



HAL
open science

Generative Adversarial Networks in Cardiology

Youssef Skandarani, Alain Lalande, Jonathan Afilalo, Pierre-Marc Jodoin

► **To cite this version:**

Youssef Skandarani, Alain Lalande, Jonathan Afilalo, Pierre-Marc Jodoin. Generative Adversarial Networks in Cardiology. Canadian Journal of Cardiology, 2022, 38 (2), pp.196-203. 10.1016/j.cjca.2021.11.003 . hal-03612001

HAL Id: hal-03612001

<https://hal.science/hal-03612001>

Submitted on 22 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Generative Adversarial Networks in Cardiology

Youssef Skandarani, Engr, MSc ^{a,*}, Alain Lalande, Ph.D ^{a,b}, Jonathan Afilalo, MD ^{c,**} and Pierre-Marc Jodoin, Ph.D ^{d,**}

^aLaboratoire ImVIA, Université de Bourgogne, 64 rue Sully, 21000 Dijon, France

^bMedical Imaging Department, University Hospital of Dijon, 1 Bld Jeanne d'Arc, 21079, Dijon, France

^cJewish General Hospital, McGill University, 3755 Côte Ste-Catherine Road, Montreal, Qc, Canada, H3T 1E2

^dUniversité de Sherbrooke, 2500 Boul. de l'Université, Sherbrooke, Qc, Canada, J1K 2R1

SUMMARY

Since their inception, Generative Adversarial Networks (GANs) enabled new frontiers in data synthesis in multiple fields. Cardiology is not an exception. In this work we present an overview of GANs, their applications and clinical impact in cardiology as well as the future perspectives they enable for cardiology and the challenges yet to be solved. In addition we present a showcase of their capabilities on cardiac imaging data.

ABSTRACT

Generative Adversarial Networks (GANs) are state-of-the-art neural network models used to synthesize images and other data. GANs brought a considerable improvement to the quality of synthetic data, quickly becoming the standard for data generation tasks. In this work, we summarize the applications of GANs in the field of cardiology, including generation of realistic cardiac images, electrocardiography signals, and synthetic electronic health records. The utility of GAN-generated data is discussed with respect to research, clinical care, and academia. Moreover, we present illustrative examples of our GAN-generated cardiac magnetic resonance and echocardiography images, showing the evolution in image quality across six different models, which has become almost indistinguishable from real images. Finally, we discuss future applications, such as modality translation or patient trajectory modeling. Moreover, we discuss the pending challenges that GANs need to overcome, namely their training dynamics, the medical fidelity or the data regulations and ethics questions, to become integrated in cardiology workflows.

1. Introduction

Generative adversarial networks (GANs) are a type of neural network dedicated to the creation of synthetic data that is similar to a real, reference dataset [1]. For healthcare applications, GANs have been primarily used to synthesize or post-process medical images; and secondarily to synthesize electronic health data [2]. In recent years, the quality of GANs has surpassed that of other generative models (e.g. auto-encoders and energy-based models) such that GANs are now regarded as the *de facto* solution for synthetic data generation.

Regarded as one of the most notable artificial intelligence discoveries of the past decade, colloquially birthed around the table of a Montreal brewery circa 2014, GANs have been incorporated in a variety of computer vision and data analysis applications ranging from simulated radiographic images to synthetic medical records.

GANs were introduced by Ian Goodfellow et al.[1] as a new class of probabilistic generative models that overcome many of the drawbacks of the previously prevalent generative models such as Gaussian mixtures, Markov chains and, more recently, variational autoencoders (VAE). GANs substantially improved the quality of the generated data while maintaining a straightforward sampling approach.

Initially showcased on general datasets such as that of human faces or objects, GANs achieved high quality samples compared to previous approaches. Moreover, GANs provide a way to generate more diverse data samples, such as faces with blue or brown eyes. These powerful models have also been shown surprisingly effective at generating computer artworks, some being even exposed and sold [3] !

The aforementioned success of GANs, prompted researchers to investigate their capability in more

fields that would benefit from their generative modeling performance.

These applications have been used mainly in research projects, while their use in clinical practice still faces some challenges. This popularity of GANs can be attributed to the following three characteristics:

1. Unlike most neural networks, GANs are unsupervised models that do not necessitate (labor-intensive) human annotations to be trained;
2. GANs are simple, generic models that accommodate a wide variety of datasets, including medical data;
3. The high resolution and quality of data produced by the latest GAN models easily confound human observers, including medical experts.

Here, we review the main characteristics of GANs, discuss their applications and limitations in cardiology, and compare six widely cited models which we trained to synthesize images based on two publicly available datasets, namely the Automated Cardiac Diagnosis Challenge (ACDC) dataset [4], a cardiac magnetic resonance (CMR) dataset, and the Cardiac Acquisitions for Multi-structure Ultrasound Segmentation (CAMUS) dataset [5], an echocardiography dataset. The models used for this purpose are the *Deep Convolutional GAN* [6], the *Least Squares GAN* [7], the *WassersteinGAN* [8], the *HingeGAN* [9], the *StyleGAN* [10], and the *SPatially Adaptive DENormalization (SPADE) GAN* [11].

Overall, this work is intended as an overview of the different usages that GANs have in cardiology and their potential applications and future directions. In addition, we give a brief look at the capability of GANs in the generation of cardiac images and the evolution of their generative powers throughout the years.

2. Generative Adversarial Networks

A GAN is composed of two connected deep neural networks: a Generator and a Discriminator. The Generator receives random data as its input and attempts to generate pseudo-real synthetic data as its output (e.g. a fake CMR image).

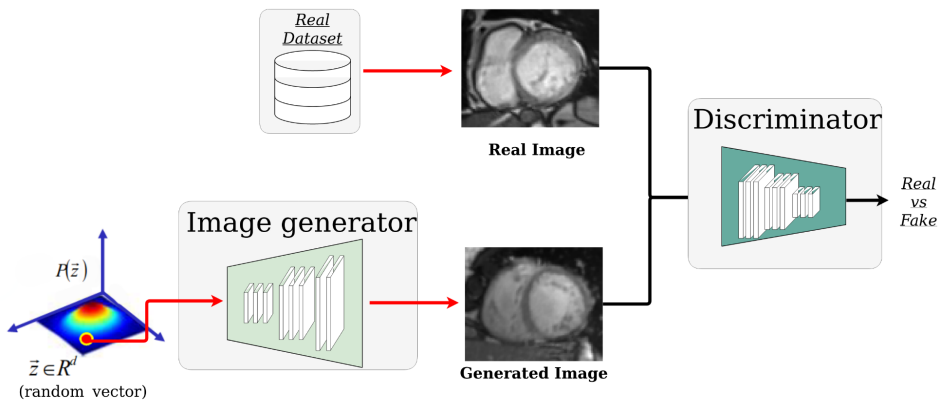


Figure 1: GANs are neural networks dedicated to the synthesis of highly-realistic data (e.g. cardiac MRIs). During the training of the GAN, two neural networks (Generator and Discriminator) compete against each other. The Generator produces synthetic images and the Discriminator attempts to differentiate real images from generated images.

The Discriminator receives either the generated or real data as its input and attempts to predict whether it is fake or real as its output (Figure 1). The real data (e.g. a collection of actual CMR images) is provided *a priori* and plays a major role in training the GAN and shaping the final appearance of the generated synthetic data.

The goal of the Discriminator is to differentiate real data from synthetic data. The goal of the Generator is to produce synthetic data whose characteristics are as close to that of the real data as possible, inasmuch to "fool" the Discriminator into classifying it as real. These adversarial goals are analogous to the art forger (i.e. the Generator) and art authenticator (i.e. the Discriminator). Training of the Generator's neural network is carried out through a gradient descent algorithm to minimize a loss function representing the inaccuracies of the synthetic data relative to the real data. This loss function is

ingenuously back-propagated from the Discriminator's loss function (cross-entropy), such that the Generator can be trained without its own separate loss function. Through this process, the Generator progressively learns to produce samples with characteristics matching those of the real dataset, thereby making them almost indistinguishable.

Practically, GANs are trained in a two step fashion. First, the generator is trained by feeding it random vectors that would make it generate samples of data (images typically). This in turn is fed to the discriminator which is optimized to predict whether it is real or fake and the error signal is back-propagated to train the generator. Next, the discriminator is trained to discriminate between real and fake data. This is done by feeding the discriminator network with real and synthetic data and optimizing it to correctly predict the real and the fake data. This process is repeated across the dataset at hand until the discriminator is not able to distinguish between fake and real data.

Later, the trained generator is used to sample a number of new data points by varying the input noise and collecting the generated output.

3. Applications in Cardiology

In this section, we review the foremost applications of GANs encompassing different types of data in cardiology, namely: cardiac imaging data, electrocardiography (ECG) signal data, and observational health data.

3.1. Cardiac Imaging

Since their inception, GANs have primarily been used to generate images [12]. In cardiology, GANs have been used to generate images of the cardiovascular organs successfully mimicking modalities such as CMR, computed tomography, and echocardiography. For example, Amirrajab *et al.* [13] combined a heart simulator [14] and a GAN to generate synthetic short-axis cine CMR images at various slice locations. Skandarani *et al.* [15] combined a cardiac segmentation map and a VAE-GAN system to generate synthetic short-axis cine CMR images with accompanying anatomical labels. Furthermore, inputting disease labels alongside training images enabled GANs to generate synthetic images of specified cardiovascular diseases, also known as conditional image synthesis.

While the clinical utility of image synthesis is not yet well defined, GANs have been useful to generate large numbers of synthetic scans in order to train (data-hungry) machine learning models to predict cardiac diagnoses or automatically segment cardiac chamber volumes. In turn, these machine learning models may yield tangible benefits for workflow efficiency and patient care. Creating such synthetic data is far less costly and time-consuming than acquiring and manually labelling more real scans. Skandarani *et al.* [15] demonstrated improvements in their CMR segmentation neural network model after augmenting their training dataset with 10,000 synthetic images.

Conditional image synthesis with GANs has enabled exciting demonstrations of modality translation, effectively converting any given image from its original modality to another modality. For example, Zhang *et al.* [16] trained a cycleGAN [17] to convert CMR images to computed tomography (and vice versa) for enhanced visualization of calcified structures that were otherwise difficult to see on CMR. It has also enabled disease simulation, altering any given image to reflect the expected phenotypical progression of a specified disease over time (using similar algorithms to those used for face aging simulations). Moreover, Chen *et al.* [18] proposed a GAN based approach to generate virtual patient images that manage to be visually realistic while maintaining a feature space that is pathophysiologically interpretable.

GANs have been useful in post-processing to generate improved images with less noise, less artifacts, and greater resolution (termed *super resolution*) than the raw acquired images. For example, Xia *et al.* [19] used a GAN to increase the inter-slice resolution of cine CMR images from $1.8 \times 1.8 \times 10 \text{ mm}^3$ (slice thickness 8 mm and slice gap 2 mm) to $1.8 \times 1.8 \times 0.625 \text{ mm}^3$. This process drastically improves the evaluation of cardiac measures that include ventricular and myocardium surface delineation, ventricular volumes and ejection fraction estimation. GANs have also been useful in image reconstruction to more efficiently generate images with shorter acquisition times or reduced radiation doses.

Other applications of GANs for surgery planning are also being explored. For example, Zhang *et al.* [20] showed that GANs can be used to predict the optimal stenosis repair patch location, shape and size for patients suffering from a tetralogy of Fallot.

3.2. Electrocardiography

GANs can be applied to synthesize other data types, including electrocardiography (ECG) tracings. For example, Zhu *et al.* [21] designed a GAN composed of a bidirectional LSTM (Long Short-Term Memory) as its Generator and a CNN (Convolutional Neural Network) as its Discriminator, to ultimately yield synthetic ECG tracings with characteristics mimicking cardiac diseases. The GAN performed best when compared to other generative models using recurrent neural networks and auto-encoders. Wulan *et al.* [22] compared three different models, of which the GANs performed best to generate diversified and realistic ECG tracings. As above, GANs can also be applied for conversion and post-processing. Sarkar *et al.* [23] developed a GAN model to convert photoplethysmograms to synthetic ECG tracings, leading to more accurate computations of heart rate. Singh *et al.* [24] developed a GAN model to de-noise raw ECG signals, leading to higher ratings of ECG tracing quality than other state-of-the-art denoising methods.

3.3. Observational Health Data

Electronic Health Records (EHRs) contain a wealth of patient data with the potential to fuel unprecedented discoveries, however, the completeness of this data is limited by clinical indications and its accessibility is tightly restricted by regulatory agents. GANs can be applied to synthesize alpha-numerical data as found in EHRs [25] such as age, sex, race, height, weight, comorbid conditions, vital signs, laboratory values, test results, complications, and more. This synthetic data can be inserted to either "fill in the blanks" for real patients that have missing or not-yet-observed data points, or create characteristics for imaginary patients that have similar distributions and relationships to those observed in real cohorts but cannot be identified. While earlier iterations such as *medGAN* struggled to synthesize non-continuous variables with discrete values, newer iterations overcame this limitation and evolved to accommodate tabular and time-series data architectures capable of mirroring full EHRs [26].

Use-cases have highlighted several benefits of GANs. Firstly, to improve the performance of predictive models by imputing missing data points or augmenting sparse datasets. Che *et al.* [27] showed that the addition of GAN generated synthetic data to real patient data resulted in improved performance of their CNN-based risk prediction model. Secondly, to improve the broad-scale accessibility of data by sharing privacy-preserving synthetic versions of actual EHRs. GANs can be used to generate synthetic versions of EHRs or other patient-level repositories in order to bypass the privacy concerns that accompany the use of real patient information. To this end, both Yoon *et al.* [2] and Zhang *et al.* [28] proposed GAN models to generate datasets endowed with similar properties to those of real patients. Thirdly, to improve the personalized delivery of care by creating simulations of patient outcomes that are easily interpretable for clinicians. Within the latter, potential therapies can be added or removed from simulations, yielding individualized treatment effects for particular patients.

4. Empirical Example of GANs in Cardiac Imaging

To showcase the capabilities of GANs at generating realistic cardiac images, we trained six different types of GAN models using two publicly-available datasets for CMR and echocardiography, respectively. The methodology and results are presented below, accompanied by illustrative images.

4.1. Datasets

Automated Cardiac Diagnosis Challenge (ACDC) dataset [4]. This dataset contains short-axis cine CMR sequences from 150 patients acquired at the University Hospital of Dijon. The dataset is divided into 5 evenly distributed subgroups (4 pathological groups and 1 healthy group). The images were acquired using two MRI scanners of different magnetic strengths (1.5T and 3T). The pixel size ranges from 0.7 mm to 1.9 mm. The slice spacing ranges from 5 mm to 10 mm. A clinical expert selected the end-diastolic and end-systolic frames. The same expert completed manual annotations of the endocardial and epicardial borders of the left ventricle and the endocardial border of the right ventricle.

Cardiac Acquisitions for Multi-structure Ultrasound Segmentation (CAMUS) dataset [5]. This dataset contains apical two-chamber and four-chamber cine echocardiography sequences from 500 patients, half of which were considered to be pathological on the basis of left ventricular ejection fraction <45%. The acquired sequences were resampled to a Cartesian coordinate system using a constant grid resolution of $\lambda/2 = 0.31 \text{ mm}$ on the x-axis parallel to the probe, and $\lambda/4 = 0.15 \text{ mm}$ on the z-axis perpendicular to the probe. A clinical expert selected the end-diastolic and end-systolic frames and

completed manual annotations of the endocardial and epicardial borders of the left ventricle and atrium.

4.2. Selected GANs

The literature on GANs has substantially grown over the past 8 years [29]. Among the multitude of published models, we selected six fitting the following criteria: 1) widely adopted, 2) high-quality generated images, 3) not exceeding the capabilities of a single 12GB GPU, and 4) representative of the evolution of GANs since 2016.

DCGAN-2016 Deep Convolutional GAN [6] is the first GAN to use convolutional layers as opposed to original GANs that only used fully connected layers. Due to its simplistic design, DCGAN is often used as the *de facto* baseline GAN. Nevertheless, there has been a recent and notable increase in its image quality and training stability. To our knowledge, DCGAN is among the most widely implemented GANs to-date.

LSGAN-2017 Least Squares GANs [7] improves sample quality by replacing the cross entropy loss function of the Discriminator in the original GAN by the mean squared error loss function. This reduces the problems encountered with vanishing gradients and creates a more stable training process.

WGAN-2017 Wasserstein GANs [8] is considered a significant breakthrough in GAN research. It uses the Wasserstein earth-mover distance as the GAN loss function, and in doing so, reduces the problems encountered with mode collapse and creates a more stable training process.

HingeGAN-2017 Introduced by Lim *et al.* [9], HingeGAN uses the margin maximization loss function, which theoretically converges to a Nash equilibrium between the Generator and the Discriminator. Like WGAN and LSGAN, HingeGAN stabilizes the optimization process of GANs.

SPADE GAN-2019 Spatially Adaptive DEnormalization (SPADE) GAN [11] is a type of GAN for *image-to-image* translation, improving upon the previously published pix2pix [30] model. SPADE GAN produces high-quality results on a wide range of datasets, perfectly matching a semantic input mask.

Style GAN-2020 StyleGAN [10] is considered the most state-of-the-art generative neural network that builds on previous models, such as progressive GAN [31], training the GAN with increasing resolution, leading to better image quality and a more stable training process. StyleGAN has a highly modified Generator capable of synthesizing high quality and diverse images. Generator modifications include: adaptive instance normalization blocks (AdaIN), noise injection at each network level and the use a 8-layer MLP mapping network on the input latent vector \vec{z} .




Figure 2: Examples of images generated from six GANs. *Top row*: cardiac MRIs; *bottom row*: ultrasound images.

4.3. Cardiac Image Synthesis

To optimize results with the six GAN models, several hyperparameter values were tested and tuned for each model. Following hundreds of training rounds and up to 200 GPU days of processing, the best performing configuration for each model was retained. Of note, the majority of training rounds diverged and produced degenerated models, highlighting the sensitivity of GANs.

Figure 2 shows a representative set of synthetic images generated by each GAN using the CMR and echocardiography datasets. Incremental improvements in image quality are observed from the older GANs (*DCGAN*, *LSGAN*, *WGAN*, *HingeGAN*) to the more recent ones (*StyleGAN*, *SPADE GAN*) as a result of the model's sophistication and deeper training. Indeed, StyleGAN and SpadeGAN required 30 days and 10 days of training, respectively.

4.4. Subjective Image Evaluation

Visual inspection of the images revealed that CMR images produced by the older GANs lacked some anatomical details and suffered from blurry organs, rendering them less realistic. Echocardiography images produced by the older and newer GANs were more similar in part due to the fact that ultrasound images are inherently less crisp. Overall, StyleGAN was the best performing model, producing images that were indistinguishable from their real counterparts (Figure 3).

To allow cardiac imaging experts to inspect and validate the quality of GAN-generated images, we developed an online web application (<https://gans.devpods.com>) that displays real images from the ACDC dataset alongside synthetic images from StyleGAN. When presented with these images in a blinded fashion and asked to identify those images that are real and those that are synthetic, cardiac imaging experts typically achieve a ceiling of 60% accuracy.

5. Present status and future directions

GANs are exciting new tools with a variety of clinical and research applications in the field of cardiology. Their ability to synthesize realistic images, ECG tracings, and EHR data is already being used by researchers to efficiently augment the available sample sizes to train predictive models (producing models with higher accuracy which in turn improve patient care), especially when working with small datasets or rare pathologies. They are also used by various stakeholders in the healthcare setting to mine privacy-preserving synthetic versions of EHRs to gain insights into patient trajectories and care delivery, especially when access to the real data would otherwise be restricted. GANs are increasingly used in imaging post-processing platforms to produce enhanced versions of raw acquired images with super-resolution and fewer artifacts.

Future implementations of GANs will exploit their ability to conditionally synthesize data based on inputted source data and user-defined parameters. One exciting example is imaging modality translation, whereby inputted source images can be converted into another modality in order to highlight those features better represented by the latter modality which could, for example, prevent exposing the patient to ionizing radiation in case of synthesis of CT images from other modalities such as MRI or ultrasound. Another example with great potential in clinical practice is patient trajectory modeling, whereby a given patient's phenotype can be extrapolated to a future state in order to reflect the effects of a specified disease or treatment. The simulated patient has been dubbed a *digital twin*, allowing for easily interpretable representations of counterfactual conditions [32], [33].




Figure 3: Real or Fake? Real CMR images [B, C, E, G], StyleGAN-generated CMR images [A, D, F, H], real echocardiography images [2, 3, 5, 7], StyleGAN-generated echocardiography images [1, 4, 6, 8].

For example, how might this patient's heart look on CMR if he developed severe mitral regurgitation, and how might it remodel after a mitral valve replacement procedure? The technology underlying this type of GAN is similar to that employed in face aging software that addresses how a person's face might look if he aged 20 years. A different use-case for GAN-generated images is in medical education, enabling learners at all levels to gain "on-demand" access to imaging cases of any specified disease or combination of pathologies, even if rare and otherwise inaccessible in their respective centres, similar to what was proposed by Diller et al. for congenital heart disease [34]. The generative *GAN teaching file* could, in the absence of patient privacy concerns and data storage needs, be disseminated broadly and become a powerful educational tool.

5.1. Challenges

Despite their generative capabilities, GANs have associated drawbacks that may hinder their adoption in data synthesis software. It is, therefore, imperative to acknowledge their limitations to employ them judiciously and effectively. The limitations of GANs are especially apparent when working with small or imbalanced dataset, in which case the model may amplify biases within the original data. Some of these limitations are outlined below:

- Training Dynamics.** The adversarial framework of GANs renders the training process unstable and unpredictable. *Mode collapse* is a common problem wherein the generator of a GAN learns to only synthesize a subcategory of the reference dataset, thereby restricting the utility of the trained model; for example, learning to only synthesize echocardiographic images of healthy cases even though the reference dataset contains some abnormal cases. Another issue is the excessive time

required to train advanced GAN models, requiring weeks of training even with contemporary graphics processing units.

- **Model Evaluation.** The quality of generated images is difficult to objectively measure, leading to uncertainty about model convergence and added challenges in training dynamics. Measures such as the Inception Score [35], Frechet Inception Distance [36], and Kernel Inception Distance [37] represent attempts to approximate subjective human evaluations; nevertheless, these measures often correlate inadequately with expert raters.
- **Medical Fidelity.** Clinicians must be cautious about adopting new technologies as their decisions affect the well-being of patients. Since the neural networks that form the building blocks of GANs are a class of *black-box* probabilistic models, their behaviour is not fully understood. Consequently, this lack of understanding hinders the widespread acceptance of data derived from neural networks. Moreover, GAN-generated images may not fully encompass all of the subtle features of an organ or pathology. These issues should be addressed for GANs to be accepted as a reliable clinical tool.
- **Data Regulations and Ethics.** Medical data is highly sensitive as it is deeply tied to a person's health information. As such, it is still fuzzy whether synthetic data generated by GANs trained with real data related to real patients is covered by data laws and regulations. Looking at the landscape of the various regulatory frameworks around the world, there is still debate whether the generated data can be relied on as authentic medical data. Likewise, GANs may also amplify the biases that can be found in the original datasets which poses an issue about the ethics of relying on GAN generated data as a source of truth.

6. Conclusion

In the present work, we summarized the applications of GANs in cardiology, which include the generation of realistic cardiac images and ECG tracings, conversion of cardiac images from one modality to another, sophisticated post-processing of images and signals, generation of privacy-preserving synthetic EHR data, imputation of missing data, and simulation of patient treatment responses. These applications are already being used to catalyze exciting research projects, and are beginning to permeate clinical workflows and medical education. Clinical uptake is sure to accelerate as the quality of GAN-generated data continues to improve and as model training becomes more robust, interpretable, and efficient. Nonetheless, the usage of GAN-sourced data still remains shrouded by ethical and data regulations questions that need to be addressed before GANs become part of the toolset available for a clinical usage.

Funding Sources

The authors have no funding sources to declare.

Disclosures

The authors have no conflict of interest to disclose.

References

- [1] I. Goodfellow *et al.*, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, 2014, vol. 27.
- [2] J. Yoon, L. N. Drumright, and M. van der Schaar, "Anonymization Through Data Synthesis Using Generative Adversarial Networks (ADS-GAN)," *IEEE J. Biomed. Health Inform.*, vol. 24, no. 8, pp. 2378–2388, Aug. 2020, doi: 10.1109/JBHI.2020.2980262.
- [3] J. Vincent, "A never-ending stream of AI art goes up for auction," *The Verge*, Mar. 05, 2019. <https://www.theverge.com/2019/3/5/18251267/ai-art-gans-mario-klingsmann-auction-sothebys-technology> (accessed Oct. 26, 2021).
- [4] O. Bernard *et al.*, "Deep Learning Techniques for Automatic MRI Cardiac Multi-Structures Segmentation and Diagnosis: Is the Problem Solved?," *IEEE Trans. Med. Imaging*, vol. 37, no. 11, pp. 2514–2525, Nov. 2018, doi: 10.1109/TMI.2018.2837502.
- [5] S. Leclerc *et al.*, "Deep Learning for Segmentation Using an Open Large-Scale Dataset in 2D Echocardiography," *IEEE Trans. Med.*

Imaging, vol. 38, no. 9, pp. 2198–2210, Sep. 2019, doi: 10.1109/TMI.2019.2900516.

- [6] A. Radford, L. Metz, and S. Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. Accessed: Oct. 26, 2021. [Online]. Available: <http://arxiv.org/abs/1511.06434>
- [7] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least Squares Generative Adversarial Networks," in *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct. 2017, pp. 2813–2821. doi: 10.1109/ICCV.2017.304.
- [8] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein Generative Adversarial Networks," in *Proceedings of the 34th International Conference on Machine Learning*, Jul. 2017, pp. 214–223. Accessed: Oct. 26, 2021. [Online]. Available: <https://proceedings.mlr.press/v70/arjovsky17a.html>
- [9] J. H. Lim and J. C. Ye, "Geometric GAN," *ArXiv170502894 Cond-Mat Stat*, May 2017, Accessed: Dec. 18, 2020. [Online]. Available: <http://arxiv.org/abs/1705.02894>
- [10] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and Improving the Image Quality of StyleGAN," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020, pp. 8107–8116. doi: 10.1109/CVPR42600.2020.00813.
- [11] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic Image Synthesis With Spatially-Adaptive Normalization," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019, pp. 2332–2341. doi: 10.1109/CVPR.2019.00244.
- [12] X. Yi, E. Walla, and P. Babyn, "Generative adversarial network in medical imaging: A review," *Med. Image Anal.*, vol. 58, p. 101552, Dec. 2019, doi: 10.1016/j.media.2019.101552.
- [13] S. Amirrajab *et al.*, "XCAT-GAN for Synthesizing 3D Consistent Labeled Cardiac MR Images on Anatomically Variable XCAT Phantoms," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, vol. 12264, A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M. A. Zuluaga, S. K. Zhou, D. Racoceanu, and L. Joskowicz, Eds. Cham: Springer International Publishing, 2020, pp. 128–137. doi: 10.1007/978-3-030-59719-1_13.
- [14] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, "4D XCAT phantom for multimodality imaging research," *Med. Phys.*, vol. 37, no. 9, pp. 4902–4915, Sep. 2010, doi: 10.1118/1.3480985.
- [15] Y. Skandarani, N. Painchaud, P.-M. Jodoin, and A. Lalonde, "On the effectiveness of GAN generated cardiac MRIs for segmentation," *ArXiv200509026 Cs Eess*, May 2020, Accessed: Oct. 26, 2021. [Online]. Available: <http://arxiv.org/abs/2005.09026>
- [16] Z. Zhang, L. Yang, and Y. Zheng, "Translating and Segmenting Multimodal Medical Volumes with Cycle- and Shape-Consistency Generative Adversarial Network," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun. 2018, pp. 9242–9251. doi: 10.1109/CVPR.2018.00963.
- [17] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks," in *2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, Oct. 2017, pp. 2242–2251. doi: 10.1109/ICCV.2017.244.
- [18] J. Chen *et al.*, "Generative Invertible Networks (GIN): Pathophysiology-Interpretable Feature Mapping and Virtual Patient Generation," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Cham, 2018, pp. 537–545. doi: 10.1007/978-3-030-00928-1_61.
- [19] Y. Xia, N. Ravikumar, J. P. Greenwood, S. Neubauer, S. E. Petersen, and A. F. Frangi, "Super-Resolution of Cardiac MR Cine Imaging using Conditional GANs and Unsupervised Transfer Learning," *Med. Image Anal.*, vol. 71, p. 102037, Jul. 2021, doi: 10.1016/j.media.2021.102037.
- [20] G. Zhang, Y. Mao, M. Li, L. Peng, Y. Ling, and X. Zhou, "The Optimal Tetralogy of Fallot Repair Using Generative Adversarial Networks," *Front. Physiol.*, vol. 12, p. 613330, 2021, doi: 10.3389/fphys.2021.613330.
- [21] F. Zhu, F. Ye, Y. Fu, Q. Liu, and B. Shen, "Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network," *Sci. Rep.*, vol. 9, no. 1, p. 6734, May 2019, doi: 10.1038/s41598-019-42516-z.
- [22] N. Wulan, W. Wang, P. Sun, K. Wang, Y. Xia, and H. Zhang, "Generating electrocardiogram signals by deep learning," *Neurocomputing*, vol. 404, pp. 122–136, Sep. 2020, doi: 10.1016/j.neucom.2020.04.076.
- [23] P. Sarkar and A. Etemad, "CardioGAN: Attentive Generative Adversarial Network with Dual Discriminators for Synthesis of ECG from PPG," *ArXiv201000104 Cs Eess*, Dec. 2020, Accessed: Oct. 26, 2021. [Online]. Available: <http://arxiv.org/abs/2010.00104>
- [24] P. Singh and G. Pradhan, "A New ECG Denoising Framework Using Generative Adversarial Network," *IEEE/ACM Trans. Comput. Biol. Bioinform.*, vol. 18, no. 2, pp. 759–764, Mar. 2021, doi: 10.1109/TCBB.2020.2976981.
- [25] J. Georges-Filteau and E. Cirillo, "Synthetic Observational Health Data with GANs: from slow adoption to a boom in medical research and ultimately digital twins?," *ArXiv200513510 Cs Q-Bio Stat*, Nov. 2020, doi: 10.22541/au.158921777.79483839/v2.
- [26] M. K. Baowaly, C.-C. Lin, C.-L. Liu, and K.-T. Chen, "Synthesizing electronic health records using improved generative adversarial networks," *J. Am. Med. Inform. Assoc.*, vol. 26, no. 3, pp. 228–241, Mar. 2019, doi: 10.1093/jamia/ocy142.
- [27] Z. Che, Y. Cheng, S. Zhai, Z. Sun, and Y. Liu, "Boosting Deep Learning Risk Prediction with Generative Adversarial Networks for Electronic Health Records," in *2017 IEEE International Conference on Data Mining (ICDM)*, New Orleans, LA, Nov. 2017, pp. 787–792. doi: 10.1109/ICDM.2017.93.
- [28] Z. Zhang, C. Yan, D. A. Mesa, J. Sun, and B. A. Malin, "Ensuring electronic medical record simulation through better training, modeling, and evaluation," *J. Am. Med. Inform. Assoc. JAMIA*, vol. 27, no. 1, pp. 99–108, Jan. 2020, doi: 10.1093/jamia/ocz161.
- [29] L. Gong and Y. Zhou, "A Review: Generative Adversarial Networks," in *2019 14th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, Xi'an, China, Jun. 2019, pp. 505–510. doi: 10.1109/ICIEA.2019.8833686.
- [30] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017, pp. 5967–5976. doi: 10.1109/CVPR.2017.632.
- [31] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," *ArXiv171010196 Cs Stat*, Feb. 2018, Accessed: Oct. 26, 2021. [Online]. Available: <http://arxiv.org/abs/1710.10196>
- [32] H.-C. Shin *et al.*, "Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks," in *Simulation and Synthesis in Medical Imaging*, Cham, 2018, pp. 1–11. doi: 10.1007/978-3-030-00536-8_1.
- [33] F. Lau, T. Hendriks, J. Lieman-Sifry, S. Sall, and D. Golden, "ScarGAN: Chained Generative Adversarial Networks to Simulate Pathological Tissue on Cardiovascular MR Scans," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, vol. 11045, D. Stoyanov, Z. Taylor, G. Carneiro, T. Syeda-Mahmood, A. Martel, L. Maier-Hein, J. M. R. S. Tavares, A. Bradley, J. P. Papa, V. Belagiannis, J. C. Nascimento, Z. Lu, S. Conjeti, M. Moradi, H. Greenspan, and A. Madabhushi, Eds. Cham: Springer International Publishing, 2018, pp. 343–350. doi: 10.1007/978-3-030-00889-5_39.
- [34] G.-P. Diller *et al.*, "Utility of deep learning networks for the generation of artificial cardiac magnetic resonance images in congenital heart disease," *BMC Med. Imaging*, vol. 20, no. 1, p. 113, Oct. 2020, doi: 10.1186/s12880-020-00511-1.
- [35] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, Dec. 2016, pp. 2234–2242.
- [36] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Red Hook, NY,

USA, Dec. 2017, pp. 6629–6640.

- [37] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton, “Demystifying MMD GANs,” *ArXiv180101401 Cs Stat*, Mar. 2018, Accessed: Dec. 18, 2020. [Online]. Available: <http://arxiv.org/abs/1801.01401>