CycleGAN for style transfer in X-ray angiography

Oleksandra Tmenova [1,2] · **Rémi Martin** [1] ·
Luc Duong [1]

Abstract

Purpose We aim to perform generation of angiograms for various vascular structures as a mean of data augmentation in learning tasks. The task is to enhance the realism of vessels images generated from an anatomically-realistic cardiorespiratory simulator to make them look like real angiographies.
Methods The enhancement is performed by applying the CycleGAN deep network for transferring the style of real angiograms acquired during percutaneous interventions into a data set composed of realistically simulated arteries.
Results The cycle-consistency was evaluated by comparing an input simulated image with the one obtained after two cycles of image translation. An average structural similarity (SSIM) of 0.948 on our data sets has been obtained. The vessel preservation was measured by comparing segmentations of an input image and its corresponding enhanced image using Dice coefficient.
Conclusions We proposed an application of the CycleGAN deep network for enhancing the artificial data as an alternative to classical data augmentation techniques for medical applications, particularly focused on angiogram generation. We discussed success and failure cases, explaining conditions for the realistic data augmentation which respects both the complex physiology of arteries and the various patterns and textures generated by X-ray angiography.

1 Introduction

Deep learning methods have quickly gained acceptance in the field of medical image analysis. However, in order to provide satisfying results, machine learning

✉ Oleksandra Tmenova, oleksandra.tmenova@gmail.com

[1] École de technologie supérieure. Department of Software and IT Engineering, 1100 Notre-Dame W., Montréal, Canada

[2] Taras Shevchenko National University of Kyiv, Volodymyrska St, 60, Kyiv, Ukraine

techniques require a large amount of significant data. In the medical community, it is difficult to obtain medical data for a number of reasons. First of all, the large resource needed to collect and manage a database is a major hurdle. Secondly, obtaining expert annotations of imaging data for supervised and semi-supervised learning is a rather expensive and time-consuming process. Thirdly, legal considerations regulate the sharing of medical data, which varies from country to country and are subject to ethical committees and patient's approval.

For natural images research, data augmentation is widely used to increase the number of samples. However, standard data augmentation (mainly data warping) should be used with caution in a medical context. Indeed, in the case of angiographies, vessels must not be unrealistically wide or narrow, as structural and anatomical information could be used for further analysis and diagnosis of a patient [7]. This is achieved by regulating the maximum amounts of translation, rotation and non-rigid deformation applied to obtain plausible physical variations of the medical images [11]. However, this often comes down as a lengthy and costly process that requires a medical expert's intervention.

In this paper, we are proposing a new approach to address the problem of an insufficient amount of training data in medical image analysis applications, such as vessel segmentation, detection, and classification from X-ray angiography. Such method has already been proven successful for the enhancement of realism of synthetic eye images [13], where the authors show an improvement of learning accuracy. We propose a solution of the given problem by applying CycleGAN [15] to perform a style transfer from real images obtained using X-ray based techniques to artificial images generated by the simulator [12] (Figure 1). The proposed approach has the potential to become a more affordable and convenient way of generating more training data while conserving the anatomical integrity of arteries and realistic X-ray patterns. Furthermore, because the generation process is entirely controlled, the annotations of those images can also be easily generated.

The contributions of this work are the following:

1. Proposition of an anatomy-based data-augmentation technique for X-ray angiography.
2. Evaluation of the CycleGAN method for images containing low contrast and a high amount of noise.
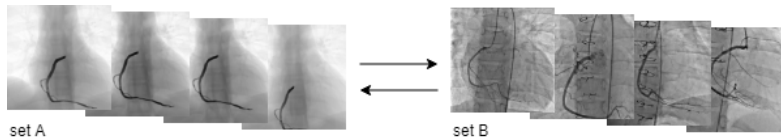

Fig. 1 Image translation between two domains: simulated data (set A) and real angiograms (set B).

2 Related work

2.1 Standard data augmentation

Deep learning approaches require a large amount of data to train a model that not only provides good results on the test at hand but also avoids over-fitting. However, many domains, such as medical imaging, suffer from a shortage of data. A common approach is to take a smaller but labeled data set and augment it in a way that improves the effectiveness of models [9]. Amongst them, image augmentation [4, 5, 8, 14] and data augmentation guided by expert knowledge [14] are some of the most popular techniques. The most generic approach for image augmentation is data warping. It consists of performing geometric transformations such as translation, reflection, and cropping, as well as changing its color scheme.

In [4] elastic deformations are the main source of variation in mammography at a lesion level. In [5, 28] data warping is used to obtain more data samples of fundus images depicting retinal blood vessels for further segmentation. Same is applied in [8] for augmenting datasets of brain tumors. Tumors do not have a definite shape and therefore may exhibit large spatial and structural variability, which is not the case for other objects, such as heart vessels which have definite unchangeable structure.

In [14] geometric and color augmentation is combined with specialist knowledge about skin lesion symmetry (anti-symmetry). Lesion segmentation is approximated to ellipse whose axes are then used for controlling varied distortions of the original lesion sample. The segmentation mask is used so that symmetry is preserved and further classification of lesions could be done.

2.2 GAN-based methods for data augmentation

Another approach is to synthesize new data using generative adversarial networks (GANs) [2]. The idea is to use two competing neural networks: a generator and a discriminator (or adversary). A generator is generating a sample learning from the distribution of training data. A discriminator is evaluating how similar is the output of the generator to true data samples, by calculating the adversarial loss and acting like a two-class classifier: generated image either belongs to a target distribution or it does not. For medical image synthesis, GANs have become a more generic solution [45] as they allow creating more variety of data (appearance of a specific pathology, different imaging protocols) comparing to traditional data augmentation. GAN-based methods include unconditional and conditional medical image synthesis.

The first group of methods implies that image is generated from noise without any additional information. Liver lesions [16], lung nodules [1], and brain MR images [17] generation are some of many successful cases. In [29], unconditional GANs are used for obtaining synthetic fundus patches to obtain more data for segmentation task.

The second group (conditional GANs) are methods where additional information is provided to generate images with desired properties. An image is the most common condition used for such supervision, which was first offered in "pix2pix" framework [3]. The task of mapping is done using paired examples for training

when every image from a source domain A has a corresponding pair from a target domain B. During training, a generator takes an input image, *realA* and tries to map it to target domain B, measuring success by comparing a generated image, *fakeB*, to ground truth, *realB*, which serves as a strict condition. Successful applications of conditional GANs for vessels include those for synthesizing retinal fundus images, all using vessel map as conditional information for generator [18–21]. In [22] a task of generating synthetic X-rays of the lumbar spine is tackled using conversion from sagittal to coronal radiographic projections of the trunk.

2.3 Data augmentation for vessels

Mentioned applications of both standard and GAN-based approaches for data augmentation of vessels images come down to tasks done on publicly available datasets: retinal fundus imaging (datasets DRIVE [23], STARE on which most of the vessels computer vision tasks are trained and tested so far); lung vessels CT imaging (VESSEL12 challenge [24] for vessel segmentation); coronary arteries (CAVAREV [25], CoronARe challenge [26] for 3-D reconstruction, this dataset has projections of thorax with left coronary tree only; IVUS dataset [27] of ultrasound images for vessel segmentation). In recent work [30], cerebral vasculatures data is synthesized and coupled with real clinical data to achieve a state-of-the-art performance of a deep learning model. To our best knowledge, there are no solutions for tackling the problem of insufficient data when it comes to computation of biophysical measurements in pediatric interventional cardiology where a little amount of data is available. With this work, we, therefore, want to start closing this gap and initiate further research into generating cardiac angiography data.

2.4 Enhancing realism of angiograms

While standard data augmentation techniques imply producing more data by warping given samples to achieve variety, our idea is to first simulate varied cardiac angiograms and then to enhance them for obtaining realistic images. This puts our approach closer to GAN-based methods which allow generating data with additional conditions. For our work, we wish to generate realistic images from simulated images, where the vessel's geometry is properly preserved, while also introducing radiographic noise and angiography-related artifacts. An input image is artificial data and an output image is realistic-looking data.

This modification can be interpreted as an image-to-image translation task and is a core of many cross-modality image synthesis applications of GANs. Some of them require paired data, like in the "pix2pix" framework [3]: PET (positron emission tomography) to MR (magnetic resonance) translation of brain images [31], T1- and T2-weighted MRI to MRA (magnetic resonance angiography) translation of brain images [32], CT (computed tomography) to PET translation of thorax [33] and liver images [34]. Paired data is obtained by co-registering different image modalities.

We want to avoid expensive and time-consuming labeling. Frameworks like CycleGAN [15] and similar methods [35, 36] relax this labeling constrain. CycleGAN does not use any labeled examples and instead of building image-to-image

1 mapping, creates a translation between domains. In order to maintain a meaning-
2 ful mapping between images from source and target domains, a second GAN is
3 introduced which maps an output of the first generator back to the original input.
4 The intuition behind this is similar to language translation: to check whether a
5 French sentence has been correctly translated to Spanish we can translate it back
6 to French and compare with the initial sentence to evaluate the "quality" of the
7 translation. In CycleGAN model (Figure 2), this role is performed by an addi-
8 tional component in the loss function called cycle-consistency-loss, which ensures
9 forward and backward cycle consistencies:

$$L_{cyc}(G, D_A, B, A) = E_{a \sim p_{data}(a)}[|F(G(a)) - a|_1] + E_{b \sim p_{data}(b)}[|G(F(b)) - b|_1].$$

10 Other two components are adversarial losses for two generators. Thus, the
11 objective of CycleGAN includes three losses: adversarial loss for the first generator,
12 adversarial loss for the second generator and cycle-consistency loss.

$$L(G, F, D_A, D_B) = L_{GAN_1} + L_{GAN_2} + \lambda_{cyc}(G, D_A, B, A),$$

13 where λ controls the relative importance of the objectives.

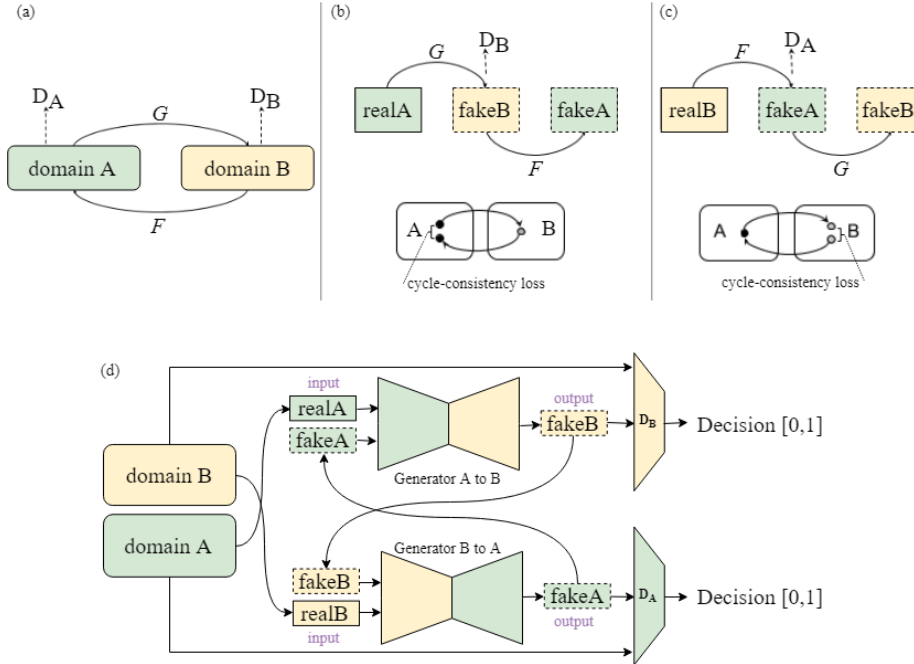**Fig. 2** (a) CycleGAN model contains two mapping functions $G : A \to B$ and $F : B \to A$, and associated adversarial discriminators D_B and D_A. D_B encourages G to translate A into outputs indistinguishable from domain B, and vice versa for F. (b) forward cycle-consistency loss: for each image from the domain A: $realA \to fakeB = G(realA) \to fakeA = F(G(realA)) \approx realA$;(c) backward cycle-consistency loss: for each image from domain B: $realB \to fakeA = F(realB) \to fakeB = G(F(realB)) \approx realB$. (d) CycleGAN architecture which includes two generator and two discriminator neural networks.

CycleGAN framework, within the last years, became already widely used for various cross-modality translations tasks, such as: MR to CT translation of myocardium [37], cardio-vascular volumes [38], musculoskeletal [39], brain [40] images; T1 to T2 translation of brain images [41–43]; Real to Synthetic translation of endoscopy images for monocular depth estimation [44].

Theoretical ability to do cross-domain translations in an unsupervised fashion, empirical proofs of accomplishing this task for various medical image applications encouraged us to refer to CycleGAN model. We, therefore, have an objective of expanding the quality of simulated cardiac angiography data by enhancing its realism using CycleGAN framework.

3 Methodology

3.1 Data description

X-ray angiography sequences - or angiograms, are two-dimensional X-ray moving sequences. They are used in percutaneous interventions for the treatment of vascular disease. 2D X-ray angiography allows for the visualization of the arteries in real-time during the intervention [7] and the guidance of catheters.

Data acquisition. The angiographies used in our proposed work were acquired from different patients with congenital heart disease from the Sainte-Justine's Hospital (Montreal, Canada). These data were collected offline and anonymized after approval by the Sainte-Justine's Institutional Ethics Review Board. The angiograms were acquired by a C-arm Infinix-CFI BP (Toshiba Medical). Samples are given in Figure 3. The angiograms comprise different vascular structures such as pulmonary, aorta, right coronary artery (RCA) and left coronary artery (LCA).
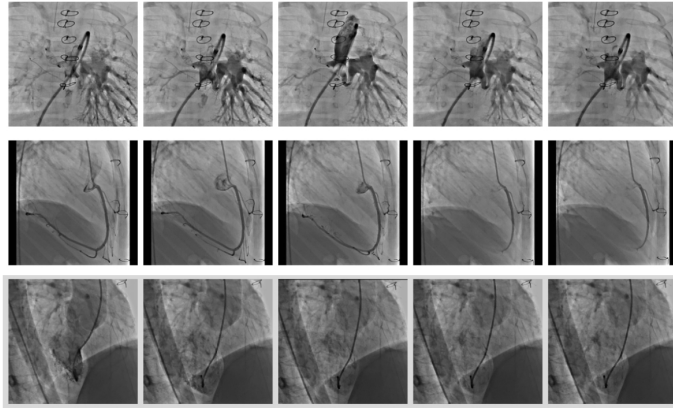

Fig. 3 Examples of angiograms (5 frames of each sequence only).

Grouping the data. To prove the proposed approach works, we needed to experiment first with smaller data sets, separating X-ray scans of different vessels. Thus, we grouped all angiography sequences into two categories: RCA and LCA. Frames of different sequences of each category can be seen in Figure 4.
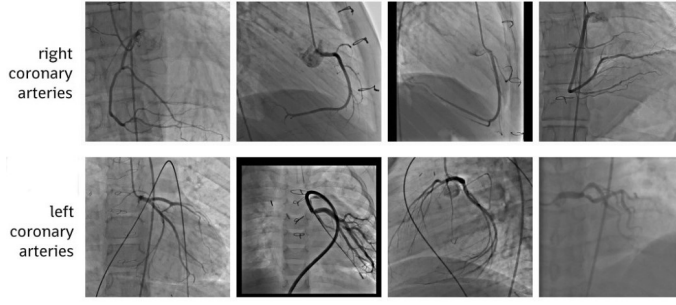
Fig. 4 Result of grouping angiographies into two different categories (4 randomly selected frames of each category only): right coronary arteries, left coronary arteries.

Analyzing the data. X-ray angiography is known for different imaging artifacts. Some sequences depict sternal sutures which means the patient had previous operations. While this can alter the result of segmentation, for our goal it is a realistic-looking-data style we want to achieve. We can see guiding wires and catheter as these are tools used during catheterization to introduce the contrast agent. Visibility of vessels depends on the diffusion of the contrast agent in the blood flow. Therefore, the vessels are not always entirely visible on the angiograms [7]. This results in many *empty* frames containing only background information which are not useful for our applications. Hence, we removed these frames from sequences.

3.2 Simulating artificial data

The data domain set is composed of artificial data which we want to modify in order to make it look like the target domain, i.e. real angiograms.

XCAT simulator. Artificial images were obtained using *XCAT 2.0* motion simulator [12]. It is a CT simulation tool based on the 4D extended cardiac-torso phantom, a whole-body computer model of the human anatomy and physiology based on NURBS surfaces. XCAT provides an accurate representation of the complex human anatomy and has the advantage, that the shape of its organs can be altered to realistically model anatomical variations and patient motion [12]. Those biological constraints make this tool particularly adapted to our process of obtaining more realistic data. Using XCAT, we obtained projection images of different vessel types. They already can be directly used as training data e.g. for segmentation and reconstruction tasks [6]. However, despite realistic anatomy, the textures of projections are still not realistic enough and a learning model trained solely on simulated data might fail when applied to real data during validation or testing, given all the artifacts present in real data which the simulator does not create.

Simulator parameters. The main advantage of XCAT is its ability to simulate anatomically diverse patient models as we choose the gender of a patient, indicate whether a patient is healthy or has any (e.g. heart) lesions, change the motion mode (respiratory, cardiac motion, or both), etc. For our proposed work, we used only one set of parameters (default), as the calculation of CT projections

is time-consuming due to mathematical complexities behind the simulator [12]. Apart from the simulator parameters, we can modify the projection parameters which affect the acquisition process, e.g. the angle of projection. Even though the shape of the vessel is preserved when rotating, the final images will serve as completely different samples for training any model. Examples of simulated images are presented in Figure 5.
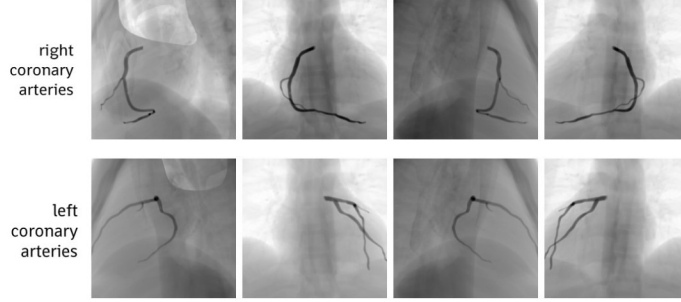

Fig. 5 Examples of simulated images of two vessel types: first frames of right coronary arteries and left coronary arteries sequences, with different angiographic views (lateral, cranial,caudal,left anterior oblique, right anterior oblique).

3.3 Data sets and training details

The CycleGAN framework requires two data sets: set A and set B. In our case, set A, or source domain, is the artificial data from the simulator and set B, or target domain, is real, clinical, data.

We formed small data sets by grouping together sequences of real data comprising up to three angiographic view projections, and sequences of artificial data from similar views, e.g. 0-40 degree angulation for RCA and 0-20 degree angulation for LCA. As we can see in Figures 4 and 5, real and simulated images have many differences in textures. Additionally, real data contain specific structures such as catheters, wires, balloons, and stents unseen in simulations. We limited our data sets by choosing RCA and LCA with fewer objects in images. We first grouped together training data: $trainA$, $trainB$, and then selected a few images from each of the sets as testing data: $testA$ and $testB$ accordingly, and removed them from training data. RCA data set: the training set size was 130 for simulated, 154 for real data. LCA data set: 138 for simulated, 16 for real.

We used the implementation of CycleGAN by X. Hu [1]. All the networks were trained from scratch with a learning rate of 0.0002. The weight for the cycle consistency loss λ was 20 for RCA and 30 for LCA (default is 10). Each model was trained for 200 epochs. After every 10 epochs, we launched the testing phase, i.e. translating of testing images into real style and back, and saved the results.

---

[1] Xiaowei Hu, "Tensorflow implementation for learning an image-to-image translation without input-output pairs", https://github.com/XHUJOY/CycleGAN-tensorflow, 2017

1 **4 Results and discussion**

2 4.1 Visual examination

3 Based on the convergence of losses, we chose the optimal number of epochs for
4 training each model: for RCA it was 50, for LCA it was 160 epochs. Each model
5 was still trained for 200 epochs (Figure 6) for further analysis of how this and
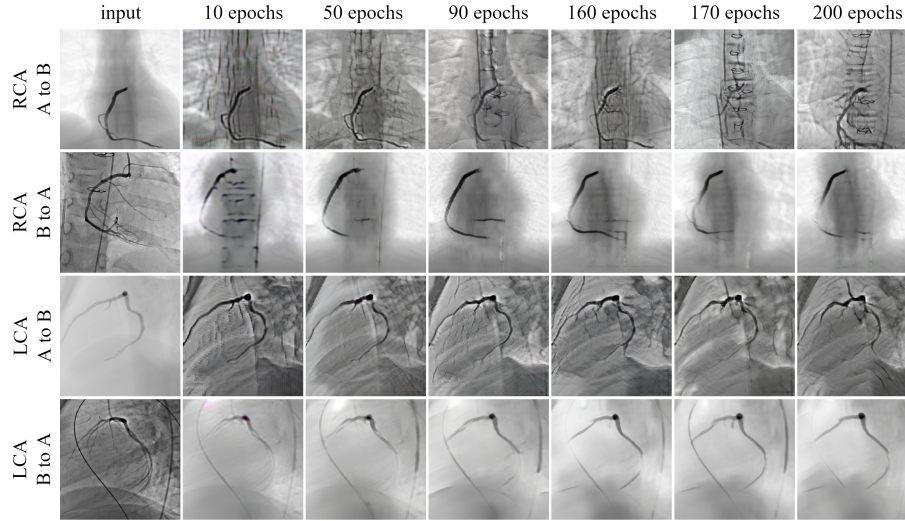6 other hyper-parameters (e.g. λ coefficient) are influencing training.

Fig. 6 Results of tests after different number of epochs.

7 Firstly, we observe that the longer the training is, the better is the quality of
8 output images, i.e. they are more smooth and don't have chalkboard artifacts like
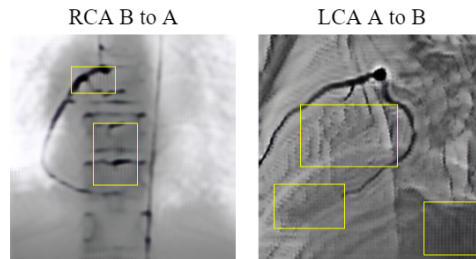9 outputs on early stages of training (Figure 7).

Fig. 7 Chalkboard artifacts that are especially visible on images produced by models trained
for a small amount of epochs.

10 Furthermore, we see that over the time of training the RCA model starts
11 deforming vessel shapes, while LCA model displays only slight distortion. This is
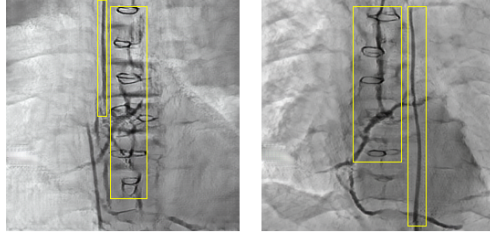
Fig. 8 Clinical objects that are visible in the later RCA results: models trained for 150 and 170 epochs. Highlighted are wires and catheters generated by CycleGAN trained on RCA dataset.

due to λ which is s a key hyperparameter for preserving the structure of vessels. A small λ would fail to prevent big distortions, while a too large value would suppress fine details in the resulting image. Additionally, we see more various clinical objects, such as catheters and wires (Figure 8), appearing in RCA results (even if a vessel shape itself may get deformed). This is possible because the number of available training images from domain B (real angiographies) is much larger for RCA data set than the one for LCA (154 against 16). Therefore, the more variety in training data, the more variety is possible in resulting images (given, domains are still corresponding).

We observe that CycleGAN model is able to transfer radiographic noise and texture of real X-ray angiography images onto simulated ones, often together with additional clinical objects (if they are present in training data). Quality of translation depends on the amount of training data available, the number of epochs and λ weight of cycle-consistency loss.

4.2 Quantitative evaluation

Quantitative evaluation of generative models is often a challenging task, especially when experimental setting doesn't imply having ground truth images for testing and validation. So, for analyzing the performance of our model beyond the visual judgments, we utilized two ideas: cycle-consistency and vessel segmentation (Figure 9).

During its training, CycleGAN model is using cycle-consistency loss as the difference between $real A$ and $fake A$ which it aims to minimize. To evaluate quantitatively and show which level of cycle consistency our model reached, we measured the similarity between input images and output images after one cycle of translations, as shown in Figure 9. We used NRMSE (normalized root mean squared error), SSIM (structure similarity), and PSNR (peak signal-to-noise ratio) metrics, which are the most common metrics for measuring image similarity. Results are given in Table 1. According to all the metrics, RCA performs better at A to B translation than LCA, having lower error (average of 0.029), higher structural similarity (average of 0.960) and higher PSNR value (31.303). LCA, on the other hand, has the highest error (0.27) while for other models this metric is almost equal to 0.1 (and even less, 0.03, for RCA).

Second criteria we used for quantitative evaluation of the obtained model is segmentation of input and output images to measure how well vessel shapes are
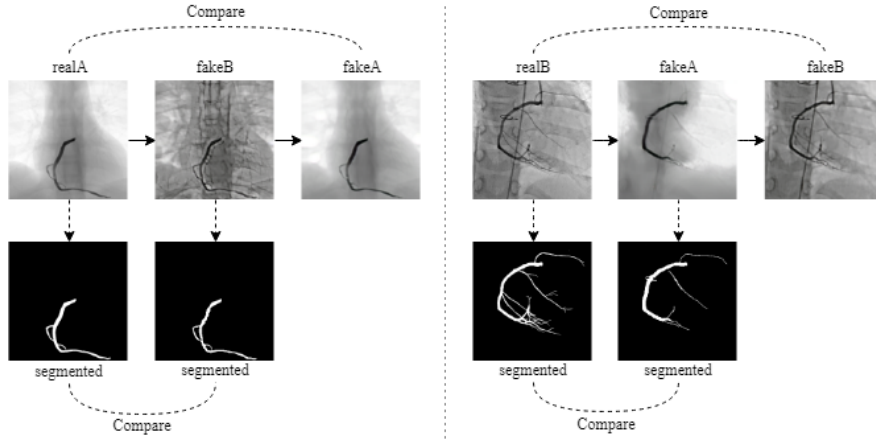
Fig. 9 Visual representation of our evaluation approach. (a) in translation $A \rightarrow B \rightarrow A$ we compare *simulated* (*realA*) with *fake_simulated* (*fakeA*) images for measuring cycle consistency, and compare segmented *simulated* (*realA*) with segmented *enhanced* (*fakeB*) images for measuring shape preservation; (b) vice versa for translation $B \rightarrow A \rightarrow B$

DATASET	NRMSE		SSIM		PSNR		DICE	
and translation direction	mean	st.dev	mean	st.dev	mean	st.dev	mean	st.dev
RCA A->B	**0,029**	0,010	**0,960**	0,003	**31,303**	2,927	**0,718**	0,015
RCA B->A	0,088	0,028	0,738	0,018	21,960	2,787	0,589	0,032
LCA A->B	0,270	0,192	0,935	0,022	17,474	7,200	0,637	0,074
LCA B->A	0,085	0,008	0,761	0,007	21,486	0,795	0,454	0,030

Table 1 Different similarity metrics (NRMSE, SSIM, PSNR) for evaluation of cycle-consistency loss, and Dice scores for evaluation of vessel preservation given for each data set and each translation direction.

preserved (Figure 9). First, manual segmentation of vessels was done for input testing images and corresponding output images produced by CycleGAN (with the white color marking vessels, and black color marking background). Then, an overlap of obtained binary images was calculated using the Dice metric. Results of segmentation evaluation are presented in Table 1 (last column) and examples of segmentation are given in Figure 10. Both images of overlays and Dice scores values display that RCA and LCA models are performing better at A to B translation (RCA A to B, LCA A to B) than at the opposite one. In the enhanced images, original physiology of vessels is preserved, though their shapes seem to be slightly thinner which have caused Dice scores to be 0.718 and 0.637 for RCA and LCA data sets, respectively.

Recommendations. Based on the different results explored above, authors advise to take into consideration the following:

1. Sizes of training sets for GANs. The number of 100-150 real images should be sufficient for training a good model, while a small number (10-50) may lead to slow convergence and overfitting.
2. Weight of the cycle-consistency component in the loss function (λ). A large λ (e.g. 30) suppresses fine details as well as extra variations in the resulting image, while a too small value (e.g. 10) fails to prevent deformations. An average value (e.g. 20, as used for RCA data set) should be optimal for our task.
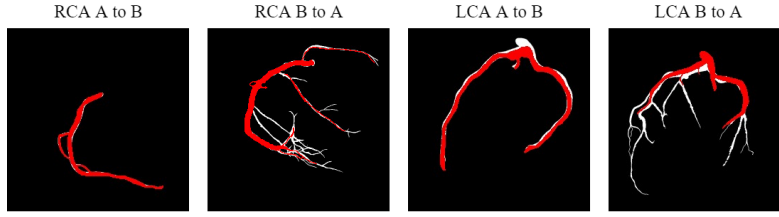
Fig. 10 Segmented images overlayed on top of each other for visual comparison. White color marks segmentations of original input images, and red color marks segmentations of images produced by CycleGAN.

3. Number of epochs. As we see on Figure 6, the more epochs we train our model with, the fewer image artifacts we obtain (chalkboard ones, white spots, found at early stages), but the more overfitting to a training set our models are. Providing this, we recommend evaluating results obtained after a different number of epochs to choose the best model for further style transfer.

5 Conclusions and future works

We introduced a new application of unsupervised image translation for generating realistic-looking angiograms from simulated ones. Our approach can be used with the purpose of obtaining more data for training and evaluating medical image processing methods for vessels specifically. We used CycleGAN and tuned its parameters to achieve the best results for RCA and LCA data sets.

One of the promising applications of our pipeline is generating more data and compiling it into a public dataset of coronary artery 2D images with corresponding annotation information.

Another application of the method is implementing a cross-modality translation but specifically for angiograms. The objective is to train models that learn style (e.g. noise and artifacts distribution) specific to different acquisition systems so that available angiography data could be used together for training and testing of other learning models.

Finally, results of B to A translation (real → simulated) suggested another area of application: unsupervised vessel segmentation where source domain is real angiograms and target domain is segmentation masks of simulated images, which can be obtained by applying Hessian-based filters.

Future works include experimenting with more data sets (other angiographic views, types of vessels i.e. aorta and pulmonary, bones, etc.); using obtained models as pre-trained models for other data sets; keeping *empty* background frames in training data; adding catheters, wires as a style; improving evaluation of method by implementing a "real vs fake" perceptual study with human observers; training similar models with different modality vessel images and with segmented images.

Research involving Human Participants and/or Animals

Ethical approval: All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed consent

Informed consent: Informed consent was obtained from all individual participants included in the study.

References

1. Chuquicusma, M. J., Hussein, S., Burt, J., & Bagci, U. (2018, April). How to fool radiologists with generative adversarial networks? a visual turing test for lung cancer diagnosis. In 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018) (pp. 240-244). IEEE.
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In Advances in neural information processing systems (pp. 2672-2680).
3. Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 1125-1134).
4. Kooi, T., Litjens, G., Van Ginneken, B., Gubern-Mérida, A., Sánchez, C. I., Mann, R., ... & Karssemeijer, N. (2017). Large scale deep learning for computer aided detection of mammographic lesions. Medical image analysis, 35, 303-312.
5. Liskowski, P., & Krawiec, K. (2016). Segmenting retinal blood vessels with deep neural networks. IEEE transactions on medical imaging, 35(11), 2369-2380.
6. Martin, R., Vachon, E., Miró, J., & Duong, L. (2017, April). 3D reconstruction of vascular structures using graph-based voxel coloring. In 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017) (pp. 1032-1035). IEEE.
7. M'hiri, F., Duong, L., Desrosiers, C., Leye, M., Miró, J., & Cheriet, M. (2016). A graph-based approach for spatio-temporal segmentation of coronary arteries in X-ray angiographic sequences. Computers in biology and medicine, 79, 45-58.
8. Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumor segmentation using convolutional neural networks in MRI images. IEEE transactions on medical imaging, 35(5), 1240-1251.
9. Wang, Jason, and Luis Perez. "The effectiveness of data augmentation in image classification using deep learning." *Convolutional Neural Networks Vis. Recognit* (2017).
10. Falcao, A., & Forkert, N. D. (2018, October). Four-Dimensional ASL MR Angiography Phantoms with Noise Learned by Neural Styling. In Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis: 7th Joint International Workshop, CVII-STENT 2018 and Third International Workshop, LABELS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings (Vol. 11043, p. 131). Springer.
11. Roth, H. R., Lee, C. T., Shin, H. C., Seff, A., Kim, L., Yao, J., ... & Summers, R. M. (2015). Anatomy-specific classification of medical images using deep convolutional nets. arXiv preprint arXiv:1504.04003.
12. Segars, W. P., Mahesh, M., Beck, T. J., Frey, E. C., & Tsui, B. M. (2008). Realistic CT simulation using the 4D XCAT phantom. Medical physics, 35(8), 3800-3808.
13. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., & Webb, R. (2017). Learning from simulated and unsupervised images through adversarial training. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 2107-2116).
14. Vasconcelos, Cristina Nader, and Bárbara Nader Vasconcelos. "Increasing deep learning melanoma classification by classical and expert knowledge based image transforms." *CoRR, abs/1702.07025* 1 (2017).

15. Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE international conference on computer vision (pp. 2223-2232).
16. Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. Neurocomputing, 321, 321-331.
17. Bermudez, C., Plassard, A. J., Davis, L. T., Newton, A. T., Resnick, S. M., & Landman, B. A. (2018, March). Learning implicit brain MRI manifolds with deep learning. In Medical Imaging 2018: Image Processing (Vol. 10574, p. 105741L). International Society for Optics and Photonics.
18. Zhao, He, Huiqi Li, and Li Cheng. "Synthesizing filamentary structured images with gans." arXiv preprint arXiv:1706.02185 (2017).
19. Guibas, John T., Tejpal S. Virdi, and Peter S. Li. "Synthetic medical images from dual generative adversarial networks." arXiv preprint arXiv:1709.01872(2017).
20. Costa, P., Galdran, A., Meyer, M. I., Niemeijer, M., Abràmoff, M., Mendonça, A. M., & Campilho, A. (2018). End-to-end adversarial retinal image synthesis. IEEE transactions on medical imaging, 37(3), 781-791.
21. Appan, Pujitha, and Jayanthi Sivaswamy. "Retinal image synthesis for cad development." International Conference Image Analysis and Recognition. Springer, Cham, 2018.
22. Galbusera, F., Niemeyer, F., Seyfried, M., Bassani, T., Casaroli, G., Kienle, A., & Wilke, H. J. (2018). Exploring the potential of generative adversarial networks for synthesizing radiological images of the spine to be used in in silico trials. Frontiers in Bioengineering and Biotechnology, 6, 53.
23. Staal, J., Abràmoff, M. D., Niemeijer, M., Viergever, M. A., & Van Ginneken, B. (2004). Ridge-based vessel segmentation in color images of the retina. IEEE transactions on medical imaging, 23(4), 501-509.
24. VESsel SEgmentation in the Lung 2012, IEEE International Symposium on Biomedical Imaging (ISBI 2012) https://vessel12.grand-challenge.org/
25. Rohkohl, C., Lauritsch, G., Keil, A., & Hornegger, J. (2010). CAVAREV—an open platform for evaluating 3D and 4D cardiac vasculature reconstruction. Physics in Medicine & Biology, 55(10), 2905.
26. Çimen, S., Unberath, M., Frangi, A., & Maier, A. (2017). CoronARe: A Coronary Artery Reconstruction Challenge. In Molecular Imaging, Reconstruction and Analysis of Moving Body Organs, and Stroke Imaging and Treatment (pp. 96-104). Springer, Cham.
27. Balocco, S., Gatta, C., Ciompi, F., Wahle, A., Radeva, P., Carlier, S., ... & Kovarnik, T. (2014). Standardized evaluation methodology and reference database for evaluating IVUS image segmentation. Computerized medical imaging and graphics, 38(2), 70-90.
28. Yan, Zengqiang, Xin Yang, and Kwang-Ting Tim Cheng. "A Three-stage Deep Learning Model for Accurate Retinal Vessel Segmentation." IEEE journal of biomedical and health informatics (2018).
29. Lahiri, A., Ayush, K., Kumar Biswas, P., & Mitra, P. (2017). Generative adversarial learning for reducing manual annotation in semantic segmentation on large scale miscroscopy images: Automated vessel segmentation in retinal fundus image as test case. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (pp. 42-48).
30. Tetteh, G., Efremov, V., Forkert, N. D., Schneider, M., Kirschke, J., Weber, B., ... & Menze, B. H. (2018). Deepvesselnet: Vessel segmentation, centerline prediction, and bifurcation detection in 3-d angiographic volumes. arXiv preprint arXiv:1803.09340.
31. Choi, Hongyoon, and Dong Soo Lee. "Generation of structural MR images from amyloid PET: Application to MR-less quantification." Journal of Nuclear Medicine59.7 (2018): 1111-1117.
32. Olut, S., Sahin, Y. H., Demir, U., & Unal, G. (2018, September). Generative adversarial training for MRA image synthesis using multi-contrast MRI. In International Workshop on PRedictive Intelligence In MEdicine (pp. 147-154). Springer, Cham.
33. Bi, L., Kim, J., Kumar, A., Feng, D., & Fulham, M. (2017). Synthesis of positron emission tomography (PET) images via multi-channel generative adversarial networks (GANs). In Molecular Imaging, Reconstruction and Analysis of Moving Body Organs, and Stroke Imaging and Treatment (pp. 43-51). Springer, Cham.
34. Ben-Cohen, A., Klang, E., Raskin, S. P., Soffer, S., Ben-Haim, S., Konen, E., ... & Greenspan, H. (2019). Cross-modality synthesis from CT to PET using FCN and GAN networks for improved automated lesion detection. Engineering Applications of Artificial Intelligence, 78, 186-194.

35. Liu, Ming-Yu, Thomas Breuel, and Jan Kautz. "Unsupervised image-to-image translation networks." Advances in Neural Information Processing Systems. 2017.
36. Kim, T., Cha, M., Kim, H., Lee, J. K., & Kim, J. (2017, August). Learning to discover cross-domain relations with generative adversarial networks. In Proceedings of the 34th International Conference on Machine Learning-Volume 70 (pp. 1857-1865). JMLR. org.
37. Chartsias, A., Joyce, T., Dharmakumar, R., & Tsaftaris, S. A. (2017, September). Adversarial image synthesis for unpaired multi-modal cardiac data. In International Workshop on Simulation and Synthesis in Medical Imaging (pp. 3-13). Springer, Cham.
38. Zhang, Zizhao, Lin Yang, and Yefeng Zheng. "Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2018.
39. Hiasa, Y., Otake, Y., Takao, M., Matsuoka, T., Takashima, K., Carass, A., ... & Sato, Y. (2018, September). Cross-modality image synthesis from unpaired data using CycleGAN. In International Workshop on Simulation and Synthesis in Medical Imaging (pp. 31-41). Springer, Cham.
40. Wolterink, J. M., Dinkla, A. M., Savenije, M. H., Seevinck, P. R., van den Berg, C. A., & Išgum, I. (2017, September). Deep MR to CT synthesis using unpaired data. In International Workshop on Simulation and Synthesis in Medical Imaging (pp. 14-23). Springer, Cham.
41. Dar, S. U., Yurt, M., Karacan, L., Erdem, A., Erdem, E., & Çukur, T. (2019). Image synthesis in multi-contrast MRI with conditional generative adversarial networks. IEEE transactions on medical imaging.
42. Yang, Q., Li, N., Zhao, Z., Fan, X., Chang, E. C., & Xu, Y. (2018). Mri image-to-image translation for cross-modality image registration and segmentation. arXiv preprint arXiv:1801.06940.
43. Welander, Per, Simon Karlsson, and Anders Eklund. "Generative adversarial networks for image-to-image translation on multi-contrast MR images-A comparison of CycleGAN and UNIT." arXiv preprint arXiv:1806.07777 (2018).
44. Mahmood, Faisal, Richard Chen, and Nicholas J. Durr. "Unsupervised reverse domain adaptation for synthetic medical images via adversarial training." IEEE transactions on medical imaging 37.12 (2018): 2572-2581.
45. Yi, X., Walia, E., & Babyn, P. (2018). Generative adversarial network in medical imaging: A review. arXiv preprint arXiv:1809.07294.