UltraGAN: Ultrasound Enhancement through Adversarial Generation

Maria Escobar[1⋆][0000−0002−5880−762X](✉), Angela Castillo[1⋆], Andrés Romero[2], and Pablo Arbeláez[1][0000−0001−5244−2407]

[1] Center for Research and Formation in Artificial Intelligence, Universidad de los Andes, Bogotá, Colombia
{mc.escobar11, a.castillo13, pa.arbelaez}@uniandes.edu.co
[2] Computer Vision Lab, ETHZ, Zürich, Switzerland
roandres@ethz.ch

Abstract. Ultrasound images are used for a wide variety of medical purposes because of their capacity to study moving structures in real time. However, the quality of ultrasound images is significantly affected by external factors limiting interpretability. We present UltraGAN, a novel method for ultrasound enhancement that transfers quality details while preserving structural information. UltraGAN incorporates frequency loss functions and an anatomical coherence constraint to perform quality enhancement. We show improvement in image quality without sacrificing anatomical consistency. We validate UltraGAN on a publicly available dataset for echocardiography segmentation and demonstrate that our quality-enhanced images are able to improve downstream tasks. To ensure reproducibility we provide our source code and training models.

Keywords: Generative Adversarial Networks · Echocardiography · Ultrasound images · Image quality enhancement.

1 Introduction

Echocardiography is one of the most commonly used tests for the diagnosis of cardiovascular diseases as it is low-cost and non-invasive [13]. The quality of echocardiogram images is highly related to three main factors: (i) intrinsic characteristics of the sonograph, (ii) sonograph manipulation [5], and (iii) the manual configurations of the operator. As the intrinsic configurations of each device are hard to control in practice, factors such as the expertise of the user and the contrast adjustment on the machine [23] may potentially affect the final diagnosis. Fundamentally, the criteria to assess the presence of some pathologies or the correct functionality of the left ventricle depends entirely on the quality of the resulting images on this test [11]. Therefore, the interpretability of ultrasound images is limited by their quality [4].

Once physicians acquire echocardiograms, they perform manual segmentation [21] to determine diagnostic measurements of the heart's chambers [13].

⋆ Both authors contributed equally to this work.

The quality of the resulting measurements depends directly on the precise outline of the chambers' boundaries [21]. Nevertheless, if the ultrasound image has a low quality, low resolution, or low contrast, the segmentation will be harder because of the lack of a clear difference between two adjacent structures. With an image that is not entirely intelligible, manual segmentation is purely done under the visual subjectivity of the physician to recognize the heart's boundaries.

Generative Adversarial Networks (GANs) [12] are a type of generative models that learn the statistical representation of the training data. During GAN training, the generator network alternates with the discriminator, so the generator can produce new data that resembles the original dataset. GANs have successfully tackled image-to-image transformation problems, including but not limited to: image colorization [8], super resolution [28], multi-domain and multimodal mappings [26], and image enhancement [7]. In the medical field, several works have introduced GANs into their approach for tasks that include data augmentation [3], medical images attacks [6] and image synthesis [29,2,1].

Recently, some automated methods have been developed for ultrasound image quality enhancement [9,25,15,16]. Liao *et al.* [20] proposed a method to enhance ultrasound images using a quality transfer network. The algorithm was tested in echo view classification, showing that quality transfer helps improve the performance. Also, Lartaud *et al.* [18] trained a convolutional network for data augmentation by changing the quality of the images to create contrast and non-contrast images for segmentation. The augmented data allowed the improvement of the segmentation method. Jafari *et al.* [17] trained a model to transform quality between ultrasound images. The approach introduced a segmentation network in the training of the GAN to provide an anatomical constraint added by the segmentation task. Nevertheless, these methods were developed and evaluated on private data, complicating the possibility of a direct comparison.

In this work, we present UltraGAN, a novel framework for ultrasound image enhancement through adversarial training. Our method receives as input a low-quality image and performs high quality enhancement without compromising the underlying anatomical structures of the input. Our main contributions can be summarized in the two following points: (1) We introduce specific frequency loss functions to maintain both coarse and fine-grained details of the original image. (2) We guide the generation with the anatomical coherence constraint by adding the segmentation map as input in the discriminator.

As our results demonstrate, UltraGAN transfers detailed information without modifying the important anatomical structures. Hence, our method could be useful for a better interpretability in clinical settings. Moreover, for the quantitative validation of our method, we compare the resulting segmentations with and without our enhancement, and we report an increase in performance in this downstream task by adding our enhanced images to the training data. We make our source code and results publicly available [3].

[3] https://github.com/BCV-Uniandes/UltraGAN

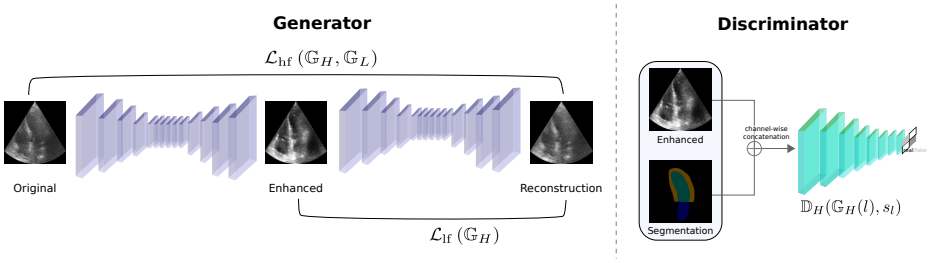

Fig. 1. Overview of our generation scheme. We add a frequency consistency loss to preserve fine details and coarse structures. We concatenate the segmentation map along with the input image for the discriminator to classify as real or enhanced. This particular case corresponds to an enhanced input.

2 Methodology

Our method consists of a Generative Adversarial Network designed to enhance the quality of ultrasounds without compromising the underlying anatomical information contained in the original image. The power of GANs relies on a minimax two-player game, where two different networks are trained in an adversarial fashion. The generator ($\mathbb{G}$) translates images from one domain to another, while the discriminator ($\mathbb{D}$) is in charge of determining whether the image is a real example from the dataset or a generated example coming from $\mathbb{G}$.

2.1 Problem Formulation

Given a set of low-quality ultrasounds $\{l_i\}_{i=1}^{N} \in L$ with a data distribution $l \sim p_{\text{data}}(l)$ and a set of high quality ultrasounds $\{h_i\}_{i=1}^{N} \in H$ with a data distribution $h \sim p_{\text{data}}(h)$, our goal is to learn mapping functions that translate from low to high quality domain and vice versa. Accordingly, we have two generators $\mathbb{G}_H : L \rightarrow H$ and $\mathbb{G}_L : H \rightarrow L$. We also have two discriminators: $\mathbb{D}_H$ distinguishes between real high quality images h_i and generated high quality images $\mathbb{G}_H(l_i)$, and $\mathbb{D}_L$ distinguishes between real low-quality images l_i and generated low-quality images $\mathbb{G}_L(h_i)$. Since we want our mapping functions to be aware of the original structural information, we also have the segmentation of the anatomical regions of interest present in the ultrasound image s_h or s_l.

It is important to note that at inference time, we only use the generator trained for the translation from low to high quality, even though we optimized for both generators in the training phase.

2.2 Model

We start from the generator architecture of CycleGAN [30] which consists of a series of down-sampling layers, residual blocks and up-sampling layers. Fig. 1

shows the training scheme for our model. For the discriminator, we build upon PatchGAN [30,14]. Since we want our model to learn how to produce anatomically coherent results, our discriminator has two inputs: the ultrasound image (whether real or generated) and the corresponding segmentation of the anatomical regions of interest.

2.3 Loss functions

To enforce the task of quality translation during training, we use an identity loss and we alter the traditional adversarial and cycle consistency losses to create an anatomically coherent adversarial loss and frequency cycle consistency losses.

Anatomically Coherent Adversarial Loss. The goal of the adversarial loss is to match the generated images to the corresponding real distribution. Inspired by the idea of conditional GANs [22] and pix2pix [14], we modify the adversarial loss to include as input the segmentation of the anatomical regions of interest. For the networks $\mathbb{G}_H$ and $\mathbb{D}_H$ our anatomically coherent adversarial loss is defined as:

$$\begin{aligned}\mathcal{L}_{\text{adv}}\left(\mathbb{G}_H, \mathbb{D}_H\right) = &\mathbb{E}_{h\sim p_{\text{data}}(h)}\left[\log \mathbb{D}_H(h, s_h)\right] \\ &+ \mathbb{E}_{l\sim p_{\text{data}}(l)}\left[\log\left(1 - \mathbb{D}_H(\mathbb{G}_H(l), s_l)\right)\right]\end{aligned} \tag{1}$$

By adding the segmentation as an input to the discriminator, we make sure that the networks learn the underlying relationship between the anatomical regions of interest and the structures in the generated image. Furthermore, the segmentation is not necessary at test time, since we only use the generator. In the final objective function, we also consider the adversarial loss for $\mathbb{G}_L$ and $\mathbb{D}_L$.

Frequency Cycle Consistency. The cycle consistency loss [30] ensures a direct one-to-one mapping from an image of one domain to another. However, the cycle consistency constraint is a pixel-wise L1 norm between the original image and the reconstruction, which enforces the output to have similar intensities. Yet, during the process of quality enhancement, it is more useful to think of the image in terms of frequency rather than intensity [10]. Low frequencies contain the structural information of an image, while high frequencies contain the fine details. With this concept in mind, we create two types of frequency consistency losses enforcing our training scheme to improve quality enhancement.

During quality translation, we aim to preserve the structural information present in the low frequencies of the original image. To extract low frequencies, we pass the images through a Gaussian pyramid [24] ϕ at K scales, then compute the L1 norm between the structural information of the original and the generated image (Eq. 2). We also want our generators to transfer image details of the corresponding quality in the form of high frequencies. Therefore, we obtain those frequencies through a Laplacian pyramid [24] γ at K scales and calculate the L1 norm between the high frequencies of the original image and the high frequencies of the reconstruction (Eq. 3). The loss concept is better illustrated in Fig. 1.

$$\mathcal{L}_{\text{lf}}\left(\mathbb{G}_H\right) = \sum_{k=1}^{K} \left\|\phi_k(l) - \phi_k\left(\mathbb{G}_H(l)\right)\right\|_1 \tag{2}$$

$$\mathcal{L}_{\text{hf}}\left(\mathbb{G}_H, \mathbb{G}_L\right) = \sum_{k=1}^{K} \left\|\gamma_k(l) - \gamma_k\left(\mathbb{G}_L(\mathbb{G}_H(l))\right)\right\|_1 \tag{3}$$

Identity Loss. The identity loss introduces a new constraint ensuring that the generator does not modify images from the same domain. This loss is particularly useful for quality enhancement in real clinical applications. In practice, we do not have a quality label but still we would like to transform all images to high quality without modifying the image if it already has a high quality.

$$\mathcal{L}_{\text{idt}} = \left\|h - \mathbb{G}_H(h)\right\|_1 \tag{4}$$

Overall Loss. For simplicity, we show only one of the pathways in the loss formulations, but our overall loss is defined as the weighted sum of the losses in both pathways $H \rightarrow L$ and $L \rightarrow H$, where each λ represents the relative importance of each loss function in the system: $\mathcal{L}_{\text{UltraGAN}} = \lambda_{\text{adv}}\mathcal{L}_{\text{adv}} + \lambda_{\text{lf}}\mathcal{L}_{\text{lf}} + \lambda_{\text{hf}}\mathcal{L}_{\text{hf}} + \lambda_{\text{idt}}\mathcal{L}_{\text{idt}}$.

3 Experiments

3.1 Dataset

To validate our method, we use the publicly available "Cardiac Acquisitions for Multi-structure Ultrasound Segmentation" (CAMUS) dataset [19]. The CAMUS dataset contains 2D Ultrasound images and multi-structure segmentations of 500 patients. This dataset is particularly relevant because it includes ultrasound images from three different qualities labeled by expert physicians. Besides, CAMUS includes patients with different left ventricle ejection fraction, making it a realistic problem with healthy and pathological subjects. The task in the CAMUS dataset is to segment the left ventricular endocardium (LV_{Endo}), left ventricular epicardium (LV_{Epi}) and left atrium (LA) in two chamber (2CH) and four chamber (4CH) views for End of Diastole (ED) and End of Systole (ES).

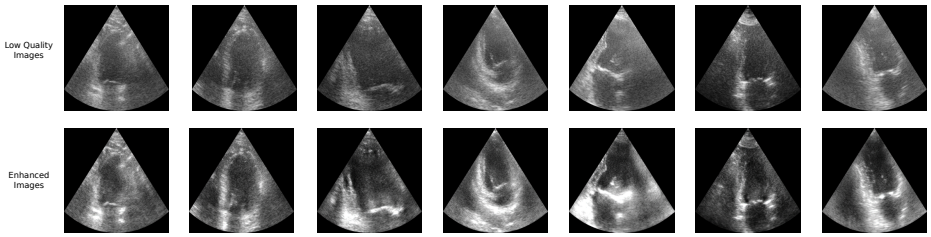

Fig. 2. Qualitative comparison of the low-quality and enhanced images using UltraGAN. Our method is able to enhance ultrasound images, improving the interpretability of the heart structures regardless of the view.

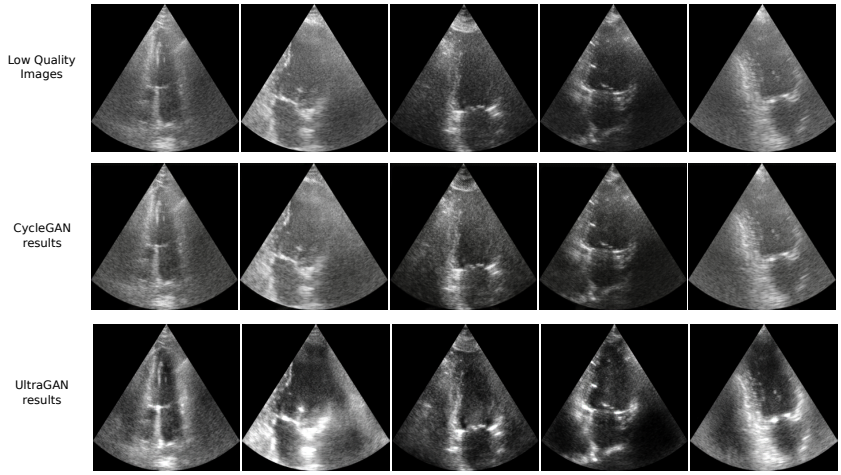

Fig. 3. Qualitative comparison between CycleGAN result and UltraGAN. The images generated by CycleGAN are perceptually similar to the original Low quality images. In contrast, images enhanced by UltraGAN show a clear difference between anatomical structures.

3.2 Experimental setup

We first assess the importance of each of the components of our method in the task of quality enhancement. For simplicity, we divide the ultrasound images into high and low quality, classifying the medium quality images as low quality during our enhancement process. We train our method with 80% of the images and evaluate on the remaining images. We enhance the training data with three variants of our system: without the anatomically coherent adversarial loss, without the frequency cycle consistency losses, and with the original CycleGAN losses. Nevertheless, the evaluation of image quality in an unpaired setup is a subjective process and performing perceptual studies would require an extensive effort by expert physicians. However, we take advantage of the publicly available segmentation masks to provide multi-structure segmentation as a down-stream quantitative metric, in which the right global anatomical structure is required for a good performance.

As a baseline, we first train a simple U-Net model [27] using the standard 10 fold cross-validation split of the CAMUS dataset. Then, we use UltraGAN to enhance the quality of all the training images and train the same U-Net with the original images as well as the enhanced augmentation.

4 Results

4.1 Image Enhancement

The enhancement in image quality provided by UltraGAN is noticeable even for untrained eyes. Fig. 2 shows the comparison between the low-quality images and

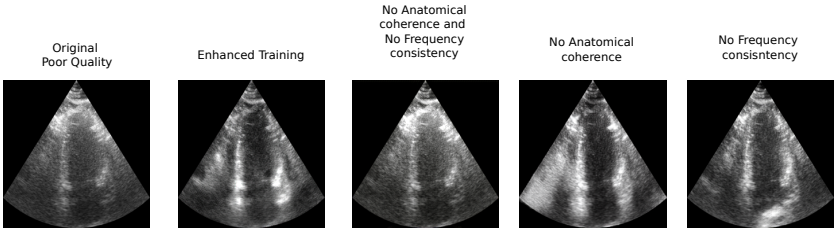

Fig. 4. Ablation examples of our enhancement method. We show the results obtained for every stage of the generation.

the enhanced images we produce. In the enhanced images, the heart's chambers are recognizable and their boundaries are easy to identify. The examples illustrate the preservation of anatomical consistency in the enhancement process for both 2CH and 4CH views.

Furthermore, in Figs. 4 and 3 we demonstrate that the training choices for our method improve over the baseline, in which we just consider the standard loss function. We compare the original image against our enhancement showing that there is a better definition of the structure of the heart with more defined walls. We also compare the difference between having just the anatomical coherence or the frequency consistency. The images enhanced using merely frequency consistency maintain finer details, yet the system tends to hallucinate high frequencies in the left part of the image. Conversely, considering just the anatomical coherence, the structure is preserved but there is not a well definition of heart regions. Overall, with UltraGAN we are able to create an image quality enhancement that takes into account frequency and structural information.

4.2 Multi-structure segmentation

In Fig. 5 we show that the segmentations obtained by using the standard data have artifacts, while training with UltraGAN-enhanced images improves the resulting segmentation. Also, for quantitative results, Table 1 shows the Dice Scores for this experiment. Here we find evidence of the previous results, showing that for each of the structures present in the ultrasound image, augmenting

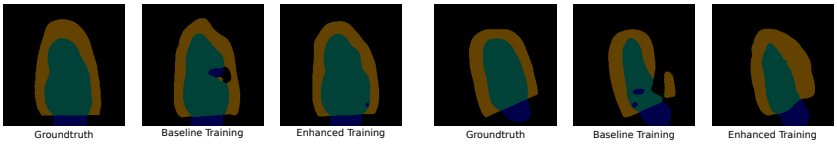

Fig. 5. Qualitative results for heart segmentation in the CAMUS dataset by using our enhanced images as data augmentation in the training stage. We present two different test examples showing the groundtruth (columns 1 and 4), the baseline results (columns 2 and 5) and the improved segmentation (columns 3 and 6).

Table 1. Segmentation results for 10-fold cross-validation set comparing standard training vs training with quality enhancement.

Method	High (%)			Medium (%)			Low (%)		
	LV_{Endo}	LV_{Epi}	LA	LV_{Endo}	LV_{Epi}	LA	LV_{Endo}	LV_{Epi}	LA
Baseline	93.07	86.61	88.99	92.02	85.32	88.13	90.76	83.10	87.52
Our method	**93.78**	**87.38**	**89.48**	**92.66**	**86.20**	**88.38**	**91.55**	**83.75**	**87.84**

Table 2. Segmentation results for 10-fold cross-validation comparing the state-of-the-art vs our quality enhanced training.

Image quality	Method	ED (%)		ES (%)	
		LV_{Endo}	LV_{Epi}	LV_{Endo}	LV_{Epi}
High	Ours	**94.40±0.7**	86.54±1.2	**92.04±1.1**	87.05±1.4
+ Medium	Leclerc *et al.*	93.90±4.3	**95.40 ±2.3**	91.60±6.1	**94.50±3.9**
Low	Ours	**93.00±1.1**	83.57±1.9	**90.10±1.3**	83.93±2.7
	Leclerc *et al.*	92.10±3.7	**94.70±2.3**	89.80±5.7	**93.67±3.2**

the training data with UltraGAN improves the segmentation results. This improvement is also consistent across all of the image qualities, suggesting that the baseline with enhanced training data preserves correctly the anatomical structures present in the ultrasound images. We also evaluate separately the segmentation of our enhanced images in a subset of the CAMUS dataset consisting of patients at pathological risk with a left ventricle ejection fraction lower than 45%. We find that for pathological cases, the average Dice score (89.5) is as good as for healthy patients (89.7). Thanks to the global consistency enforced by the other heart structures, UltraGAN is able to extract accurately atypical left ventricles.

Table 2 shows the comparison between the state-of-the-art method in the CAMUS dataset and our quality enhanced method for the High+Medium and Low qualities in the 10 fold cross-validation sets. We do not include the comparison for Left atrium segmentation since the authors do not report their performance on that class. [19] uses a modified U-Net network to achieve the results. Here we demonstrate that, even with a simpler network with less amount of parameters, by enhancing the quality of the training images we are able to outperform state-of-the-art approaches in left ventricular endocardium segmentation, and obtain competitive results in left ventricular epicardium segmentation. Thus, demonstrating that the inclusion of quality enhanced images during training can benefit a model's generalization.

5 Conclusions

We present UltraGAN, a Generative Adversarial Network designed for quality enhancement of ultrasound images. We achieve image enhancement of 2D

echocardiography images without compromising the anatomical structures. By using multi-structure segmentation as a downstream task we demonstrate that augmenting the training data with enhanced images improves the segmentation results. We expect UltraGAN to be useful in other ultrasound problems to push forward automated ultrasound analysis.

Acknowledgements

The present study is funded by MinCiencias, contract number 853-2019 project ID 120484267276.

References

1. Abdi, A.H., Jafari, M.H., Fels, S., Tsang, T., Abolmaesumi, P.: A study into echocardiography view conversion. In: Workshop of Medical Imaging Meets NeurIPS (2019)
2. Abdi, A.H., Tsang, T., Abolmaesumi, P.: Gan-enhanced conditional echocardiogram generation. In: Workshop of Medical Imaging Meets NeurIPS (2019)
3. Abhishek, K., Hamarneh, G.: Mask2lesion: Mask-constrained adversarial skin lesion image synthesis. In: Burgos, N., Gooya, A., Svoboda, D. (eds.) Simulation and Synthesis in Medical Imaging. pp. 71–80. Springer International Publishing, Cham (2019)
4. Armstrong, A.C., Ricketts, E.P., Cox, C., Adler, P., Arynchyn, A., Liu, K., Stengel, E., Sidney, S., Lewis, C.E., Schreiner, P.J., et al.: Quality control and reproducibility in m-mode, two-dimensional, and speckle tracking echocardiography acquisition and analysis: the cardia study, year 25 examination experience. Echocardiography **32**(8), 1233–1240 (2015)
5. Cahalan, M.K., Stewart, W., Pearlman, A., Goldman, M., Sears-Rogan, P., Abel, M., Russell, I., Shanewise, J., Troianos, C., et al.: American society of echocardiography and society of cardiovascular anesthesiologists task force guidelines for training in perioperative echocardiography. Journal of the American Society of Echocardiography **15**(6), 647–652 (2002)
6. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Intelligent image synthesis to attack a segmentation cnn using adversarial learning. In: Burgos, N., Gooya, A., Svoboda, D. (eds.) Simulation and Synthesis in Medical Imaging. pp. 90–99. Springer International Publishing, Cham (2019)
7. Chen, Y.S., Wang, Y.C., Kao, M.H., Chuang, Y.Y.: Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6306–6314 (2018)
8. Deshpande, A., Lu, J., Yeh, M.C., Jin Chong, M., Forsyth, D.: Learning diverse image colorization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6837–6845 (2017)
9. Duarte-Salazar, C.A., Castro-Ospina, A.E., Becerra, M.A., Delgado-Trejos, E.: Speckle noise reduction in ultrasound images for improving the metrological evaluation of biomedical applications: An overview. IEEE Access **8**, 15983–15999 (2020)
10. Fritsche, M., Gu, S., Timofte, R.: Frequency separation for real-world super-resolution. ICCV Workshop (2019)

11. Gardin, J.M., Adams, D.B., Douglas, P.S., Feigenbaum, H., Forst, D.H., Fraser, A.G., Grayburn, P.A., Katz, A.S., Keller, A.M., Kerber, R.E., et al.: Recommendations for a standardized report for adult transthoracic echocardiography: a report from the american society of echocardiography's nomenclature and standards committee and task force for a standardized echocardiography report. Journal of the American Society of Echocardiography **15**(3), 275–290 (2002)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)
13. Gottdiener, J.S., Bednarz, J., Devereux, R., Gardin, J., Klein, A., Manning, W.J., Morehead, A., Kitzman, D., Oh, J.K., Quinones, M., et al.: American society of echocardiography recommendations for use of echocardiography in clinical trials: a report from the american society of echocardiography's guidelines and standards committee and the task force on echocardiography in clinical trials. Journal of the American Society of Echocardiography **17**(10), 1086–1119 (2004)
14. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on (2017)
15. Jafari, M.H., Girgis, H., Abdi, A.H., Liao, Z., Pesteie, M., Rohling, R., Gin, K., Tsang, T., Abolmaesumi, P.: Semi-supervised learning for cardiac left ventricle segmentation using conditional deep generative models as prior. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 649–652. IEEE (2019)
16. Jafari, M.H., Girgis, H., Van Woudenberg, N., Moulson, N., Luong, C., Fung, A., Balthazaar, S., Jue, J., Tsang, M., Nair, P., et al.: Cardiac point-of-care to cart-based ultrasound translation using constrained cyclegan. International Journal of Computer Assisted Radiology and Surgery pp. 1–10 (2020)
17. Jafari, M.H., Liao, Z., Girgis, H., Pesteie, M., Rohling, R., Gin, K., Tsang, T., Abolmaesumi, P.: Echocardiography segmentation by quality translation using anatomically constrained cyclegan. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. pp. 655–663. Springer International Publishing, Cham (2019)
18. Lartaud, P.J., Rouchaud, A., Rouet, J.M., Nempont, O., Boussel, L.: Spectral ct based training dataset generation and augmentation for conventional ct vascular segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 768–775. Springer (2019)
19. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., et al.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. IEEE transactions on medical imaging **38**(9), 2198–2210 (2019)
20. Liao, Z., Jafari, M.H., Girgis, H., Gin, K., Rohling, R., Abolmaesumi, P., Tsang, T.: Echocardiography view classification using quality transfer star generative adversarial networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 687–695. Springer (2019)
21. Mazaheri, S., Sulaiman, P.S.B., Wirza, R., Khalid, F., Kadiman, S., Dimon, M.Z., Tayebi, R.M.: Echocardiography image segmentation: A survey. In: 2013 International Conference on Advanced Computer Science Applications and Technologies. pp. 327–332. IEEE (2013)
22. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)

23. Ng, A., Swanevelder, J.: Resolution in ultrasound imaging. Continuing Education in Anaesthesia Critical Care & Pain **11**(5), 186–192 (2011)
24. Oliva, A., Torralba, A., Schyns, P.G.: Hybrid images. ACM Transactions on Graphics (TOG) **25**(3), 527–532 (2006)
25. Ortiz, S.H.C., Chiu, T., Fox, M.D.: Ultrasound image enhancement: A review. Biomedical Signal Processing and Control **7**(5), 419–428 (2012)
26. Romero, A., Arbeláez, P., Van Gool, L., Timofte, R.: Smit: Stochastic multi-label image-to-image translation. In Proceedings of the IEEE International Conference on Computer Vision Workshops (2019)
27. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
28. Wang, X., Chan, K.C., Yu, K., Dong, C., Loy, C.C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2019)
29. Yang, H., Sun, J., Carass, A., Zhao, C., Lee, J., Xu, Z., Prince, J.: Unpaired brain mr-to-ct synthesis using a structure-constrained cyclegan. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. pp. 174–182. Springer International Publishing, Cham (2018)
30. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networkss. In: Computer Vision (ICCV), 2017 IEEE International Conference on (2017)