The Bioscan
AN INTERNATIONAL QUARTERLY JOURNAL OF LIFE SCIENCES
20(3): 51-56, 2025 www.thebioscan.com

Camera-Based Machine Learning for Skin Disease Detection: An On-Device Diagnostic Approach

Krishiv Garg (Bhupindra International Public School, Patiala)

Email: krishivgarg2008@gmail.com

DOI: 10.63001/tbs.2025.v20.i03.pp51-56

KEYWORDS
Skin disease detection, convolutional neural networks, mobile health, on-device inference, TensorFlow Lite, dermatology AI, image classification, explainable AI, DermaScan app, lightweight models.

Received on:

10-05-2025

Accepted on:

07-06-2025

Published on:

07-07-2025

ABSTRACT

One of the biggest problems in the Indian medical system is the lack of healthcare workers. The ever-increasing population of India has exerted significant pressure not just on the medical infrastructure but has also overwhelmed healthcare workers. This situation has caused the wait times for patients to increase significantly; speeding up this process requires the use of artificial intelligence, especially on the diagnostic side. Dermatology is one branch of medical science where diagnosis can be aided with camera image-based machine learning. Diagnosis of surface-level skin diseases that are easily visible to the naked eye can be made much quicker using image-based machine learning. It will enable doctors to expedite their diagnosis and also reduce the likelihood of an incorrect diagnosis. For the purpose of this paper, such an app called Dermascan has been created to test the real-world applications of this hypothesis.

INTRODUCTION

India's population is increasing at an unprecedented rate. It is well above 1.31 billion people, and this has exerted great stress on every industry, from human resources to medicine [1]. This ever-increasing population has caused the medical infrastructure to become overwhelmed, negatively affecting patient treatment, as there is a lack of healthcare workers. Some outpatient doctors(OPD) from Delhi have claimed that due to this overwhelming demand, they had to diagnose over 750 patients in one day, whereas providing highly accurate diagnoses and treatment would mean that they would be able to address the needs of a maximum of 250 patients a day[2]. Dermatology is one such profession that lacks the needed healthcare workers. This problem gets worse in rural areas where patients have no access or minimal access to dermatologists [3]. One highly qualified dermatologist interviewed for this paper claims that many patients have to travel across cities and even states to access dermatologists, which causes an influx of patients with skin diseases in one particular area, whereas others don't have a dermatologist, exerting great stress on both their doctors and staff, causing the demand to significantly exceed the dermatologist's physical ability to tend to patients. The Dermascan app has been created to test the abilities of camera image-based machine learning to decrease the speed of diagnosis. This experiment will allow us to gauge the problems with implementing such systems on a wider scale.

2. Literature review

This research area hasn't received the appropriate amount of attention due to the research on image-based machine learning skin disease detection being sparse.

2.1 Evolution of Image-Based Dermatology Diagnostics

Traditionally, computer vision-aided systems relied on colour and texture descriptions given by the programmers, but the recent advent of deep learning, particularly convolutional neural networks (CNNs), has revolutionised diagnostic accuracy. Esteva et al. in their research demonstrated that a single end-to-end convolutional neural network trained on 129,450 photographs could almost reach dermatologist-level performance in distinguishing benign from malignant lesions [4]. Haenssle et al. managed to prove that CNNs can outperform large cohorts of dermatologists. [5]

2.2 Public Datasets and Benchmark Tasks

Large, well-annotated datasets are key to creating efficient and highly accurate models that can detect and diagnose skin conditions. The HAM10000 collection (Tschandl et al., Scientific Data, 2018) introduced 10,015 dermatoscopic images across seven lesion types and remains a standard benchmark for multi-class skin lesion classification [6]. The popular open-source dataset platform Kaggle also has a comprehensive skin-disease classification dataset, which can classify up to 10 different kinds of skin diseases, like eczema, warts and ringworms [7]. Kaggle datasets and other open-source, high-quality datasets are essential to creating models and conducting research. Some images from this dataset were used in our model creation, which will be addressed further in the paper.

2.3 Mobile Deployment and Lightweight Models

With the increasing usage of smartphones, researchers have focused on on-device inference. Howard et al.'s MobileNet family paved the way for lightweight models, while recent work employs TensorFlow Lite to embed classifiers directly inside Android or iOS apps [8]. TensorFlow Lite offers a high inference speed, a small model size, and compatibility with various mobile platforms. These features are crucial for real-time applications like DermaScan, where latency, model size, and user privacy are paramount concerns. Another notable point is that embedding a TensorFlow Lite model in the app allows it to remain functional without significantly increasing the app size.

2.4 Problems with Machine Learning Models

A major issue with machine learning models is bias: skin disease detection models tend to produce inaccurate results when used by darker-skinned people as compared to lighter-skinned people, as the models are trained primarily on the images of lighter-skinned people [9]. This issue can be remedied by synthetic augmentation of patient data to balance skin tone distribution to include a larger variety of skin tones.

Another issue with using cameras to capture images is the model's dependence on environmental factors and technical factors: the model's prediction can be swayed by non-ideal lighting conditions or a worse camera quality with lower resolution images. This was a problem that we faced in the pilot testing of DermaScan, which led to lower model confidence, but the results were much less affected.

2.5 Explainability and Clinical Trust

To foster clinician adoption, explainable-AI (XAI) techniques such as Grad-CAM, Score-CAM and LIME have been integrated into dermatology models. Chanda, T et al. showed that saliency-map explanations aligned with dermatologists' reasoning and improved their diagnostic confidence [10]. Another challenge that needs to be addressed is patient trust; this was also addressed by the doctors interviewed for this paper: they stated that patients will be averse to an AI diagnosis and tend to prefer experienced human doctors.

2.6 Recent developments

Explainability has become a key research priority, with newer methods like Integrated Gradients and SHAP gaining popularity for clinical interpretability. Mobile-first deployments have also evolved, with real-time edge processing models optimised using quantisation and pruning techniques to minimise computational load without sacrificing accuracy. Recently, machine learning software like Teachable Machine has made it easier for researchers who lack a background in advanced machine learning to upload their data and create models for testing. Dermatologists, particularly in India, are yet to incorporate advanced artificial intelligence models for diagnostic assistance, but the future seems optimistic for artificial intelligence being used in disease diagnosis [11].

3. Methodology

3.1 Data Collection

Skin condition images across seven categories (Acne and Rosacea, Nail Fungus, Melanoma, Seborrhoeic Keratosis, Cutaneous Larva Migrans, Hairloss, Normal) were collected. A total of 983 images formed the dataset for training and validation. Only open-source, high-quality datasets from Kaggle were used. The data from the 'Skin-Disease-Dataset' was used for training data on nail fungus and cutaneous larva migrans [12]. The dataset 'Skin Diseases image dataset' was used, which included images of seborrhoeic keratosis, melanoma and some images for acne and rosacea [13]. 'The Various Skin Diseases Dataset' was used for nail fungus, acne, hair loss, and normal skin [14].

The overall training data comprised the following number of samples for each category:

1. Acne and Rosacea – 139 samples
2. Nail Fungus – 41 samples
3. Seborrhoeic Keratosis –278 samples
4. Cutaneous Larva Migrans – 15 samples
5. Hair loss – 15 samples
6. Normal – 24 samples

3.1 Model Training

Using Google Teachable Machine, images were preprocessed, resized to have a consistent size, and augmented to enhance the dataset's variability and robustness. A convolutional neural network model was trained over 100 epochs with a learning rate of 0.0098. The model was trained with a batch size of 16 and a validation split of 80:20 (80% of the samples were used for training the model, and 20% of the samples were used for testing). The model was trained to use the device camera to take input. The model resulted in output in the form of a probability on the scale from 0 to 100. The model was exported as a TensorFlow Lite model.

The CNN architecture used in this study, exported from Google Teachable Machine and visualised via the TensorFlow Lite visualising tool Netron, comprises multiple repeating blocks of convolutional operations followed by activation and pooling layers. Specifically, the model includes:

- **9 convolutional layers** using 2D filters to extract hierarchical features from input images
- **ReLU (Rectified Linear Unit)** activation functions are applied after each convolutional layer to introduce non-linearity.
- **MaxPooling layers** follow most convolutional blocks to reduce spatial dimensions and overfitting.
- **Flatten the layer** to convert the 2D feature maps into a 1D vector form for the dense layer.
- **2 fully connected (dense) layers** for final feature learning
- **Softmax output layer** to classify input images into one of the seven skin disease categories

The model structure supports hierarchical feature learning and efficient parameter reduction, making it suitable for mobile inference. Exported in TensorFlow Lite format, it is optimised for use within the DermaScan mobile application.

The TensorFlow Lite model was preferred over the standard TensorFlow method due to its lightweight nature. It allows the app to be fast, efficient and accurate without causing it to have a large app size or require large processing power. This allows the app to run on lower-end devices without any issues with either the app or the device.

4. App Development

Figure 1: The DermaScan App Snapshot

The trained TensorFlow Lite model was integrated into the DermaScan Android application; it enables real-time skin condition classification directly from smartphone camera images. It has a total of 4 pages, including one homepage, a page to scan and diagnose diseases, the second page to check when every image was taken and which disease the app detected, and an AI search page for the user to gain basic information about various diseases.

4.1 Scan page

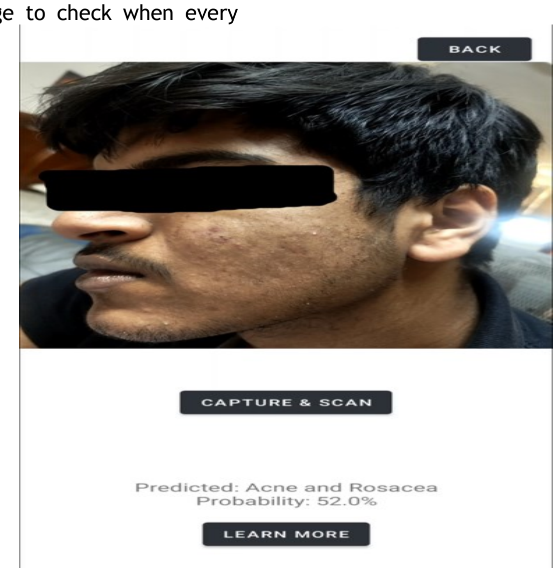

Figure 2: The DermaScan Scan page snapshot

The DermaScan App uses the device camera to capture images of patients and diagnose skin conditions. After the image is captured, it is used as input by the machine learning model, and the app outputs the condition along with the probability of the patient having that condition. There is also a learn more button that allows the user to learn more about the skin condition from either the Cleveland Clinic or the Mayo Clinic. Pressing the learn more button takes the user to the respective disease page.

4.2 Patient History Page

Figure 3: The DermaScan Patient History Page Snapshot

Doctors are required to maintain a record of patient history, and the patient history page records an entry every time the capture and scan button is pressed on the scan page. The history is stored locally on the device; this prevents an electricity outage in the server from deleting all the records. There is also a clear history button that allows you to delete every single record stored locally.

4.3 AI Search Page

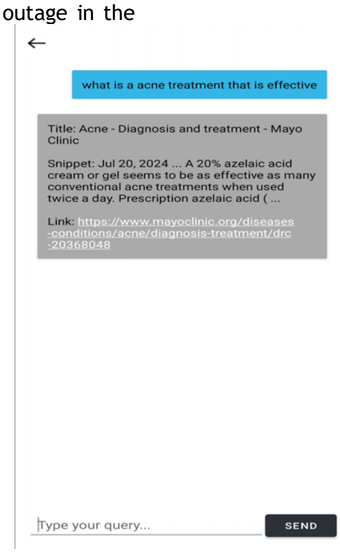

Figure 4: The DermaScan AI Search Page

The app has a built-in AI search page that allows the user to get basic information about any disease from the Cleveland Clinic and the Mayo Clinic. The user gets a preview of some information and then a link at the bottom, which, upon clicking, takes the user to the page. The search uses the data from the entirety of the two websites for this chatbot.

5. Results

5.1 Accuracy

The CNN model trained using Google Teachable Machine yielded positive results. According to our parameters, 20% of the samples were used for testing, and the results were as follows:

Class	Accuracy	Number of Samples
Acne and Rosacea	0.83	139
Nail Fungus	0.95	41
Melanoma	1.00	471
Seborrhoeic Keratosis	0.93	278
Cutaneous Larva Migrans	1.00	15
Hair loss	0.87	15
Normal	0.75	24

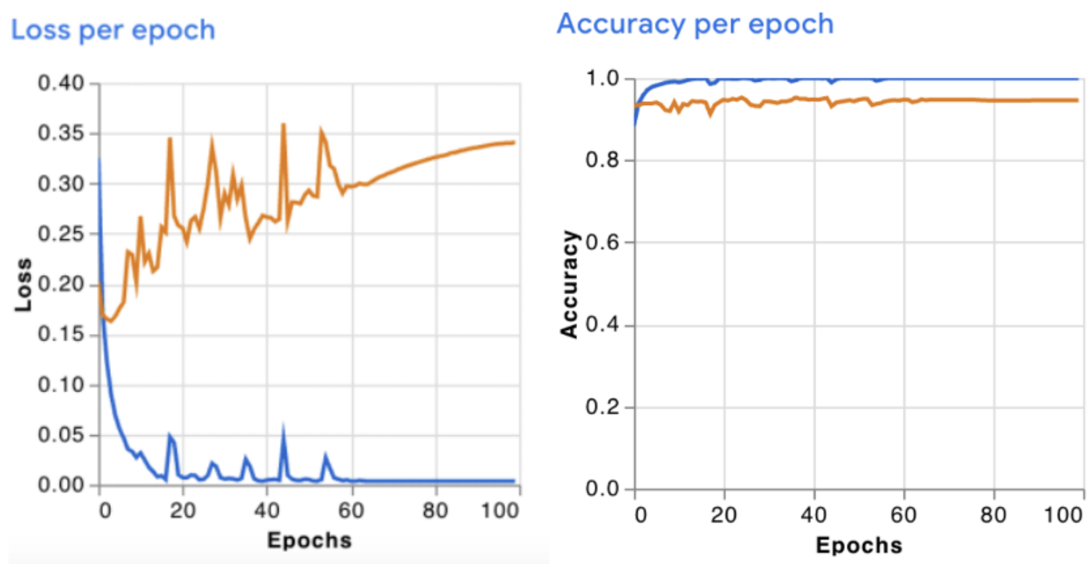

Figure 5: Loss Per Epoch and Accuracy Per Epoch

The loss per epoch graph reveals a potential issue. While training loss continuously decreases toward zero, validation loss initially drops but begins rising steadily after approximately epoch 40. This divergence suggests the model started to overfit to the training data beyond that point, memorising patterns rather than learning generalised features.

The accuracy per epoch graph shows that the training accuracy (blue) quickly converges to nearly 100% within the first 20 epochs and remains stable thereafter. The validation accuracy (orange) also improves early and plateaus around 90%, indicating strong generalisation initially.

5.2 Confusion Matrix Results

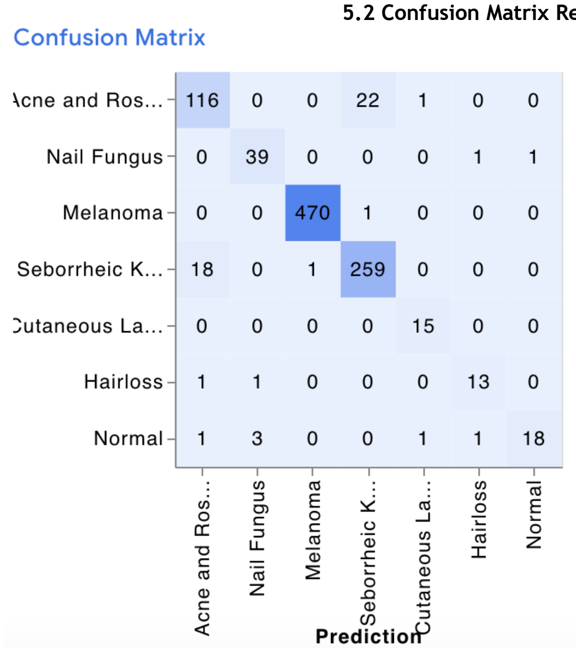

Figure 6: Confusion Matrix for CNN Model

Figure 5 shows the confusion matrix for the validation dataset, summarising the performance of the CNN model across seven skin disease classes. Melanoma and cutaneous larva migrans achieved perfect classification with 470 and 15 correct predictions, respectively. The model also performed strongly on Nail Fungus (39/41) and Seborrhoeic Keratosis (259/278).

Most misclassifications occurred between acne and rosacea and seborrhoeic keratosis, with 22 acne samples misclassified as seborrhoeic keratosis and 18 errors in the reverse direction. The normal class was the most difficult to classify, with predictions spread across nail fungus, acne, and even cutaneous larva migrans. This may be due to overlapping features or limited representation in the training dataset.

Overall, the confusion matrix highlights the model's excellent performance in high-priority conditions like melanoma while also identifying areas needing improvement, such as distinguishing visually similar or under-represented categories.

CONCLUSION

The research in this paper demonstrates the potential of camera-based deep learning models for accurate skin disease diagnosis. By leveraging convolutional neural networks capable of distinguishing between seven dermatology conditions with high accuracy, the resulting TensorFlow model that was integrated into the Dermascan mobile application offered a practical, real-time diagnostic tool. The most commendable breakthrough was that the model achieved perfect accuracy for high-priority skin conditions, such as melanoma and cutaneous larva migrans, while still performing well across other classes. The confusion matrix and accuracy trends highlight the effectiveness of this model. Overall, it can represent a significant step towards providing more accessible and faster dermatology diagnostics.

Future Work

Although Dermascan has demonstrated very strong classification accuracy, which is essential for our use case, there is still work that can be done for future research. Firstly, future efforts should focus on collecting datasets that have a more balanced representation across all skin tones, age groups, and disease subtypes. This will help reduce bias. Secondly, researchers will need to fine-tune models as they continue to increase the number of diseases that can be diagnosed. Thirdly, incorporating explainable AI (XAI) tools like Grad-CAM can increase clinical trust and help dermatologists interpret predictions more effectively. Dermascan is undergoing pilot testing, and results seem to be positive, with dermatologists reporting higher efficiency and occasional inaccuracies. This data coincides with the theoretical results yielded by the model.

REFERENCES

- Agarwal, A., Lubet, A., Mitgang, E., Mohanty, S., & Bloom, D. E. (2020). Population ageing in India: facts, issues, and options. In *New frontiers in regional science: Asian perspectives* (pp. 289-311). https://doi.org/10.1007/978-981-10-0230-4_13
- Saikia, N., Sehgal, S., Bora, J. K., & Diamond-Smith, N. (2017). Experiences of neonatal deaths among urban poor in Metropolitan Delhi: commination, poor quality, and overwhelmed system. *Journal of Social Inclusion Studies*, *3*(1–2), 1-19. https://doi.org/10.1177/2394481120170101
- Duniphin, D. (2023). Limited access to dermatology specialty care: barriers and teledermatology. *Dermatology Practical & Conceptual*, e2023031. https://doi.org/10.5826/dpc.1301a31
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, *542*(7639), 115-118. https://doi.org/10.1038/nature21056
- Haenssle, H., Fink, C., Schneiderbauer, R., Toberer, F., Buhl, T., Blum, A., Kalloo, A., Hassen, A. B. H., Thomas, L., Enk, A., Uhlmann, L., Alt, C., Arenbergerova, M., Bakos, R., Baltzer, A., Bertlich, I., Blum, A., Bokor-Billmann, T., Bowling, J., . . . Zalaudek, I. (2018). Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. *Annals of Oncology*, *29*(8), 1836-1842. https://doi.org/10.1093/annonc/mdy166
- Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, *5*(1). https://doi.org/10.1038/sdata.2018.161
- *Skin diseases image dataset*. (2021, August 16). Kaggle. https://www.kaggle.com/datasets/ismailpromus/skin-diseases-image-dataset
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017, April 17). *MobileNets: efficient convolutional neural networks for mobile vision applications*. arXiv.org. https://arxiv.org/abs/1704.04861
- Fliorent, R., Fardman, B., Podwojniak, A., Javaid, K., Tan, I. J., Ghani, H., Truong, T. M., Rao, B., & Heath, C. (2024). Artificial intelligence in dermatology: advancements and challenges in skin of color. *International Journal of Dermatology*, *63*(4), 455–461. https://doi.org/10.1111/ijd.17076
- Chanda, T., Hauser, K., Hobelsberger, S., Bucher, T., Garcia, C. N., Wies, C., Kittler, H., Tschandl, P., Navarrete-Dechent, C., Podlipnik, S., Chousakos, E., Crnaric, I., Majstorovic, J., Alhajwan, L., Foreman, T., Peternel, S., Sarap, S., Özdemir, İ., Barnhill, R. L., . . . Brinker, T. J. (2024). Dermatologist-like explainable AI enhances trust and confidence in diagnosing melanoma. *Nature Communications*, *15*(1). https://doi.org/10.1038/s41467-023-43095-4
- Sengupta, D. (2023). Artificial intelligence in diagnostic Dermatology: challenges and the way forward. *Indian Dermatology Online Journal*, *14*(6), 782-787. https://doi.org/10.4103/idoj.idoj_462_23
- *Skin-Disease-Dataset*. (2023, October 14). Kaggle. https://www.kaggle.com/datasets/subirbiswas19/skin-disease-dataset
- *Skin diseases image dataset*. (2021b, August 16). Kaggle. https://www.kaggle.com/datasets/ismailpromus/skin-diseases-image-dataset
- *Various skin diseases dataset*. (2021, March 1). Kaggle. https://www.kaggle.com/datasets/akshitmadan/various-skin-diseases-dataset?resource=download