Evolution of conditional-DCGANs for the synthesis of chest X-ray images.

Juan-Antonio Rodríguez-de-la-Cruz[a], Héctor-Gabriel Acosta-Mesa[a], Efrén Mezura-Montes[a], Fernando Arámbula Cosío[b], Borís Escalante-Ramírez[c], and Jimena Olveres Montiel[c]

[a]Instituto de Investigaciones en Inteligencia Artificial, Universidad de Veracruz, Campus Sur, Paseo 112, Col. Nueva Xalapa, C.P. 91097 Xalapa, Veracruz, México
[b]Instituto de Investigaciones en Matemáticas Aplicadas y en Sistemas, Universidad Nacional Autónoma de México, Circuito Escolar S/N, Ciudad Universitaria, Cd. Mx., C.P. 04510, México
[c]Facultad de Ingeniería, Universidad Nacional Autónoma de México, Universidad 3000, Ciudad Universitaria, Coyoacán, Cd. Mx., C.P. 04510, México

ABSTRACT

Deep learning (DL) is now widely used to perform tasks involving the analysis of biomedical imaging. However, the small amounts available of annotated examples of these types of images make it difficult to use DL-based systems, since large amounts of data are required for adequate generalization and performance. For this reason, in recent years, Generative Adversarial Networks (GANs) have been used to obtain synthetic images that artificially increase the amount available. Despite this, the usual training instability in GANs, in addition to their empirical design, does not always allow for high-quality results. Through the neuroevolution of GANs it has been possible to reduce these problems, but many of these works use benchmark datasets with thousands of images, a scenario that does not reflect the real conditions of cases in which it is necessary to increase the data due to the limited amount available. In this work, cDCGAN-PSO is presented, an algorithm for the neuroevolution of GANs that adapts the concepts of the DCGAN-PSO to a conditional-DCGAN that allows the synthesis of three classes of chest X-ray images and that is trained with only 600 images of each class. The synthetic images obtained from evolved GANs show good similarity with real chest X-ray images.

Keywords: GANs Neuroevolution, Biomedical Imaging Synthesis, Chest X-Ray Images, conditional-DCGAN, DCGAN-PSO

1. INTRODUCTION

Systems based on Deep Learning (DL) used in multiple tasks that involve the use of biomedical images have made it possible to obtain the best performances in the *state of the art*.[1] Activities such as classification, segmentation of areas of interest, noise elimination and registration, to name a few, have benefited from the use of these complex systems, which in turn support the biomedical area. However, the use of DL-based models implies the need for high amounts of data for the purpose of adequate training and generalization, since the very complexity of these is what can cause problems, e.g. overfitting, if there is a very limited amount of data. This problem is present in biomedical images, where their limited quantity is due to factors such as data privacy, the risk of some tests, e.g. exposure to radiation, or even the prohibitive cost of some tests.[2]

Further author information: (Send correspondence to J-A.R-C.)
J-A.R-C.: E-mail: juanantonio_2604@hotmail.es
H-G.A-M.: E-mail: heacosta@uv.mx
E.M-M.: E-mail: emezura@uv.mx
F.A.C.: E-mail: fernando.arambula@iimas.unam.mx
B.E-R.: E-mail: boris@unam.mx
J.O.M: E-mail: jolveres@cecav.unam.mx

The use in recent years of Generative Adversarial Networks (GANs), for the synthesis of images with a high similarity to the real ones, has allowed to increase the amounts available to train biomedical systems.[3] However, GANs usually have a training instability that ends up affecting the quality of the synthetic images obtained. Therefore, multiple neuroevolution systems, an area of Evolutionary Computation (EC) focused on improving and automating the design and/or training mechanisms of Neural Networks (NNs), focused on GANs have been developed, which has allowed progress to overcome or reduce classic GANs problems.[4] Most of these algorithms have been tested using popular benchmark datasets with thousands of images, none of them being tested in real applications or scenarios with little amount of data available, being DCGAN-PSO, to the best of the authors' knowledge, the first algorithm to be used not only with a low amount of data but also in the area of biomedical images, a field that can be highly benefited with the increase of data through GANs. However, by using DCGAN-PSO, only GANs that can synthesize a single class of images are obtained. In this paper, that proposal is extended and adapted to use conditional-GANs, a variant that allows the synthesis of multiple classes of images conditioned by the user, which allows fewer algorithm executions to be carried out to obtain different classes of synthetic images. Like the original version of the algorithm, Chest X-Ray (CXR) images were used as a case study, handling three different classes: pneumonia due to COVID-19, non-COVID-19 pneumonia and healthy.

To validate the results obtained using this new version, the quality was evaluated using FID and compared against the original version in the three classes of CXR images.

The rest of the paper is organized as follows: Section 2 contains the introduction to the key concepts related to GANs and their neuroevolution and CXR images, as well as related works. Section 3 contains the presentation and novelty of our proposal. Section 4 defines the experimentation performed to evaluate our version, as well as its results, while Section 5 contains the corresponding discussion. Section 6 presents the conclusions obtained and the proposals for future work.

2. BACKGROUND AND RELATED WORK

GANs[5] are Deep Learning models that belong to the set of generative models, a branch of unsupervised learning algorithms in charge of mapping how the data was generated. The training of a GAN is described as a zero-sum game (also called minimax) between two players with opposite objectives; these are two Neural Networks; the Generator and the Discriminator. Taking a vector of random noise sampled from prior distributions, e.g. normal or uniform, $(z \sim p_z(z))$ as input, called *latent vector*, the generator outputs samples from a more complex distribution $(G(z))$ whose goal is to be equal to the distribution of the real dataset. Meanwhile the discriminator has the task of distinguishing between the real samples$(x \sim p_{data}(x))$ and the generated samples $(G(z) \sim p_g(G(z)))$. In the case of real data the goal of the discriminator output $(D(\cdot))$ is to be near to 1. In the fake data scenario, the discriminator output's goal is to be close to 0 meanwhile the generator will try to make near to 1 i.e. fool the discriminator to classify his creations as real.

The training cost function of GANs, also called Minimax loss, is reflected by the following equation:

$$\min_G \max_D \mathop{\mathbb{E}}_{x \sim p_{data}} [log(D(x))] + \mathop{\mathbb{E}}_{z \sim p_z} [log(1—D(G(z)))] \tag{1}$$

The networks are trained simultaneously, encouraging both models towards continuous improvement and adaptation. For every iteration a gradient step with backpropagation is made to reduce the cost function of each network, optimizing their internal weights.

Conditional-GAN (cGAN)[6] is a variant of the classic GAN, which uses class labels to condition the class of the generated image from among all the classes available in the training set and thus obtain greater control over the synthesized images, contrary to the original GAN, where the synthesis of different classes from the same training set is completely random. The cost function of the cGAN is the same as the original GAN with the addition of the class labels (y):

$$\min_G \max_D \mathop{\mathbb{E}}_{x \sim p_{data}} [log(D(x|y))] + \mathop{\mathbb{E}}_{z \sim p_z} [log(1—D(G(z|y)))] \tag{2}$$

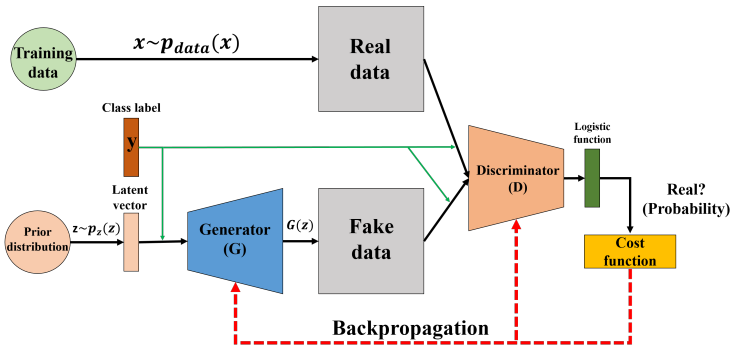

Figure 1. **General structure of a cGAN.** The green arrow represents the concatenation of the y label to the GAN's inputs.

In Fig. 1 the representation of the general structure of cGAN as well as its training process is shown.

Multiple GANs have been developed focused on the synthesis of various biomedical images, with Chest X-Ray (CXR) images being highlighted in the last year due to the need to increase the available quantities that represent COVID-19 pneumonia[7–9] to train systems that support in the diagnosis of this disease. However, other works have mainly used CXR images of the pneumonia (bacterial and viral) and healthy classes.[10–12] Nonetheless, all these works use empirically hand-designed networks, which are usually not generalizable to other applications or types of images. In addition, GANs commonly present training instability.

GANs training is complicated because there must be a balance between the skills of the generator and the discriminator. If there is a supremacy from one of the networks, training instability problems may cause the following problems: **(1)**Mode collapse: The situation in which the generator can only synthesize a small subset of data of the complete distribution since the training did not allow to generalize the richness of variants of the original distribution; **(2)** Vanishing gradient: Originated when the discriminator or the generator becomes powerful enough to cause an irreversible imbalance in training that does not make possible to the opposite network to improve its performances, thus causing a stalemate, resulting in poor visual quality synthetic results.

The use of neuroevolution to carry out the design and training of GANs has been addressed in multiple works recently.[4] Evolutionary Computation takes inspiration on the mechanism found in nature to evolve a population of potential solutions on the production of better outcomes for a given problem, being Neuroevolution the branch of EC in charge of the evolution of neural networks.[13] Among the works in GANs neuroevolution, DCGAN-PSO[14] stands out, which, unlike the previous works, uses a set of biomedical images as a real case of application. This algorithm uses a low number of images (compared to benchmark datasets with thousands of images used previously) for the evolution of DCGANs,[15] a variant of GAN focused on image synthesis. DCGAN-PSO can be extended and improved through the use of conditional-DCGANs to handle the creation of multiple classes of images through the use of a single evolved DCGAN, without the need to perform as many runs as classes are needed, as is necessary in DCGAN-PSO.

3. OUR APPROACH

Our approach to using conditional-DCGANs as an architecture to evolve in a neuroevolution algorithm is presented in this section. For this, we use the previously introduced method called DCGAN-PSO, adapting it to our proposal.

3.1 DCGAN-PSO

DCGAN-PSO[14] presents the use of a neuroevolution algorithm for the search for architectures and training of DCGANs, a variant of the GANs in which both networks that compose it are Convolutional Neural Networks (CNN), which allows obtaining networks capable of generating a single class of CXR images with a high similarity in quality and diversity with respect to the real set, an indicator of a more stable training.

DCGAN-PSO is based on Pro-GAN,[16] an approach carried out in the progressive growth of GANs, increasing the resolution of the images generated through the addition of layers, pre-established by the researchers, at specific moments of the GANs training. In addition, the variation operators of the CNN neuroevolution algorithm called psoCNN[17] are used.

The representation of the evolved DCGANs is by means of lists, where each slot represents a layer and the sequence of these the network architecture. Each of these architectures represents *particles* of a *swarm*, which when interacting can generate complex search behaviors.

For the progressive growth of the networks, a *swarm* of *particles* is initialized (random architectures) where each particle represents a DCGAN that can synthesize images at 4x4 pixels resolution. Each particle is then modified having the influence of the best architecture found so far by that *particle* (*pBest*) and the best obtained by the entire *swarm* (*gBest*). Through these better architectures, variation operators update the particles in order to find better networks. Each potential network is trained and evaluated by means of the fitness function, if the network results to improve the performance of the particle's *pBest*, its architecture and its weights are saved. At the end of a number of cycles (*generations*) determined for a resolution, the best architectures of each particle are fixed and new layers are added from these that allow doubling the resolution of the images generated. The new added layers are evolved by the same number of cycles as the previous resolution. These steps are repeated, doubling the resolution of the images until a final resolution of 256x256 pixels is obtained. Fig. 2 shows the general scheme of progressive growth carried out by this algorithm.

The fitness function used is the Frechet Inception Distance (FID).[18] FID uses the pre-trained Inception-v3 CNN[19] for the feature extraction of the real (x) and synthetic (g) images, thus obtaining a feature-vector of 2048 numerical values. This feature space is interpreted as a continuous multivariate Gaussian distribution. Therefore, from the features obtained, the Fréchet distance between both distributions is calculated using their estimated mean (μ) and covariance (Σ) by means of the following formula:

$$FID(x, g) = \|\mu_x — \mu_g\|_2^2 + Tr(\Sigma_x + \Sigma_g — 2(\Sigma_x \Sigma_g)^{\frac{1}{2}}) \tag{3}$$

The lower this metric is, the more similar the two sets of images are, being zero when they are equal.

The results obtained by this algorithm show an improvement in the quality of the synthetic images compared to handcrafted DCGANs for the CXR images of pneumonia due to COVID-19. In addition to verifying that

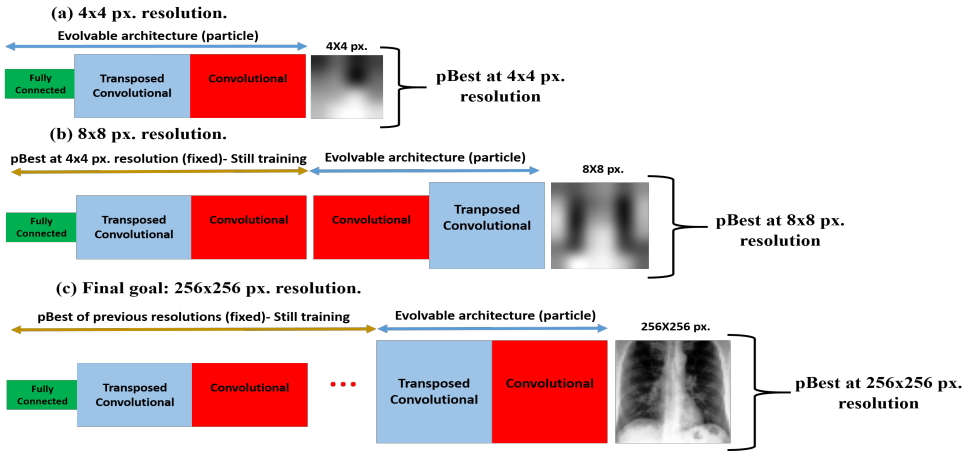

Figure 2. Progressive growth of a DCGAN represented in a particle. Here the growth of the generator is represented but the equivalent procedure is used in the discriminator. (a): At the beginning particles generate CXR images in 4^2 px. and *pBests* are selected. (b): When the resolution increases, the architecture and the trained weights of the particle's *pBest* in the previous resolution are used. *pBest*'s layers remains fixed (not evolved), but still training, while a new population is generated and evolved with particles that will be concatenated to the already known *pBest* part. (c): The process stops when 256^2 px. resolution is reached. Obtained from the original paper.[14]

the intelligent search carried out by the algorithm allows to improve the trained networks as the cycles advance, until having an early convergence. In the original paper, a brief experimentation was carried out improving the performance of a CNN for the binary classification of CXR images using the synthetic images of the evolved DCGANs, thus testing the usefulness of these synthetic biomedical images obtained through the neuroevolution of GANs.

3.2 cDCGAN-PSO

Our proposal focuses on the use of conditional-DCGANs instead of DCGANs for the neuroevolutionary algorithm, which allow generating different classes of images, controlled by the user, with only one evolved architecture. This new variant would allow the evolution of GANs architectures that generate multiple classes of biomedical images at the same time, thus reducing the number of necessary executions compared to obtain multiple classes with the original version, which only handles one class at a time per evolved network.

The changes necessary to adapt DCGAN-PSO to our proposal are the following:

- The architecture decoded from the particles is a cDCGAN instead of a DCGAN. Since the difference between architectures only concerns concatenating the class labels in the inputs of the generator and the discriminator to condition their creation and criticism, respectively, the representation of the architectures does not need any addition, since the change is made only in the implementation of the networks for their training. Therefore, those architectures that are generated from the evolved particles are also usable for the cDCGANs.
- The fitness function adopted is FID (as in DCGAN-PSO), but due to the handling of multiple classes it is required to measure the FID for each class. Therefore, we resort to using the FID averaged for each of the classes that synthesize the evolved GANs.

4. EXPERIMENTAL RESULTS

4.1 Experimental Setup

The parameters used in all experiments for both versions of the neuroevolution algorithm are the same used in the reference of the original version. These are detailed in Tab. 1.

Table 1. Experimental parameters of DCGAN-PSO and cDCGAN-PSO used.

Parameter	Value	Parameter	Value
Particle Swarm Optimization			
Swarm size	15	N° generations per resolution	10
C_g	0.5	Resolutions list	[4^2 px., 8^2 px., 16^2 px., 32^2 px., 64^2 px., 128^2 px., 256^2 px.]
Ranges of particle parameters			
N° Convolutional layers	[0,2]	N° Fully connected layers	[1,1] (4^2 px.) ; [0,0] (otherwise)
Filter size	[2,5]	N° Transposed convolutional layers	[0,0] (4^2 px.) ; [1,1] (otherwise)
N° neurons	[1,300]	N° filters	4^2:64^2 px. = [1,256]; 128^2 px. = [1,64]; 256^2 px. = [1,32]
DCGAN & cDCGAN training			
CXR images classes	COVID-19, Pneumonia, and Healthy	N° epochs per particle training	200
N° images used by class	600	Batch size	4^2:128^2 px. = 16; 256^2 px. = 14
Optimizer	Adam[20]	Learning rate (**G** and **D**)	2×10^{-4}
β_1,β_2 (optimizer)	0.5 , 0.999	LeakyReLU slope (**D**'s activations)	0.2
Weight's initializer (**G** and **D**)	$\mathcal{N}(0, 0.02)$	Noise distribution (p_z): $\mathbb{R}^{100}$	$\mathcal{N}(0, 1)$

4.2 CXR Images Datasets

The sets of CXR images of pneumonia and COVID-19 used for the cDCGAN-PSO experimentation were collected from the set available in Ref. 21, which has images of pneumonia (bacterial and viral) in addition to healthy cases. 918 images of each class were selected for their quality and processed in the same way as in the former version of the algorithm: **(1)**: grayscale conversion; **(2)**: image resizing to the different used resolutions (one dataset per resolution); **(3)**: histogram equalization to increase contrast. The COVID-19 image set is the same as the one used in the original version article. Only 600 images of each class were used in the GANs evolution.

4.3 Technical Details

Details about the optimizer and initialization of GANs used can be found in Tab. 1. Python 3.6.9 programming language with the PyTorch framework[22] were used for the implementation on the free access online platform Google Colaboratory* to have access to GPUs for the NNs training.

4.4 Fitness Evolution

The execution of cDCGAN-PSO, using the CXR images belonging to the COVID-19, pneumonia and healthy classes, was performed six times, this low number is due to computational limitations. Each run lasted approximately one week and half and with each change of resolution the time was doubled with respect to the previous resolution.

The fitness values of the *gBest* solution were monitored during different generations and evolution stages of the cDCGAN-PSO runs. The FID optimization is shown in Fig. 3.

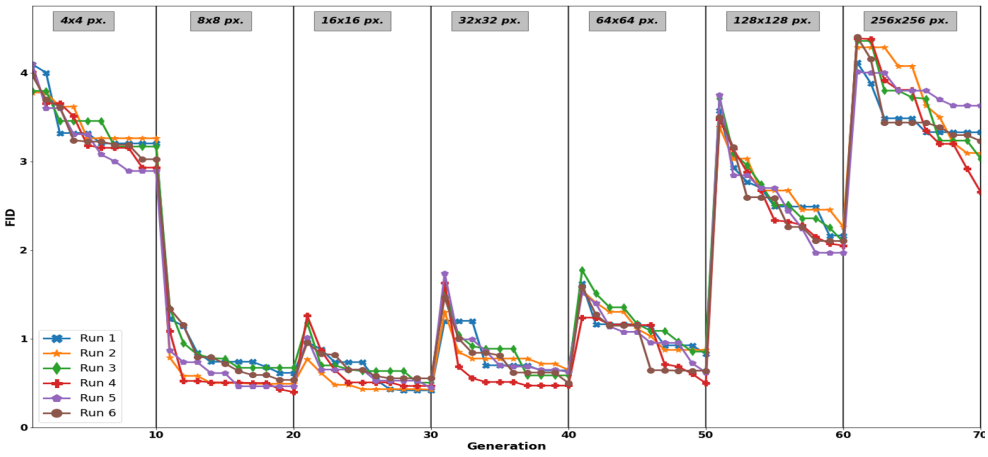

Figure 3. **cDCGAN-PSO *gBest* FID evolution.** COVID-19, healthy, and pneumonia classes (6 runs).

4.5 Qualitive Evaluation

Samples of the real images belonging to each class used as well as synthetics can be seen in Fig. 4.

4.6 FID Evaluation Comparison

The values of the FID evaluation obtained for the COVID-19 class of the original version were compared with the FID evaluation of the evolved networks of our proposal. In addition, two executions of the first version were performed per class pneumonia and healthy, and their FID was evaluated with the complete real set (918 for each class) and an equal size batch of synthetic images, the same was done for the networks obtained by our version, in order to compare the performance of each version for each of the CXR classes used in the present work. The low number of executions of the first version is due to computational restrictions. The results of these evaluations are shown in Tab. 2. The 95%-confidence Wilcoxon rank-sum (WRS) test was used to validate the statistical significance of the results.

*https://colab.research.google.com/

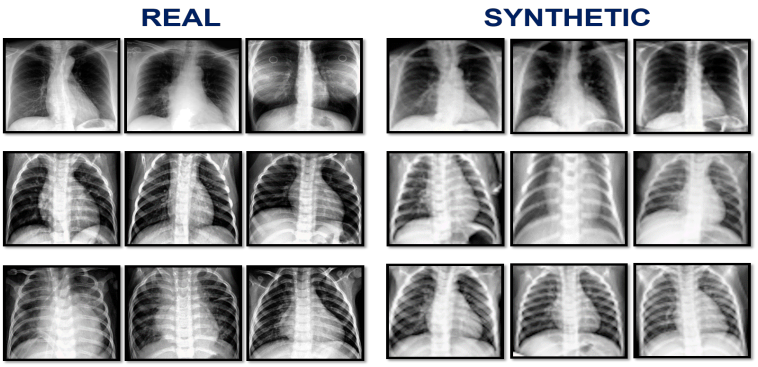

Figure 4. Sample of synthesized CXR images from cDCGAN-PSO in 256x256 pixels. Rows from top to bottom: COVID-19, pneumonia, and healthy classes.

Table 2. Results and comparison of FID evaluation values of both version of the algorithm. (=) means that the two sets of data compared have the same performance.

Class	Average FID (DCGAN-PSO)	Average FID (cDCGAN-PSO)	p-value	WRS test
COVID-19	3.052 ± 0.773	2.988 ± 0.631	0.8282	(=)
Pneumonia	3.2506 ± 0.666	3.111 ± 0.709	0.7388	(=)
Healthy	3.5916 ± 0.3242	3.345 ± 0.991	1	(=)

5. DISCUSSION

From Fig. 3 it can be observed that the fast convergence shown by the original version of the algorithm (for more details go to the article of the original version[14]) is also obtained by this new version. The behavior with sudden increases or decreases in each resolution change is as expected since the FID is a metric highly dependent on the level of detail shown by the different resolutions, then it has different ranges of values in each resolution. However, the descending behavior after each resolution change shows that the evolved networks are achieving better performances thanks to the intelligent search carried out by the neuroevolution algorithm.

Regarding the visual qualities of the synthetic images, shown in Fig.‘4, these managed to mimic the general morphological features of the real images, e.g. shape of the lungs and ribs, with only small errors (blank areas). A diversity such as that seen in the original set can also be observed. Giving indications of not existing mode collapse in the cDCGANs evolved what can be derived from more stable training. The visual quality of the results is slightly better than those obtained by DCGAN-PSO, noting an increase in the quality of the fine details of the CXR images, e.g. the more defined ribs. The reason for this is because the size of the training set was tripled by adding two new classes of CXR images. This allowed the evolved cDCGANs to use this extra amount of data to learn more finely the morphological details common to the three types of images. Hence, a visual improvement over the use of a single class was obtained when compared with the former version of the algorithm.

As can be seen in Tab. 2 , the average results of the final quality of FID obtained by each class is slightly better in the cDCGAN-PSO. The reason for this may be because, as discussed previously, the greater amount of training data allowed for better results than simply using a single class of CXR images as performed in the DCGAN-PSO.

Despite the previous observation, the results obtained by the Wilcoxon rank-sum test indicate that the performances of both algorithms are equal (i.e., there are not significant statistical differences). In this case, the advantage is obtained by the new version, the cDCGAN-PSO, because it can handle multiple classes in a single execution, then decreasing the total processing times by avoiding the sequential need the former version requires. Finally, considering the low number of executions due to technical restrictions, further experiments with more single runs need to be carried out so as to analyze the robustness of the new proposed approach.

6. CONCLUSIONS AND FUTURE WORK

In this work we present cDCGAN-PSO, a new version of DCGAN-PSO, a neuroevolutionary algorithm of GANs for the synthesis of biomedical images. Our version uses conditional-DCGANs which have the ability to synthesize multiple classes from CXR images, unlike the original version which can only handle one class at a time. The results obtained show that our version obtains similar performances compared to the original version of the algorithm, however, being able to synthesize multiple classes at the same time without a large increase in search times, i.e. about a week for the first version and a week and a half for our version, provides an attractive quality for use in scenarios with the need to balance or increase multiple classes with synthetic images for use in DL-based biomedical applications.

The experimentation carried out also verified the good quality of the synthetic images verified by their low FID values, statistically equal to those of the original version of the algorithm.

As future work, the proposed version can be used for the synthesis of images in higher resolution that allow to improve various tasks with CXR images of the biomedical area and that can be analyzed in detail by radiologist experts as a means of further corroborating their fidelity to real images.

ACKNOWLEDGMENTS

The first author, with CVU number 1001447, acknowledges support from the Mexican National Council of Science and Technology (CONACyT) through a scholarship to pursue graduate studies at the University of Veracruz. All UNAM authors gratefuly acknowledge the financial support of DGAPA, UNAM, under grants PAPIIT TA101121 and IV100420.

REFERENCES

[1] Mohapatra, S., Swarnkar, T., and Das, J., "Deep convolutional neural network in medical image processing," in [*Handbook of Deep Learning in Biomedical Engineering*], 25–60, Elsevier (2021).

[2] Guibas, J. T., Virdi, T. S., and Li, P. S., "Synthetic Medical Images from Dual Generative Adversarial Networks," *arXiv e-prints* (Sept. 2017).

[3] Kazeminia, S., Baur, C., Kuijper, A., van Ginneken, B., Navab, N., Albarqouni, S., and Mukhopadhyay, A., "Gans for medical image analysis," *Artificial Intelligence in Medicine* **109**, 101938 (2018).

[4] Costa, V., Lourenço, N., Correia, J., and Machado, P., "Neuroevolution of generative adversarial networks," in [*Deep Neural Evolution*], 293–322, Springer (2020).

[5] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., "Generative Adversarial Networks," *arXiv e-prints* (2014).

[6] Mirza, M. and Osindero, S., "Conditional Generative Adversarial Nets," *arXiv e-prints* (Nov. 2014).

[7] Loey, M., Smarandache, F., and M Khalifa, N. E., "Within the lack of chest covid-19 x-ray dataset: A novel detection model based on gan and deep transfer learning," *Symmetry* **12**(4), 651 (2020).

[8] Shams, M., Elzeki, O., Abd Elfattah, M., Medhat, T., and Hassanien, A. E., "Why are generative adversarial networks vital for deep neural networks? a case study on covid-19 chest x-ray images," in [*Big Data Analytics and Artificial Intelligence Against COVID-19: Innovation Vision and Approach*], 147–162, Springer (2020).

[9] Waheed, A., Goyal, M., Gupta, D., Khanna, A., Al-Turjman, F., and Pinheiro, P. R., "Covidgan: Data augmentation using auxiliary classifier gan for improved covid-19 detection," *IEEE Access* **8**, 91916–91923 (2020).

[10] Salehinejad, H., Valaee, S., Dowdell, T., Colak, E., and Barfett, J., "Generalization of deep neural networks for chest pathology classification in x-rays using generative adversarial networks," in [*2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*], 990–994, IEEE (2018).

[11] Khalifa, N. E. M., Taha, M. H. N., Hassanien, A. E., and Elghamrawy, S., "Detection of Coronavirus (COVID-19) Associated Pneumonia based on Generative Adversarial Networks and a Fine-Tuned Deep Transfer Learning Model using Chest X-ray Dataset," *arXiv e-prints* (Apr. 2020).

[12] Kora Venu, S. and Ravula, S., "Evaluation of deep convolutional generative adversarial networks for data augmentation of chest x-ray images," *Future Internet* **13**(1), 8 (2021).

[13] Miikkulainen, R., Liang, J., Meyerson, E., Rawal, A., Fink, D., Francon, O., Raju, B., Shahrzad, H., Navruzyan, A., Duffy, N., et al., "Evolving deep neural networks," in [*Artificial Intelligence in the Age of Neural Networks and Brain Computing*], 293–312, Elsevier (2019).

[14] Rodríguez-de-la Cruz, J.-A., Acosta-Mesa, H.-G., and Mezura-Montes, E., "Evolution of generative adversarial networks using pso for synthesis of covid-19 chest x-ray images," in [*2021 IEEE Congress on Evolutionary Computation*], 1–8, IEEE Press (2021).

[15] Radford, A., Metz, L., and Chintala, S., "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," *arXiv e-prints* (Nov. 2015).

[16] Karras, T., Aila, T., Laine, S., and Lehtinen, J., "Progressive Growing of GANs for Improved Quality, Stability, and Variation," *arXiv e-prints* (Oct. 2017).

[17] Junior, F. E. F. and Yen, G. G., "Particle swarm optimization of deep neural networks architectures for image classification," *Swarm and Evolutionary Computation* **49**, 62–74 (2019).

[18] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S., "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in [*Advances in neural information processing systems*], 6626–6637 (2017).

[19] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z., "Rethinking the inception architecture for computer vision," in [*Proceedings of the IEEE conference on computer vision and pattern recognition*], 2818–2826 (2016).

[20] Kingma, D. P. and Ba, J., "Adam: A Method for Stochastic Optimization," *arXiv e-prints* (Dec. 2014).

[21] Mooney, P., "Chest x-ray images (pneumonia)," *kaggle, Marzo* (2018).

[22] Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A., "Automatic differentiation in pytorch," (2017).