding an interactive shell to the source code editor, which lets the user run the code conveniently, analyse the results with the help of adding plots right in the notebook and make changes in the code over and over again. These features make it a good fit for introducing and optimizing the DermXNet model, which is a model, which requires data pre-processing, model training, and post processing. The implementation conducted on a Windows operating system having Intel® Core™ i7 processor 14650HX 30 MB's cache maximum boost and turbo frequency 5.20 GHz. This efficient CPU guarantees fast performance of such operations as neural network training or feature extraction, or real-time data augmentation.

The GPU incorporated in the system has 6GB of RAM, which is not sufficient for most current deep learning tasks; however, computational memory management strategies such as batch processing during training avoid reaching the hardware's memory bounds. Additional libraries as TensorFlow and Keras optimize the utilization of resources, and make sure that the model will run properly in lower end machines that have less RAM for example. The approach for skin disease identification and diagnosis involves a standardized mode of mathematical operations that include data pre-processing, feature extraction, and classification for improved and sustainable performance. Every stage towards the construction of the model caters to the factors known to complicate the analysis of medical images and so optimise the model for action in an actual diagnosis.

The first step provides data pre-processing, and the input images are further refined for analysis. This common step is to shrink or expand the images so that all images can be in the same size as the rest of the data set, say 1024×1024. Grayscale conversion reduces complexity because it provides mainly necessary color data and retains important textural and structural characteristics. Filtering and inpainting are other forms of pre-processing that try to eliminate any items on the image that may distract the model from learning good features such as hair or shadows. It arises out of the need to eliminate noise and inconsistencies in the input data to guide the model directly target features that symbolize skin diseases. The current methodology steps it to feature extraction after data pre-processing which done by another subsection known as Artificial Neural Network (ANN) module. This ANN takes the input images through its layers of abstraction to learn powerful features of skin texture, lesion margin, and coloration. All these extracted features contribute to the basis of discriminating between melanoma, non-melanoma and normal skin condition in the proposed model. Figure 4 shows the data distribution in a dataset and Figure 5 displays the results of image pre-processing.

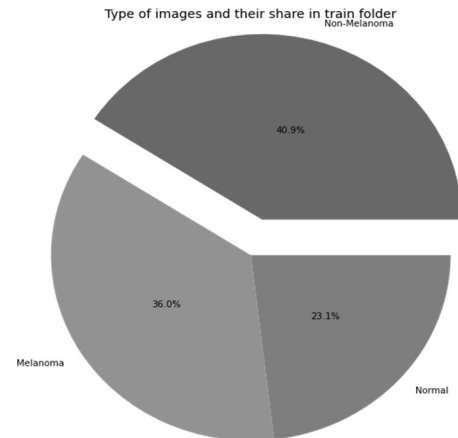


Figure 4 Data Distribution in Dataset




Figure 5 Image Pre-processing

Table 1 and Figure 6 gives a performance comparison of density with various deep learning model for a particular classification problem using important performance measures such as accuracy, precision, recall, F1-score and AUC-ROC. When it comes to standard architectures, Xception, which yielded 96.5% accuracy ranked top, followed by EfficientNetB0 which results in 94% accuracy. However, Custom CNN and the proposed DermXNet perform better than all the other techniques with accuracy of 97% and 98.38% respectively. In additional, quantitative evaluation proves that DermXNet yields higher precision (97.8%), recall (98%), F1-score (98.1%), and AUC-ROC (98.7%) as compared to the other models. On the other hand, it started models such as AlexNet and VGG16 show

comparatively lower efficiency in terms of accuracy, which stands at 87.5% and 88.2% respectively. The progressive changes across the models replicate the developmental changes in the architecture – deeper layers, feature extraction, and efficient use of parameters. The results show that DermXNet has better overall performance than previous methods emphasizing the fact that DermXNet has designed optimally for a particular dataset or a problem at hand and can be used to establish a new baseline in this domain.

preferred algorithm, the eXtreme Gradient Boosting (XGBoost), utilized. Handling of non-linear relationships and class imbalances makes XGBoost useful for this medical application. Based on the feature extracted from ANN, XGBoost has developed a cascade of decision trees, which sorts the image into its respective category. This interpolated form of combining ANN and XGBoost takes advantage of strength of either approach where deep learning learns intricate patterns while XGBoost is highly responsive and easy to interpret.

Table 1 Comparison of Various Model based on Performance Metrics

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC (%)
AlexNet	87.5	85.3	84.7	84.9	88.2
VGG16	88.2	86.5	85.8	86	89.1
InceptionV3	89.9	87.6	87.2	87.4	90.3
ResNet50	90.75	89	88.5	88.8	91.5
DenseNet121	92.3	90.2	90.1	90.3	92.7
MobileNet	93.1	91.5	91.2	91.4	93.8
EffNetB0	94	92.7	93	93.2	94.9
Xception	96.5	95.4	95.8	96	96.8
Custom CNN	97	96.1	96.5	96.8	97.3
Proposed DermXNet	98.38	97.8	98	98.1	98.7

Table 2 Training and Testing Time

Model	Training Time (s)	Testing Time (ms/image)
AlexNet	150	20.5
VGG16	140	19.8
InceptionV3	135	19
ResNet50	130	18.5
DenseNet121	125	17.8
MobileNet	120	17
EfficientNetB0	110	16.5
Xception	105	16
Custom CNN	100	15.8
Proposed DermXNet	95	15

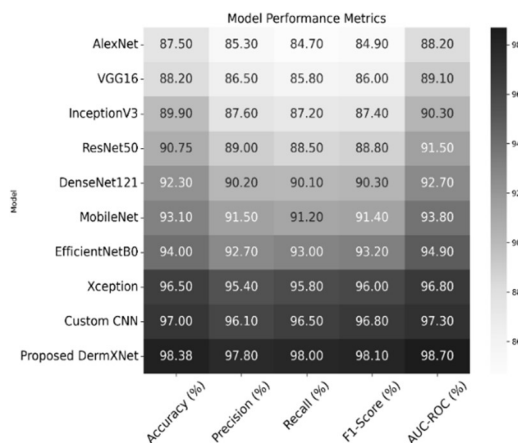


Figure 6 Model Performance Metrics

The extracted features then fed to the classification phase of the model and that a

Table 2 and Figure 7 shows the training and testing time taken by different models, which aims to highlight on the computational time. The architecture of the proposed DermXNet is also optimized and very efficient, as indicated by the shortest training time of 95s and among the shortest testing time per image as 15ms. Custom CNN comes closely second with 100 seconds training time and testing time of 15.8ms, which are rather proof of the model's simplified structure. When comparing to standard architectures, EfficientNetB0 and Xception had the lowest training time at 110 and 105sec, respectively, and less than 17ms per image in the test phase. Even the basic architectures such as AlexNet and VGG16 are much slower with training times of 150 and 140 seconds and a testing time of 20.5ms per image and 19.8ms per image respectively. The trend is due to recent developments in deep learning with modern architectures of deep learning requiring fewer FLOPs while succeeding in predictions that are

more accurate. Our results also indicate that as a classifier, DermXNet is highly efficient in terms of time and boasts a very high accuracy when compared with other methods, making it ideal for real-world applications. Figure 8 shows the train and validation accuracy. Figure 9 shows the train and validation accuracy

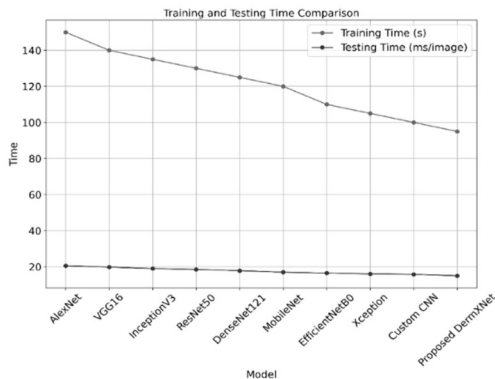


Figure 7 Training and Testing Time Comparison

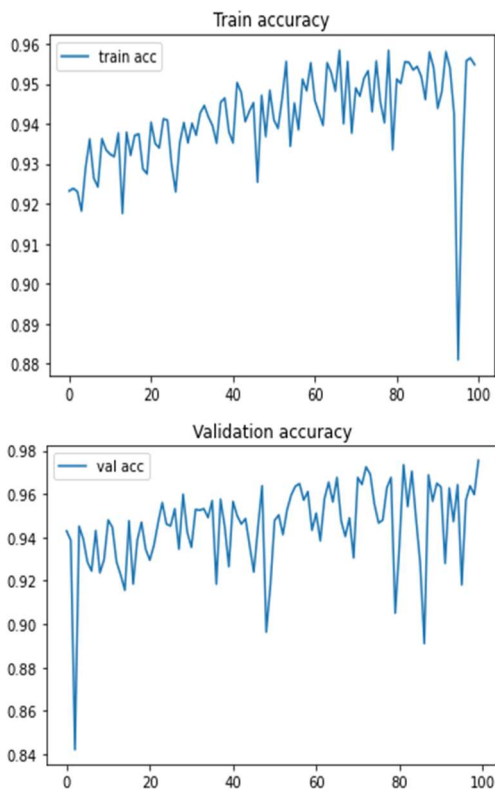


Figure 8 Training and Validation Accuracy

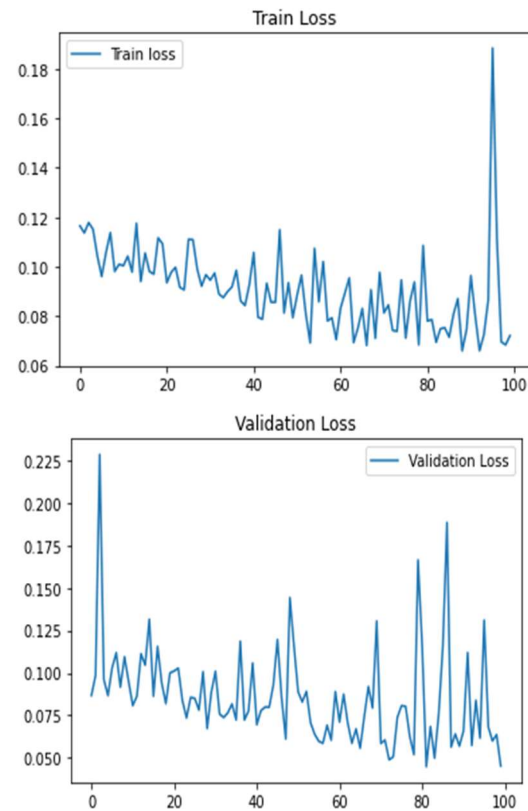


Figure 9 Training and Validation Loss

Table 3 Model Complexity (Number of Parameters)

Model	Parameters (Millions)
AlexNet	60
VGG16	138
InceptionV3	23.8
ResNet50	25.6
DenseNet121	8
MobileNet	4.2
EfficientNetB0	5.3
Xception	22.9
Custom CNN	6.5
Proposed DermXNet	4

Table 3 and Figure 10 shows the complexity of models depending on their size measured by the number of weights and the relationship between model size and performance. Among all, the proposed DermXNet has the least number of parameters, even 4 million that defines its relatively compact structure. Even this minimal complexity exalts DermXNet above its counterparts for accuracy and computational

efficiency as elicited in Table 1 and 2. MobileNet with 4,186,818 parameters and ENetB0 5,328,346 parameters show low complexity but reasonable performance thus is suitable for the resource-constrained case. Compared to AlexNet and VGG16, these models are much smaller in terms of parameters, at 60 million for AlexNet, and 138 million for VGG16, again due to their design of training on much larger parameter sets. Out of these models, InceptionV3, ResNet50 and Xception are comparatively more modern and efficient with parameters lying between 22.9 and 25.6 million. Custom CNN has moderate 6.5 million parameters making it reasonable. The design of DermXNet is quite simple and this fact confirms the aptitude of the model to achieve high performance with low computation load.

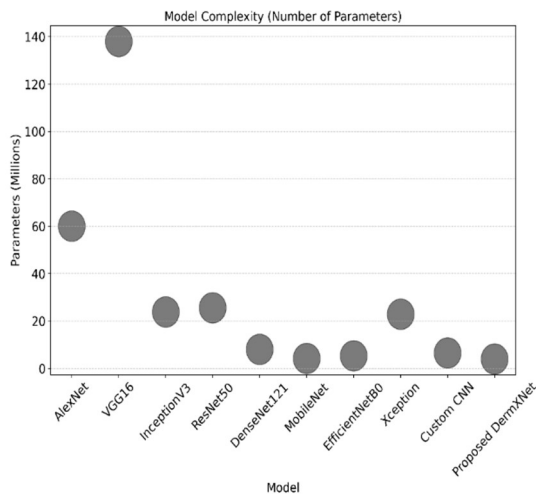


Figure 10 Model Complexity (Number of Parameters)

Table 4 Performance Metrics Comparison based Various Optimizer

Optimizer	Accuracy (%)	Final Loss
SGD	91.5	0.258
RMSprop	94.7	0.185
Adagrad	93.2	0.21
Adam	98.38	0.072
Adamax	96.4	0.135
Nadam	97.8	0.09

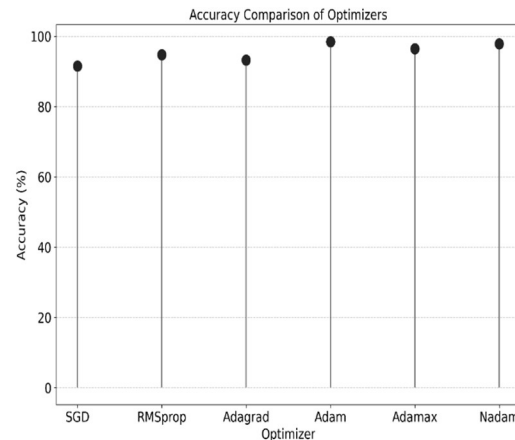


Figure 11 Accuracy Comparison of Optimizers

In

Table 4, Figure 11 and Figure 12 selected optimizers presented and compared in terms of accuracy and final loss. In terms of accuracy and loss, Adam can be seen as the leading optimizer of the tested ones with accuracy 0.9838 and minimal final loss 0.072, that proves the efficiency of the optimizer in achieving the best result as for precision as for convergence. Nadam almost performs similarly to Adam with 97.8% accuracy as well as final loss of 0.09 and thus it might be a good optimizer for those tasks, which need a smooth optimization. The other Adam variant, Adamax, gave slightly comparable performance with an accuracy of 96.4% and a final loss of 0.135. The basic versions of Ratios optimizers, SGD and Adagrad, demonstrate lower accuracy (91.5% and 93.2%), and higher loss (0.258 and 0.21), which points to the problem of slow convergence. RMSprop learn the accuracy of 94.7% with the loss of 0.185; however, the times are changed, and now we have many Adam-based optimizers. Overall, we can conclude that skills such as accuracy and speed of work confirm that optimizer is a key to the highest results for complex tasks, where Adam is successfully used.

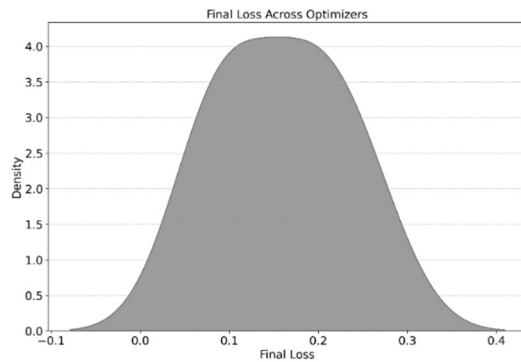


Figure 12 Final Loss across Optimizers

Table 5 Comparison Based on Various Activation Functions

Activation Function	Accuracy (%)	Final Loss
ReLU	96.5	0.098
Leaky ReLU	97.8	0.082
Sigmoid	92.3	0.214
Tanh	94	0.192
ELU	97.5	0.089
Swish	98.38	0.072

Table 5, Figure 13 and Figure 14 shows that different activation functions have an effect on the overall accuracy and final loss of a model, which determines its efficiency. Swish wins the accuracy battle with the highest accuracy of 98.38 % and the lowest final loss of 0.072 showing the best performance of recognizing complex patterns for great convergence. Leaky ReLU and ELU, which have been known to be robustly tackling the dying neuron problem, also yielded a very good accuracy of 97.8% and 97.5% with small final loss of 0.082 and 0.089 as well. These functions help to meet the drawbacks of the standard ReLU, for example, the “dying neuron” problem that in turn gives a better flow of gradients. ReLU one of the most used functions give very good results, it has 96.5% accuracy but the loss is slightly higher 0.098. Comparing with other traditional functions such Sigmoid and Tanh, their accuracy and final loss of are only 92.3 % and 0.214, 94 % and 0.192 respectively, confirms the problem of gradient saturation. In sum, the analysis indicates that Swish results in the most substantial improvement in model performance.

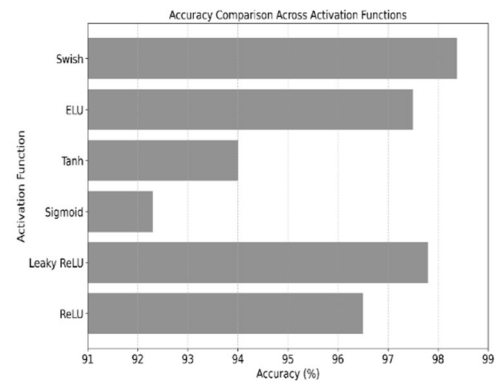


Figure 13 Accuracy Comparison across Activation Functions




Figure 14 Final Loss Distribution across Activation Functions

Table 6 Convergence Rate

Activation Function	Epochs to Reach 95% Accuracy
ReLU	30
Leaky ReLU	28
Sigmoid	50
Tanh	45
ELU	27
Swish	25

Table 6 and Figure 15 shows the Convergence of Related activation function of epoch 95% accuracy degree on the bases of epoch Swish shows the shortest convergence time, which is 25 epochs, and thus the algorithm proves its effectiveness when training deep learning models. ELU consequently comes next with the 27 epochs indicating the ability of the method to capture complex architectures effectively without causing instabilities of gradients. The next

activation function is the Leaky ReLU that produced a good result with convergence in 28 epochs, better than epoch 30 of the traditional ReLU algorithm. This shows that the function of Leaky ReLU to accept small gradients for negative inputs enhances learning. However, the conventional activation function such as Sigmoid and Tanh take much more time, 50 and 45 epochs respectively. These slower rates are because of its limitations; gradient saturation, and slow learning in deep networks. In this case, Swish converges very fast, has high accuracy, and lower loss; thus, making it highly suitable for training. Figure 16 shows the confusion matrix.

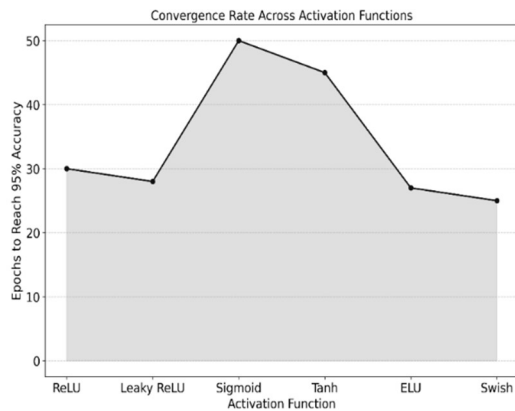


Figure 15 Convergence Rate across Activation Functions




Figure 16 Confusion Matrix

Table 7 Comparative Analysis of DermXNet vs. Existing Models (PMI Approach)

Aspect	Strengths of DermXNet	Weaknesses of DermXNet	Interesting Insights
Accuracy	Achieves 98.38%, outperforming CNN-based models.	Limited to binary classification (melanoma vs. non-melanoma).	Hybrid models can outperform deep CNNs with fewer parameters.
Computational Efficiency	Requires only 4M parameters, making it lightweight.	May need fine-tuning for real-world scalability.	Smaller models can be optimized for mobile applications.
Classification Robustness	Handles class imbalances better than CNNs.	Dataset-dependent; biased training data can affect performance.	Combining ANN with XGBoost enhances decision-making.

Training & Inference Time	Fast training (95s) & inference (15ms per image), ideal for real-time use.	May need further optimization for larger datasets.	Hybrid approaches can reduce overfitting while improving efficiency.
Clinical Interpretability	XGBoost improves classification stability.	Lacks built-in interpretability for medical practitioners.	Potential for explainable AI (XAI) integration in the future.

5. CONCLUSION AND FUTURE WORK

This work presented a hybrid deep learning model named DermXNet that ensures efficient and accurate skin disease detection through the integration of ANN and eXtreme Gradient Boosting (XGBoost). This being considered, class imbalance, overfitting, and efficiency in computation in medical imaging were all addressed effectively through the incorporation of deep ANN features into an XGBoost classifier. DermXNet attained a stunning accuracy of 98.38%, at a complexity of only 4 million parameters, and came out to be better than normal deep learning models such as ResNet50, VGG16, and Xception, and of course, in less costly computational resources. The fast training time (95s) and inference speed (15ms. per image), are attractive attributes for real-time applications concerned with mobile health and resource-limited clinical settings. DermXNet increases quality in inputs via different pre-processing strategies known to include artifact removal, image normalization, and augmentation that promise maintained high diagnostic performance. The findings demonstrate the prospects of hybrid AI models with respect to the early detection of skin cancer, telemedicine, and AI-based diagnosis, proffering a scalable solution for remote screening and clinical decision-making support. DermXNet will serve as a basis for future research into multi-class skin disease classification, incorporation of clinical metadata for improved interpretability, and deployment optimization on mobile and edge devices. In conjunction with Explainable AI (XAI) techniques, model transparency will help in improved acceptance in real-life health applications.

REFERENCES

- [1] Matheus B. Rocha, et al., "Skin cancer diagnosis using NIR spectroscopy data of skin lesions in vivo using machine learning algorithms", *BBE*, Volume 44, Issue 4, 2024, Pages 824-835, ISSN 0208-5216, DOI: 10.1016/j.bbe.2024.10.001
- [2] Ahmad Naeem, et al., "SNC_Net: Skin Cancer Detection by Integrating Handcrafted and Deep Learning-Based Features Using Dermo copy Images", *Mathematics*, 12(7), 2024, 1030, DOI: 10.3390/math12071030
- [3] A. Bindhu, et al., "Segmentation of skin cancer using Fuzzy U-network via deep learning", *MS*, Volume 26, 100677, 2024, ISSN 2665-9174, DOI: 10.1016/j.measen.2023.100677
- [4] Mohamad Abou Ali, et al., "Naturalize Revolution: Unprecedented AI-Driven Precision in Skin Cancer Classification Using Deep Learning", *BioMed Informatics* 4, 2024, no. 1: 638-660, DOI: 10.3390/biomedinformatics4010035
- [5] Li Zhang, et al., "A deep learning outline aimed at prompt skin cancer detection utilizing gated recurrent unit networks and improved orca predation algorithm", *BSPC*, Volume 90, 2024, 105858, ISSN 1746-8094, DOI: 10.1016/j.bspc.2023.105858
- [6] Suhendro Y. Irianto, et al., "Early Identification of Skin Cancer Using Region Growing Technique and a Deep Learning Algorithm", *HIJ*, Vol 5, 2024, No 3, DOI: 10.28991/HIJ-2024-05-03-07
- [7] HossamMagdyBalaha, et al., "An aseptic approach towards skin lesion localization and grading using deep learning and harris hawks optimization", *MTA* 83, 2024, 19787-19815, DOI: 10.1007/s11042-023-16201-3
- [8] TurkerTuncer, et al., "A lightweight deep convolutional neural network model for skin cancer image classification", *ASC*, Volume 162, 2024, 111794, ISSN 1568-4946, DOI: 10.1016/j.asoc.2024.111794

- [9] Tjahjamoorniarasih, et al., "Skin Cancer Classification from Dermatoscopy Images Using Deep Neural Network", *AJ*, Vol 16, Issue 1, 2024, p 245, DOI: 10.15849/IJASCA.240330.15
- [10] Mabrook S. Al-Rakhani, et al., "Effective Skin Cancer Diagnosis Through Federated Learning and Deep Convolutional Neural Networks", *AAI*, 2024, 38(1), DOI: 10.1080/08839514.2024.2364145
- [11] Omneya Attallah, et al., "Skin cancer classification leveraging multi-directional compact convolutional neural network ensembles and Gabor wavelets", *SR* 14, 2024, 20637, DOI: 10.1038/s41598-024-69954-8
- [12] G. Renith, et al., "Automated Skin Cancer Diagnosis and Localization Using Deep Reinforcement Learning", *IETE JR*, 2024, 70(4), 3631–3645, DOI: 10.1080/03772063.2023.2291805
- [13] Nasser A. AlSadhan, et al., "Skin Cancer Recognition Using Unified Deep Convolutional Neural Networks", *Cancers* 16, 2024, no. 7: 1246, DOI: 10.3390/cancers16071246
- [14] Anna M. SmakGregoor, et al., "An artificial intelligence based app for skin cancer detection evaluated in a population based setting", *NPJDM*, 2023, 6, 90, DOI: 10.1038/s41746-023-00831-w
- [15] Vipin Venugopal, et al., "A deep neural network using modified EfficientNet for skin cancer detection in dermoscopic images", Volume 8, 2023, 100278, ISSN 2772-6622, DOI: 10.1016/j.dajour.2023.100278
- [16] Tehseen Mazhar, et al., "The Role of Machine Learning and Deep Learning Approaches for the Detection of Skin Cancer", *Healthcare* 11, 2023, no. 3: 415, DOI: 10.3390/healthcare11030415
- [17] Noor ul Huda, et al., "Skin Cancer Malignancy Classification and Segmentation Using Machine Learning Algorithms", *JOM* 75, 2023, 3121–3135, DOI: 10.1007/s11837-023-05856-w
- [18] Sumit Kumar Singh, et al., "Fuzzy Logic with Deep Learning for Detection of Skin Cancer", *AS*, 13(15), 2023, 8927, DOI: 10.3390/app13158927
- [19] Jitendra V. Tembhurne, et al., "Skin cancer detection using ensemble of machine learning and deep learning techniques", *MTA* 82, 2023, 27501–27524, DOI: 10.1007/s11042-023-14697-3
- [20] Jaisakthi S M, et al., "Classification of skin cancer from dermoscopic images using deep neural network architectures", *MTA* 82, 2023, 15763–15778, DOI: 10.1007/s11042-022-13847-3