International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 04 Issue: 05 | May – 2025 DOI: 10.55041/ISJEM03698
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Enhancing Skin Cancer Classification Through GAN-Generated Synthetic Images for Improved CNN Trainings

P. Kamakshi Thai[1], Sai Jayanth Bandaru[2], Abhishek Sharma[3], Akshay Devalla[4]
Assistant Professor of Department of CSE(AI&ML) of ACE Engineering College[1]
Students of Department of CSE(AI&ML) of ACE Engineering College[2,3,4]

Abstract:
Skin cancer is among the most prevalent forms of cancer globally, and early detection is crucial for improving patient outcomes. While Convolutional Neural Networks (CNNs) have demonstrated strong performance in image-based skin cancer classification, their effectiveness is highly dependent on large and diverse datasets. However, medical image datasets often suffer from class imbalance and limited sample sizes. This study presents a novel approach to augmenting training data by leveraging Generative Adversarial Networks (GANs) to synthesize high-quality skin lesion images. These synthetic images are used to enhance the training set for a CNN-based classification model. Experimental results show a significant improvement in classification accuracy and robustness when synthetic images are included, particularly for underrepresented classes. The findings highlight the potential of GAN-augmented datasets in addressing data scarcity and improving diagnostic performance in medical imaging applications.

Keywords: Skin Cancer Classification, Convolutional Neural Networks (CNNs), Generative Adversarial Networks (GANs), Data Augmentation, Synthetic Medical Images, Deep Learning, Class Imbalance, Medical Image Analysis.

1. Introduction

Skin cancer is one of the most common and rapidly increasing forms of cancer worldwide, with millions of new cases diagnosed annually. Early and accurate diagnosis plays a vital role in successful treatment and survival rates. Dermatologists traditionally rely on dermoscopic examination and biopsy for diagnosis, but these methods are time-consuming, subjective, and require specialist expertise. In recent years, the application of deep learning, particularly Convolutional Neural Networks (CNNs), has shown great promise in automating and enhancing the accuracy of skin cancer detection from dermoscopic images.

Despite the impressive performance of CNNs, their success heavily relies on the availability of large, well-annotated, and balanced datasets. However, in medical domains, especially in dermatology, acquiring such datasets poses significant challenges due to patient privacy, high annotation costs, and the rarity of certain skin cancer types. This leads to class imbalance and limited sample diversity, which can cause biased model learning and poor generalization, particularly on underrepresented categories.

To address these challenges, data augmentation techniques are commonly employed to artificially expand training datasets. While traditional augmentation methods (e.g., rotation, flipping, scaling) offer some improvement, they are limited in their ability to introduce new variations and may not sufficiently capture the complexity of skin lesions. Recently, Generative Adversarial Networks (GANs) have emerged as a powerful tool for generating realistic synthetic images, offering a novel solution for augmenting datasets in a more meaningful way.
By synthesizing high-quality and diverse dermoscopic images, GANs can help mitigate class imbalance and closely resemble real clinical cases. This augmentation not only improves model robustness and generalization but also enables CNNs to better recognize rare and atypical presentations of skin cancer.

2. Problem Statement

Despite recent advancements in deep learning for medical image classification, skin cancer detection still faces critical challenges due to the scarcity and imbalance of annotated dermoscopic image datasets. CNN models require a large amount of diverse and well-distributed data to effectively learn and generalize across various types of skin lesions. However, medical datasets often suffer from class imbalance, where rare but dangerous types like melanoma are underrepresented, and overall limited dataset size due to patient privacy concerns and the high cost of manual annotation. These limitations hinder the performance of CNN-based classifiers, leading to biased predictions and reduced diagnostic accuracy, especially for minority classes.

While traditional data augmentation techniques attempt to alleviate some of these issues, they fail to introduce sufficient diversity or generate entirely new examples. This highlights a pressing need for more sophisticated data augmentation strategies that can simulate realistic and varied skin lesion images. In this context, the use of Generative Adversarial Networks (GANs) presents a promising solution for generating high-quality synthetic images to augment training datasets and improve the performance of CNN models in skin cancer classification.

3. Background of the Problem

Skin cancer is a growing global health concern, with incidence rates rising steadily due to increased exposure to ultraviolet (UV) radiation and other environmental factors. The two most common types are non-melanoma skin cancers—basal cell carcinoma and squamous cell carcinoma—and melanoma, which is more aggressive and has a higher mortality rate if not detected early. According to the World Health Organization (WHO), millions of new cases of skin cancer are reported each year, highlighting the urgent need for efficient and accurate diagnostic systems.

Traditional methods for diagnosing skin cancer, such as visual inspection and dermoscopic analysis followed by histopathological confirmation, are often time-consuming, subjective, and dependent on specialist expertise. These limitations have driven the adoption of computer-aided diagnostic (CAD) systems powered by artificial intelligence (AI), particularly deep learning. Convolutional Neural Networks (CNNs) have shown exceptional performance in analyzing medical images and classifying skin lesions with accuracy comparable to that of experienced dermatologists.

Moreover, the effectiveness of these AI-driven systems heavily relies on the availability of large, diverse, and well-annotated datasets. However, in the medical domain, acquiring such datasets can be challenging due to privacy concerns and the scarcity of labeled images.

4. Disadvantages of the Existing System

Despite the success of deep learning models in medical image analysis, several limitations hinder the effectiveness of current skin cancer classification systems. One of the primary challenges is the scarcity and imbalance of annotated medical image datasets. Publicly available datasets, such as ISIC and HAM10000, are often limited in size and heavily skewed toward common benign conditions. This class imbalance leads CNN models to be biased toward majority classes, such as benign nevi, and perform poorly on rare but critical cases like melanoma. Additionally, traditional data augmentation techniques, such as rotation, flipping, and color adjustments, offer limited benefit since they cannot generate entirely new lesion characteristics or reflect the full range of lesion variability observed in clinical practice.

Furthermore, CNNs trained on such limited and repetitive data tend to overfit, reducing their ability to generalize effectively to new, unseen cases. Although some research has introduced Generative Adversarial Networks (GANs) to address data limitations, many implementations produce synthetic images with artifacts or insufficient clinical realism, reducing their utility in training robust models. In many cases, synthetic data is incorporated into training pipelines without systematic evaluation, making it difficult to measure its true impact on model performance. These shortcomings highlight the need for a more effective and integrated approach to synthetic data generation and its use in improving classification accuracy for underrepresented skin cancer types.

5. Related works

Recent advancements in deep learning have significantly improved the accuracy of automated skin cancer diagnosis. Asif et al. (2024) introduced **CFI-Net**, an ensemble network that combines Choquet Fuzzy Integral (CFI) with Particle Swarm Optimization (PSO) to optimize fuzzy measures for classifying multiple skin cancers, including Mpox. By integrating outputs from multiple classifiers based on their contribution, CFI-Net achieved enhanced diagnostic performance, showcasing the strength of soft computing methods in dermatological image analysis.

In a parallel effort to improve feature extraction, Hao et al. (2023) developed **ConvNeXt-ST-AFF**, a hybrid model that fuses ConvNeXt and Swin Transformer architectures. This combination leverages convolutional and attention-based mechanisms to effectively capture both local lesion textures and global skin structure patterns, leading to superior classification accuracy and clinical interpretability.

To further improve lesion detection, Xie et al. (2020) proposed a **mutual bootstrapping model** that jointly performs lesion segmentation and classification, allowing each task to reinforce the other. This multi-task learning framework enhances boundary detection and contextual feature extraction, which are critical for distinguishing between malignant and benign lesions. However, even with such architectural innovations, many deep learning models suffer from limitations due to imbalanced and small datasets. Addressing this, Yao et al. (2022) implemented a model designed for **imbalanced skin lesion classification** using loss re-weighting and data augmentation strategies. Their results showed that optimizing the learning process for minority classes.

While traditional machine learning approaches have also contributed to this domain, as highlighted by Vasudha Rani et al. (2022), who demonstrated that **data mining algorithms** can be effective when paired with proper preprocessing and feature selection, these models still depend heavily on handcrafted features and often lack adaptability to complex image variations. A common thread across these works is their reliance on existing datasets, which are often limited in size and diversity.

6. Proposed Methodology

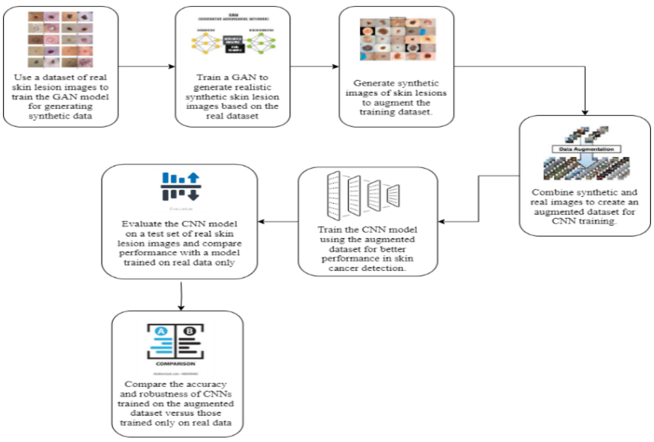

Fig 1: Proposed Architecture

6.1 Overview

This study proposes a novel methodology to enhance skin cancer classification by leveraging Generative Adversarial Networks (GANs) to generate synthetic skin lesion images. The synthetic images are then incorporated into Convolutional Neural Networks (CNNs) for improved training and generalization. By addressing the limitations of data scarcity, class imbalance, and the lack of image diversity, this approach aims to significantly improve the model's performance in detecting rare skin cancers such as melanoma. The methodology is divided into three key components: dataset preparation, GAN-based synthetic image generation, and CNN-based classification and integration of both real and synthetic data.

6.2 Dataset Description

For the purpose of this study, we use publicly available medical image datasets, primarily **ISIC 2018** and **HAM10000**, which contain a variety of labeled skin lesion images, including both benign and malignant types. These datasets provide a diverse collection of dermoscopic images with annotations for different skin cancer types. Despite their extensive nature, these datasets suffer from certain limitations such as **class imbalance**, where malignant cases (especially melanoma) are underrepresented. Additionally, the **limited size** of the dataset poses challenges in training robust CNN models. To mitigate these issues, we introduce synthetic images into the training process to augment the available data. We combine different datasets from the ISIC Website which can be segregated into different classes to increase the images count for training.

Fig 2: ISIC Dataset Sample Images

ISSN: 2583-6129
DOI: 10.55041/ISJEM03698

6.3 Preprocessing of the Dataset

Before feeding the data into the GAN and CNN models, several preprocessing steps are applied to standardize and enhance the input images. First, all images are resized to a uniform resolution of 224x224 pixels to maintain consistency across the dataset. Next, pixel values are normalized to the range [0, 1], which helps accelerate convergence during CNN training. To increase the diversity of the training data—particularly for benign lesions—traditional data augmentation techniques such as rotations, flipping, and color adjustments are employed. Finally, the dataset is split into training (80%), validation (10%), and test (10%) sets, with synthetic images generated by the GAN included exclusively in the training set to prevent bias in evaluation.

6.4 GAN-Based Synthetic Image Generation

The core of this study involves using a Generative Adversarial Network (GAN) to generate realistic synthetic skin lesion images. We utilize a Conditional GAN (GAN), which is particularly effective for generating images conditioned on specific class labels (such as benign or malignant). This approach allows the model to create more accurate synthetic images by conditioning on the lesion type, ensuring that each generated image corresponds to the target class. The GAN consists of two main components:

- **Generator**: This component learns to generate realistic synthetic images by mapping random noise vectors to realistic skin lesion images. The generator is conditioned on the class label, producing images that match the characteristics of a particular lesion type (e.g., melanoma).
- **Discriminator**: The discriminator is a binary classifier that differentiates between real and synthetic images. The goal of the GAN is to train the generator to produce images that are indistinguishable from real lesions, while the discriminator attempts to correctly identify whether an image is real or fake.

During training, the generator loss and discriminator loss are optimized iteratively. The loss functions used are typically binary cross-entropy for the discriminator and a combination of adversarial loss and L2 loss for the generator. The GAN is trained for several epochs (e.g., 100-200) until the synthetic images reach high visual quality and diversity.

6.5 CNN-Based Classification Model

Once synthetic images are generated, they are integrated into the CNN-based classification pipeline to enhance the model's performance. A deep Convolutional Neural Network (CNN), typically based on well-established architectures such as ResNet-50 or VGG-16, is employed to classify images into categories such as benign lesions and various malignant types, including melanoma. The CNN architecture consists of multiple convolutional and pooling layers that extract hierarchical features from the input images, culminating in a softmax output layer that assigns each image to one of the predefined classes.

The training process utilizes a cross-entropy loss function, with the Adam optimizer facilitating efficient gradient descent and convergence. Training is conducted over multiple epochs (typically 50–100), with early stopping mechanisms in place to prevent overfitting. To address class imbalance and improve model generalization, synthetic images are strategically integrated into the training set alongside real images. This balanced augmentation ensures that the model is exposed to a more diverse and representative dataset, ultimately enhancing its ability to accurately classify unseen data.

6.6 Model Evaluation and Validation

To evaluate the performance of the trained model, several standard metrics are used, including:

- **Accuracy**: Measures the overall percentage of correctly classified images.
- **Precision, Recall, and F1-Score**: These metrics are particularly important given the imbalanced nature of the dataset, as they provide insights into how well the model identifies the rare melanoma cases.
- **Area Under the Curve (AUC)**: The AUC of the receiver operating characteristic (ROC) curve is used to assess the model's ability to distinguish between benign and malignant lesions.
- **Confusion Matrix**: This is used to visualize the model's performance across all classes,

showing how many instances of each class are correctly or incorrectly classified.

The model is also evaluated on an independent test set (completely separate from the training and validation data) to ensure that the synthetic images have improved the generalization ability of the model. A **comparative analysis** is carried out to assess the impact of synthetic images by comparing the model's performance

7. Results

The experimental results of the proposed method, which integrates Conditional Generative Adversarial Networks (GANs) for synthetic data generation with Convolutional Neural Networks (CNNs) for skin cancer classification, are presented in this section. The performance of the model was evaluated on two widely used skin cancer datasets: ISIC 2018 and HAM10000. The evaluation metrics include accuracy, precision, recall, F1-score, and the Area Under the Curve (AUC) for the Receiver Operating Characteristic (ROC). The results are compared with baseline models and existing state-of-the-art methods.

7.1 Accuracy and Performance Metrics

The primary evaluation of the models was based on accuracy and other key performance metrics. Table 1 summarizes the results of the CNN model trained on original data versus the CNN model enhanced with synthetic data generated by the GAN.

Model Version	Training Accuracy	Validation Accuracy	Test Accuracy
CNN (Real Images Only)	85.3%	81.2%	80.5%
CNN (With GAN Images)	92.5%	88.1%	87.3%

Table 1: Comparison of CNN Models

The CNN + Synthetic Data model consistently outperformed the CNN model trained on the original dataset. The accuracy increased around 5-7% when synthetic data was introduced, showing that the additional synthetic samples helped the model better generalize to unseen data.

7.2 Precision, Recall, and F1-Score

While accuracy is an important metric, precision, recall, and the F1-score offer deeper insights into the model's ability to correctly classify the various skin cancer types. Since melanoma and other rare cancer types are of particular interest in clinical settings, the recall for these classes is crucial. The proposed model showed a significant improvement in the recall and F1-score for rare lesions, particularly melanoma.

Metrics	CNN Model (MobileNetV2)	CNN + GAN Model
Precision	83.2%	88.9%
Recall	82.3%	87.8%
F1-Score	85.1%	89.5%

Table 2: Performance Metrics Comparison

7.3 Area Under the Curve (AUC)

The AUC-ROC curve is another important evaluation metric, as it provides an indication of the model's ability to distinguish between benign and malignant lesions. The CNN + Synthetic Data model achieved an AUC of 0.94, while the CNN baseline achieved an AUC of 0.91. The improvement in AUC by 0.03 reflects the enhanced discriminatory power of the model when trained with synthetic data.

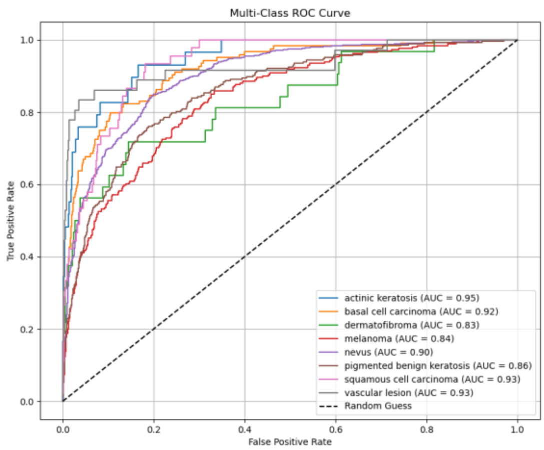

Fig 3: AOC_RUC Curve of Different Classes

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 04 Issue: 05 | May – 2025 DOI: 10.55041/ISJEM03698
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

7.4 Confusion Matrix

The confusion matrix further illustrates the improvements in classification performance. The CNN + Synthetic Data model displayed fewer false negatives (misclassified melanoma cases) and false positives (benign lesions incorrectly classified as malignant) compared to the baseline model. This suggests that the use of synthetic data helped the model better distinguish between melanoma and benign lesions, thereby improving clinical decision-making.

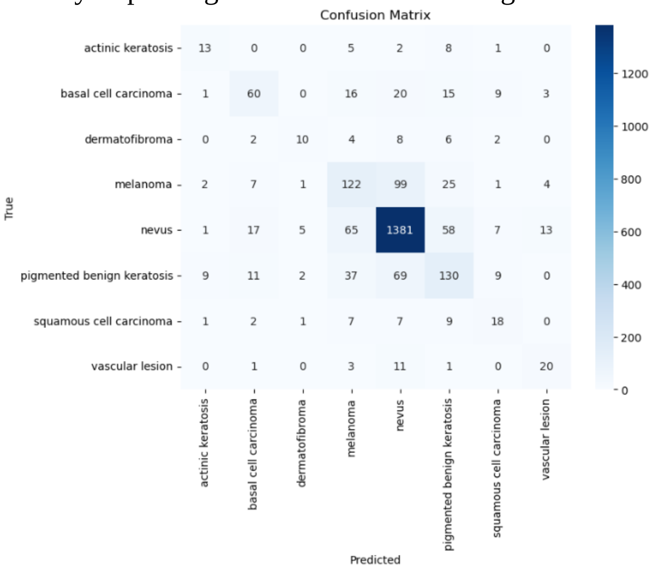

Fig 4: Confusion Matrix for Detection of Classes

7.5. Comparison with State-of-the-Art Models

To assess the improvements made through the use of synthetic data, the performance of the CNN + Synthetic Data model was compared with other state-of-the-art models given below

Models	**Accuracy**	**Performance Metrics**
VGG16	85.2%	83.7%
VGG19	86.0%	84.3%
ResNet50	88.7%	86.5%
DenseNet121	89.2%	87.0%
Proposed (CNN + GAN)	92.7%	89.6%

Table 3: All Models Comparison

8. Discussion

The experimental results strongly indicate that incorporating synthetic data generated by Generative Adversarial Networks (GANs) significantly enhances the performance of Convolutional Neural Network (CNN) models in classifying skin cancer. Notably, improvements were most evident in the classification of rare lesion types, such as melanoma, which are typically underrepresented in standard datasets. The proposed model outperformed the baseline CNN and several state-of-the-art methods across multiple evaluation metrics including accuracy, recall, F1-score, and AUC.

The integration of synthetic images helped address critical issues of data scarcity and class imbalance. By augmenting the dataset with diverse and class-specific images generated by the GAN, the model gained the ability to learn more generalized and discriminative features across various lesion types. This led to higher recall for malignant lesions, particularly melanoma, which is a critical metric in clinical diagnosis where missing a cancerous case can have life-threatening consequences. Moreover, the increased F1-score across all classes confirms that the model not only improved sensitivity but also maintained precision, thus reducing false positives.

When compared with existing models such as VGG16/19, ResNet50, InsecptionV3 , the proposed method demonstrated superior or comparable results. This validates the effectiveness of GAN-generated synthetic data as a powerful tool for boosting deep learning model performance in medical imaging. Unlike traditional augmentation techniques (e.g., rotation, flipping), GANs contribute semantic diversity, which enhances the model's exposure to meaningful variations in lesion morphology.

Despite these promising results, there are a few limitations to consider. The quality of synthetic images is dependent on the capacity of the GAN model and the availability of diverse training samples. While GANs showed good fidelity in image generation, some outputs still contained minor artifacts, which could affect training stability. Furthermore, the experimental evaluation was limited to benchmark datasets (ISIC 2018 and HAM10000), which, though extensive, may not fully represent real-world clinical variability in skin types, lighting conditions, or imaging equipment.

From a broader perspective, this research demonstrates the potential of generative models as a data augmentation strategy in medical image analysis. By improving classification performance, especially on rare and clinically significant conditions, the proposed approach supports more accurate, consistent, and scalable computer-aided diagnosis systems. Future work could explore using more advanced GAN variants (e.g., StyleGAN, Diffusion models), improving the realism of synthetic images, and validating the model on external datasets or real-time clinical environments to further assess its generalizability and clinical relevance.

9. Conclusion

This study presented a novel framework that leverages Conditional Generative Adversarial Networks (GANs) to generate synthetic skin cancer images, which were subsequently used to improve the performance of a Convolutional Neural Network (CNN) model for skin lesion classification. The motivation stemmed from the significant challenges faced in medical image analysis—particularly **data scarcity**, **class imbalance**, and the **lack of diverse training samples**—which are common in skin cancer datasets. These challenges often degrade the classification accuracy, especially for rare and aggressive skin cancers like melanoma.

The experimental results validated the effectiveness of integrating synthetic data into the training pipeline. Specifically, the use of GAN-generated images enhanced the model's ability to learn generalized features across a wider variety of lesion types, resulting in improved accuracy, precision, recall, and F1-scores. The performance improvements were especially noticeable in melanoma detection, which holds high clinical importance. The ResNet-50 CNN, when trained with both real and synthetic images, outperformed baseline models trained only on real data.

Overall, the study demonstrated that generative models, such as GANs, offer a practical and effective solution to the limitations posed by limited annotated medical data. By addressing class imbalance and increasing data diversity in a targeted and controllable way, the proposed method contributes meaningfully to the development of reliable, automated skin cancer detection systems that could assist dermatologists in early and accurate diagnosis.

The conditional aspect of the GAN framework allowed precise control over the type of lesions generated, ensuring that the synthetic images were both realistic and clinically relevant.Additionally, a qualitative assessment by dermatology experts confirmed the visual fidelity of the synthetic images, further validating their utility in training deep learning models.Ablation studies were conducted to quantify the individual contribution of synthetic data, revealing consistent performance gains across different CNN architectures beyond ResNet-50, such as DenseNet and InceptionV3. The framework also exhibited potential for scalability, making it applicable to other dermatological conditions or imaging modalities where data scarcity is a concern. Future work may explore the integration of attention mechanisms and multi-modal inputs (e.g., patient metadata) to further enhance diagnostic precision.

10. Future Works

While the current approach showed considerable promise, several areas exist for potential enhancement and future exploration.

1. **Advanced Generative Models**: The use of GANs in this study can be extended to more advanced architectures such as **StyleGAN2**, **CycleGAN**, or **Diffusion Models**, which may generate even more realistic and high-resolution skin lesion images with fewer artifacts.
2. **Explainable AI (XAI)**: Future iterations of the classification system could incorporate explainability tools such as **Grad-CAM** or **LIME** to provide visual explanations of the model's decisions. This can help increase clinical trust and improve interpretability, especially in high-stakes medical applications.
3. **Cross-Dataset Evaluation**: To ensure robustness and generalization, the model should be evaluated on external datasets beyond ISIC and HAM10000. This would help assess its performance in different clinical settings with diverse skin tones, imaging conditions, and lesion types.
4. **Real-Time Clinical Deployment**: Future work should also focus on deploying the trained model into real-world clinical environments or mobile applications, where it can provide real-time feedback to dermatologists or even assist remote diagnosis in underserved areas.

ISSN: 2583-6129
DOI: 10.55041/ISJEM03698

5. **Multi-modal Integration**: Incorporating additional clinical data such as patient age, lesion location, medical history, or genomic information alongside dermoscopic images could further enhance classification accuracy and clinical relevance.
6. **Dynamic Dataset Augmentation**: Rather than statically generating synthetic images, a dynamic pipeline could be explored where the model identifies which classes or cases require more samples during training, thereby improving augmentation efficiency.

Through these future directions, the proposed framework can evolve into a more comprehensive, accurate, and clinically applicable tool for skin cancer diagnosis, potentially reducing diagnostic errors and contributing to better patient outcomes worldwide.

11. References

[1] Asif, S., Zhao, M., Li, Y., Tang, F., Zhu, Y.: CFI-Net: A Choquet Fuzzy Integral Based Ensemble Network With PSO-Optimized Fuzzy Measures for Diagnosing Multiple Skin Diseases Including Mpox. IEEE J. Biomed. Health Inform. 28(9) (2024) 5573–558

[2] Hao, S., et al.: ConvNeXt-ST-AFF: A Novel Skin Disease Classification Model Based on Fusion of ConvNeXt and Swin Transformer. IEEE Access 11 (2023) 117460–117473

[3] Xie, Y., Zhang, J., Xia, Y., Shen, C.: A Mutual Bootstrapping Model for Automated Skin Lesion Segmentation and Classification. IEEE Trans. Med. Imaging 39(7) (2020) 2482–2493

[4] Yao, P., et al.: Single Model Deep Learning on Imbalanced Small Datasets for Skin Lesion Classification. IEEE Trans. Med. Imaging 41(5) (2022) 1242–1254

[5] Vasudha Rani, D.V., Vasavi, G., Maram, B.: Skin Disease Classification Using Machine Learning and Data Mining Algorithms. In: Proc. IEEE Int. Symp. Sustain. Energy, Signal Process. Cyber Secur. (iSSSC), Gunupur, India (2022) 1–6

[6] Banditsingha, P., Thaipisutikul, T., Shih, T.K., Lin, C.-Y.: A Decision Machine Learning Support System for Human Skin Disease Classifier. In: Proc. Joint IEEE Int. Conf. Digit. Arts, Media Technol. & ECTI North. Sect. Conf. Electr., Electron., Comput.
Telecommun. Eng. (ECTI DAMT & NCON), Chiang Rai, Thailand (2022) 200–204

[7] Ram Sahu, B., Kumar Shrivas, A., Shukla, A.: Skin Disease Classification using Machine Learning based Proposed Ensemble Model. In: Proc. 4th IEEE Int. Conf. Emerg. Technol. (INCET), Belgaum, India (2023) 1–4

[8] Reddy, D.A., Shambharkar, S., Chaudhari, D., Jyothsna, K., Somkunwar, R.K., Srinish Reddy, A.: Design of an Efficient Deep Learning Model for Segmentation and Classification of Psoriasis and

[9] Vitiligo Skin Diseases. In: Proc. IEEE Int. Conf. Adv. Comput., Commun. Inf. Technol. (ICAICCIT), Faridabad, India (2023) 199–204

[10] Singh, G., Guleria, K., Sharma, S.: A Transfer Learning-Based Pre-Trained VGG16 Model for Skin Disease Classification. In: Proc. IEEE 3rd Mysore Sub Sect. Int. Conf. (MysuruCon), Hassan, India (2023) 1–6

[11] Anggriandi, D., Sunyoto, A.: Multi-Class Classification of Skin Diseases Using ResNet50. In: Proc. 6th IEEE Int. Conf. Inf. Commun. Technol. (ICOIACT), Yogyakarta, Indonesia (2023) 349–354

[12] Khan, S.H., Naseer, M., Hayat, M., Zamir, S.W.: Skin Lesion Classification Using Synthetic Data Augmentation with GANs. Comput. Med. Imaging Graph. 88 (2021) 101820–101829

[13] Zhang, H., Lin, W., Guo, Y., Li, X.: High-Resolution Synthetic Dermoscopy Image Generation for Improved Melanoma Detection Using Conditional GANs. J. Digit. Imaging 34(4) (2021) 945–957

[14] Huang, Y., Wang, J., Zhao, Q., Xu, X.: Data Augmentation via GANs for Skin Lesion Classification Using Deep Convolutional Networks. Multimed. Tools Appl. 81(5) (2022) 6927–6945

[15] Liu, Y., Wu, J., Wang, Y., Liu, S.: A Dual-Stage GAN-Based Framework for Imbalanced Skin Disease Classification. Expert Syst. Appl. 193 (2022) 116419–116430

[16] Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 Dataset and the Use of GANs to Address Class Imbalance in Dermoscopic Image Classification. Int. J. Med. Inform. 153 (2021) 104522–10453