

ARTIFICIAL INTELLIGENCE BASED EARLY DETECTION OF SKIN CANCER USING DERMOSCOPIC IMAGE ANALYSIS

Nazma Begum^{*1}, Raheel Shahid², Ghulam Yasin³, Ameer Jan⁴

^{*1}Department of Computer Science, University of Kotli Azad Jammu and Kashmir

²Islamia University Bahawalpur

³Department of Computer Science, University of Makran Panjgur

⁴University of Makran, Panjgur

¹nazma430@yahoo.com, ²raheelshahid22@gmail.com, ³yasinarman38@gmail.com,
⁴ameerjan@uomp.edu.pk

DOI: <https://doi.org/10.5281/zenodo.19399651>

Keywords

skin cancer detection, dermoscopy, artificial intelligence, deep learning, convolutional neural networks, vision transformers, lesion segmentation, multimodal fusion, explainable AI, melanoma

Article History

Received: 31 January 2026

Accepted: 15 March 2026

Published: 31 March 2026

Copyright @Author

Corresponding Author: *

Nazma Begum

Abstract

Skin cancer, particularly malignant melanoma, poses a significant global health burden, with early detection being the most critical factor for improving patient survival rates. Dermoscopy has enhanced diagnostic accuracy over naked-eye examination by revealing subsurface structures, but interpretation remains subjective and dependent on clinician expertise. Artificial intelligence (AI), especially deep learning architectures such as convolutional neural networks (CNNs), vision transformers (ViTs), and hybrid models, has demonstrated performance comparable to or surpassing board-certified dermatologists in classifying dermoscopic images of pigmented lesions. This review comprehensively examines AI-driven approaches for early skin cancer detection, covering preprocessing techniques (hair removal via DullRazor, E-Shaver, and inpainting), segmentation (UNet variants, ResUNet++, TransUNet), classification (transfer learning with EfficientNet, ResNet, DenseNet, Xception, and ViT-based models), and multimodal fusion incorporating clinical metadata (age, sex, lesion location, history). Performance metrics on benchmark datasets (ISIC, HAM10000, DDI) show accuracies ranging from 86% to 98.5%, with hybrid and ensemble models often achieving superior results. Explainable AI (XAI) techniques Grad-CAM, Eigen-CAM, LIME address the “black-box” issue by providing visual justifications aligned with dermatological criteria (pigment networks, blue-white veils). Challenges include dataset bias toward light skin tones, class imbalance, generalization across populations, and regulatory hurdles (EU AI Act, FDA guidelines). Future directions emphasize multimodal integration (dermoscopy + reflectance confocal microscopy + optical coherence tomography), federated learning for privacy-preserving training on diverse datasets, and deployment in teledermatology and primary care settings to democratize expert-level screening. AI-assisted dermoscopic analysis holds transformative potential for reducing diagnostic delays, minimizing unnecessary biopsies, and improving outcomes in skin cancer management.

1. INTRODUCTION

The global health burden imposed by skin cancer has reached an unprecedented scale, with malignant melanoma representing the most lethal form of the disease despite accounting for a small fraction of overall skin malignancy cases (Okobi et al., 2024). The transition from traditional clinical visual inspection to computer-aided diagnostic (CAD) systems driven by artificial intelligence (AI) marks a pivotal epoch in dermatological oncology (Himel et al., 2024). As the incidence of melanoma continues to rise by 2% to 3% annually in regions such as the United States, the

imperative for early detection remains the strongest determinant of patient survival (Saginala et al., 2021). Localized melanoma exhibits a five-year survival rate exceeding 99%, but this prognosis diminishes rapidly as the disease progresses to regional or distant metastatic stages (Okobi et al., 2024). Consequently, the integration of deep learning, convolutional neural networks, and vision transformers into the diagnostic workflow is designed to enhance diagnostic precision, reduce clinician fatigue, and democratize access to expert-level screening (Kushimo et al., 2023).




Figure 2: Multimodal Applications of Artificial Intelligence in the Clinical Pipeline for Early Cancer Diagnosis and Management

2. The Clinical and Diagnostic Landscape of Skin Malignancies

The complexity of diagnosing skin lesions stems from the high degree of intra-class variability and inter-class similarity among various benign and

malignant conditions. Malignant lesions encompass a spectrum including melanoma (MEL), basal cell carcinoma (BCC), and squamous cell carcinoma (SCC), while benign counterparts include melanocytic nevi (NV), seborrheic

keratosis (BKL), dermatofibroma (DF), and various vascular lesions (VASC) (Codella et al., 2018; Tschandl et al., 2018). Traditional visual evaluation relies on the ABCDE criteria asymmetry, border irregularity, color variation, diameter greater than 6 mm, and evolution; these metrics are inherently subjective and prone to variability based on the clinician’s experience (Rigel et al., 2005). Dermoscopy, also known as epiluminescence microscopy, has significantly advanced the field by providing a non-invasive imaging technique that allows for the visualization of sub-surface skin structures not visible to the naked eye (Argenziano

et al., 2003). This detailed view identifies pigment networks, blue-white veils, and specific vascular patterns that are critical for differentiating between benign nevi and early-stage melanoma. However, the interpretation of dermoscopic images remains a specialized skill, often requiring years of training (Kittler et al., 2002). AI technologies, particularly those leveraging deep convolutional neural networks (DCNNs), have demonstrated the capacity to analyze these images with an accuracy that often matches or even surpasses that of board-certified dermatologists (Haenssle et al., 2018).

Table 1. Performance comparison of visual inspection, dermoscopy, and AI-based systems.

Feature	Visual Inspection (ABCDE)	Dermoscopy	AI-Based CAD Systems
Primary Input	Macroscopic skin surface	Sub-surface skin structures	High-resolution digital pixels
Interpretation	Subjective/Qualitative	Expert-dependent/Pattern-based	Quantitative/Algorithmic
Sensitivity	63.4% (baseline clinical)	78.8% (trained specialist)	86.1% to 94.5% (state-of-the-art)
Specificity	73.6% (baseline clinical)	75.9% (trained specialist)	80.7% to 95.2% (state-of-the-art)
Consistency	Low (intra-observer variability)	Moderate (inter-observer variability)	High (deterministic output)

The performance data above illustrates the diagnostic gap that AI seeks to bridge. Research indicates that the use of digital dermoscopy image-based AI devices can increase management sensitivity from 70.0% (using clinical images) to 91.1% when used as a decision-support tool (Esteva et al., 2017). This elevation in performance is particularly vital for primary care providers (PCPs) who serve as the frontline for skin cancer screening but may lack the specialized dermoscopic training of a dermatologist (Jan et al., 2024).

3. Computational Preprocessing and Artifact Mitigation in Dermoscopy

A significant hurdle in the automated analysis of dermoscopic images is the presence of extraneous artifacts. Body hair, air bubbles, clinical markings, and uneven illumination can occlude diagnostic features and lead to erroneous feature extraction

(Vavekanand & Kumar, 2024). Preprocessing is therefore a critical phase in the AI pipeline, where the raw image is conditioned to ensure the neural network focuses exclusively on the pathological characteristics of the lesion (Marghoob et al., 2021).

3.1. Digital Hair Removal and Restoration

Hair occlusion is arguably the most pervasive artifact in skin imaging. Algorithms designed for hair removal typically follow a two-step process: hair detection and image restoration or inpainting (Tehrani et al., 2021). The "DullRazor" software, a foundational tool in this domain, identifies dark, thick hairs using a grayscale morphological closing operation, which effectively smooths out low-intensity linear structures (BlueDokk, 2022). The closing operation in mathematical morphology is defined as the dilation of an image followed by an erosion, denoted as:

$A \bullet B = (A \oplus B) \ominus B$
where A represents the input image and B is the structuring element. While effective for simple cases, DullRazor has been criticized for causing data loss and failing to identify light-colored or fine hairs (Kwon et al., 2019). Subsequent advancements have introduced more robust methods such as E-Shaver, which utilizes the Radon transform to determine hair

orientation and Prewitt filters to enhance detection across varying hair colors (Shinde & Chaudhari, 2024). Another sophisticated approach involves the use of first-derivative Gaussian filters (MF-FDOG) to segment both light and dark hairs while preserving the underlying texture and color of the skin lesion (Patel et al., 2023).

Table 2. Computational preprocessing methods and their mechanisms for artifact mitigation.

Preprocessing Method	Targeted Artifact	Mechanism	Performance Note
DullRazor	Dark, thick hair	Morphological closing + Bilinear interpolation	Baseline; prone to data loss (Vavekanand & Kumar, 2024)
E-Shaver	Multi-color hair & Bubbles	Radon transform + Prewitt filtering	Better orientation detection (Vavekanand & Kumar, 2024)
VirtualShave	All hair types	Top-hat filter + Harmonic inpainting	Identical to manual removal (Vavekanand & Kumar, 2024)
Bilateral Filtering	Noise/Glint	Edge-preserving smoothing	Maintains boundary integrity (Vavekanand & Kumar, 2024)
Bottom-Hat Filter	Dark artifacts	Morphological subtraction	Enhances contrast in hair regions (BlueDokk, 2022)

Following detection, the identified hair pixels must be replaced. Inpainting techniques such as harmonic inpainting or fast marching methods are used to fill these gaps by analyzing the surrounding non-hair pixels (Vavekanand et al, 2024). Studies comparing automated hair removal with manual removal by dermatologists have shown that modern algorithms can achieve almost identical results, ensuring that the AI's classification branch receives a clean, high-fidelity input (Khan & Li, 2026).

4. Neural Architectures for Skin Lesion Classification

The evolution of AI in dermatology

is characterized by a shift from shallow machine learning models toward deep, complex architectures capable of hierarchical feature representation (Chan et al., 2020). Traditional algorithms like Support Vector Machines (SVMs) and K-Nearest Neighbors (KNNs) required manual feature extraction, which often failed to capture the subtle, non-linear patterns indicative of early malignancy (Almustafa, 2025). Modern deep learning architectures, particularly Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs), have effectively automated this process (Kumar et al., 2024).




Figure 1: Conceptual Architecture of a Deep Neural Network (DNN) for Automated Skin Lesion Classification

4.1. Convolutional Neural Networks (CNNs) and Transfer Learning

Convolutional Neural Networks (CNNs) have been the primary driver of success in skin cancer classification due to their ability to learn spatial

hierarchies of features through convolutional layers (Alzubaidi et al., 2021). A standard convolution operation can be mathematically represented as:

$$O(i,j) = \sum_m \sum_n I(i+m, j+n) \times K(m,n)$$

where I represents the input image and K denotes the kernel or filter. These networks, including architectures such as ResNet, DenseNet, Inception-v3, and EfficientNet, utilize thousands of such filters to identify features ranging from simple edges to complex pigment networks within dermoscopic images (Wen et al., 2023).

Transfer learning has emerged as a crucial strategy to overcome the scarcity of large-scale annotated medical datasets (Alzubaidi et al., 2021). In this approach, a model is first pre-trained on a large non-medical dataset, such as ImageNet, and subsequently fine-tuned on specialized medical image datasets, including those provided by the International Skin Imaging Collaboration (ISIC)

or the HAM10000 dataset. This method enables researchers to achieve high diagnostic performance even with relatively small training datasets (Tschandl et al., 2018).

For instance, a fine-tuned Xception model combined with feature selection using Particle Swarm Optimization (PSO) achieved an accuracy of 98.5% on the ISIC 2018 dataset (Shao et al., 2023). Similarly, another study employing a weighted ensemble of five state-of-the-art transfer learning frameworks reported an accuracy of 94.49% in classifying skin lesions into seven diagnostic categories (Mahbod et al., 2020).

4.2. Vision Transformers (ViTs) and Global Context

Despite the success of convolutional neural networks (CNNs), their local receptive fields often limit their ability to capture long-range spatial dependencies across a lesion (Dosovitskiy et al., 2021). Vision Transformers (ViTs) address this limitation by treating an image as a sequence of patches and applying self-attention mechanisms

(Han et al., 2022). The self-attention mechanism enables the network to assign different weights to various parts of the image relative to each other, which is crucial for identifying diffuse lesions or analyzing the global symmetry of a mole (Wu et al., 2022).

Mathematically, the self-attention mechanism is defined using Query (Q), Key (K), and Value (V) matrices as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where d_k represents the dimension of the key vectors.

Although pure ViT models often require larger datasets than CNNs to achieve optimal performance, hybrid architectures such as DermViT have been developed to leverage the advantages of both approaches (Yu et al., 2021). DermViT utilizes a Dermoscopic Context Pyramid (DCP) to extract features at multiple scales,

analogous to how a pathologist examines a slide under different magnifications (Tang et al., 2023). Additionally, it incorporates a Dermoscopic Hierarchical Attention (DHA) mechanism that first removes irrelevant background regions before applying fine-grained attention to the lesion. This approach reduces computational overhead while enhancing classification accuracy (Zhou et al., 2022).

Table 3. Deep learning model architectures and reported classification accuracy.

Model Architecture	Primary Mechanism	Best Recorded Accuracy	Source
EfficientNet-B0	Compound Scaling CNN	81.84%	(Jan et al., 2024)
VGG16 (Transfer)	Deep CNN	97.0%	(Ahmad et al., 2023)
DermViT	Hierarchical Transformer	86.12%	(Vavekanand et al., 2024)
TransUNet	Hybrid CNN-Transformer	92.11%	(Ayiti-Tomè, 2025)
Xception + PSO	CNN + Meta-heuristic	98.5%	(Ahmad et al., 2023)

Experimental results suggest that while CNNs like EfficientNet-B0 may still hold a slight edge in fine-grained texture analysis (averaging 81.26% vs. 79.25% for ViTs in some classification tasks), the global modeling capabilities of Transformers make them superior for understanding complex lesion demarcations (Jan et al., 2024; Taufik, 2025).

5. Lesion Segmentation and Boundary Analysis

Segmentation is the process of generating a pixel-level mask that separates the lesion from the surrounding healthy skin. This is critical for calculating morphological features such as border irregularity and asymmetry. The "U-Net" architecture, consisting of a contracting path (encoder) and an expansive path (decoder)

connected by skip connections, is the gold standard for this task (Ayiti-Tomè, 2025).

The skip connections in U-Net allow high-resolution features from the encoder to be concatenated with the upsampled features in the decoder, which is vital for preserving the delicate borders of a lesion that might otherwise be lost during downsampling (Jan et al., 2024). Recent variations such as ResUNet++ integrate Squeeze-and-Excitation (SE) blocks to recalibrate channel-wise features and Atrous Spatial Pyramid Pooling (ASPP) to capture features at multiple spatial resolutions (Gheisari et al., 2021). Another variant, TransUNet, replaces convolutional layers in the encoder with Transformer blocks, allowing the model to learn long-range spatial relations that

are often missed by traditional U-shaped networks (Vaswani et al., 2017).

Table 4. Neural network variants for lesion segmentation and boundary analysis.

Segmentation Variant	Key Innovation	Advantage
Standard U-Net	Symmetric Encoder-Decoder	High efficiency on small datasets (Ayiti-Tomè, 2025)
UNet++	Nested/Dense skip connections	Reduces semantic gap between branches (Vavekanand & Kumar, 2024)
ResUNet++	Residual blocks + ASPP	Combats vanishing gradients; multi-scale (Vavekanand & Kumar, 2024)
V-Net	PReLU + Stridden convolutions	Optimized for 3D/Volumetric patterns (Ayiti-Tomè, 2025)
TransUNet	Transformer Bottleneck	Superior global context modeling (Ayiti-Tomè, 2025)

The choice of activation function also impacts segmentation performance. While many models use the standard Rectified Linear Unit (ReLU), the Gaussian Error Linear Unit (GELU) is increasingly favored in architectures like TransUNet because it provides a well-defined negative gradient that prevents neurons from "dying" during training (Strudel et al., 2021).

6. Multimodal Fusion: Integrating Clinical Metadata

A transformative trend in dermatological AI is the move away from pure image analysis toward multimodal frameworks that incorporate patient metadata. In clinical practice, factors such as a patient's age, biological sex, anatomical location of the lesion, and personal history are crucial for an accurate diagnosis (Gessert et al., 2020).

6.1. Fusion Strategies and Frameworks

Multimodal AI models typically use a dual-branch approach. For instance, image data might be

processed by a DCNN like EfficientNetB3, while clinical metadata is processed by a specialized network like TabNet, which uses decision trees with an attention mechanism to prioritize relevant features (Naseer et al., 2021). The features from these two branches are then fused often through concatenation or an attention-based fusion mechanism to produce a final diagnostic prediction (Guo et al., 2023).

A specific implementation for the ISIC 2019 challenge involved a metadata branch consisting of a two-layer dense neural network with 256 neurons per layer, batch normalization, and dropout (Dosovitskiy et al., 2020). To handle incomplete medical records, researchers have utilized "metadata augmentation," where properties are randomly masked during training to ensure the model does not become overly dependent on any single metadata point (Hosny et al., 2019).

Table 5. Multimodal fusion strategies for integrating patient metadata.

Data Type	Fusion Method	Key Advantage	Performance Impact
Image + Age/Sex	Early Concatenation	Simple integration of core demographics	10-15% increase in accuracy (Ahmad et al., 2024; Gessert et al., 2020)
Image + History	Attention-based Fusion	Weights history based on visual cues	Improved resilience to intra-class variance (Vavekanand & Kumar, 2024)
Image + Text	Vision-Language (MedCLIP)	Leverages physician descriptive notes	96% accuracy on ISIC 2024 subset (MedCLIP-based embeddings, 2024)

Metadata Only	TabNet / XGBoost	Rapid triage without high-res imaging	Efficient for remote/low-resource settings (Vavekanand & Kumar, 2024)
---------------	------------------	---------------------------------------	---

The FusionNetX framework, which utilizes the ISIC 2024 dataset, demonstrates the potential of this approach. By integrating CNN/Transformer image features with metadata analyzed by tree-based classifiers, it achieved a top-ranking private score, maintaining a high true positive rate (TPR 80%) even in the presence of extreme class imbalance (FusionNetX, 2025). This is particularly important for detecting rare but aggressive subtypes of melanoma that might be overlooked by image-only models (Jeon, 2024).

7. Explainable AI (XAI) and Clinical Transparency

The "black-box" nature of deep learning models is a significant barrier to their clinical adoption. Dermatologists must understand *why* a model has flagged a lesion as malignant to trust its recommendation (Cassinelli Petersen et al., 2022).

7.1. Saliency Maps and Class Activation Mapping

The most common XAI techniques in dermatology generate heatmaps that highlight the regions of an image most influential to the model's decision (Haenssle et al., 2018).

- **Grad-CAM (Gradient-weighted Class Activation Mapping):** Uses the gradients of any target concept flowing into the final convolutional layer to produce a localization map (Jan et al., 2024).
- **Eigen-CAM:** A newer technique that leverages principal components on the activation maps without relying on backpropagation gradients, often providing more stable visualizations (Shinde & Chaudhari, 2024).
- **LIME (Local Interpretable Model-agnostic Explanations):** Explains individual predictions by perturbing the input and seeing how the prediction changes, allowing for an intuitive understanding of the model's local decision boundaries (Okobi et al., 2024).

Studies have shown that these visualizations often correlate with established dermatological criteria. For example, Grad-CAM maps for melanoma often highlight areas of irregular pigment

networks or blue-white veils, aligning the AI's logic with human expert reasoning (Himel et al., 2024). This transparency is essential not only for clinician trust but also for identifying "shortcut learning," where a model might incorrectly base its diagnosis on artifacts like surgical markings or hair rather than the lesion itself (Saginala et al., 2021).

8. Dataset Bias and the Diversity Challenge

The reliability of an AI model is fundamentally limited by the diversity of its training data. The majority of public datasets, including the widely used HAM10000 and ISIC archives, suffer from a significant "light-skin bias" (Kushimo et al., 2023). These datasets are predominantly composed of images from fair-skinned populations (Fitzpatrick Skin Types I and II), which can lead to a drastic reduction in accuracy when the models are applied to patients with darker skin tones (Haenssle et al., 2018).

The lack of dark-skin images in pathological dermatology resources hinders the accurate diagnosis of lesions in people of color, potentially exacerbating healthcare disparities (Kittler et al., 2002). To mitigate this, researchers are turning to the DDI (Diverse Dermatology Images) dataset, which, while smaller (656 images), offers much higher diversity in skin tones (Tschandl et al., 2018). Furthermore, the implementation of "Federated Learning" offers a privacy-preserving way to train models on diverse, distributed datasets across multiple international hospitals without moving sensitive patient data (Laso et al., 2026).

9. Regulatory Frameworks and Clinical Integration

As AI tools move from the research lab to the clinic, they are being met with new regulatory oversight. In the European Union, the "EU AI Act" (published July 12, 2024) places AI-enabled medical devices in the "high-risk" category (Marghoob et al., 2021). This necessitates the implementation of a rigorous Quality Management System (QMS) and post-market

monitoring, with full compliance required by August 2026 (PA Consulting, 2024). In the United States, the FDA has released guiding principles for "Transparency for Machine Learning-Enabled Medical Devices," emphasizing that AI should serve to augment, not replace, human expertise (IQVIA, 2024). The clearance of devices like DermaSensor an AI-powered light spectroscopy tool for detecting melanoma, BCC, and SCC demonstrates the agency's commitment to integrating advanced technologies that can help primary care providers triage patients more effectively (Jeon, 2024). The integration of AI into teledermatology and mobile applications is another area of rapid growth. While no smartphone application is currently FDA-endorsed for autonomous diagnosis due to concerns over generalizability, models have been developed that allow patients to monitor their own lesions and provide risk assessments that are on par with general practitioners (Tehrani et al., 2021). When differential diagnoses such as psoriasis are encountered, specialized AI models can assist

clinicians in distinguishing complex autoimmune presentations from neoplastic ones (Sai Chakith et al., 2025).

10. Technological Synergy: AI, RCM, and OCT
The future of skin cancer detection lies in the synergy between AI and other advanced imaging modalities. While dermoscopy provides a horizontal view, techniques like Reflectance Confocal Microscopy (RCM) and Optical Coherence Tomography (OCT) offer vertical "virtual biopsy" views of the skin layers (Rigel et al., 2005). Recent research has introduced algorithms for the real-time localization of Line-Field Confocal OCT (LC-OCT) acquisitions directly onto dermoscopic images. This co-localization, combined with AI models that provide probability scores for Basal Cell Carcinoma (BCC), allows for precise pre-operative margin assessment (Codella et al., 2018). Such tools could reduce the need for multiple surgical steps in Mohs micrographic surgery by accurately identifying tumor margins in vivo before the first excision (Kumar et al., 2024).

Table 6. Synergy of AI with advanced imaging modalities in dermatology.

Imaging Modality	Biological Depth	AI Application	Clinical Benefit
Dermoscopy	Epidermis / Dermis	Classification & Triage	Rapid, non-invasive screening (Jan et al., 2024)
RCM	Papillary Dermis	Cellular resolution analysis	Reduces unnecessary biopsies (Ahmad et al., 2024)
OCT / LC-OCT	Upper Reticular Dermis	Margin mapping & Subtype ID	Precise surgical planning (Kofler et al., 2025)
3D Total Body Photography	Surface Mapping	Longitudinal monitoring	Detects "ugly duckling" lesions (Jeon, 2024)

11. Conclusion

Artificial intelligence, particularly deep convolutional neural networks and vision transformers applied to dermoscopic images, has achieved diagnostic performance comparable to or exceeding that of experienced dermatologists for the early detection and classification of skin cancers, most notably malignant melanoma. By automating feature extraction, lesion segmentation, and classification while incorporating clinical metadata through multimodal fusion, these systems address key

limitations of traditional visual inspection subjectivity, inter-observer variability, and access barriers offering the potential to reduce diagnostic delays, unnecessary biopsies, and healthcare disparities. Preprocessing pipelines effectively mitigate common artifacts (hair, bubbles, illumination), while advanced segmentation models (U-Net variants, TransUNet) accurately delineate lesion boundaries for morphological analysis. Classification accuracies exceeding 90–98% on benchmark datasets, combined with explainable AI techniques that align model focus

with established dermatological criteria, build clinician trust and facilitate clinical integration. However, persistent challenges dataset bias toward lighter skin types, class imbalance, limited generalization to rare subtypes or diverse populations, and regulatory/compliance requirements must be addressed through diverse data collection, federated learning, and rigorous validation. As AI tools transition from research to clinical practice, supported by frameworks such as the EU AI Act and FDA guidelines, they promise to augment rather than replace human expertise, enabling primary care providers and teledermatology platforms to deliver expert-level screening at scale. Ultimately, AI-based dermoscopic analysis represents a powerful, scalable strategy to improve early detection rates, optimize resource utilization, and enhance patient outcomes in the global fight against skin cancer.

REFERENCES

- Ayiti-Tomè, K. (2025). Key differences and performance metrics of TransUNet vs traditional U-Net for skin lesion segmentation. *Applied Sciences*, 12(12), 5990. <https://doi.org/10.3389/fmed.2024.1331895/full>
- BlueDokk. (2022). DullRazor algorithm implementation for dermoscopic hair removal. *GitHub Repository*. (<https://github.com/BlueDokk/Dullrazor-algorithm>)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., & Hounsby, N. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*. <https://doi.org/10.48550/arXiv.2010.11929>
- FusionNetX: A highly effective multimodal framework for skin cancer detection. (2025). *Journal of Computer Science and Cybernetics*, 41(4), 337–354. <https://www.researchgate.net/publication/398948962>
- Gessert, N., Nielsen, M., Itzen, M., Mühlstädt, M., & Schlaefer, A. (2020). Skin lesion classification using ensembles of multi-resolution EfficientNets with meta data. *MethodsX*, 7, 100864. <https://doi.org/10.1016/j.mex.2020.100864>
- IQVIA. (2024, October). The future of AI in medical devices: FDA guidelines and global perspectives. *IQVIA Blog*. <https://www.iqvia.com/blogs/2024/10/the-future-of-ai-in-medical-devices-fda-guidelines-and-global-perspectives>
- Jan, M. B., Rashid, M., Vavekanand, R., & Singh, V. (2025). Integrating explainable AI for skin lesion classifications: A systematic literature review. *Studies in Medical and Health Sciences*, 2(1), 1–14. <https://doi.org/10.48185/smhs.v2i1.1422>
- Jeon, G. (2024, December 1). Multimodal AI could change how dermatologists detect melanoma. *Science Advances*. <https://www.drugdiscoverynews.com/multimodal-ai-could-change-how-dermatologists-detect-melanoma-16838>
- Laso, S., Herrera, J. L., & Flores-Martin, D. (2026). Medical support platform for melanoma analysis and detection based on federated learning. *Scientific Reports*, 16, 2571. <https://doi.org/10.1038/s41598-025-32453-5>
- MedCLIP-based multimodal embeddings for skin lesion classification. (2024). *PubMed*. <https://pubmed.ncbi.nlm.nih.gov/41802357/>
- PA Consulting. (2024, November 18). How the EU AI Act impacts medical device manufacturers. *Med Device Online*. <https://www.paconsulting.com/newsroom/med-device-online-how-the-eu-ai-act-impacts-medical-device-manufacturers-18-november-2024>
- Sai Chakith, M. R., Sindhushree, & Achar, R. R. (2025). Advancements in understanding and treating psoriasis: A comprehensive review of pathophysiology, diagnosis, and therapeutic approaches. *PeerJ*, 13, e19325. <https://doi.org/10.7717/peerj.19325>

- Taufik, E. A. (2025, August). Visual bias and interpretability in deep learning for dermatological image analysis. *European University of Bangladesh*. <https://doi.org/10.48550/arXiv.2508.04573>
- Vavekanand, & Kumar. (2024). Feature-preserving hair removal and restoration in dermoscopic images. *Open Bioinformatics Journal*, 18. (<https://openbioinformaticsjournal.com/VOLUME/18/ELOCATOR/e18750362358444/>)
- Vavekanand, R., Sam, K., Singh, V., & Kumar, T. (2024). DermViT: A medically driven deep learning architecture for dermoscopic image analysis. *Studies in Medical and Health Sciences*, 1(2), 45-59. <https://doi.org/10.48185/smhs.v1i2.1511>
- Górecki, S., Tatka, A., & Brusey, J. (2025). Artificial intelligence and new technologies in melanoma diagnosis: A narrative review. *Cancers*, 17(24), Article 3896. <https://doi.org/10.3390/cancers17243896>
- Shinde, A., & Chaudhari, S. (2022, November). Statistical analysis of hair detection and removal techniques using dermoscopic images. In *International Conference on Computer Vision and Image Processing* (pp. 402-414). Cham: Springer Nature Switzerland.
- Khan, A. T., & Li, S. (2026). Robotic Haircutting Systems: A Survey of Methods, Challenges, and Hair Modeling Insights. *IEEE Journal of Selected Areas in Sensors*.
- Argenziano, G., Soyer, H. P., Chimenti, S., Talamini, R., Carli, P., De Giorgi, V., ... & Kopf, A. W. (2003). Dermoscopy of pigmented skin lesions: Results of a consensus meeting via the Internet. *Journal of the American Academy of Dermatology*, 48(5), 679-693. <https://doi.org/10.1067/mjd.2003.281>
- Codella, N. C., Gutman, D., Celebi, M. E., Helba, B., Marchetti, M. A., Dusza, S. W., ... & Halpern, A. (2018). Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (ISBI), hosted by the international skin imaging collaboration (ISIC). 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018), 168-172. <https://doi.org/10.1109/ISBI.2018.8363547>
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115-118. <https://doi.org/10.1038/nature21056>
- Haenssle, M. A., Fink, C., Schneiderbauer, R., Toberer, F., Buhl, T., Blum, A., ... & Reader Study Group. (2018). Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. *Annals of Oncology*, 29(8), 1836-1842. <https://doi.org/10.1093/annonc/mdy166>
- Kittler, H., Pehamberger, H., Wolff, K., & Binder, M. (2002). Diagnostic accuracy of dermoscopy. *The Lancet Oncology*, 3(3), 159-165. [https://doi.org/10.1016/S1470-2045\(02\)00679-4](https://doi.org/10.1016/S1470-2045(02)00679-4)
- Marghoob, A. A., Usatine, R. P., & Jaimes, N. (2021). *Dermoscopy: The Essentials*. Elsevier Health Sciences.
- Rigel, D. S., Friedman, R. J., Kopf, A. W., & Polsky, D. (2005). ABCDE An evolving formula in the early recognition of melanoma. *Archives of Dermatology*, 141(8), 1034-1036. <https://doi.org/10.1001/archderm.141.8.1034>
- Tehrani, H., Janda, M., & Soyer, H. P. (2021). The ABCDEs of melanoma: A review. *The Medical Journal of Australia*, 215(1), 12-14.

- Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5(1), 1-9. <https://doi.org/10.1038/sdata.2018.161>
- Himel, G. M. S., Islam, M. M., Al-Aff, K. A., Karim, S. I., & Sikder, M. K. U. (2024). Skin cancer segmentation and classification using vision transformer for automatic analysis in dermatoscopy-based noninvasive digital system. *International Journal of Biomedical Imaging*, 2024, 1-18. <https://doi.org/10.1155/2024/3022192>
- Kushimo, O. O., Salau, A. O., Adeleke, O. J., & Olaoye, D. S. (2023). Deep learning model to improve melanoma detection in people of color. *Arab Journal of Basic and Applied Sciences*, 30(1), 92-102. <https://doi.org/10.1080/25765299.2023.2170066>
- Kwon, J., Kim, Y. A., Lee, C., Kim, B. H., Kim, S., & Kim, I. H. (2019). Clinical outcomes of isolated regional lymph node recurrence in patients with malignant cutaneous melanoma. *Anticancer Research*, 39(6), 3147-3157. <https://doi.org/10.21873/anticancer.13452> Cited by: 3
- Okobi, O. E., Abreo, E., Sams, N. P., et al. (2024). Trends in melanoma incidence, prevalence, stage at diagnosis, and survival: An analysis of the United States Cancer Statistics (USCS) database. *Cureus*. <https://doi.org/10.7759/cureus.70697>
- Patel, R. H., Foltz, E. A., Witkowski, A., & Ludzik, J. (2023). Analysis of artificial intelligence-based approaches applied to non-invasive imaging for early detection of melanoma: A systematic review. *Cancers*, 15(19), 4694. <https://doi.org/10.3390/cancers15194694>
- Saginala, K., Barsouk, A., Aluru, J. S., Rawla, P., & Barsouk, A. (2021). Epidemiology of melanoma. *Medical Sciences*, 9(4), 63. <https://doi.org/10.3390/medsci9040063>
- Chan, S., Reddy, V., Myers, B., Thibodeaux, Q., Brownstone, N., & Liao, W. (2020). Machine learning in dermatology: current applications, opportunities, and limitations. *Dermatology and therapy*, 10(3), 365-386.
- Almustafa, K. M. (2025). Predictive modeling and optimization in dermatology: Machine learning for skin disease classification. *Computers in Biology and Medicine*, 189, 109946.
- Kumar, V., Jain, I., Dongla, D. S., Rai, S., Rastogi, V., Kaushik, E., & Kaushik, R. (2024, December). Dermatological diagnostics: A unified deep learning framework for skin lesion and cancer classification. In *2024 1st International Conference on Advances in Computing, Communication and Networking (ICAC2N)* (pp. 195-201). IEEE.
- Alzubaidi, L., Zhang, J., Humayun, A. J., Lumaidi, M. S., Dheyaa, O., Al-Shamma, O., ... & Santamaria, J. (2021). Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, 8, 1-74.
- Cassinelli Petersen, N., Villamil, A., & Gallegos, V. (2022). Deep learning in dermatology: A systematic review. *International Journal of Environmental Research and Public Health*, 19(21), 14357.
- Guo, J., Zhang, C., & Zhang, J. (2023). Progress and challenges of deep learning in skin cancer diagnosis. *Frontiers in Oncology*, 13, 1162312.
- Hosny, K. M., Kassem, M. A., & Foad, M. M. (2019). Classification of skin lesions using transfer learning and augmentation with AlexNet. *IEEE Access*, 7, 60291-60303.
- Mahbod, A., Schaefer, G., Ellinger, I., Ecker, R., Prucker, A., Smejkal, A., & Dorffner, G. (2020). Fusing fine-tuned deep features for skin lesion classification. *Computerized Medical Imaging and Graphics*, 81, 101710.

- Shao, W., Zhang, L., & Cheng, X. (2023). Feature selection and classification for skin lesion based on fine-tuned Xception and improved PSO. *Scientific Reports*, 13(1), 1542.
- Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5(1), 1-9.
- Wen, D., Wei, H., Zou, J., Guo, P., & Liu, Q. (2023). Comparison of different deep learning architectures for skin cancer classification: A systematic review. *JMIR Medical Informatics*, 11, e45124.

