



XAI-DermNet: A Dual-Modality Deep Learning and Explainable AI Fusion Framework for Transparent Skin Lesion Diagnosis from Dermoscopic Images

Sukkirtha K.^{1,*} Anbuchelian S.¹ John Prakash A.¹

¹ Ramanujan Computing Centre, Anna University, Chennai, Tamil Nadu, India

Emails: sukkirtha05au@gmail.com · anbuchelian@annauniv.edu · johnprakash@annauniv.edu

Received: January 17, 2026 Revised: February 15, 2026 Accepted: March 19, 2026 ★ Corresponding author

ABSTRACT

The advancement of trustworthy diagnostic tools in dermatological automation is hindered by the limited transparency of current deep learning systems, which function as opaque models and impede clinical acceptance. This research presents a novel intelligent framework for skin lesion analysis that unites deep learning methodologies with explainable artificial intelligence (XAI) principles to address this interpretability deficit. The proposed approach utilizes a transfer-learned ResNet50 architecture for robust image classification, coupled with Local Interpretable Model-agnostic Explanations (LIME) to furnish clear, visual justifications for the model's outputs. Performance assessment on the HAM10000 benchmark yielded a classification accuracy of 94.3%, with a validation accuracy of 91.8%. Concurrently, the LIME framework effectively identified and visualized diagnostically critical features in the lesion images, thereby elucidating the model's reasoning process for medical practitioners. These findings confirm that augmenting high-performance deep learning with post-hoc explanatory techniques yields a credible and understandable diagnostic instrument, thereby promoting clinician trust and facilitating data-informed medical judgments. Subsequent developments will prioritize scalable cloud implementation, interoperability with healthcare information systems, extension to underrepresented lesion categories, and rigorous evaluation in diverse clinical environments.

Keywords: Skin lesion diagnosis ▪ Deep learning ▪ ResNet50 ▪ Explainable artificial intelligence ▪ LIME

1. INTRODUCTION

Skin diseases are among the most prevalent health conditions worldwide, affecting individuals of all ages and skin types. Early and accurate diagnosis plays a crucial role in providing effective treatment and preventing the progression of severe dermatological conditions such as melanoma and other malignant skin lesions. Traditionally, dermatologists rely on visual inspection and dermoscopic examination to identify skin abnormalities. However, this diagnostic approach can sometimes be subjective and depends heavily on the expertise and experience of the clinician. In addition, manual examina-

tion can be time-consuming when large numbers of patients require screening. As a result, there is increasing interest in developing automated diagnostic systems that can assist dermatologists in detecting and classifying skin diseases more efficiently and consistently. Advances in medical imaging and computational techniques have created opportunities to improve diagnostic accuracy through intelligent computer-aided systems. [1]

2. RELATED WORKS

Nguyen and Vo (2025) proposed lightweight CNN architectures tailored for early melanoma detection, emphasizing computational efficiency without compromising diagnostic performance. Their approach utilized compact convolutional blocks and parameter-reduction strategies to enable deployment in resource-constrained clinical settings. Experimental results demonstrated that the proposed lightweight models achieved competitive classification accuracy while significantly reducing model complexity and inference time compared to conventional deep CNN architectures. The study highlights the feasibility of efficient deep learning frameworks for real-time and portable melanoma screening applications.

Ahmed et al. (2024) proposed an enhanced residual network-based framework for multi-class skin lesion classification by refining pretrained ResNet architectures using transfer learning. Their approach focused on improving discriminative feature extraction while addressing class imbalance issues in dermoscopic datasets. Experimental results indicated that the optimized ResNet variants outperformed conventional CNN baselines, achieving classification accuracy of approximately 90%. The findings demonstrate that deep residual learning improves model robustness, generalization capability, and diagnostic reliability in automated skin lesion analysis.

Altini et al. (2024) proposed a comprehensive comparative analysis between deep learning models and healthcare professionals for disease detection using medical imaging data. The study reported that deep learning systems achieved a pooled sensitivity of approximately 87.0% and specificity of 92.5%, closely comparable to clinicians, who demonstrated 86.4% sensitivity and 90.5% specificity. These results indicate that advanced deep learning approaches can reach near expert-level diagnostic performance, highlighting their potential as effective clinical decision-support systems.

Singh et al. (2024) proposed a hybrid framework that integrates ResNet and Vision Transformer (ViT) architectures to enhance skin lesion classification performance. The model combines convolutional feature extraction with transformer-based global attention mechanisms to capture both local texture patterns and long-range contextual dependencies in dermoscopic images. Experimental evaluation demonstrated improved accuracy and robustness compared to standalone CNN or transformer models. The study indicates that CNN-ViT fusion strategies can effectively enhance discriminative representation and overall diagnostic performance in automated skin lesion analysis.

Patel et al. (2023) proposed an interpretable deep learning framework for skin cancer detection by integrating LIME and SHAP with convolutional neural network models to enhance model transparency. Their system achieved a classification accuracy of approximately 92% on benchmark dermoscopic datasets while providing localized feature importance and visual explanations for predictions. The results demonstrate that incorporating explainable AI techniques can maintain high diagnostic performance while improving clinical trust and interpretability in automated skin lesion analysis.

Sun et al. (2023) proposed a robust deep learning framework for skin lesion classification designed to handle noisy and mis-

labeled training data. The study incorporated noise-tolerant learning strategies and adaptive loss functions to mitigate the impact of incorrect annotations in dermoscopic datasets. Experimental evaluation demonstrated that the proposed method maintained stable performance under varying noise levels, achieving classification accuracy of approximately 91%, outperforming conventional CNN models trained without noise-robust mechanisms. The findings emphasize the importance of resilient training strategies to improve reliability and generalization in real-world clinical datasets.

Although existing studies demonstrate strong performance in skin lesion classification, several limitations persist. Many approaches focus primarily on accuracy while neglecting interpretability, which limits their clinical adoption. In addition, issues such as class imbalance, lack of comprehensive evaluation metrics, and insufficient generalization across datasets remain largely unaddressed. Furthermore, most methods do not incorporate explainable AI techniques to justify model predictions, reducing trust among medical practitioners. In contrast, the proposed XAI-DermNet framework integrates deep learning with explainable AI while addressing class imbalance and providing detailed performance evaluation, thereby improving both reliability and transparency.

3. METHODOLOGY

3.1 Data Collection

Dermoscopic images used in this study were obtained from the HAM10000 Dataset, a widely used benchmark dataset for skin lesion analysis. It contains over 10,000 labeled images across seven skin lesion categories, including melanoma, melanocytic nevi, basal cell carcinoma, actinic keratoses, benign keratosis-like lesions, dermatofibroma, and vascular lesions. These images were used to train and evaluate the proposed XAI-DermNet model.

The HAM10000 dataset provides a diverse collection of dermoscopic images collected from different sources, ensuring variability in lesion types, skin tones, and imaging conditions. This diversity is important for training robust machine learning models capable of generalizing across different dermatological scenarios. The dataset is widely adopted in skin lesion classification research due to its high-quality annotations and standardized dermoscopic imaging format, which facilitates reliable benchmarking and model comparison.

3.2 Data Preprocessing

The images were preprocessed to ensure uniform input for the model. This included resizing images to a fixed dimension, normalizing pixel values, and applying enhancement techniques such as grayscale conversion, thresholding, and noise reduction. Data augmentation methods, including rotation, flipping, and zooming, were also applied to increase dataset diversity and reduce overfitting.

Preprocessing plays a crucial role in improving the quality and consistency of input data before it is fed into the deep learning model. Resizing ensures that all images have the same spatial resolution compatible with the input requirements of the neural network. Normalization standardizes pixel intensity values, enabling stable and efficient training. Additionally, augmentation techniques help simulate

real-world variations in image orientation and scale, allowing the model to learn invariant features and improving its generalization ability. To address the class imbalance present in the HAM10000 dataset, multiple strategies were employed. Data augmentation techniques such as rotation, flipping, and zooming were applied to underrepresented classes to increase sample diversity. Additionally, class weighting was incorporated during model training to assign higher importance to minority classes. A balanced sampling strategy was also utilized to ensure equal representation of all lesion categories. These techniques collectively improved classification performance and reduced bias toward dominant classes.

3.3 Model Building and Training

A deep learning model based on the ResNet50 architecture was used for skin lesion classification. Transfer learning with pretrained ImageNet weights was applied and fine-tuned on the dataset. Additional fully connected and softmax layers were added for multi-class classification. The model was trained using the Adam optimizer with categorical cross-entropy loss.

ResNet50 is a deep residual network architecture that utilizes skip connections to address the vanishing gradient problem commonly encountered in deep neural networks. By leveraging transfer learning, the pretrained convolutional layers capture general visual features such as edges, textures, and shapes, which can then be adapted to the specific task of skin lesion classification. Fine-tuning the network allows the model to learn domain-specific features relevant to dermoscopic images while reducing the training time and computational requirements.

3.4 Model Interpretation

Explainable AI was incorporated using LIME to provide visual explanations for model predictions. LIME highlights important regions of the image influencing the classification, improving interpretability and trust in the system.

The integration of explainable AI techniques is particularly important in medical applications where transparency and reliability are critical. LIME generates local explanations by approximating the behavior of the trained model around a specific prediction. By highlighting influential regions of the dermoscopic image, the system allows clinicians to understand the reasoning behind the model's classification results. This interpretability helps bridge the gap between automated AI predictions and clinical decision-making.

3.5 User Interface and Chatbot Integration

A web-based interface was developed using Streamlit to allow users to upload images and receive classification results along with explanation maps. A chatbot was also integrated to provide basic dermatology information and assist users with related queries.

The user interface was designed to provide a simple and interactive platform for accessing the diagnostic capabilities of the proposed system. Through the interface, users can upload dermoscopic images, initiate automated analysis, and visualize the classification results together with explainable AI outputs. The integrated chatbot provides additional support by answering common dermatology-related questions and

guiding users through the analysis process. This integration enhances usability and ensures that both healthcare professionals and non-expert users can interact effectively with the system.

3.6 System Architecture




Figure 1. System architecture of the proposed XAI-DermNet framework.

Figure 1 depicts the System Architecture diagram of the proposed system.

3.7 System Specifications

The development of the proposed XAI-DermNet framework requires specific hardware and software configurations to support model training and implementation. These system specifications ensure that the deep learning model can be trained efficiently and that the preprocessing and visualization tasks can be executed without performance limitations.

3.7.1 Hardware Requirements

The system used for development includes a 12th Gen Intel Core i5-1235U processor with a speed of 1.30 GHz, 8GB RAM, and 512 GB storage. It operates on a 64-bit architecture with a minimum network speed of 20 Mbps to support cloud-based training and data access. This hardware configuration provides adequate computational capability for dermoscopic image preprocessing, model training, and result visualization during experimentation.

3.7.2 Software Requirement

The implementation was carried out using Python as the primary programming language. The deep learning model was developed using TensorFlow, with supporting libraries such as OpenCV and Matplotlib. Development and experimentation were conducted using Google Colab and Visual Studio Code on the Windows 11 platform. These tools provided the necessary environment for model training, image processing, and visualization of results.

The model training was accelerated using GPU-enabled computation. Specifically, training was performed on an

NVIDIA RTX 3060 GPU, requiring approximately 4–6 hours for convergence. The ResNet50 model contains approximately 25 million parameters, providing an optimal balance between computational efficiency and classification performance. These implementation details improve reproducibility and allow future researchers to replicate the proposed framework.

3.8 Tools and Technologies used

This study utilizes several tools and technologies to develop and implement the proposed XAI-DermNet framework. Python was used as the primary programming language due to its simplicity and extensive support for machine learning and data science libraries. The deep learning model was developed using TensorFlow, an open-source framework widely used for building and training neural networks, particularly convolutional neural networks for image classification tasks. Image processing and visualization were performed using Python libraries such as OpenCV and Matplotlib. Model development and experimentation were conducted on Google Colab, which provides GPU-enabled cloud computing for efficient training.

4. RESULTS AND DISCUSSION

The proposed XAI-DermNet system was implemented to perform automated skin lesion classification and provide interpretable results. The implementation ensures that the system functions efficiently and supports user interaction for image analysis and prediction.

The system is composed of the following key modules:

Data Collection and Preprocessing: Acquisition of dermoscopic images and preparation through resizing, normalization, and augmentation.

Model Building and Training: Development and training of the deep learning model for skin lesion classification.

Model Interpretation: Application of explainable AI techniques to visualize and explain model predictions.

Data Visualization and Insights: Presentation of model performance and analysis results through visual representations.

User Interface and Chatbot Integration: An interactive interface for image upload, prediction display, and user guidance.

These modules work together to support the implementation and evaluation of the proposed diagnostic framework.

4.1 Data Collection and Preprocessing Module

This module capture skin lesion images from the input source, such as digital cameras or image files. It preprocesses the images to ensure consistent quality and reduce noise, including tasks like resizing and normalization to meet the input requirements of the skin lesion analysis model. These preprocessing operations are essential for improving the reliability of the deep learning model by ensuring that all input images maintain a uniform resolution and standardized pixel distribution.

In addition to resizing and normalization, preprocessing techniques help remove unwanted artifacts and illumination variations that may affect the accuracy of the classification process. By improving the clarity and consistency of dermo-

scopic images, the system enables the deep learning model to focus on relevant lesion characteristics such as texture, color variation, and structural patterns.

This process generates a binary image in which the regions of interest are highlighted in white against back-ground, facilitating improved feature extraction for skin lesion analysis. The segmentation step plays a crucial role in isolating the lesion area from surrounding skin regions, thereby allowing the model to concentrate on diagnostically significant features during the training and prediction stages.



Figure 2. Dermoscopic images from the HAM10000 dataset.

Figure 2 presents the original dermoscopic images representing the fourteen classes of skin lesions before the data preprocessing stage. These images contain variations in lighting conditions, background artifacts, and color intensity, which necessitate preprocessing before they can be effectively used for machine learning-based analysis.

Figure 3 depicts the fourteen classes of skin lesions after the data preprocessing stage. The processed images exhibit enhanced clarity and reduced noise, enabling the model to capture meaningful visual patterns associated with different skin lesion categories. This preprocessing step significantly contributes to improving classification performance by providing structured and consistent input data for the deep learning framework.

4.2 Model Building and Training

The trained convolutional neural network (CNN), based on the ResNet50 architecture, is utilized for skin lesion classification. The model processes each preprocessed image to detect relevant features and classify the skin lesion into the corresponding category. The base layers of the pretrained model are partially frozen, while additional layers such as dense, dropout, and output layers are incorporated to adapt the network for accurate skin lesion classification.

The model compilation process is performed using the Adam optimizer. The compiled model architecture includes various layer types, output shapes, parameter counts, and the total number of trainable and non-trainable parameters. This con-




Figure 3. Pre-processed dermoscopic images from the HAM10000 dataset.

figuration ensures efficient learning and optimization during the training phase.

The training process of the CNN model shows a gradual increase in accuracy and a decrease in loss across training epochs. These trends indicate effective learning and good generalization performance during both training and validation phases.

To justify the selection of ResNet50, a comparative evaluation was conducted with other architectures, including EfficientNet, DenseNet, and Vision Transformer (ViT). Experimental observations indicate that ResNet50 achieves a favorable trade-off between accuracy, computational complexity, and training stability. While EfficientNet demonstrated slightly improved accuracy, it required significantly higher computational resources. DenseNet showed competitive results but increased memory usage, whereas ViT required larger datasets for optimal performance. Therefore, ResNet50 was selected as the most suitable architecture for the proposed system.




Figure 4. Performance metrics of the ResNet50 model.

4.3 Model Interpretation

Explainable AI techniques, such as LIME, are integrated to enhance transparency in the model's decision-making process.

This approach generates interpretable explanations for individual classification results, enabling users to understand the key factors influencing the model's predictions. By providing visual explanations, the system helps bridge the gap between complex deep learning models and human understanding.

Figure 4 illustrates the performance evaluation of the proposed XAI-DermNet framework using the ResNet50 model on the HAM10000 dataset. The bar chart presents key evaluation metrics, including Accuracy (95.73%), Specificity (94.80%), Precision (95.84%), Sensitivity (94.90%), and F1 Score (93.44%). These results indicate that the model achieves high classification performance with a balanced trade-off between sensitivity and specificity, which is critical for reliable medical diagnosis. The high precision and accuracy further confirm the model's effectiveness in correctly identifying skin lesion categories while minimizing misclassification.


During the prediction phase, the trained model classifies the test images into the predefined set of skin lesion classes. The system is capable of identifying conditions such as Actinic keratoses and distinguishing them from healthy skin images based on the learned features extracted during the training process. The explainability mechanism highlights the important regions of the input image that contribute most significantly to the classification result, thereby improving interpretability and user trust in the system. This transparency allows medical practitioners to better evaluate the reliability of the automated diagnosis.

4.4 Quantitative Evaluation of LIME

A quantitative evaluation of the LIME explanation technique was performed to assess its effectiveness in identifying clinically relevant regions. The overlap between LIME-highlighted regions and ground truth lesion areas was measured using Intersection over Union (IoU). The model achieved an average IoU score of approximately 0.72, indicating a strong correlation between highlighted regions and actual lesion boundaries. This result confirms that the explainable AI module effectively captures diagnostically important features and enhances model transparency.

4.5 Data Visualization and Insights

This module generates real-time visualizations of skin lesion analysis results by producing charts and graphs that illustrate key metrics such as lesion count, classification accuracy, and overall model performance. These visualizations enable continuous monitoring of the model's learning behavior during both training and evaluation phases. By clearly presenting trends and variations in performance metrics, the module



(a) Metric loss.

(b) Metric accuracy.

Figure 5. Training and validation metric loss and metric accuracy.

facilitates deeper analysis and easier interpretation of results. Additionally, it supports early identification of issues such as overfitting or performance degradation, thereby contributing to more reliable and informed diagnostic decision-making.

Figures 5(a) and 5(b) illustrate the training and validation loss and accuracy across epochs, respectively. The results show a consistent decrease in both training and validation loss, accompanied by a steady increase in accuracy, indicating effective learning and convergence of the model. The close alignment between training and validation curves suggests minimal overfitting and good generalization capability. Overall, these trends demonstrate that the model is learning meaningful feature representations and achieving stable performance as training progresses.

4.6 User interface and Chatbot integration

A user-friendly interface is developed using Streamlit, along with an integrated chatbot to provide user guidance and support throughout the interaction process. The interface enables users to upload dermoscopic images, initiate automated analysis, and view the classification results along with explanatory visualizations. The integrated chatbot assists users by answering basic dermatology-related queries and guiding them through the steps required to perform image analysis effectively. This interactive environment improves accessibility and allows users to easily interpret the diagnostic outputs generated by the system. It also enhances user engagement by providing a simple and intuitive platform for interacting with the intelligent skin lesion detection system.




Figure 6. User interface of the Intelligent Skin Lesion Detection System.

Figure 6 illustrates the user interface of the Intelligent Skin Lesion Detection System, highlighting its stream-lined design for ease of use. The interface allows users to upload skin lesion images, initiate automated analysis, and receive real-time diagnostic feedback. It also integrates visualization tools and interactive controls that facilitate efficient and accurate examination of dermoscopic images. The intuitive layout of the interface ensures that users can easily navigate the system and interpret the analysis results effectively.

Figure 7 presents the LIME (Local Interpretable Model-agnostic Explanations) visualization for the Intelligent Skin Lesion Detection System. This figure demonstrates how the model highlights the most influential regions of the input image that contributed to its prediction, providing interpretable insights into the decision-making process of the deep learning classifier. By visually mapping feature importance, LIME facilitates model transparency and helps users understand and trust the system's diagnostic outputs. This interpretability is particularly important in medical applications, where understanding the reasoning behind automated predictions supports reliable clinical decision-making.

Figure 8 illustrates the user guidance module of the Intelli-




Figure 7. LIME explanation for the Intelligent Skin Lesion Detection System.




Figure 8. User guidance module of the Intelligent Skin Lesion Detection System.

gent Skin Lesion Detection System. This module provides interactive assistance by offering step-by-step instructions for image upload, analysis execution, and interpretation of results. It enhances user experience by simplifying system navigation and reducing the complexity of the diagnostic workflow. Additionally, the module ensures that users can perform the analysis accurately with minimal effort, thereby improving usability and supporting effective interaction with the system.

5. CONCLUSION

This study presents an intelligent skin lesion analysis framework that combines deep learning with explainable AI to enable accurate and transparent dermatological diagnosis. The proposed XAI-DermNet utilizes a ResNet50-based CNN for effective classification of dermoscopic images, while LIME provides visual explanations to enhance interpretability and clinical trust. Experimental results confirm reliable performance, supported by preprocessing, transfer learning, and a user-friendly interface with chatbot assistance. Future work will focus on improving robustness through advanced architectures, ensemble methods, and more diverse datasets.

Additionally, extending the system to real-time clinical use, mobile platforms, and healthcare integration will enhance its practical applicability. Overall, the framework offers a reliable and interpretable solution for early skin disease detection and improved patient care.

REFERENCES

- [1] N. Codella, V. Rotemberg, P. Tschandl, M. E. Celebi, S. Dusza, M. A. Marchetti, and A. Halpern, "Skin lesion analysis toward melanoma detection: A challenge at the International Skin Imaging Collaboration (ISIC)," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 2, pp. 519–526, 2019.
- [2] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.
- [3] 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, p. 180161, 2018.
- [4] H. L. Gururaj, N. Manju, A. Nagarjun, V. M. Aradhya, and F. Flammini, "DeepSkin: A deep learning approach for skin cancer classification," *IEEE Access*, vol. 11, pp. 50 205–50 214, 2023.
- [5] Y. C. Lee, S. H. Jung, and H. H. Won, "WonDerM: Skin lesion classification with fine-tuned neural networks," arXiv preprint arXiv:1808.03426, 2018.
- [6] A. Mahbod, G. Schaefer, C. Wang, R. Ecker, and I. Ellinge, "Skin lesion classification using hybrid deep neural networks," in *ICASSP 2019 – IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 1229–1233.
- [7] J. Saeed and S. Zeebaree, "Skin lesion classification based on deep convolutional neural network architectures," *Journal of Applied Science and Technology Trends*, vol. 2, no. 1, pp. 41–51, 2021.
- [8] A. S. Shete, A. S. Rane, P. S. Gaikwad, and M. H. Patil, "Detection of skin cancer using CNN algorithm," *International Journal*, vol. 6, no. 5, pp. 1–4, 2021.
- [9] M. Ahmed, S. Sadek, and A. Mohammed, "Advancements in ResNet-based architectures for multi-class skin lesion classification," *Journal of Medical Imaging and Health Informatics*, vol. 14, no. 2, pp. 210–225, 2024.
- [10] F. Altini, D. Cascio, and G. Taormina, "A comparison of deep learning performance against healthcare professionals in detecting diseases from medical imaging," *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 1, pp. 112–124, 2024.
- [11] D. Brown and S. Green, "Fairness and bias in AI-driven dermatology: A review of current datasets," *Science Advances*, vol. 10, no. 4, p. eadj2024, 2024.
- [12] L. Chen, J. Huang, and K. Smith, "Improving clinical trust: A comparative study of LIME vs. Grad-CAM in dermatological AI," *Journal of Biomedical Informatics*, vol. 145, p. 104456, 2023.
- [13] R. Das and M. Sen, "A hybrid deep learning framework for skin lesion classification on imbalanced datasets," *Expert Systems with Applications*, vol. 213, p. 118897, 2023.
- [14] V. Gupta, P. Sharma, and R. Kumar, "Real-time skin cancer detection on edge devices using optimized ResNet50," *Sensors*, vol. 24, no. 1, p. 312, 2024.
- [15] A. Ibrahim, M. Hassan, and S. Ali, "Assessing the reliability of LIME for skin lesion explanations in clinical settings," *JMIR Medical Informatics*, vol. 11, p. e45892, 2023.
- [16] Z. Khalid, A. Rehman, and M. A. Khan, "Future of explainable AI in personalized healthcare: From theory to practice in dermatology," *Artificial Intelligence Review*, vol. 58, no. 1, pp. 45–68, 2025.
- [17] A. Khan, M. Rahim, and S. Hussain, "Deep learning-based automated skin lesion segmentation and classification: A 2024 perspective," *IEEE Reviews in Biomedical Engineering*, vol. 17, pp. 88–102, 2024.
- [18] B. Lopez, M. Garcia, and A. Fernandez, "Multi-modality data fusion for skin cancer diagnosis: Deep learning meets clinical history," *Medical Image Analysis*, vol. 91, p. 103012, 2024.
- [19] H. Malik, S. Farooq, and A. Jamil, "A novel ResNet-50 based approach for basal cell carcinoma detection," *Biomedical Signal Processing and Control*, vol. 85, p. 104832, 2023.
- [20] P. Mishra, R. Gupta, and A. Pratap, "Ensemble learning with ResNet50 and InceptionV3 for high-accuracy skin cancer screening," *Journal of Digital Imaging*, vol. 36, no. 4, pp. 1450–1465, 2023.
- [21] T. Nguyen and P. Vo, "Lightweight CNN architectures for early melanoma detection: A 2025 outlook," *Pattern Recognition Letters*, vol. 182, pp. 12–20, 2025.
- [22] K. Patel, A. Kumar, and J. Singh, "Beyond black-box models: Applying LIME and SHAP for interpretable skin cancer detection," *Diagnostics*, vol. 13, no. 4, p. 725, 2023.
- [23] M. Roberts, D. Driggs, and M. Thorpe, "Common pitfalls and recommendations for deep learning in skin cancer: A 2023 update," *The Lancet Digital Health*, vol. 5, no. 2, pp. e94–e105, 2023.
- [24] S. Sharma and R. Kumar, "Explainable AI for dermatology: A systematic review of XAI techniques in skin cancer diagnosis," *Artificial Intelligence in Medicine*, vol. 148, p. 102765, 2024.
- [25] S. Singh, A. K. Mondal, and A. Kumar, "Fusion of ResNet and Vision Transformers (ViT) for enhanced skin lesion analysis," *Computers in Biology and Medicine*, vol. 168, p. 107784, 2024.

-
- [26] J. R. Smith and L. White, “The role of transfer learning in rare skin disease identification,” *Nature Machine Intelligence*, vol. 5, no. 3, pp. 220–231, 2023.
 - [27] K. Sun, L. Zhang, and H. Wang, “Robust skin lesion classification under noisy labels using deep learning,” *Scientific Reports*, vol. 13, p. 11452, 2023.
 - [28] Y. Wang, X. Liu, and Z. Zhang, “Explainable AI (XAI) for medical diagnosis: Challenges and future directions in dermatology,” *IEEE Transactions on Medical Imaging*, vol. 43, no. 3, pp. 890–905, 2024.
 - [29] G. Zhang and F. Li, “Self-supervised learning for skin lesion classification with limited labeled data,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 35, no. 6, pp. 7800–7812, 2024.
 - [30] X. Zhao, Y. Wang, and Y. Sun, “Attention-guided ResNet for fine-grained skin lesion classification,” *Computerized Medical Imaging and Graphics*, vol. 112, p. 102334, 2024.