JOURNAL OF
EMERGING INVESTIGATORS

Article

Mitigating skin color bias in dermatology AI using CycleGAN-based data augmentation

Srivatsav Kannan[1], Vani Ramasamy[1]
[1] The Indian Public School, Coimbatore, Tamil Nadu, India

SUMMARY

Dermatological AI systems frequently underperform on darker skin tones due to severe data scarcity, with publicly available datasets overwhelmingly skewed toward light skin. This underrepresentation results in diagnostic inaccuracies that contribute to healthcare disparities. In this study, we investigated whether augmenting a dermatology dataset using a Cycle-consistent generative adversarial network (CycleGAN) to generate synthetic dark-skinned images from light-skinned counterparts could improve classification performance on real dark skin lesions. We hypothesized that a model trained on both real light skin images and synthetic dark skin images would achieve higher diagnostic accuracy on dark skin than a model trained solely on light skin data. The CycleGAN was trained on images from the Stanford Diverse Dermatological Images (DDI) and Fitzpatrick17k datasets to generate synthetic dark skin images. To evaluate the hypothesis, we compared the performance of two ResNet50 classifiers, ResNet A and ResNet B, in distinguishing benign and malignant lesions. ResNet A was trained on light skin images from the Stanford DDI and Fitzpatrick17k datasets, while ResNet B was trained on a balanced dataset of real light skin images from these datasets and CycleGAN-generated dark skin images. When evaluated on real dark skin images, ResNet B achieved a higher classification accuracy (71.95%) than ResNet A (65.16%), along with higher precision, recall, and F1 scores. These results supported our hypothesis that training on synthetic dark skin images can enhance dermatological classifier performance on underrepresented skin tones. This approach offers a promising solution for addressing bias in dermatological AI systems when real-world datasets lack diversity.

INTRODUCTION

Skin tone has long been recognized as a significant factor influencing dermatological diagnosis and patient outcomes. Individuals with darker skin tones experience disproportionately poorer outcomes across multiple skin conditions, particularly melanoma, where delayed detection is common (1). In fact, the five-year survival rate for melanoma is approximately 94% in White individuals but decreases to nearly 70% in Black individuals, largely due to delays and inaccuracies in early diagnosis (1). These disparities arise from a combination of factors, including differences in lesion appearance on darker skin, limited representation of dark skin in medical education and textbooks, and reduced clinical exposure among physicians to melanoma presentations in patients with skin of color (2, 3).

As artificial intelligence (AI) systems increasingly achieve expert-level performance in dermatological diagnosis, they risk inheriting the same diagnostic disparities between light and dark-skinned populations observed in clinical practice (4). Model performance is fundamentally constrained by the diversity and representativeness of training data, and existing imbalances in dermatological image datasets therefore directly translate into biased predictions (5). Notably, fewer than 10% of images in commonly used public dermatology datasets represent darker skin tones, limiting these models' ability to generalize equitably across skin types (6).

One promising solution to address this data imbalance lies in the use of generative models, particularly generative adversarial networks (GANs), to synthetically augment datasets with dark-skinned images. GANs consist of two neural networks trained in opposition, one generating images and the other evaluating their realism, allowing the system to learn how to produce images that closely resemble real clinical data (7). Despite this promise, it remains unclear whether synthetic dark-skin images can meaningfully improve diagnostic accuracy when incorporated into model training.

We hypothesized that augmenting a dermatological training dataset with synthetic dark-skinned images, generated using a Cycle-consistent generative adversarial network (CycleGAN), would improve the accuracy of classifying dark-skinned images into benign or malignant lesions compared to training solely on light-skinned images. Our results supported this hypothesis, demonstrating a 6.79% increase in diagnostic performance on real dark-skinned images when synthetic augmentation was incorporated. These findings indicate that augmentation with synthetic dermatological images may provide a scalable method for improving fairness in dermatological AI, with the potential to contribute to improved healthcare outcomes for patients with underrepresented skin tones.

RESULTS

To evaluate the hypothesis that CycleGAN-generated synthetic dark skin images improve diagnostic performance on dark skin dermatological images, we conducted a controlled comparison between two ResNet50-based classifiers, convolutional neural networks built on a 50-layer deep residual learning architecture, trained under different data conditions. The Stanford Diverse Dermatological Images (DDI) and Fitzpatrick17k datasets, publicly available collections of dermatological images, were used for all training and testing (8, 9).

ResNet A was trained exclusively on light-skinned images, while ResNet B was trained on a dataset of equal size composed of real light-skinned images and CycleGAN-generated dark-skinned images **(Figure 1)**. Both models

https://doi.org/10.59720/25-159

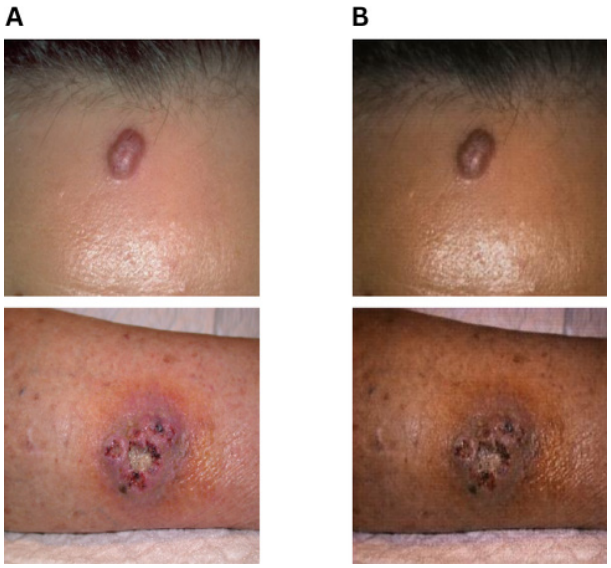

Figure 1. Paired examples of CycleGAN inputs and outputs. (A) Real light-skinned dermatological images from the Stanford Diverse Dermatology Images (DDI) and Fitzpatrick17k datasets (8, 9). (B) Corresponding synthetic dark-skinned translations generated using the trained CycleGAN model. The images illustrate image-to-image translation in which skin tone is modified while lesion morphology and structural features are preserved. Representative samples shown (n=2).

were evaluated on the same test set of 221 real dark-skinned dermatological images for their ability to classify lesions as benign or malignant.

ResNet A achieved an accuracy of 65.16%, with a precision of 50.54%, sensitivity of 60.26%, specificity of 67.83%, negative predictive value (NPV) of 75.78%, and F1-score of 54.97% **(Table 1)**. These results are consistent with expectations, as models trained exclusively on light-skinned data often generalize poorly to darker skin tones, resulting in high false positive rates (reflected in low precision) (6).

Model	Accuracy	Precision	Sensitivity	Specificity	NPV	F1-Score
ResNet A	65.16%	50.54%	60.26%	67.83%	75.78%	54.97%
ResNet B	71.95%	60.00%	61.54%	77.62%	78.72%	60.76%

Table 1. Performance on real dark-skinned images. Accuracy, precision, sensitivity, specificity, negative predictive value (NPV), and F1-score for ResNet A, a ResNet50-based classifier trained exclusively on real light-skinned images, and ResNet B, a ResNet50-based classifier trained on both real light-skinned images and CycleGAN-generated synthetic dark-skinned images. Metrics were calculated using predictions on test set of 221 real dark-skinned images (143 benign, 78 malignant).

In contrast, ResNet B, achieved markedly improved performance across all metrics: an accuracy of 71.95%, precision of 60.00%, sensitivity of 61.54%, specificity of 77.62%, NPV of 78.72%, and F1-score of 60.76% **(Table 1)**.

A closer inspection of the confusion matrices further contextualizes these results **(Figure 2)**. Compared to ResNet A, ResNet B demonstrated improved classification performance across both lesion classes, with correctly classified benign cases increasing from 97 to 111 and false positives decreasing from 46 to 32. Malignant lesion detection also showed a modest improvement, with false negatives reduced from 31 to 30 and true positives increasing from 47 to 48. These shifts indicate gains in both specificity and sensitivity, contributing to higher overall accuracy and precision. The more pronounced reduction in benign misclassification suggests that synthetic augmentation primarily enhanced detection of benign lesion features on darker skin tones, while malignant feature representations

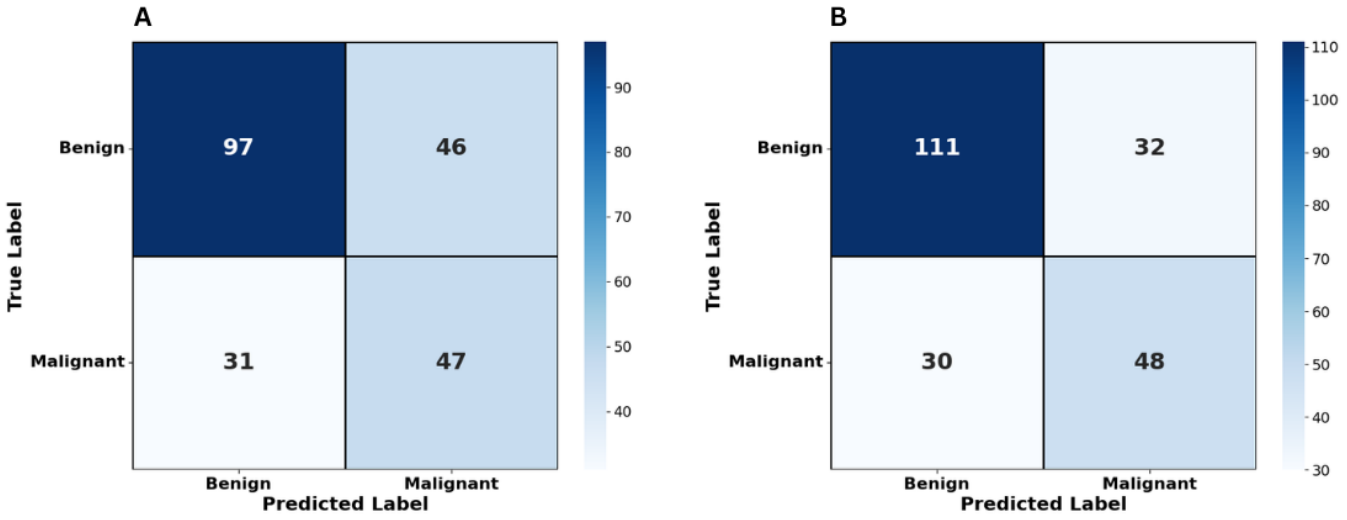

Figure 2. Confusion matrices displaying diagnostic performance on real dark-skinned images. (A) Confusion matrix of ResNet A, a ResNet50-based classifier trained exclusively on real light-skinned images. (B) Confusion matrix of ResNet B, a ResNet50-based classifier trained on both real light-skinned images and CycleGAN-generated synthetic dark-skinned images. Matrix entries represent absolute counts of benign and malignant predictions compared to ground truth labels from a test set of 221 real dark-skinned images (143 benign, 78 malignant). Darker shades of blue indicate higher counts, while lighter shades represent lower counts. True labels are displayed on the y-axis, predicted labels on the x-axis.

https://doi.org/10.59720/25-159

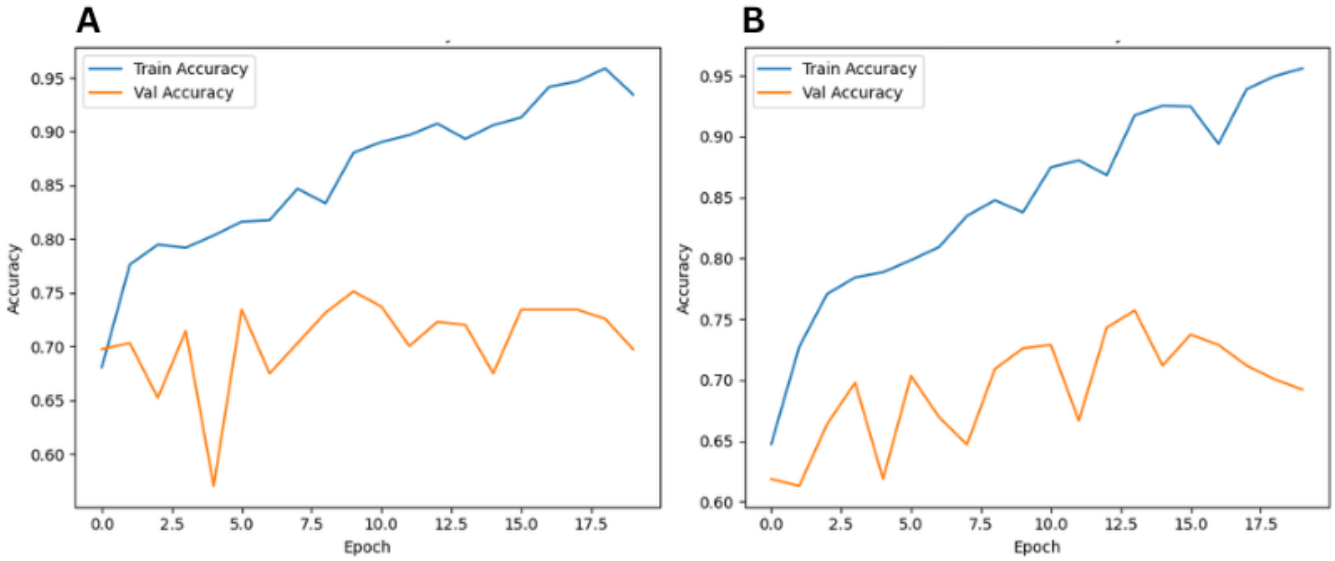

Figure 3. Training and validation accuracy curves. (A) Accuracy curves for ResNet A, a ResNet50-based convolutional neural network trained exclusively on real light-skinned images. (B) Accuracy curves for ResNet B, a ResNet50-based convolutional neural network trained on both real light-skinned images and CycleGAN-generated synthetic dark-skinned images. Training and validation accuracy were computed at the end of each epoch as the proportion of correctly classified images (benign or malignant) relative to total samples in the dataset.

were only marginally improved.

An examination of the training and validation curves indicates mild overfitting in both ResNet A and ResNet B, reflected by a consistent gap between training and validation accuracy and loss on the light-skinned validation set **(Figures 3 and 4)**. Importantly, however, the overall learning trajectories of the two models were highly similar. This suggests that the inclusion of CycleGAN-generated dark-skinned images did not degrade generalization to light-skinned data. The observed improvements in diagnostic performance on dark-skinned test images were achieved without compromising model behavior on light-skinned validation data.

A statistically significant difference between the two models' predictions was observed (McNemar's test, χ^2 = 15.0, p = 0.036), supporting the hypothesis that incorporating synthetic dark-skinned images during training improves classification performance on real dark-skinned skin lesion images.

DISCUSSION

This study set out to evaluate whether augmenting dermatological training datasets with synthetic dark skin images generated by a CycleGAN could improve classification performance on real dark skin lesion images.

The observed improvements in performance metrics strongly support the hypothesis. The model trained with both light skin and synthetic dark skin images, ResNet B, achieved an accuracy of 71.95% on a test set containing real dark-skinned images, compared to just 65.16% for the baseline model, ResNet A, trained only on the light-skinned images. This improvement was observed across all evaluated metrics. In particular, the increase in F1-score from 54.97% to 60.76% indicates that classification performance improved for both benign and malignant lesions in a manner consistent with the overall gain in accuracy.

Despite these gains, several important limitations warrant careful consideration. Although GAN-based augmentation provides a practical mechanism for increasing skin tone diversity, it cannot fully substitute for large-scale, real-world clinical data (10). In particular, malignant lesions such as melanoma exhibit clinically meaningful differences across skin tones that extend beyond pigmentation alone. Melanomas in individuals with darker skin are more likely to present in acral locations and may display different morphological patterns (11). By design, CycleGAN primarily performs a domain-level appearance transfer, altering global pigmentation while preserving underlying lesion geometry and spatial context. As a result, synthetic images may fail to capture deeper phenotypic differences between light skin and dark skin lesions. Without expert dermatological validation, it therefore remains uncertain whether generated images fully reflect the spectrum of malignant presentations observed clinically.

The limited availability of real dark-skinned images constrained the statistical power of performance comparisons and makes it difficult to fully rule out the influence of randomness on the observed improvements in model performance. This scarcity also precluded a direct comparison between models trained on synthetic versus real dark-skinned data, as allocating real dark-skinned images to training would have further diminished an already limited test set.

Additionally, due to dataset size constraints, some degree of data reuse was necessary between CycleGAN training and ResNet evaluation. Although no identical images appeared across ResNet training and testing splits, the possibility of indirect information transfer cannot be entirely excluded. Larger, independent datasets would mitigate this concern and enable more rigorous separation between generative and classifier learning stages. These limitations reflect the broader and well-documented scarcity of diverse dermatological imaging data.

Looking ahead, several extensions of this work are possible.

https://doi.org/10.59720/25-159

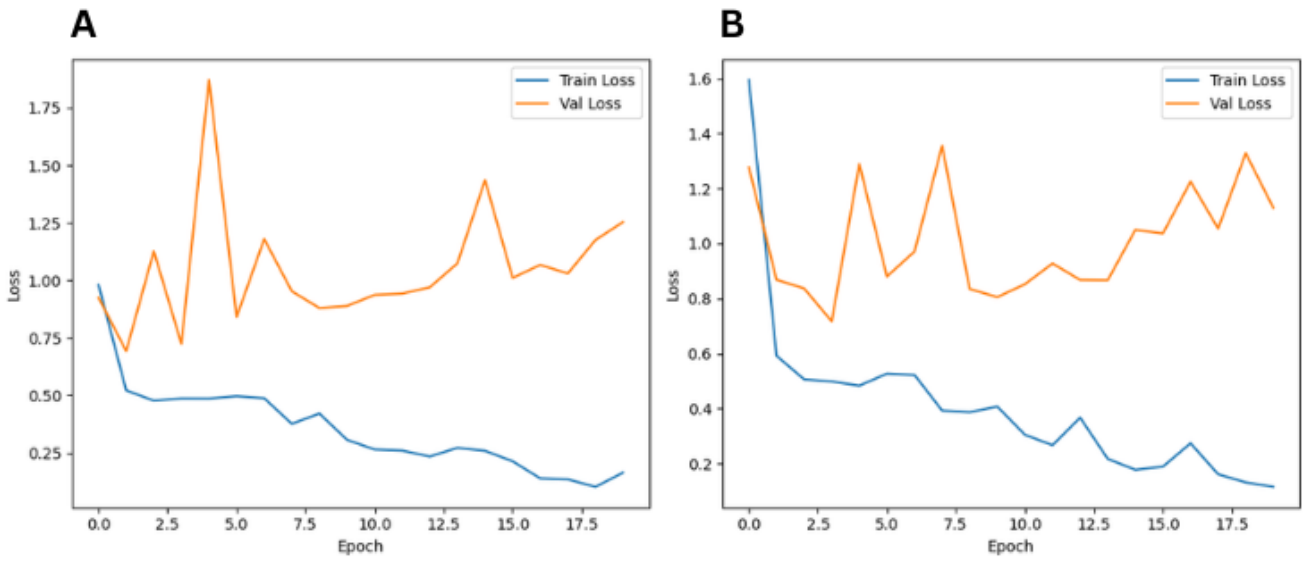

Figure 4. Training and validation loss curves. (A) Loss curves for ResNet A, a ResNet50-based convolutional neural network trained exclusively on real light-skinned images. (B) Loss curves for ResNet B, a ResNet50-based convolutional neural network trained on both real light-skinned images and CycleGAN-generated synthetic dark-skinned images. Training and validation loss computed at the end of each epoch using categorical cross-entropy loss.

Alternative generative architectures, such as StyleGAN2 or diffusion-based models, may offer improved synthesis quality and reduced visual artifacts, potentially strengthening downstream classification performance. Future studies could also stratify results by disease category to assess whether synthetic augmentation disproportionately benefits certain diagnostic classes. Importantly, expert dermatological review of generated images would be a critical next step to validate clinical fidelity and guide model refinement.

Prior studies have explored synthetic image generation as a means of increasing diversity in dermatological datasets, with evidence that such augmentation can improve classification performance. For instance, Rezk et al. used style transfer–based methods to modify skin tone and demonstrated improved accuracy and area under the receiving operating characteristics curve (AUC - ROC) when synthetic images were incorporated into training, alongside qualitative validation of image realism (12). Similarly, Ghorbani et al. introduced DermGAN, a GAN-based framework for generating dermatological images with controllable pathology and skin tone, showing that synthetic images could match real images in visual fidelity and support classification performance (13). While style transfer methods primarily modify higher-level visual attributes such as color and texture without explicitly preserving disease presentation, and DermGAN generates synthetic images from latent representations without direct correspondence to a specific real input image, the CycleGAN used in our study performs image-to-image translation with cycle-consistency constraints, enabling targeted skin tone modification while maintaining accurate and realistic lesion morphology derived from real clinical images.

Moreover, these prior studies primarily evaluated performance on mixed test sets dominated by light-skinned images and did not explicitly assess generalization to real dark-skinned samples. To our knowledge, no prior study has isolated the impact of CycleGAN-based skin tone translation on classification performance specifically for real dark-skinned test images, nor examined whether performance gains arise from improved dark skin representation or increased dataset size.

The results of this study addressed this gap and support the use of synthetic dark-skinned images as a viable strategy for mitigating performance disparities in dermatological AI. Training a ResNet model on a balanced combination of real light skin images and CycleGAN-generated dark skin images resulted in a 6.79% percentage point increase in accuracy, accompanied by consistent gains in precision, recall, and F1-score. While data limitations necessitate cautious interpretation, these findings suggest that generative augmentation can serve as a practical and scalable step toward more equitable AI systems, particularly in clinical domains where access to diverse real-world data remains constrained.

MATERIALS AND METHODS

CycleGAN Training and Image Synthesis

To generate synthetic dark-skinned dermatological images, we trained a Cycle-consistent generative adversarial network (CycleGAN) using a curated subset of clinical images from the Stanford Diverse Dermatology Images (DDI) and Fitzpatrick17k datasets (8, 9). These sources were selected for their annotation of images using the Fitzpatrick Skin Type Scale, a dermatological classification system that categorizes skin into six types (I to VI) based on its response to ultraviolet exposure (14). This scale provided a standardized and clinically relevant framework for separating training domains and evaluating the impact of skin tone on model performance. In this study, Types I–IV were grouped as “light-skinned” and Types V–VI as “dark-skinned,” reflecting both the relative pigmentation differences across the scale and the

https://doi.org/10.59720/25-159

typical overrepresentation of Types I–IV compared to the underrepresented darker skin categories in dermatological datasets (6).

All images used for training and testing were public and appropriately licensed. The CycleGAN architecture was chosen over other GAN architectures as it is commonly used for domain transfer tasks like this study's (7). The original datasets were cleaned by removing samples lacking visible lesions or containing obstructive text or markers. All images were manually cropped to isolate the pathological area and resized to a resolution of 256×256 pixels. The training set was curated to increase performance by excluding low-quality images with poor lesion visibility. A total of 1625 curated images, including 1404 light and 221 dark, were used to train the CycleGAN **(Table 2)**.

The CycleGAN architecture employed Unet-based generators and PatchGAN discriminators. Images were normalized to the range [−1, 1] prior to training. Generator G learned to map light skin images (Domain X) to the dark skin images (Domain Y), while generator F performed the inverse. Two discriminators, DX and DY, evaluated the realism of generated images in their respective domains. The model was optimized using three loss functions. The adversarial loss function trained the generators to create images that closely resembled actual images from the target domain (7; Equation 1).

$$L_{GAN}(G, D_Y, X, Y) = E_{x\sim P_x(x)}[\log D_Y(G(x))] \quad \text{..... (Equation 1)}$$

Cycle Consistency Loss ensured that translating an image to the target domain and back yielded a reconstruction of the original image (7; Equation 2).

$$L_{cyc}(G, F) = E_{x\sim P_X(x)}\left[\left|\left|F(G(x)) - x\right|\right|_1\right] + E_{y\sim P_y(y)}\left[\left|\left|F(G(y)) - y\right|\right|_1\right]$$

.....(Equation 2)

Identity Loss prevented unnecessary transformation when the input already belonged to the target domain (7; Equation 3).

$$L_{identity}(G, Y) = E_{y\sim P_Y(y)}\left[\left|\left|G(Y) - y\right|\right|_1\right] \text{.....(Equation 3)}$$

The full generator loss was defined as:

$$L^G_{total} = \lambda \cdot L_{GAN}(G, D_Y, X, Y) + L_{cyc}(G, F) + \frac{\lambda}{2} \cdot L_{identity}(G, Y)$$

.....(Equation 4)

where λ was a weighting coefficient, set at 10 in this study in alignment with the original CycleGAN implementation (7). The model was trained for 50 epochs using Adam optimizers for both generators and discriminators. After training, synthetic dark skin images were generated by passing light skin images from the Stanford DDI and Fitzpatrick17k datasets through the trained generator G and saved for later use in ResNet model training (8, 9).

ResNet diagnostic classifier training

To evaluate the impact of synthetic data on classification performance, two binary skin lesion classifiers were trained to distinguish between benign and malignant lesions. Both classifiers were based on the ResNet50 architecture, pretrained on the ImageNet dataset. This architecture was chosen as it has been found to perform well for diagnosis of dermatological conditions (15).

Dataset	Skin Tone of image		Total
	Light	Dark	
Stanford DDI	267	113	380
Fitzpatrick 17k	1137	108	1245
Total	1404	221	1625

Table 2. Final curated training image counts for CycleGAN. Number of light-skinned (Fitzpatrick Types I–IV) and dark-skinned (Fitzpatrick Types V–VI) dermatological images used to train the CycleGAN model, sourced from the Stanford Diverse Dermatology Images (DDI) and Fitzpatrick17k datasets (8, 9).

ResNet50 is a deep convolutional neural network comprising 50 layers with residual connections, which allow gradients to flow through shortcut paths during training and alleviate vanishing gradient issues (16). The final classification head was removed, and a new head composed of a dense layer (512 neurons, ReLU activation), a global average pooling layer, and a sigmoid output layer were added for binary classification. All input images were resized to 224×224 pixels.

Training data for ResNet A consisted exclusively of curated light skin images from the Fitzpatrick17k and Stanford DDI datasets (8, 9). A total of 701 benign and 703 malignant images were used for training, with an additional 179 benign and 175 malignant light skin images used as the validation set **(Table 3)**. No dark skin images, real or synthetic, were included in the training or validation of ResNet A.

ResNet B was trained on a dataset of identical size and class composition to ResNet A. Specifically, ~50% of the original light-skinned training images from each class were retained unchanged, while the remaining ~50% were replaced with CycleGAN-generated dark-skinned translations. As a result, ResNet B was trained on 701 benign images (350 real light-skinned and 351 synthetic dark-skinned) and 703 malignant images (351 real light-skinned and 352 synthetic

Dataset	Benign	Malignant	Total
Training Dataset	701	703	1404
Validation Dataset	179	175	354
Dark-Skinned Testing Dataset	143	78	221

Table 3. Datasets used for ResNet A and ResNet B. Number of benign and malignant images used for training, validating and testing ResNet A, a ResNet50-based classifier trained exclusively on real light-skinned images, and ResNet B, a ResNet50-based classifier trained on both real light-skinned images and CycleGAN-generated synthetic dark-skinned images.

https://doi.org/10.59720/25-159

dark-skinned), resulting in an approximately even split between real light-skinned and synthetic dark-skinned samples. No real dark-skinned images were included in training. This design ensured that any observed performance differences arose from the inclusion of synthetic dark skin representations rather than increased data volume, while preserving baseline performance on light-skinned images.

Both models were trained using the Adam optimizer with binary cross-entropy loss and a batch size of 16. The validation set of real light skin images was used to monitor the training progress, with early stopping applied to mitigate overfitting and retrieve the best-performing model. The training history was saved and visualized to evaluate convergence trends.

Performance was measured using accuracy, precision, sensitivity, specificity, NPV, and F1-score on a test set consisting of 143 benign and 78 malignant real dark-skinned images from the Stanford DDI and Fitzpatrick17k datasets (8, 9). To assess whether differences in classification performance between ResNet A and ResNet B were statistically significant, McNemar's test was applied to paired prediction outcomes on the same test set. We constructed a 2×2 contingency table to identify cases misclassified by one model but correctly classified by the other (discordant pairs) and calculated the chi-squared statistic based on these discordant counts. A two-sided significance threshold of $p < 0.05$ was used to determine statistical significance.

Received: July 9, 2025
Accepted: January 26, 2026
Published: June 24, 2026

REFERENCES

1. Gupta, Alpana K., et al. "Skin Cancer Concerns in People of Color: Risk Factors and Prevention." *Asian Pacific Journal of Cancer Prevention*, vol. 17, no. 12, 2016, pp. 5257–5264. https://doi.org/10.22034/APJCP.2016.17.12.5257
2. Rizvi, Zehra, et al. "The Bias of Physicians and Lack of Education in Patients of Color with Melanoma as Causes of Increased Mortality: A Scoping Review." *Cureus*, vol. 14, no. 11, 19 Nov. 2022, p. e31669. https://doi.org/10.7759/cureus.31669
3. LaMonica, Lauren C., et al. "Skin of Colour Dermatology Education in US Primary Care Residency Programmes: A Nationwide Cross-Sectional Survey of Programme Directors." *Skin Health and Disease*, vol. 5, no. 1, 23 Jan. 2025, pp. 9–13. https://doi.org/10.1093/skinhd/vzae001
4. Brinker, Titus J., et al. "A Convolutional Neural Network Trained with Dermoscopic Images Performed on Par with 145 Dermatologists in a Clinical Melanoma Image Classification Task." *European Journal of Cancer,* vol. 111, Apr. 2019, pp. 148–154 https://doi.org/10.1016/j.ejca.2019.02.005
5. Parikh, Ravi B., et al. "Addressing Bias in Artificial Intelligence in Health Care." *JAMA*, vol. 322, no. 24, 24 Dec. 2019, p. 2377, https://doi.org/10.1001/jama.2019.18058.
6. Alipour, Neda, et al. "Skin Type Diversity in Skin Lesion Datasets: A Review." *Current Dermatology Reports*, vol. 13, 14 Aug. 2024. https://doi.org/10.1007/s13671-024-00440-0
7. Zhu, Jun-Yan, et al. "Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks." *ArXiv.org*, 2017. https://doi.org/10.48550/arXiv.1703.10593
8. "DDI - Diverse Dermatology Images." *Center for Artificial Intelligence in Medicine & Imaging, Stanford University*. https://aimi.stanford.edu/datasets/ddi-diverse-dermatology-images. Accessed 3 Mar. 2026.
9. Groh, Matthew, et al. "Towards Transparency in Dermatology Image Datasets with Skin Tone Annotations by Experts, Crowds, and an Algorithm." *ArXiv*, 6 July 2022. https://doi.org/10.48550/arXiv.2207.02942
10. Izu-Belloso, Rosa Maria, et al. "Generative Adversarial Networks in Dermatology: A Narrative Review of Current Applications, Challenges, and Future Perspectives." *Bioengineering*, vol. 12, no. 10, 16 Oct. 2025, p. 1113. https://doi.org/10.3390/bioengineering12101113
11. Brunsgaard, Elise, et al. "Melanoma in Skin of Color: Part I. Epidemiology and Clinical Presentation." *Journal of the American Academy of Dermatology*, vol. 89, no. 3, May 2022. https://doi.org/10.1016/j.jaad.2022.04.056
12. Rezk, Eman, et al. "Improving Skin Color Diversity in Cancer Detection: Deep Learning Approach." *JMIR Dermatology*, vol. 5, no. 3, 19 Aug. 2022, p. e39143. https://doi.org/10.2196/39143
13. Ghorbani, Amirata, et al. "DermGAN: Synthetic Generation of Clinical Skin Images with Pathology." *ArXiv.org*, 2019. https://doi.org/10.48550/arXiv.1911.08716
14. Oliveira, Rita, et al. "An Overview of Methods to Characterize Skin Type: Focus on Visual Rating Scales and Self-Report Instruments." *Cosmetics*, vol. 10, no. 1, 11 Jan. 2023, pp. 14–14. https://doi.org/10.3390/cosmetics10010014
15. Abd ElGhany, Sameh, et al. "Diagnosis of Various Skin Cancer Lesions Based on Fine-Tuned ResNet50 Deep Network." *Computers, Materials & Continua*, vol. 68, no. 1, 2021, pp. 117–135. https://doi.org/10.32604/cmc.2021.016102
16. He, Kaiming, et al. "Deep Residual Learning for Image Recognition." *ArXiv*, 2015. https://doi.org/10.48550/arXiv.1911.08716

Copyright: © 2026 Srivatsav Kannan and Vani Ramasamy. All JEI articles are distributed under the Creative Commons Attribution Noncommercial No Derivatives 4.0 International License. This means that you are free to share, copy, redistribute, remix, transform, or build upon the material for any purpose, provided that you credit the original author and source, include a link to the license, indicate any changes that were made, and make no representation that JEI or the original author(s) endorse you or your use of the work. The full details of the license are available at https://creativecommons.org/licenses/by-nc-nd/4.0/deed.en.