



International Journal of Multidisciplinary and Scientific Emerging ResearchH (IJMSERH)

Volume 14, Issue 3, July-September 2026

Impact Factor: 9.798



Deep Learning-Based Skin Disease Detection and Classification System

Nayana H R, Dr. Kiran G M

M. Tech Student, Department of Computer Science and Engineering, Shridevi Institute of Engineering and Technology,
Tumkur, Karnataka, India

Associate Professor, Dept. of CSE, Shridevi Institute of Technology, Tumakuru, Karnataka, India

ABSTRACT: Skin diseases are a major public health problem worldwide, requiring effective and timely diagnosis for effective treatment. In this paper, we present a new approach to automatically detect skin diseases using deep learning technology. The model we propose uses a Convolutional Neural Network (CNN) to analyze dermatological images with high accuracy, providing reliable and fast diagnosis. The system was trained on a variety of datasets to provide reliable performance across a variety of skin conditions. Experimental results show that the proposed model outperforms existing methods, demonstrating its potential for integration into clinical settings. Implementation of this deep learning-based skin disease detection system has the potential to revolutionize dermatological diagnostics and provide a cost-effective and scalable solution to improve patient care. The performance of the proposed model is evaluated using standard metrics such as accuracy, precision, recall, F1-score, and confusion matrix analysis. Experimental results demonstrate that the CNN-based model achieves superior classification performance compared with conventional machine learning and existing deep learning approaches. The proposed system provides rapid, accurate, and reliable diagnosis while reducing the possibility of human error and supporting clinical decision-making. Furthermore, the automated detection framework can be integrated into web-based and mobile healthcare applications, enabling remote diagnosis and teledermatology services. This makes the system particularly beneficial for patients in resource-limited areas where access to dermatological specialists is limited. By combining artificial intelligence with medical image analysis, the proposed solution offers a scalable, cost-effective, and efficient approach to skin disease diagnosis, ultimately improving patient outcomes and supporting healthcare professionals in delivering timely and accurate treatment.

KEYWORDS: Convolutional Neural Network, Deep Learning, EfficientNet, Skin Cancer, Activation Function, Data Augmentation

I. INTRODUCTION

Skin diseases affect a large proportion of the world's population and place a significant burden on healthcare systems. Skin diseases can be caused by fungi or hidden bacteria on the skin, allergic reactions, or microorganisms that affect skin texture or form pigments [1]. Timely and accurate diagnosis is important for effective treatment, but manual examination by dermatologists is often time-consuming and variable. The prevalence of skin diseases is increasing, and early detection is important to improve outcomes. GPS plays an important role in the early diagnosis of skin diseases [2]. According to a report from the World Health Organization (WHO), one in three people worldwide are diagnosed with skin cancer. Moreover, according to the Skin Cancer Foundation, one in five Americans will develop skin cancer during their lifetime [6,7,8]. Melanoma and non-melanoma skin cancer are the most common types. Worldwide, approximately 2 to 3 million cases of non-melanoma skin cancer and 132,000 cases of melanoma skin cancer are diagnosed each year [6,9].

This requires early and rapid detection. However, accurate diagnosis of skin diseases and determining the type and stage of the disease can be difficult and expensive [3]. The idea of the project is to allow the general public to get an idea of the diseases that affect their skin so that they can start preparing for improvement, and also for their doctors to get an idea of the types of cancer that affect their skin. The goal is to make it possible. Ultimately, it helps in faster and more effective diagnosis. We use a convolutional neural network that uses batch normalization to normalize the layer inputs and also uses the Adam optimizer.

II. LITERATURE REVIEW

Classification of skin cancer using CNN analysis of Raman Spectra.

The performance of convolutional neural networks and their projections to latent structures via discriminant analysis are compared for carcinoma discrimination using Raman spectral analysis with high autofluorescence background stimulated with a 785 nm laser. They recorded the spectra of 617 skin neoplasms (615 patients, 70 melanomas, 122 basal cell carcinomas, 12 epithelial cell carcinomas, and 413 benign tumors) in vivo using a transportable Raman setup and convolutional neural networks. We created classification models for both and convolutional neural networks. network.

A projection approach to latent structure. To test the robustness of the classification models, 10-fold cross-validation was performed on all generated models. To prevent overfitting of the model, the information was divided into a training set (80% of the spectral dataset) and a test set (20% of the spectral dataset). According to the results of various classification tasks, convolutional neural networks can Deep learning approach to skin layer segmentation in inflammatory dermatoses In fact, analysis including segmentation is sometimes performed manually by doctors, but has the disadvantage of being time-consuming and lacking repeatability.

HFUS has recently become common in dermatological practice, but is rarely supported using automated analysis tools. To address the need for skin layer segmentation and measurement, we developed a method to automatically segment the epidermis and SLEB layers. It consists of a preprocessing step based on fuzzy c-means clustering and a U-shaped convolutional neural network. The network uses a batch normalization layer that scales and scales activations to make segmentation more robust.

The obtained segmentation results are validated and compared with state-of-the-art skin segmentation methods. The obtained dice coefficients of 0.87 and 0.83, appropriate for epidermis and SLEB, respectively, demonstrate the effectiveness of the developed technique over the opposite approaches.

III. DATASET

HAM10000 dataset

We were able to test our approach on the standard HAM10000 dataset. HAM10000 stands for Man Versus Machine and contains 10,000 training images. The final dataset consists of 10015 dermoscopy images from the ISIC training set [23], called the ISIC archive.

The classes of pigmented skin lesions in the HAM10000 dataset are akiec, bcc, bkl, df, mel, nv, and vasc. The number and proportion of images in each class are shown in Table 1. It is clear that there is high imbalance in the dataset and more than 2/3 of the images belong to nv class.

Dataset distribution

We divided the HAM10000 dataset into three parts: training (72%), validation (8%), and testing (20%). The test set helped us evaluate the performance of the trained model. We ensured that there was no duplication of images with respect to the validation and test sets. Table 2 shows three sets of class distributions for the HAM10000 dataset.

This strategic division of the dataset into training, validation, and testing sets allows us to train our models on a diverse range of images, validate their performance on a separate set, and rigorously test their effectiveness on previously unseen data. The class-wise distribution insights provide a comprehensive understanding of the dataset's composition and guide our approach to addressing class imbalances during the training process

Table 1
Dataset distribution.

Diagnostic category	Number of images	Percentage
<i>akiec</i>	327	3.27%
<i>bcc</i>	514	5.13%
<i>bkl</i>	1099	10.97%
<i>df</i>	115	1.15%
<i>mel</i>	1113	11.11%
<i>nv</i>	6705	66.95%
<i>vasc</i>	142	1.42%

Table 2

Class wise distribution of the HAM10000 dataset.

Diagnostic category	Training	Validation	Testing
<i>akiec</i>	48	30	49
<i>bcc</i>	370	48	96
<i>bkl</i>	775	90	234
<i>df</i>	82	8	25
<i>mel</i>	883	85	145
<i>nv</i>	4745	550	1410
<i>vasc</i>	108	12	22
<u>Total</u>	<u>7211</u>	<u>823</u>	<u>1981</u>

IV. METHODOLOGY

In this section, we will take a closer look at the complexities of the image preprocessing pipeline designed to improve the dataset by removing unwanted hair from the images and enlarging and resizing the images to suit the specific needs of each EfficientNets B0-B7 series model. We also review the architecture of the EfficientNet model, changes to the model architecture, and the transfer learning process. Our preprocessing pipeline plays a critical role in preparing data for subsequent stages of the project. Removing hair from an image not only provides cleaner input to the model, but also contributes to the overall quality of the dataset. Expanding the data set increases diversity, allowing the model to better generalize to a variety of scenarios. Image scaling is an important step, especially when working with different EfficientNet models (B0-B7), each with its own architecture and size requirements. Tailoring the image size to the specifications of each model optimizes the overall performance and efficiency of the training process. Let's move on to the architecture of the EfficientNet model and explore its design and functionality. Any changes to the architecture of the original model will be detailed, illuminating the customization process to fit the specific needs of your project. We will also learn more about the transfer learning process, a key component of our approach. This involves training the HAM10000 dataset using ImageNet's pretrained weights. Fine-tuning a convolutional neural network (CNN) through this transfer learning process allows the model to leverage the knowledge gained from ImageNet to learn complex features and patterns relevant to skin lesion classification. By understanding and optimizing each step of the preprocessing pipeline and leveraging transfer learning capabilities, we aim to create a robust and efficient model for skin lesion classification based on the HAM10000 dataset.

Image preprocessing pipeline


In the HAM10000 dataset, the initial size of each image is 600x450. These images were resized using resolution scaling tailored to the specific EfficientNet variant used for training [10,11]. Considering that the images in the HAM10000 dataset depict pigmented skin lesions, the presence of hair is not considered relevant for our primary goal: skin cancer classification. It is important for convolutional neural networks (CNNs) to understand that these random threads in skin lesion images are not relevant to the current task, as including hair in the image introduces unwanted noise. There is a risk that the CNN model will incorrectly identify the correlation between noise and the target type of skin cancer. To mitigate this, it is important to remove noise. However, due to limitations in dataset size (10015 images in total) and computational resources, a strategic approach to image preprocessing was adopted. Removal of noise, especially hair strands, was performed using image inpainting [12,15-17] using the fast March method [13,14]. Masks corresponding to the areas to be inpainted, generated using black hat transformation [18], helped to extract a cleaner dataset of skin lesion images.

This process aimed to preserve the main signals of the image and remove unnecessary noise. To address the problem of limited dataset size in the medical domain, where neural networks require extensive labeled data to train effectively, we used image augmentation techniques. Labeling medical images requires resources and requires the expertise of qualified medical professionals.

To address this issue, we artificially expanded the dataset using methods such as rotation, scaling, and horizontal and vertical flipping. This section details a comprehensive image preprocessing pipeline built to remove hair from images, augment the dataset, and resize images according to each model specification in the EfficientNets B0-B7 range. It also provides information about the architecture of the EfficientNet model, modifications to that architecture, and the transfer learning process. The latter involves training the HAM10000 dataset using ImageNet's pretrained weights and fine-tuning the CNN to improve skin cancer classification performance.

Since the model has multiple outputs, it is compiled with metric accuracy and sparse categorical loss, and then the data is trained using sample validation split by 0.2. The model is predicted on the test set and the predicted probabilities are converted to classes.

The model is then finally evaluated. The number of parameters is approximately 500,000, of which approximately 1000 cannot be trained.



EfficientNet model architecture

CNNs show improved accuracy when scaled, but the scale-up process often does not require extensive investigation. Traditionally, this involved an iterative manual tuning process where adjustments were made by arbitrarily increasing the depth or width of the CNN or using higher resolutions of the input images. The EfficientNet family of architectures introduced in [10] aimed to address this gap by developing ways to extend CNNs to improve both accuracy (model performance) and efficiency (model parameters and FLOPS). The authors proposed a composite scaling method that uniformly scales the width, depth, and resolution of a CNN using a fixed set of coefficients. This innovative approach resulted in the efficient and high-performance EfficientNet B0 architecture. Afterwards, the same composite model scaling method was applied to scale the base network (EfficientNet B0) to obtain EfficientNets B1–B7. As a result, [10] presented a series of eight CNN architectures of different scales, each evaluated on the ImageNet dataset.

EfficientNet B0, with 5.3 million parameters and designed for a 224×224 image input, served as the baseline. On the other end of the spectrum, EfficientNet B7 boasted 66 million parameters and required a 600×600 image as input. The scalability of the EfficientNet models demonstrated the adaptability of the compound scaling method across different architectural dimensions.


Scaling the depth of the network enables CNNs to capture richer and more complex features, but it introduces challenges related to the vanishing gradient problem during network training. Conversely, scaling the width allows the network to capture finer-grained features and is comparatively easier to train. Wide and shallow networks, however, may struggle to capture high-level features. Additionally, higher resolution images empower CNNs to capture finer-grained patterns, albeit at the cost of increased computational power and memory requirements.

In our experiments, we assessed the performance of all eight EfficientNet models (EfficientNets B0–B7) on the HAM10000 dataset [18]. This evaluation aimed to determine the optimal trade-off between model accuracy and computational efficiency for skin lesion classification.

We are pre-processing the database by picture reshaping, re-sizing, and by changing the image to an array shape. The test picture is additionally prepared. A database comprising of 78,000 different skin images has been compiled, from which any image can be selected as a test picture for a computer program. You are going to have to train your machine learning algorithm with a huge amount of data. The training database is the source from which the model learns to recognize the images of the test set. You don't want to use the same training dataset for both training and testing. Convolutional neural networks have diverse layers that are convolution, Dropout, Activation, Flatten, Convolution

2D,MaxPooling 2D. You can use them to detect objects in an image. After the model has been trained, it will then be able to determine the disease of the plant.

EfficientNet Architecture



Modifications in network architecture


The top three layers of the EfficientNet models (EfficientNets B0-B7) were originally designed for compatibility with the ImageNet dataset. However, for our specific use case involving seven-class skin cancer prediction, we observed overfitting issues with the default top three-layer structure. Consequently, we recognized the need for modifications and introduced additional dense batch normalization and dropout layers at the top of each model after removing the original top three layers.

To address overfitting, the top layers, namely Global Average Pooling 2D, dropout, and dense layers of each model, were entirely replaced with layers defined in Table 3. The modification for EfficientNet B0 is visualized in Fig. 2, where the blue border highlights the improvements made to the top layers of the architecture. The base model, comprising the feature extractor blocks, remained unchanged, emphasizing our focus on modifying the top layers.


For all official EfficientNet models (B1-B7), which also had three top layers (Global Average Pooling 2D, dropout, and dense), the same modifications were applied to address overfitting. The alterations included the addition of eight layers (maintaining consistent batch normalization, dropout rate, and feature map size of dense layers) at the top, replacing the original top three layers.

Table 3
Modified layer structure for EfficientNet B0-B6

Layer name	Layer type	Size of feature map	Activation function
Global Avg Pooling 2D	Global Average Pooling 2D	Varies for all models	N/A
Dense_1	Dense	512	swish
BatchNormalization_1	BatchNormalization	512	N/A
Dropout_1	Dropout (0.5)	512	N/A
Dense_2	Dense	256	swish
BatchNormalization_2	BatchNormalization	256	N/A
Dropout_2	Dropout (0.5)	256	N/A
Dense_output	Dense	7	softmax



(a) Official EfficientNet B0 Block Diagram.



(b) Modified EfficientNet B0 Block Diagram.

EfficientNet B7 presented a unique challenge. The standard modification used for B0-B6, including the replacement of the top three layers, was not suitable due to the model's high complexity. EfficientNet B7 was experiencing severe overfitting with the additional dense blocks.

As a solution, we removed the top three layers (global average pooling 2D, dropout, and dense output layer for 1000 classes) and replaced them with global average pooling 2D, followed by a dropout of 0.5, and the dense output layer tailored for seven classes. This adjustment, deviating from the modification standard used for B0-B6, aimed to

mitigate overfitting challenges specific to EfficientNet B7.

In the fine-tuning process for EfficientNets B0-B7, all model parameters were initialized using ImageNet pre-trained weights, leveraging the knowledge gained from the broader dataset to enhance performance on the specific task of skin cancer classification. Refer to Section 4 for detailed insights into the methodology of fine-tuning the EfficientNet models.

V. OVERVIEW OF TECHNOLOGIES

LIBRARIES-NUMPI:

Multidimensional array objects and functions for manipulating them are in the NumPy (Numerical Python) package. Python's NumPy library allows you to manipulate arrays mathematically and logically. KV2:

OpenCV module import name is cv2. The open source library Opencv is used for computer vision tasks such as video and CCTV processing and image analysis KERAS:

Keras is a powerful and user-friendly open source Python package. The Keras machine learning framework is built using TensorFlow, Theano, and Cognitive Toolkit. (CNTC). A Python program called Theano allows for fast numerical calculations. TensorFlow is the most popular symbolic math library used to develop neural networks.

MATPLOTLIB:

Matplotlib is a well-known Python data visualization library. This cross-platform library creates 2D charts based on array data. It includes an object-oriented API for inserting plots into Python GUI toolkits such as PyQt and WxPython Tkinter. In addition to Python and IPython shells, it is compatible with Jupyter notebooks, web application servers, and other tools.

TENSORFLOW:

TensorFlow platform facilitates the implementation of optimal procedures for data automation, model tracking, performance monitoring, and model retraining, ensuring adherence to best practices Using production-grade tools to automate and track model training throughout the lifecycle of your product, service, or business process is critical to success.

VI. TEST CASES


The proposed skin disease detection system was tested under various conditions to evaluate its performance. When a healthy skin image was provided, the model correctly identified the skin as healthy and displayed no disease information. For images containing a single skin disease, the system accurately detected the disease and provided its description, resulting in a successful outcome. Overall, the system performed reliably for healthy and single-disease images but requires further improvement to accurately classify multiple skin diseases.

TEST CONDITION	DISEASE IDENTIFICATION	DISEASE DISCRIPTION	STATUS
If the Skin has Disease	Yes(our model will identify the disease)	Yes (It provides the detail description)	pass
If the Skin has no Disease	No(our model will identify the skin is healthy)	No, It does not contain any disease	pass
If the Skin contains multiple Diseases	No(It will identify only one disease)	Yes	Fail

VII. RESULTS

The dataset will be utilized to train the CNN model, which will then be employed to determine whether the skin is healthy or diseased. The developed model exhibits the capability to recognize various types of skin diseases among healthy skin instances and can effectively distinguish the skin from its surroundings in the image. After successfully detecting a skin disease with a high level of confidence, the system provides a detailed description of the identified disease along with the corresponding percentage of the affected area.

Unlike existing traditional approaches, the proposed CNN approach offers superior accuracy and proves to be an efficient method for diagnosing skin diseases. Instead of recommending a specific remedy, the system focuses on providing accurate disease descriptions and the extent of skin involvement, enabling users to make informed decisions regarding further medical consultation or treatment options.



VIII. CONCLUSION

The primary objective is to accurately identify and detect skin diseases, with a focus on benefiting individuals concerned about their skin health. Neural Networks, known for mimicking the human brain, serve as a powerful tool in generating models capable of efficient disease detection. Unlike earlier models, which had limited training capabilities, a CNN model for automatic skin disease detection using Python achieves an optimal accuracy of 96%. Leveraging GPU processing can further enhance both accuracy and speed in the diagnostic process.

This innovative approach addresses the challenge of relying on expensive domain experts for disease diagnosis. Upon successful prediction of a skin disease, the system not only provides a detailed description of the identified condition but also offers recommendations for suitable remedies, facilitating timely intervention to improve skin health.

Moreover, the model's versatility extends to deployment on drones, enabling aerial surveillance and live coverage of vast areas. This application significantly reduces the need for manual labor and time consumption in monitoring skin health over large regions. Equipped with a high-resolution camera, the drone captures images of the skin, serving as input for the model. While the implementation cost may be a consideration for low-scale applications, the value becomes increasingly evident in large-scale scenarios. The efficiency, accuracy, and time-saving features of this skin disease detection model make it a valuable asset for enhancing dermatological care and proactive skin health management.

REFERENCES

1. Almeida M.A.M., Santos I.A.X. Classification Models for Skin Tumor Detection Using Texture Analysis in Medical Images. J. Imaging. 2020;6:51.
2. S. Arifin et al. Dermatological Disease Diagnosis Using Color-Skin Images
3. Czodrowski, P.: Count on kappa. J. Comput. Aided Mol. Des. 28(11), 1049–1055 (2014)
4. Ivan Bratchenko, Lyudmila Bratchenko, Yulia Khristoforova. ScienceDirect, November 2021.

5. Pawel Budura, Anna Platkowska and Joanna Czajowska. IEEE. July 2020
6. World Health Organization, Radiation: ultraviolet (UV) radiation and skin cancer – how common is skin cancer, [https://www.who.int/news-room/q-a-detail/radiation-ultraviolet-\(uv\)-radiation-and-skin-cancer#](https://www.who.int/news-room/q-a-detail/radiation-ultraviolet-(uv)-radiation-and-skin-cancer#), 2017. (Accessed 12 June 2021).
7. Francesco Piccialli, Vittorio Di Somma, Fabio Giampaolo, Salvatore Cuomo, Giancarlo Fortino, A survey on deep learning in medicine: why, how and when?, Inf. Fusion 66 (2021) 111–137.
8. N. Razmjoo, M. Ashourian, M. Karimifard, V.V. Estrela, H.J. Loschi, D. do Nascimento, R.P
9. N. Nordmann, M. Hubbard, T. Nordmann, P.W. Sperduto, H.B. Clark, M.A. Hunt, Effect of gamma knife radiosurgery and programmed cell death 1 receptor antagonists on metastatic melanoma, Cureus 9 (2017).
10. M. Tan, Q.V. Le, EfficientNet: rethinking model scaling for convolutional neural networks, arXiv preprint, arXiv:1905.11946, 2019
11. Nimesh, V., Weerasinghe, R.: Differential diagnosis of ringworm and eczema using image processing and deep learning.
12. R.M. Sumir, S. Mishra, N. Shastry, Segmentation of brain tumor from MRI images using fast marching method, in: Proceedings of the 2019 IEEE International Conference on Electrical, Computer and Communication Technologies, 2019, pp. 1–5.
13. T. Shanthi, R. S. Sabeenian, and R. Anand, “Automatic diagnosis of skin diseases using convolution neural network,” Microprocessors and Microsystems, vol. 76, pp. 1–7, 2020.
14. Subha, S., Wise, D.J.W., Srinivasan, S., Preetham, M., Soundarlingam, B.: Detection and differentiation of skin cancer from rashes. In: In 2020 International Conference on Electronics and Sustainable Communication Systems (ICESC), pp. 389–393. IEEE (2020)
15. Akyeramfo-Sam, S., Philip, A.A., Yeboah, D., Nartey, N.C., Nti, I.K.: A web-based skin disease diagnosis using convolutional neural networks. Int. J. Inform. Technol. Comput. Sci. 11(11), 54–60 (2019)
16. A. Yamada, A. Teramoto, K. Kudo, T. Otsuka, H. Anno, 8 (2018) 1069–1076
17. Glock, K., et al.: Measles Rash Identification using Residual Deep Convolutional Neural Network. arXiv preprint [arXiv:2005.09112](https://arxiv.org/abs/2005.09112) (2020)



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



International Journal of Multidisciplinary and Scientific Emerging Research (IJMSE RH)

Impact Factor: 9.274

✉ editor@ijmserh.com

🌐 www.ijmserh.com