Research report

CNN-Based Method for Benign and Malignant Skin Lesion Classification

✉, **Mehedi Hasan Naim**

[1] mehedi15-14721@diu.edu.bd

Abstract

Skin cancer is one of the most common forms of cancer, with its diagnosis typically involving manual screening and visual examination by dermatologists. However, these traditional methods are often complex, time-consuming, and prone to errors. With the growing incidence and significant mortality rates associated with skin cancer, early detection is crucial for effective treatment. In this study, we present an optimized deep convolutional neural network (DCNN) model specifically designed to classify skin lesions as either benign or malignant, thereby improving diagnostic accuracy. Our approach utilizes a dataset of 3,297 dermoscopic images and incorporates rigorous data preprocessing along with the softmax activation function. We evaluated multiple optimizers—Adam, RMSProp, and SGD—each configured with a learning rate of 0.0001. Our findings indicate that the Adam optimizer outperforms others, achieving an accuracy of 84%, a loss rate of 32%, a recall of 85%, a precision of 85%, an F1-score of 85%, and a ROC-AUC of 83% on the ISIC dataset. This CNN-based approach offers a promising tool for enhancing early skin cancer detection and diagnosis.

Keywords: Skin Cancer; Convolutional Neural Network; Deep Learning; Medical Images

1. Introduction

Skin cancer ranks among the most common types of cancer globally, with millions of new diagnoses occurring annually. This condition arises when normal skin cells experience mutations, resulting in uncontrolled proliferation and tumor formation, which can be classified as either malignant or benign. Malignant tumors possess the ability to metastasize to other regions of the body, while benign tumors remain localized. Timely detection of skin cancer is vital for successful treatment, which may include topical therapies or surgical procedures.

Typically, the diagnosis of skin cancer initiates with a visual examination, often supplemented by dermoscopic evaluation, which can achieve nearly complete curability when the disease is identified early. Nevertheless, accurately diagnosing skin cancer poses significant challenges, even for seasoned dermatologists, due to the wide range of disease presentations. The differences in appearance among various skin cancer types and the constraints of individual expertise further complicate the diagnostic process.

To overcome these obstacles, it is imperative to develop automated systems that facilitate the early and precise detection of skin cancer. Computer vision methodologies, particularly those utilizing deep neural networks such as Convolutional Neural Networks (CNNs), have demonstrated substantial potential in the realm of medical image analysis. These techniques, in conjunction with hybrid machine

learning models, have proven effective in classifying skin lesions and supporting the early identification of skin cancer.

The primary objective of this research is to establish automated tools that aid dermatologists in the precise diagnosis of skin lesions. We have created a CNN model that has been trained and evaluated using the International Skin Image Collaboration (ISIC) archive dataset, enabling the classification of skin lesions into benign and malignant categories.

Related work

Melanoma, a lethal variant of skin cancer, results in the loss of thousands of lives worldwide annually. In recent years, significant efforts have been made to enhance early detection methods to improve survival outcomes. A variety of machine learning approaches have been introduced to improve the segmentation and classification of skin lesions, thereby aiding in the prompt identification of melanoma at its most treatable stage. Ogundokun et al. presented a novel strategy for skin cancer detection by creating a hybrid deep convolutional neural network architecture that combines the advantages of the Xception and MobileNetV2 models. They tackled the issue of imbalanced datasets through data augmentation techniques and utilized transfer learning to address the challenge of insufficient labeled data. Their approach yielded remarkable results, achieving an accuracy of 97.56%, an area under the curve of 97.00%, 100% sensitivity, 93.33% precision, a 96.55% F1 score, and a minimal false positive rate of 0.0370. Nigar et al. designed a CNN model capable of identifying six distinct skin conditions, including actinic keratosis, benign keratosis, melanoma, basal cell carcinoma, insect bites, and acne. Their model attained an impressive accuracy of 97%, with precision, recall, and F1 scores all reaching 91%. In a similar vein, Jaisakthi et al. introduced an EfficientNet-based architecture for transfer learning, concentrating on the automatic differentiation between melanoma and non-melanoma skin cancers. By optimally fine-tuning EfficientNet-B6 with the Ranger optimizer, they achieved an AUC-ROC score of 0.9681.

Nigar et al. developed a convolutional neural network (CNN) model capable of identifying six distinct skin conditions, which include actinic keratosis, benign keratosis, melanoma, basal cell carcinoma, insect bites, and acne. Their model demonstrated a remarkable accuracy rate of 97%, with precision, recall, and F1 scores all reaching 91%. In a similar vein, Jaisakthi et al. presented an EfficientNet-based architecture aimed at transfer learning, specifically targeting the automatic differentiation between melanoma and non-melanoma skin cancers. By optimally fine-tuning EfficientNet-B6 with the Ranger optimizer, they achieved an AUC-ROC score of 0.9681. Hameed et al. introduced a classification technique designed to categorize skin lesions into seven unique classes. Their methodology incorporated data augmentation and image preparation techniques, resulting in an impressive accuracy rate of 92.5%. Their comparative analysis underscored the effectiveness and potential advantages of their approach over previously published methods utilizing the same dataset. Additionally, Bozkurt investigated the influence of data augmentation on the performance of a pre-trained Inception-Resnet-v2 model for skin lesion classification. The study reported a peak accuracy of 95.09% with an augmented dataset, in contrast to an accuracy of 83.59% with the original, non-augmented dataset, thereby highlighting the substantial benefits of data augmentation.

Saifan et al. developed a system that employs a pre-trained deep convolutional neural network (CNN) to distinguish among six different skin conditions: acne, athlete's foot, chickenpox, eczema, skin cancer, and vitiligo. They assembled a dataset comprising 3,000 images sourced from various online platforms, designating 90% for training purposes and 10% for evaluating accuracy on new data. To enhance user accessibility, they also created an Android application that serves as an interface for the system. The system demonstrated an accuracy of up to 81.75%, indicating its effectiveness in classification tasks.

Gangatharan et al. introduced an automated method for classifying skin cancers, including Actinic Keratosis, Basal Cell Carcinoma, Dermatofibroma, and Melanoma, utilizing deep learning techniques applied to the ISIC dataset. They employed both Adam and SGD optimizers during the training process, achieving an impressive accuracy of 98.4% with a loss rate of 7.7% after 500 epochs. The model was subsequently integrated into a Python extension called Streamlit, which offers a user-friendly interface for image classification.

Albawi et al. concentrated on optimizing the architecture of CNNs and identifying suitable parameters for skin cancer prediction. Their methodology resulted in an outstanding accuracy rate of 98.5%, surpassing other existing techniques. In a similar vein, Kahia et al. adapted the VGG16 and InceptionV3 architectures for the classification of skin lesions, utilizing a dataset categorized into three classes: melanoma, nevus, and seborrheic keratosis. Their modified VGG16 model achieved an accuracy rate of 73.33%, underscoring the potential of their approach in skin lesion classification.

Xinrong Lu and colleagues devised an automated technique for the diagnosis of skin cancer utilizing dermoscopy images sourced from the HAM10000 dataset. Their methodology, which is founded on an enhanced XceptionNet model featuring the swish activation function and depth-wise separable convolutions, yielded remarkable performance metrics, including a perfect accuracy of 100%, a sensitivity of 94.05%, a precision of 97.07%, and an F1 score of 95.53%. Additionally, Xikun Wang investigated a hybrid approach that integrated multiple models, such as CNN, GANs, KNN, and SVM, to classify skin lesions as benign or malignant using the ISIC dataset. The findings revealed that deep learning models, particularly CNN and GANs, achieved notably superior accuracy, with CNN attaining approximately 92% accuracy in distinguishing between benign and malignant lesions. GANs also exhibited commendable accuracy, ranging from 85% to 90%, whereas KNN and SVM models demonstrated lower accuracy levels, with KNN achieving between 70% to 80% and SVM around 80%.

Materials and Methods

A. Dataset

In this research, we employed the ISIC (International Skin Image Collaboration) archive dataset, which encompasses a diverse range of skin cancer types. Our primary focus was on the classification of benign versus malignant lesions. The dataset comprised 1,800 benign images, with 1,440 designated for training purposes and 360 reserved for testing. Additionally, there were 1,497 malignant images, of which 1,197 were utilized for training and 300 for testing. This distribution of data is depicted in Figure 1. All images were resized to a standardized dimension of (224x224x3) pixels in RGB format and subsequently transformed into NumPy arrays. Given that the images lacked pre-existing

labels, we generated and assigned labels as illustrated in Figure 2. These labeled images were then amalgamated into a comprehensive training set, which was thoroughly shuffled for further analysis and model training.

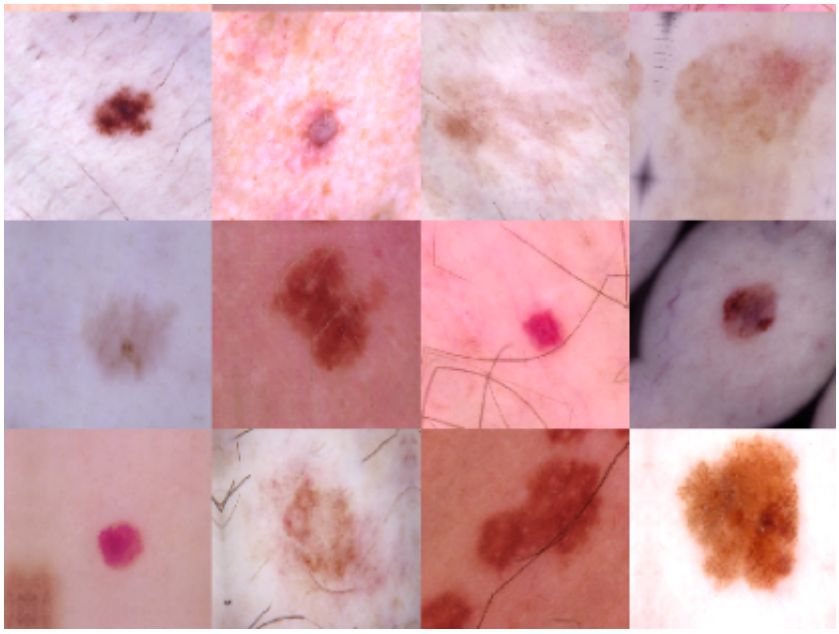


Figure 1. Sample images

B. Preprocessing

Prior to the development and training of our deep convolutional neural network (CNN), we executed several essential preprocessing steps:

Resizing Images: Each image in the dataset was resized to a consistent dimension of 224x224 pixels with three color channels (RGB). The dataset was then divided, with approximately 80% allocated for training and the remaining 20% reserved for testing. This division facilitated effective training of the CNN and allowed for performance evaluation on previously unseen data.

Normalization of Dataset: Normalization is a vital preprocessing procedure in machine learning and data analysis, as it standardizes the data scale and enhances model efficacy. Various normalization techniques can be applied, including min-max normalization, z-score normalization, and decimal scaling. In this study, we normalized the dataset by dividing each pixel value by 255, which corresponds to the maximum grayscale value of the images. This approach ensured that all pixel values were confined within the range of 0 to 1, fostering consistency throughout the dataset and enhancing the performance of the CNN model.

Data Augmentation:

Data augmentation was a crucial step in expanding the size and diversity of our dataset, ultimately improving the performance and robustness of our model. We employed a variety of augmentation techniques, including color adjustments and affine transformations. These transformations included

operations such as shearing, zooming, rotation, and horizontal flipping. These methods are widely recognized and effective for enhancing the generalization capability of models by simulating variations that may occur in real-world data.

Augmentation Setting	Range/Value
Shear Range	0.2
Zoom Range	0.2
Width Shift Range	0.2
Height Shift Range	0.2
Rotation Range	40 degrees
Horizontal Flip	True

Table 1. Augmentation setting

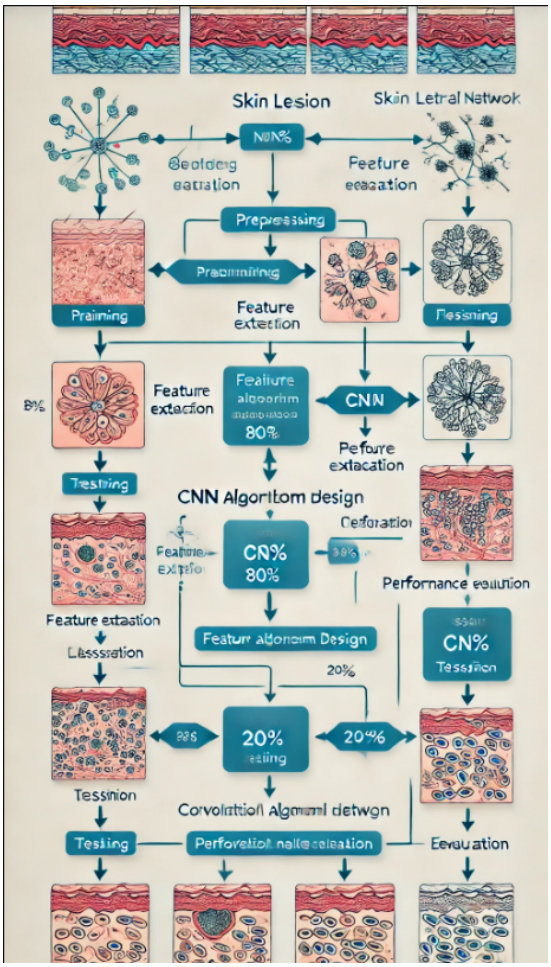


Figure 2. Flowchart of the proposed method

The Convolutional Neural Network (CNN) model

We have developed is tailored for the classification of images into "benign" or "malignant" categories, utilizing markers derived from the image metadata. The dataset employed comprises a combination of both cancerous and non-cancerous images.

Our model consists of three fundamental layers, each fulfilling a specific role:

- Convolutional Layer: This layer serves as the essential component of the CNN, tasked with scanning the input images through the application of filters or kernels via convolution operations. By modifying the filter size and stride, this layer generates an output known as a feature map or activation map, which emphasizes significant features within the image. In our implementation, we utilized 2D convolution with a 3x3 kernel and 64 filters, which were instrumental in detecting critical image features (refer to Figure 4 for an illustration of a convolution operation).

- Pooling Layer: Following the convolutional layer, the pooling layer is employed to diminish the dimensionality of the feature maps, thereby streamlining the subsequent processing stages. Common techniques for pooling include max pooling and average pooling, where either the maximum or average value is selected, respectively. In our model, we implemented MaxPooling2D with a 2x2 kernel, effectively reducing the spatial resolution while retaining essential features (see Figure 5 for an example of a max-pooling operation).

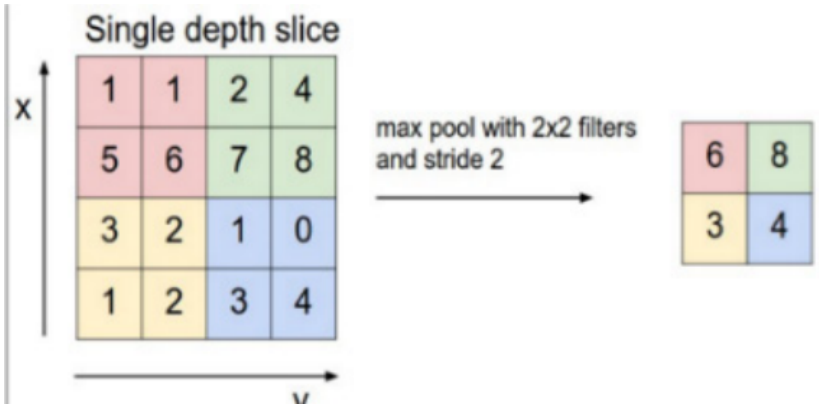


Figure 3: Example of max pooling operation

- Fully Connected Layer: This layer is positioned at the conclusion of the CNN architecture. Once the feature maps are transformed into one-dimensional data, the fully connected layer links each input to all neurons within the layer. The softmax activation function is subsequently applied to produce a probability distribution across N dimensions, where N represents the number of classes to be classified.

These three layers collaborate to extract, process, and classify relevant features from the input images, ultimately facilitating precise classification based on the patterns learned.

In order to reduce errors throughout the training phase, we employed the Rectified Linear Unit (ReLU) activation function, which functions by converting all negative inputs to zero, as expressed in the following formula:

$f(x) = 0$ for $x < 0$

$f(x) = x$ for $x \geq 0$ (1)

This activation function permits positive inputs to remain unchanged, thereby enhancing the model's capacity to learn intricate patterns. Furthermore, a dropout layer was integrated to mitigate the risk of overfitting during the training process. This layer operates by randomly "dropping out" a designated proportion of neurons in each training epoch, ensuring that the network does not become overly dependent on any individual neuron. All the layers mentioned were implemented using the sequential model of the CNN architecture.

Training and Evaluation

In the development of our Convolutional Neural Network (CNN) model, we employed Keras, a Python-based deep learning API. All methodologies utilized in this research were executed through Keras. The dataset for this investigation comprised 2,637 images designated for training and 660 images allocated for testing.

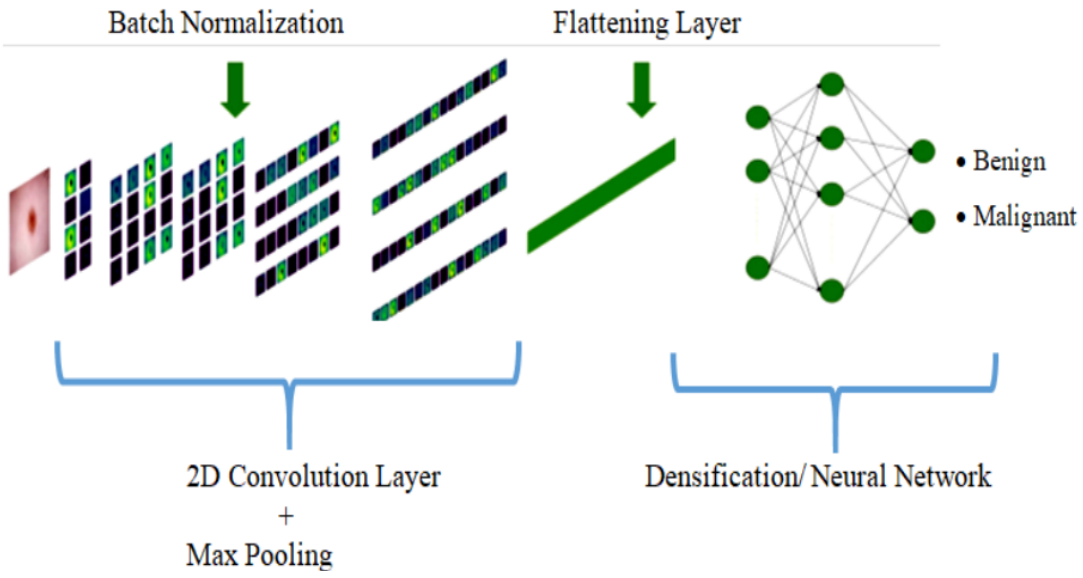


Figure 4: Proposed model for training & testing

Throughout the training phase, we meticulously assessed a range of hyperparameters to identify the optimal combination that would enhance the model's performance. This process included testing various configurations to refine the model's functionality and improve its predictive accuracy. A crucial element of this endeavor was the selection of the most effective hyperparameters, which are outlined in the table below:

Parameter	Value
Epochs	50
Batch size	32
Optimizer	Adam, RMSProp, SGD
Activation function	Softmax
Learning rate	0.0001
Drop out	0.1
Verbose	2
Callbacks	ReduceLROnPlateau

Table 2: Parameter & value list

In our analysis, employing a limited number of epochs may not allow the model to adequately grasp the underlying patterns, which could lead to underfitting. On the other hand, training the model for an excessive number of epochs may result in overfitting, where the model captures not only significant patterns but also irrelevant noise present in the data. It is essential to find an appropriate balance in the number of epochs to ensure optimal model performance. Training the model for 50 epochs facilitates multiple iterations over the complete dataset, thereby increasing the likelihood of acquiring knowledge from a varied set of samples and potentially guiding the model towards an optimal solution. The batch size, which specifies the number of training samples processed in each iteration, plays a crucial role in influencing the training process and the performance of the resulting model.

To assess and visualize the effects of various optimizers on model performance, we implemented three different optimizers: Adam, RMSProp, and SGD. The Softmax activation function was utilized in the output layer for classification purposes, transforming the input vector into a probability distribution across the various classes. Furthermore, a dropout layer was incorporated to mitigate the risk of overfitting.

The callback ReduceLROnPlateau was employed to dynamically modify the learning rate during training based on specific criteria, primarily aimed at enhancing training stability and convergence.

Metrics for Performance Evaluation

To measure the effectiveness of the models, we applied the following metrics: accuracy, precision, recall, and F1-score. These metrics were calculated using the respective formulas:

- **Accuracy**: (TP+TN)/(TP+FP+TN+FN)
- **Recall/Sensitivity**: TP/(TP+FN)
- **Precision**: TP/(TP+FP)
- **F1-Score**: 2×((Precision×Recall)/(Precision+Recall))

Where:

- **TP (True Positives)**: Correct predictions of the positive class.

- **FN (False Negatives)**: Incorrect predictions where the actual class is positive but predicted as negative.
- **FP (False Positives)**: Incorrect predictions where the actual class is negative but predicted as positive.

Results and Discussion

The dataset comprised 2,637 training images, consisting of 1,440 benign and 1,197 malignant samples, along with 660 test images, which included 360 benign and 300 malignant samples. These images were obtained from the International Skin Imaging Collaboration (ISIC) archive. The model was trained utilizing the Adam, RMSProp, and SGD optimizers, with a learning rate set at 0.0001 over a span of 50 epochs.

During the training phase, a consistent reduction in the loss function was noted for both the training and validation datasets as the epochs progressed. Concurrently, there was a gradual enhancement in both training and validation accuracy, signifying an improvement in the model's efficacy over time.

An analysis of the accuracy and loss metrics for the three optimizers indicated that the Adam optimizer outperformed the others in our binary skin disease classification task. By the conclusion of epoch 50, the Adam optimizer achieved a training accuracy of 90% and a validation accuracy of 84%, with training and validation losses recorded at 22% and 34%, respectively. The overall accuracy and loss of the network were determined to be 84% and 32%.

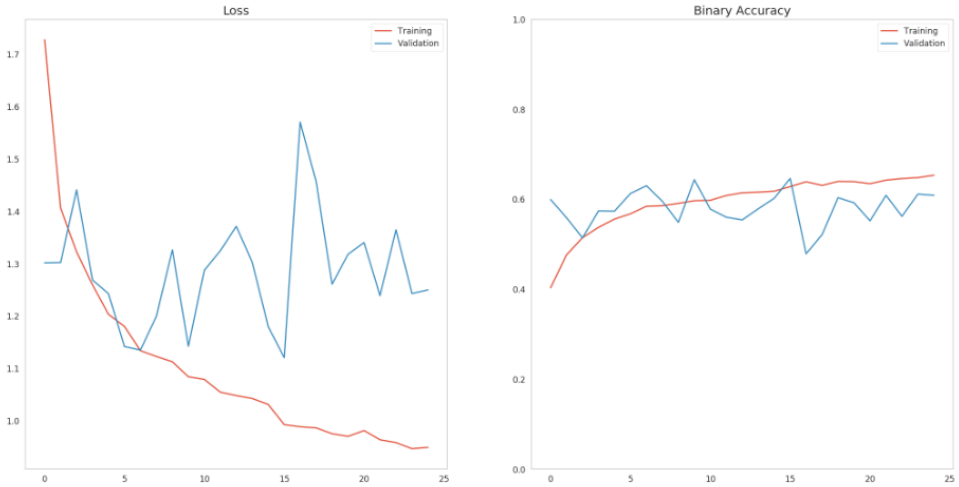


Figure 5: Loss & binary accuracy

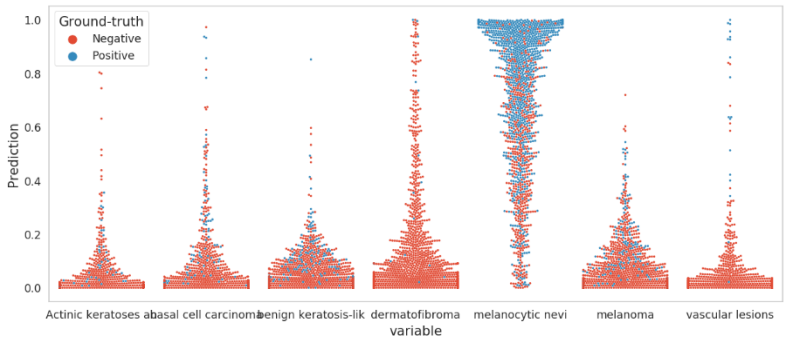


Figure 6: Detailed performance by group.

In contrast, the RMSProp optimizer reached its maximum training accuracy of 85% by the 50th epoch, while the peak validation accuracy was recorded at 79% by the 5th epoch. However, following this point, the validation accuracy experienced a decline, ultimately stabilizing at 63%, which suggests overfitting. The training loss for the model was noted at 30%, and the validation loss reached a plateau of 50% around the 35th epoch, culminating in an overall accuracy of 68% with a corresponding loss of 47%. The SGD optimizer, on the other hand, demonstrated an overall accuracy of 71% with a loss of 59%, showing no indications of overfitting.

Confusion Matrix

The confusion matrix illustrates the distribution of correct and incorrect predictions made by the model:

- Adam Optimizer: Accurately identified 307 benign and 247 malignant images, while misclassifying 53 benign and 53 malignant images.

- RMSProp Optimizer: Correctly classified 356 benign and 94 malignant images, with 4 benign and 206 malignant images misclassified.

- SGD Optimizer: Successfully predicted 250 benign and 217 malignant images, but misclassified 110 benign and 83 malignant images.

Recall, Precision, and F1-score

Recall evaluates the model's capability to accurately identify all relevant instances of a specific class, with a high recall indicating effectiveness in capturing positive cases. Precision measures the accuracy of the model's positive predictions, where a high precision signifies that positive class predictions are generally correct. The F1-Score integrates both precision and recall into a singular metric, offering a balanced assessment of performance, especially in scenarios with imbalanced class distributions.

ROC Curve

The Receiver Operating Characteristic (ROC) curve assesses the classification model's performance by plotting the true positive rate against the false positive rate across various thresholds. Among the

optimizers, the Adam optimizer achieved the highest Area Under the Curve (AUC) score of 0.83, surpassing RMSProp, which recorded an AUC of 0.65, and SGD, which had an AUC of 0.70.

Predictions on Test Data

The Adam optimizer demonstrated a well-rounded performance in distinguishing between benign and malignant classes. In contrast, the RMSProp optimizer excelled in the classification of benign images, while the SGD optimizer proved to be more effective in identifying malignant images.

Conclusion

Skin cancer represents a widespread and critical health issue, making early detection essential for successful treatment. The integration of machine learning (ML) has profoundly influenced dermatology, particularly in facilitating early detection and enhancing the diagnostic precision of skin cancer. This study aims to develop an effective system that serves as a real-time clinical diagnostic tool for skin cancer identification.

In this investigation, a tailored Convolutional Neural Network (CNN) model was utilized to categorize skin cancer images into benign and malignant groups. The model was trained and evaluated using a dataset comprising 3,297 images sourced from the International Skin Image Collaboration archive.

The findings indicate that the CNN model, optimized with the Adam optimizer, achieved commendable performance, attaining an accuracy of 84% and a loss rate of 0.32. Additionally, it surpassed models utilizing the RMSProp and SGD optimizers, exhibiting superior metrics in Recall, Precision, F1-score, and ROC-AUC.

In the future, we intend to refine these models to facilitate classification across various skin cancer types, thereby enhancing their clinical applicability.

References

1. Aveiro-Róbalo TR, Pérez-Del-Vallín V. Gamification for well-being: applications for health and fitness. Gamification and Augmented Reality 2023;1:16–16. https://doi.org/10.56294/gr202316.
2. Castillo JIR. Aumented reality im surgery: improving precision and reducing ridk. Gamification and Augmented Reality 2023;1:15–15. https://doi.org/10.56294/gr202315.
3. « Convolutional Neural Network Tutorial [Update] », Simplilearn.com. Disponible sur: https://www.simplilearn.com/tutorials/deep-learning-tutorial/convolutional-neural-network.
4. N. Chuchu, J. Dinnes, Y. Takwoingi, R.N. Matin, S.E. Bayliss, C. Davenport, J.F. Moreau, O. Bassett, K. Godfrey, C. O'Sullivan, F.M. Walter, R.Motley, J.J. Deeks, and H.C. Williams, “Teledermatology for diagnosing skin cancer in adults”. Cochrane Database of Systematic Reviews, Dec. 2018, doi:10.1002/14651858.cd013193.
5. K. Team, « Keras documentation: RMSprop ». Disponible sur: https://keras.io/api/optimizers/rmsprop/.

6. Hameed, M. Umer, U. Hafeez, H. Mustafa, A. Sohaib, M. Abubakar Siddique and H. Ahmad Madni, “Skin lesion classification in dermoscopic images using stacked convolutional neural network,” Journal of Ambient Intelligence and Humanized Computing, vol. 14, no. 4, pp. 3551–3565, Apr. 2023, doi: 10.1007/s12652-021-03485-2.
7. K. Team, « Keras documentation: ReduceLROnPlateau ». Disponible sur: https://keras.io/api/callbacks/reduce_lr_on_plateau/.
8. F. Bozkurt, “Skin lesion classification on dermatoscopic images using effective data augmentation and pre-trained deep learning approach”, Multimedia Tools and Applications, vol. 82, pp. 8985–19003, Nov. 2022, doi: 10.1007/s11042-022-14095-1.
9. M.S. Ali, M.S. Miah, J. Haque, M.M. Rahman, and M.K. Islam, “An enhanced technique of skin cancer classification using deep convolutional neural network with transfer learning models”, Machine Learning with Applications, vol. 5, article no. 100036, Sept. 2021, doi: 10.1016/j.mlwa.2021.100036.
10. « Keras », Keras. Disponible sur: https://keras.io/.
11. Maglogiannis and C.N. Doukas, “Overview of Advanced Computer Vision Systems for Skin Lesions Characterization”, IEEE Transactions on Information Technology in Biomedicine, vol. 13, no. 5, pp. 721–733, Sept. 2009, doi:10.1109/titb.2009.2017529.
12. R.O. Ogundokun, A. Li, R.S. Babatunde, C. Umezuruike, P.O. Sadiku, A.T. Abdulahi, and A.N. Babatunde, “Enhancing Skin Cancer Detection and Classification in Dermoscopic Images through Concatenated MobileNetV2 and Xception Models”. Bioengineering, vol. 10, no. 8, article no. 979, Aug. 2023, doi: 10.3390/bioengineering10080979.
13. N. Nigar, A. Wajid, S. Islam, and M.K. Shahzad, “SKIN CANCER CLASSIFICATION: A DEEP LEARNING APPROACH ”, Pak. J. Sci., vol. 75, no. 02, article no. 02, juill. 2023, doi: 10.57041/pjs.v75i02.851.
14. « ISIC | International Skin Imaging Collaboration », ISIC. Consulté le: 23 octobre 2023. [En ligne]. Disponible sur: https://www.isic-archive.com.
15. Castillo-Gonzalez W, Lepez CO, Bonardi MC. Augmented reality and environmental education: strategy for greater awareness. Gamification and Augmented Reality 2023;1:10–10. https://doi.org/10.56294/gr202310.
16. « Softmax Activation Function: Everything You Need to Know | Pinecone ». Disponible sur: https://www.pinecone.io/learn/softmax-activation/.
17. « Adam », Keras. Disponible sur: https://keras.io/api/optimizers/adam.
18. K. Mandal, P. K. D. Sarma, and S. Dehuri, “Image-based Skin Disease Detection and Classification through Bioinspired Machine Learning Approaches”. International Journal on Recent and Innovation Trends in Computing and Communication, vol. 12, no. 1, Sept. 2023, pp. 85–94, doi:10.17762/ijritcc.v12i1.7914.
19. « Convolutional Neural Networks (CNN): Step 1- Convolution Operation - Blogs - SuperDataScience | Machine Learning | AI | Data Science Career | Analytics | Success ». Disponible sur: https://www.superdatascience.com/blogs/convolutional-neural-networks-cnn-step-1-convolution-operation.
20. S. Albawi, M. H. Arif, and J. Waleed, “Skin cancer classification dermatologist-level based on deep learning model”, Acta Scientiarum. Technology, vol. 45, pp. e61531-e61531, 2023, doi: 10.4025/actascitechnol.v45i1.61531.

21. Auza-Santiváñez JC, Díaz JAC, Cruz OAV, Robles-Nina SM, Escalante CS, Huanca BA. Bibliometric Analysis of the Worldwide Scholarly Output on Artificial Intelligence in Scopus. Gamification and Augmented Reality 2023;1:11–11. https://doi.org/10.56294/gr202311.
22. Kahia, A. Echtioui, F. Kallel, and A. Ben Hamida, “Skin Cancer Classification using Deep Learning Models”, Proc. The 14th International Conference on Agents and Artificial Intelligence, Vienna, Austria: SCITEPRESS - Science and Technology Publications, 2022, pp. 554-559. doi: 10.5220/0010976400003116.
23. G. N, M. S, R. R, S. V, S. R. K. V, and S. K. B, “Enhancing Skin Cancer Diagnosis with Deep Learning-Based Classification”, Int. J. Recent Innov. Trends Comput. Commun, vol. 11, no 5s, pp. 105-111, Mai. 2023, doi: 10.17762/ijritcc.v11i5s.6634.
24. J. S M, M. P, C. Aravindan, and R. Appavu, “Classification of skin cancer from dermoscopic images using deep neural network architectures”, Multimedia Tools and Applications, vol. 82, no 10, pp. 15763-15778, 2023, doi: 10.1007/s11042-022-13847-3.
25. WHO. Cancer. Available https://www.who.int/health-topics/cancer#tab=tab_1)
26. R. Saifan and F. Jubair, “Six skin diseases classification using deep convolutional neural network”, Int. J. Electr. Comput. Eng. IJECE, vol. 12, no 3, article no. 3, juin 2022, pp. 3072-3082, doi: 10.11591/ijece.v12i3.pp3072-3082.
27. X. Wang, “Deep Learning-based and Machine Learning-based Application in Skin Cancer Image Classification”, J. Phys. Conf. Ser., vol. 2405, article no. 012024, Déc. 2022, doi: 10.1088/1742-6596/2405/1/012024.
28. X. Lu et F. Abolhasani Zadeh, “Deep Learning-Based Classification for Melanoma Detection Using XceptionNet”, J. Healthc. Eng., vol. 2022, article no .2196096, Mars 2022, doi: 10.1155/2022/2196096.
29. Joel, « When is max pooling exactly applied in convolutional neural networks? », Artificial Intelligence Stack Exchange. Disponible sur: https://ai.stackexchange.com/q/17857.
30. K. Team, « Keras documentation: SGD ». Disponible sur: https://keras.io/api/optimizers/sgd/.