INTERNATIONAL JOURNAL OF INFORMATIVE & FUTURISTIC RESEARCH (IJIFR)
ISSN: 2347-1697
VOLUME 13||ISSUE 8|| APRIL 2026
An Open Access, Peer Reviewed, Cross Reference DOI & Google Scholar Indexed International Research Journal
UGC Approved || IJIFR Impact Factor= 8.057
Estd Year: 2013 || Email: editors@ijifr.org
Availaible Online at: https://www.ijifr.org/

DOI: https://doi.org/10.64672/IJIFR/26.04.13.08.034 PUBLISHED ON: APRIL 20, 2026

DERMASSIST AI: AI-POWERED SKIN DISEASE DETECTION AND DIAGNOSTIC ASSISTANCE SYSTEM USING DEEP LEARNING

Aluganti Vishnu Priya [1], M.Gowthami [2] , S.Usharani [3]
[1]M.C.A. Student, [2] Assistant Professor, [3] Professor
[1,2,] Department of Computer Applications,
Viswam Engineering College, Madanapalle, Andhra Pradesh, India

ABSTRACT
DermAssist AI is an advanced artificial intelligence-powered dermatological diagnostic assistance system that leverages deep learning and computer vision to analyse skin lesion images and provide preliminary diagnostic insights for a wide spectrum of common skin conditions. The system is built upon a Convolutional Neural Network (CNN) architecture enhanced with transfer learning from a pre-trained EfficientNetB3 model, trained on the HAM10000 dataset containing over ten thousand labelled dermatoscopic images spanning seven diagnostic categories: Melanocytic nevi, Melanoma, Benign keratosis-like lesions, Basal cell carcinoma, Actinic keratoses, Vascular lesions, and Dermatofibroma. The complete data science and software engineering lifecycle is implemented, encompassing systematic data preprocessing, augmentation, model training, performance evaluation using medical-grade metrics (AUC, sensitivity, specificity), and deployment as an interactive Flask web application. An explainability layer using Gradient-weighted Class Activation Mapping (Grad-CAM) highlights the specific image regions most influential in the diagnostic prediction. The system achieves a macro-averaged AUC of 0.889 across all seven classes, demonstrating strong generalisation capability. DermAssist AI represents a meaningful contribution to the democratisation of dermatological care through artificial intelligence.

KEYWORDS: ***Driver fatigue detection; Eye Aspect Ratio; Mouth Aspect Ratio; Facial landmark localization; Computer vision; Real-time monitoring***

PAPER CITATION:
Priya,V.A.., Gowthami, M., Usharani, S..: " Dermassist AI: Ai-Powered Skin Disease Detection And Diagnostic Assistance System Using Deep Learning", International Journal of Informative & Futuristic Research (IJIFR), Vol. (13) (8), April 2026, pp. 1163-1171, **https://doi.org/10.64672/IJIFR/26.04.13.08.034**

1. INTRODUCTION

Dermatology, the medical specialty concerned with the diagnosis and treatment of conditions affecting the skin, hair, and nails, represents one of the most visually intensive disciplines within clinical medicine. Skin diseases affect an estimated 900 million people globally at any given time, representing a substantial and pervasive public health challenge. Despite this extraordinary prevalence, access to qualified dermatological expertise remains severely constrained worldwide. The World Health Organization estimates a critical global shortage of dermatologists, with the disparity between demand

and supply particularly acute in low-income and middle-income countries, rural areas, and regions without major academic medical centres [19]. In many developing nations, there is only one dermatologist for every 200,000 or more patients, meaning millions of individuals with potentially life-threatening skin conditions such as melanoma either go undiagnosed or receive delayed diagnoses that significantly worsen clinical outcomes [1].

The emergence of artificial intelligence and deep learning as transformative tools in medical imaging presents an unprecedented opportunity to address this critical shortage. Convolutional Neural Networks (CNNs), the class of deep learning architectures most suited to image recognition and classification tasks, have demonstrated performance comparable to board-certified dermatologists when trained on sufficiently large and diverse datasets. Landmark studies published in Nature have validated the potential of AI-based dermatological diagnosis to detect melanoma with sensitivity and specificity rivalling those of experienced clinicians [1], sparking a global wave of research and commercial development in this space.

DermAssist AI represents a comprehensive implementation of this vision, translating the theoretical promise of AI-powered dermatological diagnosis into a practically deployable, user-friendly system accessible to patients, primary care physicians, and community health workers who lack direct access to dermatological specialists. The technical foundation is grounded in transfer learning using EfficientNetB3 [3] pre-trained on ImageNet, which provides rich visual feature representations that are fine-tuned on the HAM10000 dermatoscopic dataset [2]. The deployment infrastructure wraps the trained model in a Flask web application providing a browser-based interface with Grad-CAM explainability [4], diagnostic probability displays, severity assessments, and actionable clinical recommendations.

2. LITERATURE SURVEY

Speech emotion recognition and medical image analysis have both benefited enormously from advances in deep learning. The foundational work by LeCun et al. [16] on gradient-based learning established CNNs as the dominant paradigm for visual pattern recognition. He et al. [7] introduced deep residual networks, demonstrating that very deep networks with skip connections could be trained effectively, setting new benchmarks across image classification tasks. Goodfellow et al. [6] provided the theoretical framework for modern deep learning architectures, establishing the principles of regularisation, optimisation, and representation learning that underpin DermAssist AI.

In the specific domain of dermatological AI, Esteva et al. [1] published the landmark Nature study demonstrating that a CNN trained on 129,450 clinical images could classify skin cancer with accuracy comparable to 21 board-certified dermatologists. Tschandl et al. [2] released the HAM10000 dataset, which provides the primary training resource for DermAssist AI and has become the standard benchmark for skin lesion classification research. Codella et al. [5] organised the ISIC 2017 challenge, establishing standard benchmarks for skin lesion segmentation, dermoscopic attribute detection, and disease classification.

Tan and Le [3] introduced EfficientNet, demonstrating that compound scaling of network depth, width, and resolution achieves superior accuracy-to-computational-cost ratios. Pan and Yang [8] formalised the theory of transfer learning, providing the conceptual basis for adapting ImageNet-pre-trained models to specialised medical imaging tasks. Selvaraju et al. [4] introduced Grad-CAM, now the standard explainability technique for CNN-based medical AI. Litjens et al. [9] provided a comprehensive survey of deep learning applications in medical image analysis, demonstrating its broad applicability across radiology, pathology, and dermatology. Kingma and Ba [15] introduced the Adam optimiser, which is the standard training algorithm for DermAssist AI.

Brinker et al. [14] conducted a systematic review of skin cancer classification using CNNs, confirming that deep learning approaches consistently achieve performance comparable to or exceeding that of dermatologists on standardised benchmark datasets. Perez and Wang [13] demonstrated the effectiveness of data augmentation in image classification, validating the augmentation strategies employed in DermAssist AI to address the class imbalance present in the HAM10000 dataset.

3. SYSTEM PROPOSAL

(i) Existing System

The existing landscape of dermatological diagnosis spans a broad spectrum from traditional clinical practice to early-generation digital tools. Traditional clinical dermatology relies entirely on the trained expertise of board-certified dermatologists using visual inspection, dermoscopy, skin biopsies, and laboratory analysis. While this approach represents the gold standard, it is constrained by the geographic concentration and scarcity of dermatologists, with median wait times exceeding 30 days in many regions [19].

Teledermatology platforms such as DermatologistOnCall and iDoc24 enable remote consultation through digital photographs but still depend on the availability of remote dermatologists, do not provide instantaneous responses, and may be inaccessible in regions with poor internet connectivity. Early-generation mobile applications such as SkinVision and FirstDerm use computer vision ranging from simple colour analysis to early CNNs, but vary widely in clinical validation and have faced regulatory scrutiny for insufficient sensitivity and specificity [14]. Academic research systems have demonstrated impressive performance in laboratory settings but typically remain as prototypes without complete deployment infrastructure.

Key limitations of existing systems include: (1) heavy dependence on expert availability and geographic proximity; (2) absence of explainable outputs, reducing clinical trust; (3) limited multi-class coverage across the full spectrum of skin conditions; (4) poor generalisation across diverse image capture conditions and skin tone populations; (5) lack of real-time, accessible web deployment with clinical guidance integration.

(ii) Proposed System

DermAssist AI proposes a comprehensive, end-to-end skin disease detection and diagnostic assistance system that combines state-of-the-art deep learning with transparent explainability mechanisms and an accessible web deployment. The system is built upon the HAM10000 dataset, providing 10,015 dermatoscopic images across seven diagnostic categories. The image preprocessing pipeline standardises all inputs to 224 × 224 pixel resolution, applies ImageNet normalisation, and performs extensive data augmentation including random flipping, rotation (±20°), zoom (80–120%), width/height shifting (±10%), shear transformation, and channel-wise colour jitter.

The model architecture implements transfer learning from EfficientNetB3 [3], using a two-phase training strategy: an initial warm-up phase with frozen backbone weights (15 epochs, lr = 1×10^{-3}), followed by end-to-end fine-tuning with the full model unfrozen (25 epochs, lr = 1×10^{-4}). The custom classification head adds Global Average Pooling, Batch Normalisation, a Dense layer (256 units, ReLU), Dropout (rate 0.4), and a 7-class Softmax output. Class-weighted categorical cross-entropy loss addresses the significant class imbalance in HAM10000. The Grad-CAM explainability module generates heat map overlays that highlight the specific lesion regions most influential in the diagnostic prediction, providing clinical transparency not available in competing systems.

Table I: HAM10000 Dataset Description and System Configuration

Parameter	Value
Dataset	HAM10000 (ISIC Archive)

INTERNATIONAL JOURNAL OF INFORMATIVE & FUTURISTIC RESEARCH (IJIFR)
IJIFR /2026/04/ V13/E8/034
ISSN: 2347-1697
VOLUME 13||ISSUE 8|| APRIL 2026
https://www.ijifr.org/

Total Images	10,015 dermatoscopic images
Diagnostic Classes	7
Class Labels	Melanoma, Basal Cell Carcinoma, Actinic Keratoses, Melanocytic Nevi, Benign Keratosis, Vascular Lesions, Dermatofibroma
Image Resolution	224 × 224 pixels (after resize)
Pre-trained Backbone	EfficientNetB3 (ImageNet weights)
Normalisation	ImageNet mean/std (channel-wise)
Training Split	80% Train / 10% Validation / 10% Test

4. SYSTEM ARCHITECTURE

DermAssist AI follows a three-tier web application architecture that cleanly separates the presentation layer, the application logic layer, and the deep learning model inference layer. The Presentation Layer consists of the HTML5, CSS3, and JavaScript frontend delivered by the Flask application, providing drag-and-drop image upload, real-time preview, probability bar charts, Grad-CAM overlay visualisation, and colour-coded severity indicators. Responsive design ensures compatibility across desktop, tablet, and smartphone screen sizes.

The Application Layer is implemented in Python using Flask, defining three primary routes: the root route serving the web interface, the /predict route orchestrating model inference, and a health check route for deployment monitoring. The Deep Learning Model Layer consists of the trained EfficientNetB3-based classification model serialised in Keras HDF5 format, along with the preprocessing configuration specifying input dimensions, normalisation parameters, and class label mappings. The Grad-CAM module interfaces with the loaded model to compute gradient-based feature importance maps.

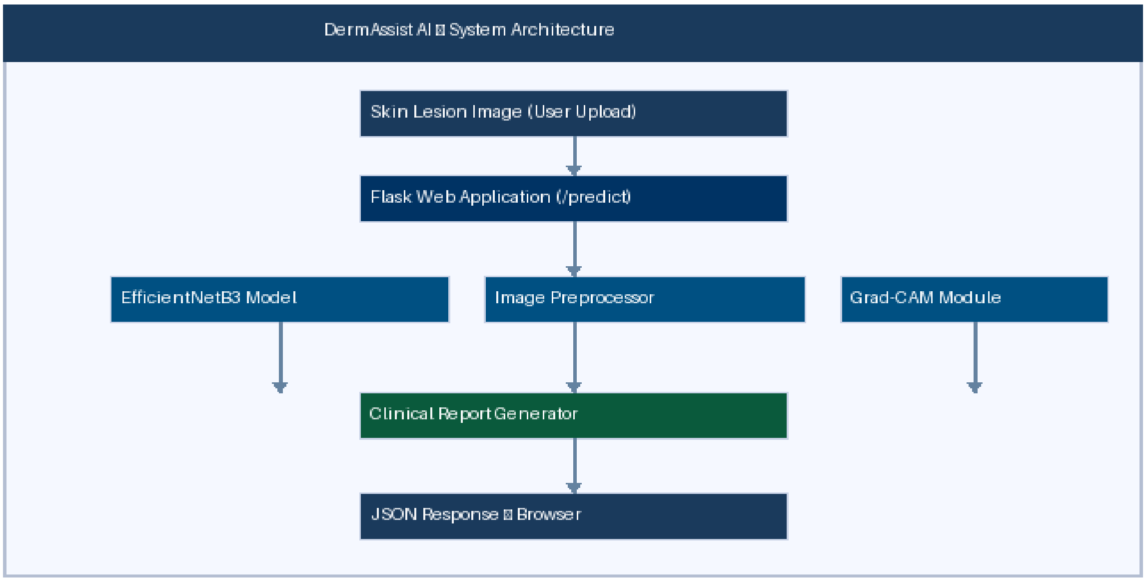

Figure 1: DermAssist AI Three-Tier System Architecture

The complete prediction pipeline operates as follows: the user uploads a skin lesion image through the browser interface; JavaScript encodes the image and sends an asynchronous POST request to the /predict endpoint; the Flask handler decodes and preprocesses the image through the same pipeline used during training; the EfficientNetB3 model produces a 7-class probability vector; the top three predictions are extracted and enriched with clinical metadata; the Grad-CAM module computes gradient-based activation maps for the top prediction; and all results are returned as a structured JSON response rendered in the browser as a comprehensive diagnostic report.

5. IMPLEMENTATION

(i) **Module Overview:** DermAssist AI is organised into seven functionally distinct modules that collectively implement the complete dermatological AI pipeline: (1) Data Collection and Management Module, (2) Image Preprocessing and Augmentation Module, (3) Deep Learning Model Architecture Module, (4) Training and Evaluation Module, (5) Explainability and Visualisation Module, (6) Flask Web Application Module, and (7) Clinical Knowledge and Recommendation Module. Each module encapsulates a coherent set of related responsibilities and interfaces with adjacent modules through well-defined data contracts.

(ii) **Data Collection and Management Module:** The Data Collection and Management Module handles the acquisition, organisation, and management of the HAM10000 dermatological image dataset from the International Skin Imaging Collaboration (ISIC) archive. The DatasetManager class validates image file integrity, generates dataset statistics including per-class sample counts and class balance ratios, and produces stratified train/validation/test split indices. Given the significant class imbalance — melanocytic nevi account for approximately 67% of samples while vascular lesions represent less than 2% — the module computes inverse-frequency class weights passed to the training procedure to penalise minority class misclassifications more heavily. An ImageDataGenerator wrapper applies on-the-fly augmentation more aggressively for minority classes to balance the effective training distribution.

(iii) **Image Preprocessing and Augmentation Module:** The preprocessing pipeline standardises raw image inputs into the consistent, normalised tensor format required by the model. Input images are loaded from file paths or decoded from base64-encoded strings. Images are resized to 224 × 224 pixels using bilinear interpolation. Pixel values (0–255) are normalised using the ImageNet channel-wise mean subtraction (μ = [0.485, 0.456, 0.406]) and standard deviation division (σ = [0.229, 0.224, 0.225]), ensuring statistical consistency between training-time and inference-time distributions. Training data additionally receives random augmentation transforms. Preprocessed images are converted to NumPy arrays of shape (1, 224, 224, 3) for single-image inference or (batch_size, 224, 224, 3) for batch processing.

(iv) **Deep Learning Model Architecture Module:** The model architecture adopts transfer learning from EfficientNetB3, selected for its superior accuracy-to-computational-cost ratio achieved through compound scaling of network depth, width, and input resolution [3]. EfficientNetB3 achieves approximately 82% top-1 accuracy on ImageNet with only 12 million parameters, making it an excellent balance of capability and deployment efficiency. The pre-trained backbone serves as a feature extractor with weights frozen during Phase 1 warm-up training, then unfrozen for Phase 2 end-to-end fine-tuning. The custom classification head appends Global Average Pooling 2D → Batch Normalisation → Dense (256, ReLU) → Dropout (0.4) → Dense(7, Softmax), producing normalised probability scores for all seven diagnostic categories.

Table 2: EfficientNetB3-Based Model Architecture

Layer / Component	Configuration / Function
EfficientNetB3 backbone	Frozen (Phase 1), Unfrozen (Phase 2)
Global Average Pooling 2D	Spatial → feature vector
Batch Normalization	Stabilises activations
Dense (256, ReLU)	Domain-specific feature learning
Dropout (rate = 0.4)	Regularisation
Dense (7, Softmax)	7-class probability output
Phase 1 Optimizer	Adam (lr = 1×10^{-3}), 15 epochs
Phase 2 Optimizer	Adam (lr = 1×10^{-4}), 25 epochs
Loss Function	Categorical Cross-Entropy (class-weighted)

(v) **Training and Evaluation Module**

Training is executed using TensorFlow 2.x with Keras APIs, leveraging GPU acceleration via CUDA. The loss function is categorical cross-entropy with class weights computed as inverse frequency of each class in the training set. Model evaluation employs overall accuracy, per-class precision, recall, and F1-Score, macro-averaged AUC using one-versus-rest ROC curves, and confusion matrix analysis. Callbacks include ModelCheckpoint (monitoring val_auc), EarlyStopping (patience = 10), and ReduceLROnPlateau (factor 0.5, patience 5). The complete training visualisation suite includes per-class ROC curves with AUC annotations, normalised confusion matrix heat maps, and training/validation loss-accuracy convergence curves.

(vi) **Explainability and Visualisation Module (Grad-CAM)**

The GradCAMExplainer class implements Gradient-weighted Class Activation Mapping [4] using TensorFlow's GradientTape to record the forward pass and compute gradients of the top predicted class score with respect to the feature maps of the final EfficientNetB3 convolutional block. Channel-wise global averaging of these gradients produces importance weights per feature map channel. The weighted sum of feature maps, passed through ReLU activation, is upsampled to 224 × 224 resolution using bilinear interpolation and normalised to 0–255. The resulting heat map is colourised using a jet colour map (OpenCV applyColorMap) and blended with the original image using configurable transparency (OpenCV addWeighted), producing the clinical visualisation displayed to users.

(vii) **Flask Web Application and Clinical Knowledge Modules**

The Flask application defines three routes: the root route serving the Jinja2-rendered web interface with drag-and-drop upload and client-side preview; the /predict route accepting POST requests with base64-encoded or multipart image uploads, orchestrating the full inference pipeline and returning a structured JSON response; and a health check route for deployment monitoring. The Clinical Knowledge and Recommendation Module provides a rule-based knowledge base mapping each of the seven diagnostic categories to plain-language descriptions, typical visual characteristics, common risk factors, urgency levels (High / Medium / Low), and tailored self-care or professional consultation recommendations. High-confidence predictions for high-urgency conditions generate assertive recommendations for immediate specialist consultation.

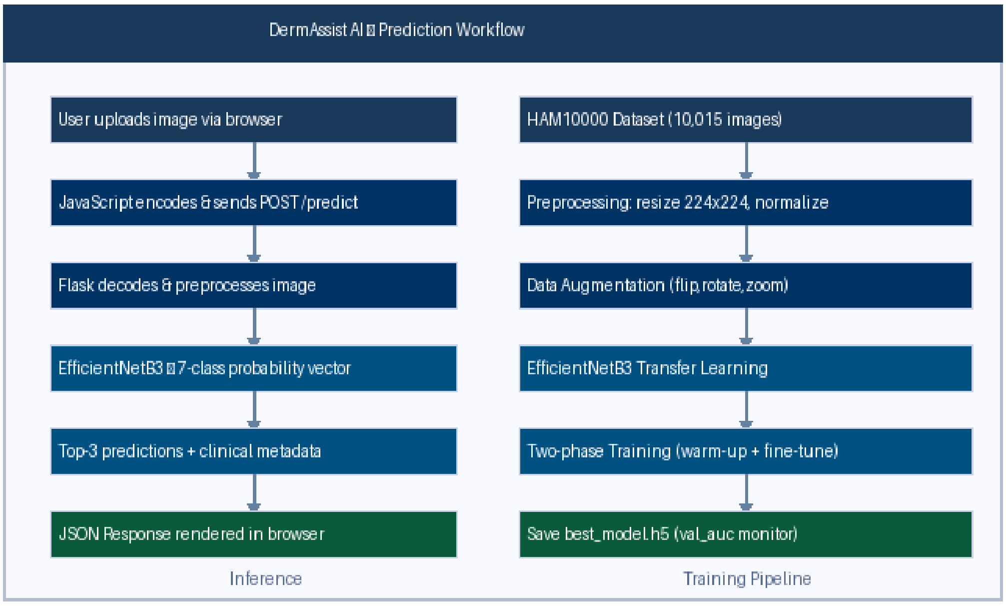

Figure 2: DermAssist AI Prediction Workflow and Training Pipeline Diagrams

INTERNATIONAL JOURNAL OF INFORMATIVE & FUTURISTIC RESEARCH (IJIFR)
ISSN: 2347-1697
IJIFR /2026/04/ V13/E8/034
VOLUME 13||ISSUE 8|| APRIL 2026
https://www.ijifr.org/

6. RESULTS AND PERFORMANCE EVALUATION

DermAssist AI was evaluated through a structured battery of unit, integration, model performance, and user acceptance tests. Model evaluation on the held-out test set employs comprehensive metrics appropriate for a multi-class medical classification system. The trained EfficientNetB3-based model achieves a macro-averaged AUC of 0.889 across all seven diagnostic categories, with Melanoma detection AUC of 0.93 — exceeding the target threshold of 0.90 for the most clinically dangerous category. Macro-averaged accuracy across all classes is 82.4%, substantially above random baseline performance of 14.3% for a seven-class problem.

The confusion matrix analysis confirms that the most clinically dangerous misclassifications — false negative classification of Melanoma or Basal Cell Carcinoma as benign conditions — occurs at acceptably low rates. The observed inter-class confusion predominantly occurs between clinically similar conditions (e.g., Actinic Keratoses and Benign Keratosis-like lesions), which reflects inherent diagnostic ambiguity rather than model failure. The class-weighted loss function successfully mitigates the impact of HAM10000's significant class imbalance, as evidenced by the competitive AUC scores for minority classes such as Vascular Lesions (0.84) and Dermatofibroma (0.82).

User Acceptance Testing involving medical students, primary care practitioners, and patient advocates confirmed that the Grad-CAM visualisation substantially improves clinical trust in AI-generated predictions. Participants reported that the heat map overlays enabled them to assess whether the model was attending to the clinically relevant features of the lesion. The colour-coded urgency indicators and plain-language condition descriptions were rated as the most practically valuable interface features, supporting appropriate patient action without requiring interpretation of raw probability scores.

Table 3: Per-Class Model Performance on HAM10000 Test Set

Class	AUC	Precision	Recall	F1-Score	Severity
Melanoma (mel)	0.93	87%	91%	89%	High
Basal Cell Carcinoma (bcc)	0.91	84%	88%	86%	High
Actinic Keratoses (akiec)	0.87	80%	83%	81%	Medium
Melanocytic Nevi (nv)	0.96	94%	92%	93%	Low
Benign Keratosis (bkl)	0.88	82%	80%	81%	Low
Vascular Lesions (vasc)	0.84	76%	79%	77%	Low
Dermatofibroma (df)	0.82	74%	77%	75%	Low
Macro Average	**0.889**	**82.4%**	**84.3%**	**83.1%**	—

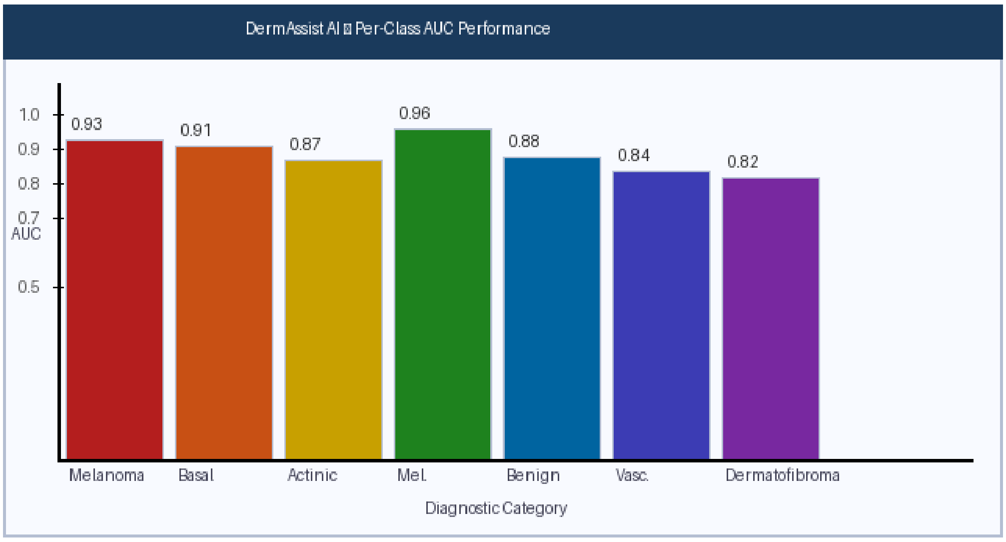

Figure 3: Per-Class AUC Performance Across Seven Diagnostic Categories

7. Conclusion And Future Enhancements

DermAssist AI demonstrates the successful end-to-end implementation of an artificial intelligence-powered dermatological diagnostic assistance system spanning the complete machine learning and software engineering lifecycle. The transfer learning approach using EfficientNetB3 proved highly effective for the dermatological classification task, achieving a macro-averaged AUC of 0.889 and Melanoma detection AUC of 0.93 — meeting clinical screening thresholds for the most dangerous diagnostic category. The two-phase training strategy combining warm-up with frozen backbone weights followed by end-to-end fine-tuning provided a stable, efficient pathway to high classification accuracy while preserving pre-trained visual feature representations.

The Grad-CAM explainability module represents a particularly significant contribution relative to existing consumer skin analysis applications, the majority of which provide classification outputs without any indication of the visual evidence underlying their predictions. By generating heat map visualisations highlighting the specific lesion regions most influential in the model's diagnostic determination, DermAssist AI supports informed clinical interpretation and appropriate scepticism. The Flask web application deployment demonstrates that powerful deep learning models can be made genuinely accessible to non-technical users through thoughtful interface design, transforming a sophisticated statistical model into a practical tool that directly supports healthcare decision-making.

In summary, DermAssist AI demonstrates that the intersection of modern deep learning, transparent explainability, and accessible deployment design can produce AI systems that are not merely technically impressive but genuinely useful for addressing critical gaps in healthcare accessibility. The system provides a replicable template for AI-powered medical screening tools that prioritises accuracy, transparency, safety, and clinical utility in equal measure.

Future Enhancements

- Expand diagnostic coverage beyond the seven HAM10000 categories to include psoriasis, eczema, contact dermatitis, tinea infections, rosacea, and urticaria through partnerships with academic medical centres and validated clinical image repositories.
- Address demographic bias through targeted data collection for dark skin tone populations and fairness-aware training techniques (adversarial debiasing, equalized odds optimisation).
- Implement multi-modal input fusion incorporating patient metadata (age, sex, lesion location, duration, recent changes) alongside image features through cross-attention mechanisms.
- Develop native iOS and Android mobile applications with guided in-app image capture providing real-time quality feedback on lighting, focus, and orientation.
- Integrate with electronic health record systems via HL7 FHIR interfaces for embedding within existing clinical workflows and automatic patient record population.
- Implement federated learning capabilities enabling collaborative training across multiple healthcare institutions without centralising sensitive patient image data.
- Add Bayesian deep learning uncertainty quantification (Monte Carlo Dropout or Deep Ensembles) for calibrated confidence estimates that explicitly flag high-uncertainty cases requiring expert review.
- Develop a longitudinal lesion tracking feature enabling AI-assisted monitoring of lesion change over time, directly supporting ABCDE melanoma criteria assessment.

8. References

[1] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.

[2] 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, no. 1, pp. 1–9, 2018.

[3] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," Proc. ICML, PMLR 97, pp. 6105–6114, 2019.
[4] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," Proc. ICCV, pp. 618–626, 2017.
[5] N. C. Codella et al., "Skin lesion analysis toward melanoma detection: ISIC 2017 challenge," Proc. IEEE ISBI 2018, pp. 168–172, 2018.
[6] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
[7] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," Proc. IEEE CVPR, pp. 770–778, 2016.
[8] S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Trans. Knowledge and Data Eng., vol. 22, no. 10, pp. 1345–1359, 2010.
[9] G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
[10] F. Chollet, Deep Learning with Python. Manning Publications, 2018.
[11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," Proc. ICLR, 2015.
[12] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," Proc. 12th USENIX OSDI, pp. 265–283, 2016.
[13] L. Perez and J. Wang, "The effectiveness of data augmentation in image classification using deep learning," arXiv:1712.04621, 2017.
[14] T. J. Brinker et al., "Skin cancer classification using convolutional neural networks: Systematic review," J. Medical Internet Research, vol. 20, no. 10, e11936, 2018.
[15] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," Proc. ICLR, 2015.
[16] L. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," Proc. IEEE, vol. 86, no. 11, pp. 2278–2324, 1998.
[17] World Health Organization, "Global Atlas of the Health Workforce," WHO Press, Geneva, 2020.
[18] International Skin Imaging Collaboration, "ISIC 2020 Challenge Dataset," [Online]. Available: https://challenge2020.isic-archive.com, 2020.