

Application of GANs for Augmentation of the Mammography Repository

Yaneth Reyes-Hernández^{1,*}, Fernando Perez-Tellez², Jorge A. Ruiz-Vanoye¹,
Jazmín Rodríguez Flores¹, Eric Simancas Acevedo¹, Ocotlán Diaz-Parra¹,
Jaime Aguilar Ortiz¹, Francisco Rafael Trejo-Macotela¹

¹ Universidad Politécnica de Pachuca,
Mexico

² Technological University of Dublin,
Faculty of Computing, Digital and Data,
Ireland

yane.rehyes@gmail.com

Abstract. Generative adversarial networks (GANs) offer an innovative approach to synthetic image generation. They have significantly impacted the creation of images that would otherwise be difficult to obtain. In this study, we examine several GAN architectures to determine whether they can generate synthetic mammography images to enrich an existing repository, thereby improving AI training for breast-cancer detection and supporting research into this disease with a more diverse dataset.

Keywords. GAN, synthetic data, synthetic images, mammograms, generative AI.

1 Introduction

García (2023) highlights the importance of generative artificial intelligence (generative AI), as it can produce unique and original images that did not previously exist, enabling their application across different projects.

Amazon Web Services, Inc. (2024) states that generative AI can be used to create synthetic medical data, which is useful for training AI models, simulating clinical trials and studying rare diseases in the absence of large real-world datasets.

Díaz et al. (2021) discuss the paucity of large public databases and how this has constrained the application of artificial intelligence (AI) in the clinical field, despite the vast wealth of information provided by current image-archiving systems.

In the article by Goodfellow et al. (2014), generative adversarial networks (GANs) were introduced for the creation of synthetic data.

Image generation using generative artificial intelligence can have a profound impact in the medical field. Mammograms have been selected for this study, as the World Health Organization OPS/OMS, (2024) reports that breast cancer is the most common cancer and the leading cause of cancer-related death worldwide, primarily affecting women.

This project examines several GAN architectures that may prove useful for generating synthetic medical images from mammograms. We analyze the performance of these models and evaluate the resulting images to ascertain their suitability for enriching an existing repository.

1.1 Problem Statement

Raraz-Vidal (2023) emphasizes the importance of datasets for training and developing AI models, noting that as AI applications become more varied and sophisticated, the demand for data grows exponentially.

McNulty et al. (2024) note that, in the field of medical imaging, access to data is often limited by patient-privacy restrictions and the challenge of acquiring sufficient data for rare diseases.

Arceo-Martínez et al. (2021) note that breast cancer is the most commonly diagnosed cancer

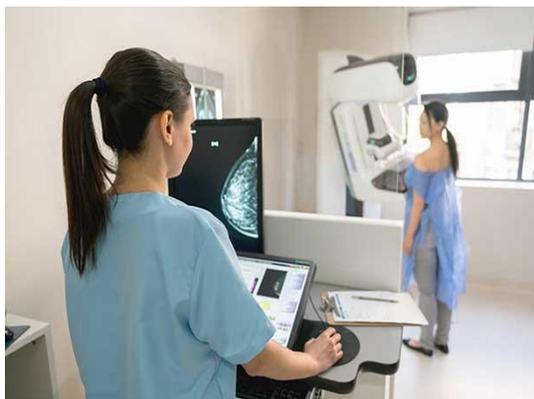


Fig. 1. Breast cancer screening exam (Instituto Nacional del Cáncer, 2012)

worldwide, with an incidence rate exceeding 45 cases per 100 000; women are the most affected population.

As shown in Fig. 1, breast-cancer screening is carried out through medical-imaging techniques.

According to Freire Hidalgo (2021), various radiological screening techniques are currently used for the early detection and diagnosis of breast cancer. Among the most commonly used is mammography, due to its low cost and minimal radiation exposure for the patient.

Although several repositories contain medical images of breast cancer, the available data remain insufficient, limiting their use in AI models designed to identify key patterns for disease detection. This limitation may result in biased and less robust models. For this reason, generation of mammographic images via generative image-based AI is considered a potential solution, as it could expand existing repositories and strengthen classification models, making them more reliable for breast-cancer diagnosis.

2 State of the Art

Within the scope of this research, several related projects have been identified that aim to generate synthetic images using generative artificial intelligence. These studies inform this project by guiding decisions on model selection and the interpretation of results.

In the study by Costa et al. (2018), an adversarial autoencoder is employed to generate retinal images. The visual and quantitative results demonstrate that the synthesized images differ substantially from those in the training set, while remaining anatomically consistent and exhibiting acceptable visual quality.

Akpinar et al. (2025) review how synthetic-data generation can be applied to train deep-learning models in the healthcare domain, thereby enabling the effective use of small datasets.

L. Wang et al. (2020) examine applications of GANs in medical contexts, highlighting their ability to synthesize high-quality images when data are scarce.

Wickramaratne & Mahmud (2021) similarly review GAN applications in medicine, emphasizing their capacity to synthesize high-quality images under data-scarce conditions.

The study by Zhang et al. (2023) explores the applications of GANs in medicine and emphasizes their importance for clinical medical research, particularly in the areas of privacy protection and medical diagnostics; however, it also highlights the need to consider ethical and legal aspects, as well as validation by expert radiologists.

In the research by Zhu et al. (2024), advanced models such as Transformers, Graph Neural Networks (GNNs) and Generative Adversarial Networks (GANs) are integrated to optimize sports training and enhance injury prevention. The generated motion sequences were more realistic and diverse, contributing to improved outcomes.

Ali & Shah (2022) reviewed GAN-based methods for combating COVID-19. Their findings indicate that GANs play a significant role in data augmentation by generating synthetic CT and chest X-ray images from limited existing datasets, thereby directly contributing to improved diagnostic performance.

The study by (Beers et al., 2018) highlights how the progressive-growing technique in GANs enables the generation of high-resolution medical images, which is especially valuable for tasks requiring a high level of detail.

In the study by Mardani et al. (2019), the application of GANs for magnetic resonance imaging (MRI) reconstruction is shown to be effective in enhancing both spatial resolution and

overall image quality, thereby contributing to more accurate diagnostics.

Saboo et al. (2021) explore the use of StyleGAN to generate realistic synthetic chest X-rays; they also present a generator–encoder system that enables organ-scale editing of radiographic images, marking a significant advancement towards disease-aware image generation and editing.

3 Generative Adversarial Networks (GANs)

GANs are generative models primarily designed to produce synthetic data that closely resemble real-world datasets. Fig. 2 illustrates how this generative model functions: it comprises two neural networks—one that generates images and another that discriminates between real and synthetic images. The discriminator's goal is to determine whether an image is genuine or fake, while the generator seeks to learn the distribution of real images so effectively that the discriminator cannot distinguish its outputs from actual samples.

The operation of GANs is characterized by a minimax game in which the generator and discriminator are trained concurrently. The discriminator is trained to maximize the probability of correctly labelling images, whereas the generator is trained to fool the discriminator:

$$\min_G \max_D V(D, G) = E_{x \sim P_{data}(x)} [\log D(x)] + E_{z \sim P_z(z)} [\log(1 - D(G(z)))].$$

3.1 Selected GAN Models

GANs have several variants designed to improve neural-network training. For this study, we selected DCGAN and StyleGAN.

3.1.1 DCGAN (Deep Convolutional GAN)

Radford et al. (2016) note that DCGAN can generate images of high resolution and quality.

- DCGAN is a variant of GAN that employs convolutional neural networks (CNNs).

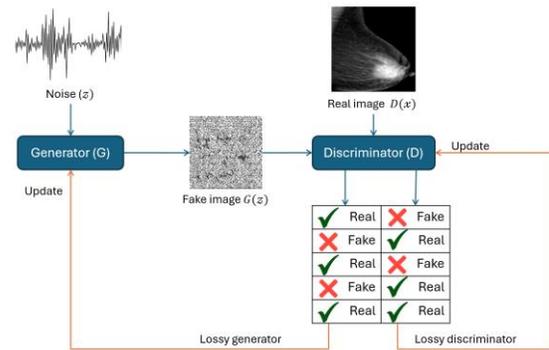


Fig. 2. GAN Training Process

- The discriminator comprises convolutional layers with strides, batch-normalization layers and LeakyReLU activations.
- The generator consists of transposed convolutional layers, batch-normalization layers and ReLU activations.

The loss function for the generator (G) is defined as:

$$L_G = E_{z \sim P_z(z)} [\log(1 - D(G(z)))].$$

The loss function for the discriminator (D) is defined as:

$$L_D = E_{x \sim P_{data}(x)} [\log D(x)] + E_{z \sim P_z(z)} [\log(1 - D(G(z)))].$$

3.1.2 StyleGAN

In the study by Karras et al. (2019), StyleGAN is described as a GAN architecture for generating high-resolution, realistic images and draws on the style-transfer literature:

- It generates high-quality images.
- Unlike conventional generators that receive the latent code as input in the first layer, the style-based architecture completely omits the input layer and begins with a learned constant input.

- It employs a mapping network, which, instead of feeding a random latent vector z directly into the generator, maps z into another vector w using a multilayer perceptron (MLP). This intermediate vector w controls various image features such as texture and lighting.
- The vector w is transformed through learned affine transformations into specific "styles" $y=(y_s, y_b)$. These styles control adaptive instance normalization (AdaIN) operations, which are applied after each convolutional layer in the synthesis network g . AdaIN normalizes each feature map individually, then scales and shifts it using the corresponding style components.
- As shown in Fig. 3, the architecture includes explicit noise inputs at each layer of the synthesis network. These are single-channel uncorrelated Gaussian noise images that are scaled by learned per-feature factors and added to the output of the corresponding convolution:

$$AdaIN(x_i, y) = y_{s,i} \frac{x_i - \mu(x_i)}{\sigma(x_i)} + y_{b,i}.$$

4 Experimentation

In the reviewed studies, several control factors were identified that affect the generation of synthetic images through generative AI models. These images have been assessed using various metrics, with the Fréchet Inception Distance (FID) as the most commonly used quantitative metric:

- **Learning Rate:** Controls how much the model's parameters are updated.
- **Batch Size:** Number of images used per training iteration.
- **Epochs:** Number of training cycles (iterations over the full dataset).

For this research project, we adopted the Taguchi method, which emphasizes selecting appropriate levels of control factors to minimize variability caused by noise factors and thereby

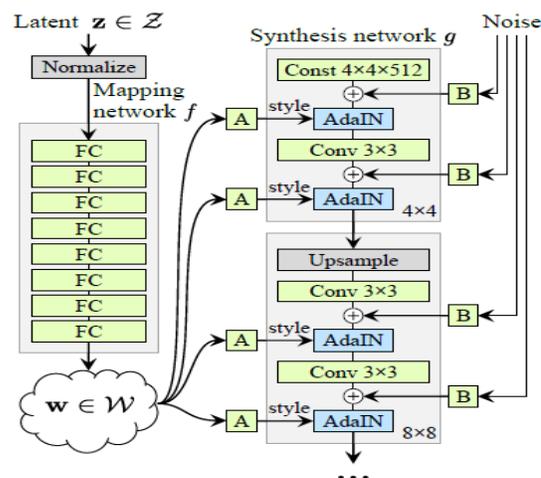


Fig. 3. Style-based generator (Karras et al., 2019)

Table 1. Factors and levels for experimentation.

Factor	Level 1	Level 2	Level 3
Learning Rate	0.0001	0.0002	0.0003
Batch size	16	32	64
Epoch	250	500	750

produce a robust process or output (Naranjo-Palacios et al., 2020).

This experimental design enables the exploration of different control-factor configurations to achieve the desired outcome, allowing the most effective settings to be identified and analyzed.

Table 1 presents the factors and levels used in this experiment, resulting in an L9 orthogonal array.

4.1. Model Implementation

Python was chosen for the implementation of the selected models, owing to its widespread popularity as a programming language and its open-source nature.

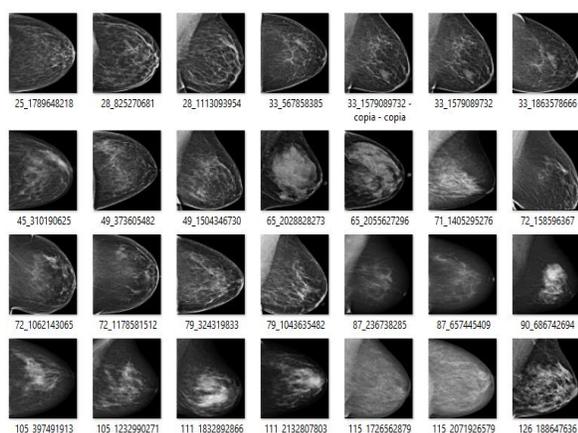
Given the computational demands of GAN algorithms and based on initial tests across various devices and platforms, a system equipped with an 8 GB GPU was selected as the primary hardware for model training. The specifications of the device are listed in Table, which were considered

Table 2. Laptop Specifications

Specification	Description
Device Type	Laptop
Operating System	Windows 11
Processor	AMD Ryzen 7
RAM	16GB
GPU	8GB
GPU Model	NVIDIA RTX 4060
Storage	500GB

Table 3. DCGAN Matrix

Exp	Learning Rate	Batch size	Epoch	FID
1	0.0001	16	250	0.448
2	0.0001	32	500	0.002
3	0.0001	64	750	0.161
4	0.0002	16	250	0.025
5	0.0002	32	500	0.006
6	0.0002	64	750	2.472
7	0.0003	16	250	0.012
8	0.0003	32	500	0.024
9	0.0003	64	750	0.02

**Fig. 4.** Original dataset (RSNA-ROI-Mammography, 2023)

sufficient for executing and training the generative models.

4.1 Selected Repository

The RSNA-ROI-Mammography dataset, obtained from Kaggle, was selected for model training. Fig. 4 shows an example from the original dataset, from which image-data segmentation was performed, resulting in a total of 1,821 left-sided images used in the experiments. Note that these images underwent preprocessing, being resized and standardized to a resolution of 512x512 pixels for compatibility with the generative models.

5 Results

Following the implementation of the algorithms in Python and completion of the required model tests, the Taguchi experiments were conducted, comprising nine runs per model, each with a distinct configuration.

5.1 DCGAN Results

The experimentation was conducted using the proposed L9 orthogonal array. Table 3 presents the configurations used, as well as the results obtained for each experiment.

As can be observed, Experiment 2 yielded the best result among all trials, achieving an FID of 0.0017, with a total execution time of 8,735.06 seconds. Fig. 5 shows the synthetic images generated in this experiment.

Fig. 6 visually illustrates the behavior of the GAN networks. Initially, the generator network exhibits high loss values, but as training progresses, particularly in the middle epoch, the generator's loss begins to decrease while the discriminator's loss starts showing some spikes. This indicates that the generator is improving. Although a perfect balance is not achieved, the results suggest that the images generated are visually similar to those in the original dataset.

5.2 StyleGAN Results

The experimentation was carried out using the proposed L9 orthogonal array. Table 4 presents

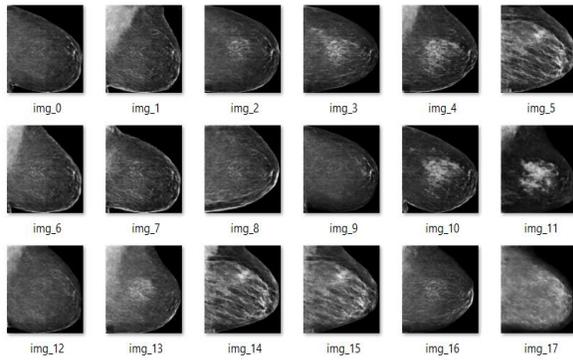


Fig. 5. DCGAN Synthetic Images

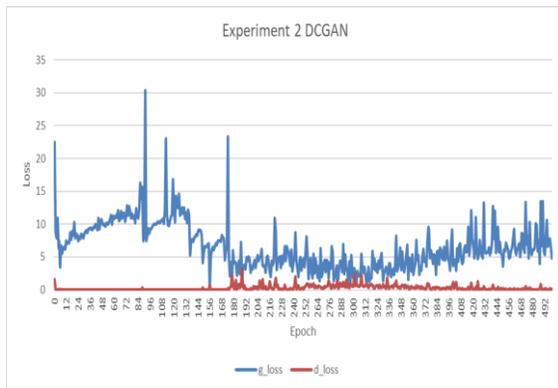


Fig. 6. DCGAN Loss Graph

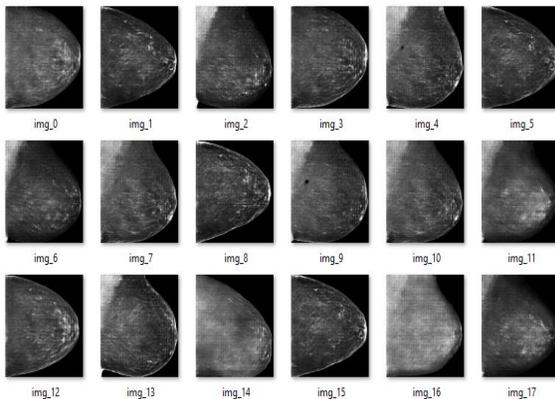


Fig. 7. StyleGAN Synthetic Images

the configurations used, along with the results obtained from each experiment.

As can be seen, Experiment 5 produced the best result, achieving an FID of 0.0031 with a total

Table 1. StyleGAN Matrix

Exp	Learning Rate	Batch size	Epoch	FID
1	0.0001	16	250	0.1224
2	0.0001	32	500	0.045
3	0.0001	64	750	0.0335
4	0.0002	16	250	0.1603
5	0.0002	32	500	0.0031
6	0.0002	64	750	0.6647
7	0.0003	16	250	0.0061
8	0.0003	32	500	0.4959
9	0.0003	64	750	0.0335

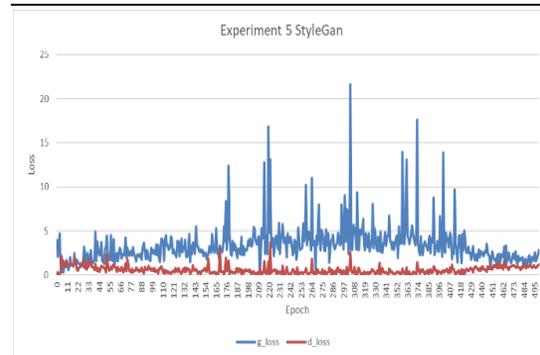


Fig. 8. StyleGAN Loss Graph

execution time of 4,662.08 seconds. Fig. 7 shows the images generated in this experiment.

Fig. 8 visually displays the behavior of the GAN networks. It can be observed that both the **discriminator** and **generator** maintain low loss values. Although the generator shows moderate spikes in some epochs, it recovers towards the end of the training, improving its performance. This results in higher-quality generated images and suggests that the networks are approaching convergence.

6 Conclusions

Following a search for various mammography image datasets, it was found that most available datasets were relatively small. Since image data are being used, it is important to emphasize the need for preprocessing—specifically, resizing

images to fixed dimensions and converting them to NumPy arrays. This ensures that the algorithm handles the data properly and helps to prevent potential errors during execution.

Based on the research and analysis of related projects, the most common generative-AI models were evaluated. This process helped identify the computational resources required and the quality of the results obtained. Consequently, GANs were selected for this study, as they are well suited to the available resources and the project's specific requirements.

The application of Taguchi experimentation enabled the structured implementation of various GAN algorithms under specific configurations. This approach identified the factors that delivered optimal performance in generating synthetic images. The most effective configuration comprised a batch size of 32, 500 training epochs and a learning rate of 0.0001–0.0002.

For the DCGAN algorithm, Experiment 2 yielded the best result, achieving an FID of 0.0017, and the generated images were visually similar to real images. However, for the StyleGAN model, although Experiment 5 produced a relatively favorable FID of 0.0031, the visual quality of the images did not meet expectations. Despite the FID score being close to optimal, it is important to note that the images have not yet been reviewed by a medical expert to confirm their clinical reliability.

References

1. Akpınar, M.H., Sengur, A., Salvi, M., Seoni, S., Faust, O., Mir, H., Molinari, F., Acharya, U.R. (2025). Synthetic Data Generation via Generative Adversarial Networks in Healthcare: A Systematic Review of Image and Signal Based Studies. *IEEE Open Journal of Engineering in Medicine and Biology*, Vol. 6, pp. 183–192. DOI: 10.1109/OJEMB.2024.3508472.
2. Ali, H., Shah, Z. (2022). Combating COVID-19 Using Generative Adversarial Networks and Artificial Intelligence for Medical Images: Scoping Review. *JMIR Medical Informatics*, Vol. 10, No. 6, pp. e37365. DOI: 10.2196/37365.
3. Amazon Web Services, Inc. (2024). ¿Qué es la IA generativa? Explicación de la inteligencia artificial generativa AWS. Amazon Web Services, Inc.
4. Arceo-Martínez, M.T., López-Meza, J.E., Ochoa-Zarzosa, A., Palomera-Sanchez, Z., Arceo-Martínez, M.T., López-Meza, J.E., Ochoa-Zarzosa, A., Palomera-Sanchez, Z. (2021). Estado actual del cáncer de mama en México: Principales tipos y factores de riesgo. *Gaceta mexicana de oncología*, Vol. 20, No. 3, pp. 101–110. DOI: 10.24875/j.gamo.21000134.
5. Beers, A., Brown, J., Chang, K., Campbell, J.P., Ostmo, S., Chiang, M.F., Kalpathy-Cramer, J. (2018). High-resolution Medical Image Synthesis Using Progressively Grown Generative Adversarial Networks. DOI: 10.48550/arXiv.1805.03144.
6. Costa, P., Galdran, A., Meyer, M.I., Niemeijer, M., Abràmoff, M., Mendonça, A.M., Campilho, A. (2018). End-to-End Adversarial Retinal Image Synthesis. *IEEE Transactions on Medical Imaging*, Vol. 37, No. 3, pp. 781–791. DOI: 10.1109/TMI.2017.2759102.
7. Díaz, O., Rodríguez-Ruiz, A., Gubern-Mérida, A., Martí, R., Chevalier, M. (2021). ¿Son los sistemas de inteligencia artificial una herramienta útil para los programas de cribado de cáncer de mama? *Radiología*, Vol. 63, No. 3, pp. 236–244. DOI: 10.1016/j.rx.2020.11.006.
8. Freire-Hidalgo, J.M. (2021). Machine learning y clustering para detección de cáncer de mama a través de imágenes de mamografía.
9. García, E. (2023). Generación de imágenes a través de Inteligencia Artificial. OBS Business School.
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y. (2014). Generative Adversarial Nets.
11. Instituto Nacional del Cáncer (2012). Aspectos generales de los exámenes de detección del cáncer.

12. **Karras, T., Laine, S., Aila, T. (2019).** A Style-Based Generator Architecture for Generative Adversarial Networks. DOI: 10.48550/arXiv.1812.04948.
13. **Mardani, M., Gong, E., Cheng, J.Y., Vasanawala, S.S., Zaharchuk, G., Xing, L., Pauly, J.M. (2019).** Deep Generative Adversarial Neural Networks for Compressive Sensing MRI. *IEEE Transactions on Medical Imaging*, Vol. 38, No. 1, pp. 167–179. DOI: 10.1109/TMI.2018.2858752.
14. **McNulty, J.R., Kho, L., Case, A.L., Fornaca, C., Johnston, D., Slater, D., Abzug, J.M., Russell, S.A. (2024).** Synthetic Medical Imaging Generation with Generative Adversarial Networks for Plain Radiographs. DOI: 10.48550/arXiv.2403.19107.
15. **Naranjo-Palacios, F., Rios-Lira, A.J., Pantoja-Pacheco, Y.V., Tapia-Esquivias, M., Naranjo-Palacios, F., Rios-Lira, A.J., Pantoja-Pacheco, Y.V., Tapia-Esquivias, M. (2020).** Diseños ortogonales de Taguchi fraccionados. *Ingeniería, investigación y tecnología*, Vol. 21, No. 2, pp. 1–12. DOI:10.22201/ft.25940732e.2020.21n2.011.
16. **Pan American Health Orgnaization / World Health Organization (2024).** Cáncer de mama. <https://www.paho.org/es/temas/cancer-mama>
17. **Radford, A., Metz, L., Chintala, S. (2016).** Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks DOI: 10.48550/arXiv.1511.06434.
18. **Raraz-Vidal, J. (2023).** La importancia de las bases de datos para el entrenamiento en inteligencia artificial. *Revista Peruana de Investigación en Salud*, Vol. 7, No. 3, pp. 121–122. DOI: 10.35839/repis.7.3.1970.
19. **RSNA ROI Mammography (2023).** <https://www.kaggle.com/datasets/quachnam/rsna-roi-mammography>.
20. **Saboo, A., Ramachandran, S.N., Dierkes, K., Keles, H.Y. (2021).** Towards Disease-Aware Image Editing of Chest X-rays. DOI:10.48550/arXiv.2109.01071.
21. **Wang, L., Chen, W., Yang, W., Bi, F., Yu, F.R. (2020).** A State-of-the-Art Review on Image Synthesis with Generative Adversarial Networks. *IEEE Access*, Vol. 8, pp. 63514–63537. DOI: 10.1109/ACCESS.2020.2982224.
22. **Wickramaratne, S.D., Mahmud, M.S. (2021).** Conditional-GAN Based Data Augmentation for Deep Learning Task Classifier Improvement Using fNIRS Data. *Frontiers in Big Data*, Vol. 4, pp. 1–12. DOI: 10.3389/fdata.2021.659146.
23. **Zhang, F., Wang, L., Zhao, J., Zhang, X. (2023).** Medical Applications of Generative Adversarial Network: A Visualization Analysis. *Acta Radiologica*, Vol. 64, No. 10, pp. 2757–2767. DOI: 10.1177/02841851231189035.
24. **Zhu, J., Ye, Z., Ren, M., Ma, G. (2024).** Transformative Skeletal Motion Analysis: Optimization of Exercise Training and Injury Prevention through Graph Neural Networks. *Frontiers in Neuroscience*, Vol. 18. DOI: 10.3389/fnins.2024.1353257.

*Article received on 12/02/2025; accepted on 27/04/2025.
Corresponding author is Yaneth Reyes-Hernández.