



Semi-supervised Virtual Regression of Aortic Dissections Using 3D Generative Inpainting

Antonio Pepe^{1,2,3} , Gabriel Mistelbauer^{2,4} , Christina Gsaxner^{1,3,5} ,
Jianning Li^{1,3} , Dominik Fleischmann² , Dieter Schmalstieg¹ ,
and Jan Egger^{1,3,5} 

¹ Institute of Computer Graphics and Vision, Graz University of Technology,
Graz, Austria

{antonio.pepe,egger}@tugraz.at

² School of Medicine, Department of Radiology, Stanford University,
Stanford, CA, USA

a pepe@stanford.edu

³ Computer Algorithms in Medicine Laboratory, Graz, Austria

⁴ Department of Simulation and Graphics, University of Magdeburg,
Magdeburg, Germany

⁵ Department of Oral and Maxillofacial Surgery, Medical University of Graz,
Graz, Austria

Abstract. Aortic dissection (AD) is a condition of the main artery of the human body, resulting in the formation of a new flow channel, or false lumen (FL). The disease is usually diagnosed with a computed tomography angiography (CTA) scan during the acute phase. A better understanding of the causes of AD requires knowledge of aortic geometry prior to the event, which is available only in very rare circumstances. In this work, we propose an approach to reconstruct the aorta before the formation of a dissection by performing 3D inpainting with a two-stage generative adversarial network (GAN). In the first stage of our two-stage GAN, a network is trained on the 3D edge information of the healthy aorta to reconstruct the aortic wall. The second stage infers the image information of the aorta to reconstruct the entire dataset. We train our two-stage GAN with 3D patches from 55 non-dissected aortic datasets and evaluate it on 20 more non-dissected datasets, demonstrating that our proposed 3D architecture outperforms its 2D counterpart. To obtain pre-dissection aortae, we mask the entire FL in AD datasets. Finally, we provide qualitative feedback from a renowned expert on the obtained pre-dissection cases.

Keywords: Inpainting · Aortic dissection · Deep learning · Edge reconstruction · Generative adversarial networks

1 Introduction

The aortic dissection (AD) is an uncommon, but life-threatening condition of the aorta (Fig. 1), caused by a diseased medial layer inside the aortic wall [9]. Characteristic is the sudden development of a so-called *primary entry tear*, a defect on the luminal side of the aorta [11]. This allows blood to enter the teared aortic wall, further dissect it, and form a new ‘false’ blood flow channel inside the aortic wall. As the dissected aortic wall is thinner and of lower mechanical strength, it may rupture, leading to a typically lethal event [19]. When the newly formed blood channel reconnects with the original ‘true’ blood flow channel of the aorta, two separate flow channels are generated, referred to as true lumen (TL) and false lumen (FL), as the interior of a blood vessel is called *lumen* [11].

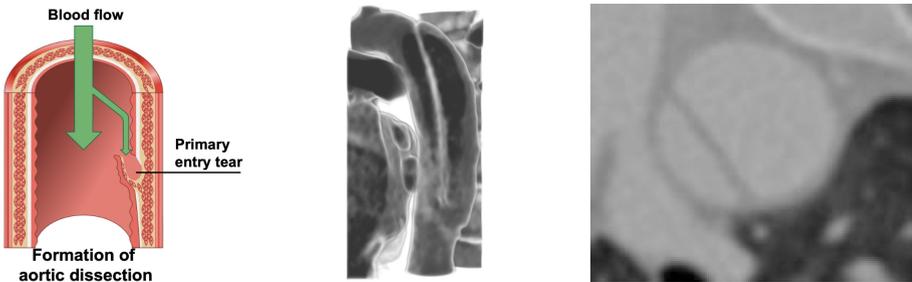


Fig. 1. Illustration of an AD. Left, in color: The formation of a FL with a primary entry tear. It is clearly visible that the FL grows within the medial layer of the dissected aortic wall, i.e. the central layer in yellow. In gray: actual CTA images of AD showing a rendered sagittal view (center) and a cropped axial view (right). (Color figure online)

Diagnosing ADs is typically done by acquiring CTA that provides detailed anatomic information of therapeutically relevant features, such as the anatomy of TL and FL [5]. A more in-depth discussion is provided in two recent reviews on the biomechanics of the aortic wall [22] and on the adopted techniques in medical imaging [19]. Once AD has been diagnosed, patients undergo life-long surveillance, comprising CTA scans after three months, six months and annually thereafter to monitor the progression of the aortic growth [7, 15]. Investigating the onset of the disease remains rather unexplored, since acquiring data before the diagnosis of AD is highly unlikely [19]. Such data would only exist from an earlier examination that required a thorax CTA. Having such data would help simulate and understand the causes of ADs as well as visually convey the progression over time [15]. The aortic diameter is currently used for risk assessment, as this increases with the growth of the FL [15, 19].

In this work, we address this shortage using image processing via generative adversarial networks (GANs) [6], which have found applications in different research areas, such as image artifact reduction [24]. By using a free-form

3D inpainting approach we are able to generate plausible pre-dissection CTA datasets. We consider the FL of an AD as an image ‘artifact’ to be masked for removal.

1.1 Background

A large body of recent research in image processing employs GAN, including free-form inpainting problems [17, 26]. A GAN consists of two distinguished deep networks, a generator and a discriminator [6]. The two networks are trained in a zero-sum game, where the generator learns to generate realistic fake data, the discriminator, to distinguish it from real data.

Due to the availability of large databases like ImageNet [1], GAN inpainting has been widely and intensively tested on 2D images with rectangular masks. In the field of computer vision, Liu et al. [12] claim to be the first to address the problem of free-form inpainting with deep learning, using partial convolution and a U-Net as generator. Yu et al. [26] use GAN with gated convolution to deal with free-form masks. They also introduce the idea of adding manual sketches as guidelines to reinforce particular shapes in the inpainted images. Nazeri et al. [17] automate the creation of these guidelines using adversarial edge learning. They suggest a two-stage GAN with 2D inputs, which initially hallucinates the edges inside the masked area, successively it uses this information as a guideline and performs the actual inpainting task. Free-form inpainting approaches have been less intensively investigated in the medical imaging community. Mirsky et al. [14] qualitatively evaluated the effectiveness of CT inpainting using 3D GANs. The authors used a rectangular ROI to add or remove a tumor in thoracic CT images. In their evaluation, they show how the method can *fool* expert radiologists in different cases. Armanious et al. [2] evaluated the role of perceptual losses, as considered also by Nazeri et al. [17], in CT and MRI inpainting.

1.2 Contribution

Using a semi-supervised approach, we apply a free-form mask to exclude the region occupied by the FL and train a two-stage 3D GAN to first hallucinate the edges of the masked volume and then, secondly, perform an edge-guided inpainting. Our contributions are:

- a two-stage, semi-supervised 3D GAN network for the reconstruction of corrupted or missing content in CTA images,
- the virtual regression of AD via prediction of the pre-dissection aortic shape,
- the evaluation of the generated data by a clinical expert.

2 Data Acquisition and Preprocessing

Since before-AD/with-AD image pairs are unavailable, we use public datasets of healthy aortae for training. A total of 75 CTA datasets were acquired, 40 from

the CAD-PE challenge (www.cad-pe.org) and 35 from Masoudi et al. [13]. Both data sources contain images of pulmonary embolism, which usually does not affect the original shape of the aorta. For a qualitative evaluation of the virtual AD regression, we use 52 non-public datasets with Type-B AD. Type-B indicates that the dissection only involves the descending aorta, after the arch [4].

To contain the computational cost, the datasets were manually cropped to a variable size of approximate $100 \times 100 \times D$ voxels, with the number of slices, D , depending on the length of the aorta. The extent of the smallest cropped volume is $100 \times 100 \times 133$, whereas the extent of the largest is $104 \times 104 \times 473$. All images contain the descending aorta and its immediate surroundings.

3 Method

Aortic dissection (AD) can extend along the whole aorta. For this reason, inpainting techniques can easily fail when considering single-view, 2D approaches. Extending the idea from Nazeri et al. [17], we define a 3D inpainting method based on a two-stage GAN. Each stage consists of a 3D-GAN with spectral normalization [16]. Each layer is initialized using the Xavier method [23]. A scheme of the overall architecture is shown in Fig. 2. The network input is obtained from a 3D slab, \mathbf{V} , of size $I = 100 \times 100 \times 12$; this can be interpreted as a 3D sliding window, \mathbf{V} , which slides along the descending aorta. The sliding window with overlap allows the information to be transmitted to the following slabs.

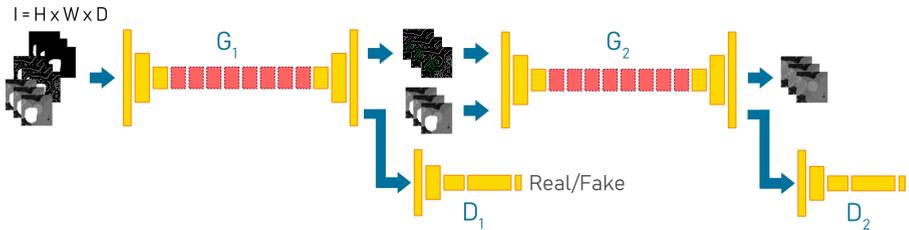


Fig. 2. Architecture of the two-stage 3D network. Each input volume has size I . In yellow, the convolution and deconvolution layers. In red, the residual blocks. The layers of the discriminators follow the Patch-GAN architecture [17]: each layer consists of a 3D convolution (Kernel size: 3, stride: 2, channels, in order: $\{64, 128, 256, 512, 1\}$), spectral normalization and LeakyRELU activation with threshold 0.2. For the generators, the encoding layers consist of 3d convolution (cubic kernel size, in order: $\{7, 4, 4\}$, stride: $\{1, 2, 2\}$, channels: $\{64, 128, 256\}$), spectral normalization, instance normalization and ReLU activations. The residual blocks are similar to the encoding layers with 256 channels and dilation factor 2. The decoding layers mirror the encoding layers. (Color figure online)

Given the 3D slab \mathbf{V} , its edges $\mathbf{C} = \text{edge}(\mathbf{V})$ and a mask \mathbf{M} ; we define the following inputs and output for the first generator G_1 :

$$\mathbf{C}_M = \mathbf{C} \odot (1 - \mathbf{M}), \quad (1)$$

$$\mathbf{V}_M = \mathbf{V} \odot (1 - \mathbf{M}), \quad (2)$$

$$\mathbf{C}_G = G_1(\mathbf{V}_M, \mathbf{C}_M, \mathbf{M}), \quad (3)$$

where \odot is the Hadamard product; \mathbf{C}_M , the masked edge volume; \mathbf{V}_M , the masked original volume, and \mathbf{C}_G , the output of the first generator, which contains the reconstructed edges. We provide both the edge volume and the original volume as we empirically found that this reduces the training time. The output volume \mathbf{C}_G is then used to evaluate the adversarial loss, \mathcal{L}_{adv1} , and the feature matching loss, \mathcal{L}_{FM} , used in the objective function. In particular, having defined with \mathbb{E} the expectation operator,

$$\mathcal{L}_{adv1} = \mathbb{E}_{(\mathbf{C}, \mathbf{V})} [\log D_1(\mathbf{C}, \mathbf{V})] + \mathbb{E}_{(\mathbf{V})} \log [1 - D_1(\mathbf{C}_G, \mathbf{V})] \quad (4)$$

defines a metric of the average distance between the two distributions $(\mathbf{C}_G, \mathbf{V})$ and (\mathbf{C}, \mathbf{V}) .

$$\mathcal{L}_{FM} = \mathbb{E} \left[\sum_{i=1}^L \frac{1}{N_i} \left\| D_1^{(i)}(\mathbf{C}) - D_1^{(i)}(\mathbf{C}_G) \right\|_1 \right] \quad (5)$$

defines a metric in the feature space; where L is the number of convolution layers in the discriminators, $D_1^{(i)}$ is the activation of layer i . Similar volumes should therefore have similar activation maps. This information is then combined in the adversarial objective function:

$$\min_{G_1} \max_{D_1} \mathcal{L}_{G_1} = \min_{G_1} \left[\lambda_{adv1} \max_{D_1} (\mathcal{L}_{adv1}) + \lambda_{FM} \mathcal{L}_{FM} \right], \quad (6)$$

where λ_{adv1} and λ_{FM} are regularization parameters. We extend this evaluation also to the second generator, G_2 :

$$\mathbf{C}_I = \mathbf{C}_M + \mathbf{C}_G \odot \mathbf{M}, \quad (7)$$

$$\mathbf{V}_G = G_2(\mathbf{V}_M, \mathbf{C}_I), \quad (8)$$

$$\mathbf{V}_I = \mathbf{V}_M + \mathbf{V}_G \odot \mathbf{M}, \quad (9)$$

where \mathbf{C}_I represents the inpainting of \mathbf{C}_M ; \mathbf{V}_G , the output of the second generator, and \mathbf{V}_I , the final inpainting of \mathbf{V}_M .

Ad-hoc loss functions have been defined for 2D image outputs. These include the style loss and the perceptual loss [10], which provide a metric based on the activation maps of a pre-trained VGG-19 network. Practical experiments revealed that the distances in the feature space are not significantly affected by

the dataset used to train the network. Nonetheless, VGG-19 is trained on 2D images; therefore we introduced an averaged perceptual loss and style loss,

$$\mathcal{L}_{perc,av} = \mathbb{E} \left[\sum_{i,j} \frac{1}{N_i} \|\phi_i(\mathbf{I}_j) - \phi_i(\mathbf{I}_{G_j})\|_1 \right], \quad (10)$$

$$\mathcal{L}_{style,av} = \mathbb{E}_{k,j} \left[\|G_k^\phi(\mathbf{I}_j) - G_k^\phi(\mathbf{I}_{G_j})\|_1 \right], \quad (11)$$

where we compute the loss values on a per-slice level j and average them over the volume. This considers the fact that medical experts evaluate the volume by analyzing single 2D axial views. Each ϕ_i is the activation map of the layer i in the VGG-19 network; G_k^ϕ is a Gram matrix derived from the activation map ϕ_k [10]. Thus, we define a composite objective function:

$$\min_{G_2} \max_{D_2} \mathcal{L}_{G_2} = \min_{G_2} \max_{D_2} [\lambda_{adv2} \mathcal{L}_{adv2} + \lambda_{l_1} \mathcal{L}_{l_1} + \lambda_p \mathcal{L}_{perc,av} + \lambda_s \mathcal{L}_{style,av}] \quad (12)$$

where λ_{adv2} , λ_{l_1} , λ_p , and λ_s are regularization parameters; \mathcal{L}_{adv2} is the adversarial loss (Eq. 4), and \mathcal{L}_{l_1} is the L1 loss.

Previously, 2D networks were trained with the masks provided by Liu et al. [12] or with rectangular masks [12, 17]. In this work, we generate random 3D masks by filling each of the 12 slices with a mask from Liu et al., a random super-ellipse (SE), or a segmentation mask. During training, the masks and the patch volumes are selected randomly. During testing, the patch volumes are consecutively extracted with a sliding window, together with the relative masks, with a pace of four slices.

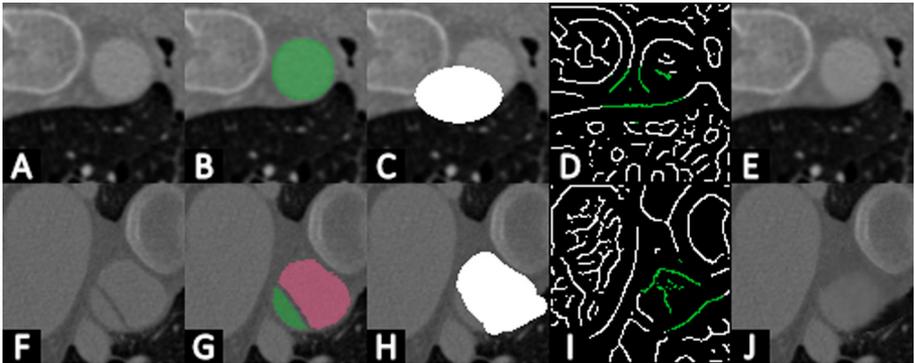


Fig. 3. Example of two inpainted slices. A) A healthy aorta, F) A dissected aorta, B) G) Segmentation of true lumen (green) and false lumen (red), C) H) masked image, D) I) hallucinated edges (green), E) J) inpainting result on the aorta and surroundings. (Color figure online)

Table 1. Quantitative results of the inpainting, compared to the previous 2D version [17]. Precision (Prec.) and Accuracy (Acc.) refer to the edge reconstruction.

Input	Mask	Prec	Acc	\mathcal{L}_{l_1}	$\mathcal{L}_{perc,av}$	$\mathcal{L}_{style,av}$	HDD	PSNR	MAE
2D	SE	32.26	28.64	0.349	0.048	0.065	0.610	30.465	0.017
2D	SE/FF	10.32	10.34	0.192	0.099	0.102	1.248	17.049	0.186
3D	SE	58.89	55.44	0.261	0.045	0.020	0.082	45.898	0.003
3D	SE/FF	88.67	86.53	0.201	0.049	0.027	0.168	41.81	0.005

4 Experiments and Results

To evaluate our approach, we train the network on cropped regions of interest (ROI) from 55 non-dissected aortae, which we augment through reflections, as the aorta can be situated on both sides of the spine. We evaluate its performance over a separate pool of 20 images of non-dissected aortae, to be able to quantify the results. Moreover, we test our approach with after-AD images and ask medical experts to evaluate the results.

The training was performed using PyTorch v1.2 on a desktop computer (CPU: Intel i7-8700, 64 GB RAM, GPU: NVIDIA Titan RTX 24 GB) and with a batch size of 16. The learning rate was set to 10^{-4} , λ_{adv1} and λ_{adv2} to 0.1, λ_{FM} to 10, λ_s to 250, λ_{l_1} to 1 and λ_p to 0.1.

Initially, we only train the first stage, where the volume edges are reconstructed. First, we train with stacks of random SE to provide the training with more degrees of freedom than the usual rectangles. We continue training by mixing the SE slices with the free-form masks from Liu et al. [12] and further segmentation masks of medical images. We refer to this combination as FF.

Table 1 summarizes the measured precision and accuracy for this stage. The 3D information enhances the performances considerably, especially, for FF. Probably due to size and shape variety of masks, the 3D approach outperformed the 2D approach. Convolution in 3D can use information from adjacent slices, resulting in a network capable of inpainting in larger planar areas.

We repeat the same experiments for the second stage. Here, precision and recall may not be reliable measures, as small differences are acceptable. Therefore, we evaluate the performance of the inpainting stage using the loss metrics, together with the peak signal-to-noise ratio (PSNR), the mean absolute error (MAE) and the Hausdorff distance (HDD).

The measures in Table 1 show a trend comparable with the first stage. In both the experiments, the 3D approach provides results which are visibly more accurate. Nonetheless, a detailed comparison between the two trends shows that there is still space for improvements in the second stage. Figure 3 provides a visual overview of the results.

From a qualitative point of view, we further saw that the edge reconstruction step guarantees more natural edges for the aorta reconstruction, compared to the second-stage network trained without the edge information. The images of

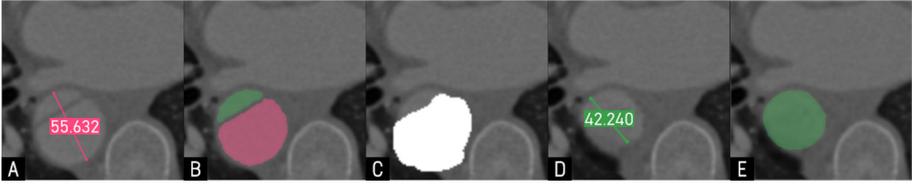


Fig. 4. Example of size change during AD regression. Values in millimeter. A) Original image. The aorta presents a higher diameter due to the presence of the FL, B) Segmentation of TL (green) and FL (red), C) Masked image, D) Reconstructed pre-dissection state. The diameter is now lower. E) Segmentation of the pre-dissection aorta. (Color figure online)

the resulting pre-dissection case were reviewed by a cardiovascular radiologist with more than 20 years of experience in the interpretation of CT angiograms of patients with aortic diseases. The expert was asked to evaluate if the virtual pre-dissection aorta reasonably resembled the expected appearance of non-diseased aorta. Specifically, the expert was asked to assess if the following key features of a dissection were suppressed in the virtual pre-dissection dataset: (a) Presence of a dissection flap, (b) difference in contrast enhancement between TL and FL, (c) shape of the outer contour of the aorta, and (d) introduced artifacts or unexpected findings.

- (a) The dissection flap – the diagnostic hallmark of aortic dissection in CT imaging – was successfully eliminated in almost all images.
- (b) The contrast opacification of the simulated pre-dissection aorta appeared homogeneous, with elimination of the typical attenuation difference between TL versus FL blood. There is a contrast heterogeneity without an abrupt transition caused by a dissection flap. The contrast heterogeneity exhibits commonly observed mixing artifacts.
- (c) The shape of the virtual pre-dissection aorta was considered remarkably similar to the expected appearance of a normal aorta, with circular or slightly oval cross section on transverse CT images, with coronal and sagittal reformats. This is in contradiction to the typically more eccentric shape of a dissected aorta, where the outer wall of the FL stretches and dilates due to thinning and loss of elastic tissue.
- (d) Even though, subjectively, the appearance of the virtual regression was remarkably similar to a non-dissected aorta, the interface between the inserted aorta and its surrounding, typically a few millimeters outside of the aortic contour, was noticeable (see Fig. 4). However, these artifacts did not distract from the evaluation of the important features of the aorta.

5 Conclusion

We have evaluated the potential of 3D image inpainting for visually regressing images of aortic dissection. In particular, we define a two-stage 3D-GAN for

medical image inpainting and train it with patch volumes containing healthy aortae and free-form masks obtained with different approaches. We mask only the dissected traits of the aorta and use the GAN to reconstruct their original healthy state, before the development of AD. We obtain quantitative results by inpainting a separate image pool of healthy aortae, and compare them with a 2D approach. Moreover, we perform a qualitative evaluation of the inpainting operation on AD cases under the guidance of a senior radiologist specialized in cardiovascular diseases. To the best of our knowledge, this is a unique application of GAN and the first evaluation of AD regression leveraging the strengths of neural networks. The results show how the suggested 3D inpainting approach performs considerably better than the state-of-the-art 2D counterpart. Although expected, this underlines the necessity to process the aorta as a whole or, at least, as a union of overlapping segments. Furthermore, the network is trained with randomly located masks, and provides enough generalization to reconstruct also other objects near the aorta, such as the spine. The 2D version has been currently integrated in the online open-science platform www.studierfenster.at [25] and a video tutorial is available on YouTube¹ [21]. Future work sees the extension to the whole aorta as well as the segmentation of these pairs [20] and their use in biomechanical simulations [3]. Additionally, we will evaluate long-short term memory approaches to remove the need for an overlap between slabs [8].

Acknowledgement. This work received fundings from the TU Graz LEAD Project *Mechanics, Modeling and Simulation of Aortic Dissection* (<https://www.tugraz.at/projekte/biomechaorta/BioMechAorta>), the Austrian Science Fund (FWF) KLI 678-B31 *enFaced* and the COMET K-Project 871132 *CAMed* of the Austrian Research Promotion Agency (FFG). Antonio Pepe was also supported by an Austrian Marshall Plan Foundation Scholarship (Scholarship n. 94212102222019) [18]. Additionally, the authors would like to thank Professor Dr. Gerhard A. Holzapfel (Graz University of Technology) for his advice.

References

1. ImageNet. www.image-net.org. Accessed 10 July 2020 07:10:23
2. Armanious, K., Mecky, Y., Gatidis, S., Yang, B.: Adversarial inpainting of medical imaging modalities. In: Proceedings of International Conference on Acoustics, Speech and Signal Processing (2019). <https://doi.org/10.1109/ICASSP.2019.8682677>
3. Bäumlner, K., Vedula, V., Sailer, A.M., et al.: Fluid-structure interaction simulations of patient-specific aortic dissection. *Biomech. Model. Mechanobiol.* (2020). <https://doi.org/10.1007/s10237-020-01294-8>
4. Daily, P.O., Trueblood, H.W., Stinson, E.B., et al.: Management of acute aortic dissections. *Ann. Thoracic Surg.* **10**(3), 237–247 (1970). [https://doi.org/10.1016/S0003-4975\(10\)65594-4](https://doi.org/10.1016/S0003-4975(10)65594-4)

¹ Deep Learning Inpainting of Aortic Dissections with Studierfenster: <https://www.youtube.com/watch?v=c85qV-CDOX4>.

5. Egger, J., Gunacker, S., Pepe, A., et al.: A comprehensive workflow and framework for immersive virtual endoscopy of dissected aortae from CTA data. In: SPIE Medical Imaging 1131531 (2020). <https://doi.org/10.1117/12.2559239>
6. Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al.: Generative adversarial nets. In: Proceedings of Advances in Neural Information Processing Systems, vol. 27, pp. 2672–2680 (2014). <https://doi.org/10.5555/2969033.2969125>
7. Hahn, L.D., Mistelbauer, G., Higashigaito, K., et al.: Ct-based true- and false-lumen segmentation in type b aortic dissection using machine learning. *Radiol. Cardiothor. Imaging* **2**(3), 1–10 (2020)
8. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
9. Howard, D.P., Banerjee, A., Fairhead, J.F., et al.: Population-based study of incidence and outcome of acute aortic dissection and premorbid risk factor control: 10-year results from the oxford vascular study. *Circulation* **127**(20), 2031–2037 (2013). <https://doi.org/10.1161/circulationaha.112.000483>
10. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
11. LeMaire, S.A., Russell, L.: Epidemiology of thoracic aortic dissection. *Nat. Rev. Cardiol. Nat. Publ. Group* **8**(2), 103–113 (2011). <https://doi.org/10.1038/nrcardio.2010.187>
12. Liu, G., Reda, F.A., Shih, K.J., Wang, T.-C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11215, pp. 89–105. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01252-6_6
13. Masoudi, M., Pourreza, H.R., Saadatmand-Tarzjan, M., et al.: A new dataset of computed-tomography angiography images for computer-aided detection of pulmonary embolism. *Sci. Data* **5**, 180180 (2018). <https://doi.org/10.1038/sdata.2018.180>
14. Mirsky, Y., Mahler, T., Shelef, I., Elovici, Y.: CT-GAN: Malicious tampering of 3D medical imagery using deep learning. In: Proceedings of the 28th USENIX Security Symposium (2019). <https://doi.org/10.5555/3361338.3361371>
15. Mistelbauer, G., Schmidt, J., Sailer, A.M., et al.: Aortic dissection maps: comprehensive visualization of aortic dissections for risk assessment. In: Proceedings of Eurographics Workshop on Visual Computing for Biology and Medicine, pp. 143–152 (2016). <https://doi.org/10.2312/vcbm.20161282>
16. Miyato, T., Kataoka, T., Koyama, M., et al.: Spectral normalization for generative adversarial networks. In: Proceedings of International Conference on Learning Representations (2018)
17. Nazeri, K., Ng, E., Joseph, T., et al.: Edgeconnect: generative image inpainting with adversarial edge learning. In: Proceedings of International Conference on Computer Vision Workshops (2019)
18. Pepe, A., Fleischmann, D., Schmalstieg, D., Egger, J.: Visual computing of dissected aortae. In: Technical Report for the Austrian Marshall Plan Foundation, pp. 1–32 (2020, to appear)
19. Pepe, A., Li, J., Rolf-Pissarczyk, M., et al.: Detection, segmentation, simulation and visualization of aortic dissections: a review. *Med. Image Anal.* (2020). <https://doi.org/10.1016/j.media.2020.101773>

20. Pepe, A., Schussnig, R., Li, J., et al.: Iris: interactive real-time feedback image segmentation with deep learning. In: *Medical Imaging 2020: Biomedical Applications in Molecular, Structural, and Functional Imaging*, vol. 11317 (2020). <https://doi.org/10.1117/12.2551354>
21. Prutsch, A., Pepe, A., Egger, J.: Design and development of a web-based tool for inpainting of dissected aortae in angiography images. In: *Proceedings of Central European Seminar on Computer Graphics*, pp. 1–8 (2020)
22. Sherifova, S., Holzapfel, G.A.: Biomechanics of aortic wall failure with a focus on dissection and aneurysm: a review. *Acta Biomaterialia* **99**, 1–17 (2019). <https://doi.org/10.1016/j.actbio.2019.08.017>
23. Sun, W., Su, F., Wang, L.: Improving deep neural networks with multi-layer maxout networks and a novel initialization method. *Neurocomputing* **278**, 34–40 (2018). <https://doi.org/10.1016/j.neucom.2017.05.103>
24. Wang, J., Zhao, Y., Noble, J.H., Dawant, B.M.: Conditional generative adversarial networks for metal artifact reduction in CT Images of the Ear. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11070, pp. 3–11. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_1
25. Wild, D., Weber, M., Egger, J.: Client/server based online environment for manual segmentation of medical images. In: *41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3463–3467. IEEE (2019). <https://doi.org/10.1109/EMBC.2019.8856481>
26. Yu, J., Lin, Z., Yang, J., et al.: Free-form image inpainting with gated convolution. In: *Proceedings of International Conference on Computer Vision* (2019). <https://doi.org/10.1109/ICCV.2019.00457>