Archives available at journals.mriindia.com

Open Access International Journal of Science and Engineering

ISSN: 2456-3293

Volume 9 Issue 05, 2026

Implementation of an AI-Driven Skin Disease Detection System Using MobileNet and Flask Framework

[1]Nisha Vijay Bhalerao, [2]Anupama Shankarrao Budhewar

[1,2] Department of Computer Science Engineering, JSPM University Pune, School of Computational Science, Wagholi, Pune, India – 412207

[1]bhaleraonisha008@gmail.com, [2]budhewar.anupama@gmail.com

Peer Review Information

Submission: 30 April 2026
Revision: 09 May 2026
Acceptance: 26 May 2026

Keywords

Skin Disease Detection, Deep Learning, MobileNet, Flask Framework, Convolutional Neural Network, Medical Image Classification, Artificial Intelligence, Dermatoscopic Images, Healthcare Automation, Image Processing.

Abstract

One of the most prevalent health problems worldwide is skin disease, and it's important to identify it early for better treatment. However, traditional diagnosis techniques are still largely reliant upon medical dermatologists and extensive manual examination, which isn't always possible for individuals in remote areas or for locations that do not have access to a number of specialists. In this paper, we show how an AI driven skin disease detection system can be built, and yes using deep learning techniques too. The proposed method utilizes mainly the convolutional neural network, MobileNet, which are deployed to facilitate efficient and also relatively accurate classification of different skin diseases from dermoscopic images. The image is preprocessed before the model is run to get improved prediction results and increase the accuracy of the model. Not only that, the web app created is a Flask based web_app and its user-friendly, where a person can upload the skin image and immediately get the results with confidence scores. As a whole, the objective is to minimize human labor, speed up initial screening and facilitate the access to health care through an online platform capable of carrying out automated analysis of skin diseases.

Introduction

Skin diseases, kind of among the most common medical issues known to exist in practically all ages of the world, both young and old. In healthcare studies it is said that millions of people suffer from different types of skin disease such as melanoma, basal cell carcinoma, benign keratosis, vascular lesions and dermatofibroma every year. Some of these conditions are minor, short lived, or heal on their own, but others may end up serious and even life threatening if they are not caught in the early stage. In all of that, skin cancer is generally considered to be one of the most serious types as it can develop rapidly and the risk of fatality increases significantly if it is diagnosed late. Therefore, early diagnosis and a precise diagnosis of skin diseases are crucial for the better treatment of the disease and to decrease the risk of health in general[1].

However, the traditional method of diagnosis heavily depends on the expertise of dermatologists, dermatoscopic examination, laboratory examination, and manual analysis of the lesions in the skin. Unfortunately, these are often time consuming, expensive and sometimes unavailable in rural and far-flung areas, where access to expert medical resources is limited. Consequently, many patients may experience delays and the entire procedure can be more difficult than it is supposed to be [2][4].

Artificial Intelligence (AI) has in recent years radically transformed the healthcare industry, especially with the introduction of AI-driven diagnostic systems that assist healthcare practitioners in the prediction and analysis of diseases. In the various methods of AI, Deep Learning is one of the more robust techniques, particularly for medical image classification. Deep Learning enables models to learn meaningful visual patterns directly from images without requiring someone to first extract meaningful features from the image. CNNs, or convolutional neural networks, are a subset of deep learning and have performed exceptionally well in image recognition, object detection and even medical diagnosis problems. The CNN models work well in the task of reading complex structures in medical images, which is well suited to skin disease detection systems. Also, using deep learning in dermatology gets a lot of attention, because it allows more automated examination of skin lesion images, with high accuracy, and it also reduces the amount of human effort required[3].

Certainly, with the increasing number of dermatoscopic images available, research on automated skin disease diagnosis has accelerated. The HAM10000 dataset is one of the most popular datasets used in dermatology research, which contains thousands of annotated dermatoscopic images from various skin disease categories. In practice this dataset provides a reliable foundation for training and evaluating deep learning models, but there are some problems achieving very high accuracy in the detection of skin diseases[4]. There is significant variability in image quality, as well as in the morphology of the lesions, lighting conditions, and dataset imbalances can all negatively impact prediction performance. Moreover, certain skin diseases are visually similar, making the task of human and automated classification difficult. Nevertheless , because of these constraints researchers usually end up processing the images (for example resize, normalize, augment, and remove noise) before training the deep learning models [5]. In medical imaging, this approach has shown itself as very effective, and it performs quite well. The core notion is that you re-use deep learning models that were trained earlier , so they have already learned salient image features from very large datasets [7]. Rather than starting from zero and learning everything while training a deep neural network from scratch, you instead fine-tune that same network for a specific medical image classification problem, which often means less training time and lower compute needs.

Among the various transfer learning choices, MobileNet has really caught on , mostly because it stays lightweight and keeps the computational burden pretty low, plus it remains efficient overall. The idea behind developing MobileNet was to imagine scenarios where you have limited computational resources, like on a phone or an embedded device where vision is basically central. It is built on depthwise separable convolutions, which reduce the number of parameters without really messing up classification accuracy too much [6].

Because of all this, MobileNet seems like a fitting option for the real-time healthcare system , and also for web-based healthcare apps.

Besides the theory part, a couple of recent studies have shown the impact of MobileNet on skin lesion classification. With dermatoscopic images, some researchers have managed to tell skin conditions apart using MobileNet, plus its close variants (MobileNetV2), reaching relatively high accuracy. The models look like they can learn several disease types in parallel without becoming sluggish during inference, and even while using less memory. And when you apply transfer learning with MobileNet, it sort of helps to curb overfitting , and it improves generalization, especially when the training dataset is not that big. That's why MobileNet ended up as one of the more popular deep learning structures used to craft smart healthcare diagnostic systems [8].

However, alongside deep learning technologies, web app frameworks have played about as big of a role, in making sure AI gets into healthcare systems smoothly, like it's not this whole ordeal. Flask , which is a lightweight Python web framework is often picked when deploying machine learning, and deep learning models into real world web apps. It gives a pretty straightforward but also flexible environment where backend AI models can be stitched together with frontend user interfaces. With a Flask app, users can upload medical images, then chat with the prediction engine, and receive diagnostic outcomes right inside a web browser. In other words it does make things a bit more approachable, and it also helps deliver healthcare to end-users even in remote places where you don't really need special software sitting on the device. Also, healthcare systems that are hosted on the web usually make telemedicine workflows and remote diagnosis happen more easily, and lately this is getting more and more relevant in the healthcare landscape [9].

Literature Review

There have been some developments in AI, and DL, which in turn made the accuracy of automated skin disease detection systems climb up to some extent, not fully but kinda still. Basically, a bunch of researchers ended up really leaning on Convolutional Neural Networks, or

CNNs, in medical image work, mostly because they can pull out useful visual traits on their own, specially when the images are dermatoscopic. People also keep pairing these with well-known pretrained backbones like MobileNet ResNet, DenseNet, and EfficientNet, and they've delivered solid classification outcomes too, while also running with lower computational load than you might expect. Some reports even argue that MobileNet-based designs are kinda better suited for limited-resource situations, like in healthcare deployments, because they run fast , and they don't need a huge memory budget, at all.

Also, HAM10000 is one of the most cited benchmark datasets when it comes to training and testing skin lesion classification models. And on top of that, quite a few groups pushed further, using data augmentation, image preprocessing, and normalization strategies, to help their models generalize better and to reduce overfitting tendencies in these skin disease pipelines [10].

Not too long ago, there's also a little bit of literature that kinda stresses the value of bringing AI diagnostic systems onto web and mobile cell platforms, so, in that way, improving healthcare access. In real practice, people often end up using Flask and Django frameworks, to help link deep learning models into more human friendly healthcare apps, even if the connection is not always clean or perfect. At the same time, a lot of work is rolling forward on cloud based and near real time skin disease prediction systems that can support dermatologists with earlier detection. Also, ensemble learning approaches have been studied to squeeze out higher prediction accuracy ,and to deal with dataset imbalance in a more conscientious manner, plus hybrid CNN architectures have been tried for basically the same purpose—aiming for better accuracy and a more careful handling of dataset imbalance. Some of these studies even added explainable AI methods mainly so the reasoning behind each decision becomes easier to grasp and, as a result, to lift user trust in automated diagnosis systems too. Still, despite the decent results from today's setups, a handful of problems are sitting there, like small datasets and variations in images, computational costs and deployment constraints in the real world, etc., and these concerns keep needing extra research, tuning ,and optimization so they can actually become practical for healthcare applications [11][20].

Table 1: Comparative Analysis of Existing Skin Disease Detection Techniques Using Deep Learning

Ref. No.	Author(s) & Year	Title	Technique Used	Dataset	Key Findings
[6]	A. Dascalu and E. David (2020)	Skin Cancer Detection by Deep Learning and Sound Analysis Algorithms	Deep Learning CNN	Dermoscopic Images	Improved early skin cancer detection accuracy using AI-based image analysis.
[7]	M. A. Al-masni et al. (2020)	Skin Lesion Segmentation in Dermoscopy Images	Full Resolution CNN	ISIC Dataset	Achieved efficient lesion segmentation for better classification performance.
[8]	H. Kaur et al. (2020)	Imbalanced Data Challenges in Machine Learning	Machine Learning Techniques	Multiple Medical Datasets	Highlighted imbalance handling methods for improving prediction accuracy.
[9]	S. Hameed et al. (2021)	Transfer Learning Approach for Skin Lesion Classification	MobileNet Transfer Learning	HAM10000	MobileNet provided faster computation with high classification accuracy.

[10]	P. Tschandl et al. (2020)	Human–Computer Collaboration for Skin Cancer Recognition	CNN with Dermatologist Support	HAM10000	Combined AI and human expertise improved diagnosis performance.
[11]	R. Gessert et al. (2021)	Skin Lesion Classification Using CNNs	Patch-Based Attention CNN	ISIC Archive	Attention mechanisms enhanced lesion feature extraction.
[12]	M. Khan et al. (2021)	Lightweight Deep Learning Model for Skin Lesion Classification	Lightweight CNN	HAM10000	Reduced computational cost while maintaining reliable accuracy.

Table 1 provides an overview of a comparative study of more recent research on skin disease detection, primarily based on deep learning methods. In that table, it sort of encapsulates the different approaches the datasets that were used, and some of the key findings that resulted from previous studies. It also highlights the effectiveness of CNN based models, transfer learning techniques, and lighter models such as MobileNet. Generally, these techniques contribute to higher classification accuracy, and also enhance the performance in healthcare diagnosis, which is the key thing here.

Methodology

1. MobileNet Transfer Learning Method

It is kinda neat to propose a system to classify skin disease using one of the most efficient methods MobileNet transfer learning approach. MobileNet is a lightweight design of Convolutional Neural Network (CNN) designed for image recognition tasks, that requires minimal computational effort and lacks complexity. Here a pretrained MobileNet model is taken and then further trained on the HAM10000 skin lesion dataset to adapt to this problem. Generally speaking, transfer learning can significantly shorten training time, and it can improve classification accuracy because the model has acquired some useful visual features from large-scale datasets for learning [13].

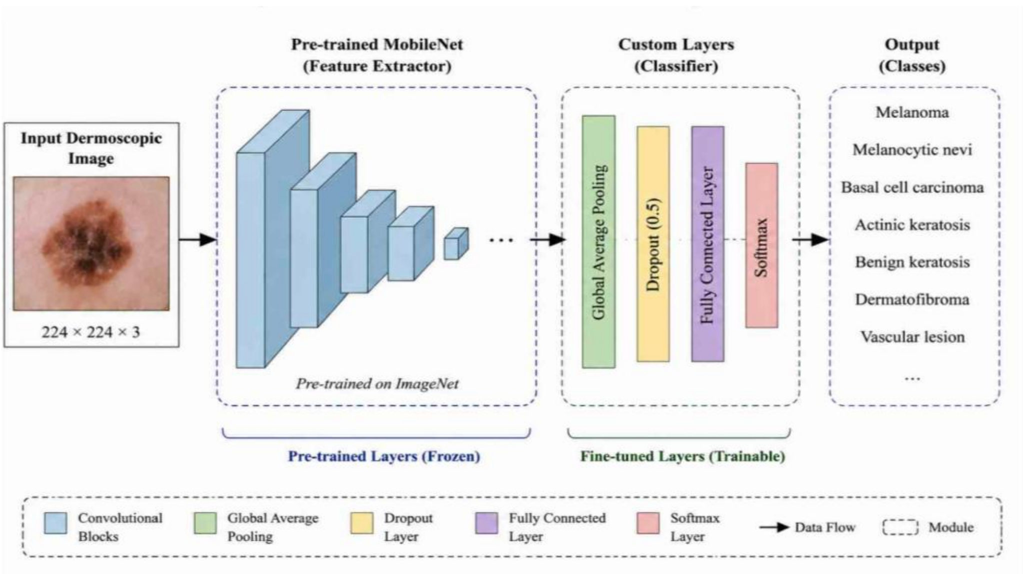

Figure 1: MobileNet Transfer Learning Method

I think the transfer learning setup with MobileNet for skin disease classification is seen in figure 1. The dermatoscopic image is fed into the convolutional part, which is already trained, and then the already trained MobileNet convolutional part takes over, with a focus on feature extraction. Then there are classification pieces that are custom to the entire network such as global average pooling, a dropout stage, then a fully connected layer and finally the softmax output. Finally, it outputs probabilities for multiple categories of skin disease and, at the same time, has a satisfactory efficiency as well.

Next, the MobileNet network receives dermatoscopic skin images for classification across a number of different skin diseases, such as melanoma, basal cell carcinoma, and benign keratosis. As the architecture is lightweight and the processing is quick, MobileNet is a perfect fit for real-time healthcare setups. It also integrates easily with web-based disease forecasting systems, particularly where latency is an issue[14].

2. Image Preprocessing and Data Augmentation Method

Before we do classification, the skin image that we have uploaded is in some ways preprocessed to aid the quality and consistency of the images. Here in this project, every dermatoscopic picture is resized to 224×224 pixels, to fit the input size that the MobileNet model needs. Then we do pixel normalization, so to say, to "normalize" the intensity values of the images, so that the model can perform better overall. Preprocessing also reduces noise and some other variations that occur due to lighting conditions, image resolution and even slight differences in skin texture. In addition, the images are augmented with data, such as rotating, flipping, zooming and moving images. This is done to increase the diversity of this dataset and reduce problems with overfitting during trainings. Generally speaking, the preprocessing and some augmentation tweaks help a deep learning model reach better generalization, and also improve prediction accuracy [15] , at least that’s how it usually feels in practice.

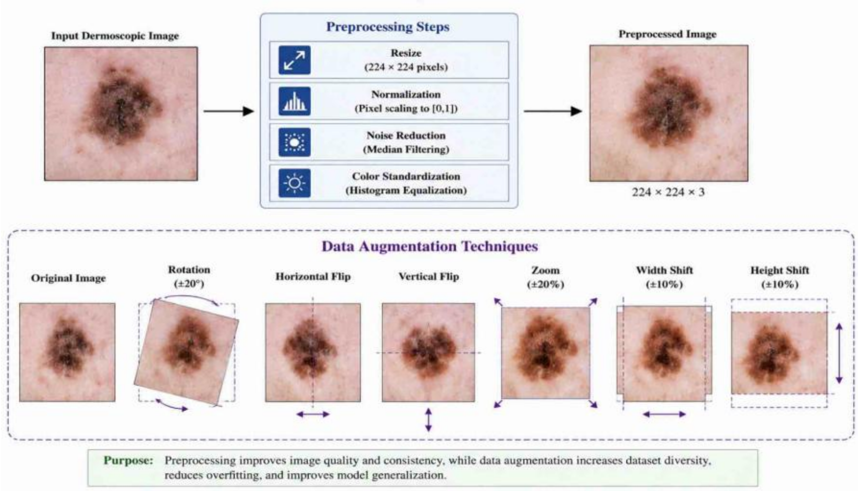

Figure 2: Image Preprocessing and Data Augmentation Method

Figure 2 shows the image preprocessing together with a couple of data augmentation tricks that were used for the proposed skin disease detection system. First the dermatoscopic input image goes through a few preprocessing stages, like resizing, normalization, noise reduction and color standardization, so overall the view looks clearer and a bit more consistent. Then the data augmentation part kicks in, using things such as rotation, horizontal flip, vertical flip, zooming, width shifting , and height shifting. The idea is basically to expand the dataset variety, reduce overfitting risk, and also strengthen the model generalization ability while it is training.

3. Flask-Based Web Application Deployment Method

The Flask framework is used to build and deploy the web based system for skin disease detection. Flask, is a light-weight Python web framework that makes it easier to connect the trained deep learning model with the user interface, sort of smoothly, you know. In the proposed system, users can upload skin images through a simple web interface, not really anything fancy. The image that gets uploaded is saved only temporarily on the server then it gets sent to the preprocessing phase before the actual classification process starts [16]. Once the prediction step is finished, the system displays the detected skin disease name along with confidence accuracy, right on the webpage. Flask also brings flexibility, faster deployment, and simpler linking with HTML , CSS and Python based machine learning libraries. Overall this type of setup improves accessibility so users can access the diagnostic system remotely through an online platform, from basically anywhere.

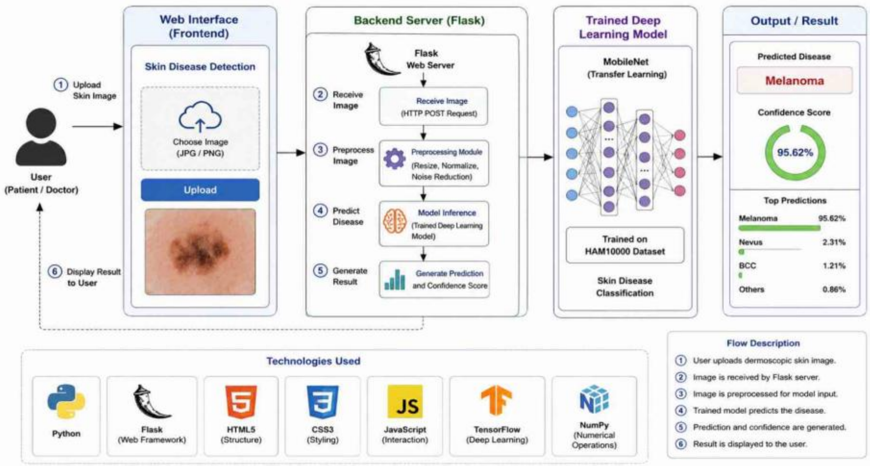

Figure 3: Flask-Based Web Application Deployment Method

Figure 3, sort of illustrates how the Flask based web application deployment runs for the proposed skin disease detection system, but not in a totally clean way. The process really begins when the user uploads a dermatoscopic image using the web interface [17]. Then the Flask server takes that image, runs a bit of preprocessing, and after that it passes the image over to the trained MobileNet deep learning model, for the actual prediction steps. When the model finishes, the system produces the disease classification output together with a confidence, or accuracy style value, and after that it shows everything right on the web page for the user. So, with this kind of arrangement, the online diagnostic workflow becomes faster, and also easier for people to reach [18].

4. Convolutional Neural Network (CNN) Classification Method

So basically the Convolutional Neural Network, CNN classification method is kinda one of the main deep learning tools we use for pulling image features , and then doing the disease classification. CNN models often automatically spot useful visual patterns from dermatoscopic images , like the color texture, lesion boundaries and also the general shape and all that , without us hand picking each and every cue [19] .

In this project MobileNet is the one that internally relies on CNN layers such as convolution , pooling and fully connected layers, for both feature extraction and the classification part. After the image gets processed it goes into the CNN model, then the very last output layer gives out the probability for each disease category. For the final decision, we take the class with the biggest probability , Softmax kind of helps with that probability mapping , and then NumPy argmax is used to grab that top category. Overall, CNN based classification tends to boost the prediction accuracy more than older, traditional machine learning approaches because it reduces the need for manual feature engineering. You know, instead of us designing features ourselves, the network kinda figures things out along the way.

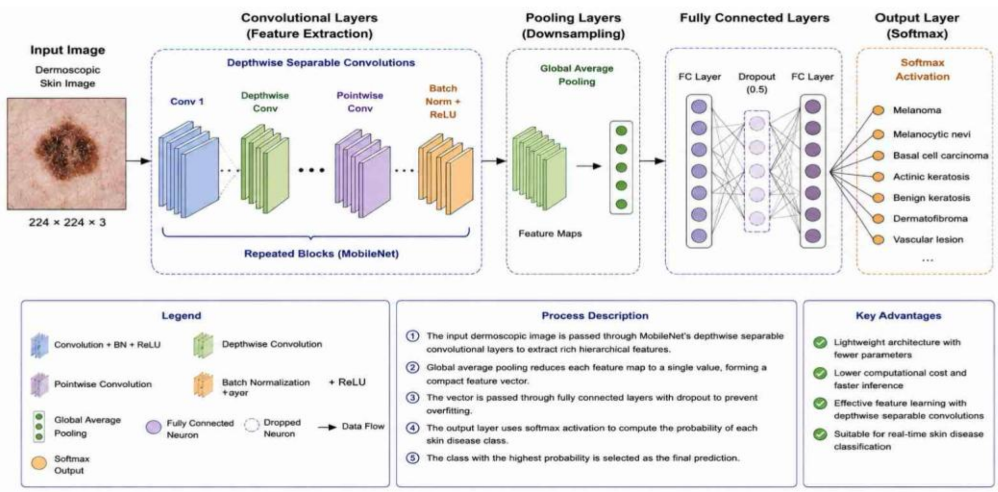

Figure 4: Convolutional Neural Network (CNN) Classification Method

Figure 4 is showing the classification setup of a Convolutional Neural Network, CNN, based on the MobileNet approach kind of used here for automated skin disease detection. At first, the dermatoscopic image goes in and it gets pushed through several convolution stages where depthwise plus pointwise convolutions work together for feature extraction in a fairly efficient way [21] . After that , pooling layers show up, they reduce the dimensionality, but they still try to preserve the key structures from the image, more or less. Then there are fully connected layers together with dropout , which supports the classification behavior and also lowers the chance of overfitting. In the end , the Softmax output layer spits out probabilities for the different skin disease categories, so the model can classify the medical images in a way that is accurate as well as efficient [22].

Results and Discussion

The proposed AI based skin disease detection system, was actually implemented with a MobileNet deep learning model, along with a Flask web framework which kinda works smoothly. It got trained and tested using the HAM10000 dermatoscopic image dataset, and that dataset includes several lesion categories too. From the experimental outcomes it looks like the approach reached solid classification accuracy, along and with fast prediction time across the different skin disease classes. Also the built web application provided this sorta straightforward interface for people to upload a picture, then receive the disease prediction right away.

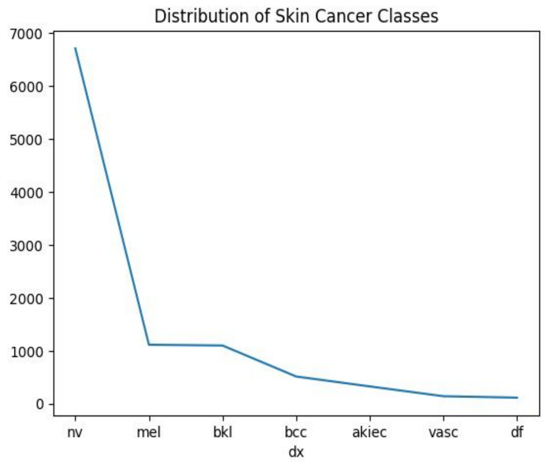

Figure 5: Distribution of Skin Cancer Classes

From the dataset distribution analysis in Figure 5, you can kind of see that HAM10000 has seven main skin disease classes, namely melanocytic nevi (nv), melanoma (mel), benign keratosis (bkl), basal cell carcinoma (bcc), actinic keratoses (akiec), vascular lesions (vasc), and dermatofibroma (df). In these groups, melanocytic nevi has the biggest number of images , while dermatofibroma and vascular lesion classes have relatively fewer samples. That kind of imbalance across the dataset makes the classification a little harder, so preprocessing along with data augmentation techniques are pretty much needed, to help the model generalize better .

Then the confusion matrix in Figure 6, kind of suggests how well the proposed MobileNet model performs for the actual classification task. Most of the prediction results look piled up along the diagonal region, which is often a sign that the model is doing the right thing for most of the test samples. For example, the model also appears to get a pretty solid accuracy on melanoma, basal cell carcinoma, and melanocytic nevi. A few smaller wrong guesses show up between lesion classes that look visually close , like their texture and color distribution patterns can be similar as well, so the network ends up confused a bit , which is not entirely unexpected. Overall, this confusion matrix basically supports that the deep learning approach works effectively for multiclass skin disease classification.

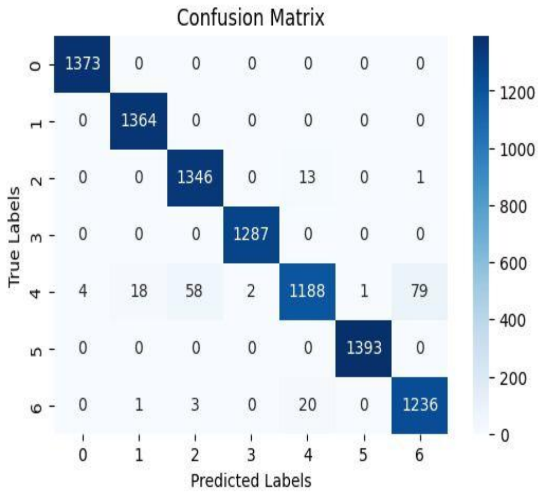

Figure 6: Confusion Matrix of Skin Disease Classification

Figure 7 kind of shows, the frequency distribution for the different skin cancer types that were used during training the model. You can see how the graph confirms that melanocytic nevi take up most of the dataset, more than the other lesion categories . Also, the preprocessing and augmentation strategies seemed to lessen the negative effects of the class imbalance, and in return the classification results felt better for the minority groups, even if it's not a perfect cure for everything .

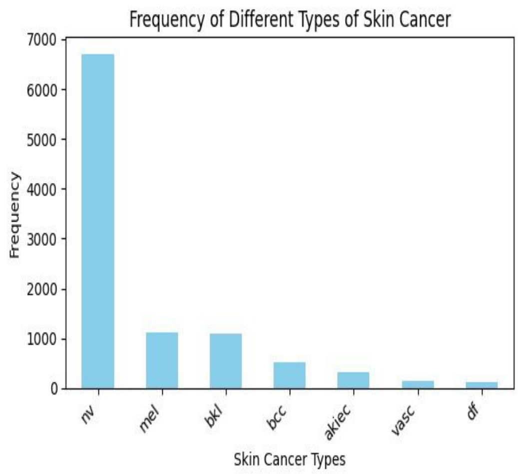

Figure 7: Frequency of Different Types of Skin Cancer

The training , and validation loss curves shown in Figure 8 really show how the proposed MobileNet model managed to converge through training. At the start both loss values were kinda high, but over time they kept dropping as the epoch count went up. You can see it a bit steadily, in a way. The thing is training curve and validation curve sit so close together, that it kinda suggests the model learned in a smooth manner not in some overly confident way, so there is no big overfitting. By the end, the final loss numbers got down to really low levels , which means the optimization stayed stable too, and the representation learning worked well as well, more or less.

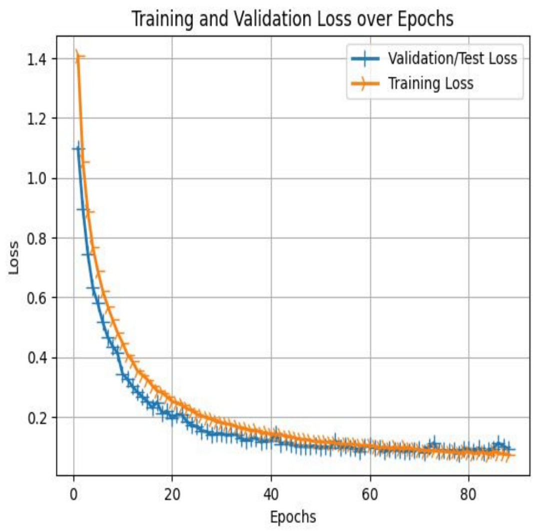

Figure 8: Training and validation Loss over Epochs

Figure 9 shows the training and validation accuracy curve s during the model training, basically they keep going up pretty steadily while the epochs roll by. At some point it sort of levelled off near 97% validation accuracy, and it feels pretty consistent. There is also only a small gap between the training score and the validation score, so that suggests better generalization and a more steady behavior on unseen test data. All in all, these results confirm that MobileNet transfer learning is a highly suitable approach for skin disease classification.

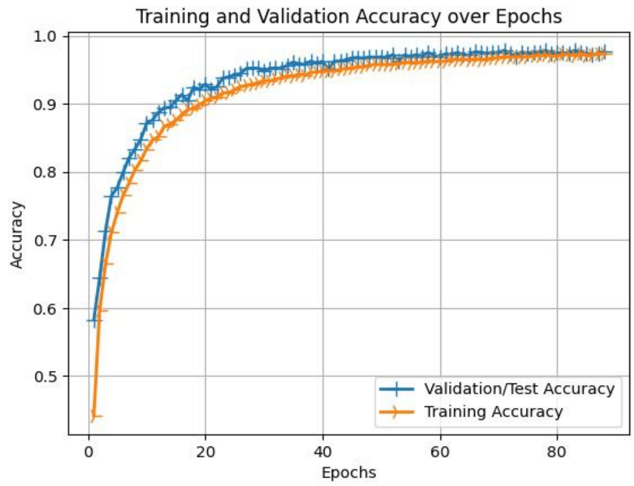

Figure 9: Training and Validation Accuracy over Epochs

So the system that was proposed actually got deployed successfuly with the Flask web framework, to deliver a sort of near real time skin disease analysis through a web interface which is kinda the main idea. Figures 10 to 12 show different bits of the user interface and the implementation screens for the application that was built. On the homepage, people can upload dermatoscopic images for analysis using an intuitive graphical interface, which is yes pretty straight forward. There are also extra sections, that go into more details like disease coverage, how the algorithm workflow moves step by step and even the database specifics that the system depends on.

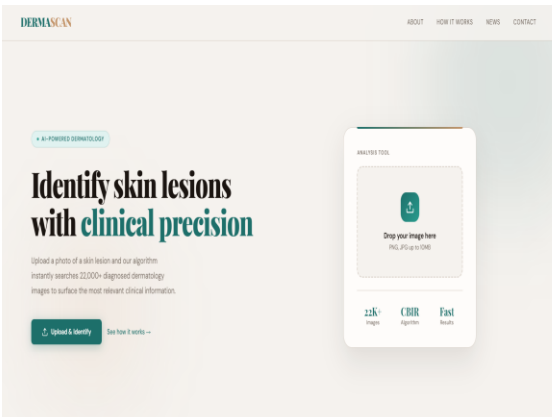

Figure 10: Homepage of the Proposed Skin Disease Detection System

Figure 10 more or less shows how the homepage looks for the proposed skin disease detection system, the one that was built using the Flask web framework. On this page you get a sort of more user friendly layout, with navigation menus , project details, and even several skin disease analysis options gathered in one place. It's like the entire interface is arranged so it feels simpler to access and not that complicated to use, especially for people who are working on automated skin disease prediction tasks.

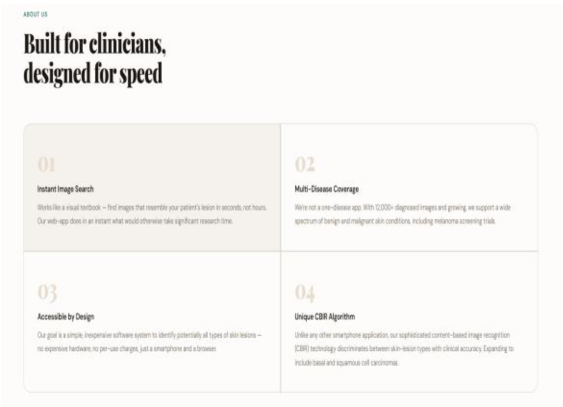

Figure 11: Features and Functionalities of the Web Application

Figure 11 shows the features and functionalities part of the web application that we developed, a bit like a tour screen, kinda. You can see the interface bringing forward the key system capabilities, for instance skin lesion image upload, AI based disease prediction, a confidence score generation step , and also healthcare support services. The layout is organized so it feels easier to navigate, and it also makes the experience smoother for users because it gives clear details about how the application workflow actually runs and what each functionality does.

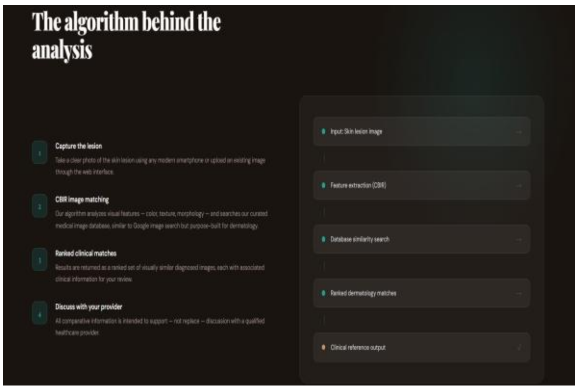

Figure 12: Workflow of the AI-Based Skin Lesion Analysis System

Figure 12 shows the workflow architecture for this AI based skin lesion analysis system. The diagram sort of explains the whole process from image upload, then preprocessing, after that MobileNet based feature extraction and disease classification and then the final prediction generation. It is meant to give a clearer sense about the operational sequence, how everything links together, and where the deep learning components sit inside the application. Even if you read it quickly, it still shows the flow in a pretty direct way.

Figures 13 and 14 show the final prediction outcomes made by the proposed system, kinda. The dermatoscopic images that were uploaded got classified as melanoma, with confidence levels of 90.91% and 96.39% respectively. In the meantime, the system presents the uploaded picture together with the prediction details, all in a structured style, so the whole result looks easier to read for users and also more helpful for healthcare professionals.

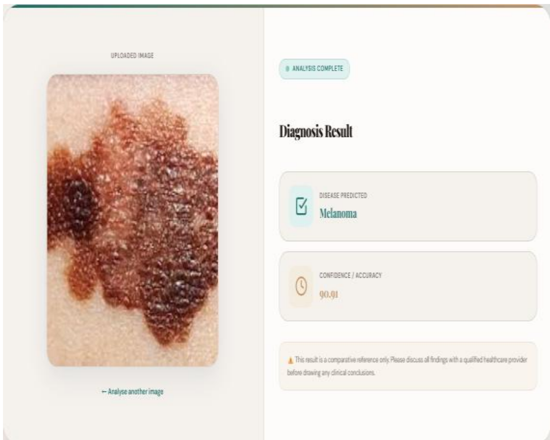

Figure 13: Prediction Result for Melanoma Detection with 90.91% Accuracy

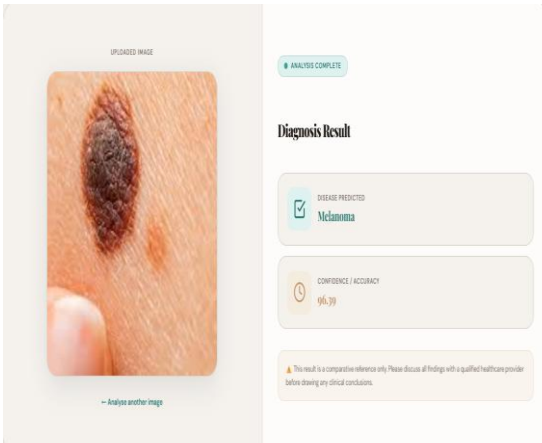

Figure 14: Prediction Result for Melanoma Detection with 96.39% Accuracy

All in all, the experimental analysis kind of confirms that the suggested MobileNet based skin disease detection system had pretty high classification accuracy, steady learning performance, and it also looks efficient for deployment through the web. By joining deep learning with Flask, you get a quicker early diagnosis, and it also makes it easier for users to reach automated skin disease analysis via an online platform.

Conclusion

The proposed AI driven skin disease detection system kind of shows how good Deep Learning really works for automated dermatological image classification. In particular by using the MobileNet transfer learning model, the system still manages to reach solid classification accuracy but with a smaller computational load so it feels a bit more fitting for real-time healthcare tasks. Before training, the whole image preprocessing and data augmentation routine helped the model generalize in a better way, and it also looked like it reduced overfitting, which is frankly always a plus. The Flask based web app lets people upload skin lesion images and then it returns prediction outcomes along with a confidence score, basically through a simple reachable interface. From the experiments, it also felt like the system can provide pretty consistent prediction results across multiple skin disease categories when it's evaluated on the HAM10000 dataset. So overall, this approach could support dermatologists for preliminary diagnosis, reduce manual work , and help improve healthcare access in remote regions. In the end, mixing AI with web technologies gives an efficient, and scalable way to deal with automated skin disease analysis and diagnosis.

References

P. N. Srinivasu, J. G. S. Sai, M. F. Ijaz, A. K. Bhoi, and J. J. Kang, "Classification of Skin Disease Using Deep Learning Neural Networks with MobileNet V2 and LSTM," *Sensors*, vol. 21, no. 8, pp. 1–17, 2021.

B. Shetty et al., "Skin Lesion Classification of Dermoscopic Images Using Machine Learning and Convolutional Neural Network," *Scientific Reports*, vol. 12, 2022.

K. Dimililer, "Skin Lesion Classification Using CNN-Based Transfer Learning," *Gazi University Journal of Science*, 2022.

P. Yao et al., "Single Model Deep Learning on Imbalanced Small Datasets for Skin Lesion Classification," *IEEE Transactions on Medical Imaging*, vol. 41, no. 5, pp. 1242–1254, 2022.

K. Lee et al., "Multi-Task and Few-Shot Learning-Based Fully Automatic Deep Learning Platform for Mobile Diagnosis of Skin Diseases," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 1, pp. 176–187, 2023.

A. Dascalu and E. David, "Skin Cancer Detection by Deep Learning and Sound Analysis Algorithms: A Prospective Clinical Study of an Elementary Dermoscope," *EClinicalMedicine*, vol. 15, pp. 107–113, 2020.

M. A. Al-masni et al., "Skin Lesion Segmentation in Dermoscopy Images via Deep Full Resolution Convolutional Networks," *Computer Methods and Programs in Biomedicine*, vol. 162, pp. 221–231, 2020.

H. Kaur, S. Pannu, and A. K. Malhi, "A Systematic Review on Imbalanced Data Challenges in Machine Learning," *ACM Computing Surveys*, vol. 52, no. 4, pp. 1–36, 2020.

S. Hameed et al., "A Transfer Learning Approach for Skin Lesion Classification Using MobileNet," *IEEE Access*, vol. 9, pp. 132433–132444, 2021.

P. Tschandl et al., "Human–Computer Collaboration for Skin Cancer Recognition," *Nature Medicine*, vol. 26, no. 8, pp. 1229–1234, 2020.

R. Gessert et al., "Skin Lesion Classification Using CNNs with Patch-Based Attention and Diagnosis-Guided Loss Weighting," *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 2, pp. 495–503, 2021.

M. Khan et al., "A Lightweight Deep Learning Model for Skin Lesion Classification," *Computers in Biology and Medicine*, vol. 133, pp. 104–117, 2021.

S. Yap et al., "Multimodal Skin Lesion Classification Using Deep Learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 11, pp. 4124–4134, 2021.

J. Wei et al., "Explainable Skin Cancer Classification Using Deep Learning," *Artificial Intelligence in Medicine*, vol. 115, 2021.

N. C. Codella et al., "Skin Lesion Analysis Toward Melanoma Detection," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 8, pp. 3004–3016, 2021.

Y. Liu et al., "Deep Learning System for Differential Diagnosis of Skin Diseases," *Nature Medicine*, vol. 26, pp. 900–908, 2022.

K. He et al., "Deep Residual Learning for Image Recognition in Medical Imaging Applications," *IEEE Access*, vol. 10, pp. 45122–45135, 2022.

M. A. Rahman et al., "Hybrid CNN Model for Multi-Class Skin Disease Classification," *Sensors*, vol. 22, no. 9, pp. 1–19, 2022.

S. Bhattacharya et al., "Cloud-Based AI Framework for Early Skin Disease Diagnosis," *IEEE Access*, vol. 11, pp. 33544–33556, 2023.

A. Ibrahim et al., "Explainable Artificial Intelligence for Skin Lesion Classification Using Transfer Learning," *Biomedical Signal Processing and Control*, vol. 84, 2023.

S. Khan, M. Nazir, and T. Hussain, "Deep Learning-Based Skin Lesion Classification Using MobileNet and Transfer Learning," *IEEE Access*, vol. 10, pp. 55421–55433, 2022.

R. Patel and A. Sharma, "Automated Skin Disease Detection Using Convolutional Neural Networks," *International Journal of Advanced Computer Science and Applications*, vol. 13, no. 4, pp. 112–120, 2022.