AI-BASED TOOL FOR PRELIMINARY DIAGNOSIS OF DERMATOLOGICAL MANIFESTATIONS

Project Reference No.: *48S_BE_2858*

College : Atria Institute of Technology, Bengaluru
Branch : Computer Science and Engineering
Guide(s) : Dr. Pradeep Kumar
Student(s): Mr. Shashank Anand
Mr. Kishore Jat
Ms. Shaima Farhath
Ms. Alisha Khalid

Keywords:

Skin Disease Diagnosis, Machine Learning (ML), Deep Learning (DL), Convolutional Neural Networks (CNN), EfficientNet-B0, HAM10000 Dataset, Skin Lesion Classification, Data Augmentation, Image Preprocessing, Transfer Learning, Medical Imaging, Dermatology, Real-time Analysis**.**

Introduction:

This project develops an AI-based diagnostic tool to classify dermatological diseases from dermoscopic images, addressing the global challenge of early skin cancer detection (e.g., melanoma). Leveraging the HAM10000 dataset (10,015 images across 7 lesion types) and EfficientNet-B0, the tool achieves 92.31% accuracy through optimized pre-processing, data augmentation, and transfer learning. Designed for clinicians and non-experts, it offers a scalable, real-time preliminary screening solution to bridge gaps in dermatological care, reduce dependency on specialists, and improve patient outcomes.

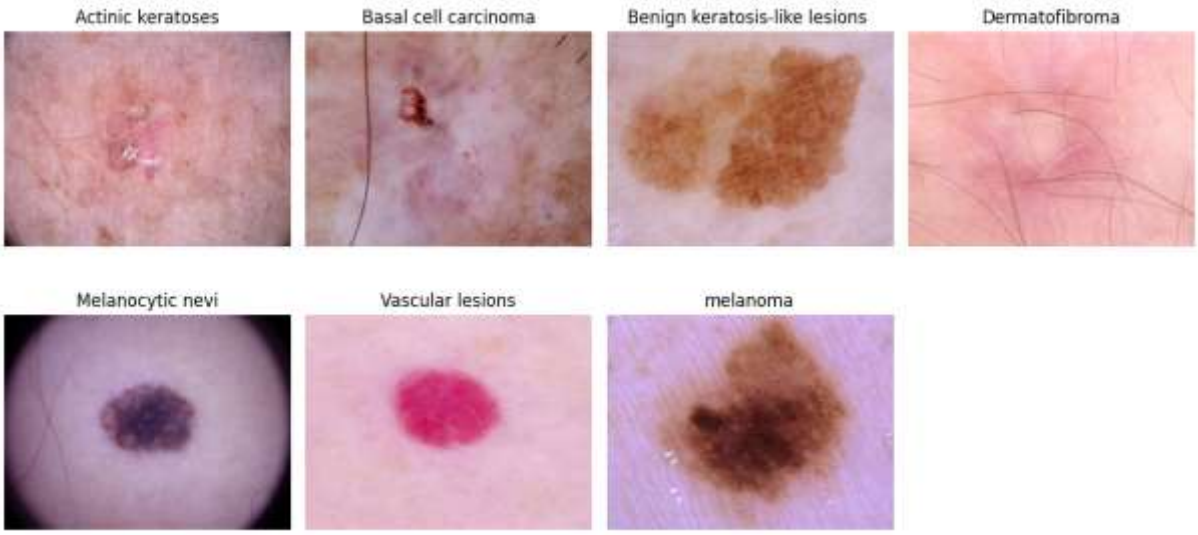

Figure 1: Skin Legion Categories

Objectives:

1. Develop AI tool for skin disease diagnosis using dermoscopic images.
2. Classify 7 lesion types accurately with EfficientNet-B0 model.
3. Improve dataset balance using augmentation and pre-processing.
4. Ensure user-friendly access via a simple web interface.
5. Achieve >90% accuracy for reliable preliminary screening.

Methodology:

1. Dataset and Preprocessing

The HAM10000 dataset, consisting of 10,015 dermoscopic images across 7 lesion classes, serves as the foundation. To address class imbalance, techniques like oversampling and data augmentation (including rotation, flipping, and zoom) are applied. Each image undergoes preprocessing steps such as resizing to 224x224 pixels, RGB channel normalization, and pixel value standardization to ensure consistency for model input.

2. Model Architecture and Training

The EfficientNet-B0 model, pretrained on ImageNet, forms the core of the classification system. Modifications include the addition of a Global Average Pooling (GAP) layer to minimize overfitting, followed by a dense layer with 512 units and ReLU activation. A

dropout rate of 30% is applied before the final softmax output layer for 7-class classification. Training employs the Adam optimizer with an initial learning rate of 0.001, dynamically reduced to 1e-6, and uses weighted categorical cross-entropy to prioritize underrepresented classes. The model trains over 20 epochs with a batch size of 32.

3. Performance Evaluation

The model's effectiveness is assessed using standard metrics such as accuracy, precision, recall, and F1-score. Additional insights are derived from AUC-ROC curves and a confusion matrix, which highlight the model's ability to discriminate between classes and identify areas for improvement, such as distinguishing melanoma from benign keratosis.

4. System Deployment

A user-friendly web interface, built with HTML, CSS, and JavaScript, allows seamless image uploads. The backend, powered by Django, processes images through the trained model via REST APIs and returns the predicted lesion class along with a confidence score, ensuring accessibility for both medical professionals and non-experts.

5. Validation and Future Enhancements

The dataset is split into 80% training and 20% testing sets using stratified sampling to maintain class distribution. Post-deployment analysis reveals specific challenges, such as confusion between similar-looking lesions, prompting plans to integrate explainability tools like Grad-CAM in future iterations.

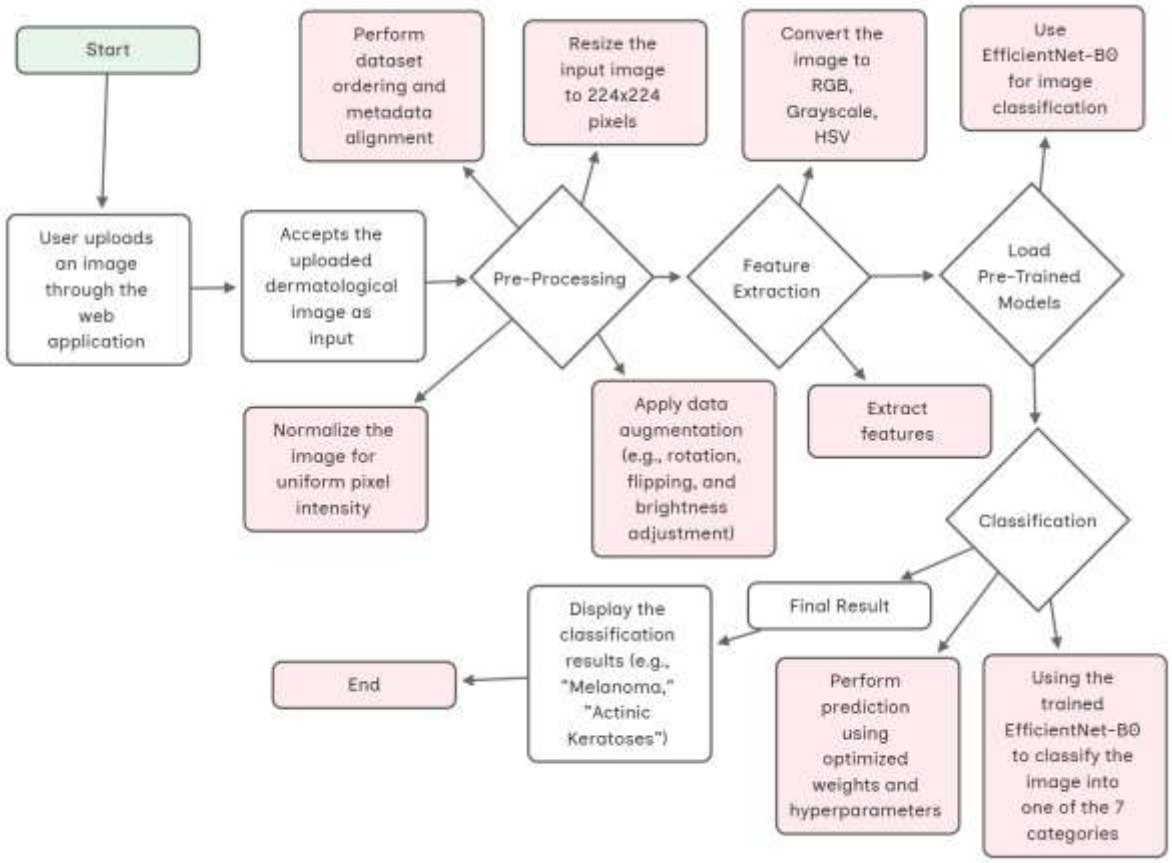

Result and Conclusion:

The AI-based diagnostic tool demonstrated robust performance in classifying dermatological lesions, achieving an overall accuracy of 92.31% on the test set. The model showed balanced precision and recall scores averaging 92%, indicating reliable classification across most lesion types. While common conditions like melanocytic nevi were detected with excellent performance (F1-score: 0.94), rarer classes such as dermatofibroma showed slightly lower recall (0.83) due to limited training samples. Analysis of the confusion matrix revealed that most misclassifications occurred between morphologically similar lesions, particularly between melanoma and benign keratosis, suggesting opportunities for future refinement. The training process exhibited strong stability, with validation accuracy stabilizing at 92.31% and minimal overfitting, as evidenced by the low training loss (0.02) and controlled validation loss (0.41). The model's discriminative capability was further confirmed by AUC-ROC scores ranging from 0.97 to 1.00 across all classes.

In conclusion, this project successfully developed an accessible and accurate AI tool for preliminary dermatological diagnosis using the EfficientNet-B0 architecture trained on the HAM10000 dataset. The tool's high accuracy, combined with its real-time analysis capability, makes it a practical solution for early disease detection, particularly in resource-constrained healthcare settings. The lightweight model design, enhanced by strategic data augmentation, effectively addressed class imbalance while maintaining computational efficiency. The user-friendly web interface further ensures that the technology is accessible to both medical professionals and non-specialists. However, limitations remain, including the need for more diverse datasets to improve generalizability and the potential integration of explainability tools like Grad-CAM to enhance clinical trust. Future work could explore advanced architectures such as EfficientNet-B7 or ensemble methods to further boost performance on minority classes. Ultimately, this tool exemplifies the transformative potential of AI in democratizing healthcare by enabling timely, cost-effective preliminary diagnostics and reducing dependence on specialized infrastructure.

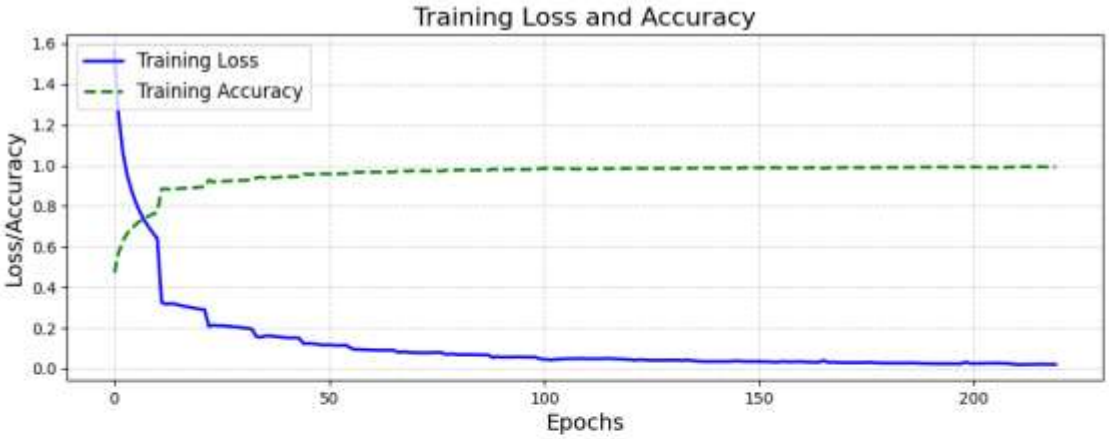
Training Loss and Accuracy
Training Loss
Training Accuracy
Loss/Accuracy
Epochs

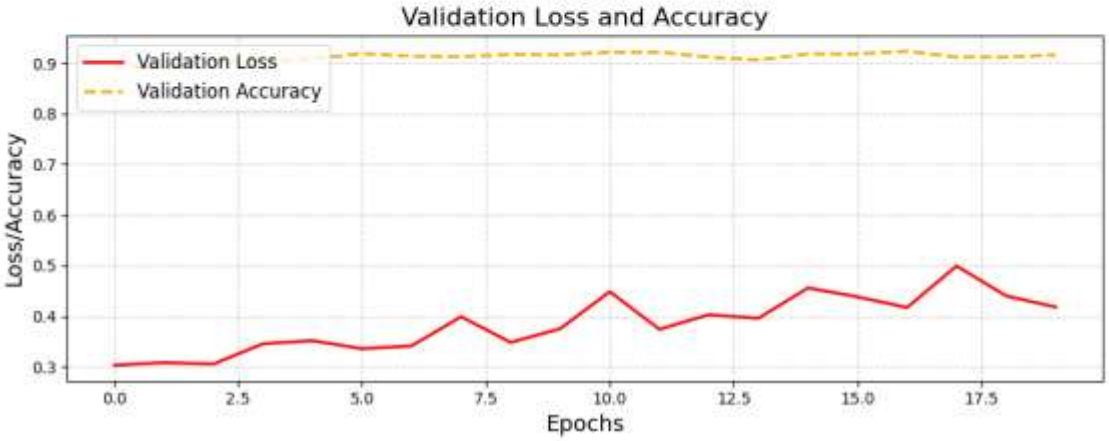
Validation Loss and Accuracy
Validation Loss
Validation Accuracy
Loss/Accuracy
Epochs

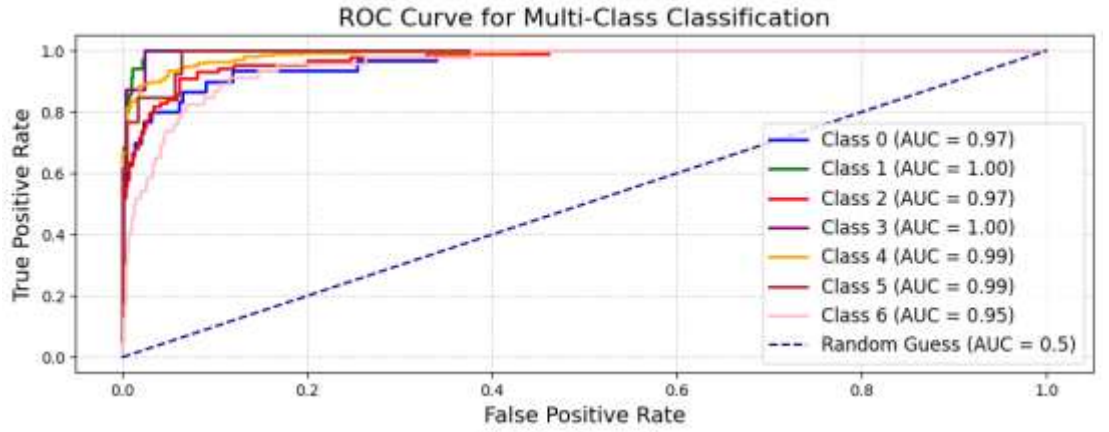
ROC Curve for Multi-Class Classification
Class 0 (AUC = 0.97)
Class 1 (AUC = 1.00)
Class 2 (AUC = 0.97)
Class 3 (AUC = 1.00)
Class 4 (AUC = 0.99)
Class 5 (AUC = 0.99)
Class 6 (AUC = 0.95)
Random Guess (AUC = 0.5)
True Positive Rate
False Positive Rate

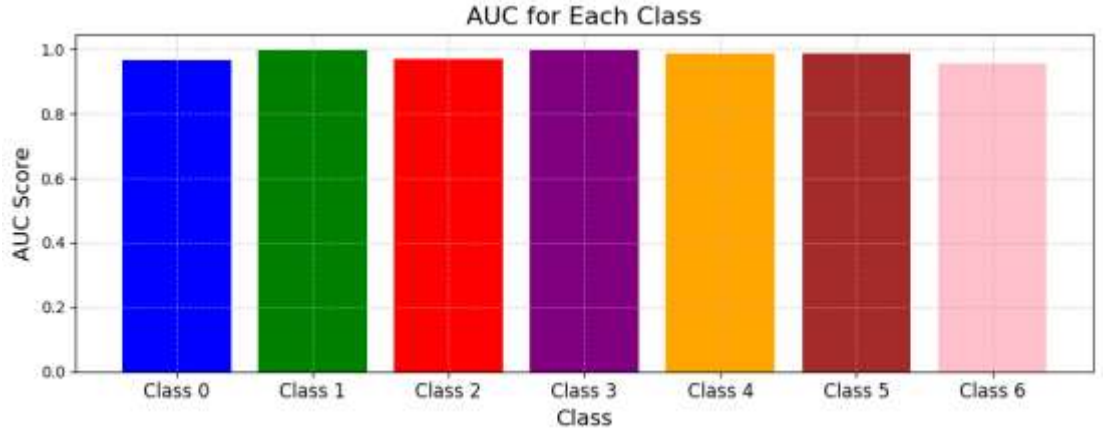
AUC for Each Class
AUC Score
Class 0
Class 1
Class 2
Class 3
Class 4
Class 5
Class 6
Class

Future Scope:

The future scope of this project includes:

1. Mobile app development to enable smartphone-based diagnosis for remote healthcare access.
2. Expansion of dataset diversity including more skin types and rare conditions to improve generalizability.
3. Implementation of vision transformers to enhance classification accuracy beyond current CNN-based approaches.
4. Integration of real-time video analysis for dynamic lesion monitoring and progression tracking.
5. Cloud-edge computing deployment to balance processing speed with data privacy requirements.
6. Development of severity grading and treatment recommendation systems for comprehensive clinical support.
7. Telemedicine platform integration to facilitate adoption in existing healthcare systems.
8. Federated learning implementation to enable continuous model improvement while protecting patient data.
9. Clinical validation studies to measure impact on early detection rates and treatment outcomes.
10. Framework adaptation for other medical imaging domains like radiology and pathology diagnostics.