Undergraduate Conference on Intelligent Computing and Systems (UCICS 2025)
26-27 February, 2025; Varendra University, Rajshahi, Bangladesh

Skin Cancer Detection Using Deep Learning

Amimoy Kumar Biswas
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
amimoysammo@gmail.com

Samila Tasrin
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
samilatasrin@gmail.com

Nafisa Anjum
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
n.anjum.bd@gmail.com

Md. Nahid Hasan
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
nahid7532140@gmail.com

Tasfiah Islam
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
tasfiahislam18@gmail.com

Md. Taufiq Khan
Computer Science and Engineering
Varendra University
Rajshahi, Bangladesh
khantaufiq2001@gmail.com

***Abstract*—Skin cancer detection has become a crucial area of research due to its increasing prevalence and the need for early diagnosis. In this study, we leverage deep learning techniques to classify skin cancer images, focusing on a comparative analysis of feature extractors and classifiers. Models such as DenseNet121, ResNet50, MobileNetV2, and DenseNet169 were paired with various classifiers, including Logistic Regression, XGBoost, Random Forest, and custom neural networks. DenseNet121 combined with a custom classifier achieved the highest performance with an accuracy of 93.5% and an AUC of 0.9842, demonstrating its potential for practical applications.**

***Index Terms*—Skin Cancer, Image Classification, Transfer Learning, Deep Learning.**

I. Introduction

Skin cancer is one of the most common cancers worldwide, with melanoma being the deadliest subtype. Early detection is critical for effective treatment. Advances in deep learning have enabled the development of automated systems capable of accurately identifying cancerous lesions.Melanoma, while accounting for a smaller percentage of skin cancer cases, is responsible for the majority of skin cancer-related deaths. In 2024, it's estimated that about 8,290 people will die from melanoma in the U.S. [1]. Early detection of melanoma significantly improves survival rates. The five-year relative survival rate for localized melanoma is over 99% in the U.S. [1].Deep learning has revolutionized medical imaging, with models like DenseNet, ResNet, and MobileNet excelling in feature extraction and classification tasks. Complementary techniques, such as CLAHE for preprocessing and SMOTE for data balancing, further enhance model performance.This research provides a comprehensive evaluation of deep learning-based feature extractors and classifiers for skin cancer detection. DenseNet121, combined with a custom classifier, achieved state-of-the-art performance. The study also highlights the importance of preprocessing techniques and class balancing in enhancing model accuracy.This study aims to evaluate the performance of various feature extractors and classifiers in detecting skin cancer from dermoscopic images.

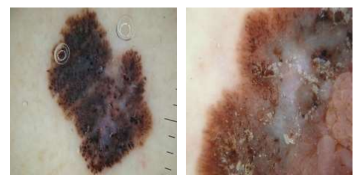

Fig. 1. Malignant

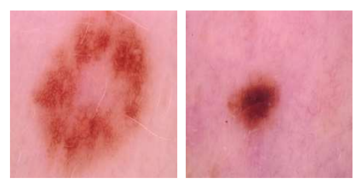

Fig. 2. Benign

A. Related Work

Majtner et al. [2] detected skin lesions using CNN techniques for feature extraction and preprocessing instead of classification. A publicly available ISIC dataset with 900 training samples and 379 test samples (benign and malignant classes) was used. Preprocessing steps included grayscale conversion and image downsampling. AlexNet with binary masking and bounding box methods was employed for feature extraction, followed by LDA techniques for feature reduction. The KNN classifier achieved the best precision (86%) and specificity (99.9%).

Vipin et al. [3] developed a two-stage system for melanoma diagnosis to minimize human interference, making it less error-prone and time-efficient. The system used a reduced ISIC dataset of 7,353 images. The segmentation stage utilized a symmetric U-Net architecture, comprising contracting/downsampling, bottleneck, and expanding/upsampling paths. The classification stage employed a deep residual network integrating CNN and RNN techniques. The system achieved an accuracy of 88.7% and recall of 91%, using weighted binary cross-entropy as the loss function. Nasr-

Esfahani et al. [4] developed a custom CNN for preprocessing, feature extraction, and classification. Using 170 non-dermoscopic images from the UMCG archive, the dataset was augmented to 6,120 images. Preprocessing included illumination correction, mask generation, and Gaussian filtering. The CNN had alternate convolutional and max-pooling layers followed by a fully connected layer. The system classified melanoma and nevus with an accuracy of 81% and specificity of 80%.

Attia et al. [5] employed a fully CNN with autoencoder-decoder architectures like FCN and SegNet. Recurrent neural networks and LSTM architecture were integrated to improve results. Training used 900 lesion images, and testing involved 375 samples from the ISBI 2016 challenge dataset. The model achieved an accuracy of 98% and specificity of 94%.

Mukherjee et al. [6] proposed the CMLD (CNN Malignant Lesion Detection) architecture, combining two datasets: MEDNODE (1,700 images) and Dermofit (13,000 images). The architecture included three 2D convolutional layers, two max-pooling layers, and a softmax layer for prediction. Accuracies of 90.14% and 90.58% were achieved on individual datasets, and 83.07% on the combined dataset.

II. Methodology

A. Dataset

The dataset used in this study was sourced from the Melanoma skin cancer dataset, a publicly available dataset of 1000 images from Kaggle, comprising labeled images of skin lesions. The images were preprocessed to ensure uniformity and quality, including resizing to 224×224 pixels and applying contrast enhancement techniques such as CLAHE. Table 1 presents the dataset, which comprises 1,000 samples, evenly divided into two categories: 500 healthy data points and 500 unhealthy data points. This balanced distribution ensures a fair evaluation of models for classification tasks.

TABLE I
SUMMARY OF HEALTHY AND UNHEALTHY DATA IN THE DATASET.

Category	**Count**
Healthy Data	500
Unhealthy Data	500
Total	**1000**

B. Preprocessing

The images were passed through preprocessing steps to improve feature extraction:

- **CLAHE (Contrast Limited Adaptive Histogram Equalization):** It provides local histogram equalization to improve the exposure of the image.
- **Laplacian Filtering:** The Laplacian kernel is used to sharpen the image.
- **Normalization:** Pixel intensity values are scaled down to the range [0, 1].
- **Data Balancing:** SMOTE is used to generate synthetic data to compensate for class imbalance.

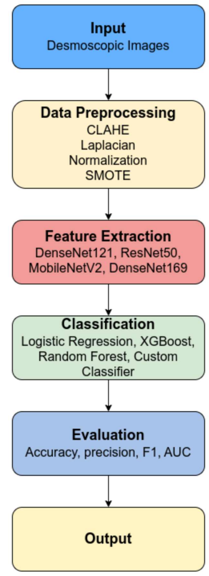

Fig. 3. Method Diagram

C. Feature Extractors

We utilized the following pre-trained SOTA CNN models as feature extractors:

- DenseNet121
- ResNet50
- MobileNetV2
- DenseNet169

D. Classifiers

Extracted features were classified using the following algorithms:

- Logistic Regression
- XGBoost
- Random Forest
- Custom Neural Network Classifier

E. Evaluation Metrics

Performance was assessed using several key metrics: **Accuracy**, which is the ratio of correctly predicted instances to the total instances, providing a general measure of the model's correctness; **Precision**, which measures the proportion of positive predictions that are actually correct, is important when false positives are costly; **Recall**, which measures the ability of the model to correctly identify all actual positive instances, is crucial when false negatives are costly; **F1 Score**, the harmonic mean of precision and recall, offering a balanced measure of performance, especially in cases of class imbalance; and **Area Under the Curve (AUC)**, which reflects the model's ability to distinguish between classes, with a higher AUC indicating better classification performance.

III. Results

Table 2 presents the comparative performance of feature extractors paired with different classifiers. DenseNet121 with a custom classifier achieved the best results across all metrics.

TABLE II
PERFORMANCE METRICS FOR DIFFERENT FEATURE EXTRACTORS AND CLASSIFIERS

Feature Extractor	Classifier	Accuracy	Precision	Recall	F1 Score	AUC
DenseNet121	Custom Classifier	0.935	0.9354	0.9350	0.9350	0.9842
	Logistic Regression	0.94	0.9888	0.8900	0.9368	0.9706
	XGBoost	0.935	0.97808	0.8900	0.9319	0.9675
	Random Forest	0.93	0.9886	0.8700	0.9255	0.9657
ResNet50	Custom Classifier	0.6600	0.7051	0.5500	0.6180	0.7120
	Logistic Regression	0.8850	0.8811	0.8900	0.8855	0.9454
	XGBoost	0.8700	0.8700	0.8700	0.8700	0.9543
	Random Forest	0.8200	0.8330	0.8000	0.8163	0.9177
MobileNetV2	Custom Classifier	0.9100	0.9659	0.8500	0.9043	0.9614
	Logistic Regression	0.8950	0.9247	0.8600	0.8911	0.9566
	XGBoost	0.9300	0.9255	0.8700	0.8969	0.9600
	Random Forest	0.8850	0.9052	0.8600	0.8820	0.9489
DenseNet169	Custom Classifier	0.9300	0.9583	0.9200	0.9388	0.9852
	Logistic Regression	0.9400	0.9680	0.9100	0.9381	0.9797
	XGBoost	0.9400	0.9583	0.9200	0.9388	0.9852

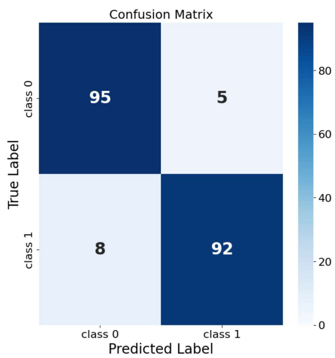

Fig. 4. Confusion Matrix

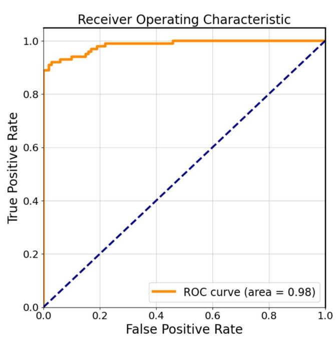

Fig. 5. Receiver Operating Characteristic

IV. DISCUSSION

DenseNet121 outperformed other feature extractors due to its efficient handling of spatial hierarchies in image data. DenseNet169 and MobileNetV2 also delivered robust results, with the latter being suitable for resource-constrained environments. ResNet50 showed relatively low performance, suggesting that it may require further fine-tuning or augmentation strategies.

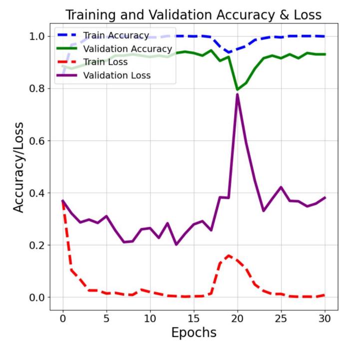

Fig. 6. Training and Validation Accuracy & loss

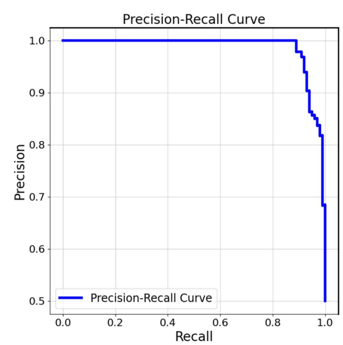

Fig. 7. Precision-Recall Curve

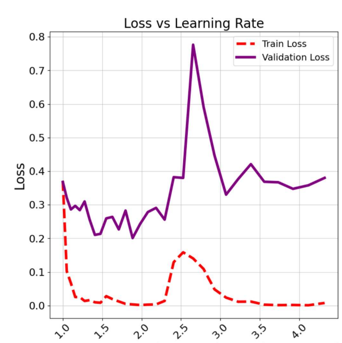

Fig. 8. Loss vs Learning Rate

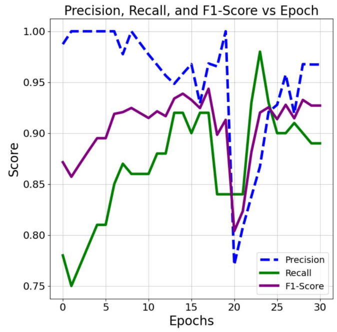

Fig. 9. Precision, Recall, and F1-Score Epoch

V. Conclusion

This study highlights the potential of deep learning for skin cancer detection. DenseNet121 combined with a custom classifier provides a reliable solution, achieving high accuracy and AUC. Future work will focus on integrating explainability tools like SHAP and LIME to improve model transparency and trustworthiness.

References

[1] Cancer.org, "Melanoma survival rates," https://www.cancer.org/cancer/types/melanoma-skin-cancer/detection-diagnosis-staging/survival-rates-for-melanoma-skin-cancer-by-stage.html, 2025.

[2] H. J. Majtner T, Yildirim-Yayilgan S, "Optimised deep learning features for improved melanoma detection," *Multimed Tools Appl*, vol. 78, no. 9, pp. 11 883–11 903, 2018. [Online]. Available: https://link.springer.com/article/10.1007/s11042-018-6734-6

[3] V. Vipin, M. Nath, V. Sreejith *et al.*, "Detection of melanoma using deep learning techniques: A review," *Proceedings of the 2021 International Conference on Communication, Control and Information Sciences (ICCISc)*, 2021. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9484861

[4] E. Nasr-Esfahani, S. Samavi, N. Karimi *et al.*, "Melanoma detection by analysis of clinical images using convolutional neural networks," *Proceedings of the 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2016. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7590963

[5] M. Attia, M. Hossny, S. Nahavandi *et al.*, "Skin melanoma segmentation using recurrent and convolutional neural networks," *Proceedings of IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, 2017. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7950522

[6] S. Mukherjee, A. Malu, A. Balamurali *et al.*, "Twisent: A multistage system for analyzing sentiment in twitter," *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM '12)*, 2012. [Online]. Available: https://dl.acm.org/doi/abs/10.1145/2396761.2398684