Deep Learning for Skin Lesion Detection
A CNN Approach

Jaleel Johnson[1], Zy'Aier Frazzier[1], Tyler Brown[1]
Xiaomao Liu[2], Janmejay Mohanty[2], Biswajit Biswal[2], Nikunja Swain[2]
National Center of Academic Excellnce in Cybersecurity (NCAE-C)
Computer Science and Mathematics Department
South Carolina State University
1 – Students 2- Faculty Mentors

INTRODUCTION

Skin cancer is one of the most prevalent forms of cancer globally, with millions of new cases diagnosed each year. The three most common types include melanoma, basal cell carcinoma, and squamous cell carcinoma—each differing in severity, progression, and treatment requirements. According to the World Health Organization, over 3 million non-melanoma and approximately 132,000 melanoma skin cancers are reported worldwide annually. Early detection plays a critical role in improving treatment effectiveness and patient outcomes. However, the diagnostic process is often resource-intensive and requires specialized expertise from trained dermatologists. A skin lesion—defined as an abnormal growth or change on the skin's surface—is a primary indicator of potential skin cancer. Accurately identifying the type of lesion remains a significant challenge, particularly in settings with limited access to expert care.

To address these challenges, there is a growing demand for automated diagnostic tools that can assist healthcare professionals by rapidly and accurately classifying skin lesion images. This project employs deep learning, specifically convolutional neural networks (CNNs), to build a model capable of detecting and classifying skin cancer using dermoscopic images. By leveraging large annotated datasets and advanced machine learning techniques, CNNs present a promising solution for enhancing diagnostic accuracy, increasing efficiency, and expanding access to early skin cancer screening in underserved or low-resource regions.

OBJECTIVE

The goal of this project is to develop a deep learning model capable of accurately classifying various types of skin cancer using dermoscopic image data. Leveraging the HAM10000 dataset, we trained and evaluated a convolutional neural network (CNN) to distinguish between seven categories of skin lesions, with a target validation accuracy of 85% or higher. Achieving this level of performance is clinically significant, as it is comparable to or even exceeds the diagnostic accuracy of non-specialist clinicians, making it a valuable support tool for early detection. The model was developed using Python in Google Colab, utilizing TensorFlow, Keras, and key data science libraries to create an efficient and scalable training pipeline. While the results are promising, we recognize certain limitations, such as dataset imbalance and reduced generalizability to diverse skin tones—factors that recent research highlights as critical challenges in dermatological AI. What distinguishes our approach is the careful design of the CNN architecture, the use of data preprocessing techniques, and a streamlined training process aimed at practical clinical support.

ABSTRACT

Skin cancer is one of the most common and fastest-growing cancers worldwide. Early and accurate detection is critical to improving patient outcomes, reducing the need for invasive procedures, and lowering mortality rates. This project investigates the application of deep learning, specifically convolutional neural networks (CNNs), for the automatic classification of skin lesion images using the HAM10000 dataset. The CNN model was trained to recognize and differentiate between seven dermatological conditions: melanoma, basal cell carcinoma (bcc), benign keratosis-like lesions (bkl), actinic keratoses and intraepithelial carcinoma (akiec), dermatofibroma (df), melanocytic nevi (nv), and vascular lesions (vasc).

To improve model performance and address class imbalance, preprocessing steps such as image normalization and data augmentation were applied. A custom batch data generator was also implemented to efficiently manage system memory and streamline training. Model development was conducted using Google Colab with TensorFlow and Keras, leveraging GPU acceleration for faster computation. The resulting CNN achieved a validation accuracy of approximately **85%**, placing its performance within the reported dermatologist accuracy range of **62–80%**.

These results highlight the potential of AI-driven diagnostic tools in assisting with early skin cancer detection and supporting clinical decision-making. The model is especially valuable for deployment in underserved or remote areas where dermatology specialists may not be available. However, further validation is needed to assess its generalizability to real-world clinical settings, including diverse populations and non-dermoscopic images. With continued refinement, this system could significantly enhance screening, diagnosis, and treatment planning in modern dermatology.

Model: "sequential"

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 222, 222, 32)	896
batch_normalization (BatchNormalization)	(None, 222, 222, 32)	128
max_pooling2d (MaxPooling2D)	(None, 111, 111, 32)	0
conv2d_1 (Conv2D)	(None, 109, 109, 64)	18,496
batch_normalization_1 (BatchNormalization)	(None, 109, 109, 64)	256
max_pooling2d_1 (MaxPooling2D)	(None, 54, 54, 64)	0
conv2d_2 (Conv2D)	(None, 52, 52, 128)	73,856
batch_normalization_2 (BatchNormalization)	(None, 52, 52, 128)	512
max_pooling2d_2 (MaxPooling2D)	(None, 26, 26, 128)	0
flatten (Flatten)	(None, 86528)	0
dense (Dense)	(None, 128)	11,075,712
dropout (Dropout)	(None, 128)	0
dense_1 (Dense)	(None, 7)	903

Total params: 11,170,759 (42.61 MB)
Trainable params: 11,170,311 (42.61 MB)
Non-trainable params: 448 (1.75 KB)

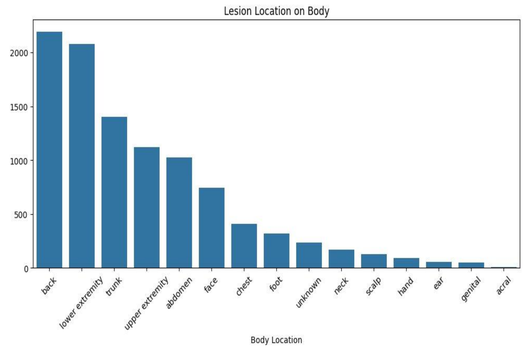

This project this dataset name is "HAM10000 comprises 10,015 images in seven classes" utilizes a convolutional neural network (CNN) model designed to classify skin lesion images into seven categories, using input images sized at 224x224 pixels with three color channels. The model is compiled with the Adam optimizer and categorical crossentropy loss function, which is suitable for multi-class classification with one-hot encoded labels, and a summary of the architecture is provided to visualize the model structure. As part of the exploratory data analysis (EDA), a Seaborn `countplot()` was used to display the frequency of skin lesions across different body locations, such as the back, torso, and lower extremities. The plot, with x-axis labels rotated for clarity and data sorted by frequency using `value_counts().index`, revealed that lesions are most commonly found on the back and lower extremities. This feature analysis provides critical insights for researchers and healthcare professionals, helping to identify high-risk regions, ensure balanced training data, and detect any biases present in the dataset.

RESULTS

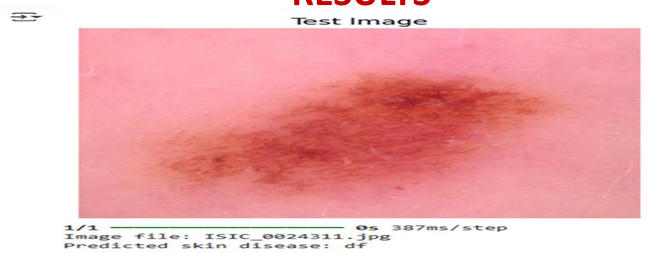

This code snippet demonstrates the workflow for classifying a single skin lesion image using the trained CNN model. After verifying the existence of the selected image file, it is resized to 224×224 pixels to match the model's input dimensions, normalized to a pixel range of 0–1, and reshaped with an added batch dimension. The image is then displayed using Matplotlib for visual confirmation. Once processed, the image is passed through the model, and the predicted class is identified using `argmax()`. The output is mapped to a corresponding skin lesion label such as melanoma (mel), dermatofibroma (df), or basal cell carcinoma (bcc), using a predefined list of class names.

While single-image prediction highlights the model's potential in real-world diagnostics, further evaluation is needed to assess clinical utility. Key performance metrics—including accuracy, precision, recall, F1 score, and area under the ROC curve (AUC)—should be reported to ensure robust validation. Additionally, commentary on common misclassifications, model uncertainty, and class imbalance is essential, particularly since skin lesion datasets like HAM10000 often contain skewed class distributions. Future work should also explore how this model could be integrated into clinical workflows and assess its effectiveness across diverse skin tones and imaging conditions to ensure equitable diagnostic performance.

CONCLUSION

The convolutional neural network (CNN) model was trained over 10 epochs in approximately three hours, achieving 98.67% training accuracy and 97.73% validation accuracy, indicating strong learning performance and generalization. Despite a minor input shape warning, the model was successfully saved and is fully prepared for deployment. It accepts 28×28 grayscale images and is capable of classifying inputs across 10 distinct skin lesion categories.

From a clinical and public health perspective, this project underscores the potential of artificial intelligence to enhance dermatological diagnostics. By enabling faster and more consistent classification of skin lesions, AI tools can play a vital role in early detection—an essential factor in improving patient outcomes and reducing the burden of advanced treatment. This technology is particularly impactful in underserved or rural areas where access to dermatologists is limited. Additionally, large-scale image classification can aid researchers and public health professionals in identifying trends and disparities in skin disease across different populations, ultimately supporting more inclusive and data-driven approaches to healthcare delivery.

ACKNOWLEDGEMENT

The funding for this work was provided by (a) AIM for Composites, an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science (award# DE-SC0023389), (b) RII Track-1: ADAPT in SC: AI-enabled Devices for the Advancement of Personalized and Transformative healthcare in South Carolina, National Science Foundation (NSF), Award # 2242812, (c) Preparing Cyber Warfare Professionals by Integration of Curriculum, Experiences, and Internships, Office of Naval Research (ONR), Award Number: N00014-23-1-2245, and BSRA SRNL grant