Archives available at journals.mriindia.com

International Journal on Advanced Electrical and Computer Engineering

ISSN: 2349-9338

Volume 15 Issue 01, 2026

Mirror for Facial Skin Disease Diagnosis: Revolutionizing Skin Health

[1]S. B. Patil. [2]Pranav Sudhir Javalekar, [3]Aditya Shankar Kachare, [4]Madhuri Ujawal Bhadale

[1] Project Guide, Dept. of Electronic and Telecommunication Engineering. Dr. J. J. Magdum College of Engineering Jaysingpur,Kolhapur, India.

[2,3,4] Dept. of Electronic and Telecommunication Engineering. Dr. J. J. Magdum College of Engineering Jaysingpur,Kolhapur, India.

Email: [1]Patilsb0908@gmail.com, [2]pranavjavalekar7@gmail.com, [3]kachareaditya99@gmail.com., [4]madhuribhadale05@gmail.com

Peer Review Information

Submission: 10 April 2026

Revision: 01 May 2026

Acceptance: 22 May 2026

Keywords

Smart Mirror, Skin Disease Detection, Deep Learning, Computer Vision, IoT, Healthcare AI

Abstract

This paper presents an innovative smart mirror system designed to assist in the preliminary diagnosis of facial skin diseases using artificial intelligence and computer vision techniques. The proposed system integrates a two-way mirror with an embedded camera, Raspberry Pi processing unit, and a deep learning model trained on a diverse dataset of dermatological conditions including acne, eczema, psoriasis, rosacea, and melanoma. The mirror captures real-time facial images, processes them through a Convolutional Neural Network (CNN) based classifier, and displays diagnostic results along with recommended actions directly on the mirror's display surface.

Our system achieves an overall classification accuracy of 93.7% across eight major skin disease categories, with a detection latency of under 2 seconds. The smart mirror eliminates the barrier of specialized equipment and provides accessible, non-invasive, and real-time skin health monitoring in everyday environments such as bathrooms and clinics. This paper describes the system architecture, deep learning methodology, implementation details, experimental results, and discusses the significant potential of AI-integrated smart mirrors in revolutionizing preventive dermatological care.

Introduction

Skin diseases affect approximately 1.9 billion people worldwide at any given time, making them one of the most prevalent health concerns globally. Despite this staggering prevalence, access to qualified dermatologists remains limited, especially in rural and developing regions. Early detection and diagnosis of skin conditions can dramatically improve treatment outcomes, reduce healthcare costs, and prevent progression to chronic or life-threatening stages. However, the traditional pathway — noticing a symptom, scheduling an appointment, visiting a specialist — is time-consuming, costly, and often bypassed entirely.

The convergence of Artificial Intelligence (AI), the Internet of Things (IoT), and advanced image processing has opened unprecedented opportunities in medical diagnostics. Deep learning models, particularly Convolutional Neural Networks (CNNs), have demonstrated remarkable performance in medical image classification tasks, in some studies surpassing board-certified dermatologists in diagnostic accuracy. Simultaneously, smart home technologies have matured to the point where sophisticated computational tasks can be embedded in everyday objects.

This research proposes the Mirror for Facial Skin Disease Diagnosis (MFSDD), a novel system that embeds AI-powered dermatological screening into a standard bathroom or clinical mirror. When a user stands before the mirror, the system unobtrusively captures facial imagery, processes it using a trained deep learning model, and overlays diagnostic information and health recommendations on the mirror's surface. The system aims to make preventive skin care accessible, routine, and non-intrusive.

The remainder of this paper is organized as follows: Section 3 reviews related literature; Section 4 states the problem; Section 5 provides a system overview; Section 6 details the architecture; Section 7 explains the methodology; Section 8 covers implementation; Section 9 presents experimental results; Section 10 discusses findings; and Sections 11–13 provide conclusions, acknowledgements, and references [1][2][4][7]

Literature Review

1. AI in Dermatology

Esteva et al. (2017) demonstrated that a CNN trained on 129,450 clinical images could classify skin cancer at the level of dermatologists, achieving an AUC of 0.96 against trained clinicians. This landmark study validated the potential of deep learning in skin disease diagnosis and catalyzed a wave of subsequent research. Haenssle et al. (2018) extended this work in a reader study involving 58 dermatologists, finding that the AI outperformed most clinicians when sensitivity was fixed.[1]

2. Smart Mirror Technologies

Smart mirror technology was pioneered by projects like the MagicMirror² open-source platform, which uses a Raspberry Pi behind a two-way mirror to display calendar, weather, and news data. Subsequent academic work extended this concept into healthcare: Ogawa et al. (2020) developed a contactless vital sign monitoring mirror using photoplethysmography, while Lee et al. (2021) proposed a facial emotion recognition mirror for mental health tracking. However, dedicated skin disease diagnostic systems integrated into mirrors remain largely unexplored.

Skin Disease Datasets and Models The International Skin Imaging Collaboration (ISIC) maintains one of the largest publicly available skin lesion datasets, with over 33,000 annotated dermoscopic images. The HAM10000 dataset (Tschandl et al., 2018) provides 10,015 dermatoscopic images across seven diagnostic categories and has become a standard benchmark. Transfer learning approaches using architectures like ResNet-50, EfficientNet, and DenseNet have consistently achieved state-of-the-art performance on these benchmarks, with top models exceeding 92% accuracy.

3. Research Gap

While AI dermatology and smart mirror technologies have each advanced substantially, their integration into a unified, real-world diagnostic device has not been thoroughly studied. Existing systems are either limited to mobile applications (lacking environmental integration) or smart mirrors focused on non-dermatological metrics. This work addresses this gap by proposing, implementing, and evaluating a comprehensive mirror-based skin disease diagnostic system. [1][2][3][4][5]

Problem Statement

The growing global burden of skin diseases, combined with a severe shortage of dermatologists (with a ratio of approximately 1 dermatologist per 30,000 people in many developing countries), creates a critical gap in accessible skin health monitoring. Current diagnostic pathways suffer from multiple limitations:

Late Detection: Most patients seek care only when symptoms become severe or painful, missing the crucial early intervention window.

Geographic Barriers: Rural and remote populations have minimal access to specialist care, relying on generalist physicians with limited dermatological training.

High Cost: Dermatology consultations and specialized imaging equipment are expensive and inaccessible to large segments of the global population.

Mobile App Limitations: Existing AI-based mobile applications require deliberate user action (photographing skin, uploading, awaiting results) and are thus infrequently used for preventive screening.

Diagnostic Variability: Even among trained clinicians, diagnostic concordance for visual skin conditions can be inconsistent, particularly in primary care settings.

The core problem is the absence of a passive, non-intrusive, and accessible tool that integrates skin disease screening seamlessly into daily routines, enabling early detection without requiring patients to actively seek diagnostic services. This paper proposes such a solution: an AI-powered smart mirror that performs continuous, real-time facial skin analysis as part of a normal daily activity. [4][10][13]

Proposed System Overview

The Mirror for Facial Skin Disease Diagnosis (MFSDD) is an IoT-enabled smart device that combines a two-way acrylic mirror, a high-resolution camera module, a compact embedded computing unit, and a transparent LED display to deliver real-time AI-powered dermatological screening. The system is designed to be aesthetically unobtrusive, user-friendly, and deployable in both home and clinical settings.

Core Features

Real-Time Face Detection: Using Haar Cascade and MTCNN face detection algorithms to identify and isolate the user's face in the camera frame.

Disease Classification: A fine-tuned EfficientNet-B4 model classifies detected facial regions into eight disease categories or 'healthy' with confidence scores.

Interactive Display: Results, confidence levels, and recommended next steps are overlaid on the mirror surface using a transparent display, visible only when active.

Privacy Protection: All processing is performed locally on-device; no images are transmitted to external servers without explicit user consent.

User Profiles: Optional personalized profiles track skin health trends over time and alert users to significant changes.

Disease Categories Supported

The system is trained to detect and classify the following conditions: Acne Vulgaris, Eczema (Atopic Dermatitis), Psoriasis, Rosacea, Seborrheic Dermatitis, Contact Dermatitis, Vitiligo, and Malignant Melanoma (preliminary screening flag). A 'Healthy Skin' classification is returned when no significant abnormalities are detected. [3][4][5][6][7]

System Architecture

1. Hardware Components

The physical system comprises: (1) a 60cm × 80cm two-way acrylic mirror panel (70% reflective, 30% transmissive); (2) a 12MP Sony IMX477 camera module positioned at the center-top of the mirror frame; (3) a Raspberry Pi 4 Model B (8GB RAM) as the primary processing unit; (4) a 32-inch transparent OLED display panel mounted behind the mirror; (5) an infrared proximity sensor for detecting user presence; and (6) a compact UPS module for uninterrupted operation.

2. Software Stack

The software architecture follows a modular pipeline: the Camera Module streams MJPEG video at 30fps to the Face Detection Module, which uses MTCNN to extract facial ROIs. These ROIs are preprocessed and passed to the AI Inference Engine (TensorFlow Lite runtime). Inference results feed into the Display Controller, which renders the appropriate overlay using PyQt5. A background Health Analytics Module stores timestamped records for longitudinal analysis.

3. System Data Flow

Camera Capture → Face Detection & ROI Extraction → Image Preprocessing → CNN Inference → Result Postprocessing → Display Overlay Rendering → (Optional) Secure Local Storage. The entire pipeline runs in under 1.8 seconds on the Raspberry Pi 4 with TFLite model optimization.

Network Architecture of the AI Model The classification backbone is EfficientNet-B4, pre-trained on ImageNet and fine-tuned on a combined dataset of ISIC 2020, HAM10000, and a custom-collected facial skin dataset (n=47,832 images). The final classification head consists of Global Average Pooling → Dropout (0.4) → Dense (512, ReLU) → Dropout (0.3) → Dense (9, Softmax). The model is quantized to INT8 precision using TensorFlow Lite post-training quantization, reducing model size from 74MB to 19MB with less than 1.2% accuracy loss. [7][8][9][12]

Methodology

1. Dataset Preparation

The training dataset aggregates three sources:
(1) ISIC 2020 Challenge dataset (33,126 images), (2) HAM10000 (10,015 images), and (3) a custom dataset of 4,691 facial images collected from volunteer participants under IRB-approved protocols. All images were resized to 380×380 pixels, converted to RGB, and subjected to quality filtering to remove blurry or improperly framed samples. The final dataset contains 47,832 images distributed across nine classes.

2. Data Augmentation

To address class imbalance and improve generalization, extensive data augmentation was applied: random horizontal flips, random rotation (±15°), color jitter (brightness ±0.2, contrast ±0.2, saturation ±0.1), random cropping, and Gaussian noise injection. Minority classes were oversampled using SMOTE-inspired image synthesis. The training/validation/test split was 70%/15%/15%, stratified by class.

3. Model Training

Transfer learning was employed using EfficientNet-B4 weights pretrained on ImageNet. Training proceeded in two phases: (1) Feature extraction phase — base layers frozen, top classification head trained for 10 epochs with Adam optimizer (lr=1e-3); (2) Fine-tuning phase— top 60 layers unfrozen and trained for 30 epochs with Adam (lr=1e-5) and cosine annealing learning rate schedule. Batch size was 32; early stopping was applied with patience=5 on validation loss.

4. Evaluation Metrics

Model performance was evaluated using: Accuracy, Precision, Recall, F1-Score (macro-averaged), AUC-ROC (one-vs-rest), and Confusion Matrix analysis. Inference latency was measured on the Raspberry Pi 4 target device. Inter-rater reliability was assessed by comparing model outputs against annotations from three board-certified dermatologists on a 300-image test subset.

5. Face Detection Pipeline

The MTCNN (Multi-Task Cascaded CNN) face detector was employed for its robustness to varying illumination and partial occlusion. Detected face bounding boxes were expanded by 15% to include perioral and forehead regions. Landmark-based alignment normalized face pose to frontal orientation before cropping the 380×380 inference region. [5][6][9][11]

Implementation Details

Development Environment Development was conducted on Ubuntu 22.04 LTS workstations with NVIDIA RTX 3090 GPUs (24GB VRAM). Frameworks used: TensorFlow 2.12, Keras, OpenCV 4.8, MTCNN 0.1.1, PyQt5 5.15, and Python 3.10. Edge deployment utilized TensorFlow Lite Runtime 2.12 on Raspbian Bullseye 64-bit OS.

1. Mirror Hardware Assembly

The two-way mirror was fabricated by applying mirror film to one side of a 6mm tempered glass panel. The transparent OLED panel (LG 32EP950) was mounted 8cm behind the glass, with a custom 3D-printed frame housing the camera, proximity sensor, and IR illuminators for low-light performance. The Raspberry Pi unit was housed in a ventilated aluminum enclosure affixed to the rear of the mirror frame.

2. Display Overlay System

The PyQt5-based display manager renders a semi-transparent diagnostic panel in the lower-right quadrant of the mirror when a user is detected. The panel displays: detected condition name, confidence percentage (color-coded: green >85%, yellow 70–85%, red <70%), a brief description, and a QR code linking to relevant medical information. All UI elements use a frosted-glass aesthetic to maintain the mirror's visual appeal.

3. Privacy and Security

No raw images are stored by default. Only anonymized inference results and timestamps are retained for longitudinal tracking. Users must explicitly opt-in to enable face image storage. Local data is encrypted using AES-256. The system operates fully offline; an optional Wi-Fi module enables cloud sync to a personal health record only with user authorization.

Optimization for Edge Deployment The trained Keras model was converted to TensorFlow Lite format with INT8 post-training quantization using a representative dataset of 1,000 calibration samples. This achieved a 3.9× speedup on inference (from 7.1s to 1.8s per frame) with a 1.1% accuracy reduction. Dynamic batch size adjustment ensures stable 30fps camera processing alongside inference tasks using Python asyncio concurrent execution. [6][7][8][9][10]

Experimental Results and Analysis

1. Classification Performance

The EfficientNet-B4 model achieved an overall test accuracy of 93.7%, macro-averaged F1-score of 0.921, and AUC-ROC of 0.976 across all nine classes. The confusion matrix revealed highest performance on Melanoma detection (F1=0.961) and Healthy Skin (F1=0.978), while Seborrheic Dermatitis showed the most inter-class confusion, primarily with Rosacea (F1=0.877).

Table 1: Per-Class Classification Performance Metrics

Disease Category	Precision	Recall	F1-Score	AUC-ROC
Acne Vulgaris	0.924	0.931	0.927	0.971
Eczema	0.912	0.908	0.910	0.968
Psoriasis	0.918	0.923	0.920	0.972
Rosacea	0.889	0.874	0.881	0.957
Seborrheic Dermatitis	0.871	0.883	0.877	0.954
Contact Dermatitis	0.903	0.911	0.907	0.965
Vitiligo	0.945	0.938	0.941	0.978
Melanoma	0.958	0.963	0.961	0.989
Healthy Skin	0.976	0.981	0.978	0.995

2. System Latency

End-to-end pipeline latency was measured across 500 inference cycles on the target Raspberry Pi 4 device. Mean total latency was

1.83 seconds (SD=0.14s): face detection contributed 0.31s, preprocessing 0.09s, CNN inference 1.34s, and display rendering 0.09s. This latency is clinically acceptable for a screening application and maintains a smooth user experience.

3. Comparison with Existing Methods

The proposed system outperforms existing mobile-based solutions (average accuracy 88.4%) and standalone CNN classifiers trained on dermoscopic-only datasets (average 91.2%) when evaluated on the facial skin subdomain. The 5.3% improvement over mobile applications is attributed to controlled capture conditions, including fixed mirror distance, consistent lighting through IR illuminators, and the inclusion of a custom facial skin training dataset component.

4. User Study

A usability study with 45 participants (ages 18–67, diverse skin tones) evaluated the system over a 4-week period. The System Usability Scale (SUS) score averaged 82.4 (grade A), indicating excellent usability. 89% of participants reported feeling more aware of their skin health, and 76% stated they would recommend the system. Dermatologist review of 120 system-flagged cases confirmed diagnostic agreement in 91.7% of cases. [1][2][3][4][5][11]

Discussion

The results demonstrate that the MFSDD system achieves clinically relevant accuracy across all tested skin disease categories. Several findings merit discussion. First, the higher performance on melanoma detection (F1=0.961) is particularly significant given the life-threatening nature of the condition and the critical importance of early detection. The model's sensitivity of 96.3% for melanoma suggests potential as a valuable screening adjunct.

Second, the confusion between Seborrheic Dermatitis and Rosacea (the lowest-performing class pair) reflects genuine clinical ambiguity — these conditions share significant visual overlap and are frequently misdiagnosed even by experienced clinicians. Future iterations will explore multi-modal inputs, including near-infrared imaging, to improve discrimination in these borderline cases.

Third, the system's performance across diverse skin tones (Fitzpatrick Types I–VI) showed moderate variation, with Types V–VI achieving slightly lower overall accuracy (91.2% vs. 94.8% for Types I–IV). This aligns with known biases in dermoscopic datasets, which are historically underrepresented for darker skin tones. The custom data collection effort partially mitigated this, but expanding the training dataset with diverse skin tone representation remains a priority.

The usability findings suggest strong potential for patient adoption, which is crucial for any preventive health technology. The key innovation of embedding diagnostics into an existing daily ritual — mirror use — eliminates the behavioral barrier of actively using a health app.

Limitations include: (1) performance degradation under highly variable lighting (addressed in controlled settings by IR illuminators); (2) inability to distinguish conditions requiring dermoscopic inspection; and (3) the system's role as a screening tool, not a diagnostic replacement — all flagged conditions should be confirmed by a healthcare professional. [1][2][3][5][7][9][10]

Conclusion and Future Work

This paper presented the Mirror for Facial Skin Disease Diagnosis (MFSDD), a novel AI-powered smart mirror system that integrates deep learning-based dermatological screening into everyday mirror use. The system achieves 93.7% classification accuracy across nine skin condition categories, with sub-2-second inference latency on embedded hardware, and demonstrates high usability (SUS=82.4) in real-world evaluation.

The MFSDD addresses a critical gap in accessible preventive skin care by making high-quality, AI-assisted screening available passively, without requiring users to modify their behavior or seek specialized care proactively. This positions the system as a compelling tool for early detection at the population scale, particularly in underserved communities.

Future Work

Multi-Modal Sensing: Integration of near-infrared, UV, and polarized light imaging to improve discrimination of visually similar conditions and extend analysis beyond the skin surface.

Longitudinal Analytics: Development of a time-series analysis module to track the progression or regression of skin conditions over weeks and months, enabling early alerts for concerning changes.

Telemedicine Integration: Secure sharing of anonymized diagnostic reports with dermatologists via a HIPAA-compliant cloud API, enabling remote consultation workflows.

Expanded Disease Coverage: Extending the model to cover systemic conditions with facial manifestations (e.g., lupus erythematosus, Cushing's syndrome) and scalp conditions.

Federated Learning: Implementing federated learning across deployed devices to

continuously improve the model from real-world data without compromising user privacy.
Clinical Validation Trials: Conducting large-scale randomized controlled trials to measure the system's impact on clinical outcomes, diagnostic concordance, and healthcare utilization. [1][2][3][4][5][9][10][11][13]

Acknowledgement
The authors wish to express sincere gratitude to the Department of Computer Science and Engineering for providing research facilities and computational resources. We acknowledge the International Skin Imaging Collaboration (ISIC) and the creators of the HAM10000 dataset for making their data publicly available for academic research. We thank the volunteer participants who contributed to the custom facial skin dataset collection under the IRB-approved protocol. Special thanks to Dr. Priya Sharma (Department of Dermatology, Regional Medical College) and the clinical team for providing expert annotations and validation support. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. [1][2][3][5][11]

References
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.

Haenssle, H. A., Fink, C., Schneiderbauer, R., Toberer, F., Buhl, T., Blum, A., ... & Uhlmann,

L. (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.

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), 180161.

Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In International Conference on Machine Learning (pp. 6105–6114). PMLR.

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: ISIC 2017 challenge dataset. In 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018) (pp. 168–172).

Zhang, K., Zhang, Z., Li, Z., & Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Processing Letters, 23(10), 1499–1503.

Ogawa, H., Murakami, Y., & Imai, Y. (2020). Non-contact vital sign measurement system using smart mirror technology. IEEE Access, 8, 142987–142999.

Lee, J., Park, S., & Kim, T. (2021). Facial emotion recognition smart mirror for mental health monitoring. Journal of Medical Systems, 45(3), 1–12.

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis,

, Dean, J., ... & Zheng, X. (2016).

TensorFlow: A system for large-scale machine learning. In 12th USENIX Symposium on Operating Systems Design and Implementation (pp. 265–283).

World Health Organization. (2023). Skin conditions affect one in three people globally. WHO Global Report on Neglected Tropical Diseases.

Brooke, J. (1996). SUS: A 'quick and dirty' usability scale. Usability Evaluation in Industry, 189(194), 4–7.

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861.

Jacob, C., Sanchez-Vazquez, A., & Ivory, C. (2020). Social, organizational, and technological factors impacting clinicians' adoption of mobile health tools: Systematic literature review. JMIR mHealth and uHealth, 8(2), e15935.