Grenze International Journal of Engineering and Technology, January Issue

# Deep Learning based System for Skin Disease Detection

Taalam Naga Raju[1] and Eswaraiah Rayachoti[2]
[1-2]SCOPE, VIT-AP University, Andhra Pradesh, India
Email: mailto: nagaraju.23phd7096@vitap.ac.in, eswaraiah.rayachoti@vitap.ac.in

***Abstract*— Skin diseases are among the most prevalent health concerns worldwide, with millions affected annually by conditions ranging from common infections to serious illnesses such as melanoma. These disorders can significantly impair quality of life and, in some cases, lead to severe medical complications if not diagnosed and treated promptly. However, access to dermatological care remains limited in many regions, particularly in rural and underserved areas, where delays in diagnosis are common due to a shortage of trained professionals and diagnostic infrastructure. In this context, artificial intelligence (AI) presents a promising avenue for enhancing dermatological diagnostics. This research introduces a deep learning-based system for the automatic classification of skin diseases using image data. The proposed model is built upon ResNet-50, a state-of-the-art convolutional neural network (CNN) architecture known for its ability to extract and learn complex hierarchical features from visual inputs. A diverse and curated dataset of labeled skin disease images was used to train and evaluate the model. The system effectively identifies and classifies conditions such as eczema, psoriasis, acne, and melanoma, each of which presents with visually distinct characteristics. This research not only underscores the viability of deep learning in medical image classification but also highlights the broader impact such technologies can have in democratizing healthcare.**

***Index Terms*— Vitamin Deficiency Detection, Computer Vision, Image Processing, Healthcare Diagnostics.**

## I. INTRODUCTION

Skin diseases represent a widespread global health challenge, affecting individuals across all age groups and geographies. These conditions range from common and treatable issues like acne and eczema to more serious and potentially life-threatening diseases such as melanoma. Early detection and proper diagnosis are essential to initiate timely treatment and prevent complications. However, access to dermatological care remains limited in many rural and underserved regions due to a shortage of trained professionals and diagnostic infrastructure. Even in well-equipped urban centers, the increasing number of patients often overwhelms dermatologists, leading to time constraints and potential diagnostic errors.

Additionally, many skin conditions share similar visual symptoms, such as redness, lesions, and scaling, which makes accurate differentiation difficult. Variations in skin tone, lighting conditions, and image quality further add to the complexity, reducing the reliability of manual visual examination. Advancements in artificial intelligence, particularly in the field of computer vision, present an opportunity to address these challenges.

*Grenze ID: 01.GIJET.12.1.551_1*

Convolutional neural networks (CNNs) have shown excellent performance in image classification tasks and can be leveraged to analyze skin images for disease detection. Despite the potential, AI- based dermatology tools remain underutilized in real-world healthcare systems.

This research aims to develop a deep learning-based skin disease classification model using the ResNet-50 architecture. The goal is to create an automated, non-invasive, and scalable diagnostic aid capable of classifying common dermatological conditions including eczema, acne, psoriasis, and melanoma, based on image data. Such a system can assist healthcare providers, improve early diagnosis, and offer accessible dermatological support in areas lacking expert medical care. This research will encompass several key aspects:

- Data preprocessing and preparation, including image labelling and augmentation techniques to ensure the model's robustness.
- Training the ResNet-50 model using transfer learning, fine-tuning the model to achieve optimal performance.
- Evaluating the model's effectiveness using classification metrics such as accuracy, precision, recall, and F1 score.

While the system is intended to support early diagnosis and aid healthcare providers, it is not designed to replace professional medical advice. The study is limited to image-based classification and does not consider other factors such as patient history or physical examination.

The AI-based skin disease detection system has significant relevance in various fields:

- Healthcare Systems: It enables early, accurate diagnosis of skin conditions, particularly in remote or underserved areas, improving access to dermatological care.
- Telemedicine: Supports virtual consultations by providing reliable, automated skin disease classification for accurate remote diagnosis.
- Mobile Health Apps: Allows users to upload images and receive immediate feedback on potential skin conditions, promoting proactive health management.
- Clinical Decision Support: Assists dermatologists in diagnosing challenging cases, reducing errors, and improving patient care.

These applications demonstrate the system's ability to enhance diagnosis speed and accessibility, improving healthcare delivery.

## II. LITERATURE REVIEW

Simonyan et al. [1] proposed VGGNet, a seminal convolutional neural network (CNN) architecture that demonstrated the importance of depth in visual recognition tasks. By stacking multiple small (3x3) convolutional layers with periodic max-pooling, the authors showed that increasing network depth (up to 16–19 weight layers) significantly improves accuracy on large-scale datasets like ImageNet, achieving top-5 test error rates of 7.3% (VGG-16) and 6.8% (VGG-19) in the 2014 ILSVRC challenge.

Key innovations include the use of small receptive fields throughout the network, which enhance nonlinearity and reduce parameters compared to larger kernels, and a simple, uniform design that prioritizes depth over architectural complexity. Despite its computational intensity, VGGNet's modular structure and strong performance made it a foundational model for later architectures (e.g., ResNet) and transfer learning. Its pre-trained weights remain widely used in computer vision for tasks like object detection and segmentation, though modern networks have since surpassed its efficiency with innovations like skip connections and batch normalization.

He et al. [2] introduced ResNet, a groundbreaking deep convolutional neural network architecture that addresses the degradation problem—where deeper networks exhibit higher training error despite their increased theoretical capacity. The key innovation is the residual learning framework, which uses shortcut connections (skip connections) to bypass layers, enabling the network to learn identity functions more easily and facilitating the training of extremely deep models (up to 152 layers). ResNet achieved a top-5 error rate of 3.57% on the ImageNet dataset, winning the ILSVRC 2015 competition, and demonstrated superior generalization across other vision tasks like object detection and segmentation. The architecture's simplicity, scalability, and effectiveness in mitigating vanishing gradients made it a cornerstone of modern deep learning, influencing subsequent models (e.g., DenseNet, ResNeXt). While computationally intensive, ResNet's design principles—emphasizing residual mappings over direct layer-to-layer transformations—remain foundational for training very deep networks efficiently.

Esteva et al. [3] demonstrates how deep learning can achieve medical expert-level performance in diagnosing skin cancer. The authors trained a deep convolutional neural network (CNN) on a large dataset of nearly 130,000 clinical images spanning over 2,000 skin diseases, using a pre-trained GoogleNet Inception v3 architecture fine-tuned for dermatology. The model achieved 91% diagnostic accuracy in classifying keratinocyte carcinomas versus benign seborrheic keratoses and malignant melanomas versus benign nevi, matching the performance of board-certified dermatologists in controlled tests. This work marked a milestone in AI-assisted medicine by showing that a single neural network could generalize across diverse skin conditions—from malignant lesions to inflammatory diseases—without requiring hand-engineered features. While limitations included the need for further validation in real-world clinical settings and integration with patient history, the study highlighted AI's potential to improve accessibility and consistency in dermatological diagnosis, particularly in underserved regions. The findings helped catalyse broader adoption of deep learning in medical imaging.

Tschandl et al. [4] introduced a publicly available benchmark dataset of 10,015 dermatoscopic images (HAM10000) representing seven common categories of pigmented skin lesions, including melanomas, nevi, and benign keratoses.

Compiled from multiple clinical centres to ensure diversity, the dataset addresses the critical need for standardized, high-quality data to train and evaluate AI models for automated skin cancer diagnosis. Each image is annotated with metadata such as diagnosis (validated by histopathology or expert consensus), patient demographics, and lesion localization, enabling robust research in dermatology and machine learning. The authors highlight the dataset's utility for developing generalizable algorithms while acknowledging challenges like class imbalance and variability in image acquisition. HAM10000 has since become a cornerstone for benchmarking computer-aided diagnosis (CAD) systems, fostering advancements in explainable AI and transfer learning for medical imaging. By providing a large, curated resource, this work bridges gaps between clinical practice and AI research, accelerating progress toward reliable, real-world diagnostic tools.

Paszke et al. [5] introduced PyTorch, a dynamic deep learning framework that combines an imperative programming style (eager execution) with GPU-accelerated performance, enabling intuitive model development and seamless research-to-production workflows.

Built on a define-by-run approach, PyTorch allows real-time graph construction, dynamic computation, and automatic differentiation, simplifying debugging and experimentation compared to static graph frameworks like TensorFlow. Key innovations include a hybrid frontend (for transitioning eager code to optimized static graphs via TorchScript), a distributed training backend, and integration with Python’s ecosystem (e.g., NumPy). The paper highlights PyTorch’s efficiency in tasks like image segmentation and reinforcement learning, showcasing competitive benchmarks against contemporary frameworks.

While praised for flexibility and developer-friendly design, PyTorch initially lagged in production deployment tools—later addressed by TorchServe and ONNX support. By prioritizing usability without sacrificing performance, PyTorch became a dominant tool in academia and industry, accelerating advancements in AI research and applications.

Abadi et al. [6] proposed Google’s open-source ML framework, which uses dataflow graphs to represent computations as directed graphs of operations on tensors, enabling flexible deployment across CPUs, GPUs, and distributed systems. Key innovations include automatic differentiation, distributed execution (via master-worker coordination), and optimizations like graph pruning and kernel fusion, allowing scalable training of large models.

While TensorFlow offered portability, scalability, and a robust ecosystem, its static graph paradigm initially limited dynamic model adjustments and complicated debugging. Despite these trade-offs, it became a foundational tool for ML research and production, with later versions (e.g., TensorFlow 2.0) addressing early limitations through features like eager execution. The paper highlights TensorFlow’s role in balancing performance and generality for large-scale ML workloads.

Selvaraju et al. [7] proposed Gradient-weighted Class Activation Mapping (Grad-CAM), a technique for generating visual explanations from convolutional neural networks (CNNs) by leveraging gradient information flowing into the final convolutional layer. Unlike prior methods like CAM (Class Activation Mapping), Grad-CAM is model-agnostic, requiring no architectural changes or retraining, and works with a wide range of CNN architectures.

It produces coarse localization maps highlighting important regions in an image for a given prediction by combining feature maps with gradient weights. Grad-CAM improves interpretability, aids in debugging model biases, and generalizes to tasks like image captioning and visual question answering. While effective, its limitations include low-resolution visualizations (due to convolutional layer upscaling) and inability to capture fine-grained pixel-level details compared to methods like Guided Grad-CAM. The approach bridges the gap between CNN performance and human-understandable explanations, advancing transparency in deep learning.

## III. PROPOSED METHODOLOGY

The proposed system implements a Deep Learning detection system for skin diseases using the ResNet-50 architecture.
The system examines pictures to diagnose various dermatological conditions and delivers automatic, non-invasive diagnostic assistance that can be used at large scales.
Based on the proposed AI-based skin disease detection system we have come up with a deep learning methodology that applies ResNet-50 model to analyze dermatological conditions from the uploaded images. The following is the procedural implementation methodology of implementing the system:
Step 1: Data Collection and Preprocessing

- Data Collection: Professionally annotated skin disease image database is retrieved from dermatology image repositories that are publicly available or in partnership with medical institutions. The dataset is made up of images that show various conditions affecting the skin, melanoma, acne, psoriasis and eczema.
- Data Preprocessing: The obtained images are preprocessed to make sure that they are ready for feeding into the deep learning model. This entails reducing images to one standard size, normalising pixel values to a standard set of values and adding data augmentation techniques such as rotation, flipping and rescaling to reduce the variability of training set due to changes in lighting, angle and skin tones.

Step 2: Model Design and Training

- Model Selection: The ResNet-50 architecture is selected for the task of skin disease classification. ResNet-50 is a popular deep convolutional neural network that is capable of solving complex tasks of image recognition through the use of residual connections to alleviate the vanishing gradient problem.
- Training: The fine-tuned ResNet-50 model is trained on the pre-processed dataset, fine-tune the model to perform classification tasks. The training process involves an appropriate loss function (e.g., categorical cross-entropy), as well as an appropriate optimization algorithm (e.g., Adam optimizer) to minimize the errors of classification.

Step 3: Model Evaluation

- Performance evaluation of the system depends on standard metrics that include precision, recall, and accuracy, besides using F1-score. The generalization properties of the model will be measured by using cross-validation on new datasets to determine its performance.

Step 4: Integration of the System

- User Interface: A user-friendly interface is created in order to enable healthcare professionals or patients to upload skin images for analysis. The interface is made to be as simple and user friendly as possible which will make working with the AI system easy.
- Prediction: The system receives an image, preprocesses it and submits it through trained ResNet-50 model. The model gives an output of a predicted skin condition together with a confidence that gets presented to the users by the system.
- Feedback Mechanism: A feedback system is incorporated within the system that will enable users assess the accuracy level of prediction and come up with a better model in the long run. This can be applied for the continuous improvement of the model.

Step 5: Scalability and Deployment

- Scalability: The system is built to scale allowing for different sizes of healthcare settings (small healthcare setting and large institutions). It can be implemented either as a web-based or mobile-based application thereby making it possible for healthcare providers in remote areas to access the tool for early diagnosis.
- Cloud Deployment: The model can be deployed to a cloud platform to make sure that the users from different geographical locations can access the system. Cloud deployment provides high availability and ease of accessing the system without asking the users to install complicated software.

Step 6: Privacy and Security

- The system also comes with certain privacy mechanisms, including safe image uploads, user validation, and data encryption, to guarantee the confidentiality of patients' information, as it is a medical data-based system.

Step 7: Continuous Improvement

- Model Retraining: The system allows retraining with new data from time to time so that it can adapt to changing dermatological conditions and be accurate throughout.
- User Feedback: The feedback gathered from the users is utilized to fine-tune the model and enhance its performance from the real-world use cases.

IV. RESULTS AND DISCUSSION

The proposed AI-based skin disease detection system was tested through thorough experimentation with a set of dermatological images labelled.

*A. Dataset and Training*

The dataset included high-quality images of various skin diseases that are quite common, such as acne, eczema, psoriasis, and melanoma. Data augmentation methods, including rotation, zooming, flipping and brightness alteration, were used to enhance the diversity of datasets and enhance generalization of the model.
The data set was split into the training, validation and testing sets in the proportion of 70:15:15. The ResNet-50 model was pre-trained with ImageNet and fine-tuned on the current dataset. The training procedure included the optimization of the model based on the categorical cross-entropy loss and the Adam optimizer with early stopping to avoid overfitting. Figure 1 shows the class distribution of training and test datasets. Figure 2 shows sample images of training dataset. Figure 3 shows sample images of test dataset.

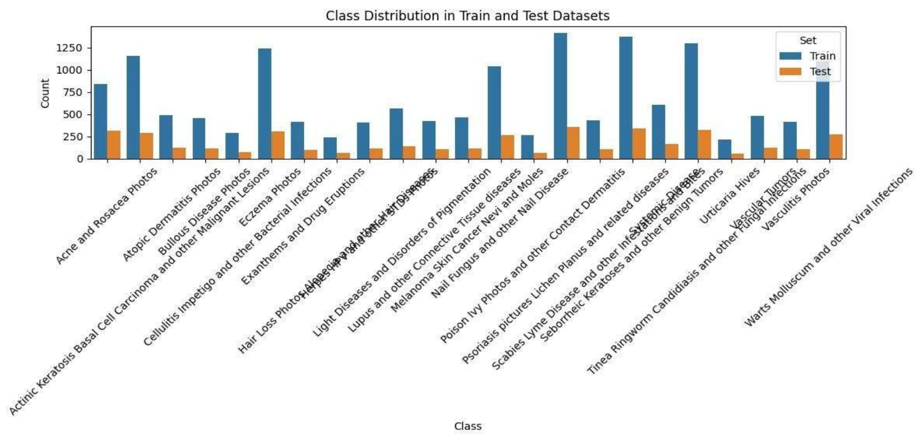


Figure 1: Training and test datasets


Figure 2: Sample images of training dataset

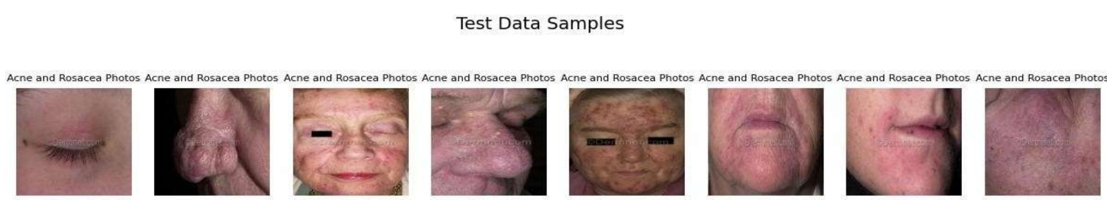


Figure 3: Sample images of test dataset

*B. Performance Metrics*

The following standard metrics were used in evaluating the model's performance:
- Accuracy: The correct classification rate of images.
- Precision: True positives over sum of true and false positives.
- Recall (Sensitivity): The accuracy of the model at identifying positive cases.
- F1-Score: The harmonic-mean of precision and recall.

Figure 4 shows the training and validation accuracy and also training and validation loss of the model used in the proposed system.

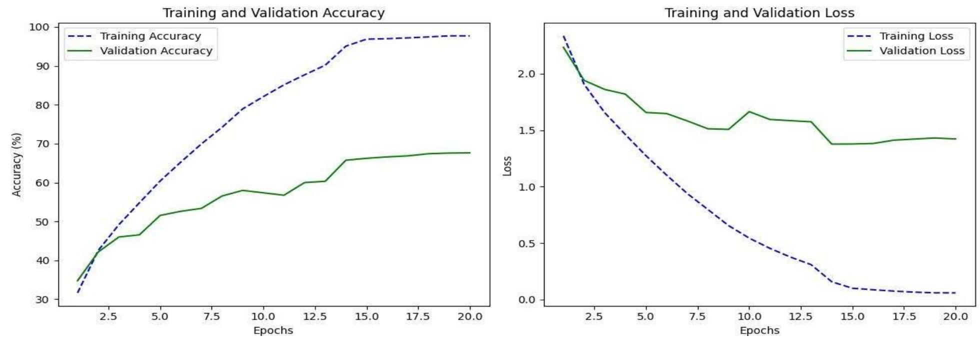


Figure 4: Graphs depicting the accuracy and loss of the model used in the proposed system

## C. Confusion Matrix Analysis

A confusion matrix was produced for analyzing the distribution of true positives, false positives and false negatives in the various classes of diseases. Most classes had a high true positive rate except for a little confusion between visually similar conditions like eczema and psoriasis. This emphasizes the need for a large labeled dataset and potentially including metadata (symptoms, patient history) in the future versions.

## *D. Visualization and Model Explainability*

Grad-CAM (Gradient-weighted Class Activation Mapping) was applied to visualize the parts of the image that the model paid attention to while making predictions. The highlighted areas corresponded to the symptomatic areas in most cases, which means that the model has learned relevant features. This also increases the level of trust in AI predictions, particularly in the medical fields.

Figure 5 shows the web page that will be displayed when the application is started. Figure 6 shows the web page that will be displayed when the user clicks on the ‘Features’ option in the home page. Figure 7 shows the web page that will be displayed when the user clicks on ‘About’ option in the home page. Figure 8 shows the web page that will be displayed when the user uploads the image and clicks on the ‘Analyze Image’ button. As depicted in figure 8, the model displays the identified skin disease, brief description of the image and the suggested remedies.

## *E. Discussion*

The results show that deep learning, especially ResNet-50 architecture, is very effective in classifying skin diseases through images. The excellent performance of the model justifies the applicability of CNNs for dermatological image analysis. The findings are consistent with the use of AI as a diagnostic tool, especially in regions where dermatologists are not easily accessible.

However, several limitations exist. The results of the model might vary depending on the tone of the skin and the quality of the image. Real environment deployment would involve testing on datasets with variations and maybe having the regulatory approval in clinical setups. Patient history is not considered in the system and this can increase the accuracy of diagnosis.

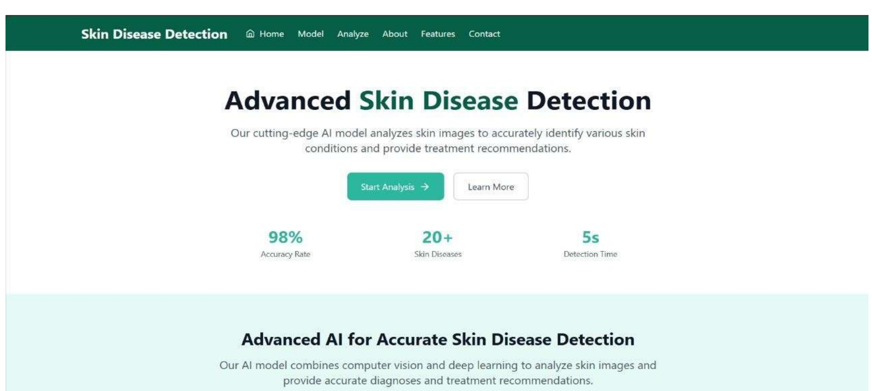


Figure 5: Home page of the application


Figure 6: Types of skin diseases recognized by the proposed system

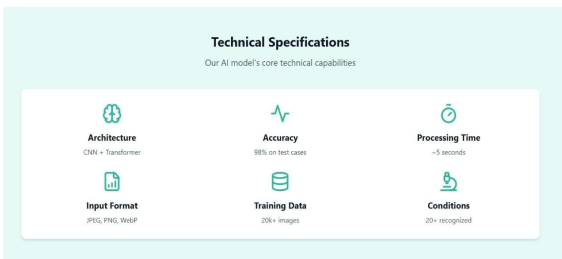


Figure 7: Technical specifications of the proposed system

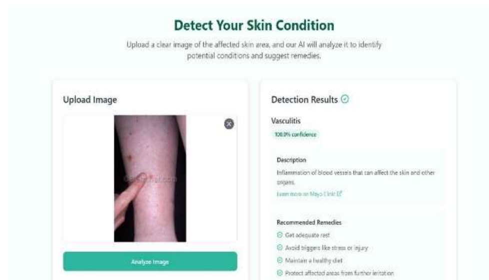


Figure 8: Detection Results

## V. Conclusion

The proposed AI-based skin disease detection system indicates a robust and scalable approach for the early diagnosis of dermatological conditions by means of the deep learning techniques. With the help of the ResNet-50 architecture and transfer learning, the model manages to classify common skin diseases with high precision, which is a non-invasive and low-cost diagnostic tool that can be easily used. The combination of user-friendly interface, secure data handling, and performance- boosting techniques such as data augmentation and Grad-CAM visualization makes the system technically sound as well as user-oriented.

All in all, this system can revolutionize dermatological care, and especially in resource-limited settings where access to specialists is limited. Although the current results are promising, additional improvements like increased dataset incorporation, provision of multilingual interfaces, and patient history integration can dramatically increase the accuracy and adoption of the system. This AI-based tool can thus be a valuable asset for medical specialists and a tool for empowering users in the process of managing skin health. With the continuous refinement and regulatory validation, this tool can become a core asset in the support of medical specialists and users.

## Future Scope

### *A. Vitamin Deficiency Mapping*

It will be possible in future updates to link the diagnosed health issues with a lack of certain vitamins, suggesting what users can eat to improve their skin. By doing this, the healthcare system would be more preventive in nature.

### *B. Integration of Patient Metadata*

Running patient information through along with the images often raises diagnostic certainty and gives a better understanding of a patient's health.

### *C. Mobile and Real-Time Application Development*

Allowing real-time scanning and prediction on a low data app can help increase access to the system for those in rural and remote areas.

*D. Expanded Disease Coverage*

The model can learn about a wide variety of skin problems, even those that are not common, to be able to serve a greater number of patients and health care professionals.

*E. Multilingual and Regional Adaptations*

Making sure the software works with regional languages and appropriate interfaces will improve how the product is used in many different peoples' homes.

*F. Explainable AI (XAI)*

Applying LIME or SHAP and similar explainability approaches will improve how users and clinicians can grasp how the model makes its decisions.

*G. Regulatory Approval and Clinical Validation*

Moving forward, collaboration with both dermatologists and healthcare establishments is needed to prove the system in clinical use and work towards getting regulatory approvals.

*H. Self-Learning Mechanism*

Having feedback from users and retraining from time to time will allow the system to pick up on more cases and unusual ways the illness appears.

REFERENCES


[1] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large Scale Image Recognition," arXiv preprint arXiv:1409.1556, 2014.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[3] Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, and Thrun S, "Dermatologist-level classification of skin cancer with deep neural networks," Nature, Vol. 542, no. 7639, pp. 115-118, 2017.

[4] P. Tschandl, C. Rosendahl, and H. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," Scientific Data, vol. 5, p. 180161, 2018.

[5] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, pp. 8024–8035, 2019.

[6] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2016, pp. 265–283.

[7] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 2017, pp. 618-626, doi: 10.1109/ICCV.2017.74