© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717366

Confidence-Aware Dual-CNN Framework for Vitiligo Detection Using Efficientnet-B0 With Freeze-Unfreeze Transfer Learning

DHEERAJ S KUMAR[1], ROHIT JOHN ALEX[2], ADHIDEV M D[3], DR. K. ARTHI[4]
[1,2,3]Dept. of Data Science and Business Systems SRM Institute of Science and Technology, Kattankulathur, India
[4]Associate Professor Dept. of Data Science and Business Systems SRM Institute of Science and Technology, Kattankulathur, India

Abstract- Vitiligo is a chronic depigmentation disorder affecting approximately 1–2% of the global population, and its accurate automated detection remains a challenging task due to variations in lesion size, skin tone, and imaging conditions. Existing deep learning-based vitiligo classifiers produce binary predictions without any measure of confidence or uncertainty, which can be clinically misleading. This paper proposes a Confidence-Aware Dual-CNN framework that employs two EfficientNet-B0 models with architecturally distinct classification heads, trained using a freeze-unfreeze transfer learning strategy. When both models produce consistent predictions, the system outputs a confident diagnosis of either Healthy or Vitiligo. When the models disagree beyond a designed disagreement threshold, the case is flagged as Uncertain and referred to a dermatologist for manual review. This is the first vitiligo detection framework to incorporate disagreement-based clinical uncertainty estimation. Evaluated on a dataset of 3,628 images, Model A achieves 94.98% accuracy, Model B achieves 96.28% accuracy, and the Dual-CNN framework achieves 97.82% accuracy on confident predictions, surpassing the existing IEEE CNN Autoencoder baseline of 90.16%. Supporting modules include an OpenCV-based lesion segmentation pipeline for progression tracking and a Random Forest model for treatment recommendation achieving 90% accuracy with a macro F1-score of 0.87.

Index Terms— Vitiligo Detection, Dual CNN, EfficientNet-B0, Transfer Learning, Uncertainty Estimation, Disagreement Threshold, OpenCV, Random Forest, Medical Image Classification

I. INTRODUCTION

Vitiligo is an autoimmune skin disorder characterized by the progressive loss of melanocytes, resulting in depigmented patches on the skin. With a global prevalence estimated between 1% and 2%, the condition significantly impacts patient quality of life, carrying considerable psychological and social burden. While vitiligo is not life-threatening, its visual manifestations necessitate early and accurate diagnosis to facilitate timely treatment and monitor disease progression. Traditional diagnosis relies on clinical examination by dermatologists, which is subjective, time-consuming, and limited in availability in resource-constrained healthcare settings.

The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has demonstrated considerable promise in automating the classification of skin conditions from dermoscopic images. However, most existing approaches for vitiligo detection treat the problem as a straightforward binary classification task, producing a single deterministic prediction without communicating any measure of confidence or uncertainty. In clinical environments, overconfident AI predictions can be dangerous, particularly when the model encounters ambiguous or borderline cases. A dermatologist does not simply declare a binary verdict — they acknowledge uncertainty and recommend further evaluation when evidence is inconclusive.

To address this critical gap, this paper introduces a Confidence-Aware Dual-CNN framework for vitiligo detection. The proposed system trains two EfficientNet-B0 models with architecturally distinct classification heads using a freeze-unfreeze transfer learning strategy. Agreement between the two models signals a confident prediction, while disagreement triggers an uncertainty flag that recommends dermatologist referral. This mechanism introduces a clinically meaningful third output — Uncertain — that no existing published vitiligo classification system provides. The framework is complemented by an OpenCV-based lesion segmentation module for progression analysis and a Random Forest classifier for treatment recommendation, creating a comprehensive end-to-end vitiligo management system.

II. LITERATURE REVIEW

Deep learning has been widely applied to dermatological image analysis, with numerous studies demonstrating the effectiveness of CNNs for skin lesion classification. Alzakari et al. [1] proposed LesionNet, combining SIFT features with a custom CNN for automated skin lesion classification, achieving competitive accuracy on benchmark datasets. Srinivasu et al. [2] demonstrated the utility of MobileNetV2 combined with LSTM for skin disease classification, achieving robust performance with reduced computational overhead. These works underscore the effectiveness of transfer learning in medical imaging applications.

In the specific context of vitiligo, Peng et al. [3] applied Generative Adversarial Networks and improved SMOTE techniques to address class imbalance in facial pigmentation disorder classification, including vitiligo. A CNN Autoencoder-based approach reported in IEEE Xplore [4] achieved 90.16% accuracy for vitiligo binary classification, representing the current baseline against which the proposed framework is compared. However, neither of these works addresses the issue of prediction uncertainty, which remains a significant limitation for clinical deployment.

Explainable AI methods such as Grad-CAM have been integrated into dermatology systems to enhance model transparency [5]. Munjal et al. [6] proposed SkinSage XAI, an explainable deep learning solution for skin lesion diagnosis that highlights clinically relevant image regions. While these approaches improve interpretability, they do not fundamentally address the problem of uncertain or ambiguous predictions. The proposed dual-model disagreement framework represents a novel approach to quantifying and communicating model uncertainty, which has not been previously explored in the vitiligo detection literature.

III. METHODOLOGY

The proposed system consists of three integrated modules: the core Confidence-Aware Dual-CNN for vitiligo classification, an OpenCV-based lesion segmentation and progression analysis pipeline, and a Random Forest model for treatment recommendation. The following subsections describe each component in detail.

A. Dataset and Preprocessing

The dataset used in this study comprises 3,628 dermoscopic images drawn from the publicly available Kaggle vitiligo dataset (shinynose/vitiligo), organized into two classes: 1,528 healthy skin images and 2,100 vitiligo-affected images. The class imbalance between categories is addressed through a Weighted Random Sampler during training, which assigns higher sampling probability to the minority class. All images are resized to 224×224 pixels to match the input requirements of EfficientNet-B0.

A stratified split strategy is employed to partition the dataset into 70% training, 15% validation, and 15% testing subsets, ensuring proportional class representation across all splits. Training images are subjected to strong data augmentation including random horizontal and vertical flips, rotations up to 25 degrees, color jitter with brightness and contrast variation of ±0.3, random affine translations, and normalization using ImageNet mean and standard deviation values.

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717366

Fig. 1. Sample vitiligo dermoscopic input image from the dataset

B. Core Novelty: Confidence-Aware Dual-CNN Architecture

The central contribution of this work is the Confidence-Aware Dual-CNN framework, which trains two EfficientNet-B0 models with intentionally different classification heads on the same dataset. The architectural divergence ensures that each model learns a distinct decision boundary, making their agreement a meaningful signal of prediction confidence.

1) Model A — Wider Classification Head

Model A employs an EfficientNet-B0 backbone with a wider classification head consisting of a linear layer mapping 1,280 feature dimensions to 256, followed by batch normalization, ReLU activation, and 0.4 dropout. A second linear layer reduces the representation to 128 dimensions, followed by ReLU, 0.3 dropout, and a final linear layer producing a 2-class output.

2) Model B — Deeper Classification Head

Model B employs the same EfficientNet-B0 backbone but uses a deeper classification head: 1,280 → 512 → 256 → 128 → 2. Each intermediate layer is followed by batch normalization, ReLU activation, and higher dropout rates of 0.5 and 0.4 respectively, providing stronger regularization and ensuring intentional architectural divergence from Model A.

3) Freeze-Unfreeze Transfer Learning Strategy

Both models are trained using a two-phase freeze-unfreeze transfer learning strategy. In Phase 1 (epochs 1–9), the EfficientNet-B0 backbone is frozen and only the classification head parameters are updated. In Phase 2 (epoch 10 onwards), the final 30 layers of Model A and the final 20 layers of Model B are selectively unfrozen, with the learning rate reduced by a factor of 10 to enable fine-grained feature adaptation. A ReduceLROnPlateau scheduler further reduces the learning rate when validation accuracy plateaus.

4) Disagreement-Based Uncertainty Detection

The core novelty of the proposed framework lies in its disagreement-based uncertainty mechanism. For each input image, both models independently produce a probability score representing the likelihood of vitiligo. A prediction is flagged as Uncertain if $|P\_A - P\_B| > 0.20$ or if the two models predict different classes. When the uncertainty condition is triggered, the system recommends dermatologist referral. This threshold was selected empirically, yielding 97.82% accuracy on confident predictions with approximately 11% uncertainty rate.

C. OpenCV Lesion Segmentation and Progression Analysis

The lesion segmentation module processes dermoscopic images using a pipeline of classical computer vision techniques implemented with OpenCV. The pipeline applies grayscale conversion, followed by Contrast Limited Adaptive Histogram Equalization (CLAHE) for local contrast enhancement, adaptive thresholding to produce a

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717366

binary mask, and contour detection using cv2.findContours to isolate individual lesion boundaries.

(a) Grayscale + CLAHE

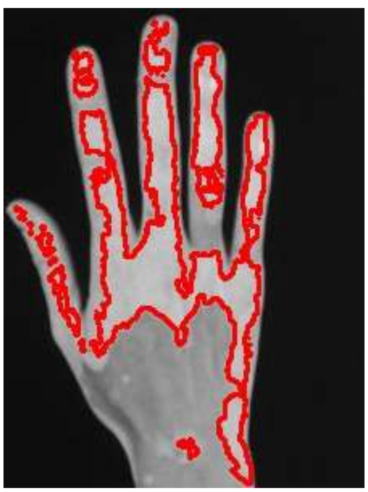

(b) Contour Detection Output

Fig. 2. OpenCV lesion segmentation pipeline: (a) grayscale with CLAHE enhancement, (b) vitiligo lesion contours detected in red

D. Random Forest Treatment Recommendation

A Random Forest classifier is employed for treatment recommendation, taking as input the vitiligo stage classification derived from the segmentation-estimated affected area percentage and the computed progression speed rate. The classifier consists of 100 decision trees trained with the class_weight='balanced' parameter to address class imbalance in the treatment recommendation dataset.

IV. SYSTEM ARCHITECTURE

The overall system architecture is organized as a sequential pipeline with three distinct functional layers. In the input layer, a dermoscopic skin image is submitted by the user. In the processing layer, the image is simultaneously passed through both trained EfficientNet-B0 models, each producing a class probability vector via softmax activation. The probabilities are compared using the disagreement detection logic: if both models agree and the probability difference is within the threshold of 0.20, the prediction is accepted as confident. If disagreement is detected, the case is flagged as Uncertain and the system recommends dermatologist referral.

In cases where a confident vitiligo prediction is made, the image is additionally processed by the OpenCV segmentation pipeline to estimate the affected skin area and classify the progression stage. The stage information and progression speed rate are then passed to the Random Forest treatment recommendation module. The complete system thus provides three levels of output: a binary classification with confidence indication, a quantitative lesion analysis, and a treatment recommendation.

V. RESULTS AND ANALYSIS

A. Experimental Setup

All experiments are conducted on Google Colab using an NVIDIA T4 GPU. The models are implemented in PyTorch using the timm library for EfficientNet-B0 with pre-trained ImageNet weights. Training is performed with a batch size of 32 and an initial learning rate of 1×10^{-3} using the Adam optimizer. CrossEntropyLoss with class weights is employed to address the 1,528 healthy versus 2,100 vitiligo class imbalance. A maximum of 25 epochs is configured, with EarlyStopping (patience=7) and ReduceLROnPlateau (factor=0.5, patience=3) callbacks.

B. Training Convergence

Fig. 3 presents the training and validation loss and accuracy curves for both Model A and Model B across 25 epochs. The dotted vertical line in each subplot marks epoch 10, at which the backbone is unfrozen. Both models exhibit a clear improvement in validation accuracy following the unfreeze point, confirming the effectiveness of the two-phase transfer learning strategy. Model A achieves stable validation accuracy above 0.97 from epoch 15 onwards, while Model B demonstrates consistent convergence with validation accuracy reaching 0.98 in later epochs.

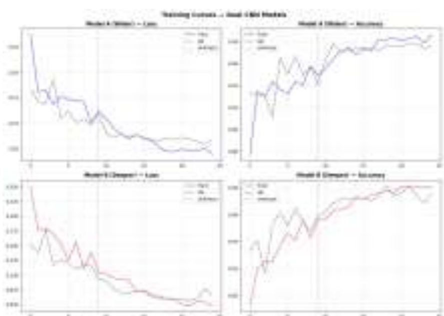

Fig. 3. Training curves for Dual-CNN models showing loss and accuracy across 25 epochs. Dotted line marks the unfreeze point at epoch 10.

C. Classification Performance

The performance of the proposed framework is evaluated on the held-out test set of 538 images comprising 230 healthy and 308 vitiligo samples. Table I presents the overall accuracy comparison between the proposed models and the existing IEEE baseline.

Model	Accuracy
IEEE CNN Autoencoder (2022) [Baseline]	90.16%
Model A — EfficientNet-B0 (Wider Head)	94.98%
Model B — EfficientNet-B0 (Deeper Head)	96.28%
Dual-CNN (Confident Predictions Only)	97.82%

Table I. Overall Accuracy Comparison

The Dual-CNN framework achieves 97.82% accuracy on confident predictions, representing a 7.66 percentage point improvement over the IEEE CNN Autoencoder baseline of 90.16%. Individually, Model A achieves 94.98% and Model B achieves 96.28%.

D. Class-Wise Performance Metrics

Table II presents the detailed precision, recall, and F1-score metrics for both models across both classification categories.

Class	Model A Prec.	Model A Rec.	Model A F1	Model B Prec.	Model B Rec.	Model B F1
Healthy	0.90	0.99	0.94	0.93	0.99	0.96
Vitiligo	0.99	0.92	0.95	0.99	0.94	0.97
Macro Avg.	0.95	0.96	0.95	0.96	0.97	0.96
Weighted Avg.	0.95	0.95	0.95	0.96	0.96	0.96

Table II. Class-wise Performance Metrics for Model A and Model B

Both models demonstrate well-balanced performance across both classes. Model A achieves a healthy recall of 0.99 and vitiligo precision of 0.99. Model B consistently exceeds Model A, with vitiligo recall improving from 0.92 to 0.94, confirming the benefit of the deeper classification head architecture.

E. Sample Prediction Output

Fig. 4 illustrates a representative output of the Dual-CNN system on a vitiligo-positive test image. Both Model A and Model B produce consistent high-confidence predictions of 99.3% and 99.7% respectively, with a probability difference of only 0.005, well within the disagreement threshold of 0.20. The system correctly classifies the image as Vitiligo Detected and proceeds to the spread analysis module.

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717366

Fig. 4. Sample Dual-CNN prediction output showing both model probabilities, difference score, and final classification decision

F. Uncertainty Analysis

Of the 538 test images, approximately 11% are flagged as Uncertain by the disagreement detection mechanism. These cases represent dermoscopic images where the two models produce sufficiently different probability outputs to trigger the uncertainty threshold of 0.20. The remaining 89% of test images receive confident predictions, on which the system achieves 97.82% accuracy. This uncertainty flagging mechanism ensures that borderline cases are directed to clinical review rather than forced into an incorrect binary prediction.

G. Supporting Module Performance

The Random Forest treatment recommendation model achieves an overall accuracy of 90% with a macro F1-score of 0.87 and a weighted F1-score of 0.90. The OpenCV lesion segmentation pipeline reliably quantifies the affected skin area percentage through CLAHE contrast enhancement and adaptive thresholding, enabling consistent progression tracking across varying image conditions.

VI. ADVANTAGES

The primary advantage of the proposed framework over existing vitiligo classifiers is its ability to communicate prediction uncertainty rather than forcing a binary output for every input. This is clinically significant because overconfident misclassifications in medical AI can lead to incorrect treatment decisions. The dual-model disagreement mechanism is inherently interpretable: when two independently trained models disagree, the case merits human review.

The freeze-unfreeze transfer learning strategy enables effective training on a relatively small dataset of 3,628 images by leveraging pre-trained ImageNet features. The Weighted Random Sampler and class-weighted loss function together ensure that the model does not collapse to predicting only the majority class, which was a critical challenge given the dataset imbalance.

VII. LIMITATIONS

The proposed framework has several limitations. First, the dataset is limited to 3,628 images from a single Kaggle source, which may not fully capture the diversity of vitiligo presentations across different skin tones and imaging conditions. Second, the disagreement threshold of 0.20 is determined empirically and may require recalibration for different clinical environments. Third, the OpenCV segmentation pipeline may produce unreliable area estimates for images with poor lighting, hair occlusion, or complex background textures. Fourth, the treatment recommendation module is trained on a synthetically constructed dataset rather than real patient outcomes.

VIII. FUTURE SCOPE

The framework can be evaluated on a larger and more diverse vitiligo dataset incorporating images from multiple clinical centers across a wider range of skin tones. The disagreement threshold can be adaptively calibrated using Bayesian uncertainty estimation or Monte Carlo Dropout. The treatment recommendation module can be strengthened by training on real patient outcome data in collaboration with dermatology clinicians. Extension to multi-class vitiligo severity staging and deployment as a web-based or mobile clinical decision support application represent natural directions for future development.

IX. CONCLUSION

This paper presented a Confidence-Aware Dual-CNN framework for vitiligo detection that addresses the fundamental limitation of existing binary classifiers: the absence of uncertainty estimation. By training two EfficientNet-B0 models with architecturally distinct classification heads using a freeze-unfreeze transfer learning strategy, and employing a designed disagreement threshold of 0.20 to flag uncertain predictions, the proposed system introduces the first vitiligo AI that communicates when it does not know.

Evaluated on 538 test images, Model A achieves 94.98% accuracy, Model B achieves 96.28%, and the Dual-CNN framework achieves 97.82% on confident predictions — a 7.66 percentage point improvement over the existing IEEE CNN Autoencoder baseline of 90.16%. Class-wise evaluation confirms balanced performance with macro F1-scores of 0.95 and 0.96 for Model A and Model B respectively. The supporting OpenCV lesion segmentation module provides quantitative progression tracking, and the Random Forest treatment recommendation module achieves 90% accuracy with a macro F1-score of 0.87.

REFERENCES

[1] S. A. Alzakari et al., "LesionNet: An automated approach for skin lesion classification using SIFT features with customized CNN," Frontiers in Medicine, vol. 11, 2024.

[2] P. N. Srinivasu et al., "Classification of Skin Disease Using Deep Learning Neural Networks with MobileNet V2 and LSTM," Sensors, vol. 21, no. 8, p. 2852, 2021.

[3] J. Peng et al., "Classification of Non-tumorous Facial Pigmentation Disorders Using GAN and Improved SMOTE," in Proc. 43rd IEEE EMBC, 2021, pp. 1–4.

[4] "Vitiligo Binary Classification Using CNN Autoencoder," IEEE Xplore, 2022. [Accuracy: 90.16%]

[5] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in Proc. IEEE ICCV, 2017, pp. 618–626.

[6] G. Munjal et al., "SkinSage XAI: An explainable deep learning solution for skin lesion diagnosis," Health Care Science, vol. 3, pp. 438–455, 2024.

[7] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. ICML, 2019.

[8] H. Liu et al., "A skin disease classification model based on multi-scale combined efficient channel attention module," Scientific Reports, vol. 15, 2025.