IFMBE Proceedings 106

Fernando Emilio Ballina
Ricardo Armentano
Rubén Carlos Acevedo
Gustavo Javier Meschino *Editors*

Advances in Bioengineering and Clinical Engineering

Proceedings of the XXIV Argentinian Congress of Bioengineering (SABI 2023), October 3–6, 2023, Buenos Aires, Argentina - Volume 1

Springer

IFMBE Proceedings 106

The IFMBE Proceedings Book Series is an official publication of *the International Federation for Medical and Biological Engineering* (IFMBE). The series gathers the proceedings of various international conferences, which are either organized or endorsed by the Federation. Books published in this series report on cutting-edge findings and provide an informative survey on the most challenging topics and advances in the fields of medicine, biology, clinical engineering, and biophysics.

The series aims at disseminating high quality scientific information, encouraging both basic and applied research, and promoting world-wide collaboration between researchers and practitioners in the field of Medical and Biological Engineering.

Topics include, but are not limited to:

- Diagnostic Imaging, Image Processing, Biomedical Signal Processing
- Modeling and Simulation, Biomechanics
- Biomaterials, Cellular and Tissue Engineering
- Information and Communication in Medicine, Telemedicine and e-Health
- Instrumentation and Clinical Engineering
- Surgery, Minimal Invasive Interventions, Endoscopy and Image Guided Therapy
- Audiology, Ophthalmology, Emergency and Dental Medicine Applications
- Radiology, Radiation Oncology and Biological Effects of Radiation
- Drug Delivery and Pharmaceutical Engineering
- Neuroengineering, and Artificial Intelligence in Healthcare

IFMBE proceedings are indexed by SCOPUS, EI Compendex, Japanese Science and Technology Agency (JST), SCImago. They are also submitted for consideration by WoS.

Proposals can be submitted by contacting the Springer responsible editor shown on the series webpage (see “Contacts”), or by getting in touch with the series editor Ratko Magjarevic.

Fernando Emilio Ballina · Ricardo Armentano ·
Rubén Carlos Acevedo ·
Gustavo Javier Meschino
Editors

Advances in Bioengineering and Clinical Engineering

Proceedings of the XXIV Argentinian Congress of Bioengineering (SABI 2023), October 3–6, 2023, Buenos Aires, Argentina - Volume 1

Springer

Editors
Fernando Emilio Ballina
Instituto de Ingeniería y Agronomía
Universidad Nacional Arturo Jauretche
Florencio Varela, Buenos Aires, Argentina

Rubén Carlos Acevedo
Facultad de Ingeniería
Universidad Nacional de Entre Ríos
Oro Verde, Argentina

Ricardo Armentano
EMBS IEEE Technical Committee on
Cardiopulmonary Systems
and Physiology-Based Engineering
New Jersey, USA

Gustavo Javier Meschino
Instituto de Investigaciones Científicas y
Tecnológicas en Electrónica
Universidad Nacional de Mar del Plata
Mar del Plata, Argentina

ISSN 1680-0737 ISSN 1433-9277 (electronic)
IFMBE Proceedings
ISBN 978-3-031-61959-5 ISBN 978-3-031-61960-1 (eBook)
https://doi.org/10.1007/978-3-031-61960-1

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

The XXIV Bioengineering Congress and XIII Clinical Engineering Conference (SABI 2023) was held in the Ciudad Autónoma de Buenos Aires (Argentine) and in the city of Florencio Varela from October 3–6, 2023. The event represents the scientific meeting of the Argentine Society of Bioengineering, and on this occasion, it was organized by the Arturo Jauretche University.

The congress covered topics such as bioinstrumentation, digital signal processing and biomedical images, rehabilitation engineering, biomaterials and tissue engineering, clinical engineering, bioinformatics, modeling and simulation of biological systems, medical informatics, education, among others.

The IFMBE organized a special session on Biomedical Engineering Education for professionals and students as well as a special session on Women in Biomedical Engineering.

As a satellite event of the congress, the so-called Student SABI was held, an event aimed especially at students in which presentations by specialists, 37 works showcased, workshops, and visits to companies were held. The objective of this event is to strengthen the bond between students from different universities and promote the exchange of experiences between them.

It is both our pleasure and honor to extend a cordial welcome to all participants actively engaging in the exploration of the proceedings of SABI 2023. The conference showcased an impressive array of over 145 research papers and ten conferences by international experts, all converging to deliberate on the challenges intrinsic to the advancement of future technologies in medicine and biology.

Conferences of this nature inherently serve the purpose of facilitating social interactions among individuals who share common interests and expertise. These gatherings provide attendees with the opportunity to extract novel insights, exchange prevailing ideas, and delve into critical aspects of healthcare. This conference, therefore, stands as an invaluable platform not only to stay updated within one's specific area of expertise but also to explore the forefront of advancements in other domains. While an attendee's specialization may extend beyond the realm of Medical and Biological Engineering, the compilation of works presented herein holds the potential to provide noteworthy insights capable of revolutionizing approaches to broader challenges.

We are confident that each of you found considerable satisfaction in the extensive opportunities offered during SABI 2023. The event proved to be a remarkable confluence of experiences and expertise spanning a wide spectrum of fields, all encapsulated under a unified roof. This collaborative endeavor has undoubtedly sparked a tangible wave of motivation and diversity, resonating not only across the Americas but also reverberating throughout the global landscape

Organization

Committees

Conference Chair

Fernando Ballina	Instituto de Ingeniería y Agronomía, Universidad Nacional Arturo Jauretche, Buenos Aires, Argentina

Scientific Chair

Ricardo Armentano	EMBS IEEE Technical Committee on Cardiopulmonary Systems and Physiology-Based Engineering, New Jersey, USA

Scientific Committee Members

Ruben Acevedo	Universidad Nacional de Entre Ríos
Enrique Avila Perona	Universidad Nacional de San Juan
Fernando Ballina	Universidad Nacional Arturo Jauretche
Ramiro Barreiro Saravia	Ministerio de Salud Pública de Salta
Maria Belluzo	Universidad Nacional de La Plata
Diego Beltramone	Universidad Nacional de Córdoba
Martin Belzunce	Universidad Nacional de San Martín
Paula Bonomini	Instituto Tecnológico de Buenos Aires
Ariel Braidot	Universidad Nacional de Entre Ríos
Diego Campana	Universidad Nacional de Entre Ríos
Victor Carmona	Universidad Nacional de San Juan
Carolina Carrere	Universidad Nacional de Entre Ríos
Mariano Casciaro	Universidad Favaloro
Paola Catalfamo Formento	Universidad Nacional de Entre Ríos
Santiago Collavini	Universidad Nacional Arturo Jauretche
Diego Comas	Universidad Nacional de Mar del Plata
Damian Craiem	Universidad Favaloro
Ronald Del Aguila Heidenreich	Universidad Nacional de Córdoba
Juan Fernández	Universidad Nacional de La Plata
Mariano Fernandez Corazza	Universidad Nacional de La Plata

Lucila Figueroa Gallo	Universidad Nacional de Tucumán
Jose Gallardo	Instituto Universitario del Hospital Italiano
Pablo Gleiser	Instituto Tecnológico de Buenos Aires
Sebastian Graf	Universidad Favaloro
Juan Graffigna	Universidad Nacional de San Juan
Alejandro Hadad	Universidad Nacional de Entre Ríos
Ramiro Irastorza	Universidad Nacional Arturo Jauretche
Sergio Liberczuk	Universidad Nacional Arturo Jauretche
Rossana Madrid	Universidad Nacional de Tucumán
Ana Carolina Maldonado	Universidad Nacional de Córdoba
Maria Carla Mantaras	Universidad Nacional de Entre Ríos
Ignacio Marolla	Universidad Nacional Arturo Jauretche
Ezequie Mazzoni	Universidad Nacional Arturo Jauretche
Gustavo Meschino	Universidad Nacional de Mar del Plata
Carlos Muravchik	Universidad Nacional de La Plata
Tamara Oberti	Universidad Nacional de La Plata
Juan Manuel Olivera	Universidad Nacional de Tucuman
Federico Paschetta	Instituto Tecnológico de Buenos Aires
María Perez	Universidad Nacional de San Juan
Pablo Peruzzo	Universidad Nacional de La Plata
Sergio Ponce	Universidad Tecnológica Nacional
Agustina Portu	Universidad Nacional de San Martín
Emiliano Ravera	Universidad Nacional de Entre Ríos
Marcelo Risk	Instituto Universitario del Hospital Italiano
Héctor Rodrigo	Universidad Nacional de San Juan
Jorge Daniel Romero	Ministerio de Salud de Tierra del Fuego
Débora Rubio	Universidad Tecnológica Nacional
Eduardo Salinas	Universidad Nacional Arturo Jauretche
Enrique Spinelli	Universidad Nacional de La Plata
Carolina Tabernig	Universidad Nacional de Entre Ríos
Emanuel Tello	Universidad Nacional de San Juan
Rosa Weisz	Universidad Nacional de Entre Ríos
Sandra Wray	Instituto Universitario del Hospital Italiano
Bonifacio Zanutto	Universidad de Buenos Aires

Contents

Artificial Intelligence and Data Science

Implementation of Generative Adversarial Neural Networks for Lung Ultrasound Image Synthesis: Quality-Based Optimal Latent Space Dimension Selection Using FID Score

Iván A. Lisman[1,2], Ricardo A. Veiga[2,3], and Fabián Acquaticci[1,3(✉)]

[1] Instituto Nacional de Tecnología Industrial, Ministerio de Economía, San Martín, Buenos Aires, Argentina
{ilisman,facquaticci}@inti.gob.ar

[2] Departamento de Electrónica, Facultad de Ingeniería, Universidad de Buenos Aires, Buenos Aires, Argentina

[3] Instituto de Ingeniería Biomédica, Universidad de Buenos Aires, Buenos Aires, Argentina

Abstract. This paper presents the optimization of generative adversarial neural networks with the objective of improving the generation of synthetic lung ultrasound images, using the Fréchet distance in Inception as a quality metric. The latent space or input noise is transformed into synthetic data, so its size is a fundamental hyperparameter for the network. An optimal latent space dimension was found for this network and the training data, improving the FID value by 19.26% in comparison with the worst case achieved. The quality of the images generated by the optimized model was quantified comparing their FID with those obtained by degrading the original images using various filters. The results showed that the quality of the generated images is similar to those that suffered low level of degradation. This study contributes to the field of lung ultrasound image synthesis since optimization based on a quality metric provides a quantitative approach to improve the quality of the generated images. These findings could be relevant to improve the generation of medical images, which could be used to perform data augmentation in order to improving the fitting of classifier models that use such images for training.

Keywords: Lung Ultrasound · Generative Adversarial Network · Fréchet Distance · Latent Space

1 Introduction

Even though chest radiography is the "golden standard" for diagnosing and monitoring pulmonary diseases, the use of lung ultrasound (LUS) in patients with COVID-19 [1] is a novel technique that presents great benefits such as repeatability and low cost. Besides, it does not require patients to be transported from their homes or moved from their beds in intensive care cases [2].

F. E. Ballina et al. (Eds.): SABI 2023, IFMBE Proceedings 106, pp. 3–15, 2024.
https://doi.org/10.1007/978-3-031-61960-1_1

With the growth of machine learning techniques, the information obtained from LUS has begun to be used in order to explore several diagnostic strategies. Those include classifier models that aim to decide, with greater accuracy, if the images associated with a patient might be from a healthy individual or one with some respiratory disease [3].

The training of these models generally requires a great amount of samples, which are often not available because of a variety of reasons in the medical field (labeling, confidentiality, among others). Data augmentation techniques, like the use of Generative Adversarial Networks [4] (GANs), are commonly employed to solve these limitations. Using GANs for this purpose, particularly on ultrasound images [5], is still a novel technique. The objective is to synthetize a sufficient number of samples, allowing the assemble of a larger dataset for posterior training of classifier models.

This paper quantitatively explores the optimization of the dimension (i.e., size) of the input noise or latent space of a GAN with the purpose of improving the quality of the images that it can generate. This is done as a function of the Fréchet Inception Distance [6] (FID), which serves as a quality metric.

FID is one of the most commonly used metrics that employs the Inception V3 network [7] in order to obtain the main features of both the original (used to train the GAN) and synthetic images (generated by the GAN); thus measuring the relative quality between them. This is mainly done by observing the level of "inheritance" the generated images have with respect to the originals. The Inception V3 network is the core of most state-of-the-art computer vision solutions, achieving a really good performance on image recognition [8]. It has 48 hidden layers, around 25 million parameters and it was created for the recognition of the Imagenet dataset [9] objects, which has more than 14 million labeled images from around 20,000 categories. Recent works indicate a correlation between FID and medical image realism [5, 10–12]. Considering the aforementioned, the Inception V3 network might be able to capture the main features and the diversity of the distribution of lung ultrasound images, even if it has not been trained with such data.

The transformation of the latent space into generated data is a crucial part of a GAN operation. However, among the diverse recent studies about these networks, only a few of them paid attention to the dimension of the latent space [13, 14].

2 Materials and Methods

2.1 Lung Ultrasound Images

In this preliminary study, the lung ultrasound dataset used is the combination of several publicly available resources, such as BC PoCUS, GrepMed, Nephropocus and The Point Care Ultrasound Atlas. These are videos from 6 patients with no pathologies, which were segmented into frames, giving a total of 304 images of healthy individuals. All of them have a resolution of 256×256 pixels and are in a grayscale of 256 values.

2.2 Framework (Hardware and Software)

The algorithms for all the experiments were programmed in Python using Keras 2.2.4-tf from the Tensorflow library (version 2.1.0). All of them were executed on a Geforce

GTX 960 with 4 GB of video memory, whose resources were used through the tools provided by NVIDIAS's CUDA platform (version 10.1.0).

2.3 DCGAN

The generation of synthetic images was carried out using an architecture variant of GANs known as Deep Convolutional GANs [15] (DCGANs), which allows for stability in training and leads to higher-resolution images.

Up or down-sampling was performed using convolutional layers of 5×5 kernels and weights randomly initialized from a normal distribution of zero-mean and standard deviation 0.02. These convolutional layers were modified in order to perform kernel strides of 2 so the resolution will be doubled in the case of the generator or decrease by half in the case of the discriminator.

Batch normalization [16] was performed at all the layers (except the output ones) on both the discriminator and generator, stabilizing the training and normalizing the input of each unit with the aim of having zero-mean and unity variance.

The use of ReLU as activation function was replaced by its variant LeakyReLU with a value of α (leakage slope) equal to 0.2, except for the generator and the discriminator outputs, in which the hyperbolic tangent and sigmoid function were used, respectively.

The LUS images for training were scaled to 128×128 pixels with a range of $[-1, 1]$ grayscale values. Stable trainings were achieved using Adam optimizer [17] with a learning rate of 0.0001 for the generator and 0.00001 for the discriminator, both with a momentum β_1 of 0.5.

The network architecture used for the discriminator is depicted in Fig. 1 and each layer details can be seen in Table 1.

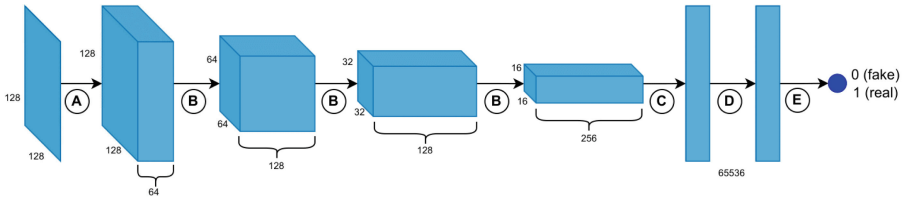

Fig. 1. DCGAN Discriminator (layer details in Table 1).

A similar dual architecture for the generator is shown in Fig. 2 and Table 2.

Some extra layers were added for the sake of functionality, such as a “reshape” in order to project the latent space, and a “dropout” to avoid overfitting, for example.

Table 1. Discriminator layers

Reference	Activation layers		
A	Convolution (5×5)	Batch Norm	LeakyReLu (0.2)
B	Convolution (5×5 & stride $= 2$)	Batch Norm	LeakyReLu (0.2)
C	Flatten		
D	Dropout (0.4)		
E	Dense/Fully Connected (sigmoid activation)		

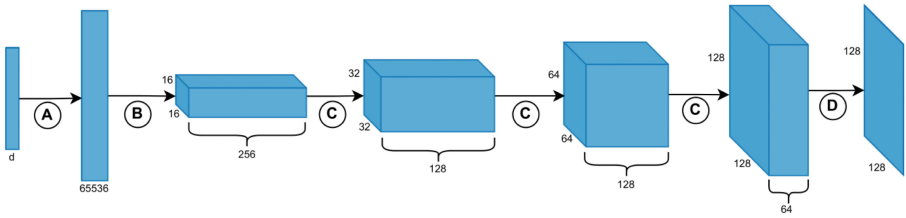

Fig. 2. DCGAN Generator (layer details in Table 2).

Table 2. Generator layers.

Reference	Activation layers		
A	Dense/Fully Connected	Batch Norm	LeakyReLu (0.2)
B	Reshape		
C	Deconvolution (5×5 & stride $= 2$)	Batch Norm	LeakyReLu (0.2)
D	Convolution (5×5) (tanh activation)		

2.4 Quality Metric (FID)

For the quality evaluation of the images using FID, a "prediction" through the Inception network must be done. Specifically, the coding layer (the last layer before the output classification of the images) is used to capture specific features from the input images. These activations are calculated for both sets of images (the synthetically generated and the originals). An activation vector of dimension 2048 was obtained for each image prediction, thus giving a 304 × 2048 activation matrix for each group to be compared (originals and generated), whose mean and covariance matrix were calculated. Finally, the FID between the original images and the synthetically generated ones can be calculated as:

$$\mathrm{FID} = \left|\overline{\mu}_{\mathrm{orig}} - \overline{\mu}_{\mathrm{gen}}\right|^2 + \mathrm{tr}\left(\sum\nolimits_{\mathrm{orig}} + \sum\nolimits_{\mathrm{gen}} - 2\left(\sum\nolimits_{\mathrm{orig}}\sum\nolimits_{\mathrm{gen}}\right)^{\frac{1}{2}}\right) \quad (1)$$

The distance between the two distributions is calculated using Fréchet distance, where the real part of any complex number was considered when necessary.

Degradation. With the purpose of interpreting the quality metric for the synthetic LUS images, the original images (304 frames from healthy patients) were compared against themselves after being degraded with different noises [6]:

- **Gaussian blur:** The image is convolved with a gaussian kernel with standard deviation $\alpha \in [0, 6]$. When α is larger, the image is smoother (larger disturbance).
- **Gaussian noise:** A matrix N with Gaussian noise scaled to [0, 255] was constructed. The noisy image is computed as $(1 - \alpha)X + \alpha N$ for $\alpha \in [0,1]$ where X is the original image. When α is larger, the image is noisier (larger disturbance).
- **Salt&Pepper noise:** Some pixels are randomly chosen to be flipped to white or black, with 50% probability, where the ratio of pixel flipped to white or black is given by the noise level $\alpha \in [0, 0.4]$. When α is larger, more noise is added by flipping pixels to white or black (larger disturbance).

For each filter applied in order to degrade the image, α took 15 equally spaced values within the indicated intervals.

2.5 Latent Space Optimization

In this paper, the latent space dimension was selected as a fundamental hyperparameter to be optimized. It was varied by taking values of $d \in \{4, 10, 16, 50, 64, 100, 128, 256, 512\}$. For each one of them, a total of 5 values of FID were obtained to account for the variability of the metric.

On the other hand, the batch size took values $batch\_size \in \{10, 30\}$, where each one was fixed for each latent space dimension training; in order to reinforce the conclusion regarding the explored hyperparameter effect on the network.

The number of epochs used in training was set to 1000. This was determined by observing during training that the loss curves no longer fluctuated abruptly. This limits the computation time (e.g., in cases where resources are limited) and, at the same time, avoids stopping the learning prematurely.

3 Results

3.1 GAN Optimization

The mean FID, for both batch sizes, from the evaluation after the training of the models is depicted in Fig. 3. Each point is associated with an error band representing the highest and lowest FIDs of the 5 trials mentioned in Subsect. 2.5.

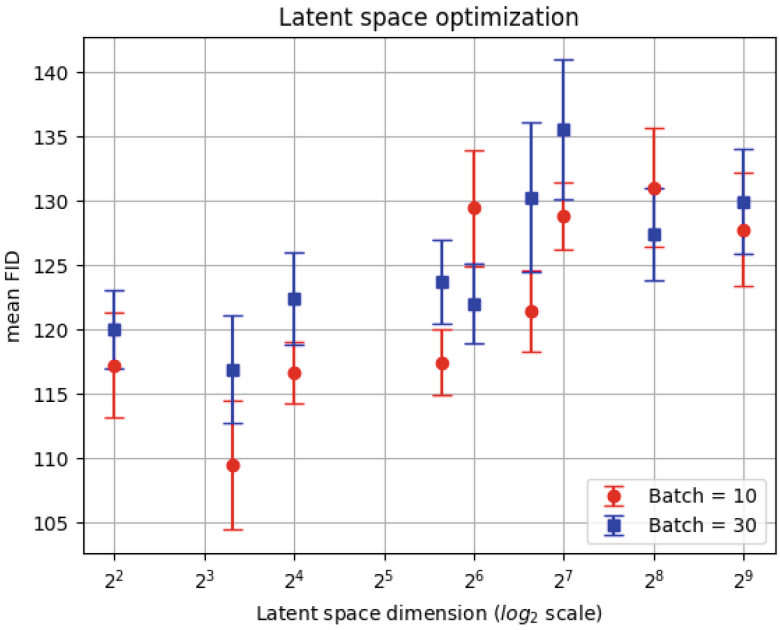

Fig. 3. Relationship between the FID associated with each model and its latent space dimension.

The minimum mean FID, for both batches, is found when the latent space dimension of the model is 10, whose value is 109.389 and with a batch size of 10. In Fig. 4 the loss of this model throughout the epochs is shown. In the case of the discriminator, the value remains stable around 0.5 (dotted line).

3.2 Image Degradation and FID

The effect of the degradation phenomenon is shown in Fig. 5 for the Gaussian blurring on a healthy patient, while in Fig. 6 and 7 the results for Gaussian noise and Salt&Pepper noise can be seen, respectively.

In Fig. 8, a quantitative comparison of the FID of lung ultrasound images, as a function of the degradation level applied to the image for different noises, is shown. In addition, the dotted line highlights the mean FID value reached by the synthetic images obtained with the optimized generator.

Finally, Fig. 9 shows the comparison between a synthetic image obtained with the GAN having its latent space optimized and an image degraded by Gaussian blurring to the point of obtaining the same FID value as the generated image. Furthermore, Fig. 10 and Fig. 11 display the results for Gaussian and "Salt&Pepper" noises, respectively.

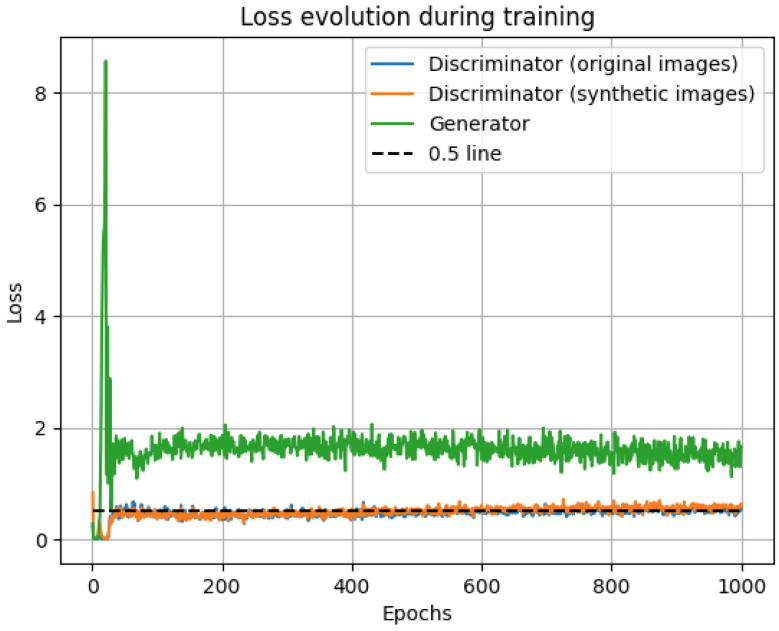

Fig. 4. Loss updates over the course of the training for the best mean FID network.

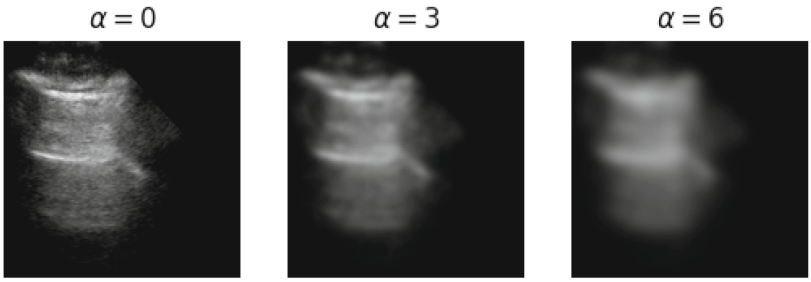

Fig. 5. Original image and its variants degraded by Gaussian blurring.

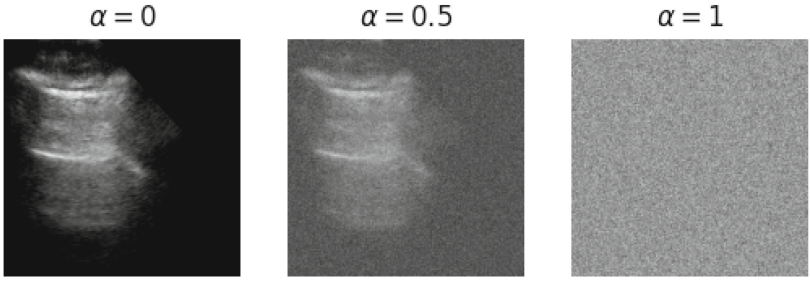

Fig. 6. Original image and its variants degraded by Gaussian noise.

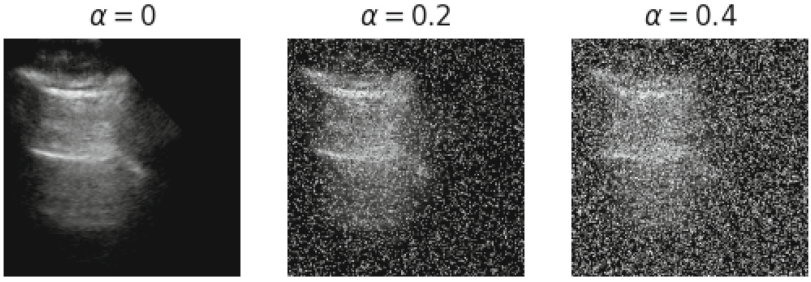

Fig. 7. Original image and its variants degraded by "Salt&Pepper" noise.

4 Discussion

In general, one of the main problems of GANs is their difficult convergence, associated with high instability when training them. To tackle the various factors that may cause this, numerous architectures have emerged over the years, DCGAN being one of them, and the one chosen for this work.

The numerical difference between the learning rates for the different actors in the network (generator and discriminator) usually leads to an improvement in convergence

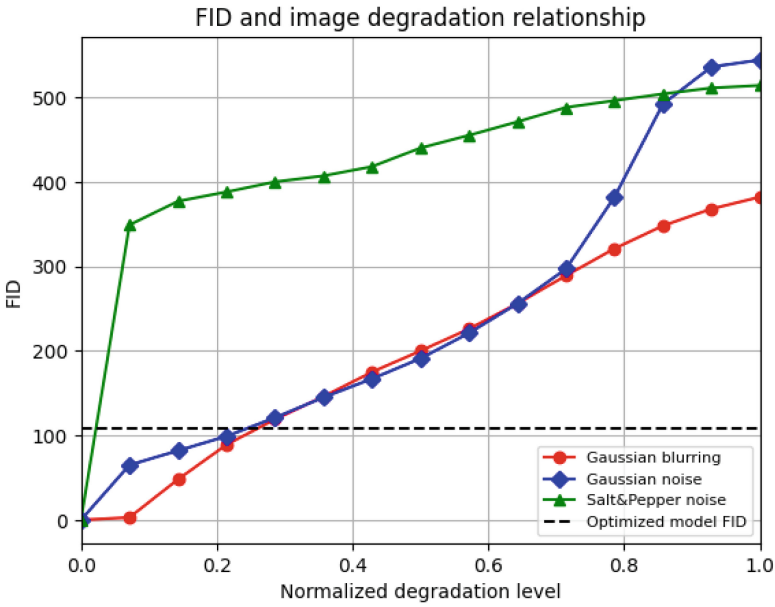

Fig. 8. FID as a function of the normalized level of image degradation for different noises, and the mean FID achieved by the optimized model.

Fig. 9. Original image, synthetic image obtained with the optimized network and a degraded image by Gaussian blur such that the last two have the same FID (around 109).

and performance [6]. In this case, a lower learning constant on the discriminator can prevent it from learning too quickly to discern between the noise initially delivered by the generator and an original image, as this can anchor the training of the generator.

In terms of the metric used, FID has inherent biases in the Inception V3 network as a feature extractor [18], and these must be taken into account when comparing different architectures. Although it is often assumed that Inception-based metrics are not suitable for evaluating medical images because their training does not include images of this

Fig. 10. Original image, synthetic image obtained with the optimized network and a degraded image by Gaussian noise such that the last two have the same FID (around 109).

Fig. 11. Original image, synthetic image obtained with the optimized network and a degraded image by Salt&Pepper noise such that the last two have the same FID (around 109).

category, FID has been found to correlate positively with human perception and to be consistent with the evaluation and realism of medical images [19, 20].

From the results obtained, it is quantitatively observed that there is a significant effect (stopping the learning process at 1000 epochs) on the image quality when the latent space dimension is changed: the worst result obtained leads to a mean FID of 135.475, while the best is 109.389, improving the image quality by 19.26%.

It is also observed that, given the training conditions, a relatively low dimension, with the lowest FID, is the optimal value for the latent space. Additionally, a lower latent space dimension means a lower computational cost: a dimension of 512 requires almost 37 million parameters, while a dimension of 10, like in the optimal case found here, requires less than 4 million. Therefore, as it increases, so does the complexity and the use of resources. That is why, for the number of epochs previously set, the "optimal latent space dimension" is not only the one that allows the lowest FID, but uses an architecture with a lower degree of complexity than those with a higher dimension.

Any modification on the latent space of the GAN can lead to significant transformations of the image, so there is a correspondence between the latent space and the image space [13, 14]. By exploring the latent space it is possible to expose whether there are

signs of memorization or whether the model is collapsed [15]. Thus, if by interpolating between two points in the latent space, semantic changes occur in the image generation (e.g. ultrasound artifacts are added or removed), it would imply that the model has learned the relevant representations of the training set of images. In Fig. 12 this phenomenon can be observed in the case of the optimized network, where the dimension of the latent space is 10.

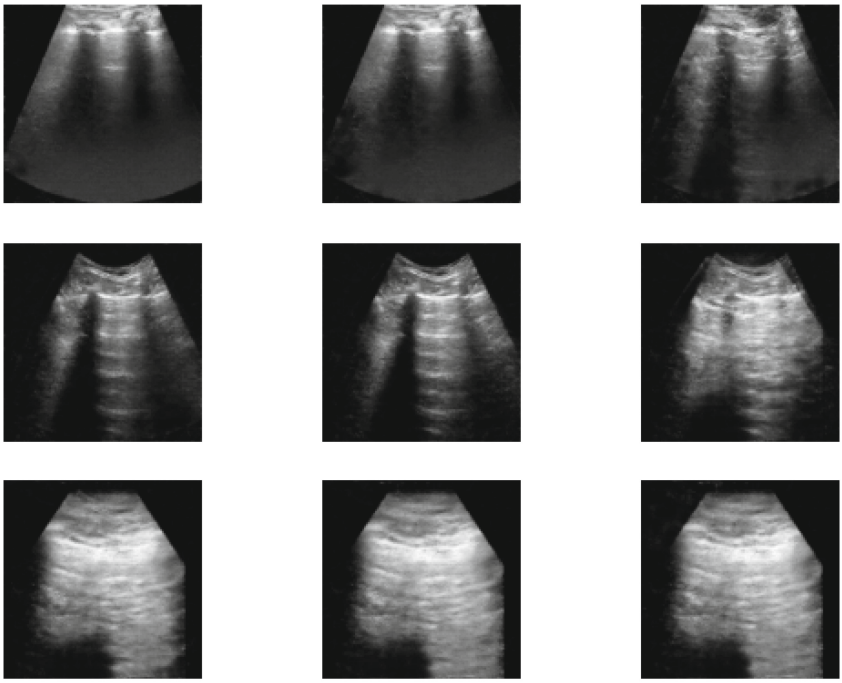

Fig. 12. Linear interpolation between 2 points of the optimal latent space found, corresponding to the first (upper left corner) and the last (lower right corner) generated image.

5 Conclusions

The effect of the latent space dimension in a GAN, designed for generating realistic lung ultrasound images, was studied in this work. Additionally, their quality was compared against noisy images, in terms of the level of image degradation.

The latent space dimension required for a GAN will depend on the complexity of the information to be captured. Thus, within the context of the experiments carried out here, it is possible to successfully train a GAN with a latent space dimension as low as 10; and even get better output quality than higher latent spaces (when stopping the training after 1000 epochs). Increasing the latent space more than necessary requires

more storage, more processing resources, and longer training times. In addition, with more parameters, more data is needed to avoid overfitting or memorization.

In future investigations, it is expected to apply the same methodology to COVID-19 images. Moreover, in order to reinforce the interpretation of the quality metric, it is planned to carry out experiments where health professionals perform visual analysis of the images. They will compare those artificially generated by the optimized network against the original images, and they also will evaluate the diagnostic value of the generated images. Finally, the intention is to evaluate whether the synthetic images obtained are useful for the training of classifier models.

Acknowledgements. This study is part of the project "Metrological evaluation of lung ultrasound using virtual vector machine for diagnosis of acute respiratory distress syndrome" (MELUS-VVM-DARDS), approved by the IADB SIM RESEARCH ENGAGEMENT OPPORTUNITY 2021–2023. This work was supported through the Research and Training grants from the National Institute of Industrial Technology.

References

1. Lichtenstein, D.A.: Lung ultrasound in the critically ill. Ann. Intensive Care **4**(1), 1–12 (2014)
2. Xue, H., Li, C., Cui, L., et al.: M-BLUE protocol for coronavirus disease-19 (COVID-19) patients: interobserver variability and correlation with disease severity. Clin. Radiol. **76**, 379–383 (2021)
3. Born, J., Wiedemann, N., Cossio, M., et al.: Accelerating detection of lung pathologies with explainable ultrasound image analysis. Appl. Sci. **11**(2), 672 (2021)
4. Goodfellow, I., Abadie-Pouget, J., Mirza, M., et al.: Generative adversarial networks. ArXiv (2014)
5. Liang, J., Yang, X., Huang, Y., et al.: Sketch guided and progressive growing GAN for realistic and editable ultrasound image synthesis. Med. Image Anal. **79**, 102461 (2022)
6. Heusel, M., Ramsauer, H., Unterthiner, T., et al.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. ArXiv (2018)
7. Szegedy, C., Liu, W., Jia, Y., et al.: Going deeper with convolutions. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, pp 1–9. IEEE (2015)
8. Szegedy, C., Vanhoucke, V., Ioffe, S., et al.: Rethinking the inception architecture for computer vision. ArXiv (2015)
9. Russakovsky, O., Deng, J., Su, H., et al.: ImageNet large scale visual recognition challenge. Int. J. Comput. Vis. **115**, 211–252 (2015)
10. Maack, L., Holstein, L., Schlaefer, A.: GANs for generation of synthetic ultrasound images from small datasets. Curr. Dir. Biomed. Eng. **8**, 17–20 (2022)
11. Middel, L., Palm, C., Erdt, M.: Synthesis of medical images using GANs. In: Greenspan, H., et al. (eds.) CLIP UNSURE 2019 2019. LNCS, vol. 11840, pp. 125–134. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32689-0_13
12. Uzunova, H., Ehrhardt, J., Jacob, F., et al.: Multi-scale GANs for memory-efficient generation of high-resolution medical images. ArXiv (2019)
13. Marin, I., Gotovac, S., Russo, M., et al.: The effect of latent space dimension on the quality of synthesized human face images. J. Commun. Softw. Syst. **17**, 124–133 (2021)
14. Bojanowski, P., Joulin, A., Lopez-Paz, D., et al.: Optimizing the latent space of generative networks. ArXiv (2019)

15. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. ArXiv (2015)
16. Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. ArXiv (2015)
17. Kingma, D.P., Ba J.: Adam: a method for stochastic optimization. ArXiv (2014)
18. Jung, S., Keuper, M.: Internalized biases of Fréchet inception distance. In: NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications (2021)
19. Woodland, M., et al.: Evaluating the performance of StyleGAN2-ADA on medical images. In: Zhao, C., Svoboda, D., Wolterink, J.M., Escobar, M. (eds.) SASHIMI 2022. LNCS, vol. 13570, pp. 142–153. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-16980-9_14
20. Bargsten, L., Schlaefer, A.: SpeckleGAN: a generative adversarial network with an adaptive speckle layer to augment limited training data for ultrasound image processing. Int. J. Comput. Assist. Radiol. Surg. **15**, 1427–1436 (2020)

An Automatized Online Platform for Left Ventricular Behavior Assessment Based on Echocardiographic Recordings

Julián F. Schiffer[1], Federico E. Bancalari Solá[1], Matías J. Gasparini[1], Marcos N. Ortiz[1], Giuliana A. Posteraro[1], Ignacio Farro[2], Ricardo L. Armentano[1,2], and Leandro J. Cymberknop[1(✉)]

[1] Grupo de Investigación y Desarrollo en Bioingeniería (GIBIO), Universidad Tecnológica Nacional, Buenos Aires, Argentina
gibio@frba.utn.edu.ar

[2] Departamento de Ingeniería Biológica, Facultad de Ingeniería, Universidad de La República, Montevideo, Uruguay

Abstract. Since its beginnings, cardiovascular (CV) medicine has relied on professionals' judgment to diagnose various types of studies carried out on patients. Today, with the support of technologies such as machine learning (ML) and artificial neural networks, medical professionals can automate, speed up or improve their diagnoses by obtaining complementary information on the detection of diseases and CV conditions. This work proposes an online platform based on ML methods, applied to the analysis of echocardiographic recordings. The aim is to obtain dynamic measurements of an individual left ventricle behavior, and contributing to the detection of different types of heart conditions, thereby improving diagnostic precision.

Keywords: Cardiovascular · Echocardiogram · Left ventricle · Deep Learning · Online Platform

1 Introduction

Heart diagnoses from echocardiograms are currently carried out by cardiologists, which introduces a margin of human error. The rate of misdiagnosis in clinical practice has been estimated at 150 out of 1,000 patients (Health Grades Patient Safety in American Hospitals Study, 2011), or 10 to 20% (in the emergency department and in autopsy discrepancy studies) [1].

Currently, artificial intelligence (AI) and machine learning (ML) methods are being incorporated to totally or partially complement a diagnosis made by a physician. This makes it possible, among other things, to reduce the margin of human error and make them more accurate. In general terms, AI is basically concerned with computational understanding of what is commonly known as intelligent behavior and the creation of artefacts that exhibit such behavior [2]. Particularly, ML models are applied to search

F. E. Ballina et al. (Eds.): SABI 2023, IFMBE Proceedings 106, pp. 16–22, 2024.
https://doi.org/10.1007/978-3-031-61960-1_2