AI-TBP: Revolutionizing Skin Cancer Detection with Hybrid AI and TBP Imaging

D Bhargava Sai Rohith[1], S Nagendra Kumar Reddy[2], N Venkata Dinesh Naidu[3] and Jawad Ahmad Dar[4]*
{bhargavasairohithdonthu@gmail.com[1], syamalanagendrareddy@gmail.com[2], dineshnaidunarni@gmail.com[3], jawadsirphysics@gmail.com[4]}

Department of Advanced Computer Science and Engineering, Vignan's Foundation for Science, Technology & Research (Deemed to be University), Vadlamudi, Guntur (Dt), 522213, Andhra Pradesh, India[1, 2, 3, 4]

Abstract. Skin cancer continues to be a significant health issue worldwide, requiring precise and effective diagnostic methods. Our research introduces a unique Hybrid Deep Learning architecture that combines three versions of the Residual Network (ResNet) ResNet-18, ResNet-34, and ResNet-50 to categorize skin lesions into eight separate types utilizing the ISIC 2019 dataset. To better leverage model performance, we deliver a comprehensive preprocessing pipeline including noise reduction, normalization, class balanced and data augmented respectively personalized for the characteristics of dermoscopic images. The proposed architecture leverages the specific benefits of each ResNet model by combining their feature representations before being fed to a customized classifier. Addressing class imbalance using a focal loss function, the model performs well on different lesion types with an overall accuracy of 91%. The proposed approach in this paper is both scalable and interpretable and hence it paves a way for medical image analysis in the future.

Keywords: Hybrid ResNet Architecture, Skin Cancer Classification, ISIC 2019 Dataset, Multi-Stage Preprocessing, Focal Loss Optimization

1 Introduction

Rising worldwide incidence of melanoma has raised the requirement of precise diagnostic methods, as the early diagnosis can help in successful treatment and survival. Dermoscopy has increasingly improved the diagnostic accuracy of skin lesions; dermoscopic assessment plays an important role in routine practice, providing valuable information to the unaided eye. The imaging is complicated and dependent on the ability of dermatologists so that it results diversity and slowing down. In the last decade, deep learning has reshaped medical imaging, and the application of convolutional neural networks [13] (CNNs) has shown great promise for the automation of skin lesion classification. CNNs are a classification of deep learning algorithm belonging to the broader family of deep neural networks. They not only identify descriptive elements from the input data via a procedure of layer-wise training. One key component of a CNN: The term 'convolution' in CNN refers to the mathematical operation of matrix multiplication. An ordinary CNN is composed of convolutional layer, non-linearity layer, pooling layer and fully connected layer, etc [2].

Routine diagnostic criteria are mainly based on clinical exploration and dermoscopic evaluation by a dermatologist, which are subjective and slow. As such, there is a growing need for

ICITSM-Part I 2025, April 28-29, Tiruchengode, India

DOI 10.4108/eai.28-4-2025.2357787

purposeful machine diagnostic systems that are able to help health workers in the classification of unlabeled skin lesions with high accuracy and reliability [17].

Methodology This work targets [12] creating a novel hybrid CNN-SVM model for skin cancer detection, where a hybrid architecture is proposed to process and analyze the ISIC 2019 dataset. The objective of the proposed method is to enhance lesion classification performance and to handle problems such as class imbalance and feature selectiveness. Using advanced data preprocessing, transfer learning, and hyperparameter search, our goal with this work is to contribute to the current work on AI-driven dermatological diagnostics. This study's results may have significant implications for early detection of skin cancer and could contribute to timely diagnosis and reduced load on the health system.

Many studies have been conducted on the detection and classification of melanoma using different methods. In 2012, an SVM based Model for melanoma cancer detection achieved an accuracy of 0.77 [7] Next year, research for melanoma early Detection using k-nearest neighbor (KNN), neural network (NN) and SVM reported accuracy range of 0.8-0.9 [19] In 2017, another research investigated the classification of melanoma through the popular method like KNN, decision trees (DT) and SVM etc. with achieved accuracy 0.78 [16] Similarly, another work of 2017 by ResNet-88 on the ISIC 2017 dataset achieved an accuracy of 0.823 [14] In addition, In 2017, a research work using statistical method, named Bayesian classifier method by presented an accuracy of 0.91 [20] while a research done using SVM on the data from Pedro Hispano Hospital achieved an accuracy of 0.93 [11].

This paper tries to create a state-of-the-art skin lesion classification system with the ISIC 2019 dataset that has approximately 25,331 images in eight classes. We suggest a hybrid approach that uses three variants of ResNet, ResNet-18, ResNet-34, and ResNet-50—to extract multi-depth features from dermoscopic images. We also suggest a multi-stage preprocessing pipeline that tries to increase image quality and mitigate dataset-specific problems such as noise and class imbalance. The research question is: Can a hybrid ResNet approach with a strong preprocessing method improve classification accuracy and robustness for the eight-class skin cancer task in the ISIC 2019 dataset compared to individual CNN approaches?

Our inspiration is to bridge the gap between computer-aided diagnostics and clinical dermatology and provide a highly efficient and accurate computational tool. By taking advantage of ResNet-18's shallow feature extraction, ResNet-34's moderate level of depth, and ResNet-50's dense representational capability, we seek to develop a well-balanced model for heterogeneous lesion detection. The preprocessing pipeline—noise filtering, normalization, class weighting, and augmentation—optimizes the ISIC 2019 dataset's utility and ensures good generalization over imbalanced classes. This work introduces: (1) a new hybrid ResNet architecture for multi-scale feature extraction, (2) a dermoscopic image preprocessing method, and (3) evidence of enhanced classification performance through rigorous testing. These advances enable computer-aided diagnostic systems to support clinicians in accurate and timely skin cancer diagnosis.

2 Literature Survey

Skin cancer classification using deep learning has seen widespread interest over recent years owing to the sheer amounts of available dermoscopic image databases and the capability of enhancing the decision-making clinical processes. Through this section, we summarize work

conducted in this space previously with regard to highlighting CNN architectures, ensemble and combination techniques, as well as preprocessing techniques, and underlining the importance of those for the ISIC 2019 dataset as well as the eight-class skin cancer classification challenge.

The work in [1] introduces a new framework for automated skin disease diagnosis using digital hair removal via morphological filtering and inpainting, followed by Gaussian filtering and Grab cut segmentation for preprocessed clinical images. It uses GLCM and statistical methods to extract the features, which are then classified using DT, SVM, and KNN. It efficiently recognizes multiple skin diseases with SVM proving to be the best performer on the ISIC 2019 and HAM10000 datasets. The method enhances objectivity and diagnostic speed over the limitations of human expert assessment relative to sophisticated modern methods.

This paper [3] uses a CNN based architecture and Residual network-50 on HAM10000 dataset which contains unbalanced data classes. To address this issue, they used augmentation techniques and their model got accuracy of 86% for CNN and 85.3% for ResNet-50 model.

Anand et al. [4] used approximately 3297 images from the ISIC dataset to use transfer learning techniques. Their model structure consisted of two dense layers with LeakyReLU activation, then a dense layer with sigmoid activation and a flatten layer by inserting VGG-16. With a training set size of 128 and training for more than 10 epochs, they obtained an accuracy of 89.09%

The study of this paper [5] uses a Conventional Neural Network model with dataset consisting of 12,378 skin lesion images from the ISIC dataset. They tested a set of 100 images with dermatologists and got a sensitivity of 74.1% while CNN model got 86.5% which have given good results compared with dermatologists.

This study [6] aims to classify melanoma skin cancer using Convolutional Neural Networks (CNN) with the ResNet architectures on the ISIC 2018 dataset. A range of ResNet variations (including ResNet-50, ResNet-40, ResNet-25, ResNet-10, and ResNet-7) were trained using data augmentation techniques and under sampling, and their performance was evaluated using the F1 score. The ResNet-50 model without augmentation achieved the best performance, having a validation accuracy of 0.83 and an F1 Score of 0.46.

Hossain et al. [9] employed about 6599 images from the ISIC dataset on Kaggle to compare various Residual Network variants. Their research yielded accuracies of 89.65%, 89.09%, 88.78%, and 86.34% for ResNet152, ResNet101, ResNet50, and ResNet18 respectively.

Dermatological diagnosis [10] is difficult because it involves a complex nature and reliance on extensive testing and practitioner skill. To counter this, the present study suggests a machine learning-based image-based method of automatic classification. Skin images are analyzed, smoothed, and classified through Convolutional Neural Networks (CNN) combined with a SoftMax classifier. The technique improves accuracy as well as efficiency, making it a reliable source for diagnosis as well as medical training.

Zhang et al. [21] compared two neural network structures, ResNet-50 and Inception-V3, on binary classification with the Kaggle ISIC archive dataset, consisting of 2637 training images and 660 test images. Their results indicate that ResNet-50 performs better than Inception-V3 with an accuracy of 88.83% versus 83.17% for Inception-V3.

Bechelli et al. [18] used machine learning and deep learning methods on the ISIC dataset of around 3297 images. Of all the models that were employed, ResNet50, VGG16, and Xception provided an accuracy rate of about 88%, while other models did not even touch 72%.

3 Dataset

The basis of this study is the International Skin Imaging Collaboration (ISIC) 2019 dataset, a representative set of dermoscopic images (see Fig [1]) designed to promote research into computer-aided skin cancer diagnosis. Provided as part of the ISIC 2019 Challenge, this dataset comprises 25,331 quality images, each of a lesion captured under controlled dermoscopic conditions. The images are labeled with ground-truth annotations across eight different classes, representing a various range of skin types found in clinical practice. The ISIC dataset contains eight classes: actinic keratosis (AK), basal cell carcinoma (BCC), benign keratosis (BKL), dermatofibroma (DF), melanoma (MEL), melanocytic nevus (NV), squamous cell carcinoma (SCC), and vascular lesions (VASC). The multi-class nature of ISIC 2019 differentiates it from previous datasets such as ISIC 2018, which had fewer categories and thus presents a robust testbed for assessing classification models across a wide range of diagnostic conditions.

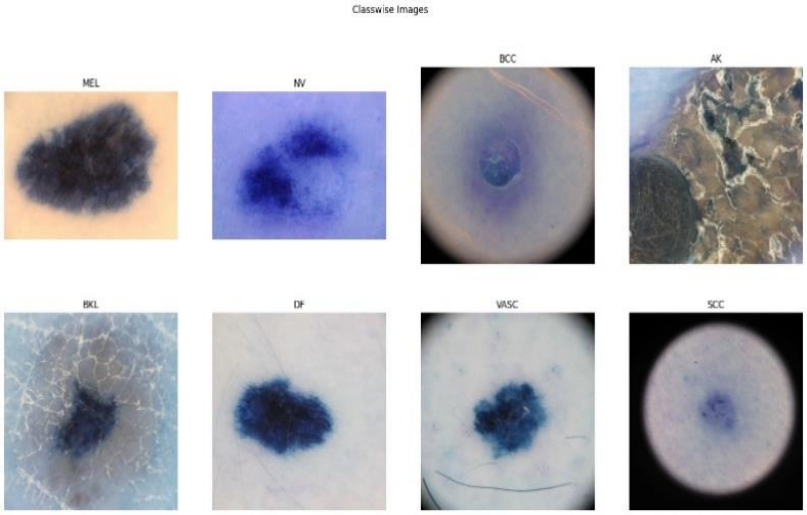

Fig. 1. Class Wise Images.

The images are divided among these classes with remarkable imbalance, reflecting patterns of real-world prevalence but posing difficulties for machine learning models. In particular, the dataset includes 4,133 images of AK, 3,323 of BCC, 2,624 of BKL, 4,761 of DF, 4,522 of MEL, 5,000 of NV, 4,372 of SCC, and 4,747 of VASC, for a total of 25,331 samples. Melanocytic nevus (NV) at 5,000 samples is the most prevalent category, and benign keratosis (BKL) at 2,624 samples is the least, exhibiting a class imbalance that will require careful handling in model training. Each image is rendered in consistent format, normally resized to 224x224 pixels to accommodate popular CNN architectures such as ResNet, and RGB color space to maintain the rich visual cues such as vascular structures and pigmentation vital to distinguishing between lesions. The dermoscopic photographs in the ISIC 2019 archive are captured at different institutions worldwide, with diversity of imaging hardware, patient populations, and lesion

types. Such diversity makes the archive more representative, but introduces variability in features like illumination, noise (e.g., hair artifacts), and resolution, which mask critical diagnostic features. To address these challenges, our contribution employs a custom preprocessing pipeline, detailed in the methodology, to normalize inputs and remove artifacts. The archive is split into training and validation subsets, with approximately 80% (20,264 images) set aside for training and 20% (5,067 images) for validation, slightly tuned to fit batch processing requirements (e.g., batch size of 32). This split enables stable model optimization and performance evaluation, so that the hybrid ResNet architecture generalizes well over the eight-class challenge.

With the ISIC 2019 dataset, this research tackles a clinically relevant issue: high-accuracy classification of a high diversity of skin lesions under conditions of intrinsic imbalances and image variability. The dataset size and diversity provide a great test bed for our hybrid method, incorporating ResNet-18, ResNet-34, and ResNet-50, and for the verification of the effectiveness of our preprocessing method in enhancing diagnostic accuracy.

4 Methodology

This section delineates the systematic approach employed to classify skin lesions into eight categories using the ISIC 2019 dataset with a hybrid ResNet model. The methodology is structured into four primary components (see Fig [2]): preprocessing pipeline (A), augmentation (B), split of data (C), and model architecture (D). Each component is designed to address the inherent complexities of dermoscopic imagery, such as noise, variability, and class imbalance, culminating in a robust framework for skin cancer classification.

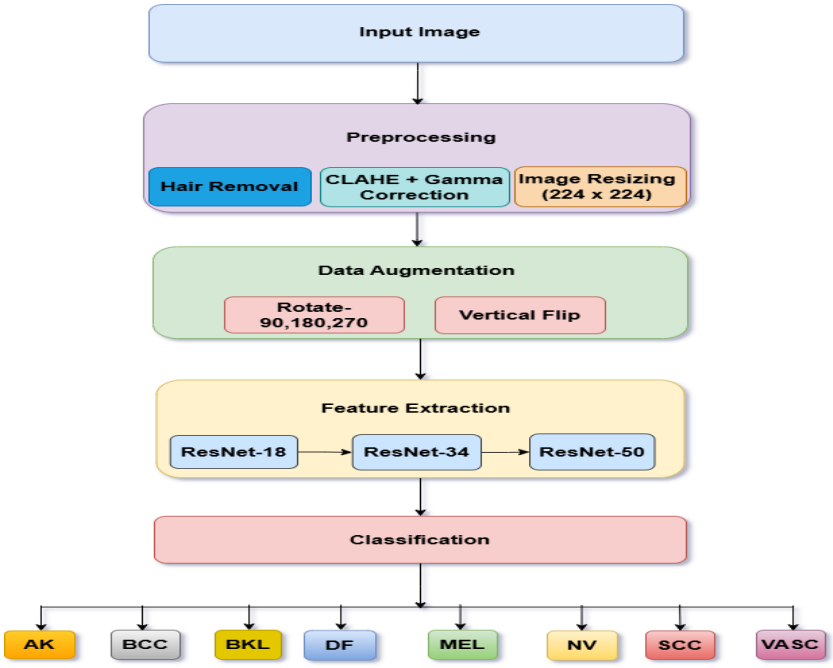

Fig. 2. Proposed Methodology.

4.1 Preprocessing

The preprocessing step concentrated on improving the quality and consistency of the ISIC 2019 dataset, which consists of dermoscopic images of skin lesions belonging to eight categories AK, BCC, BKL, DF, MEL, NV, SCC, and VASC, complemented by information from 3D Total Body Photography (TBP). These images were obtained from the ISIC archive (https://www.isic-archive.com, accessed March 2025). Three major sub-steps were utilized in the preprocessing workflow to ready the data for model training:

4.1.1 Hair Removal

Dermoscopic images frequently include hair artifacts that hide lesion characteristics, making it essential to remove them for enhanced diagnostic precision. This was accomplished through black hat morphology utilizing a 17×17 structuring element to improve dark hair visibility against the image backdrop, followed by inpainting employing the Tele algorithm (based on diffusion) to restore the impacted areas. The procedure successfully removed hair (see Fig [3a] and Fig [3b]) while maintaining the edges of lesions, confirmed by visually inspecting a selection of images.

Fig. 3(a). With Hair.

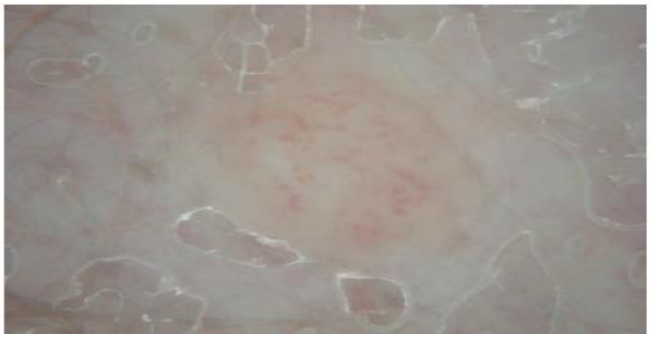

Fig. 3(b). Without Hair.

Fig. 3(a), 3(b). Images before and after Hair Removal.

4.1.2 Standardization

The process of image standardization (see Fig [4a] and Fig [4b]) begins with the conversion of the input image from its original color space format to a color space based on lightness. The image is then divided into its lightness and two channels of color. The lightness channel is augmented with a contrast adjustment method that restricts the over-amplification and affects small, local areas. Upon this augmentation, the adjusted channel of lightness is combined once again with the original color channels. Next, the image is converted back to its original color space. Then a gamma correction of 1.2 is applied to adjust brightness and contrast of the image. And then an inverse gamma value is calculated and a lookup table is created to adjust non-linear pixel intensities to improve the overall visual quality of an image. This table is then utilized to create the final standard image.

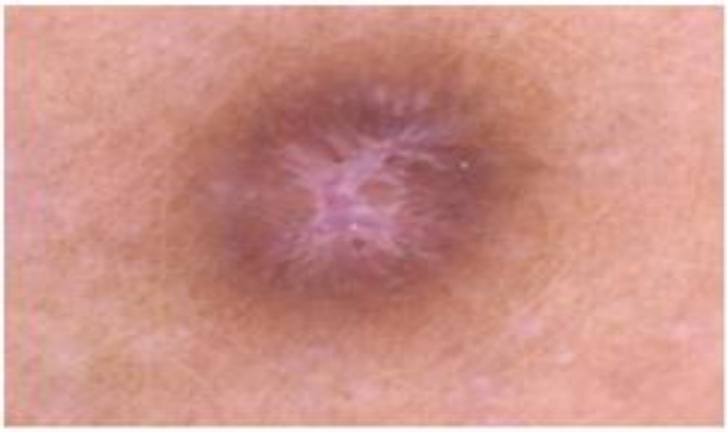

Fig. 4(a). Original Image.

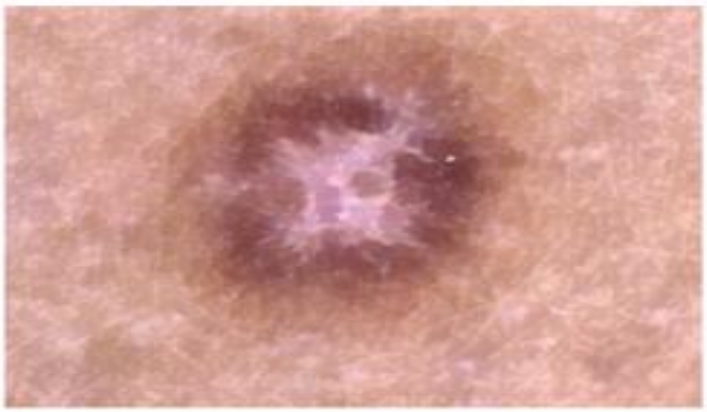

Fig. 4(b). Standardized Image.

Fig. 4(a), 4(b). Images before and after Standardized process.

4.1.3 Resizing

Uniform input dimensions are crucial for convolutional neural networks (CNNs). All images were resized to 224×224 pixels (see Fig [5a] and Fig [5b]) via bicubic interpolation to maintain image quality and ensure compatibility with the model structure. This type of pre-processing step like resolution is commonly used in pre-trained models like ResNet, to make a balance between computational efficiency and preserving lesion details. After this step, images were saved as a 4D array with dimensions (number _of _ samples, 224, 224, 3) for next processing.

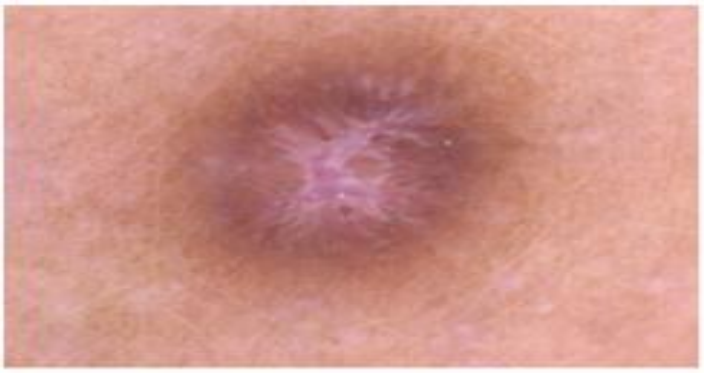

Fig. 5(a). Original Image**.**

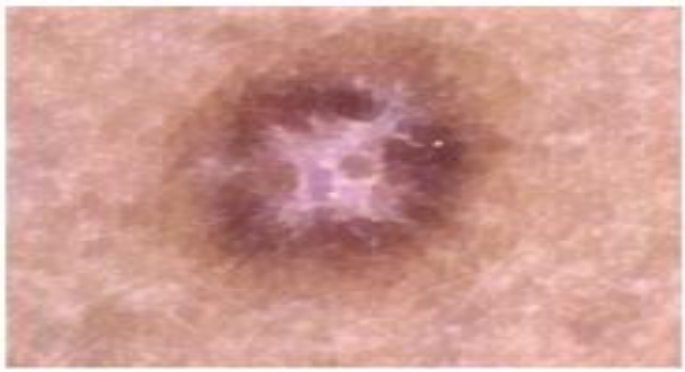

Fig. 5(b). Resized Image.

Fig. 5(a), 5(b). Images before and after Resizing process**.**

bicubic interpolation was favored over more straightforward techniques such as nearest-neighbor interpolation to reduce artifacts and preserve edge sharpness, which is vital for detecting lesion boundaries. The reason behind choosing the 224×224-pixel resolution is to support hardware acceleration, and improving training efficiency on GPU systems. After completion of resizing a random sample is taken to validate the minimal distortion and retention of lesion characteristics. This results in a uniform input format by facilitating effective model transfer across the varied ISIC 2019 dataset.

4.2 Augmentation

As the ISIC dataset contains imbalanced data classes we come across overfitting issue which would become a problem in inference phase to overcome this issue technique like augmentation is used to make dataset balanced. The dataset showed considerable imbalance, with class sizes varying from 12,875 samples for Nevus (NV) to 1,195 samples for Dermatofibroma (DF). To address this, augmentation methods (see Fig [6]) were utilized for minority classes, whereas majority classes were reduced to a target of 5,000 samples per class through random selection without replacement. The augmentation process included the following techniques:

4.2.1 Rotations

Images were rotated by 90°, 180°, and 270° to make angular variations in the samples and replicating the diverse orientations of skin lesions observed in clinical settings.

4.2.2 Flips

Horizontal and vertical mirroring was applied to the images which increases diversity by mimicking variations in lesion appearance due to patient positioning.

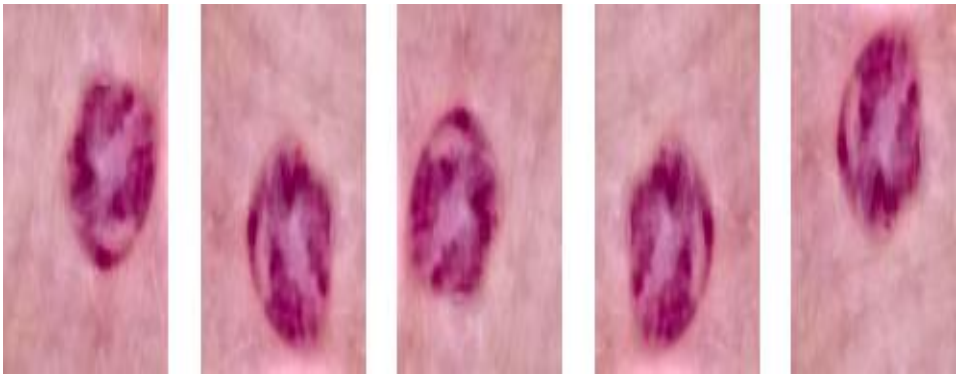

Fig. 6. Detailed View of Images in ISIC-2019 After Augmentation.

4.2.3 Optional Techniques

Additional approaches like color jitter (to change brightness, contrast, and saturation) and elastic deformations (to mimic tissue distortion) were evaluated to increase dataset variability, although their use was restricted to maintain computational efficiency. Minority classes were increased to achieve the goal of 5,000 samples for each class, leading to a balanced dataset of around 40,000 samples (5,000 samples for each of eight classes). The distribution (see Table [1]) of the classes was examined via frequency analysis to confirm uniformity, which is essential for avoiding model bias towards majority classes and enhancing generalization across all categories of skin lesions.

Table 1. Detailed View of Images in ISIC-2019 After Augmentation.

Class Name	Before Augmentation	After Augmentation
AK	867	4133
BCC	3323	3323
BKL	2624	2624
DF	239	4761
MEL	4522	4522
NV	12875	5000
SCC	628	4372
VASC	253	4747

4.3 Dataset Splitting

The skin lesion images that contain in ISIC 2019 dataset was balanced and preprocessed because we have to give equal priority among all eight diagnostic classes. After balancing of data resulting in a training set of 80% and a testing set of 20% of the total samples.

4.4 Proposed Framework

The suggested framework is a hybrid convolutional neural network (CNN) aimed at categorizing dermoscopic images from the ISIC 2019 data set into eight skin lesion types, utilizing the combined advantages of three versions of the Residual Network (ResNet), ResNet-18, ResNet-34, and ResNet-50. This architecture overcomes the constraints of single-model CNNs by integrating multi-depth feature extraction, allowing for the identification of a wide range of lesion traits from subtle textural nuances to intricate structural patterns while ensuring computational efficiency. The system combines these pre-trained models with a specialized classifier, fine-tuned using a customized training approach to address the class imbalance and diversity of the data set, providing a strong solution for the automated detection of skin cancer.

Table 2. Comparison of ResNet Models.

Model	Number of Layers	Attributes/Characteristics
ResNet-18	18	512
ResNet-34	34	512
ResNet-50	50	2048

4.4.1 ResNet-18

18-layer architecture [8], which is powerful, originally trained on ImageNet, produces a 512-dimensional feature vector after completing its final convolutional layer. It excels in extracting fundamental features like edges and textures, making it appropriate for recognizing simpler patterns in lesions, such as vascular lesions (VASC).

4.4.2 ResNet-34

A moderately deep, 34-layer model, which is also pretrained on ImageNet, generates a 512-dimensional feature vector. Also, increased depth helps us to capture mid-level features such as pigmentation gradients, which are crucial to distinguish benign keratosis (BKL) from melanoma (MEL).

4.4.3 ResNet-50

A 50-layer model with bottleneck blocks, which is trained on ImageNet that generates a 2048-dimensional feature vector. Its more complex architecture and broader feature space capture fine details like uneven borders or asymmetry, which are important to identifying complex cases such as squamous cell carcinoma (SCC). In each backbone, the original fully connected (fc) layer is replaced with an identity layer, preserving the raw feature maps for further processing. The strong general image recognition capabilities are provided by pretrained weights. Which are then fine-tuned to analyze dermoscopic images during training

4.4.4 Feature Concatenation

The feature vectors which are from ResNet-18 (512), ResNet-34 (512), and ResNet-50 (2048) combined the channel dimension to create a single 3072-dimensional representation (512 + 512 + 2048) for each input image (see Fig 7 and Table 2). This combination utilizes the varied viewpoints of the three models: ResNet-18's effectiveness, ResNet-34's equilibrium, and ResNet-50's profundity, forming a thorough feature set that captures both local and global characteristics of lesions.

4.4.5 Custom Classifier

The combined features are handled via a sequential classification module: A linear layer compresses 3072-dimensional input into 512 dimensions, succeeded by a ReLU activation to add non-linearity and improve feature discrimination. A dropout layer set to a probability of 0.5 helps regularize the network, reducing overfitting by randomly turning off neurons during training, which is especially crucial because of the dataset's moderate size (25,331 images). A concluding linear layer transforms the 512-dimensional features into 8 output logits that represent the ISIC 2019 classes.

The structure includes around 57.5 million parameters, representing the total complexity of the three ResNet backbones alongside the classifier (1.6M). (See Table 3) Comparison of Models Used for Skin Cancer Classification. Regardless of this scale, utilizing pretrained weights and a targeted classifier design guarantees computational efficiency when compared to training a model of similar size from the ground up**.**

Table 3. Comparison of Models Used for Skin Cancer Classification.

Model	Parameters	Input Size	Training Details
Custom CNN	0.6 million	224×224	Trained from scratch, Adam optimizer (lr=1e-3), early stopping, cross-entropy loss.
ResNet-50	25 million	224×224	Fine-tuned, Adam optimizer (lr=1e-3), early stopping, cross-entropy loss.
Hybrid CNN-ResNet	25.6 million	224×224	Joint training, Adam optimizer (lr=1e-3), early stopping, cross-entropy loss.

4.4.6 Loss Function

Focal Loss [15] is used to give more importance to hard-to-classify cases, which helps to address the class imbalance present in the ISIC 2019 dataset. It is defined as:

$$F\,L(p_t) = -(1 - p_t)\gamma \log(p_t) \qquad (1)$$

where p_t is the predicted probability of the true class, and $\gamma = 2.0$ helps focus the model more on difficult examples by reducing the influence of easy ones. To further handle class imbalance, class weights (α) are introduced. These weights are used to calculate as the inverse of the class frequencies (e.g., 1/5000 for NV, 1/2624 for BKL) and normalized so that their sum equals 8, ensuring all eight classes contribute equally to the training without needing to oversampling of data.

4.4.7 Optimizer

The model was trained using Stochastic Gradient Descent (SGD) with a learning rate of 0.01, momentum set to 0.9, and a weight decay factor of 1×10^{-4}. This configuration balances convergence speed and regularization, enabling efficient training of the hybrid ResNet architecture while preventing overfitting on the ISIC 2019 dataset.

4.4.8 Scheduler

The Reduce LR On Plateau scheduler lowers the learning rate by a factor of 0.1 when the validation loss remains unchanged for three consecutive epochs, enables adaptive optimization.

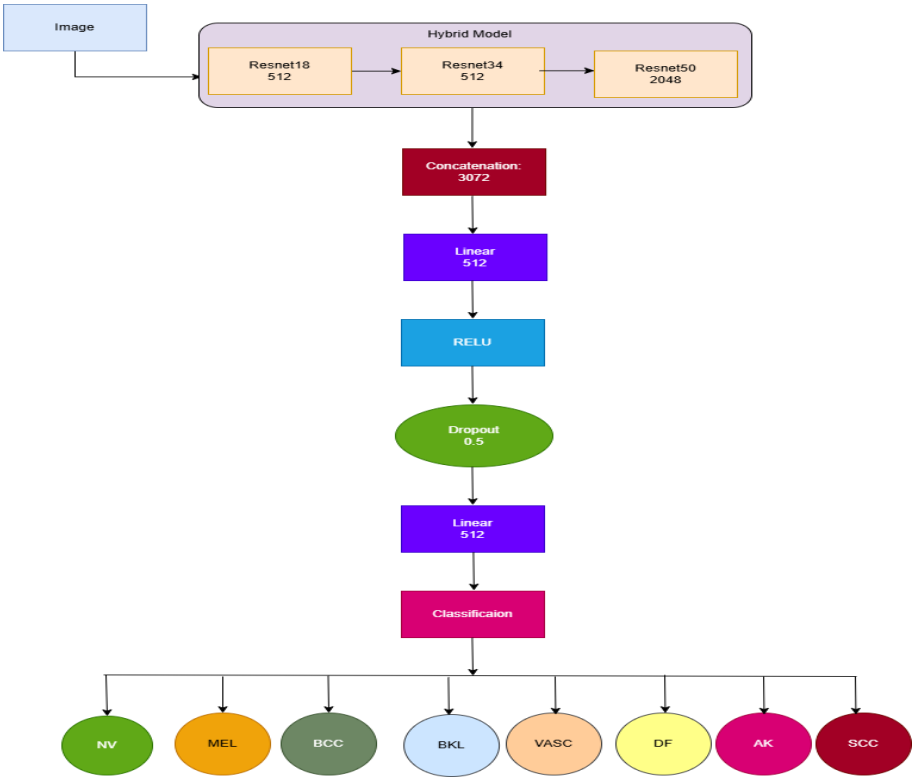

Fig. 7. Model Architecture.

5 Results

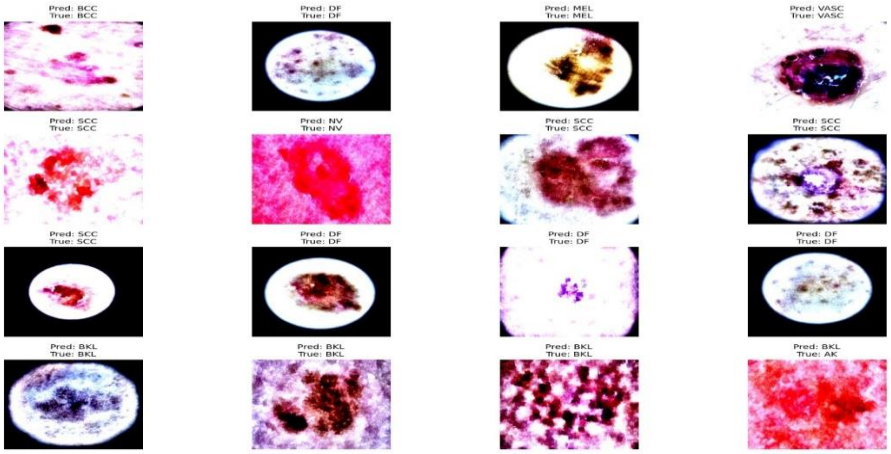

Fig. 8. Actual vs Predicted.

The suggested model for classifying skin diseases was assessed with a dataset of 6697 test samples, resulting in an overall accuracy of 0.91. (See Fig 8) Actual vs Predicted. The model exhibited impressive results across eight categories: actinic keratosis (AK) with precision 0.92, recall at 0.96, and F1-score 0.94; basal cell carcinoma (BCC) scoring 0.87, 0.85, and 0.86; benign kerato- sis (BKL) with values of 0.85, 0.77, and 0.81; dermatofibroma (DF) achieving 0.99, 1.00, and 1.00; melanoma (MEL) at 0.72, 0.70, and 0.71; melanocytic nevus (NV) with 0.77, 0.77, and 0.77; squamous cell carcinoma (SCC) scoring 0.97, 0.99, and 0.98; and vascular lesion (VASC) attaining 0.99, 1.00, and 1.00. The macro average for precision, recall, and F1-score across all classes stood at 0.88, while the weighted average achieved 0.91, indicating strong generalization across diverse class distributions, with support varying from 973 (MEL) to 19,08 (VASC). (see Table 4) Performance Metrics. These findings suggest that the model is effective, especially for classes that are well-represented, although there is potential for enhancement in melanoma Classification.

Table 4. Performance Metrics.

S. No	Paper Validation	Accuracy
1	[3]	85.3%
2	[4]	89.09%
3	[9]	89.65%
4	[21]	88.83%
5	[18]	88%
6	Proposed	91%

5.1 Confusion Matrix Analysis

The confusion matrix (see Fig 9), using the 5,056-image test set, delivers an accurate assessment of the performance of the hybrid ResNet model for eight skin lesion classes. Di-agonal represents high accuracy, with dermatofibroma (DF) at 1,917, squamous cell carcinoma (SCC) at 1,721, and vascular lesion (VASC) at 1,908 correct predictions, whereas actinic keratosis (AK) and basal cell carcinoma (BCC) deliver 1,504 and 892, respectively. Melanoma (MEL) and melanocytic nevus (NV) are less accurate with 677 and 787 correct classifications, respectively, off-diagonal errors (e.g., 154 MEL incorrectly classified as NV, 72 NV as MEL) illustrating difficulty due to similarities in pigmentation. The analysis indicates that the model's strength lies in identifying well-separated classes (DF, SCC, VASC) and that melanoma is a critical area of improvement.

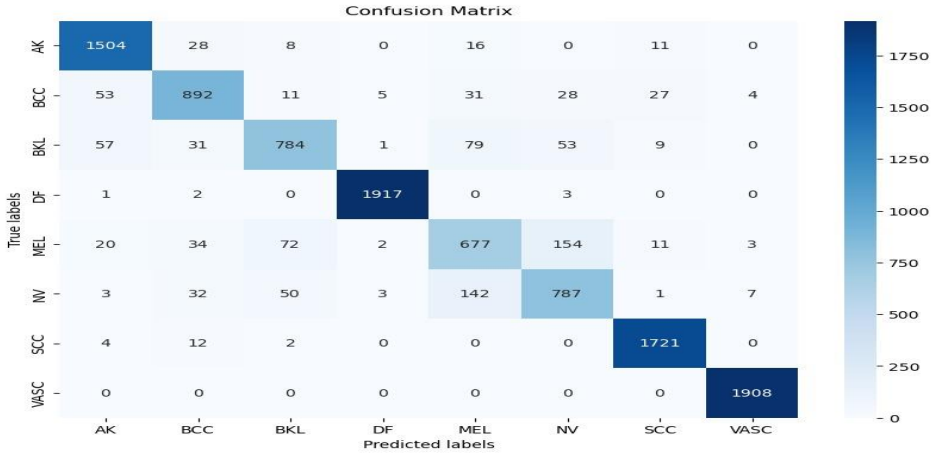

Fig. 9. Confusion Matrix.

5.2 Performance Evaluation

The performance of the hybrid ResNet model is evaluated using the ROC curve (see Fig 11) and convergence plots based (see Fig 10) on 20,256 training and 5,056 test images from the ISIC 2019 dataset. Multi-class classification ROC curve demonstrates AUC values of 0.97-1.00 for eight classes, with macro-average AUC = 0.99 and micro-average AUC = 1.00, demonstrating outstanding discriminant capability. Convergence plots indicate training loss decreasing from 0.7 to 0.1 and validation loss converging at 0.2 after 10 epochs, training accuracy increasing from 65% to 95%, and validation accuracy converging at 90%, proving strong training stability and minimal overfitting using focal loss and class balancing.

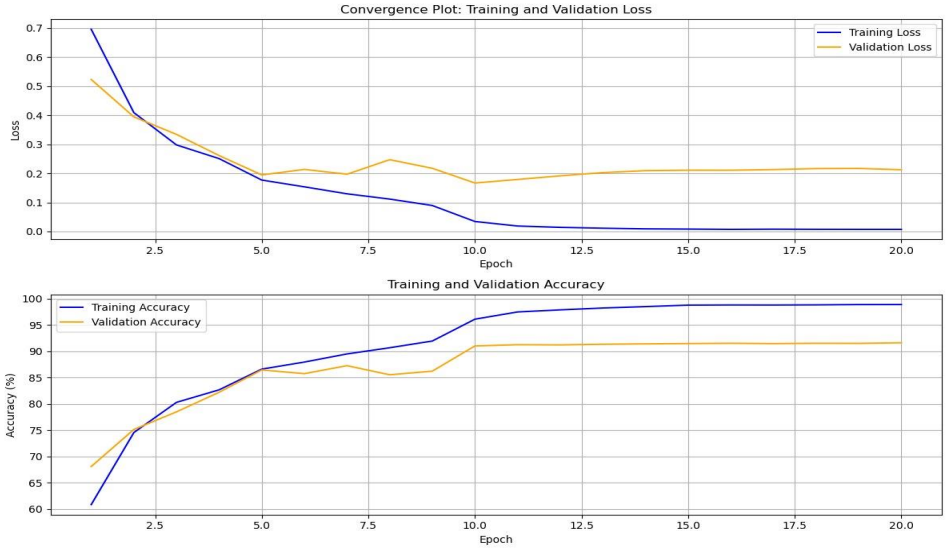

Fig. 10. Training and Validation Loss.

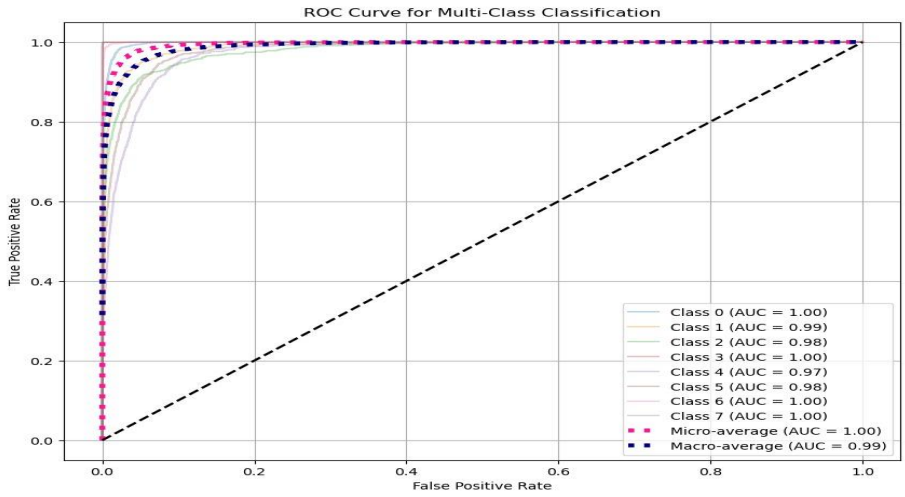

Fig. 11. ROC Curve for multi class classification.

class imbalance, attaining an overall accuracy of 0.91. These findings emphasize the framework's capability as a scalable and understandable diagnostic tool, setting the stage for future improvements, including enhancing melanoma detection and incorporating clinical validation to further progress automated skin cancer diagnosis. The hybrid ResNet model showcases robust performance in categorizing skin lesions and performing exceptionally well across various classes, as verified by the confusion matrix and ROC curve analyses.

6 Conclusion

This research showcases the effectiveness of a new hybrid deep learning model that combines ResNet-18, ResNet-34, and ResNet-50 to classify skin lesions into eight separate categories utilizing the balanced ISIC 2019 dataset. The thorough preprocessing pipeline, which includes noise reduction, normalization, class balancing, and data augmentation, along with a focal loss function, successfully manages the existing class imbalance, attaining an overall accuracy of 0.91. These findings emphasize the framework's capability as a scalable and understandable diagnostic tool, setting the stage for future improvements, including enhancing melanoma detection and incorporating clinical validation to further progress automated skin cancer diagnosis. The hybrid ResNet model showcases robust performance in categorizing skin lesions, and performing exceptionally well across various classes, as verified by the confusion matrix and ROC curve analyses.

References

[1] Mostafiz Ahammed, Md Al Mamun, and Mohammad Shorif Uddin. A machine learning approach for skin disease detection and classification using image segmentation. Healthcare Analytics, 2:100122, 2022.

[2] Saad Albawi, Oguz Bayat, Saad Al-Azawi, and Osman N Ucan. Social touch gesture recognition using convolutional neural network. Computational Intelligence and Neuroscience, 2018(1):6973103, 2018.

[3] Khalil Aljohani and Turki Turki. Automatic classification of melanoma skin cancer with deep convolutional neural networks. Ai, 3(2):512–525, 2022.

[4] Vatsala Anand, Sheifali Gupta, Ayman Altameem, Soumya Ranjan Nayak, Ramesh Chandra Poonia, and Abdul Khader Jilani Saudagar. An enhanced transfer learning-based classification for diagnosis of skin cancer. Diagnostics, 12(7):1628, 2022.

[5] Titus J Brinker, Achim Hekler, Alexander H Enk, Joachim Klode, Axel Hauschild, Carola Berking, Bastian Schilling, Sebastian Haferkamp, Dirk Schadendorf, Tim Holland-Letz, et al. Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task. European Journal of Cancer, 113:47–54, 2019.

[6] Arief Budhiman, Suyanto Suyanto, and Anditya Arifianto. Melanoma cancer classification using resnet with data augmentation. In 2019 International Seminar on Research of Information Technology and Intelligent Systems (ISRITI), pages 17–20, 2019.

[7] Charalampos Doukas, Paris Stagkopoulos, Chris T Ki-ranoudis, and Ilias Maglogiannis. Automated skin lesion assessment using mobile technologies and cloud platforms. In 2012 Annual international conference of the IEEE Engineering in Medicine and Biology Society, pages 2444–2447. IEEE, 2012.

[8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and JianSun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016.

[9] Milon Hossain, Kh Sadik, Md Musfiqur Rahman, Fa-had Ahmed, Md Nur Hossain Bhuiyan, and Moham-mad Monirujjaman Khan. Convolutional neural network-based skin cancer detection (malignant vs benign). In 2021 IEEE 12th Annual Information Technology, Electronics and Mobile Communication Conference (IEM-CON), pages 0141–0147. IEEE, 2021.

[10] Syed Inthiyaz, Baraa Riyadh Altahan, Sk Hasane Aham-mad, V Rajesh, Ruth Ramya Kalangi, Lassaad K Smi-rani, Md Amzad Hossain, and Ahmed Nabih Zaki Rashed. Skin disease detection using deep learning. Advances in Engineering Software, 175:103361, 2023.

[11] Supriya Joseph and Janu R Panicker. Skin lesion analysis system for melanoma detection with an effective hair segmentation method. In 2016 International Conference on Information Science (ICIS), pages 91–96. IEEE, 2016.

[12] Mohamed A. Kassem, Khalid M. Hosny, and Mohamed M. Fouad. Skin lesions classification into eight classes for isic 2019 using deep convolutional neural network and transfer learning. IEEE Access, 8:114822– 114832, 2020. ImageNet classification with deep convolutional neural networks. Advances in neural information processing systems, 25, 2012.

[13] Yuexiang Li and Linlin Shen. Skin lesion analysis towards melanoma detection using deep learning network. Sensors, 18(2):556, 2018.

[14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Doll´ar. Focal loss for dense object detection. In Proceedings of the IEEE international conference on computer vision, pages 2980–2988, 2017.

[15] Nay Chi Lynn and Zin Mar Kyu. Segmentation and classification of skin cancer melanoma from skin lesion images. In 2017 18th international conference on parallel and distributed computing, applications and technologies (PDCAT), pages 117–122. IEEE, 2017.

[16] Sadia Ghani Malik, Syed Shahryar Jamil, Abdul Aziz, Sana Ullah, Inam Ullah, and Mohammed Abohashrh. High-precision skin disease diagnosis through deep learning on dermoscopic images. Bioengineering,11(9):867, 2024.

[17] Sirine Mejri and Afef Elloumi Oueslati. Dermoscopic images classification using pretrained vgg-16 and resnet- 50 models. In 2024 IEEE 7th International Conference on Advanced Technologies, Signal and Image Processing (ATSIP), volume 1, pages 342–347. IEEE, 2024.

[18] HR Mhaske and DA Phalke. Melanoma skin cancer detection and classification based on supervised and unsupervised learning. In 2013 international conference on Circuits, Controls and Communications (CCUBE), pages 1–5. IEEE, 2013.

[19] RS Soumya, S Neethu, TS Niju, A Renjini, and RP Aneesh. Advanced earlier melanoma detection algorithm using colour correlogram. In 2016 International Conference on Communication Systems and Networks (ComNet), pages 190–194. IEEE, 2016.

[20] Wenbo Zhang, Xiaolin Sun, Shiyi Ma, Jinzhu Yang, Marcin Grzegorzek, Yuexi Wang, and Chen Li. Skin cancer image identification using deep convolutional neural networks. In International Conference on Image, Vision and Intelligent Systems, pages 17–27. Springer, 2023.