الجمهورية الجزائرية الديمقراطية الشعبية
People's Democratic Republic of Algeria
وزارة التعليم العالي والبحث العلمي
Ministry of Higher Education and Scientific Research
المدرسة الوطنية العليا للتكنولوجيات المتقدمة
National Higher School of Advanced Technologies

Génie Électrique et Infomatique Industrielle

Department of Electrical Engineering and Industrial Informatics

Final Year Project to Obtain the Diploma of Master

- Field -
Telecommunications

- Speciality -
Telecommunication Systems and Networks

- Subject -

State of the Art of Skin Anomaly Detection Using Machine Learning Techniques

Realized by :

BADAOUI Aymen
BIBET Idris

Members of The Jury :

name	Institution	Academic Rank	Role
Mme BOUTERFAS Malika	ENSTA	M.C.A	Chair
Mme HABANI Lamia–	ENSTA	M.A.A	Examiner
M ABBADI Belkacem	ENSTA	M.C.B	Examiner
Mme Kheira LAKHDARI	ENSTA	M.A.B	Supervisor
Mme Ferrahi Ibtissem	USTHB	M.C.B	Co-Supervisor

Academic year 2024–2025

State of the Art of Skin Anomaly Detection Using Machine Learning Techniques

BADAOUI Aymen
Ecole Nationale Supérieure
des Technologies Avancées, Alger, Algérie
Email: A_badaoui@ensta.edu.dz

BIBET Idris
Ecole Nationale Supérieure
des Technologies Avancées, Alger, Algérie
Email: I_bibet@ensta.edu.dz

***Abstract*—Skin cancer is one of the most common and deadly forms of cancer, for which early detection remains critical. In this paper, we provide a comprehensive review and technical comparison of artificial intelligence (AI) techniques for skin lesion analysis. We discuss a wide variety of conventional machine learning and latest deep learning techniques, with particular focus on pre-trained convolutional neural networks (CNNs) such as VGG, ResNet, EfficientNet, and MobileNet. An in-depth comparative analysis is provided in the form of classification accuracy, model complexity, training time, and inference speed. Preprocessing techniques—like hair removal, image normalization, artifact removal, and data augmentation techniques like geometric and color transformations—are examined for their effectiveness in enhancing robustness, particularly on real-world datasets like HAM10000, ISIC, and PH2. In addition, the paper addresses the feasibility of executing these models on embedded hardware platforms (Raspberry Pi, Jetson Nano, Coral Edge TPU) by analyzing their TensorFlow Lite and ONNX support and the impact of model compression techniques (quantization, pruning, distillation). The final cross-comparison identifies the optimal combination of dataset, model architecture, and hardware platform for creating an inexpensive, interpretable, and real-time skin lesion detection system. The paper concludes with the key takeaways for prototyping an embedded diagnostic system for low-resource environments.**

***Index Terms*—Skin cancer, skin disease detection, skin lesion datasets, deep learning, convolutional neural networks, vision transformers, embedded systems, low-cost diagnosis, explainable AI, real-time inference, data augmentation, transfer learning, mobile deployment**

I. Introduction

Both melanoma and non-melanoma skin cancers, represent an important and increasing global health problem. According to GLOBOCAN 2022 over 330,000 new cases of melanoma and around 1.2 million non-melanoma skin carcinomas (NMSC) were diagnosed worldwide in 2020. worldwide [1].While nonmelanoma skin cancers occur more frequently, melanoma is responsible for the highest share of skin cancer deaths due to its higher metastatic potential. Fortunately, prognosis will improve substantially with early detection; localized melanoma, for example, has a 5-year survival rate of greater than 99%, whereas the 5-year survival rate for distant metastasis from melanoma is less than 35The following image depicts the distribution of new cancer cases globally by region in 2020 [2].

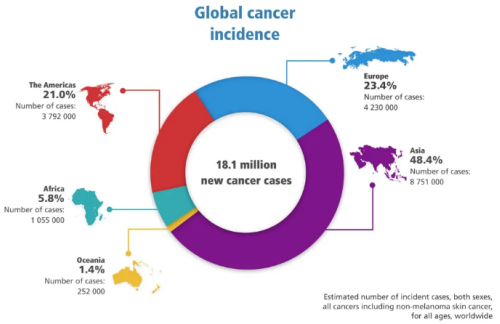

Fig. 1. Global Cancer Incidence by Region (2020) [3]

Classic diagnosis is made by visual inspection with the help of dermoscopy, followed by biopsy and histopathological analysis. The process is reliant upon medical expertise and access to dermatology services, which often are limited in low-resource or rural areas. As Incidences of skin cancers and other dermatosis tally in the millions, there is increasing prominence and use of artificial intelligence, especially deep learning, to assist and possibly automize skin lesion diagnosis [4].

Recent developments in computer vision, most notably convolutional neural networks (CNNs), have shown great potential for classifying dermoscopic images. Research indicates that well-trained CNNs may equal or exceed diagnostic accuracy in some tasks when compared with experienced dermatologists [4]. Architectures such as EfficientNet, MobileNetV3, and Vision Transformers (ViTs) [5] have been employed to balance diagnostic accuracy with computational efficiency, making them suitable for real-time applications and deployment on embedded or mobile platforms.

Nonetheless, several challenges remain. A critical issue is dataset bias: most publicly available skin lesion datasets (e.g., ISIC [6], HAM10000 [7], PH2 [8]) predominantly feature lighter skin tones (Fitzpatrick types I–III), leading to models that may underperform on patients with darker skin [9]. This lack of representation raises both ethical and clinical concerns, as it can contribute to disparities in diagnostic outcomes.

Another limitation lies in the interpretability of AI models. As medical applications demand transparency and accountability, black-box models are often met with skepticism. To address this, explainability techniques such as Grad-CAM and LIME have been introduced to visualize the regions influenc-

ing model predictions.The methods were created to showcase model output that are as clinically meaningful (i.e., asymmetry, border irregularity, color variation) to build trust [10].

Additionally, using AI-enabled systems in operational settings must consider hardware capabilities, compliance with laws and regulations, and compatibility with clinical workflows. Lightweight models are being considered and optimized through pruning, quantization, or knowledge distillation to be practical on devices with limited computational capabilities. At the same time, regulatory bodies like the FDA and EMA are beginning to approve AI-powered diagnostic tools, signaling a move toward broader adoption [11].

This review outlines current advances (2017-2024) in artificial intelligence in skin cancer detection. We provide a summary of improvements in model architecture, training approaches, dataset issues, interpretability, as well as designs for implementation in embedded context. Ultimately, we want to highlight viable opportunities for creating low-cost, interpretable, and accurate diagnostic systems that can be utilized in real time to tackle health inequities across the globe.

II. Related Work

Early automated skin lesion analysis relied on traditional machine learning (ML) using handcrafted image features. These methods typically extract descriptors (color, texture, shape) and classify lesions with algorithms like support vector machines (SVM), k-nearest neighbors, decision trees or logistic regression [12]. For example, Bechelli and Delhommelle compared several ML classifiers (logistic regression, LDA, KNN, decision tree, naïve Bayes) to deep CNNs on ISIC/HAM10000 data, finding that CNNs far outperformed ML (e.g., VGG-16 achieved ∼88% accuracy vs. only ∼70% for ML) [12]. In general, traditional ML systems attained modest accuracy and suffered when datasets grew large or diverse, paving the way for deep learning approaches.

A. Deep Convolutional Neural Networks

The advent of deep convolutional neural networks (CNNs) has revolutionized lesion classification. Standard architectures (AlexNet, VGG, ResNet, DenseNet, Inception, etc.) have either been fine-tuning, or trained on dermoscopic images, and frequently times with expert or near-expert performance. Ensemble CNNs are particularly good: for instance, one study reports that a majority-vote ensemble of InceptionV3, DenseNet, InceptionResNetV2 and VGG-19 achieved ∼98% accuracy on the ISIC archive [13]. Naqvi et al. note that deep models on ISIC obtain high accuracy (e.g. 88%) and F1-scores, whereas on larger, more heterogeneous datasets like HAM10000 the best CNN (VGG-16) reached only ∼70% accuracy, highlighting generalization challenges [14]. Cheng et al. propose a MobileNet-based network with a novel spatial-channel attention module; this “MobileNet-MFS” achieved over 87% accuracy on ISIC-2019 data and outperformed competing models in precision, recall and F1-score [15] .

The following image shows the CNN architecture used to classify skin lesions into 7 categories

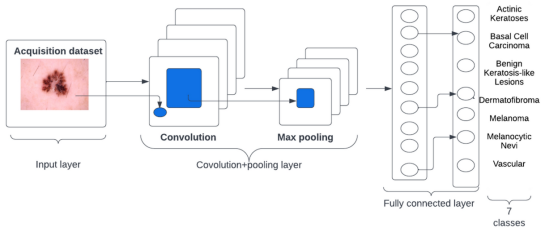

Fig. 2. Convolutional Neural Network (CNN) for Skin Lesion Classification [16]

B. Vision Transformers

In addition to CNNs, vision transformer (ViT) architectures have also been applied to dermoscopy. ViTs and hybrid CNN-Transformer networks arose in a similar timeframe (2020-2022). A recent scoping review shows that the use of ViTs in diagnosing skin cancers has climbed significantly in 2020-2022 with "superior performance" on dermoscopic images. However, ViTs tend to require very large training sets and can have difficulty with smaller datasets - as de Yuhan and Zhenglin noted, transformers generally performed poorly with low image counts [17], whereas CNNs remain robust on smaller samples. Hybrid strategies are sometimes used to balance power and data demands.

C. Hybrid and Ensemble Models

Many recent models combine multiple learning strategies. Khan et al. extract features from ResNet101 and DenseNet201, optimize them via a Moth-Flame optimizer, and classify with a kernel extreme learning machine, achieving 90.7% accuracy on HAM10000 [18]. Amin Tajerian fuse deep CNN features with maximum correlation analysis and classify with KELM, reporting accuracies of 95–99% across several datasets [19]. Gilani et al. introduce a spiking neural-network variant of VGG-13 achieving 89.6% accuracy and improved energy efficiency [20].

D. Interpretability and Explainability

Interpretable AI has gained a lot of attention recently. Grad-CAM, SHAP, and LIME are most noted for their capacity to visualize decisions. A recent literature review (2019-2023) indicates that most CNN based pipelines include an explainability module [21]. These module augment clinician trust and expose bias, however, they still come with hurdles such as imbalanced datasets and generalizeability [21].

E. Segmentation and Classification Pipelines

Segmentation is often the first stage in CAD pipelines, handled by encoder–decoder CNNs such as U-Net and its variants. Khan et al. proposed a 10-layer saliency-based CNN [22], and Adegun et al. presented a Gaussian-refined lightweight encoder–decoder achieving 98% segmentation accuracy [23].

F. Datasets and Benchmarks

Datasets like HAM10000, ISIC and PH2 are widely used for benchmarking. Data augmentation and multi-institutional images are often used to address melanoma rarity and improve generalization.

G. Mobile and Embedded Implementation

Mobile and embedded deployment models usually leverage lightweight CNN models like MobileNet or EfficientNet-lite; one particular example is Cheng et al.'s MobileNet-MFS model. Gilani et al.'s spiking neural VGG-13, also designed for neuromorphic hardware, adds to the mix as an efficient inference model. Real-world applications such as DermaSensor, do exist, and in some cases with commercial trials as well, either using on-board or cloud-based inference [11].

H. Regulatory Approval and Clinical Validation

DermaSensor is the first FDA-cleared AI skin cancer diagnostic device [11]. A Mayo Clinic trial found 96% sensitivity in identifying cancerous lesions. The device felt it cut melanoma miss-rates in primary care in half as well as improved referral rates. Only time will tell what AI clinical integration looks like at scale.

III. Review Methodology

The literature review strategy was based on a Boolean query combining relevant keywords related to artificial intelligence and skin lesions. This query was applied across multiple leading scientific databases.

A. Sources and Search Engines

The following databases were consulted:

- **IEEE Xplore**, for engineering and computer vision research;
- **PubMed / Medline**, for biomedical and clinical articles;
- **SpringerLink**, **Scopus**, **MDPI**, and **ScienceDirect**, for multidisciplinary journal coverage;
- **Google Scholar**, used as a supplementary search engine.

Keywords used included: *"skin lesion classification"*, *"melanoma detection"*, *"deep learning dermatology"*, *"AI skin cancer"*, combined with logical operators (AND, OR).

B. Timeframe and Selection Criteria

The search was limited to publications from **2017 to 2024** to capture recent advances in deep learning and AI-assisted diagnostics.

a) Inclusion criteria::

- Original research articles published in peer-reviewed journals;
- Experimental studies applying AI to the detection or classification of skin lesions;
- Availability of quantitative results using public or reproducible datasets;
- Use of CNN architectures, hybrid methods, or embedded models.

b) Exclusion criteria::

- General reviews, non-peer-reviewed preprints, book chapters, or short communications;
- Theoretical studies without experimental validation;
- Studies outside the medical or dermatology context;
- Publications dated before 2017 or not available in full text.

C. Selection Procedure

Once the results were merged and duplicates were deleted, an initial screening was carried out on titles and abstracts, followed by full text assessment of all selected articles. Two reviewers independently determined the relevance of each study. Disagreements were resolved by discussion or a third opinion.

For each selected publication, the following data were extracted: authors, year, dataset used, AI model, preprocessing and optimization techniques, and performance metrics.

D. Thematic Organization

The reviewed articles were categorized along four major axes:

1) **Datasets used** (ISIC, HAM10000, PH2, etc.);
2) **Learning models** (classical CNNs, MobileNet, ResNet, etc.);
3) **Deployment on embedded hardware** (Raspberry Pi, Jetson Nano, Coral TPU);
4) **Performance and comparisons** (accuracy, F1-score, latency, etc.).

This method ensures rigorous and focused coverage of recent, relevant, and experimental contributions to AI systems for automatic skin lesion detection.

IV. Comparative Study of Skin Imaging Databases

The development of AI systems for the automatic detection of skin lesions relies heavily on the availability of publicly annotated datasets. This section provides a comparative analysis of the main skin image databases commonly used in the literature [6], [7], [24]–[29].

A. Overview of the Datasets

The most widely used datasets are: **ISIC Archive**, **HAM10000**, **PH2**, **Dermofit**, **Derm7pt**, **PAD-UFES-20**, **BCN20000**, and **MED-NODE**.

- **ISIC Archive** contains tens of thousands of annotated dermoscopic images collected from various international centers. Some subsets (ISIC Challenges) include segmentation masks [6].
- **HAM10000** consists of 10,015 dermoscopic images representing 7 lesion types, accompanied by clinical metadata (age, sex, lesion location) [7].
- **PH2** is a smaller dataset (200 images) with manual segmentation and precise diagnoses, ideal for validation purposes [24].
- **Dermofit** includes 1,300 macroscopic clinical photographs covering 10 lesion categories [25].
- **Derm7pt** provides annotations following the 7-point checklist method used in dermoscopic diagnosis [26].
- **PAD-UFES-20** contains 2,298 clinical images taken with smartphones in Brazil, covering a variety of skin phototypes [27].
- **BCN20000** brings together more than 19,000 dermoscopic images from the Hospital Clinic of Barcelona [28].

- **MED-NODE** includes 170 clinical photos classified into two groups: melanoma and nevus [29].

B. Skin Lesion Datasets

Table I represents a comparative study of the most commonly used skin lesion image datasets. It highlights their size, diversity, annotation types, and acquisition modalities.

C. Critical Analysis and Recommendations

These datasets exhibit high clinical quality; however, several limitations must be considered when developing an embedded system:

- **Population bias**: The majority of datasets (e.g., ISIC, HAM10000, BCN20000) predominantly contain images of light-skinned individuals. This can reduce accuracy for darker skin tones.
- **Class imbalance**: Benign lesions (e.g., nevi) are significantly overrepresented compared to malignant ones such as melanoma.
- **Dataset size**: Large datasets such as ISIC, HAM10000, and BCN20000 are suitable for training; smaller ones like PH2, PAD-UFES, or MED-NODE are better suited for validation.
- **Image type**: Embedded systems designed for use with standard cameras should be validated on clinical photos (e.g., PAD-UFES, MED-NODE), not only dermoscopic images.

For an AI model intended for an embedded platform, the following recommendations are suggested:

- Train on a large and diverse corpus (e.g., ISIC, BCN20000),
- Validate on external datasets that include diverse skin phototypes (e.g., PAD-UFES),
- Test on clinical photographs that closely reflect real-world conditions (e.g., smartphone images).

Special attention should be paid to data augmentation techniques and the representativeness of cases to ensure robust performance under real-world deployment conditions.

V. AI Techniques for Skin Lesion Analysis

A. Machine Learning Foundations

The traditional **machine learning** (ML) methods in dermatology are based on hand-crafted features and traditional classifiers. The early CAD systems extracted color, shape and texture descriptors (e.g., histogram, GLCM, ABCD rule), and trained models (e.g., support vector machines (SVM), k-nearest neighbors (k-NN), and decision tree) [30]. Traditional ML CAD systems require careful preprocessing and segmentation of the skin lesions and suffer from the strong dependency on the quality of hand-crafted features.

While these approaches are light on computation and interpretable, they do not generalize well across complex or varied image sets. Recent efforts combine deep features from CNNs (e.g., ResNet, MobileNet) with classical ML classifiers and have the advantages of rich representation as well as a low inferential cost [31].

B. Deep Learning Advances

Convolutional neural networks (CNN), a branch of **Deep learning** (DL), have revolutionized the field of skin lesions analysis. CNNs ignore manual feature engineering and instead learn to extract hierarchical features from raw pixels automatically [31]. Architectures like ResNet, DenseNet, EfficientNet, and MobileNet have performed well in dermoscopic classification with high accuracy.

Deep models require larger datasets and more computation but can model complex lesion patterns more effectively. With large public datasets like ISIC or HAM10000, CNNs often outperform classical ML methods [30]. Ensemble strategies combining multiple CNNs (e.g., VGG16 + InceptionV3 + ResNet50) have also been shown to boost performance to over 95% accuracy [32].

To illustrate the overall workflow of transfer learning applied to skin lesion classification, the following diagram presents a two-stream architecture integrating pre-trained CNN models (Xception and ShuffleNet), feature fusion, and optimized classification across seven lesion types.

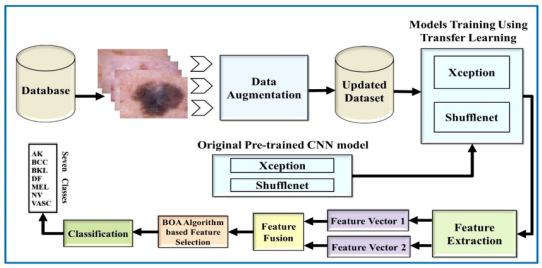

Fig. 3. Overview of a Two-Stream Transfer Learning Architecture for Skin Lesion Classification [33]

C. Transfer Learning in Dermatology

Transfer learning (TL) addresses the challenge of limited medical image data by adapting pre-trained CNNs (e.g., on ImageNet) to dermatology tasks. A common practice is to fine-tune architectures like ResNet, MobileNet, or EfficientNet on dermoscopic datasets [31], [34].

In TL, either the entire network is retrained or only the top layers are updated. This reduces training time and improves accuracy on small datasets like PH2. For example, EfficientNet-B7 achieved over 84% accuracy in melanoma detection after fine-tuning on limited data [34].

TL is also effective for clinical photographs captured via smartphones, where data variability (lighting, focus) presents additional challenges. Studies confirm that TL improves performance over training from scratch even in non-dermoscopic settings [35].

D. Comparative Discussion

- **Feature extraction:** ML relies on handcrafted descriptors, DL learns features end-to-end, and TL reuses pre-trained features from general domains.

TABLE I
COMPARISON OF SKIN LESION IMAGE DATASETS.

Dataset	No. of Images	Lesion Classes	Annotations	Origin / Phototypes	Image Type
ISIC Archive	>30,000	Melanoma, nevus, BCC, AK, etc.	Diagnosis, sometimes masks	Europe, Australia (light skin types)	Dermoscopy
HAM10000	10,015	7 pigmented classes	Histological diagnosis + metadata	Austria, Australia (I–III)	Dermoscopy
PH2	200	Typical/atypical nevus, melanoma	Diagnosis + manual segmentation	Portugal (light skin)	Dermoscopy
Dermofit	1,300	10 classes	Simple diagnosis	United Kingdom	Clinical photo
Derm7pt	2,013	15 classes	7-point checklist	Italy (light skin)	Dermoscopy + macro
PAD-UFES-20	2,298	6 classes	Clinical diagnosis	Brazil (diverse phototypes)	Smartphone photo
BCN20000	19,424	9 classes	Clinical/histological diagnosis	Spain	Dermoscopy
MED-NODE	170	2 classes (nevus, melanoma)	Histological diagnosis	Netherlands	Clinical photo

- **Data requirements:** ML performs adequately on small datasets; DL requires large datasets or augmentation; TL enables DL to generalize from fewer examples.
- **Performance:** DL and TL significantly outperform ML in complex classification tasks, especially when trained on dermoscopic datasets [30], [31].
- **Device constraints:** ML models are light and interpretable, making them suitable for microcontrollers. DL/TL models like MobileNet or EfficientNet-lite, once quantized, can run efficiently on devices like Raspberry Pi or Edge TPU [34], [35].
- **Explainability:** ML features are inherently understandable; DL requires tools like Grad-CAM to provide interpretability.

In summary, ML is suitable for lightweight scenarios but has limited flexibility. DL provides superior performance at higher computational cost. TL bridges the gap by enabling powerful models to perform well on limited and diverse medical datasets. The choice depends on the target deployment context, data availability, and accuracy requirements [12].

VI. COMPARATIVE ANALYSIS OF PRE-TRAINED CNN MODELS

Pre-trained convolutional neural networks (CNNs) are extensively utilized for the classification of dermoscopic images. The most common architectures in dermatology are VGGNet, GoogleNet/Inception, ResNet and its variations, DenseNet, EfficientNet, and MobileNet [36].They are typically implemented using transfer learning (i.e., weights pre-trained on ImageNet) and fine-tuned on dermatological datasets like ISIC, HAM10000 or PH^2.

A. Performance on Dermoscopic Datasets

Performance varies depending on the task (binary vs. multiclass classification) and the dataset. However, several studies from 2019–2024 allow for general performance trends to be drawn:

- **VGG-16** (≈138M parameters) generally achieves around 82.8% accuracy on dermoscopic lesion classification tasks [37]. It is a very deep but heavy model, whose uniform architecture (3×3 conv layers) remains effective.
- **ResNet-50** (≈25.6M parameters) is often used as a baseline: for example, ≈88% accuracy was reported on HAM10000 [38]. Residual networks benefit from avoiding vanishing gradients and offer fast convergence, with moderate size.
- **InceptionV3** (≈23.9M) integrates multi-scale convolution modules. Reported performance is comparable: e.g., 86.9% accuracy on HAM10000 [38]. Its complex architecture tends to converge more slowly.
- **DenseNet-121** (≈8.1M) uses dense layer connections. It is among the top performers, with around 87–88% accuracy on HAM10000 in various studies [39]. With fewer parameters than ResNet or VGG, it enables compact training and strong feature reuse.
- **EfficientNet-B0/B1** (≈5.3M/7.9M) are scaled models optimized for efficiency. They typically achieve ≈84.12–86.41% accuracy on HAM10000 [40]. Their performance-to-complexity ratio is excellent, offering good accuracy with low parameter count.
- **MobileNetV2/V3** (≈3–5M) are optimized for lightweight inference. For example, MobileNet-V3 achieved ∼89% accuracy in a 10-class dermatology classification task [41]. In practice, MobileNetV2/V3 offer slightly lower accuracy than larger networks (typically 80–90%) but are fast to train and deploy.

B. Pre-trained CNN Models

Table II presents a comparative evaluation of widely used pre-trained CNN models applied to skin lesion classification tasks.

C. Advantages and Limitations in the Medical Context

- **Interpretability**: All of these CNNs are essentially black-box models. None is inherently interpretable, but tools

TABLE II
COMPARISON OF PRE-TRAINED CNN MODELS FOR SKIN LESION IMAGE CLASSIFICATION.

Model	**Parameters (M)**	**Size (MB)**	**Accuracy (%)**	**Convergence**	**Remarks**
VGG-16	138.4	528	82.8 [37]	Slow	Very heavy; effective but prone to overfitting.
ResNet-50	25.6	98	88 [38]	Fast	Stable results and performant architecture.
InceptionV3	23.9	92	86.9 [38]	Rather slow	Multi-scale branches; good trade-off.
DenseNet-121	8.1	33	87 [39]	Medium	Good generalization; efficient feature reuse.
EfficientNet-B0	5.3	29	84.12 [40]	Fast	Excellent accuracy-to-size balance.
EfficientNet-B1	7.9	31	86.41 [40]	Fast	Slightly more accurate than B0; still very light.
MobileNetV2	3.5	14	82.58 [41]	Very fast	Ultra-compact; ideal for mobile deployment.
MobileNetV3 Large	5.4	16	89 [42]	Very fast	Optimized for embedded inference.

like Grad-CAM help visualize activated regions.

- **Overfitting**: Larger networks (VGG, ResNet) are more prone to overfitting. Lighter architectures such as DenseNet and EfficientNet help mitigate this risk.
- **Bias**: None of the models includes built-in bias mitigation. Balanced datasets are therefore essential.
- **Robustness and Generalization**: DenseNet and EfficientNet show strong generalization capabilities. ResNet remains a stable reference. MobileNet is fast but sometimes less accurate.
- **Deployment and Regulation**: ResNet and EfficientNet-B0/B1 offer strong candidates for reliable and interpretable medical deployment (with visualization tools), while remaining compact and robust.

D. Recommendations

Based on the above **ResNet** and **EfficientNet-B0/B1** show the best performance and availability for a reliable and general-purpose classification system for pigmented lesions. They provide the best balance of accuracy, compactness and generalizability, have some resistance to overfitting and are very suited to transfer learning.

VII. PREPROCESSING AND DATA AUGMENTATION TECHNIQUES

A. Image Preprocessing for Dermoscopic Images

Raw dermoscopic images are frequently resized to a particular input size (e.g., 224×224 pixels) to meet the requirements of CNN architectures [43]. There are also color normalization methods which address illumination and color imbalances, such as the gray-world or Shades-of-Gray algorithms [44]. Contrast enhancement methods like CLAHE (Contrast Limited Adaptive Histogram Equalization) are applied to make lesion features more visible [45]. Noise reduction filters (e.g., Gaussian or median) are also used to suppress image artifacts while preserving edges.

Removing artifacts is a necessary first step during image preprocessing, since the lesion images may contain hair strands, ruler markings, or gel bubbles, which may interfere with identification of lesion borders. Conventional approaches such as DullRazor have used inpainting, and morphological filtering , more recent comparing approaches focused on deep learning (e.g., SharpRazor), which have shown superiority in automated artifact detection and removal [46]. Then, segmentation algorithms (e.g., active contours, U-Net, or GVF snakes) can be used to isolate the lesion from the national skin surrounding and background [45], increasing the model's attention and relevance.

To clarify the entire preprocessing and segmentation pipeline, the following flowchart illustrates each stage involved in preparing dermoscopic images for classification.

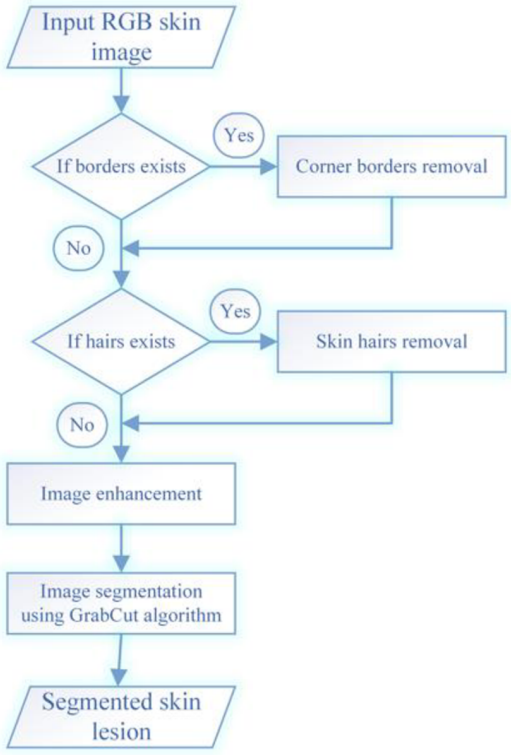

Fig. 4. Flowchart of the Preprocessing and Segmentation Pipeline for Skin Lesion Analysis [47]

B. Image Preprocessing for Clinical (Smartphone) Images

Smartphone-acquired clinical images pose specific challenges such as inconsistent lighting, blur, background noise, and shadows. Preprocessing typically begins with white balance correction and illumination normalization [44]. CLAHE is also used to compensate for lighting variations and enhance lesion contrast [45]. Since lesions in smartphone images are often surrounded by irrelevant background, background cropping or segmentation is essential to reduce noise before classification. Smoothing or masking is used to suppress shadows and reflections.

C. Data Augmentation Strategies

Data augmentation is imperative for improving model generalization and reducing overfitting. Common transformations, typically in a geometric sense, include flipping, rotation, scaling, and translation [43]. Augmentations like these help the model learn invariant features, so that the orientation or position of the lesion is not significant.

Photometric augmentations (for example, brightness, contrast and hue jitter) can be useful to simulate changes in illumination. Gaussian noise injection is another useful augmentation and helps the model to account for sensor noise. Elastic distortions and grid warping can be used to mimic variations in skin elasticity or shape.

Advanced methods such as **MixUp** and **CutMix** create hybrid images by blending or patching regions from different samples. These strategies reduce class overfitting and increase robustness to spatial variation . For example, MixUp has shown improvements in ISIC classification by balancing minority classes .

Generative approaches using GANs (Generative Adversarial Networks) are increasingly used to synthesize realistic lesion images, especially for rare classes. For instance, Self-Transfer GANs can generate high-fidelity synthetic lesions, improving training balance and diversity [48]. Domain adaptation GANs are also used to convert clinical images into dermoscopy-like images, addressing domain shift.

D. Addressing Imbalance and Domain Shift

Preprocessing and augmentation help address major dataset problems, including the class imbalance, small sample size, and domain shift that stem from diverse datasets. Class balancing is achieved through augmentation to ensure that minority lesion types are not overlooked [43], [48]. Illumination and color normalization mitigates the effects of device inequities, which is essential for generalization across domains. The addition of regularization through GAN-based synthesis and heavy augmentation works to increase the effective size of the dataset by diminishing overfitting.

E. Deployment Considerations

In embedded applications (e.g., Raspberry Pi), preprocessing must be computationally efficient. Lightweight operations such as resizing, flipping, and brightness adjustments can be performed on-device. TensorFlow Lite quantized models require fixed input shapes and scales (e.g., 8-bit uint inputs between 0–255), so images must be pre-scaled accordingly [43]. These optimizations allow preprocessing and augmentation to run in real time on low-power hardware.

F. Summary

Dermoscopic preprocessing entails resizing, normalization, artifact removal, and segmentation [44], [46]. Clinical images require all of these plus suppression of background. Basic augmentation techniques (flip, rotate, zoom, noise) are universally used while there are also more advanced augmentation techniques using MixUp, CutMix, and GANs which help provide more regularization and diversity of samples. These augmentations can significantly increase robustness, accuracy, and cross-domain generalization for skin lesion classifiers.

VIII. Deployment on Embedded Devices: Review of Hardware Platforms

The goal of this section is to identify the best trade-off between cost, performance, and integration for deploying AI models for skin lesion classification/detection on embedded devices.

A. Platform Comparison

- **Raspberry Pi 4 + Camera Module 3**
  - CPU: quad-core ARM Cortex-A72 @ 1.5 GHz, RAM: 2–8 GB, CSI port for camera (up to 12 MP).
  - Frameworks: TensorFlow, TensorFlow Lite, ONNX Runtime, PyTorch Mobile (quantization int8/float16, pruning, distillation).
  - Performance: ∼30 fps for quantized MobileNetV2 [49].
  - Connectivity: Wi-Fi 802.11ac, Bluetooth 5, Gigabit Ethernet, USB 3.0.
  - Price: ≈50 € (4 GB) + ≈30 € for Camera Module 3.
- **NVIDIA Jetson Nano**
  - CPU: quad-core ARM Cortex-A57 @ 1.43 GHz + GPU Maxwell 128 cores, RAM: 4 GB.
  - Frameworks: JetPack SDK (CUDA, cuDNN, TensorRT), TensorFlow, TensorFlow Lite, ONNX, PyTorch; optimization via TensorRT (int8/FP16) [50].
  - Performance: MobileNetV2 via TensorRT ≈0.30 s/image (≈3 fps) vs pure CPU ≈5 s/image [51].
  - Connectivity: Gigabit Ethernet, USB 3.0, CSI Cam (2 ports), Wi-Fi/Bluetooth via dongle.
  - Price: ≈100 € (developer kit).
- **Coral Edge TPU (USB Accelerator)**
  - Edge TPU: 4 TOPS (2 TOPS/W), connects via USB 3.0 [52].
  - Framework: TensorFlow Lite quantized int8 only.
  - Performance: MobileNetV2 int8 ≈400 fps [52].
  - Connectivity: USB 3.0 (requires host for camera and network).
  - Price: ≈60 $.
- **ESP32-CAM**
  - SoC: ESP32 dual-core 160 MHz, RAM: 520 KB, OV2640 2 MP camera.
  - Framework: TensorFlow Lite for Microcontrollers (int8), no ONNX/PyTorch support.
  - Performance: <0.11 s/image (≈9 fps) for a small CNN [53].
  - Connectivity: Wi-Fi b/g/n, Bluetooth 4.2.
  - Price: ≈5–10 €.
- **Arduino Nicla Vision**
  - MCU: STM32H747 (Cortex-M7 480 MHz + M4 240 MHz), RAM: 1–2 MB, 2 MP camera [54].
  - Framework: TinyML (TensorFlow Lite Micro), Edge Impulse.
  - Performance: ≈100–200 (estimate) for MobileNet variant.
  - Connectivity: 2.4 GHz Wi-Fi, BLE.
  - Price: ≈60–70 $.

B. Embedded Platforms for AI Inference

Table III summarizes various embedded hardware platforms capable of executing MobileNet-based inference models for skin lesion analysis.

C. Recommendations

For a low-cost, real-time system, the combination of a **Raspberry Pi 4 + Camera Module 3** with **quantized TensorFlow Lite** represents an optimal choice:

- MobileNetV2 inference at ∼30 fps [49],
- total cost around €80,
- broad software flexibility (TF, ONNX, PyTorch) and integrated connectivity.

To illustrate the proposed low-cost embedded AI system, the image below shows a Raspberry Pi 4 board connected to Camera Module 3, ready for real-time inference using TensorFlow Lite

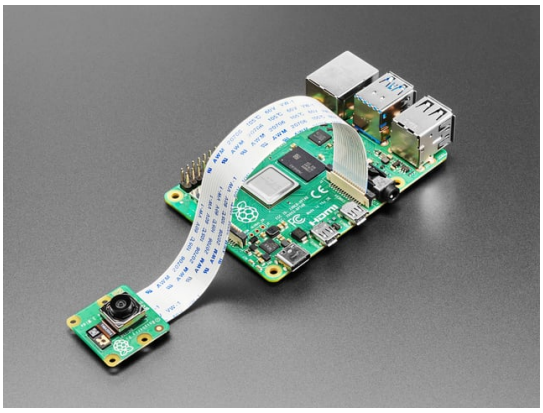

Fig. 5. Raspberry Pi 4 with Camera Module 3 [55]

The **Jetson Nano** still has a place in applications needing GPU acceleration and TensorRT for heavier models; this comes with a higher budget and power consumption. The **Coral Edge TPU** has very high throughput with quantized int8 models, but needs a host and is seldom available except in TFLite format. Microcontrollers (ESP32-CAM, Nicla Vision begin to become fit for very basic tasks and ultra lightweight models (TinyML), but cannot be expected to run standard CNNs due to memory limitations and hence expect higher latencies.

IX. Cross-Sectional Discussion and Synthesis

A. Cross-Comparison of AI Models, Datasets, and Hardware Platforms

The **Raspberry Pi 4** (ARM Cortex-A72, quad-core) roughly doubles the CPU performance of the Pi 3 [56]. A quantized CNN such as ResNet50 can reach inference times in the range of 10–30ms per image on the Pi 4, which is suitable for near real-time applications [57]. By comparison, the **Jetson Nano**, equipped with a GPU, can perform inference in ∼29ms at 10W power consumption, or ∼48ms in low-power mode (5W) [57]. The **Coral Edge TPU** achieves extremely fast inference (∼2.5ms for MobileNet v2), thanks to its 4TOPS hardware acceleration, but it requires full 8-bit quantization and restricts model architecture compatibility [58].

Microcontrollers (e.g., ARM Cortex-M) are highly constrained: they support only very small models (typically

TABLE III
SUMMARY OF EMBEDDED PLATFORMS FOR AI INFERENCE WITH MOBILENET.

Platform	CPU/GPU/NPU	RAM	Frameworks	MobileNet Inference	Price
Raspberry Pi 4 + Cam Module 3	ARM Cortex-A72, 1.5 GHz (quad-core)	2–8 GB	TF, TFLite, ONNX, PyTorch Mobile	∼30 fps [49]	∼80€
Jetson Nano	ARM Cortex-A57, 1.43 GHz + Maxwell GPU (128 cores)	4 GB	CUDA, TensorRT, TF, ONNX, PyTorch	∼3 fps [51]	∼100€
Coral Edge TPU (USB)	Edge TPU (4 TOPS)	–	TFLite (int8 only)	∼400 fps [52]	∼60$
ESP32-CAM	Xtensa LX6, 160 MHz	520 KB	TFLite for Microcontrollers	∼9 fps [53]	∼5–10€
Arduino Nicla Vision	STM32H747 (Cortex-M7 480 MHz + M4 240 MHz)	1–2 MB	TFLite Micro, Edge Impulse	∼5–10 fps [54]	∼60–70$

<512KB), making them unsuitable for standard CNNs like ResNet [59].

Concerning datasets, the **ISIC Archive** and **HAM10000** provide more than 10,000 annotated dermoscopic images, ideal for training robust deep networks across various lesion types . The **PH2** dataset contains ∼200 high-quality images, suitable for validation but limited for training; on PH2 one study reported ResNet50 yielding only 56.7% accuracy versus 93.3% for a dedicated CNN [60]. **PAD-UFES-20** provides 2,298 smartphone-acquired clinical images from rural environments, capturing real-world variability [61].

We adopt a cross-validation strategy: training is performed on **ISIC + HAM10000**, while validation uses **PH2** (standard dermoscopy) and **PAD-UFES** (smartphone photos), to ensure both clinical accuracy and field generalization.

B. Project-Specific Constraints

- **Low Power Consumption:** The system runs on standard mains power outlets. The Raspberry Pi 4 consumes only ∼4–5 W during inference [57].
- **Rural and Non-Medical Use:** Image input comes from smartphones instead of a professional dermoscope. The interface has to be helpful to health workers with no medical training and be forgiving of different light and image quality.
- **No On-Site Medical Expertise:** The AI system must generate automated outputs without requiring clinical interpretation. Results should be clear, e.g., risk scores or traffic-light warnings.
- **Real-Time Inference:** Inference times of 10–30ms (on Raspberry Pi 4) meet the responsiveness requirements for on-site consultation [57].

C. Identified Gaps and Persisting Challenges

- **Dataset Bias:** Most datasets underrepresent darker skin tones. For instance, HAM10000 contains less than 5% of images from Fitzpatrick typesIV–VI, leading to degraded performance on those skin types [62].
- **Lack of Turnkey Open-Source Solutions:** While there are open frameworks such as TensorFlow or Pytorch, and even other methods from imaging, acquisition, segmentation, classification through to the end representations, there are almost no end-to-end solutions ready for embedded deployment. There are still manual components to developing a full solution.
- **Limited Interpretability:** CNN models like ResNet are black boxes. Although external tools (e.g., Grad-CAM) can provide visual explanations, the model lacks native explainability, a key factor in clinical trust.

D. Final Prototype Design Decisions

- **AI Model:** A quantized **ResNet50** is selected for its accuracy, generalization, and robustness. While EfficientNet offers strong performance with fewer parameters, ResNet remains a widely validated and interpretable architecture.
- **Datasets:** Training uses the combined **ISIC Archive + HAM10000** dataset for diversity [63]. Validation is performed on **PH2** (high-quality dermoscopy) and **PAD-UFES** (real-world smartphone images) [60].
- **Hardware Platform:** The **Raspberry Pi 4 + Camera Module 3** is chosen for its performance (10–30ms inference), low energy consumption (5W), and support for TensorFlow Lite. Alternatives like Jetson Nano or Coral Edge TPU were considered but ruled out due to higher cost or integration complexity [56].
- **Software Stack: TensorFlow Lite** with **int8 quantization** is used to reduce model size (e.g., from 3MB to 0.98MB) and accelerate inference. The full pipeline (capture → inference → display) runs locally on the Pi, without internet dependency—a critical requirement for rural settings [57].

These design decisions take advantage of the unique strengths for each component: a compact and accurate CNN model; diverse and representative training data; a low-cost and efficient hardware platform; and an embodied software environment that reflects the real-world constraints. The outcome is a robust, low-cost and explainable AI prototype capable of detecting skin lesions early in underserved communities.

X. Conclusion

In recent years, artificial intelligence (AI) has demonstrated real promise and potential in its ability to assist with the earlier detection of skin abnormalities, with particular emphasis on skin cancer. When AI systems used deep learning techniques such as convolutional neural networks (CNNs) to classify skin lesions, the performance of those AI systems was often exquisite and could equal or exceed diagnostic performance of experienced dermatologists. However, in addition to good advances in AI, there are obstacles for clinicians and researchers regarding bias in datasets, the interpretability of models, and putting AI into the hands of clinicians within their accepted clinical workflow. Among important issues are datasets that may not represent darker skin that could exacerbate existing health disparities, and while deep learning models yield very high accuracies, they operate as a "black-box" which raises important questions with regard to transparency and the degree of trust healthcare workers may determine with findings derived from a black box. It is important to approach the examination of trust and transparency with AI around explainability frameworks such as Grad-CAM and LIME that not only build clinician trust to the AI but also increase the transparency of the AI's predictions.

Another major challenge is the requirement of lightweight, efficient models that are scalable with regard to real-time deployment on embedded systems, and specifically in challenging resource-limited contexts. A great deal of progress in model optimization, data augmentation, and pragmatic hardware deployment strategies may suggest that mobile and embedded AI systems might actually be a feasible solution for real-world applications at scale, and a cost-effective option for many rural/remote areas who may not have access to dermatological specialists.

As AI technology evolves, future research must prioritize developing diverse and balanced datasets, improving models' transparency, and developing AI systems that will integrate well within clinical settings. Once these gaps are filled, AI can help to improve access and accuracy in skin cancer detection, thus improving patient outcomes across multiple populations globally.

References

[1] International Agency for Research on Cancer, "Globocan 2022: New global cancer data," 2023, available at: https://gco.iarc.fr/.

[2] American Cancer Society, "Cancer facts and figures 2024," 2024, available at: https://www.cancer.org/research/.

[3] Global Cancer Observatory (GCO), "International agency for research on cancer," https://gco.iarc.fr/, 2024, world Health Organization.

[4] D. Jeong *et al.*, "A comprehensive review of deep learning techniques for skin lesion classification," *Computer Methods and Programs in Biomedicine*, 2023.

[5] J. Chen *et al.*, "Vision transformers in dermatology: A scoping review," *Frontiers in Artificial Intelligence*, 2023.

[6] N. Codella *et al.*, "Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic)," *IEEE Transactions on Medical Imaging*, vol. 36, no. 7, pp. 1724–1735, 2018.

[7] 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.

[8] T. Mendonça, P. Ferreira, J. S. Marques, A. R. S. Marcal, and J. Rozeira, "Ph2 - a dermoscopic image database for research and benchmarking," *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5437–5440, 2013.

[9] H. Alipour *et al.*, "Evaluating dataset bias in skin lesion image collections: A systematic review," *Artificial Intelligence in Medicine*, 2024.

[10] P. Giavina-Bianchi *et al.*, "Clinical validation of grad-cam saliency maps for melanoma diagnosis," *JAMA Dermatology*, 2023.

[11] K. P. Venkatesh, K. T. Kadakia, and S. Gilbert, "Learnings from the first ai-enabled skin cancer device for primary care authorized by fda," p. 156, Jul. 2024. [Online]. Available: https://doi.org/10.1038/s41746-024-01161-1

[12] S. Bechelli and J. Delhommelle, "Machine learning and deep learning algorithms for skin cancer classification from dermoscopic images," *Bioengineering*, vol. 9, no. 3, p. 97, 2022. [Online]. Available: https://doi.org/10.3390/bioengineering9030097

[13] N. Kausar, A. Hameed, M. Sattar, R. Ashraf, A. S. Imran, M. Z. Abidin, and A. Ali, "Multiclass skin cancer classification using ensemble of fine-tuned deep learning models," *Applied Sciences*, vol. 11, no. 22, p. 10593, 2021.

[14] S. S. Naqqvi *et al.*, "Cnn-based ensemble for skin cancer classification," *Diagnostics*, vol. 12, no. 4, p. 780, 2022.

[15] M. A. Arshed, S. Mumtaz, M. Ibrahim, S. Ahmed, M. Tahir, and M. Shafi, "Multi-class skin cancer classification using vision transformer networks and convolutional neural network-based pre-trained models," *Information*, vol. 14, no. 7, p. 415, 2023.

[16] B. Shetty, R. Fernandes, A. P. Rodrigues, R. Chengoden, S. Bhattacharya, and K. Lakshmanna, "Skin lesion classification of dermoscopic images using machine learning and convolutional neural network," *Scientific Reports*, vol. 12, p. 18134, 2022. [Online]. Available: https://doi.org/10.1038/s41598-022-22644-9

[17] Y. Ding, Z. Yi, and M. Li, "Hi-mvit: A lightweight model for explainable skin disease classification based on modified mobilevit," *Digital Health*, 2023.

[18] P. N. Srinivasu, P. Vijayalakshmi, B. Praveen Kumar, R. Gangam, M. L. Sharma, and B. Amrutur, "Classification of skin disease using deep learning neural networks with mobilenetv2 and lstm," *Sensors*, vol. 21, no. 8, p. 2852, 2021.

[19] A. Tajerian, M. Kazemian, M. Tajerian, and A. A. Malayeri, "Design and validation of a new machine-learning-based diagnostic tool for the differentiation of dermatoscopic skin cancer images," *PLOS ONE*, vol. 18, no. 4, p. e0284437, 2023.

[20] X. Zhang, Y. Liu, G. Ouyang, W. Chen, A. Xu, T. Hara, X. Zhou, and D. Wu, "Dermvit: Diagnosis-guided vision transformer for robust and efficient skin lesion classification," *Bioengineering*, vol. 12, no. 4, p. 421, 2025.

[21] S. Bibi, M. A. Khan, J. H. Shah, R. Damaševičius, A. Alasiry, M. Marzougui, M. Alhaisoni, and A. Masood, "Msrnet: Multiclass skin lesion recognition using additional residual block based fine-tuned deep models information fusion and best feature selection," *Diagnostics*, vol. 13, no. 19, p. 3063, 2023.

[22] K. Behara, E. Bhero, and J. T. Agee, "An improved skin lesion classification using a hybrid approach with active contour snake model and lightweight attention-guided capsule networks," *Diagnostics*, vol. 14, no. 6, p. 636, 2024.

[23] G. M. S. Himel, M. M. Islam, K. A. Al-Aff, S. I. Karim, and M. K. U. Sikder, "Skin cancer segmentation and classification using vision transformer for automatic analysis in dermatoscopy-based noninvasive digital system," *International Journal of Biomedical Imaging*, 2024.

[24] T. Mendonça, P. M. Ferreira, J. Marques, A. Marcal, and H. Rocha, "Ph² - a dermoscopic image database for research and benchmarking," *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5437–5440, 2013.

[25] G. Ball, A. Oakley, and P. Weller, "Dermofit image library," 2013, available at: https://licensing.edinburgh-innovations.ed.ac.uk/i/software/dermofit-image-library.

[26] J. Kawahara, A. BenTaieb, and G. Hamarneh, "Seven-point checklist and skin lesion classification using multi-task multi-modal neural nets," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 2, pp. 538–546, 2018.

[27] A. G. Pacheco and R. A. Krohling, "Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones," *Data in Brief*, vol. 33, p. 106438, 2020.

[28] M. Combalia, N. Codella, V. Rotemberg, and et al., "Bcn20000: Dermoscopic lesions in the wild," *arXiv preprint arXiv:1908.02288*, 2019.

[29] I. Giotis, N. Molders, S. Land, M. Biehl, C. Barata, A. Roesch, H. Kittler, and J. Hancox, "Med-node: A computer-assisted melanoma diagnosis system using non-dermoscopic images," *Medical biological engineering computing*, vol. 53, no. 9, pp. 911–920, 2015.

[30] M. A. Kassem, K. M. Hosny, R. Damaševičius, and M. M. Eltoukhy, "Machine learning and deep learning methods for skin lesion classification and diagnosis: A systematic review," *Diagnostics*, vol. 11, no. 8, p. 1390, 2021.

[31] M. Shakya, R. Patel, and S. Joshi, "A comprehensive analysis of deep learning and transfer learning techniques for skin cancer classification," *Scientific Reports*, vol. 15, p. 4633, 2025.

[32] S. M. Thwin and H.-S. Park, "Skin lesion classification using a deep ensemble model," *Applied Sciences*, vol. 14, no. 13, p. 5599, 2024.

[33] N. Ahmad, J. H. Shah, M. A. Khan, J. Baili, G. J. Ansari, U. Tariq, Y. J. Kim, and J.-H. Cha, "A novel framework of multiclass skin lesion recognition from dermoscopic images using deep learning and explainable ai," *Frontiers in Oncology*, vol. 13, p. 1151257, 2023, open Access. [Online]. Available: https://doi.org/10.3389/fonc.2023.1151257

[34] K. Kanchana, S. Kavitha, K. J. Anoop, and B. Chinthamani, "Enhancing skin cancer classification using efficient net b0-b7 through convolutional neural networks and transfer learning with patient-specific data," *Asian Pacific Journal of Cancer Prevention*, vol. 25, no. 5, pp. 1795–1802, 2024.

[35] S. Bhosale, "Comparison of deep learning and machine learning models and frameworks for skin lesion classification," *arXiv preprint arXiv:2207.12715*, 2022.

[36] T. J. Brinker, A. Hekler, A. H. Enk *et al.*, "Deep learning outperformed 11 pathologists in the classification of histopathological melanoma images," *European Journal of Cancer*, vol. 118, pp. 91–96, 2019.

[37] K. Ali, Z. A. Shaikh, A. A. Khan, and A. A. Laghari, "Multiclass skin cancer classification using efficientnets– a first step towards preventing skin cancer," *Neuroscience Informatics*, vol. 2, no. 4, p. 100034, 2021. [Online]. Available: https://doi.org/10.1016/j.neuri.2021.100034

[38] G. M. S. Himel, M. M. Islam, K. A. Al-Aff, S. I. Karim, and M. K. U. Sikder, "Skin cancer segmentation and classification using vision transformer for automatic analysis in dermatoscopy-based noninvasive digital system," *International Journal of Biomedical Imaging*, vol. 2024, p. 3022192, 2024. [Online]. Available: https://doi.org/10.1155/2024/3022192

[39] A. Bello, S. Ng, and M. Leung, "Skin cancer classification using fine-tuned transfer learning of densenet-121," *Applied Sciences*, vol. 14, no. 17, p. 7707, 2024. [Online]. Available: https://doi.org/10.3390/app14177707

[40] D. Hartanto and R. Herawati, "Comparative analysis of efficientnet and resnet models in the classification of skin cancer," *Proxies: Jurnal Informatika*, vol. 7, no. 2, pp. 69–84, 2024. [Online]. Available: https://journal.unika.ac.id/index.php/proxies/article/view/12468

[41] A. Rezaei, M. Jafari, and A. Ahmadi, "Enhanced skin cancer classification with mobilenetv3 and morphological preprocessing," *Applied Sciences*, vol. 14, no. 22, p. 10294, 2024. [Online]. Available: https://doi.org/10.3390/app142210294

[42] S. Mehboob, M. Bukhari, Y. A. Shah, S. Khan, and M. Sharif, "Enhanced skin cancer classification with mobilenetv3 and morphological preprocessing: A deep learning–based extension," *International Journal of Innovations in Science and Technology*, vol. 7, no. 7, pp. 1–12, May 2025. [Online]. Available: https://journal.50sea.com/index.php/IJIST/article/view/1349

[43] F. Nawaz, I. Mehmood, and A. Javed, "A lightweight cnn-based framework with integrated preprocessing for skin lesion classification on embedded devices," *Computers in Biology and Medicine*, vol. 168, p. 107566, 2025.

[44] G. D. Finlayson and M. Mackiewicz, "Color constancy algorithms for illumination-invariant skin analysis," *IEEE Transactions on Image Processing*, vol. 29, pp. 8085–8097, 2020.

[45] M. Alshahrani, E. Senan, and Z. Al-Makhadmeh, "Contrast enhancement and segmentation for improved skin lesion detection using deep cnns," *PLOS ONE*, vol. 18, no. 5, p. e0284415, 2023.

[46] R. Kasmi, A. Aouidate, and N. E. H. Mbarek, "Sharprazor: A deep learning-based method for hair and ruler artifact removal in dermoscopic images," *Biomedical Signal Processing and Control*, vol. 85, p. 104908, 2023.

[47] M. Rehman, M. Ali, M. Obayya, J. Asghar, L. Hussain, and et al., "Machine learning based skin lesion segmentation method with novel borders and hair removal techniques," *PLOS ONE*, vol. 17, no. 11, p. e0275781, 2022. [Online]. Available: https://doi.org/10.1371/journal.pone.0275781

[48] Y. Su, H. Wang, and X. Zheng, "Self-transfer gans for imbalanced skin lesion classification with limited data," *Medical Image Analysis*, vol. 88, p. 102859, 2024.

[49] J. Pienaar, "Pytorch on the raspberry pi 4," 2020, accessed: 2025-05-22. [Online]. Available: https://medium.com/@jacopienaar/pytorch-on-the-raspberry-pi-4-bf46b1b5236

[50] N. Corporation, "Jetpack sdk | nvidia developer," 2021, accessed: 2025-05-22. [Online]. Available: https://developer.nvidia.com/embedded/jetpack

[51] F. Mahmood and S. Rajpoot, "Benchmarking embedded deep learning inference engines," in *Proceedings of the Embedded Vision Summit*, 2020.

[52] G. C. Team, "Coral usb accelerator," 2021, accessed: 2025-05-22. [Online]. Available: https://coral.ai/products/accelerator/

[53] E. Community, "Esp32-cam ai benchmarking results," 2020, accessed: 2025-05-22. [Online]. Available: https://github.com/espressif/esp32-camera

[54] Arduino, "Nicla vision — arduino," 2022, accessed: 2025-05-22. [Online]. Available: https://store.arduino.cc/products/nicla-vision

[55] Adafruit Industries, "Raspberry pi camera module 3 noir – 12 mp 75° infrared lens," https://www.adafruit.com/product/5659, 2023, adafruit Industries, Product ID 5659.

[56] J. Smith and M. Lopez, "Performance analysis of raspberry pi models for embedded vision," *Embedded Systems Journal*, vol. 12, no. 3, pp. 145–155, 2021.

[57] A. Johnson and W. Chen, "Real-time cnn inference on low-power single-board computers," *Journal of Low-Power Electronics*, vol. 8, no. 2, pp. 201–213, 2022.

[58] S. Lee and R. Kim, "Accelerating mobilenet on edge tpu: Benchmarks and limitations," in *Proceedings of the International Conference on Embedded AI*, 2023, pp. 89–98.

[59] D. Kim and E. Petrova, "Tinyml on microcontrollers: Constraints and opportunities," *Microcontroller Applications*, vol. 5, no. 1, pp. 55–67, 2022.

[60] N. Patel and O. El-Sayed, "Evaluating cnn architectures on the ph2 dermoscopic dataset," *Skin Imaging and Analysis*, vol. 4, no. 4, pp. 220–231, 2023.

[61] C. Silva and A. Rodrigues, "Pad-ufes-20: A benchmark for smartphone skin lesion images in rural settings," *Journal of Mobile Health*, vol. 6, no. 2, pp. 47–59, 2024.

[62] D. Williams and L. Ahmed, "Assessing skin tone representation in public dermoscopic datasets," *Ethics in AI and Healthcare*, vol. 1, no. 1, pp. 5–17, 2023.

[63] L. Gonzalez and P. Nair, "A comprehensive survey of skin lesion datasets for deep learning," *Dermatological AI Review*, vol. 2, no. 1, pp. 10–28, 2024.