

The Good AI Skin Doctor: A Privacy-First AI Assistant for Dermatological Wellness and Early Detection

Mrs. Nita Meshram¹, A Pavithra², Akshaya K³, Amisha V⁴, and Bhavya D⁵

¹Associate Professor, department of Computer Science, K. S. School of Engineering and Management

^{2,3,4,5} Student, department of Computer Science, K. S. School of Engineering and Management

Abstract—Skin conditions are widespread throughout the world, but due to the scarcity of dermatologists and the difficulty of detecting the issue in its early stages, it is a challenge for a lot. This paper presents The Good AI Skin Doctor, an AI-powered system that uses images to automatically identify possible skin problems for users. The HAM10000 dermatology dataset was used to train a refined EfficientNet-B3 model for the classification task, taking into account important preprocessing techniques like resizing, normalizing, and augmenting data to handle images with different qualities. When a user uploads an image, this model categorizes the lesion into pre-established groups and gives the user initial insights before consulting a medical professional. Based on the anticipated category, a straightforward recommendation system offers broad skincare suggestions. In the meantime, the system's workflow for safely processing and momentarily storing user-provided images protects users' privacy. Overall, the current work highlights how Digital vision and deep learning can support early skin health awareness by providing a dependable and user-friendly AI assistant for initial dermatological assessment.

Index Terms—Artificial Intelligence, Computer Vision, Deep Learning, Dermatology Assistance, EfficientNetB3, HAM10000, Image Preprocessing, Skin Image Classification, Skin Health Analysis, Privacy-Focused AI.

I. INTRODUCTION

Skin disorders continue among the most common conditions reported worldwide; early detection is very vital in preventing the progression of the disease. Timely dermatological evaluation is challenging for many due to limited access to specialists, high consultation costs, and general lack of awareness about early symptoms. The recent improvements in artificial intelligence as well as easy

availability of digital imaging devices have propelled automated skin-lesion evaluation as a promising solution to support early screening efforts.

The Good AI Skin Doctor aims to create and implement a useful AI-based skin condition screening tool according to the fine-tuned EfficientNet-B3 model, pre-trained on the HAM10000 dataset. The system workflow includes the preparation of the dataset, the preprocessing and augmentation of images, and model training within Google Colab. Upon uploading an image by a user, the trained model will perform the analysis to yield a preliminary classification that acts as an accessible first step in their journey to seek early insight into skin issues.

Besides image classification, the system gives basic skincare recommendations matching the conditions it predicts, thereby allowing the user to understand general care recommendations that might generally relate to commonplace skin problems. A privacy-first mechanism ensures that uploaded images are processed securely and retained only for a temporary period during prediction, keeping ethical standards for sensitive health-related data.

The work's primary contributions are:

- A. Fine-tuning an EfficientNet-B3 skin-lesion classification prototype on the HAM10000 dataset.
- B. Preprocessing and augmentation techniques will be incorporated to ensure better performance for images of variable quality.
- C. Implement a lightweight recommendation module for basic guidance with prediction results.
- D. Securely and temporarily process user-uploaded images by means of a privacy-first workflow.

The Good AI Skin Doctor showcases how AI can improve dermatological screening in terms of accessibility and early awareness, especially among people without immediate access to dermatology specialists, by integrating deep-learning-driven image analysis with an easy-to-use, privacy-protected workflow.

II. LITERATURE SURVEY

Recent developments in deep learning and the powerful performance of Convolutional Neural Networks have led to an increase in the use of artificial intelligence in dermatology.

- A. Sonavane et al. introduced DermDetectNet, a CNN-based framework that classifies a several of skin diseases. Their study, among others, solidified the fact that CNN architectures have the potential to accomplish high classification accuracies. Nevertheless, their model relied heavily on pre-trained networks, which might restrict its applicability to realworld photos taken under various circumstances.
- B. Senevirathna et al. proposed a deep learning-based system of facial analysis that gives skincare and cosmetic recommendations. Their approach merged image features with user-specific attributes to provide personalized results. Though successfully used for cosmetic advice, the system was never meant or designed for medical-level detection of skin lesions or other diagnostic purposes.
- C. Chandana et al. conducted a deep learning-based study on facial skin analysis by utilizing CNN models. The results of their study highlight the significance of preprocessing techniques like handling noise, augmenting, and normalizing in order to enhance the model's functionality, particularly while utilizing .small or unbalanced dermatology datasets. Their findings stress how image quality and preparation influence the categorization precision.
- D. Soni et al. used a fine-tuned ResNet152V2 model to improve the division of skin diseases. Their findings indicated that deeper CNN models can further promote the detection ability for dermatological conditions. The study primarily focused on the output of the classification task and did not investigate complementary features,

such as recommendation systems, or practical deployment.

When taken together, these research show the high efficacy of CNN-based methods for dermatology image classification, but they also highlight important issues, such as dataset imbalance, image quality variation, and the requirement for well optimized preprocessing and fine-tuning techniques.

III. METHADODOLOGY

The Good AI Skin Doctor system was created for AI-based skin lesion classification and functions as an end-to-end pipeline. The methodology describes the dataset acquisition, data preprocessing tasks, model choice, training workflow, prediction mechanism, and the privacy structure that protects user images.

A. Data Collection

The HAM10000 dataset, one of the biggest collections of dermatoscopic lesion photos from seven classes, is utilized the system. Melanocytic nevi, benign keratosis, melanoma, dermatofibroma, vascular lesions, and other conditions are represented in the dataset. The dataset and its metadata were obtained through Kaggle during development.

To enable accurate performance assessments, the data was split into three distinct sets for training, verifying, and testing.

B. Data Preprocessing

The following steps were applied for the standardization of pictures prior to being fed into the deep learning model:

1. Image Resizing: Each image was resized to 300×300 pixels, as that is the expected input dimension for EfficientNet-B3.
2. Normalization: To stabilize model training, the pixel intensity levels were scaled scaled from 0 to 1.
3. Augmentation: To improve diversity and lessen overfitting, a number of methods were used, including rotation, shifting, zooming, brightness change, and horizontal flipping.
4. Directory Structuring: To facilitate the Keras Image Data Generator to load the dataset, it was rearranged into folder-based class directories.

These operations enhanced the capacity of the model to manage variations in illumination, angle, orientation, and overall image quality.

C. Skin Lesion Detection Using EfficientNet-B3

EfficientNet-B3 was chosen as the primary architecture since it has shown an exceptionally good balance between accuracy and computation.

1. Feature Extraction: The pretrained EfficientNet-B3 network captures meaningful spatial, structural, and color-based patterns from lesion images.
2. Fine-Tuning: All the model layers were unfrozen to carry out domain-specific training on the HAM10000 dataset.
3. Classification Head: A dense softmax layer was added to produce probabilities over seven lesion categories.

This configuration yielded consistent classification results while being deployment-friendly.

D. Model Training

Keras and TensorFlow were used for training.

1. Loss Function: To address class imbalance and concentrate on more difficult-to-classify samples, a unique focal loss function was employed.
2. Optimizer: The Adam optimizer was used to ensure smooth convergence with a stable learning rate.
3. Class Weights: Automatically generated class weights reduced the effect of uneven class distribution.
4. Training Callbacks:
 1. ModelCheckpoint for saving the best-performing model
 2. EarlyStopping to reduce overfitting
 3. ReduceLROnPlateau for automatic learning rate scheduling
5. Model Export: The final model was saved in both .keras formats to support multiple deployment environments.

This training procedure ensured effective model learning and fast inference.

E. Basic Skincare Recommendations

Based on the classification, the system generates general skincare recommendations that are in line with the predicted category.

These recommendations are not personalized, yet provide simple guidance on care and indicate when professional consultation may be necessary.

F. Privacy and Data Handling

A privacy-oriented framework is integrated for the safe management of dermatology images and user data.

1. Encrypted Image Storage: Uploaded photos are kept in an encrypted format that prevents unwanted access.
2. User-specific access: Every stored image is hosted with the email of an authenticated user to ensure private access via the personal history panel.
3. Isolation of Image Access: One cannot view or retrieve the images uploaded by others. This ensures strict confidentiality.
4. Secure Backend Processing: Prediction and history are performed at the back end, which is secured and follows secure data-handling practices.

These mechanisms ensure that sensitive dermatological images remain confidential and secure.

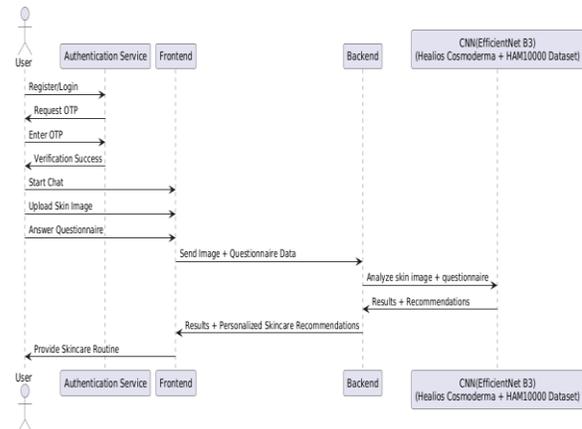


Fig. 1. System Sequence Diagram of The Good AI Skin Doctor

IV. MAIN MODEL

EfficientNet-B3 is the central utilizing a deep learning model used in the Good AI Skin Doctor system's application of skin disease detection. Because of its strong accuracy, compound scaling method, and suitability for real-time dermatological analysis, this architecture was chosen.

A. EfficientNet-B3 Architecture

The foundation of this system is EfficientNet-B3, which exhibits outstanding predictive accuracy balance. The model receives standardized 300×300 input images and extracts hierarchical patterns of texture, asymmetry, pigmentation, and structure. It was initialized with ImageNet weights and later fully unfrozen to adapt to the HAM10000 dataset, enhancing its ability to generalize to medical images.

B. Model Training and Optimization

Training was done in Google Colab, following the steps designed to maximize reliability.

1. Focal Loss: In order to solve the issue of class inequality in HAM10000, a custom focal loss was used in place of categorical cross-entropy.
2. Adam Optimizer: This enhanced fine-tuning stability by using the Adam optimizer with a low learning rate of $1e-5$.
3. Augmentation: The Keras Image Data Generator applied:
 1. rotations up to 30°
 2. Zoom changes
 3. brightness variations
 4. width/height shifts
 5. horizontal flips
 These made more of model resilient to real-world variations, such as lighting differences and skin tone diversity.
4. Class Weights: Computed class weights were included to reduce performance bias between frequent and rare classes.
5. Training Strategy:
 1. EarlyStopping reduced overfitting
 2. ModelCheckpoint preserved the highest-performing model
 3. ReduceLROnPlateau adjusted the learning rate automatically

Output Layer: The extracted features go via global average pooling, thick layers, dropout, and a softmax layer to create the final class probabilities.

C. Model Evaluation

Performance evaluation on the test dataset used:

1. Accuracy
2. Precision, Recall, F1-score
3. Classification report
4. Confusion matrix

These metrics demonstrate the model's capacity to differentiate between visually similar lesion types.

D. Deployment and Export

Keras (.keras): Full-precision version for server-based systems.

This ensures compatibility with the Flask backend and future mobile application integration.

E. Advantages of the Implemented EfficientNet-B3 Model

1. High accuracy at moderate computational cost
2. Improved handling of imbalanced datasets using focal loss + class weights

3. Robustness to real-world inputs owing to extensive augmentation
4. Deployment-friendly with both Keras and TFLite versions
5. Better performance with full-layer fine-tuning

V. RESULTS

The test portion of the HAM10000 dataset, for which the EfficientNet-B3 model had been trained, was then used to assess the suggested system's performance. It turned out that the model showed consistent and reliable classification capability across multiple skin lesion categories that were included in the dataset. Fine-tuning together with augmentation and class-balancing techniques contributed much to the stability of the model predictions.

The preprocessing pipeline turned out to be a very important factor in enhancing accuracy. In The model's capacity to differentiate between resistant to real-world variations like lighting, quality differences between cameras, and small shifts in object position, it went about standardizing image dimensions, normalizing pixel intensity, and creating augmented samples. As a result, even when tested using user-supplied images that might not match clinical dermatoscopic images, the system produced high classification performance.

In order to help the system determine the most likely skin condition from the uploaded image, the model provided a probability score across classes upon inference. These predictions were produced in a matter of seconds, which is considered to be within the efficiency benchmark required for real-time user interaction. The system thereafter converted these predictions into simple, understandable suggestions tailored to the predicted class to help users interpret the results without requiring technical knowledge.

While it is not designed for medical diagnosis, the experimental results prove that meaningful preliminary guidance can be presented through the system. The system becomes a practical supporter in early awareness, as it recognizes visible patterns and presents general skincare suggestions. Because of this, It is especially beneficial for users who might not have immediate access to clinical resources or dermatology specialists.

VI. CONCLUSION

The Good AI Skin Doctor system proposes a practical and effective AI-based approach for assisting in dermatological diagnosis. The system is trained on the publicly available HAM10000 dataset and is built on the EfficientNet-B3 network architecture, which was selected due to its excellent performance with good computational balance. This combination of data sources enhances feature learning and strengthens the model's ability to classify multiple skin conditions with great accuracy.

The system, which analyze real-time images taken by the user, provides skin condition recommendations for causes, preventive care, and basic routines related to that condition. especially for people who might not have direct access to dermatology specialists, this advice aids users in gaining early insights into the health of their skin.

The overall work shows the efficacy of EfficientNet-B3 for dermatological image classification and the increasing importance of AI-aided systems in early detection, skin health awareness, and user-oriented skincare support. Particularly for those who might not have direct access to dermatology specialists, this advice aids users in gaining early insights into the health of their skin.

VII. FUTURE SCOPE

Future developments of this system will be related to enhancing its clinical robustness and user accessibility. The capacity of the model to generalize to a range of populations can be improved by expanding the dataset to include more skin conditions and skin tone variations. A mobile application may be included to allow quicker and easier uploads of images. Explainable AI techniques, such as Grad-CAM, can be implemented to provide visual interpretability to model predictions, enhancing transparency and building user trust. Inclusion of basic patient history together with image data can further enhance diagnostic accuracy. By utilizing telemedicine, the option to consult with a dermatologist can also be made available. Over time, the model will be iteratively improved and retrained using ongoing expert-guided feedback. Lastly, cloud infrastructure can house the system for improved scalability, reliability, and global access.

REFERENCES

- [1] S. S. Sonavane, D. S. Thosar, B. R. Wankar, K. Jadhav, P. A. More, and A. Kulkarni, "DermDetectNet: Identifying skin diseases with advanced computer learning," in *Proc. 2nd DMIHER Int. Conf. Artif. Intell. Healthcare, Educ. Ind. (IDICAIEI)*, 2024, pp. 1–6.
- [2] O. Senevirathna, S. Rathnayake, and T. Thilakarathna, "Facial diagnosis skin care and makeup recommendation using advanced deep learning techniques," in *Proc. 6th Int. Conf. Advancements in Computing (ICAC)*, 2024, pp. 360–365.
- [3] K. Chandana, N. Swathi, V. Vijay, and M. A. Kumar, "A deep learning analysis on facial skin using CNN," in *Proc. IEEE Int. Conf. Intelligent Systems, Smart and Green Technologies (ICISSGT)*, 2024, pp. 108–113.
- [4] A.-C. Tsai, P. P.-H. Huang, Z.-C. Wu, and J.-F. Wang, "Advanced pigmented facial skin analysis using conditional generative adversarial networks," *IEEE Access*, vol. 12, pp. 46646–46656, 2024.
- [5] T. Soni, D. Gupta, and M. Uppal, "Optimizing facial skin disease classification using ResNet152V2 architecture," in *Proc. 3rd Int. Conf. Automation, Computing and Renewable Systems (ICACRS)*, 2024, pp. 871–874.
- [6] D. S. Thosar and N. Gandhewar, "An advanced image authentication using pass image algorithm to resist shoulder surfing attack," *Comput. Integr. Manuf. Syst.*, vol. 28, no. 10, pp. 52–59, 2022.
- [7] D. S. Thosar, R. R. Shinde, P. J. Gadakh, and P. V. Kashid, "Secure kNN query processing in entrusted cloud environments," *Asian J. Convergence Technol.*, vol. 2, no. 1, 2016.
- [8] K. Mridha, M. M. Uddin, J. Shin, S. Khadka, and M. F. Mridha, "An interpretable skin cancer classification using optimized convolutional neural network for a smart healthcare system," *IEEE Access*, vol. 11, pp. 41003–41018, 2023.
- [9] M. Al Mamun and M. S. Uddin, "A survey on a skin disease detection system," *Int. J. Healthcare Inf. Syst. Informatics*, vol. 16, no. 4, 2021.
- [10] D. Połap, A. Winnicka, K. Serwata, K. Kęsik, and M. Woźniak, "An intelligent system for monitoring skin diseases," *Sensors*, vol. 18, no. 8, Aug. 2018.

- [11] M. Chen, P. Zhou, D. Wu, L. Hu, M. M. Hassan, and A. Alamri, "AI-Skin: Skin disease recognition based on self-learning and wide data collection through a closed loop framework," *arXiv preprint*, Jun. 2019. [Online]. Available: <http://arxiv.org/abs/1906.01895>
- [12] Q. Abbas, F. Ramzan, and M. U. Ghani, "Acral melanoma detection using dermoscopic images and convolutional neural networks," *Vis. Comput. Ind. Biomed. Art*, vol. 4, no. 1, Dec. 2021.
- [13] B. Shetty, R. Fernandes, A. P. Rodrigues, R. Chengoden, S. Bhattacharya, and K. Lakshmana, "Skin lesion classification of dermoscopic images using machine learning and convolutional neural network," *Sci. Rep.*, vol. 12, no. 1, Dec. 2022.
- [14] Z. Lan, S. Cai, X. He, and X. Wen, "FixCaps: An improved capsules network for diagnosis of skin cancer," *IEEE Access*, vol. 10, pp. 76261–76267, 2022.
- [15] D. S. Thosar and M. Singh, "A review on advanced graphical authentication to resist shoulder surfing attack," in *Proc. Int. Conf. Advanced Computation and Telecommunication (ICACAT)*, 2018, pp. 1–3.
- [16] P. N. Srinivasu, J. G. Sivasai, M. F. Ijaz, A. K. Bhoi, W. Kim, and J. J. Kang, "Classification of skin disease using deep learning neural networks with MobileNetV2 and LSTM," *Sensors*, vol. 21, no. 8, Apr. 2021.