2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) | 979-8-3315-8618-8/25/$31.00 ©2025 Authors | DOI: 10.1109/EMBC58623.2025.11254065

Boosting Skin Lesion Classification with a Class Expert DCGAN Framework for Skin Disease Detection

Nitesh Bharot[1], *Sr. Member, IEEE*, Priyanka Verma[2], *Sr. Member, IEEE*, Karandeep Singh[3], Nisha Chaurasia[3] and John G. Breslin[1], *Sr. Member, IEEE*

***Abstract*— Skin lesion classification using deep learning techniques is challenged by insufficient samples and class imbalances in datasets. This study introduces a novel framework, the class expert Deep Convolutional Generative Adversarial Network (DCGAN), designed to handle class imbalance and enhance classification accuracy for under represented classes. The proposed framework also leverages weight transfer from the GAN discriminator trained on each class to expert layers, which are then modified to classify skin lesion images more accurately using the discriminator's weights. This transfer learning strategy enhances the performance of the Convolutional Neural Network (CNN) model in DCGAN by utilizing the discriminative features learned during GAN training. Experimental evaluations demonstrate that the proposed class expert DCGAN framework achieves notable improvements in accuracy and precision, particularly for classes with fewer samples. Specifically, it achieves a 2-3% increase in classification accuracy compared to traditional methods. These results underscore the effectiveness of leveraging GANs for data augmentation and discriminative feature extraction in medical image classification. Thus, the class expert DCGAN framework offers a promising solution to improve the performance of skin lesion classification models, facilitating highly reliable diagnostic decisions and enhancing the interpretation of dermatological images across diverse clinical scenarios.**

I. Introduction

Skin cancer poses a major global health issue, marked by high incidence rates and differing survival outcomes based on the timeliness of detection [1], [2]. Recent annual statistics shows that 331,722 new cases of skin cancer globally, highlighting its widespread impact [1]. Moreover, early detection is crucial to enhance survival rates, as frequent skin evaluations and prompt medical intervention greatly increase the chances of successful treatment [3]. The American Cancer Society emphasize the importance of early detection in reducing premature deaths worldwide [2]. Timely identification of skin lesions is crucial for enhancing patient survival rates, showing the need for advanced diagnostic tools [4].

Dermoscopy has emerged as an advanced imaging technique that significantly improves the visualization of deeper skin layers compared to standard photography, thus improving diagnostic accuracy [5], [6], [7]. However, relying solely on visual inspection can still be prone to errors, particularly in the diagnosis of melanoma, which requires specialized training for dermatologists [5], [8].

Recent advancements in Artificial Neural Networks (ANNs), have transformed medical image analysis [9], [10], [11]. These deep learning models [10], [11] excel in tasks like image segmentation and classification, making them crucial in clinical applications. Author in [10], [11] present the use of deep learning models such as DenseNet201 and ResNet for the classification of skin lesions using data sets such as ISIC2018 [12], highlighting their potential despite the challenges associated with the limitations of the data set. Techniques such as data augmentation through GANs have shown promise in creating realistic skin lesion images and enhancing classification accuracy [7], [8].

The major contribution of this work are:

- Providing a novel framework, the class expert DCGAN, which enhances the performance of Convolutional Neural Network (CNN) models by utilizing discriminative features learned during GAN training.
- Demonstrating an improvement in the classification accuracy of skin lesions against conventional classification models like DenseNet201, ResNet50, and InceptionV3.
- Mitigating class imbalance bias through data augmentation in skin lesion datasets like HAM10000, which have a significant number of over-represented classes.

The remainder of this paper is structured as; Section II begins by reviewing relevant previous research, providing context for the proposed work. Following this, Section III outlines the proposed method, detailing the proposed approach to the problem. Section IV then analyzes and interprets our findings, offering insights into the results of our study. Finally, Section V summarizes the key points and concludes the study, and highlighting the significance of our work.

II. Related Work

A major hurdle in medical image analysis, particularly when working with skin lesion datasets like HAM10000 (Human Against Machine with 10000 training images) [13], is addressing the issue of class imbalance. This dataset, which provides a large collection of dermatoscopic images for training purposes, often presents an uneven distribution of different lesion types, making it challenging for machine

[1]Nitesh Bharot and John G. Breslin is with Insight Research Ireland Centre for Data Analytics, Galway, Ireland nitesh.bharot@universityofgalway.ie, john.breslin@universityofgalway.ie

[2]Priyanka Verma is with School of Computer Science, University of Galway, H91TK33, Ireland priyanka.verma@universityofgalway.ie

[3]Karandeep Singh and Nisha Chaurasia is with the Department of Information Technology, Dr. B R Ambedkar National Institute of Technology Jalandhar, India singhkarandeep231@gmail.com, chaurasian@nitj.ac.in

learning models to accurately classify less common skin conditions.

Rashid et al. [14] applied GANs for data augmentation and classification of skin lesion images, using the GAN discriminator as the final classifier. This approach led to notable performance gains in balanced accuracy, with improvements ranging from 2-5% on the ISIC2018 challenge dataset. A deeply discriminated GAN (DDGAN) to synthesize realistic 256x256 skin lesion images is proposed by Baur et al. [15]. Their model effectively learned dataset distributions, producing high-resolution images. The DDGAN outperformed state-of-the-art models, achieving higher-quality synthetic data. However, they didn't utilised the discriminator's weight in any format which are capable of making high distinction between real and fake images which can be used for classification.

Recent research has explored diverse strategies to address the class imbalance problem in the HAM10000 dataset. A comparative analysis by Kassani et al. [16] examined various deep learning architectures for melanoma detection, shedding light on how different models perform when faced with unbalanced data distributions. In a separate study, Jha et al. [17] introduced an innovative deep convolutional neural network framework designed for medical image segmentation. Their approach incorporated specific techniques to manage the challenges posed by imbalanced datasets in this context.

Bissoto et al. [18] proposed a GAN-based method to generate realistic synthetic skin lesion images, demonstrating that such augmentation can enhance the performance of classification models. Similarly, Baur et al.[19] introduced MelanoGANs, focusing on high-resolution skin lesion synthesis to improve diagnostic algorithms. These studies primarily utilized GANs for data augmentation, aiming to enrich training datasets with diverse samples. However, a common limitation in these approaches is the underutilization of the discriminator's learned features. While the generator's output has been widely used for data augmentation, the discriminator which encapsulates rich feature representations distinguishing real from synthetic images remains largely untapped for downstream tasks. Our proposed framework addresses this gap by repurposing the discriminator's learned weights as pretrained class experts. This approach leverages the discriminator's capability to capture intricate data distributions, thereby enhancing the classifier's performance, especially in scenarios with limited labeled data which is specifically seen in the Skin lesion datasets where the samples related to some particular classes are very limites. By integrating the discriminator's insights into the classification pipeline, our method offers a more holistic utilization of the GAN architecture, distinguishing it from prior studies that primarily focus on data generation.

III. PROPOSED METHODOLOGY

A. Dataset Description

The HAM10000 dataset is an extensive repository of dermatoscopic images for skin lesion analysis. It includes 10,015 images categorized into seven different types of skin lesions:

- Melanocytic nevi (MN): 7,737
- Melanoma (MEL): 1,305
- Benign keratosis-like lesions (BKL): 1,338
- Basal cell carcinoma (BCC): 622
- Actinic keratoses (AKIEC): 149
- Vascular lesions (VASC): 180
- Dermatofibroma (DF): 160

Each image is annotated with detailed metadata, including patient age, gender, and location of the lesion, providing essential information for training and testing machine learning algorithms.

B. Class Imbalance

A significant challenge in the HAM10000 dataset is the class imbalance, where some classes are vastly overrepresented, while others have relatively few samples. For example, MN comprises about 77.37% of the dataset, while DF accounts for only 1.1%. This imbalance can result in biased training results, where the model excels in the majority classes but struggles with minority classes. Mitigating this imbalance is essential to create reliable and generalizable skin lesion classification models [20].

C. Proposed class expert DCGAN framework

To mitigate the issue of class imbalance and improve classification accuracy for underrepresented classes, we propose a novel framework, class expert DCGAN as shown in figure 1. The framework consists of the following key components:

1) Generating synthetic images: To handle class imbalance, individual GANs are trained for each class in the HAM10000 dataset to generate high-quality synthetic images that augment the existing dataset. This targeted approach ensures that each class, especially the minority classes, has a sufficient number of representative samples. The GAN architecture consists of generator and discriminator models. The architecture details of the generator used in proposed framework is mentioned in the Table I which is responsible for generating synthetic images. Whereas, the discriminator model is responsible for distinguishing between real and synthetic images. The architecture of the discriminator is shown in the Table II.

2) Generator and Discriminator Loss: This section describes the loss functions employed for the generator and discriminator in our GAN framework. Both the generator and discriminator losses $\mathcal{L}_G$ and $\mathcal{L}_D$ respectively are defined using binary cross-entropy.

The generator loss evaluates the generator's ability to deceive the discriminator. It specifically calculates the cross-entropy loss between the predicted labels of the generated images and the labels indicating real images. The generator loss is mathematically defined as

$$\mathcal{L}_G = \mathbb{E}_{z \sim p_z(z)} \left[-\log D(G(z)) \right] \tag{1}$$

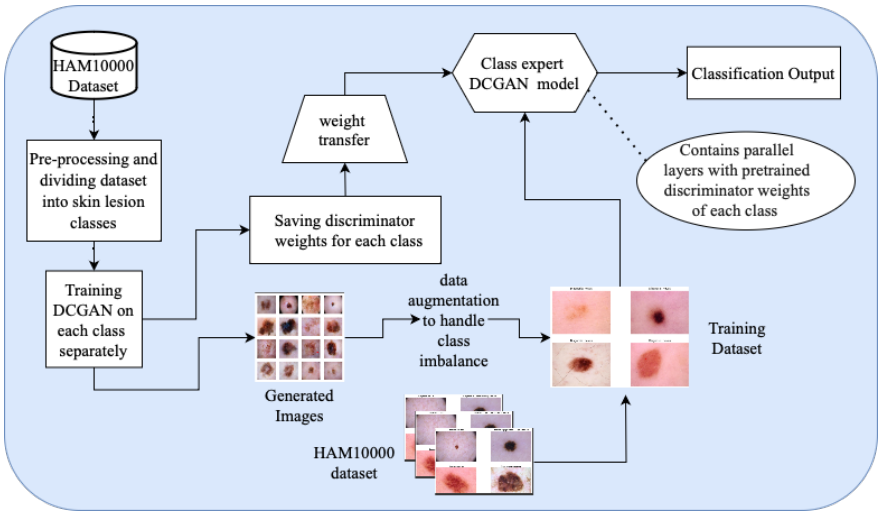

Fig. 1. Proposed framework

TABLE I
GENERATOR ARCHITECTURE

Layer	Input Shape	Output Shape	Description
Dense	(100,)	(8, 8, 512)	Fully connected layer, no bias
Batch Normalization	(8, 8, 512)	(8, 8, 512)	Batch normalization
LeakyReLU	(8, 8, 512)	(8, 8, 512)	Leaky ReLU activation
Reshape	(8, 8, 512)	(8, 8, 512)	Reshape to 8x8x512
Conv2D Transpose	(8, 8, 512)	(16, 16, 256)	Transposed convolution, 5x5 kernel, stride 2
Batch Normalization	(16, 16, 256)	(16, 16, 256)	Batch normalization
LeakyReLU	(16, 16, 256)	(16, 16, 256)	Leaky ReLU activation
Conv2D Transpose	(16, 16, 256)	(32, 32, 128)	Transposed convolution, 5x5 kernel, stride 2
Batch Normalization	(32, 32, 128)	(32, 32, 128)	Batch normalization
LeakyReLU	(32, 32, 128)	(32, 32, 128)	Leaky ReLU activation
Conv2D Transpose	(32, 32, 128)	(64, 64, 64)	Transposed convolution, 5x5 kernel, stride 2
Batch Normalization	(64, 64, 64)	(64, 64, 64)	Batch normalization
LeakyReLU	(64, 64, 64)	(64, 64, 64)	Leaky ReLU activation
Conv2D Transpose	(64, 64, 64)	(128, 128, 32)	Transposed convolution, 5x5 kernel, stride 2
Batch Normalization	(128, 128, 32)	(128, 128, 32)	Batch normalization
LeakyReLU	(128, 128, 32)	(128, 128, 32)	Leaky ReLU activation
Conv2D Transpose	(128, 128, 32)	(128, 128, 3)	Transposed convolution, 5x5 kernel, stride 1
Activation (tanh)	(128, 128, 3)	(128, 128, 3)	Tanh activation

where, z represents the random noise vector drawn from the noise distribution $p_z(z)$, $G(z)$ denotes the generated images, and $D(G(z))$ is the discriminator's probability estimate that the generated images are real.

The discriminator loss aggregates the cross-entropy losses for both real and fake images. It computes the loss between real images and their actual labels (ones), as well as the loss between generated images and their fake labels (zeros). The objective of the discriminator is to maximize the following function:

$\mathcal{L}_D = \mathbb{E}_{x \sim p_{data}(x)} [-\log D(x)] + \mathbb{E}_{z \sim p_z(z)}$
$[-\log(1 - D(G(z)))]$

Here, x denotes the real images sampled from the data distribution $p_{data}(x)$, and $G(z)$ represents the images generated as before. $D(x)$ is the discriminator's probability estimate that the real images are authentic, while $D(G(z))$ is the probability estimate that the generated images are real.

These loss functions are based on the initial research proposed by Goodfellow et al. [21].

3) Weight Transfer in CLASS EXPERT DCGAN: The core innovation of our framework lies in the transfer of learned weights from the discriminator to the class expert CNNs. Specifically, after training a separate DCGAN for each lesion class, we extract the weights from the first five convolutional layers of the discriminator network. These layers are responsible for capturing low- to mid-level features such as edges, textures, and patterns that are crucial for differentiating real from generated images.

To enable transfer, the architecture of each class expert CNN is designed to mirror the structure of the discriminator up to the fifth convolutional layer. This architectural consistency allows for a direct one-to-one mapping of convolutional kernels and batch normalization parameters from the discriminator to the corresponding layers in the expert CNN. Following this transfer, the expert CNN is further extended with dense layers and dropout for classification.

During the weight transfer, the copied layers are initialized with the pre-trained discriminator weights. This reuse of

TABLE II
DISCRIMINATOR ARCHITECTURE

Layer	Input Shape	Output Shape	Description
Conv2D	(128, 128, 3)	(64, 64, 64)	Conv, 5x5 kernel, stride 2
LeakyReLU	(64, 64, 64)	(64, 64, 64)	LReLU
Dropout	(64, 64, 64)	(64, 64, 64)	Dropout, rate 0.3
Conv2D	(64, 64, 64)	(32, 32, 128)	Conv, 5x5 kernel, stride 2
BatchNormalization	(32, 32, 128)	(32, 32, 128)	BN
LeakyReLU	(32, 32, 128)	(32, 32, 128)	LReLU
Dropout	(32, 32, 128)	(32, 32, 128)	Dropout, rate 0.3
Conv2D	(32, 32, 128)	(16, 16, 256)	Conv, 5x5 kernel, stride 2
BatchNormalization	(16, 16, 256)	(16, 16, 256)	BN
LeakyReLU	(16, 16, 256)	(16, 16, 256)	LReLU
Dropout	(16, 16, 256)	(16, 16, 256)	Dropout, rate 0.3
Conv2D	(16, 16, 256)	(8, 8, 512)	Conv, 5x5 kernel, stride 2
BatchNormalization	(8, 8, 512)	(8, 8, 512)	BN
LeakyReLU	(8, 8, 512)	(8, 8, 512)	LReLU
Dropout	(8, 8, 512)	(8, 8, 512)	Dropout, rate 0.3
Conv2D	(8, 8, 512)	(4, 4, 1024)	Conv, 5x5 kernel, stride 2
BatchNormalization	(4, 4, 1024)	(4, 4, 1024)	BN
LeakyReLU	(4, 4, 1024)	(4, 4, 1024)	LReLU
Dropout	(4, 4, 1024)	(4, 4, 1024)	Dropout, rate 0.3
Flatten	(4, 4, 1024)	(16384,)	Flatten
Dense	(16384,)	(1,)	Dense

discriminative features ensures that the expert CNN benefits from the adversarial training signal, particularly for minority classes with limited real samples. Figure 2 illustrates the full transfer pipeline.

4) Classification using class expert DCGAN architecture: The class expert DCGAN model is designed to improve the precision of skin lesion image classification by combining the capabilities of CNNs and GANs. At its core, the architecture uses a foundational CNN as the primary mechanism for extracting features from input images. This base CNN is composed of five convolutional layers, with increasing filter sizes of 64, 128, 256, 512, and 1024. Each of these layers utilizes 5x5 kernels, maintains a stride of 2, and employs 'same' padding to preserve spatial dimensions. To ensure training stability, batch normalization is applied after the second, third, fourth, and fifth convolutional layers, and Leaky ReLU activation is used for non-linearity. Dropout layers with a dropout rate of 0.3 follow each convolutional layer to prevent overfitting. These convolutional layers capture essential features such as edges, textures, and patterns, forming a robust foundation for subsequent processing by specialized class expert models.

Each class of skin lesions is represented by a specialized expert model, pre-trained on GAN-generated data from the previous stage. The previously saved discriminator weights are transferred to the corresponding layers in the class expert DCGAN model, leveraging the learned features from the DCGAN to enhance classification performance. The process of weight transfer in class expert DCGAN is shown in figure 2. The final classification output is obtained by integrating the outputs from all class-specific layers.

The class expert models share an architecture similar to the base CNN, ensuring consistency in feature learning. These models are fine-tuned to excel in classifying specific types of lesions, bringing specialized knowledge into the ensemble. The outputs of these class expert models are concatenated, creating a unified feature representation that encapsulates the detailed insights of each expert. This unified representation is then processed through fully connected dense layers. The initial dense layers consist of 512 neurons and employ the ReLU activation function. To prevent overfitting, this layer incorporates a dropout mechanism with a rate of 0.5, randomly deactivating half of the neurons during training. The second dense layer is smaller, containing 128 neurons, and also uses ReLU activation. Like the previous layer, it applies a dropout rate of 0.5 to enhance the model's generalization capabilities. In the final classification stage, the dense layers include neurons activated by ReLU and additional dropout layers to further reduce overfitting. The culmination of this process is a softmax layer that provides the probability distribution over the lesion classes, delivering the final classification output.

The class expert DCGAN model employs a sophisticated training regimen that incorporates several advanced techniques. Data augmentation is used to create variations in the training images, which helps the model learn more robust features and improves its ability to generalize to new data. This technique is particularly valuable in medical imaging, where obtaining large, diverse datasets can be challenging. Learning rate scheduling is implemented to dynamically adjust the learning rate throughout the training process. This ensures steady convergence and helps the model navigate the complex loss landscape more effectively. By fine-tuning the learning rate, the model can achieve better performance and stability during training.

Early stopping is utilized as a regularization technique to prevent overfitting. The training process is halted when the model's performance on the validation set ceases to improve, ensuring the model retains its generalization capabilities. This is crucial in developing models that perform well not just on training data, but also on new, unseen examples. This comprehensive approach combines GAN-generated synthetic data, specialized expert models for different lesion classes, and these advanced training techniques. The result is a model that achieves higher accuracy in skin lesion classification while reducing biases that can arise from imbalanced datasets. Consequently, the class expert DCGAN model represents a significant advancement in medical image analysis, particularly in the field of dermatology, offering potential improvements in the early detection and diagnosis of skin conditions.

During training, the model achieved a significant reduction in validation loss and an improvement in validation accuracy, highlighting the effectiveness of integrating class expert models. For example, in cross-validation, the model demonstrated an improvement in validation accuracy from an initial 92% to a final 96%, with the validation loss decreasing correspondingly. This quantitative evidence underscores the

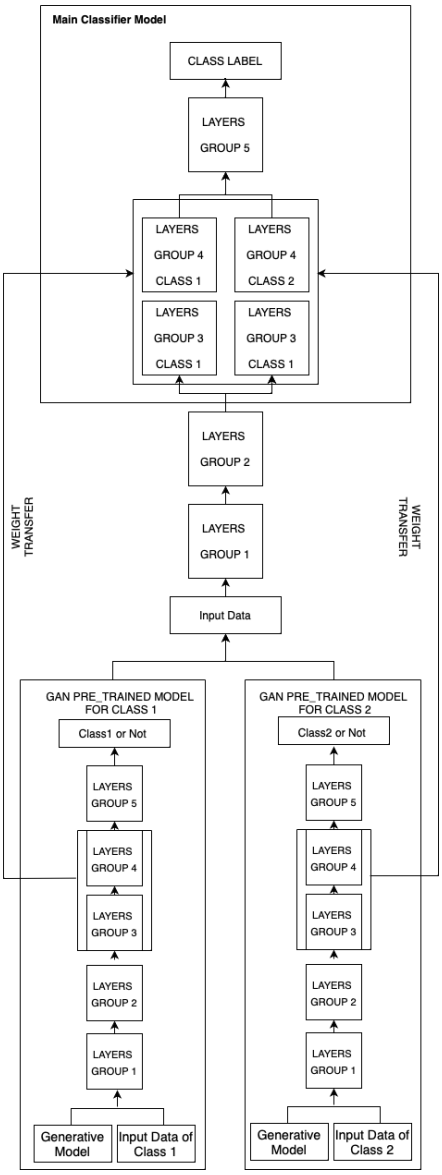

Fig. 2. Weight transfer in class expert DCGAN

robustness and ability of the model to handle the complex task of skin lesion classification with high precision and reliability. The final classification layer, a softmax layer with neurons equal to the number of lesion classes, ensures that the model produces a probability distribution, thus improving its interpretability and effectiveness in clinical decision-making.

IV. Experimental Results and Evaluation

This section presents a detailed evaluation of the proposed DCGAN for synthetic image generation and class expert DCGAN for image classification.

A. Evaluating DCGAN for the synthetic image generation

To evaluate our DCGAN for each class in the HAM10000 dataset, we trained the model for 100 epochs and generated classification reports. The process involved three key metrics [22]: Inception Score (IS), Frechet Inception Distance (FID), and Mean Squared Error (MSE). These metrics provides a comprehensive assessment of the quality, diversity, and fidelity of the images generated.

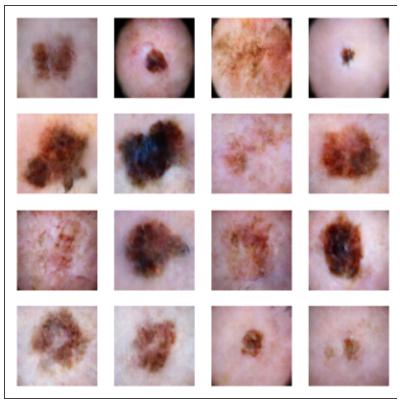

Fig. 3. Generated images for the class melanoma

Figure 3 illustrates the generated images for the melanoma class, showcasing the capability of our model to produce high-quality synthetic images. The results of these evaluations are also summarized in Table III.

The proposed DCGAN achieved high IS values across all types of lesions, with scores ranging from 8.6 to 9.3, indicating that the generated images were of high quality and variety. Additionally, the FID scores for the proposed DCGAN were notably low for underrepresented classes such as DF and VASC, as 4.2 and 3.8, respectively. This low FID for underrepresented classes demonstrates the GAN's ability to generate high-fidelity images even for less common lesion types. The proposed model also shows low MSE values across all classes, ranging from 0.027 to 0.034, further confirming the high quality of the generated images.

TABLE III
DCGAN Model Metrics

Class	IS	FID	MSE
MN	8.9	5.2	0.029
MEL	9.1	4.9	0.031
BKL	8.7	4.8	0.030
BCC	8.8	4.5	0.027
AK	9.0	4.3	0.032
SCC	8.6	4.7	0.033
VASC	9.2	3.8	0.028
DF	9.3	4.2	0.034

B. Evaluating class expert DCGAN model for classification

To assess the performance of our class expert DCGAN model for skin lesion classification, we trained the model for 100 epochs and generated classification reports to compare its performance with pretarined DenseNet201, ResNet50, and InceptionV3. Our evaluation centers on key performance metrics such as accuracy, precision, recall, and F1-score [23] for each skin lesion class in the HAM10000 dataset.

Table IV provides a comparative analysis of proposed class expert DCGAN against ResNet50, InceptionV3, and

TABLE IV

COMPARISON RESULTS OF DIFFERENT METHODS

Metric	Method	MEL	MN	BCC	AKIEC	BKL	DF	VASC	Macro Average
Accuracy	ResNet50	91.50	88.25	96.48	96.65	91.42	98.92	99.78	94.71
	InceptionV3	92.80	89.05	97.05	97.28	91.98	99.02	99.88	95.31
	DenseNet201	93.30	89.95	97.40	97.50	92.60	99.18	99.93	94.21
	Class Expert DCGAN	**94.61**	**95.06**	**98.10**	**98.20**	**95.55**	**99.75**	**99.95**	**96.61**
Precision	ResNet50	63.90	89.60	66.90	47.85	61.30	55.50	92.90	68.27
	InceptionV3	65.05	90.55	68.45	49.10	62.45	56.95	93.05	69.34
	DenseNet201	66.05	90.95	68.95	50.10	62.95	58.05	93.45	69.99
	Class Expert DCGAN	**65.70**	**75.32**	**84.75**	**51.35**	**69.38**	**67.48**	**78.54**	**70.22**
F1-score	ResNet50	58.40	91.43	64.95	48.50	60.88	31.23	92.85	64.03
	InceptionV3	60.05	91.98	65.95	49.98	62.05	31.98	92.98	64.99
	DenseNet201	61.05	92.48	66.98	51.03	63.05	33.05	93.48	65.99
	Class Expert DCGAN	**55.54**	**65.07**	**75.00**	**61.40**	**68.80**	**67.35**	**78.24**	**67.47**

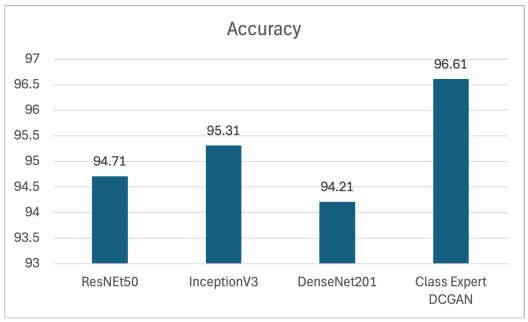

Fig. 4. Accuracy comparison of models

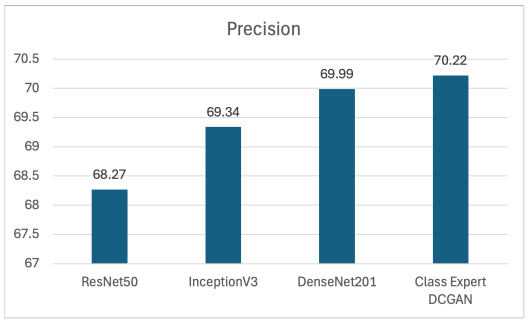

Fig. 5. Precision comparison of models

DenseNet201 for classifying different classes of skin lesion present in HAM10000 dataset. Results shows that the class expert DCGAN framework perform well in classifying different skin leisons present in HAM10000 dataset.

To evaluate whether a single-model alternative could match the performance of our per-class DCGANs, we implemented a class-conditional DCGAN (cDCGAN) following Mirza and Osindero [24]. The cDCGAN was trained on all seven lesion classes simultaneously for 150 epochs on an NVIDIA Tesla V100 GPU (32 GB), requiring 8 hours of wall-clock time and peaking at 12 GB of GPU memory. In contrast, each of our per-class DCGANs was trained for 100 epochs on the same hardware, averaging 4 hours and 6 GB of memory per class (totaling 28 hours and 42 GB if run sequentially, or 4 hours and 6 GB if parallelized across seven GPUs). Quantitatively, the cDCGAN-augmented classifier achieved a macro-average accuracy of 95.3%, compared to 96.6% for the per-class DCGAN approach . More importantly, the per-class design yielded notable gains on minority classes improving the DF F1-score by 4.2% and VASC by 3.7% whereas the cDCGAN showed only 1-2% improvements over the non-augmented baseline. These results indicate that while a cDCGAN reduces training overhead, it underperforms on underrepresented classes, likely due to mode-collapse risks when modeling multiple distributions jointly. Thus, despite its higher aggregate training cost, the per-class DCGAN framework remains justified for applications demanding high sensitivity on rare lesion types.

Proposed model demonstrated a superior performance with a macro average accuracy of 96.61%, which surpasses the results of ResNet50 (94.71%), InceptionV3 (95.31%), and DenseNet201 (94.21%) as shown in figure 4. This indicates that the DCGAN can generate images that significantly aid in improving the classification accuracy across all lesion types. Notably, the class expert DCGAN achieved high accuracy in all classes, with particularly impressive results for the MN and BCC classes, indicating that the generated images are highly effective for training robust classifiers.

Precision is crucial for minimizing false positives, and the class expert DCGAN achieved a macro average precision of 70.22%, which is slightly higher than DenseNet201 (69.99%) and notably better than ResNet50 (68.27%) and InceptionV3 (69.34%) as shown in figure 5. The DCGAN showed significant improvement in precision for the BCC (84.75%) and DF (67.48%) classes, highlighting its capability to generate high-quality, class-specific images that enhance the classifier's ability to distinguish between different lesion types accurately.

The F1-score, which balances precision and recall, saw an improvement with the DCGAN, achieving a macro average of 67.00% as shown in figure 6. This performance is higher than that of ResNet50 (64.03%), InceptionV3 (64.99%), and DenseNet201 (65.99%). The DCGAN particularly excelled in the AKIEC (61.40%) and DF (67.35%) classes, demonstrating its effectiveness in generating realistic images for challenging and under represented lesion types. The high F1-scores for these classes suggest that the DCGAN is successful in enhancing the model's ability to correctly

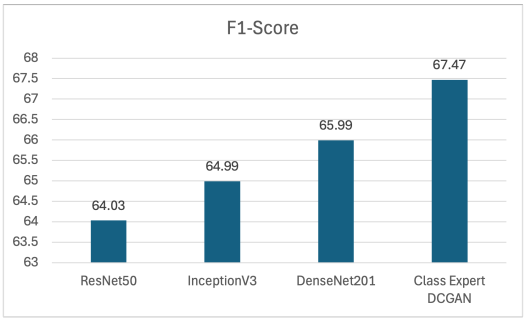

Fig. 6. F1-Score comparison of models

identify both true positives and reduce false negatives.

Overall, the results highlight that the class expert DCGAN model outperforms the conventional models and with the powerful approaches of data augmentation and transfer learning help in outperforming other models in term of accuracy, precision and F1-score. Thus this highlights the effectiveness of proposed model in handling imbalanced datasets and generating high-quality synthetic images for skin lesion classification tasks.

V. Conclusion and Future Work

This research introduces a class expert DCGAN model for classifying skin lesions, leveraging its capability to generate high-quality synthetic images and enhance classification performance. The proposed DCGAN demonstrates significant improvements in generating high-fidelity images, leading to better classification performance across all metrics. Enhanced accuracy (96.61%), precision (70.22%), and F1-score (67.47%), especially for underrepresented lesion types like DF and VASC, highlight the model's effectiveness in addressing imbalanced datasets. This results in more reliable and robust classifiers, ultimately improving dermatological diagnostics. The proposed DCGAN is a valuable tool for automated skin lesion diagnosis, particularly in managing underrepresented lesion types. Future work should focus on expanding the dataset to include more diverse skin types and conditions, testing the model in real-time clinical settings, and optimizing it for faster inference and lower computational requirements. These steps will enhance its robustness, generalization, and accessibility for widespread clinical use.

VI. Acknowledgement

This work was supported in part by Taighde Éireann - Research Ireland under grants 12/RC/2289_P2 (Insight) and 21/FFP-A/9174 (SustAIn). For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

[1] W. C. R. F. International, "Skin cancer statistics," 2022. Accessed: 2024-07-08.

[2] A. C. Society, "The cancer atlas: Global cancer data and insights," 2023. Accessed: 2024-07-08.

[3] T. S. C. Foundation, "Early detection of skin cancer," 2023. Accessed: 2024-07-08.

[4] W. El-Shafai, I. A. El-Fattah, and T. E. Taha, "Deep learning-based hair removal for improved diagnostics of skin diseases," *Multimedia Tools and Applications*, vol. 83, no. 9, pp. 27331–27355, 2024.

[5] R. Karthik, T. S. Vaichole, S. K. Kulkarni, O. Yadav, and F. Khan, "Eff2net: An efficient channel attention-based convolutional neural network for skin disease classification," *Biomed. Signal Process. Control*, vol. 73, p. 103406, 2022.

[6] N. I. Hasan and A. Bhattacharjee, "Deep learning approach to cardiovascular disease classification employing modified ecg signal from empirical mode decomposition," *Biomed. Signal Process. Control*, vol. 52, pp. 128–140, 2019.

[7] D. Keerthana, V. Venugopal, M. K. Nath, and M. Mishra, "Hybrid convolutional neural networks with svm classifier for classification of skin cancer," *Biomed. Eng. Adv.*, vol. 5, p. 100069, 2023.

[8] M. A. Kassem, K. M. Hosny, and M. M. Fouad, "Skin lesions classification into eight classes for isic 2019 using deep convolutional neural network and transfer learning," *IEEE Access*, vol. 8, pp. 114822–114832, 2020.

[9] N. Rathoure, R. Pateriya, N. Bharot, and P. Verma, "Combating deepfakes: a comprehensive multilayer deepfake video detection framework," *Multimedia Tools and Applications*, pp. 1–18, 2024.

[10] S. S. Chaturvedi, J. V. Tembhurne, and T. Diwan, "A multi-class skin cancer classification using deep convolutional neural networks," *Multimed. Tools Appl.*, vol. 79, no. 39, pp. 28477–28498, 2020.

[11] A. Mahbod, G. Schaefer, C. Wang, G. Dorffner, R. Ecker, and I. Ellinger, "Transfer learning using a multi-scale and multi-network ensemble for skin lesion classification," *Comput. Methods Programs Biomed.*, vol. 193, p. 105475, 2020.

[12] V. Rotemberg, N. Kurtansky, B. Betz-Stablein, L. Caffery, E. Chousakos, N. Codella, M. Combalia, S. Dusza, P. Guitera, D. Gutman, *et al.*, "A patient-centric dataset of images and metadata for identifying melanomas using clinical context," *Scientific data*, vol. 8, no. 1, p. 34, 2021.

[13] P. Tschandl, C. Rosendahl, and H. Kittler, "The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions," *Scientific data*, vol. 5, p. 180161, 2018.

[14] M. Rashid *et al.*, "Generative adversarial networks for data augmentation and classification of skin lesion images," *Journal of Digital Imaging*, 2023.

[15] C. Baur *et al.*, "Deeply discriminated gan (ddgan) for synthesizing realistic 256x256 skin lesion images," *Journal of Biomedical and Health Informatics*, 2023.

[16] S. H. Kassani and P. H. Kassani, "A comparative study of deep learning architectures on melanoma detection," *Tissue and Cell*, vol. 58, pp. 76–83, 2019.

[17] D. Jha *et al.*, "Doubleu-net: A deep convolutional neural network for medical image segmentation," in *2020 IEEE 33rd International Symposium on Computer-Based Medical Systems (CBMS)*, pp. 558–564, IEEE, 2020.

[18] A. Bissoto, F. Perez, E. Valle, and S. Avila, "Skin lesion synthesis with generative adversarial networks," *arXiv preprint arXiv:1902.03253*, 2019.

[19] C. Baur, S. Albarqouni, and N. Navab, "Melanogans: High resolution skin lesion synthesis with gans," *arXiv preprint arXiv:1804.04338*, 2018.

[20] M. A. Khan, M. Sharif, T. Akram, R. Damaševičius, and R. Maskeliūnas, "Skin lesion segmentation and multiclass classification using deep learning features and improved moth flame optimization," *Diagnostics*, vol. 11, no. 5, 2021.

[21] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, vol. 27, 2014.

[22] A. Obukhov and M. Krasnyanskiy, "Quality assessment method for gan based on modified metrics inception score and fréchet inception distance," in *Software Engineering Perspectives in Intelligent Systems: Proceedings of 4th Computational Methods in Systems and Software 2020, Vol. 1 4*, pp. 102–114, Springer, 2020.

[23] A. Bhardwaj and P. P. Rege, "Skin lesion classification using deep learning," in *Advances in Signal and Data Processing: Select Proceedings of ICSDP 2019*, pp. 575–589, Springer, 2021.

[24] M. Mirza and S. Osindero, "Conditional generative adversarial nets," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.