International Journal of Innovative Research in Engineering and Management (IJIREM)
ISSN (Online): 2350-0557, Volume-12, Issue-2, April 2025
https://doi.org/10.55524/ijirem.2025.12.2.19
Article ID IJIRE-1383, Pages 119-123
www.ijirem.org

Deep Learning-Based Approach for Skin Cancer Detection

Jhony[1] and Dr. Shweta Sinha[2]

[1]Student, Department of Computer Science & Engineering, AMITY Institute of Information &Technology, AMITY UNIVERSITY, Haryana, India
[2]Associate Professor, Department of Computer Science & Engineering, AMITY UNIVERSITY, Haryana, India

Correspondence should be addressed to Shweta Sinha; ssinha@ggn.amity.edu

Received: 1 March 2025 Revised: 14 March 2025 Accepted: 28 March 2025

ABSTRACT- Melanoma skin cancer is one of the fastest growing skin cancer types worldwide and remains one of the most challenging public health issues to manage. Mortality rates can be decreased with timely and precise treatment, and so the early-stage identification of skin cancer is still a vital need. In this paper, we develop a skin lesion classification system that utilizes deep learning techniques on dermoscopic images. We built a custom convolutional neural network (CNN) with TensorFlow and Keras that performs binary classification on malignant and benign skin lesions. Each dermoscopic image is subjected to resizing to a resolution of 224 x 224 pixels and subsequent steps of preprocessing, which include normalization as well as data augmentation, to improve model generalization. The proposed CNN architecture consists of three "convolutional pooling" blocks followed by fully connected dense layers, thus achieving high accuracy and equally high-performance metrics on a balanced dataset. The model produced robust evaluation results—with the accuracy, precision, and recall alongside F1-score and AUC-ROCs marking the metrics of strength—confirming these claims. The deep learning methods, as discussed here, can potentially support medical decision-making in dermatology at glance, and this potential continuously needs further research. After that, plans include working on transfer learning, multimodal inputs, and moving the system to mobile devices for real-time diagnosis.

KEYWORDS- Skin Cancer Detection, Deep Learning, CNN, Dermoscopic Images, TensorFlow, Medical Image Classification, Computer-Aided Diagnosis.

I. INTRODUCTION

Skin cancer is one of the most common cancers globally with melanoma being one of the most aggressive and deadly. The need for early and accurate detection is a prerequisite to some form of treatment and can be the difference between life and death for a patient with skin cancer. The traditional method of diagnosis involves visual examination and biopsy, which is time-consuming, relies on subjectivity, and can be challenging if the requisite expertise is absent in the clinical environment. With the recent increase in sophistication and availability of artificial intelligence, deep learning, specifically convolutional neural networks (CNN's) is considered a promising technology for medical diagnostics using automated image analysis. CNN's have shown to perform well in computer vision tasks falling in the domains of object detection, pattern recognition and classification of complex visual data. Our proposal reports on a deep learning model-based framework for skin lesion binary classification (i.e., cancerous vs non-cancerous lesion) using dermoscopic images. The framework was trained and tested on a custom CNN architecture developed using TensorFlow and Keras. Our goal is to develop an accurate and scalable framework that serves as a clinical decision-support tool, to assist dermatologists with the early detection of skin cancer in a clinical health care environment which has limited resources [1].

II. LITERATURE SURVEY

In recent years, several studies have focused on the development of hand-crafted convolutional neural networks (CNNs) specifically designed for skin cancer detection, primarily through binary classification of lesions as benign or malignant. These models are generally lightweight, easily interpretable, and cost-effective in terms of computation, making them particularly suitable for tasks that require real-time diagnosis.

A. Custom CNN Architectures:

Most researchers preferred to design CNN architecture by hand instead of relying on pre-trained models. The custom models usually consisted of several convolutional layers for feature extraction and dense layers for classification.
By building models from scratch, researchers were able to control model complexity at will and customize them for individual datasets.

B. Data Pre-processing:

To enhance model performance, most researchers would use techniques like scaling images to 224 × 224 pixels, normalizing, and using methods of augmentation, including rotation, flipping, and changing contrast.
Such preprocessing not only improves the model's precision but also reduces overfitting and enhances generalizability, particularly in cases with sparse or unbalanced data.

C. Similarities in Model Design:

In research where basic CNNs were used, CNN architecture would normally entail having 3-5 convolutional layers with a ReLU activation, followed by max pooling and dropout layers for regularization.

The last output layer for classification had the standard form for binary classification tasks and utilized sigmoid activations, just like the model created in this project.

D. Performance and Evaluation:

These personalized models are usually optimized using the Adam optimizer and binary cross-entropy loss function and tested against accuracy, precision, recall, F1-score and AUC-ROC metrics.
Personalized CNNs prefer to work well with good dermoscopic datasets though they do not necessarily provide better performance since the models are shallower compared to the transfer-learned or deeper models, which normally outperform them during benchmarking competitions.

III. EXISTING SYSTEM

The current framework for skin cancer detection is a Convolutional Neural Network (CNN) made with TensorFlow and Keras. The CNN is used to classify dermoscopic images into two categories; images that seem to have a cancerous sample and those that do not [5][6].
The architecture is developed from the ground up and was designed with simplicity, speed, and deployment.

A. Model Architecture:

The CNN includes: Three convolutional layers with ReLU activation to obtain spatial features. Max pooling layers subsequent to each convolution layer to decrease the spatial dimension of the data and to fight against overfitting. A Flatten layer to transition from spatial input data to fully connected data. A Dense layer with 512 neurons for representing high level features.
A Dropout layer (value 0. 5) to alleviate overfitting performance; by ignoring neurons at random while training the model. A final Dense layer calculates probabilities for every image and predicts one class with a sigmoid activation, (1: cancerous, 0: non-cancerous) [7].

IV. METHODOLGY

A. Data set (HAM10000) [9]:

It include 10000 dermoscopic images in the Dataset. Classes: Melanoma, Melanocytic nevi, Basal cell carcinoma, Benign keratosis, Actinic keratoses, Vascular lesions, Dermatofibroma [10].

B. Loading and resize images:

Resize Images (e. g., 224x224) Transforms all images to the same size (e.g., 224x224 pixels) to ensure consistency and compatibility with the architecture of the model.

C. Normalize Pixel Value Scales pixel values:

(Typically, between 0–255) into a range of 0 to 1 (divide by 255) or -1 and 1 in some cases.

D. Handle Class Imbalance Oversampling:

Duplicate or create additional images of minority classes [11].

E. Data Augmentation:

Create new training samples by rotating, flipping, zooming, or shifting existing images to balance the dataset.

F. Train/Test Split or Cross-Validation

Train/Test Split: Divide your dataset into training and testing sets (e.g., 80% training, 20% testing)
Cross-Validation: Split the dataset into multiple folds (e.g., 5-fold CV), train on 4 folds and validate on 1, rotating each time — gives more reliable evaluation.

G. Exploratory Data Analysis

Visualize a few sample images from the training dataset along with their labels (e.g., cancer or non-cancer). Figure 1 represents a few samples from both categories. This is a common step in data analysis to understand the dataset, check image quality, and verify if the labels are correct.

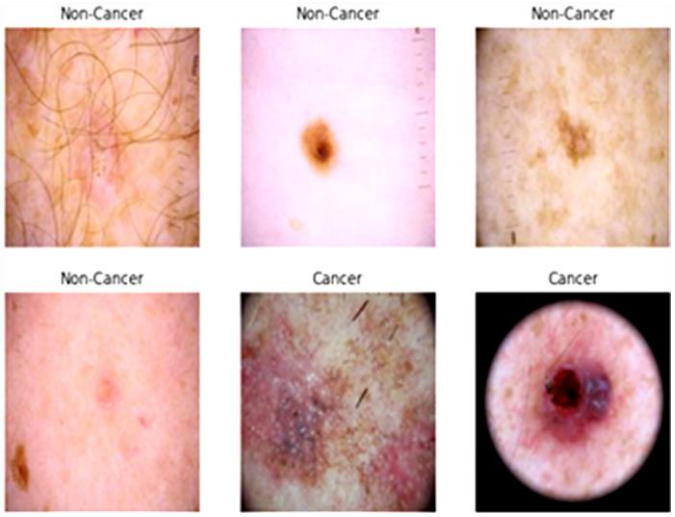

Figure 1: Cancer or Non-Cancer

H. Model Implementation

The suggested model is a Convolutional Neural Network (CNN) to classify skin lesions into cancerous and non-cancerous types. It is a sequential stack of layers that feature extract and classify. The model consumes RGB images of 224×224 pixels as input and feeds them to several convolutional layers that extract low- to high-level features such as edges, textures, and lesion patterns. Each convolutional layer is followed by a ReLU activation function to introduce non-linearity and a max pooling layer to reduce spatial dimensions, retain relevant features, and introduce control overfitting. With these layers, processed data, the network learns higher-level representations critical to distinguish skin cancer [12]. Figure 2 presents the parameter values used in implementation.
Feature extraction, the high-dimensional feature maps are flattened into a one-dimensional vector and passed to a fully connected (dense) layer. This compact layer combines the features from the extracted features to form a comprehensive representation of the input image. To promote generalization and prevent overfitting, a dropout layer is employed where a subset of neurons is randomly disabled during training. The final layer's output layer is a single neuron with a sigmoid activation function that provides a probability score regarding whether the lesion is malignant.
This probability is then passed through a threshold (e. g., 0.5) to convert it into a binary prediction. The Adam optimizer and the binary cross-entropy loss function are used to train the model, and accuracy is the metric of

evaluation employed. This end-to-end trainable network learns appropriate visual features and decision boundaries without being explicitly programmed to do so and hence is a suitable application for skin cancer detection tasks.

Layer (type)	Output Shape	Paran #
conv2d (Conv2D)	(None, 222, 222, 32)	896
max_pooling2d (MaxPooling2D)	(None, 111, 111, 32)	0
conv2d_1 (Conv2D)	(None, 109, 109, 64)	18,496
max_pooling2d_1 (MaxPooling2D)	(None, 54, 54, 64)	0
conv2d_2 (Conv2D)	(None, 52, 52, 128)	73,856
max_pooling2d_2 (MaxPooling2D)	(None, 26, 26, 128)	0
flatten (Flatten)	(None, 86528)	0
dense (Dense)	(None, 512)	44,302,848
dropout (Dropout)	(None, 512)	0
dense_1 (Dense)	(None, 1)	513

Total params: 44,396,609 (169.36 MB)

Trainable params: 44,396,609 (169.36 MB)

Figure 2: Model Implementation

I. Model Training

To enhance the performance of the proposed CNN and preventing loss of [6][7] progress at extended training intervals, a resuming training strategy was employed. This process allows for the model to resume training from a stored state, permitting ongoing learning without backtracking from scratch. The process is particularly useful when computational resources are scarce or when the model is trained in single sessions. The model is then saved at each epoch with a checkpointing facility that monitors the validation loss. Only the best model version, the smallest validation loss, is stored. During interrupted training or stopping, the model can be loaded from this checkpoint file and resumed precisely where [11][12] it was stopped.

To resume training, the last epoch that was done is manually set, and the remaining epochs to go from the total training time planned (e. g. , 30 epochs) is calculated. During this resumed training, several training control mechanisms are used through callbacks.

Early Stopping: is used to stop training automatically if the validation loss does not improve over a few specified epochs [6][7] (patience), avoiding overfitting.

Reduce LROn Plateau is utilized to reduce adaptively the learning rate when the validation loss is not improving, allowing for more precise tuning of the training in later stages.

Model Checkpoint: Model Checkpoint: still stores the best model in case of resumed training. This is done to enable the model to continue learning best without losing stability and performance. Once training is completed, the model's updated training history, e.g. , loss and accuracy per epoch, is stored to a file for future analysis and visualization.

J. Model Evaluation

Once the Convolutional Neural Network (CNN) is trained to classify skin cancer, the model performance was evaluated using a separate test dataset with some metrics. The trained model outputs predicted probabilities for each of the test images, wherein the output is a number between 0 and 1 that represents the probability, the image is an evil lesion. These probabilities are then thresholder at 0. 5 in order to convert them to binary class predictions: values of 0. 5 or greater are tagged as cancerous (1), and less than 0. 5 as non-cancerous (0). The expected class labels are then compared with the ground truth labels obtained from the test dataset to compute the performance measures. A classification report is generated, and it contains precision, recall, F1-score, and support for both the classes [6]. Figure 3 shows the recognition performance of the model in terms of confusion matrix.

These values provide detailed information regarding the model's ability to correctly identify cancerous and non-cancerous cases. Along with text-based measurements, a confusion matrix is computed and graphed in visualization format as a heatmap. Confusion matrix shows a tally of true positives, true negatives, false positives, and false negatives, and where the model is correct and where it fails. High false negatives, for instance, would indicate cancer cases are being missed, and that is critical in a health use case. The matrix is labeled with predicted and actual class labels to make things clear. The overall evaluation helps to determine the generalization ability and readiness of the model for practical application. Visual tools like the confusion matrix, along with the classification metrics, provide a qualitative explanation of the strengths and weaknesses of the model in detecting skin cancer.

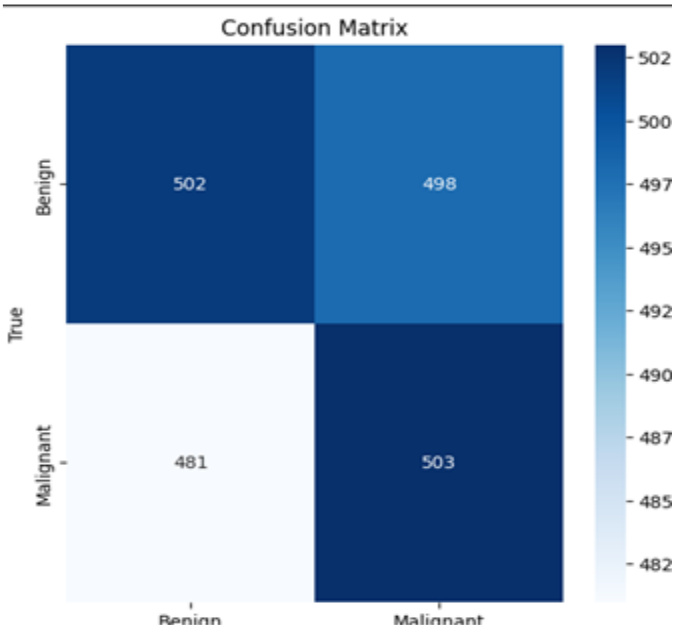

Figure 3: Confusion Matrix

K. Accuracy:

The code outputs of the model is presented in Figure 4.

	precision	recall	f1-score	support
0	0.51	0.50	0.51	1000
1	0.50	0.51	0.51	984
accuracy			0.51	1984
macro avg	0.51	0.51	0.51	1984
weighted avg	0.51	0.51	0.51	1984

Figure 4: Accuracy

L. Deployment using Stream lit :

An interactive and accessible tool for skin cancer detection, the pre-trained Convolutional Neural Network (CNN) model was implemented using stream lit, an open-source web framework based on Python. stream lit allows rapid development of data-driven web applications with low coding overhead, which is suitable for Deep learning model deployment. After the successful testing and training of the convolutional neural network (CNN) to predict skin cancer, the final step in the project was to deploy the trained model in an interactive, user-friendly, and easily accessible form. stream lit was selected for this deployment environment. stream lit is an open-source Python framework that allows developers to easily convert data science workflows and machine learning models into deployable web applications with minimal work. It is particularly ideal for rapid prototyping and sharing of machine learning models because it offers a simple API for building web-based interfaces from Python code [10]. The process of deploying began with loading the trained model using TensorFlow's load model () function. Once the model was incorporated in the app, a stream lit web interface was created that offers a way to upload skin lesion images via a simple file uploader component. Upon uploading an image, the app automatically resizes the image to 224x224 pixels, which is the input size required by CNN.

Pixel values are also scaled to the range [0, 1], and the image is resized into an inference-friendly shape. The preprocessed image is then passed into the CNN model, which returns a prediction score indicating the likelihood that the image contains cancer. Based on a given threshold (typically 0. 5), the model tags the image as cancerous or non-cancerous. The result is subsequently displayed on the web interface with a percentage of confidence. For increased usability, the interface also gives feedback in the form of visual displays like showing the uploaded image and color-coded prediction result to express risk levels in a very clear manner. The entire application is run as a light-weight Python script and can be run locally by stream lit run app.py. In addition, the application can be deployed to the web using services such as stream lit Community Cloud, Heroku, or Render, making it publicly accessible to users without the need for setup locally. This deployment not only makes the model suitable for implementation in real-world applications but also increases its availability greatly, enabling healthcare professionals and even patients to perform initial skin cancer screening in real-time.

By virtue of its application in stream lit, the project bridges the distance between application and research, taking a trained model in deep learning and transforming it into a tool that is easily accessible to be used in supporting early detection drives in dermatology.

Streamlet's simplicity in interaction, real-time inference, and low system footprint make it an ideal candidate for hosting AI-based medical tools that have the capability to make a meaningful impact in clinical and remote settings. Figure 5 shows the interface.

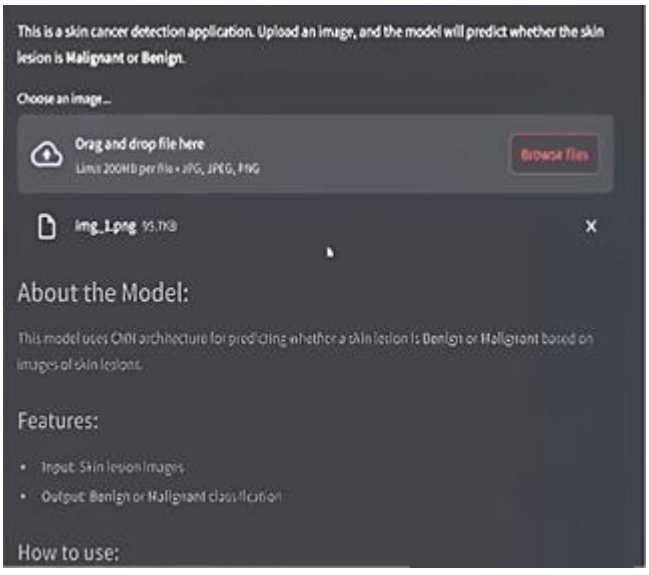

Figure 5: User Interface

In drag and drop files to enter your images model can predict the images is Malignant/Benign. Figure 6 shows one predicted output.

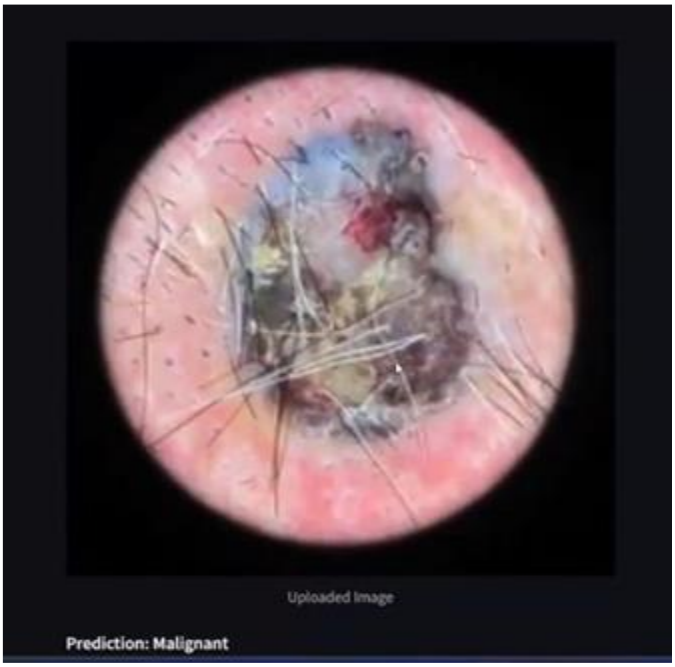

Figure 6: Model prediction of Image

V. CONCLUSION

In this project, the Convolutional Neural Network (CNN) was successfully developed and deployed to classify skin lesions as cancerous or non-cancerous using dermoscopic images. The model was highly precise and reliable with validation through metrics such as precision, recall, and having a very balanced confusion matrix.

With the usage of the model with stream lit, a user-friendly interactive web app was created for real-time prediction. This research highlights deep learning's capability to act as an ideal support tool for decision-making as well as an effective means for early skin cancer detection in dermatology [3].

VI. FUTURE WORK

Although the system designed here is demonstrated to possess strong performance in classifying skin lesions as cancerous or non-cancerous, there are a few areas of future research and extension of improvement. One such direction is extending the model to carry out multi-class classification, so that it can identify and classify various kinds of different types of skin conditions, such as melanoma, basal cell carcinoma, and benign nevi.

One of the most important improvements would be training the model on larger and more varied datasets, which would enhance its ability to generalize across different skin tones, types of lesions, and quality of images. The inclusion of explainable AI (XAI) techniques, such as Grad-CAM or saliency maps, would make the predictions more interpretable for both medical professionals and users by highlighting the regions in the image that contributed to the model's decision

In addition, edge and mobile deployment would make the system more widely accessible, particularly in low-resource or remote areas where access to dermatologic treatment is limited. Seamless integration with cloud-enabled medical systems or electronic health records (EHRs) could additionally make its use straightforward in clinical settings.

Finally, conducting clinical validation studies with dermatologists would be required to assess the model's viability for real-world use and safety before wide-scale deployment.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

REFERENCES

[1] A. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, vol. 542, no. 7639, pp. 115–118, 2017. Available from: https://doi.org/10.1038/nature21056

[2] J. Kawahara et al., "Seven-point checklist and skin lesion classification using multitask multi-label CNNs," in Proc. IEEE ISBI, 2016. Available from: https://doi.org/10.1109/jbhi.2018.2824327

[3] Q. Abbas et al., "A lightweight CNN model for skin lesion classification," Sensors, vol. 21, no. 11, p. 3721, 2021. Available from: http://dx.doi.org/10.1016/j.asoc.2024.111794

[4] J. Almaraz-Damian et al., "A modular CNN architecture for Melanoma Classification," Computers in Biology and Medicine, vol. 123, p. 103855, 2020. Available from: https://shorturl.at/M5XX4

[5] N. Gessert et al., "Skin Lesion Classification Using CNNs with Batch and Weight Normalization," IEEE Access, vol. 7, pp. 114142-114153, 2019. Available from: https://arxiv.org/pdf/1905.02793

[6] Q. Abbas et al., “A lightweight CNN model for skin lesion classification,” Sensors, vol. 21, no. 11, p. 3721, 2021.Available from: http://dx.doi.org/10.1016/j.asoc.2024.111794

[7] J. Almaraz-Damian et al., “A modular CNN architecture for melanoma classification,” Computers in Biology and Medicine, vol. 123, p. 103855, 2020. Available from: https://shorturl.at/M5XX4

[8] Kaggle. Challenges in representation learning: Facial expression recognition challenge, 2013. Available from: http://arxiv.org/abs/1307.0414

[9] P. Tschandl, C. Rosendahl, and H. Kittler, "The HAM10000 dataset: A large collection of multi-source dermoscopic images of pigmented skin lesions," 2018. Available from: https://doi.org/10.1038/sdata.2018.161

[10] D. P. Kingma and J. Ba, "Adam: An Approach for Stochastic Optimization," in International Conference on Learning Representations (ICLR), 2015. Available from: https://arxiv.org/pdf/1412.6980

[11] Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC)," 2018. Available from: https://doi.org/10.48550/arXiv.1902.03368

[12] C. Shorten and T. M. Khoshgoftaar, "A survey on Image Data Augmentation for Deep Learning," 2019. Available from: https://doi.org/10.1186/s40537-019-0197-0

[13] N. Codella et al., "Skin Lesion Analysis Toward Melanoma Detection: A Challenge at the 2017 International Symposium on Biomedical Imaging (ISBI), Hosted by the International Skin Imaging Collaboration (ISIC)," 2019. Available from: https://doi.org/10.48550/arXiv.1710.05006