Deploying Neural Network Models for Skin Cancer Recognition and Diagnosis

Dr.Bhavani.Sankar. Panda [1], Sunku.Nithin Kumar[2], Shaik. Sadiya[3], Yerninti.Pradeep[4], Gudla.Shiswami[5]

[1] *Professor, Department of CSE(AIML), Raghu Engineering College, Visakhapatnam, AP*

[2,3,4,5,] *Students, Department of CSE(AIML), Raghu Institute of Technology, Visakhapatnam, AP*

***Abstract*- Skin cancer is a significant public health issue characterized by abnormal skin cell growth, primarily in areas exposed to UV radiation. The most common types include basal cell carcinoma, squamous cell carcinoma, and melanoma, with melanoma posing severe health risks if untreated. Timely detection and accurate classification are crucial for effective treatment. This study investigates the performance of various neural network architectures with datasets which contain diverse dermatoscopic images.We use models based on classification accuracy and computational efficiency. Our findings indicate that ResNet outperforms the other models in classification accuracy, while both ResNet and the custom CNN show faster testing times compared to the MLP. This research contributes valuable insights into different neural network approaches, advancing the field of skin cancer detection and enhancing diagnostic tools in dermatology**

***Index Terms* — Skin Cancer, Classification, Deep Learning, Dermoscopy,ResNet,CNN,MLP.**

I. INTRODUCTION

Skin cancer is among the most common types of cancer, with millions of new cases diagnosed each year around the globe. Skin cancer has a much higher chance of successful treatment and survival if it is caught early. Typically visually inspecting the skin, a dermatologist will make the skin cancer diagnosis through a biopsy. These traditional methods are also time-consuming and prone to human error. It has been well documented that with the advent of deep learning, a large amount of interest has been generated for automated systems that could assist, or be used instead of, dermatologists to make the diagnosis for skin cancer.

A promising tool for image-based medical diagnostics is convolutional neural networks (CNNs) and other deep learning architectures due to its rapid advancement. The datasets employed here are: two categories of skin images and dermoscopic skin cell images to demonstrate the power of CNNs in skin cancer detection. We use and evaluate three different deep learning architectures: a custom MLP, a transfer learning on ResNet50 model and a custom CNN with residual blocks.

II. TYPES OF SKIN CANCERS

There are three main types of skin cancer :1. BCC - basal cell carcinoma , which is the most common one and it slowly spreads on the body commonly on the neck ,face etc. It can appear like a small bump on skin. 2. SCC - squamous cell carcinoma, this can appear as scaly red patches and are seen mostly on sun exposed areas of our skin while this can spread at a faster pace unlike bcc. 3. Melanoma - It's rare but threatening, appears just like a dark spot and can spread fastly as this first develops in the melanocytes of our skin.

III. BACKGROUND STUDY

Several deep learning models are used to enhance the categorization of skin cancer. By merging VGG, CapsNet, and ResNet, Azhar Imran et al. (2025) created an ensemble model that outperformed the individual models in terms of sensitivity, specificity, and precision, with an accuracy of 93.5%. CNNs are effective in medical image classification, as demonstrated by H.L. Gururaj et al. (2024) use of CNNs for skin cancer detection on the HAM10000 dataset, where DenseNet169 achieved 91.2% accuracy. A CNN-based method for classifying dermoscopic images was presented by Akash Kumar V et al. (2023). Training was optimized using RMSPROP and ADAM optimizers, resulting in a 90% accuracy rate with no CPU overhead. Yessi Jusman et al. (2024) compared Multilayer Perceptron, custom CNN, and VGG-16 on the HAM10000 dataset, revealing that VGG-16 had the highest classification accuracy and faster testing times.

IV. METHODOLOGY

The three models we use in this study are as follows:

1. CNN (Convolutional Neural Network) : CNNs are a powerful tool for image classification, and their ability to automatically learn relevant spatial features makes them well-suited for skin cancer detection. This project utilizes a 9-layer CNN architecture, comprising 3 convolutional layers, 3 max pooling layers, 1 fully connected (FC) layer, and 2 dense layers, to process both dermoscopic and "normal" skin images, aiming for robust classification.

The 9-layer CNN consists of three convolutional layers, each employing learnable filters to detect local patterns like edges and textures. These are followed by three max pooling layers, which reduce dimensionality and enhance robustness to minor image variations. A fully connected (FC) layer then combines the learned features, and finally, two dense layers act as the classifier, producing the benign/malignant predictions. The specific kernel sizes, strides, padding, and pooling window sizes are used.

CNN progressively extracts increasingly abstract features. Initial convolutional layers detect simple features, while deeper layers learn more complex, disease-specific patterns. The max pooling layers downsample, and the FC and dense layers perform the final classification. The network is designed to handle both dermoscopic and "normal" skin images, and the methodology should detail how these are processed (e.g., separate branches or shared layers).

The model is compiled using the Adam optimizer (tf.keras.optimizers.Adam()) and the categorical cross-entropy loss function (categorical_crossentropy). Accuracy is used as the evaluation metric during training.

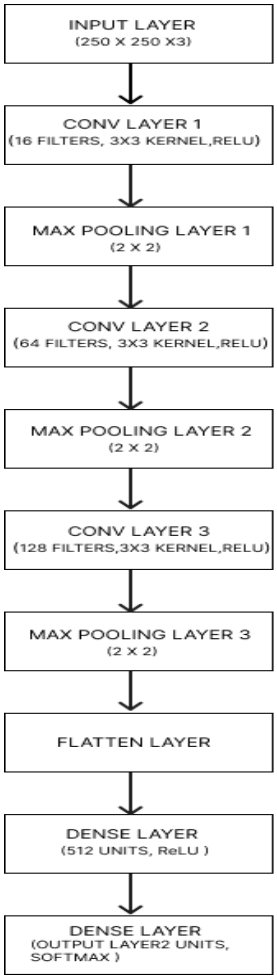

Fig 1- Architecture of custom CNN model

2. ResNet (Residual Network) : Residual blocks, used by ResNet50, enable the model to learn identity mappings through shortcut connections. By avoiding the vanishing gradient issue, these connections facilitate the efficient training of deeper networks. In the realm of deep learning, Residual Networks, or ResNets, have become a ground-breaking architecture, especially for image identification tasks.

They tackle the vanishing gradient problem, a crucial obstacle in the training of extremely deep neural networks. Gradients computed during backpropagation can get incredibly tiny as networks get deeper, which makes it more difficult for previous layers to learn effectively. ResNets get around this by adding "shortcut connections," or skip connections, which let gradients move across the network more directly.

ResNet50, a pre-trained convolutional neural network renowned for its deep residual architecture, is used at the heart of the model. The following is imported into the ResNet50 model:

1)Pre-trained Weights: To take advantage of learnt features from a sizable image dataset, use weights="imagenet".
2)Leave Out the Top Layer: Customization is possible by setting include_top=False, which eliminates the last classification layer.
3)Input Shape: To match the image dimensions in the dataset, the input shape is set to (250, 250, 3).

The ResNet50 base model's layers are all configured to be trainable (layer.trainable = True). Fine-tuning, in which the previously learned weights are modified during training to fit the new dataset, is made possible by this choice.Adapting features to the particular goal of skin cell categorization is necessary to get a higher level of model performance.

Including Personalized Layers of Classification as to adapt the ResNet50 base model to the classification task, further layers are added:
Only the global features from each channel are retained when the feature maps' spatial dimensions are reduced via the Global Average Pooling Layer (GlobalAveragePooling2D).
(Dense(128, activation="relu") is the Dense Layer. introduces non-linearity by adding a fully linked layer with 128 units and a ReLU activation function.
Layer of Dropout: (Dropout(0.5)) reduces overfitting by randomly setting 50% of the neurons to dormant during training.
(Dense(9, activation="softmax") is the output layer. produces a probability distribution for multi-class classification by implementing a final Dense layer with nine units (because there are nine classes) and a softmax activation function.

AdamW, a weight decay adaptive learning rate optimizer for improved generalization (weight_decay=1e-4).For a steady and gradual update of the model weights, use a learning rate of 1e-5.For multi-class classification using one-hot encoded labels, the loss function is categorical_crossentropy.During training and validation, performance is gauged using accuracy.

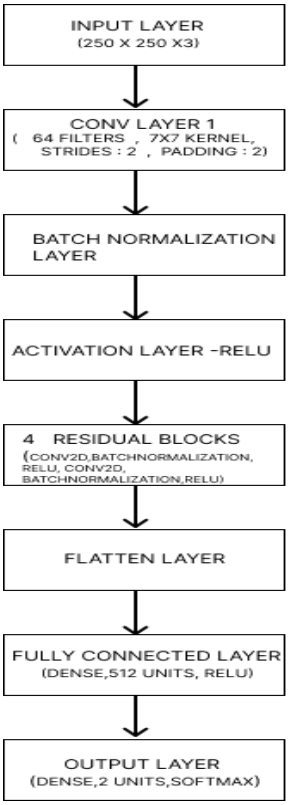

Fig 2 - Architecture of MLP

3. MLP (MultiLayer Perceptron): The Multilayer Perceptron is a basic feedforward neural network consisting of multiple layers of nodes. While traditionally not as powerful as CNNs for image data, MLPs can still be effective when used with preprocessed features or when combined with other models. In this project, MLP is used as a baseline model to compare the performance of more advanced architectures like CNN and ResNet. The insights gained from MLP's performance help refine the design and training of the deeper networks.

Residual Blocks in a Multi-Layer Perceptron (MLP).An MLP (Multi-Layer Perceptron) with convolutional layers, residual blocks, and dense layers forms the basis of the model architecture. Included in the architecture are:The first layer of convolution a 7x7 kernel, 64 filters, and stride of 2 in a Conv2D layer.In order to introduce non-linearity and stabilize the training, Batch Normalization and ReLU activation come next.

The model includes Residual Blocks, which facilitate skip connections and improve learning .In order to enable deeper network training, residual blocks solve the vanishing gradient issue and enable the model to learn identity mappings.

The first convolutional layer employs ReLU activation, Batch Normalization, and a 3x3 kernel.In order to get ready for residual addition, the second convolutional layer likewise employs a 3x3 kernel and Batch Normalization without activation.

There is now a shortcut connection:The shortcut runs through a 1x1 convolution to match the main path's geometry if the stride is not 1.Batch Normalization is also used to normalize the shortcut.After adding the output from the shortcut and the main path, ReLU is activated.

Block Setup:Four residual blocks with increasing filters and downsampling are included in the model:The stride of 64 filters is 1, the stride of 128 filters is 2, the stride of 256 filters is 2, and the stride of 512 filters is 2.

The model flattens the features and uses dense layers for classification after feature extraction via convolutional and residual layers.The flatten layer creates a 1D vector from the 2D feature maps. ReLU activation and 512 units make up this dense layer.Before the output layer, an additional layer of abstraction is added.The output layer is a dense layer with nine units, which represent the dataset's nine classifications.creates a probability distribution for multi-class classification using softmax activation.
The model is put together using - Adam is an adaptive optimizer with a 0.001 learning rate.Suitable for multi-class classification, the loss function is categorical_crossentropy.

V. DATASETS

The models are trained and evaluated on two different datasets:
Skin Cancer Image Dataset: A set of skin images classified into benign and malignant categories.
Dermoscopic Skin Cell Dataset: Images of dermoscopic skin cells used to identify various skin diseases, including cancer.
The datasets used in this study consist of:

1. Cancer Image Dataset:

This dataset comprises images of skin lesions, labeled as either "Skin Cancer" or "Not Skin Cancer" and were collected from the Kaggle platform ,dataset named “skin cancer or not skin cancer image datasets” by Md.Ismiel Hossen Abir.

Preprocessing Steps are resizing all images to 250x250 pixels and Rescaling/Normalization - Pixel values were rescaled to the range [0, 1].

Possible drawbacks are that details regarding the demographics and diversity of the Kaggle dataset(s) are essential. It is also important to take into account the original image quality and the Kaggle labeling uniformity.

2. Dermoscopic Skin Cell Dataset:

This dataset consists of 2,357 dermoscopic images of skin cells, from a kaggle dataset named ”Skin Cancer ISIC” by Andrey Katanskiy sourced from the ISIC (International Skin Imaging Collaboration) archive.

These images represent a variety of skin lesion types, the distribution of each images in these types is as follows:
Actinic Keratosis-114 images
Basal Cell Carcinoma -376 images
Dermatofibroma - 95 images
Melanoma - 438 images
Nevus - 357 images
Pigmented Benign Keratosis - 462 images
Seborrheic Keratosis - 77 images
Squamous Cell Carcinoma - 181 images
Vascular Lesion - 139 images

Preprocessing Steps are as follows: resizing - all images were resized to 250x250 pixels and Rescaling/Normalization - Pixel values were rescaled to the range [0, 1].

Possible limitations are that details regarding the demographics and diversity reflected in the ISIC dataset are essential. It is important to take into account the initial image quality, variations in dermoscopic procedures, and other biases in the data.

Combined Dataset Information:Training and Testing Split-Both datasets were split into training and testing sets using an 80-20% ratio and a stratified split was used to ensures that the class proportions are maintained in both the training and testing sets.

VI. PRE-PROCESSING

It may be useful to perform several transformations to dermoscopy images before feeding them to a model.

For MLP we do the following- ImageDataGenerator from Keras is used to load and preprocess the image data at the start of the process. For training and validation, the dataset is made up of pictures of skin cells arranged in a directory structure.
Rescaling: To standardize the input data, the image pixel values are scaled to the [0, 1] range using rescale=1/255.
The flow_from_directory technique is utilized for image loading in order to open the "Train" and "Test" directories and load pictures.
Images should be resized to 250x250 pixels.
Use three as the batch size.
For the multi-class classification task, use categorical labels.

For RESNET we do the following -
Rescaling: rescale=1/255 is used to scale pixel values to the [0, 1] range.
Rotation: Up to 30 degrees of random rotation is applied to images (rotation_range=30).
Shifting: Up to 20% of the image's size can be moved horizontally and vertically (width_shift_range=0.2, height_shift_range=0.2).
Shearing: A shear intensity of 0.2 (shear_range=0.2) is used to transform images.
Zooming: Zooms randomly between 20% of the range (zoom_range=0.2).
Flipping: when horizontal_flip=True, horizontal flipping is enabled.
Filling: The closest pixel values (fill_mode="nearest") are used to fill any empty pixels produced by transformations.
To preserve the original integrity of the test images, only rescaling is applied to the validation data.
The data is then loaded using flow_from_directory, with both datasets resized to 250x250 pixels and a batch size of 8.

For CNN we do the following - preprocess training and validation image data for a skin cell classification job using `ImageDataGenerator`.
To stabilize model training, images are rescaled from [0,255] to [0,1] using `rescale=1/255`.
The `flow_from_directory` method creates batches of three photos with one-hot encoded labels (`class_mode='categorical'`), loads images from designated folders ("skin cells\\Train" and "skin cells\\Test"), and resizes them to 250x250 pixels.
While the validation data is merely rescaled, the training data is enhanced by transformations such as rotation, shifting, and flipping. This preprocessing improves model generalization and guarantees constant input size.

Network Training:
For MLP - the fit approach is used to train the model over 10 epochs as the model makes ten iterations over the complete dataset.The model's performance is assessed on the validation set at the end of each epoch.The history object contains the accuracy of the training and validation.
Training Accuracy will be the training set's maximum recorded accuracy.Validation Accuracy will be the model's capacity for generalization is demonstrated by the maximum accuracy attained on the validation set.

For Resnet - Twenty epochs are used to train the model.
Batch Size: 8, there are eight photos in each batch. Following each epoch, a validation set is sent to track performance on unseen data.Uses AdamW (Adam with weight decay) including 1e-5 learning rate and also weight decay of 1e-4 (assists in maintaining consistency) loss function with categorical_crossentropy, appropriate for one-hot encoded labels in multi-class classification.

For CNN - The model is trained for 10 epochs.
The batch size is 3, as the weights are updated after every 3 samples.The model architecture is saved as JSON (model_architecture1.json).
The model weights are saved separately as an H5 file (model_weights.weights1.h5).
For visualizations we use accuracy plot which displays how the training and validation accuracy change over epochs.And Loss plot which displays the progression of training and validation loss.

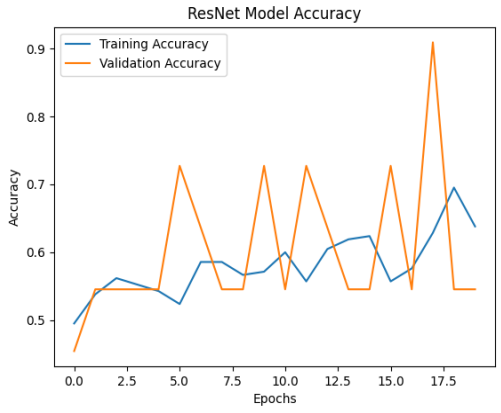

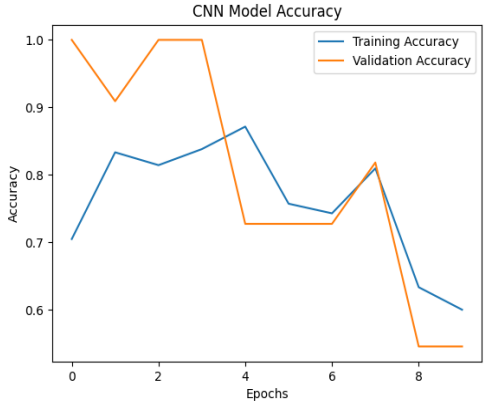

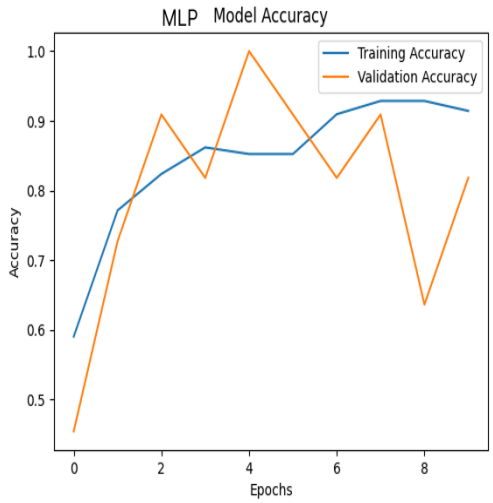

VII. RESULTS

The custom MLP model, achieved a maximum training accuracy of approx 95% and a validation accuracy of 90% over 10 epochs on both datasets of skin images and skin cells.

The custom CNN model, incorporating convolutional layers, aimed to balance complexity and generalization. It achieved a training accuracy of 85% and a validation accuracy of 95% on skin images dataset and 75% of training accuracy and 40% of validation accuracy on the skin cell images dataset.

The ResNet50 model , achieved a maximum training accuracy of 70% and validation accuracy of 90% on skin images dataset and training accuracy of 78% and validation accuracy of 60% on the dermascopic skin cell dataset.
Overall, the comparison of these models demonstrated the superior performance of Resnet50 model.

VIII. CONCLUSION

In this research, three deep learning models—a custom Convolutional Neural Network (CNN), a Multi-Layer Perceptron (MLP), and a pre-trained ResNet50 model—were developed and evaluated for skin cancer detection using dermatoscopic images and skin images. The models were trained and validated on two datasets with a focus on achieving high accuracy and.robustness in classification.
The experimental results of this study revealed that custom MLP model provides a more effective solution for skin cancer detection compared to CNN and Resnet architectures. While the custom CNN and Resnet50 models displayed reasonable performance, the MLP model consistently achieved higher accuracy and better generalization, making it a strong candidate for practical clinical use.

The study highlighted the importance of residual connections in improving learning stability in the MLP model and showcased the benefits of data augmentation and regularization techniques in all models. In conclusion, this research contributes to the development of reliable and accurate skin cancer detection systems, potentially aiding healthcare professionals in early diagnosis and improving patient outcomes. The findings support the adoption of transfer learning approaches in medical image analysis, emphasizing their role in achieving state-of-the-art performance with limited datasets.

REFERENCES

[1] C. Barata, M. E. Celebi, and J. S. Marques, “A survey of feature extraction in dermoscopy image analysis of skin cancer,” *IEEE Journal of Biomedical and Health Informatics*, 2018.

[2] A. Esteva, B. Kuprel, R. A Novoa, and et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, pp. 115–118, 2017.B. Harangi, “Skin lesion classification with ensembles of deep convolutional neural networks,” *Journal of biomedical informatics*, vol. 86, pp. 25–32, 2018.

[3] A. Menegola, M. Fornaciali, R. Pires, and et al., “Towards automated melanoma screening: Explor- ing transfer learning schemes,” *arXiv preprint arXiv:1609.01228*, 2016.

[4] K. Matsunaga, A. Hamada, A. Minagawa, and H. Koga, “Image classification of melanoma, nevus and seborrheic keratosis by deep neural

network ensemble," *arXiv preprint arXiv:1703.03108*, 2017.

[5] C. Barata, M. E. Celebi, and J. S. Marques, "Improv- ing dermoscopy image classification using color constancy," *IEEE Journal of Biomedical and Health In- formatics*, vol. 19, pp. 1146–1152, 2015.

[6] N. Srivastava, G. E. Hinton, A. Krizhevsky, and et al., "Dropout: a simple way to prevent neural networks from overfitting.," *Journal of machine learning research*, vol. 15, pp. 1929–1958, 2014.

[7] Hari Kishan Kondaveeti and Prabhat Edupuganti, "Skin Cancer Classification using Transfer Learning"

[8] Azhar Imran Abdulkareem Alzaharani, Arslan Nasir, Muhammed Bilal, Guangminsun, Abdullah Almuhaimeed "Skin Cancer Detection Using Combined Decision of Deep Learners"

[9] Nour Abuared, Alavikunhu Panthakkan, Mina Al-Saad, Saad Ali Amin and Wathiq Mansoor "Skin Cancer Classification Model Based on VGG19 and Transfer Learning"

[10] Yessi Jusman, Indah Monisa Firdiantika and Dhimas Arief Dharmawan " Performance of Multi Layer Perceptron and Deep Neural Networks in Skin Cancer Classification "

[11] H.L. Gururaj, N.manju, A.Nagarjun, V.N. Manjunath and Francesco Flammini "DeepSkin: A Deep Learning Approach for Skin Cancer Classification"

[12] Akash Kumar V, Vijaya Mishra and Monika Arora " Convolutional Neural Networks for Malignant and Benign Cell Classification using Dermatoscopic Images"

[13] Catarina Barata and Jorge S. Marques "Deep Learning for Skin Cancer Diagnosis with Hierarchical Architectures"