Jurnal Ilmu Komputer Agri-Informatika

Tersedia secara daring di:
https://jurnal.ipb.ac.id/index.php/jika

Volume 12 Nomor 2 halaman 167 – 177
eISSN: 2654-9735, pISSN: 2089-6026

DCGAN-Based Synthetic Data Augmentation for Cervical Intraepithelial Neoplasia (CIN) Images

AINI NABILAH[1]*, SITI NURMAINI[2]

Abstract

A significant challenge in medical imaging is the limited availability of high-quality datasets. Generative Artificial Intelligence (Generative AI) addresses this issue by generating synthetic medical images to augment existing datasets. This study investigates the use of Deep Convolutional Generative Adversarial Networks (DCGAN) for data augmentation in Cervical Intraepithelial Neoplasia (CIN) imaging. The dataset used in this study consists of 233 training images with a resolution of 256×256 pixels, illustrating the typical limitations of small-scale medical datasets. Two training scenarios were implemented: DCGAN with manual data augmentation and DCGAN without manual augmentation. Image quality was evaluated using the Fréchet Inception Distance (FID). The results indicate that incorporating data augmentation improves training stability and enhances the quality of generated images, achieving an FID of 2.21. In contrast, training DCGAN without manual augmentation produced a higher FID score of 2.52, indicating lower image quality. These findings highlight the effectiveness of DCGAN for medical image augmentation and its potential to enhance deep learning-based diagnostic models for cervical cancer detection and classification.

*Keywords**:** cervical intraepithelial neoplasia, data augmentation, DCGAN, generative adversarial networks*

INTRODUCTION

Medical imaging plays a vital role in disease diagnosis and treatment by providing essential insights into pathological conditions. Accurate diagnosis relies not only on the quality of image acquisition but also on effective image interpretation. However, assembling large and diverse medical image datasets remains a significant challenge, primarily due to ethical considerations, patient privacy regulations, and the limited availability of annotated data (Pinaya *et al.* 2023). This problem is especially common in *cervical intraepithelial neoplasia* (CIN), a condition where abnormal cells grow on the surface of the cervix and may lead to cervical cancer. It is often hard to collect a large and high-quality set of CIN images. CIN refers to a range of precancerous abnormalities linked to the progression of invasive cervical carcinoma. CIN is classified into three stages: CIN1, CIN2, and CIN3 (Wang *et al.* 2024). However, high-quality CIN datasets are scarce, especially for early-stage lesions. This scarcity limits the development of reliable deep learning models for automated CIN detection and classification.

Researchers attempt to address this issue by applying data augmentation techniques that transform existing images through operations such as rotation, cropping, and resizing. While these approaches can enhance model robustness, they do not substantially increase the intrinsic diversity of the dataset (Goodfellow *et al.* 2020). The introduction of redundant or slightly altered samples has become a standard procedure to improve network training performance in computer vision applications (Smaida *et al.* 2021).

This limitation highlights a key research gap: conventional augmentation alone is insufficient to address the lack of diversity in CIN datasets, underscoring the need for more advanced augmentation strategies capable of generating truly data instances. Generative Artificial Intelligence provides a promising direction for alleviating data scarcity, particularly

[1] Computer Science, Faculty of Computer Science, Universitas Sriwijaya, Palembang, South Sumatra
[2] Intelligent System Research Group, Faculty of Computer Science, Universitas Sriwijaya, Palembang, South Sumatra
*Correspondence Author: Surel: aininabilahalfatah@gmail.com

within the domain of medical image analysis. Generative AI refers to a class of artificial intelligence methods that train model to synthesize new data samples that resemble the distribution of the original dataset. Unlike traditional discriminative models, this AI focuses on generating new content, such as images, music, text, or video, rather than performing classification tasks (Musalamadugu & Kannan 2023). Recently, one of the most significant methods developed is Generative Adversarial Networks (GANs) (Goodfellow *et al.* 2020).

GANs represent a powerful approach for generating synthetic medical images. Rooted in the concept of zero-sum games from game theory, GANs consist of two primary components: a generator and a discriminator. These components are trained in opposition, with the generator aiming to create realistic images and the discriminator learning to differentiate between real and synthetic images. This adversarial training process enables the network to model the underlying data distribution and generate new, realistic samples.

Since their introduction, GANs have attracted considerable attention in medical imaging due to their ability to generate realistic synthetic data. The introduction of GANs by Goodfellow *et al.* (2020) marked a significant advancement in deep learning-based data synthesis, enabling the production of high-quality images from learned data distributions. Several studies have explored the use of GANs for medical data augmentation, addressing the challenge of the limited availability of large and diverse annotated datasets. Various techniques have been proposed across different medical fields, including Computed Tomography (CT) (Liu *et al.* 2021; Meor Yahaya & Teo 2023; Zhao *et al.* 2023), Magnetic Resonance Imaging (MRI) (Jiang *et al.* 2021; Alrashedy *et al.* 2022; Zhang *et al.* 2022), and X-Rays (Ciano *et al.* 2021; Motamed *et al.* 2021; Sundaram & Hulkund 2021). Several GAN variants have been developed to improve the quality of generated images. Methods such as CycleGAN (Zhu *et al.* 2017), DualGAN (Yi *et al.* 2017), CGAN (Mirza & Osindero 2014), and Progressive Growing GAN (PGGAN) (Karras *et al.* 2017), have demonstrated effectiveness in generating high-quality medical images for augmentation purposes. Additional GAN architectures have been proposed to enhance image synthesis quality, Motamed *et al.* (2021) introduced the Inception-Augmentation GAN (IAGAN), inspired by DAGAN (Antoniou *et al.* 2017), which enhances the performance of other GAN architectures and significantly improves model accuracy for data augmentation. Among these, Deep Convolutional GAN (DCGAN) has been recognized for its stability and effectiveness in generating high-quality medical images (Radford *et al.* 2015). Radford *et al.* introduced deep convolution layers and batch normalization in DCGAN, resulting in improved training convergence and image realism. Zhang *et al.* (2023) proposed WGAN-GP to address few-shot imbalance datasets, thereby improving the performance of classification models when training data are insufficient and unevenly distributed. Huang *et al.* (2021) explored the GAN-Based data augmentation to generate realistic 256 x 256 brain tumor MR images with diversity on a small amount of training data for more diversity and exceeding performance of GAN-based models. Zhao *et al.* (2024) demonstrated that cycleGAN-based style transfer approach can accurately convert ultrasound images from different devices into standardized format, thereby improving image quality and the performance of radiomics studies.

Among existing GAN variants, DCGAN has been widely recognized for its stability and effectiveness in generating high-quality medical images (Smaida *et al.* 2021). Smaida *et al.* (2021) used a DCGAN and GMD model and confirmed that with the help of DCGAN, the GMD model (CNN model) can yield a better accuracy compared to traditional methods. Wubineh *et al.* (2024) proposed RES_DGAN data augmentation for classification of cervical cells, showed that the performance of the model slightly improves with DCGAN in both classification problems using theResNet50V2 and Xception architectures. Similarly, Devi and Kumar (2022) demonstrates the use of DCGAN method to expand a retina image dataset, successfully generating higher-resolution and higher-quality images that significantly enhanced downstream classification accuracy.

Despite these advancements, the use of DCGAN specifically for augmenting CIN images particularly CIN1, which is frequently underrepresented remains limited. Furthermore, the comparative impact of training DCGAN with and without additional manual augmentation have not been systematically investigated. To address these gaps, the present study evaluates the effectiveness of DCGAN for augmenting CIN1 medical images and to compare the synthetic image quality produced under two training scenarios: (1) DCGAN trained with manual augmentation and (2) DCGAN trained without manual augmentation. Image quality is assessed using the Fréchet Inception Distance (FID). By generating synthetic CIN1 images that capture the underlying distribution of real lesions, DCGAN directly addresses the scarcity of CIN datasets, forming the basis for the contributions presented in this study.

METHOD

The proposed framework for CIN image augmentation using DCGAN is illustrated in Figure 1. The framework consists of five main stages: data preparation, preprocessing, augmentation, model training, and evaluation for generated images.

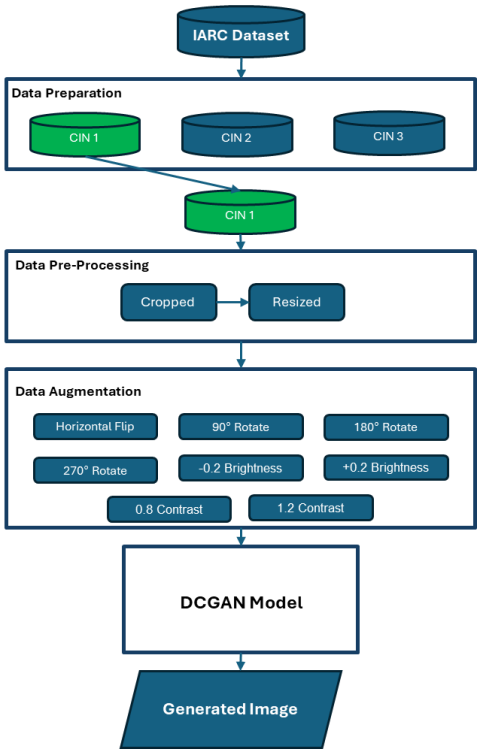

Figure 1 Research methodology

Each stage is designed to ensure that the generated images are realistic, diverse, and suitable for classification tasks. During data preparation, images are selected from the IARC dataset, with a focus on the CIN 1 category. In the preprocessing stage, images are cropped and resized to ensure consistency. Data augmentation is then applied to increase dataset variability through transformations such as rotation, flipping, brightness, and contrast adjustments. The DCGAN model is then trained to generate high-quality synthetic images. Finally, the generated images are evaluated to determine their similarity to real samples, improving their suitability for classification.

a. Data Preparation

The dataset used in this study was obtained from the International Agency for Research on Cancer (IARC) (International Agency for Research on Cancer 2021) and includes two types of cervical cancer screening examinations: Colposcopy and Visual Inspection with Acetic Acid (VIA). The colposcopy dataset includes 202 cases, with each case containing four images: an untreated image, an image after acetic acid application, an image after iodine application, and an image after saline application, yielding a total of 913 images. The VIA dataset comprises

187 cases, each containing two images: one before acetic acid application and one after, totalling 420 images.

For this study, all images were categorized into CIN 1, CIN 2, and CIN 3. Only CIN1 images were selected, as this category is typically underrepresented and clinically challenging. Image captured before and after acetic acid application were included to highlight acetowhite lesions, which indicate early precancerous changes. An example of CIN 1 images shown at Figure 2. After categorization and filtering, the final dataset used for model training consists of 233 images.

Figure 2 Sample images of CIN 1

b. Pre-Processing

The images in the dataset contained unnecessary frames, noise along the edges, and visible examination tools or lenses. To address these issues, each image was manually cropped to retain only the cervix region, as shown in Figure 3. This step ensured that irrelevant areas were removed, focusing solely on the medically significant regions.

Figure 3 Before cropping

Figure 4 After cropping

Figure 3 shows the image before cropping and figure 4 shows the image after cropping and taking only the cervix area. Following cropping, all images were resized to a uniform dimension of 512×512 pixels to ensure consistency during model training. Figure 5 shows an example of images after cropping and resizing. These preprocessing steps helped to standardized the dataset and improved input quality for the DCGAN model.

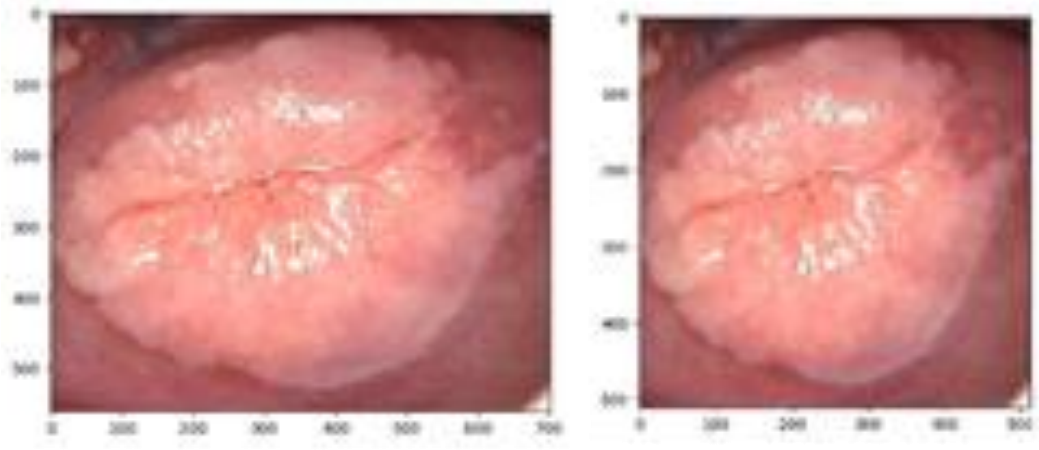

Figure 5 Before resized (left) and after resized (right)

c. Data Augmentation

To further enhance dataset diversity, nine types of augmentation were applied to each image. These augmentations included horizontal flipping, three rotational transformations (90°, 180°, and 270°), two brightness adjustments (–0.2 and +0.2), and two contrast modifications (0.8× and 1.2×). Each transformation was applied independently, resulting in nine augmented versions per image. The original images were retained, producing a total of 233 × 10 = 2330 training samples. The complete augmented dataset was then used as the input for DCGAN training. In both training scenarios (with and without augmentation), the models were optimized using the Adam optimizer with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. The augmented model used a learning rate of 1×10^{-4}, batch size 32, and latent vector dimension 256, while the non-augmented model used a learning rate of 2×10^{-4}, batch size 16, and latent dimension 128. Both models were trained for 1500 epochs. To stabilize adversarial training, label smoothing was applied (0.8 for real labels and 0.2 for fake labels), and gradient clipping in the range [–1.0, 1.0] was used to prevent exploding gradients. Although no explicit random seed was enforced, training behavior remained consistent across runs.

Quantitative evaluation was performed using the Fréchet Inception Distance (FID), computed every 500 epochs using a 1:1 ratio of real to generated images. For each FID calculation, N = 233 real CIN1 images were compared against 233 generated images, both resized to 299 × 299 pixels prior to feature extraction using InceptionV3.

d. Deep Convolutional Generative Adversarial Networks (DCGAN)

DCGAN is an improved version of the traditional GAN architecture that incorporates deep convolutional layers to enhance the quality and stability of image generation. Introduced by Radford *et al.* (2015), DCGAN addresses challenges in training standard GANs by replacing fully connected layers with deep convolutional networks, using batch normalization to stabilize training, and adopting LeakyReLU as an activation function in the discriminator. Similar to standard GAN, DCGAN consists of two primary components: a generator and a discriminator. The generator network that takes random noise z (a random number) and generates images from this noise, while the discriminator evaluates whether an image is real or generated (Liu *et al.* 2022). Figure 6 illustrates the interaction between the discriminator and generator in this study, showing how adversarial training drives both networks to improve iteratively.

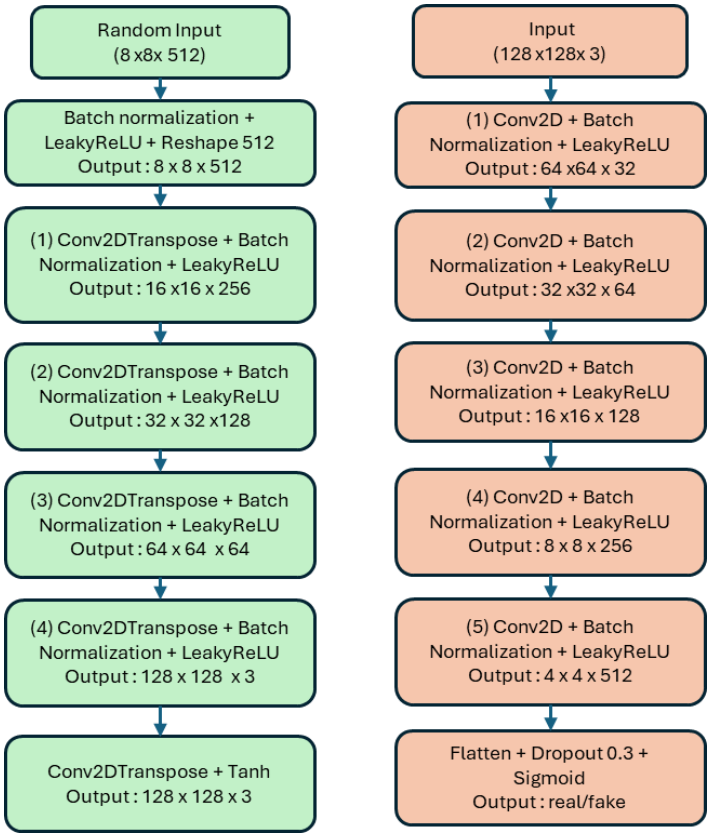

Figure 6 Generator function (G) (left) and discriminator function (d) (right)

The generator network converts random noise vectors into synthetic CIN images using a series of up-sampling layers. The architecture begins with a random noise vector of size 256, which is processed by a dense layer with 8 × 8 × 512 units and no bias parameters. The bias term is excluded because the subsequent batch normalization layer applies a learned shift (β) and scale (γ), rendering the bias redundant and preventing parameter duplication. After the dense layer, batch normalization and LeakyReLU activation ($\alpha = 0.2$) are applied, and the output is reshaped into 8 × 8 × 512 feature maps. Four transposed convolution blocks then perform progressive up-sampling as shown in Figure 6. The output is then reshaped to 8 × 8 × 512 feature maps. Four transposed convolution blocks perform progressive up-sampling, as shown at Figure 6. Each block uses a 4 × 4 kernel with a stride of 2, followed by batch normalization and LeakyReLU activation. The output layer generates a 3-channel image with tanh activation, producing RGB images in the range [-1, 1]. The discriminator network determines whether an image is real or generated. It receives a 128 × 128 × 3 RGB image as input and processes it through five convolutional blocks for progressive down-sampling as shown at Figure 6. All blocks except the first use batch normalization and LeakyReLU ($\alpha = 0.2$). The output from the convolutional layers is flattened, passed through a dropout layer for regularization, and finally through a single-unit dense layer with sigmoid activation to yield a probability score.

e. Fréchet Inception Distance (FID)

The quality of generated images was primarily evaluated using the FID score (Dowson & Landau 2003). Heusel *et al.* (2017) demonstrated that the FID is consistent in measuring increased disturbances and aligns with human judgment, also measures the similarity between the distribution of generated images and real images. Lower FID scores indicate better quality and more realistic synthetic images. The calculation process involved using a pre-trained InceptionV3 model (excluding the top classification layer) to extract features, resizing both real and generated images to 299 × 299 pixels, calculating means (μ) and covariances (Σ) of the feature distributions, and computing the FID using the formula:

$$d^2(F,G) = |\mu_1 - \mu_2|^2 + tr\left[\Sigma_1 + \Sigma_2 - 2\sqrt{(\Sigma_1\Sigma_2)}\right] \quad (1)$$

where μ_1, Σ_1 are statistics of real images and μ_2, Σ_2 are statistics of generated images. Both generator and discriminator losses were tracked throughout training to monitor convergence and stability.

RESULTS AND DISCUSSION

In this study, DCGAN model was trained under two different scenarios: (1) with data augmentation (including flipping, rotation, brightness, and contrast adjustments) and (2) without data augmentation. Both models were trained for 1500 epochs. The evaluation was conducted by analyzing the generator and discriminator loss curves as well as the FID scores. Throughout the training process, the generator loss initially exhibited high fluctuations, reflecting the difficulty of early-stage adversarial learning. However, the loss gradually stabilized around epoch 1000, indicating that the generator had become increasingly effective at producing images that could fool the discriminator. The discriminator loss showed a similar pattern of early instability followed by convergence, suggesting that both networks eventually reached a balanced adversarial state. The training results shown that at the beginning of training, the generator produced images with unclear structures. As the number of epochs increased, the quality of the resulting images improved with more realistic anatomical details as shown in Figure 7.

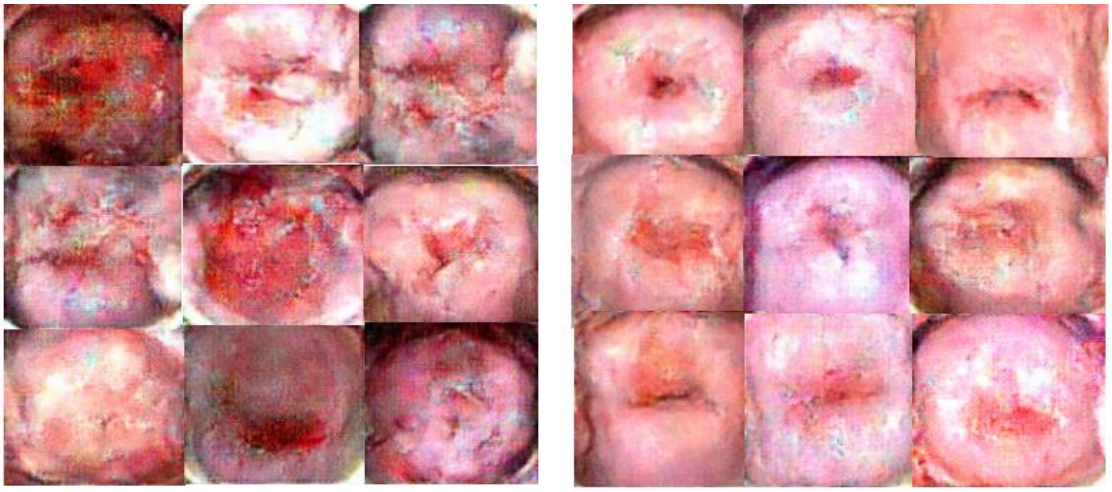

Figure 7 Sample result with data augmentation, generated data at epoch 500 (left) and generated data at epoch 1500 (right)

It can be seen from Figure 7 that there is still a lot of noise and the image is still unclear at epoch 500, but at epoch 1500, the generator successfully produced images that exhibited more coherent structure and reduced noise levels. The use of label smoothing (0.8 for real images and 0.2 for fake images) and gradient clipping ([-1.0, 1.0]) proved effective in preventing mode collapse and training instability, which are common challenges in GAN training. These stabilization strategies contributed to consistent improvements in image quality across the training process. In contrast, images produced without adding image augmentation exhibited higher noise levels, a less coherent visual structure, and more pronounced artifacts, as shown in Figure 8. In the early training stages around epoch 500, the generated images in both with data augmentation and without data augmentation contained the basic color profile of CIN cells but lacked proper cellular structure and showed significant artifacts. By epoch 1500, the model was able to generate more realistic cellular patterns, though some abnormalities in texture and shape were still present.

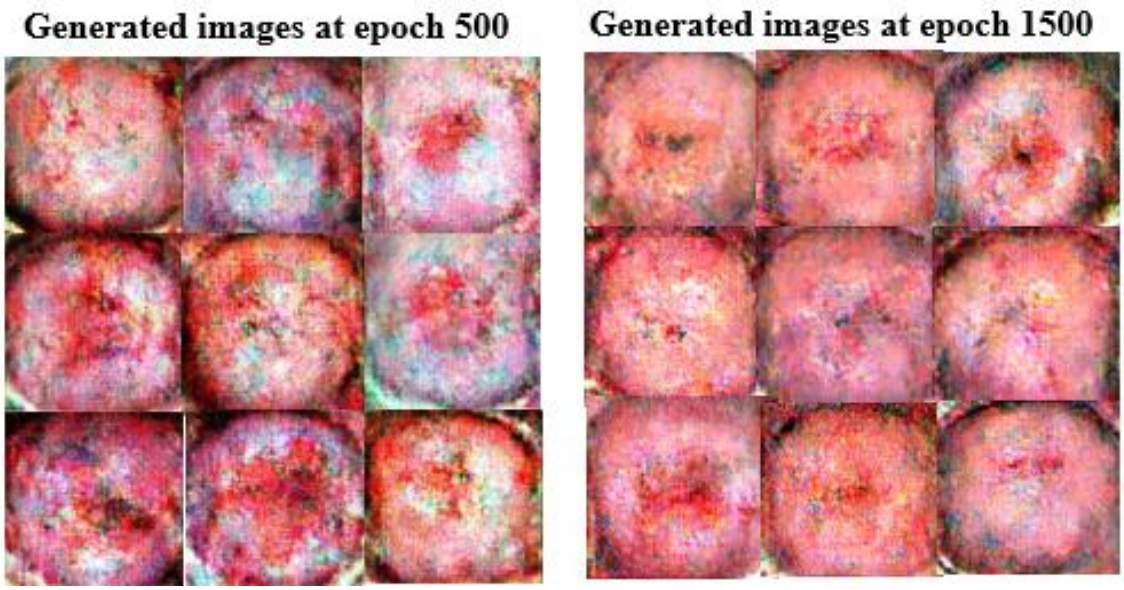

Figure 8 Result samples without data augmentation

a. Impact of Manual Data Augmentation

The improved performance observed in the augmented training scenario can be attributed to several interrelated factors that collectively enhanced the learning process. Manual data augmentation techniques effectively increased the dataset size from 233 to approximately 2097 images, providing the model with a more diverse set of examples and reducing the risk of overfitting to specific image characteristics. Moreover, transformations such as rotations and flips helped the model develop invariance to spatial orientation, which is particularly important in medical imaging, as the orientation of cervical cells should not affect classification. In addition, brightness and contrast adjustments further enhanced the model's robustness to the natural variations in staining and imaging conditions commonly encountered in cytological samples. As a result, the generator was able to learn more generalizable feature representations, leading to the synthesis of more realistic and visually consistent images.

b. Impact of Hyperparameters

Performance differences between the two models can also be attributed to differences in hyperparameter configurations. Table 1 summarizes the hyperparameter used in the DCGAN models trained with and without manual data augmentation. Several key observations can be drawn from this comparison.

Table 1 Training setup

Data Augmentation	Hyperparameters			
	Shape	Learning Rate	Batch Size	Latent Space Dimension
With	128	1e-4	32	256
Without	128	2e-4	16	128

First, the augmented model uses a larger latent space (256 compared to 128), allowing the generator to encode a richer and more diverse set of latent features. This increased representational capacity can enhance the quality and variability of generated images. Second, the batch size in the augmented model is twice as large (32 versus 16), which contributes to more stable gradient updates during training. A larger batch sizes can help stabilize training, leading to better generalization and a smoother convergence trajectory. Additionally, the learning rate for the augmented model is set to 1×10^{-4}, whereas the non-augmented model uses a higher learning rate of 2×10^{-4}. A lower learning rate can contribute to more stable training, preventing drastic updates that might destabilize the GAN.

c. FID Score Analysis

The Fréchet Inception Distance (FID) was calculated every 500 epochs to quantitatively assess the quality of generated images. Table 2 presents the FID Score obtained at different training stages for both models.

Table 2 FID score

Epoch	FID score	
	With Data Augmentation	Without Data Augmentation
500	2.52	4.01
1000	2.39	3.13
1500	2.21	2.52

As shown in Table 2, the DCGAN trained with manual data augmentation exhibited a consistent decrease in FID scores throughout training, starting from approximately 2.52 at epoch 500 and reaching 2.21 at epoch 1500. In contrast, the model trained without data augmentation began with a substantially higher FID score of 4.01 at epoch 500 and decreased to 2.52 by the final epoch. The results indicate that applying data augmentation improves the quality of generated images, as evidenced by consistently lower FID scores compared to training without augmentation. At epoch 500, the FID score with augmentation is 2.52, while the model without augmentation scores 4.01. As training progresses, both models improve, but the gap remains, with the augmented model achieving 2.21 at epoch 1500, whereas the non-augmented model scores 2.52. Furthermore, the stability of the training process was influenced by differences in hyperparameter settings, particularly batch size and learning rate. The learning rate and batch size differed between the two models may have contributed to the observed variations in performance.

d. Discussion

The combination of data augmentation and modified hyperparameters led to a more stable training process and improved image quality. The lower FID scores achieved by the augmented model indicate that the generated images are more closely resemble to real samples. While data augmentation improves diversity and generalization, it also changes the distribution of training data, requiring adjustments in the latent space dimensionality and optimization parameters. The findings demonstrate that augmenting the dataset and optimizing hyperparameters significantly enhance DCGAN's ability to generate high-quality synthetic images. Despite increasing the effective dataset size to approximately 2097 images, the dataset remains relatively small compared to those typically used for GAN training in other domains. Furthermore, *Cervical Intraepithelial Neoplasia* (CIN) images contain complex cellular structures that require high fidelity reproduction, making the generation task inherently more challenging compared to common GAN applications involving natural images or faces datasets.

CONCLUSION

The limitation of available datasets remains a major challenge in medical imaging, often leading to suboptimal classification performance due to constraints related to patient privacy, high acquisition costs, and the need for expert annotation. This study investigated the effectiveness of DCGAN for augmenting CIN1 images to address the challenge of limited medical imaging data. The experimental results showed that the DCGAN model trained with manual data augmentation achieved a lower FID score of 2.21, compared to 2.52 for the model trained without augmentation, indicating a measurable improvement in the similarity between synthetic and real images. In addition to improved image quality, the augmented training scenario exhibited more stable generator and discriminator loss curves after approximately 1000 epochs, demonstrating improved adversarial convergence. The findings further reveal that the differences in hyperparameters—such as the latent dimension (256 vs. 128), batch size (32 vs. 16), and learning rate (1×10^{-4} vs. 2×10^{-4}) - play a critical role in training stability and output quality. Collectively, these results highlight the importance of combining effective data

augmentation strategies with carefully tuned hyperparameters to enhance GAN performance in data-limited medical imaging settings.

Future work should expand on these findings by exploring more advanced GAN architectures, systematically analyzing the individual contributions of different augmentation techniques, and conducting downstream evaluations using classification models trained with the synthetic images. Such evaluations are essential to determine whether improvements in generative quality translate into tangible gains in clinical diagnostic performance.

ACKNOWLEDGEMENT

We would like to thank the Intelligent System Research Group (ISysRG), Faculty of Computer Science, Sriwijaya University, Indonesia, for providing the DL infrastructure used in this study. We also acknowledge the IARC Cervical Cancer Image Bank for supplying the dataset that supported this research.

REFERENCES

Alrashedy HHN, Almansour AF, Ibrahim DM, Hammoudeh MAA. 2022. BrainGAN: Brain MRI Image Generation and Classification Framework Using GAN Architectures and CNN Models. *Sensors*. 22 (11):4297. doi:10.3390/s22114297.

Antoniou A, StorkeyA, Edwards H. 2017. Data augmentation generative adversarial networks. arXiv preprint arXiv:1711.04340.

Ciano G, Andreini P, Mazzierli T, Bianchini M, Scarselli F. 2021. A Multi-Stage GAN for Multi-Organ Chest X-ray Image Generation and Segmentation. *Mathematics*. 9 (22):2896. doi:10.3390/math9222896.

Devi YS, Kumar SP. 2022. DR-DCGAN: A Deep Convolutional Generative Adversarial Network (DC-GAN) for Diabetic Retinopathy Image Synthesis. *Webology*. 19(2). https://api.semanticscholar.org/CorpusID:247216499.

Dowson DC, Landau B V. 2003. The Frkhet Distance between Multivariate Normal Distributions. *Journal of multivariate analysis*. 12(3): 450-455. https://api.semanticscholar.org/CorpusID:122512502.

Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. 2020. Generative Adversarial Networks. *Communications of the ACM*. 63 (11): 139-144.

Heusel M, Ramsauer H, Unterthiner T, Nessler B, Hochreiter S. 2017 Jun 26. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Advances in neural information processing systems*. 30.

Huang P, Liu X, Huang Y. 2021. Data Augmentation For Medical MR Image Using Generative Adversarial Networks. *ArXiv preprint*. abs/2111.14297. https://api.semanticscholar.org/CorpusID:244714924.

International Agency for Research on Cancer. 2021. Cervical Image Bank. *https://screening.iarc.fr/cervicalimagebank.php*.

Jiang M, Zhi M, Wei L, Yang X, Zhang J, Li Y, Wang P, Huang J, Yang G. 2021. FA-GAN: Fused attentive generative adversarial networks for MRI image super-resolution. *Comput Med Imaging Graph*. 92:101969. doi:10.1016/j.compmedimag.2021.101969.

Karras T, Aila T, Laine S, Lehtinen J. 2017 Okt 27. Progressive Growing of GANs for Improved Quality, Stability, and Variation. *arXiv preprint arXiv:1710.10196*.

Liu B, Lv J, Fan X, Luo J, Zou T. 2022. Application of an Improved DCGAN for Image Generation. *Mob Inf Syst*. 2022:1–14. doi:10.1155/2022/9005552.

Liu Y, Chen A, Shi H, Huang S, Zheng W, Liu Z, Zhang Q, Yang X. 2021. CT synthesis from MRI using multi-cycle GAN for head-and-neck radiation therapy. *Comput Med Imaging Graph*. 91:101953. doi:10.1016/j.compmedimag.2021.101953.

Meor Yahaya MS, Teo J. 2023. Data augmentation using generative adversarial networks for images and biomarkers in medicine and neuroscience. *Front Appl Math Stat*. 9. doi:10.3389/fams.2023.1162760.

Mirza M, Osindero S. 2014 Nov 6. Conditional Generative Adversarial Nets. *arXiv preprint arXiv:1411.1784*.

Motamed S, Rogalla P, Khalvati F. 2021. Data augmentation using Generative Adversarial Networks (GANs) for GAN-based detection of Pneumonia and COVID-19 in chest X-ray images. *Informatics Med Unlocked*. 27:100779. doi:10.1016/j.imu.2021.100779.

Musalamadugu TS, Kannan H. 2023 Sep 6. Generative AI for medical imaging analysis and applications. *Futur Med AI*. 1(2).

Pinaya WHL, Graham MS, Kerfoot E, Tudosiu P-D, Dafflon J, Fernandez V, Sanchez P, Wolleb J, da Costa PF, Patel A, *et al.* 2023. Generative AI for Medical Imaging: extending the MONAI Framework. *arXiv preprint arXiv:2307.15208*.

Radford A, Metz L, Chintala S. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv preprint arXiv:1511.06434*.

Smaida M, Yaroshchak S, Barg Y El. 2021. DCGAN for Enhancing Eye Diseases Classification. In: *International Workshop on Computer Modeling and Intelligent Systems*. https://api.semanticscholar.org/CorpusID:234753197.

Sundaram S, Hulkund N. 2021. GAN-based Data Augmentation for Chest X-ray Classification. *arXiv preprint arXiv:2107.02970*.

Wang L, Sun B, Xu J, Cao D, Chen Y, Xu Y, Wu D. 2024. Emerging trends and hotspots in cervical intraepithelial neoplasia research from 2013 to 2023: A bibliometric analysis. *Heliyon*. 10(11):e32114. doi:10.1016/j.heliyon.2024.e32114.

Wubineh BZ, Rusiecki A, Halawa K. 2024. Classification of cervical cells from the Pap smear image using the RES_DCGAN data augmentation and ResNet50V2 with self-attention architecture. *Neural Comput Appl*. 36(34):21801–21815. doi:10.1007/s00521-024-10404-x.

Yi Z, Zhang H, Tan P, Gong M. 2017. DualGAN: Unsupervised Dual Learning for Image-to-Image Translation. In *Proceedings of the IEEE international conference on computer vision* (pp. 2849-2857).

Zhang K, Hu H, Philbrick K, Conte GM, Sobek JD, Rouzrokh P, Erickson BJ. 2022. SOUP-GAN: Super-Resolution MRI Using Generative Adversarial Networks. *Tomography*. 8(2):905–919. doi:10.3390/tomography8020073.

Zhang Y, Wang Z, Zhang Z, Liu J, Feng Y, Wee L, Dekker A, Chen Q, Traverso A. 2023. GAN-based one dimensional medical data augmentation. *Soft Comput*. 27(15):10481–10491. doi:10.1007/s00500-023-08345-z.

Zhao B, Cheng T, Zhang X, Wang J, Zhu H, Zhao R, Li D, Zhang Z, Yu G. 2023. CT synthesis from MR in the pelvic area using Residual Transformer Conditional GAN. *Comput Med Imaging Graph*. 103:102150. doi:10.1016/j.compmedimag.2022.102150.

Zhao Z, Qin Y, Shao K, Liu Y, Zhang Y, Li H, Li W, Xu J, Zhang J, Ning B, *et al.* 2024. Radiomics Harmonization in Ultrasound Images for Cervical Cancer Lymph Node Metastasis Prediction Using Cycle-GAN. *Technol Cancer Res Treat*. 23. doi:10.1177/15330338241302237.

Zhu J-Y, Park T, Isola P, Efros AA. 2017 Mar 30. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 2223-2232).