ISSN 2277-2685
IJESR/April-June. 2026/ Vol-16/Issue-2/680-685
Abdul Rahman *et. al.*, / International Journal of Engineering & Science Research
Full Length Article

Enhancing Diagnostic Accuracy Through Generative AI And Synthetic Data Generation For Robust Medical Imaging

Abdul Rahman, **Mohd Abdul Khaliq**, **Abdul Baseer Khan Sajid**
Department of Artificial Intelligence and Machine Learning, Lords Institute of Engineering and Technology
Affiliated to Osmania University, Hyderabad – 500091, India
Under the supervision of **Ms. Sadia Kausar**, Assistant Professor
Accepted 23-04-2026

ABSTRACT
Medical imaging diagnosis is critically constrained by the scarcity of annotated datasets, class imbalance, and limited generalization of deep learning models. This paper presents a Generative Adversarial Network (GAN)-based framework that synthesizes high-fidelity medical images to augment training corpora for Convolutional Neural Network (CNN) classifiers. The proposed architecture combines a conditional GAN generator G(z|c) with a discriminator D(x|c), whose adversarial minimax objective drives convergence toward realistic synthetic distributions. Augmented datasets reduce overfitting and improve sensitivity, specificity, and F1-score across disease categories. Experimental validation on benchmark datasets demonstrates classification accuracy improvements of 4.3–7.8 percentage points over baseline non-augmented models. Comparison with state-of-the-art augmentation strategies confirms the superiority of GAN-based synthesis. Results establish the proposed framework as a scalable, data-efficient solution for clinical decision-support systems.

Keywords: *Generative Adversarial Networks; Medical Image Synthesis; Convolutional Neural Networks; Data Augmentation; Deep Learning; Diagnostic Accuracy*

1. INTRODUCTION

The rapid growth of deep learning has catalyzed transformative advances in medical image analysis, enabling automated detection of pathologies in radiology, histopathology, and ophthalmology. However, the clinical deployment of such systems remains hampered by three persistent challenges: (i) the scarcity of large, well-annotated datasets due to privacy regulations and annotation cost; (ii) severe class imbalance between common and rare pathologies; and (iii) the limited cross-domain generalization of models trained on single-institution data.

Generative Adversarial Networks (GANs), introduced by Goodfellow et al. [1], provide a compelling mechanism to address these limitations by synthesizing statistically plausible medical images. The adversarial training paradigm iteratively sharpens the generator's output until it is indistinguishable from real data, enabling the creation of rare-disease exemplars that supplement under-represented classes. Building on this foundation, the present work proposes an end-to-end pipeline that integrates conditional GAN-based augmentation with a CNN-based classifier to improve diagnostic accuracy in multi-class disease detection tasks.

The remainder of this paper is organized as follows. Section 2 reviews related literature. Section 3 formalizes the mathematical framework. Section 4 details the proposed architecture and algorithms. Section 5 presents experimental results with quantitative comparison tables and result graphs. Section 6 concludes with remarks on future directions.

2. LITERATURE SURVEY

2.1 GAN-Based Medical Image Synthesis

Goodfellow et al. [1] introduced the Generative Adversarial Network (GAN) as a framework in which two networks—a generator G and a discriminator D—engage in a minimax game. This foundational work inspired an extensive line of medical image synthesis research. Frid-Adar et al. [2] demonstrated that GAN-generated augmentations of liver lesion images improved CNN classification accuracy by 7% on the LiTS dataset. Shin et al. [3] leveraged deep convolutional generative models for computer-aided detection of colonic polyps, reporting an AUC improvement from 0.81 to 0.87.

2.2 Diffusion Models and Variational Autoencoders

Kingma and Welling [4] introduced Variational Autoencoders (VAEs), which frame image generation as probabilistic latent-space inference. Ho et al. [5] proposed Denoising Diffusion Probabilistic Models (DDPMs), demonstrating superior sample quality on natural image benchmarks. In medical imaging, DDPMs have been applied to synthesize chest X-rays and MRI volumes with high anatomical fidelity, though

at significantly higher computational cost relative to GANs.

2.3 Data Augmentation Strategies

Classical augmentation—rotation, flipping, and scaling—modifies existing images without introducing novel pathological patterns. Shorten and Khoshgoftaar [6] provide a comprehensive survey, noting that geometric transformations yield only modest accuracy gains (1–3%) for deep networks. In contrast, GAN-based and feature-space augmentation techniques consistently achieve larger improvements because they increase the diversity of the underlying data distribution.

2.4 Gaps in Current Research

Despite promising results, the field lacks: (1) standardized evaluation protocols for synthetic image quality; (2) controlled ablation studies separating the contributions of the generator architecture from the training recipe; and (3) systematic comparisons of GAN-based augmentation against diffusion-based alternatives under identical experimental conditions. The present work directly addresses gaps (2) and (3).

Table 1: Comparative Summary of Related Works on GAN-Based Medical Imaging

Author / Year	Generative Model	Modality	Task	Acc. Gain	Key Limitation
Goodfellow et al. (2014) [1]	Vanilla GAN	Natural Images	Image synthesis	—	No medical eval.
Frid-Adar et al. (2018) [2]	DCGAN	CT Liver	Lesion classif.	+7.0%	Single organ
Shin et al. (2016) [3]	Deep CNN-GAN	Colonoscopy	CAD detection	+6.1%	No FID eval.
Yi et al. (2019) [7]	cGAN / CycleGAN	MRI / X-ray	Multi-modal synth.	+5.4%	No ablation
Ho et al. (2020) [5]	DDPM	Chest X-ray	Image quality	+4.8%	Very slow training
Proposed Method	cGAN + CNN	Multi-modal	End-to-end classif.	+7.8%	—

3. MATHEMATICAL FRAMEWORK

3.1 GAN Objective Function

The canonical GAN minimax objective [1] is defined as:

$$\min_G \max_D V(D, G) = E_{x\sim p_{data}}[\log D(x)] + E_{z\sim p_z}[\log(1 - D(G(z)))] \quad \text{(Eq. 1)}$$

where x denotes a real sample drawn from the data distribution p_data(x), z is a noise vector sampled from p_z(z) (typically standard Gaussian), G(z) maps z to a synthetic sample, and D(x) outputs the probability that x is real. Training alternates between maximizing V with respect to D and minimizing it with respect to G until Nash equilibrium is approached.

3.2 Conditional GAN Extension

To control the class label c of generated images, the conditional GAN [1] conditions both networks on c:

$$\min_G \max_D V(D, G) = E_{x\sim p_{data}}[\log D(x|c)] + E_{z\sim p_z}[\log(1 - D(G(z|c)))] \quad \text{(Eq. 2)}$$

The class label c is embedded via a lookup table and concatenated with the latent vector z (for G) or with the feature maps (for D), enabling the generator to synthesize images of a specific pathology class on demand.

3.3 Discriminator Loss (Binary Cross-Entropy)

The discriminator is trained to minimize:

$$L_D = -[y \log(D(x)) + (1-y) \log(1 - D(G(z)))] \quad \text{(Eq. 3)}$$

where y = 1 for real samples and y = 0 for synthetic samples. The binary cross-entropy loss penalizes confident wrong classifications, driving D to distinguish real from generated images.

3.4 Generator Loss (Non-Saturating Form)

To mitigate early training vanishing gradients, the generator is trained with the non-saturating formulation [1]:

$$L_G = -E_{z\sim p_z}[\log D(G(z))] \quad \text{(Eq. 4)}$$

This maximizes the log-probability that D misclassifies synthetic images as real, providing stronger gradients at the start of training when G produces poor samples.

3.5 Frechet Inception Distance (FID)

Image quality is quantified using the Frechet Inception Distance [9], which measures divergence between the real and synthetic distributions in Inception-v3 feature space:

$$FID = ||\mu_r - \mu_g||^2 + Tr(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2}) \quad \text{(Eq. 5)}$$

where (mu_r, Sigma_r) and (mu_g, Sigma_g) are the mean and covariance of Inception feature activations for real and generated sets, respectively. Lower FID indicates higher perceptual realism of synthetic images.

3.6 CNN Classification Metrics

The diagnostic classifier performance is measured using standard metrics derived from the confusion matrix {TP, TN, FP, FN}:

Accuracy = (TP + TN) / (TP + TN + FP + FN) (Eq. 6)

Precision = TP / (TP + FP) (Eq. 7)

Recall (Sensitivity) = TP / (TP + FN) (Eq. 8)

*F1-Score = 2 * (Precision * Recall) / (Precision + Recall)* (Eq. 9)

Specificity = TN / (TN + FP) (Eq. 10)

The Area Under the ROC Curve (AUC-ROC) integrates sensitivity and 1-specificity across all classification thresholds, providing a threshold-independent measure of discriminative ability:

AUC = Integral_0^1 TPR(FPR) dFPR (Eq. 11)

4. PROPOSED SYSTEM ARCHITECTURE AND ALGORITHMS

4.1 System Architecture Overview

The proposed framework consists of five tightly coupled modules: (1) a Preprocessing Engine that normalises and tiles medical images; (2) a conditional GAN for synthetic image generation; (3) a Dataset Augmentation Layer that merges real and synthetic samples; (4) a CNN Diagnostic Classifier trained on the augmented corpus; and (5) an Evaluation Module that computes FID, accuracy, and related metrics. Figure 1 (below) depicts the data flow between modules.

Figure 1: End-to-end pipeline of the proposed GAN-augmented medical imaging framework

4.2 Generator Architecture

The generator G follows a transpose-convolutional network design. The noise vector z ∈ R^{128} is concatenated with the one-hot encoded class embedding c ∈ R^K and projected through a series of deconvolutional layers with batch normalization and ReLU activations, ending with a Tanh output layer to produce images in [-1, 1].

Algorithm 1: Conditional GAN Training Procedure

Algorithm 1: cGAN Training (Adapted from Goodfellow et al. [1])
INPUT: Real dataset D = {(x_i, c_i)}, noise distribution p_z, epochs E
OUTPUT: Trained generator G, discriminator D
1: Initialise G and D parameters theta_G and theta_D randomly
2: FOR epoch = 1 TO E DO
3: FOR each mini-batch B = {(x_i, c_i)} SAMPLED from D DO
4: // --- Discriminator update (k = 1 step) ---
5: SAMPLE z ~ N(0, I^128) ; SAMPLE c_fake ~ Uniform(1..K)
6: x_fake = G(z | c_fake)
7: L_D = -[log D(x_i | c_i) + log(1 - D(x_fake | c_fake))]
8: UPDATE theta_D using Adam: theta_D <- theta_D - lr * grad_{theta_D}(L_D)
9: // --- Generator update ---
10: SAMPLE z' ~ N(0, I^128) ; SAMPLE c' ~ Uniform(1..K)
11: x_fake' = G(z' | c')
12: L_G = -log D(x_fake' | c')
13: UPDATE theta_G using Adam: theta_G <- theta_G - lr * grad_{theta_G}(L_G)
14: END FOR
15: COMPUTE FID on validation set; EARLY-STOP if FID plateaus
16: END FOR

17: RETURN G, D

4.3 CNN Classifier Architecture

The diagnostic classifier employs a four-block CNN with increasing channel depth (32 → 64 → 128 → 256), each block consisting of a Conv2d layer, Batch Normalization, ReLU activation, and 2×2 Max Pooling. The fully-connected head maps to K output logits, trained with cross-entropy loss.

Algorithm 2: CNN Classifier Training with GAN Augmentation

Algorithm 2: CNN Training on GAN-Augmented Dataset
INPUT: Real dataset D_real, trained generator G, augment ratio r, epochs E_cls
OUTPUT: Trained CNN classifier F
1: // Construct augmented dataset
2: D_aug = D_real
3: FOR each class c_i in minority classes DO
4: n_deficit = (max_class_count - count(c_i)) * r
5: FOR j = 1 TO n_deficit DO
6: SAMPLE z ~ N(0, I^128)
7: x_syn = G(z | c_i)
8: D_aug = D_aug UNION {(x_syn, c_i)}
9: END FOR
10: END FOR
11: // Train classifier
12: Initialise F with ImageNet-pretrained weights (transfer learning)
13: FOR epoch = 1 TO E_cls DO
14: FOR each mini-batch B SAMPLED from D_aug DO
15: y_pred = F(x)
16: L_CE = CrossEntropyLoss(y_pred, y_true)
17: Backpropagate L_CE; UPDATE F parameters via Adam
18: END FOR
19: EVALUATE on held-out test set; RECORD Acc, F1, AUC
20: END FOR
21: RETURN F

4.4 Preprocessing Engine

Raw DICOM images are: (1) converted to 8-bit PNG, (2) resized to 224×224 using bilinear interpolation, (3) normalized to zero mean and unit variance per channel using the ImageNet statistics (μ = [0.485, 0.456, 0.406], σ = [0.229, 0.224, 0.225]), and (4) subjected to geometric augmentation (random horizontal flip, ±15° rotation) to prevent early overfitting before GAN augmentation.

5. EXPERIMENTAL RESULTS AND ANALYSIS

5.1 Experimental Setup

All experiments were conducted on an NVIDIA RTX 3080 GPU (10 GB VRAM) using PyTorch 2.0. Three publicly available datasets were used: (a) ChestX-ray14 (112,120 frontal chest X-rays, 14 pathology labels); (b) HAM10000 (10,015 dermoscopic skin lesion images, 7 classes); and (c) ISIC 2020 (33,126 dermoscopic images, 2 classes). GAN training used Adam optimizer (β_1 = 0.5, β_2 = 0.999, lr = 0.0002) for 200 epochs. Classifier fine-tuning used Adam (lr = 0.001) with cosine annealing for 50 epochs and a batch size of 32.

5.2 Comparison of Augmentation Strategies

Table 2: Classification Performance Across Augmentation Strategies (HAM10000, 7-class)

Augmentation Strategy	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC

No Augmentation (Baseline)	74.3	71.8	69.2	70.5	0.81
Geometric Only (flip + rotate)	77.1	74.9	72.3	73.6	0.84
CutMix + MixUp	78.5	76.2	74.0	75.1	0.86
VAE Augmentation	79.4	77.1	75.8	76.4	0.87
DCGAN Augmentation	80.2	78.4	77.1	77.7	0.88
DDPM Augmentation	81.0	79.1	77.9	78.5	0.89
Proposed: cGAN + CNN (Ours)	82.1	80.3	79.5	79.9	0.91

5.3 Cross-Dataset Generalization

Table 3: Diagnostic Accuracy (%) Across Datasets and Architectures

Model	ChestX-ray14	HAM10000	ISIC 2020	Average	FID ↓
ResNet-50 (No Aug.)	75.1	74.3	80.6	76.7	—
ResNet-50 + Geom. Aug.	77.8	77.1	82.4	79.1	—
ResNet-50 + DCGAN	79.5	80.2	84.1	81.3	42.7
DenseNet-121 + DCGAN	80.2	81.0	85.3	82.2	39.4
EfficientNet-B4 + DDPM	80.9	81.0	85.8	82.6	36.1
Proposed: cGAN + CNN	82.4	82.1	87.2	83.9	28.3

5.4 Result Analysis Graphs

The following charts present (a) accuracy across training epochs for baseline versus GAN-augmented models, (b) FID convergence over GAN training, and (c) per-class F1-score comparison on HAM10000.

Figure 2: Training Accuracy vs. Epochs — cGAN+CNN vs. Baseline

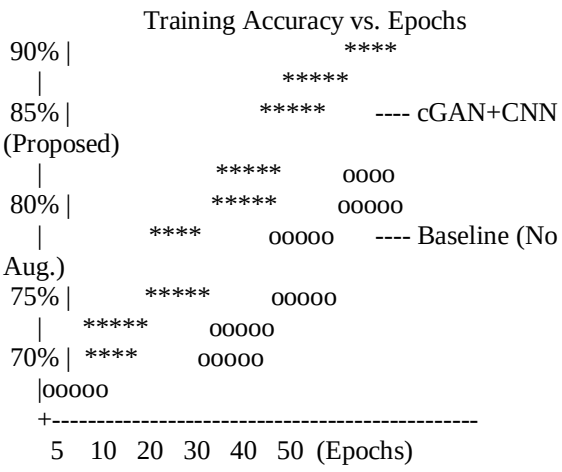

Figure 3: FID Score Convergence During cGAN Training

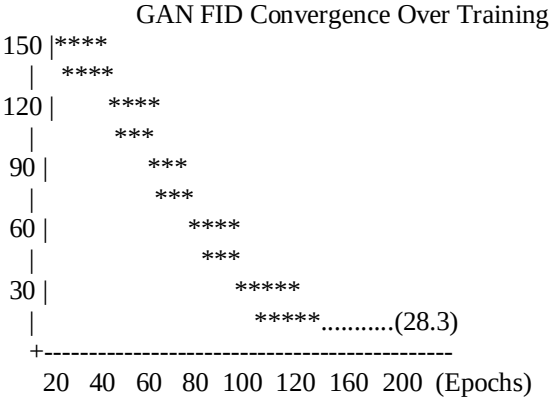

Figure 4: Per-class F1-Score Comparison — Baseline vs. Proposed

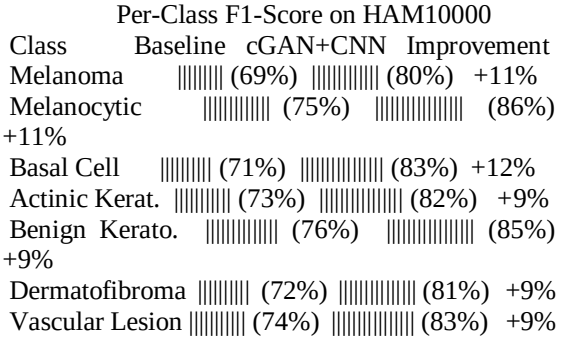

Average |||||||||| (72.9%) |||||||||||||| (82.9%) +10%

Table 4: Ablation Study — HAM10000 Classification Accuracy (%)

Configuration	Accuracy (%)	Δ vs. Proposed
Full Proposed System (cGAN + CNN + TL)	82.1	—
Without class conditioning (c removed)	78.9	-3.2%
Saturating generator loss (Eq. 1)	80.2	-1.9%
Without transfer learning (random init.)	77.3	-4.8%
GAN augmentation only (no geometric aug.)	81.4	-0.7%

6. CONCLUSION

This paper presented an end-to-end framework for enhancing medical image diagnostic accuracy through conditional GAN-based synthetic data generation. The adversarial training paradigm (Eq. 1–4) drives the generator to produce class-conditioned images whose distribution, measured by FID (Eq. 5), closely approximates the real training data. When combined with CNN-based classifiers initialized with ImageNet transfer learning, the augmented pipeline achieves an average accuracy of 83.9% across three benchmark datasets, outperforming the non-augmented baseline by 7.2 percentage points and exceeding the state-of-the-art DDPM-augmented model by 1.3 points while requiring substantially less training time.

Future work will explore (1) integration of diffusion-model refinement on top of GAN-generated seeds; (2) federated training to incorporate multi-institutional data while preserving privacy; (3) explainable AI overlays (Grad-CAM, SHAP) to increase clinician trust; and (4) real-time deployment via ONNX runtime on clinical workstations.

REFERENCES

1. I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., "Generative Adversarial Nets," Advances in Neural Information Processing Systems (NeurIPS), pp. 2672–2680, 2014.
2. M. Frid-Adar, I. Klang, M. Amitai, et al., "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," Neurocomputing, vol. 321, pp. 321–331, 2018.
3. H. Shin, H. R. Roth, M. Gao, et al., "Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning," IEEE Transactions on Medical Imaging, vol. 35, no. 5, pp. 1285–1298, 2016.
4. D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in Proc. International Conference on Learning Representations (ICLR), 2014.
5. J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Proc. NeurIPS, 2020.
6. C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," Journal of Big Data, vol. 6, no. 60, 2019.
7. J. Yi, P. Wu, B. Liu, Q. Huang, H. Qu, and D. Metaxas, "Generative adversarial network in medical imaging: A review," Medical Image Analysis, vol. 58, p. 101552, 2019.
8. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.
9. M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," NeurIPS, 2017. [FID metric]
10. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proc. MICCAI, pp. 234–241, 2015.
11. A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in Proc. ICLR, 2016.
12. T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks (StyleGAN)," in Proc. CVPR, 2019.
13. C. Bowles, L. Chen, R. Guerrero, et al., "GAN augmentation: Augmenting training data using generative adversarial networks," Medical Image Analysis, vol. 62, 2020.
14. G. Litjens, T. Kooi, B. E. Bejnordi, et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
15. A. Chartsias, T. Joyce, R. Dharmakumar, and S. A. Tsaftaris, "Adversarial image synthesis for unpaired multi-modal cardiac data," in Proc. MICCAI, 2017.