Biomedical Signal Processing and Control 85 (2023) 104886

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

journal homepage: www.elsevier.com/locate/bspc

Short communication

Translation of atherosclerotic disease features onto healthy carotid ultrasound images using domain-to-domain translation

Hazrat Ali [a], Emma Nyman [b], Ulf Näslund [b], Christer Grönlund [c,*]

[a] College of Science and Engineering, Hamad Bin Khalifa University, Doha, Qatar Foundation, Qatar
[b] Department of Public Health and Clinical Medicine, Umeå University, Umeå, Sweden
[c] Department of Radiation Sciences, Radiation Science, Biomedical Engineering, Umeå University, Umeå, Sweden

ARTICLE INFO

Keywords:
Domain-to-domain translation
Generative adversarial networks
Atherosclerosis
Ultrasound imaging
Cardiovascular disease

ABSTRACT

Objective: In this work, we evaluated a model for the translation of atherosclerotic disease features onto healthy carotid ultrasound images.
Methods: An un-paired domain-to-domain translation model – the cycle Generative Adversarial Network (cycleGAN) – was trained to translate between carotid ultrasound images of healthy arteries and images of pronounced disease. Translation performance was evaluated using the measurement of wall thickness in original and generated images. In addition, we explored disease translation in different tissue segments (subcutaneous tissue, muscle, lumen, far wall, and deep tissues), using structural similarity index measure (SSIM) maps.
Results: Features of pronounced disease were successfully translated to the healthy images (1.2 (0.33) mm vs 0.43 (0.07) mm, $p < 0.001$), while overall anatomy was retained as SSIM value was equal to 0.78 (0.02). Exploration of translated features showed that both arterial wall and subcutaneous tissues were modified in the translation, but that the subcutaneous tissue was subject to distortion of the anatomy in some cases. The image quality influenced the disease translation performance.
Conclusion: The results show that the model can learn a mapping between healthy and diseased images while retaining the overall anatomical contents. This is the first study on atherosclerosis disease translation in medical images.
Significance: The concept of translating disease onto existing healthy images may serve purposes such as education, cardiovascular risk communication in health conversations, or personalized modelling in precision medicine.

1. Introduction

Atherosclerosis leads to thickened arterial walls and focal plaques, and it is the major cause of cardiovascular disease (CVD), with a very high incidence of related deaths [1]. Ultrasound imaging of the carotid arteries is commonly used to assess disease stage, progression, and CVD risk, based on measurements of the carotid arterial wall thickness (cIMT) and the presence and composition of focal plaques [2].

In this work, we aim to evaluate a model to simulate this disease on a personalized level. We train a model to take images of a healthy person's carotid artery and generate a new image with features of pronounced disease but with preserved anatomical features. In general, such generated images could be used for educational purposes and simulations. In specific, recently, our group showed that presenting an image of a patient's carotid artery disease severity increases a patient's adherence to treatment and lowers risk factors [3,4]. Therefore, such generated images could be used as a tool in primary prevention and CVD risk communication within a health conversation [4].

The problem we want to solve is equivalent to translating features from one image domain to another. So-called domain-to-domain translation methods (deep learning models) have been proposed for similar tasks in medical imaging applications. E.g., the cycle generative adversarial network (cycleGAN) has been used to generate authentic speckle pattern features in simulated B-mode images [5], and to generate authentic musculoskeletal images [6], both from simple simulated images. GANs have also been used to translate features of super-resolution images onto images of low-resolution [7] and for the generation of different stages of tumour growth in ultrasound mammographic images

* Corresponding author at: Dept Radiation Sciences, Umeå University, S-90187 Umeå, Sweden.
E-mail addresses: haali2@hbku.edu.qa (H. Ali), emma.nyman@umu.se (E. Nyman), ulf.naslund@umu.se (U. Näslund), christer.gronlund@umu.se (C. Grönlund).

https://doi.org/10.1016/j.bspc.2023.104886
Received 27 October 2022; Received in revised form 30 January 2023; Accepted 18 March 2023
Available online 1 April 2023

[8]. Moreover, our group also recently applied a modified 3D cycleGAN model to generate authentic ultrasound image sequences of contracting skeletal muscles from simplified simulated sequences (video) [9]. To the best of our knowledge, the idea of disease injection to existing ultrasound images using deep learning is not yet explored.

In this work, we evaluate the 2D cycleGAN model [10] to solve this task, where translation between images of two domains with different feature distributions was obtained while retaining content of the original images [11–13].

2. Methods

2.1. Datasets

Longitudinal ultrasound B-mode images of the common carotid arteries were retrospectively included in this work from the VIPVIZA study [4]. The VIPVIZA study is an intervention study for primary cardiovascular disease prevention. The study enrolled 3532 healthy individuals with sub-clinical atherosclerosis of ages 40, 50, and 60 s. For each subject, multiple carotid ultrasound images were acquired at end-diastole (approximately-five consecutive heartbeats) from four different standardized projections (120-, 150-, 210- and 240-degrees Meijer arc) from both the left and right arteries. In total, approximately 140 000 (5 heartbeats × 4 projections × 3500 subjects × 2 examinations) images were available within the VIPVIZA study. The study conformed to the Declaration of Helsinki and was approved by the Swedish Ethical Review Authority.

The ultrasound images were acquired using the CardioHealth Station ultrasound system with a linear 7 MHz probe (Panasonic Healthcare Corporation of North America, Newark NJ, USA). Images had 8-bit resolution and size was 464x435 (field-of-view 40x38 mm, depth × width) and were down-sampled using cubic-splines interpolation to 400x400 for training purposes. The equivalent resolution was approximately 0.1 mm/px.

2.1.1. Selection of images

Healthy images (defined as Domain A) had thin cIMT (range 0.3–0.5 mm, mean 0.45 mm, N = 2319). Images with pronounced disease (defined as Domain B) had thick cIMT (range 1.1–2.0 mm, mean 1.2 mm, N = 2071). Images for training and testing were randomly selected from the two domains at an 80:20 ratio. The clinical characteristics of the two domain data sets are given in Table 1. Fig. 1A and B show example images of both domains.

Table 1
Subject and image characteristics of the training data domains A and B
Values represent mean and standard deviations.

	Domain A *(no disease)*	Domain B *(pronounced disease)*	*p*
Image sets			
N subjects	*210*	*170*	
N images	*2319*	*2071*	
N training images	*1855*	*1656*	
N test images	*464*	*415*	
Ultrasound measurements			
cIMT, mm	0.45 (0.9)	1.2 (0.20)	
min–max, *mm*	0.3–0.5	1.1–2.0	
Plaque present	No	Yes	
Clinical risk factors			
Low density lipo protein, (LDL)	3.51 (1.01)	3.83 (1.06)	0.003
High-density lipoprotein (HDL)	1.30 (0.39)	1.25 (0.39)	0.18
Systolic blood pressure, mmHg	127 (18)	133 (17)	0.002
Diastolic blood pressure, mmHg	83 (12)	83 (11)	0.45
Age, years	50.6 (7.6)	59.2 (2.8)	<0.001
Sex, (female:men), %	50:50	30:70	<0.001
Body-mass-index (BMI)	29 (5.6)	28 (5.0)	0.04
Framingham risk score (FRS)	10.4 (8.2)	20.7 (12.4)	<0.001

2.2. The CycleGAN model

Generative adversarial networks (GANs) have gained popularity for their ability to generate realistic-looking synthetic medical image data [10,14–16]. Domain-to-domain translation models are a branch of GANs that allow the translation of images from one domain to another with different feature distributions while retaining content [11–13,17].

When the images of two domains are not paired, the cycleGAN model offers a solution for this task. In this work, a cycleGAN model [10] was trained to translate between domain A (no atherosclerosis) and domain B (pronounced atherosclerosis) datasets.

The architecture of cycleGAN [10] comprises two GANs working in the opposite direction i.e., transformation from $G_1(x_A) \rightarrow x_B$ and transformation $G_2(x_B) \rightarrow x_A$. In this study, the cycleGAN was trained for 180 epochs (determined empirically based on a stopped reduction of the training loss). We experimented with ResNet [18] as well as U-Net architecture as the generator of the cycleGAN. We found that the ResNet architecture gave a superior performance, and hence we used a generator with the ResNet architecture.

The default setting for up-sampling in the generator part uses transpose convolution. The default method down-samples the image 4 times and then extracts the feature using the ResNet18 model. Following this, the model up-samples the features by a factor of 4 using transposed convolution. In our implementation, we replaced the transpose convolution by sub-pixel convolution layer [19]. The benefit of the sub-pixel convolution layer is that it replicates the channels of the image and shuffles the pixels in each replica. It has been demonstrated to give an advantage in the case of image super-resolution [19].

2.3. Evaluation and exploration of disease translation

The disease translation was evaluated (on the test data) by comparing the cIMT of the input and the generated output image pairs. The cIMT was manually measured in the input and output images (randomly presented and blinded to the operator). Next, features of the disease translation were explored by quantifying the structural similarity index measure (SSIM) map between the input–output image pairs [20]. The SSIM ranges from 0 to 1 (low to high similarity). The authenticity of generated images was assessed by two carotid ultrasound experts (CG and EN) who, presented blindly and in random order, classified 50 real and 50 generated images as *real* or *generated*.

Association between image quality and SSIM was quantified by Pearson's correlation coefficient. Image quality was estimated as contrast tissue-to-lumen, CTL $=(|I_{lumen} - I_{tissue}|)/|I_{tissue}|$, where the lumen and tissue intensities were computed as the mean intensities in the corresponding muscle and lumen ROIs, respectively.

3. Results

3.1. Evaluation of disease translation

The cycleGAN model was trained to translate between two carotid ultrasound image sets: domain A – healthy (thin arterial wall, carotid intima-media (cIMT) < 0.5 mm, N = 1855), and domain B – pronounced atherosclerotic disease (thick cIMT, >1.1 mm N = 1656) (Fig. 1A, B). Fig. 1C and 1D show an example of one originally healthy and the corresponding generated image with pronounced disease. The cIMT was 0.43 (0.07) mm and 1.2 (0.33) mm in the test images of the original input and generated output images (Fig. 1F, p < 0.001). Thus, the healthy images with thin cIMT were successfully translated to pronounced disease with thick cIMT (similar to that of the training data for domain B, 1.12 (0.18) mm). cIMT was not possible to measure in N = 41 (15 %) images due to poor image quality. The domain experts classified the generated images as real images in 67 % and 75 % of the cases.

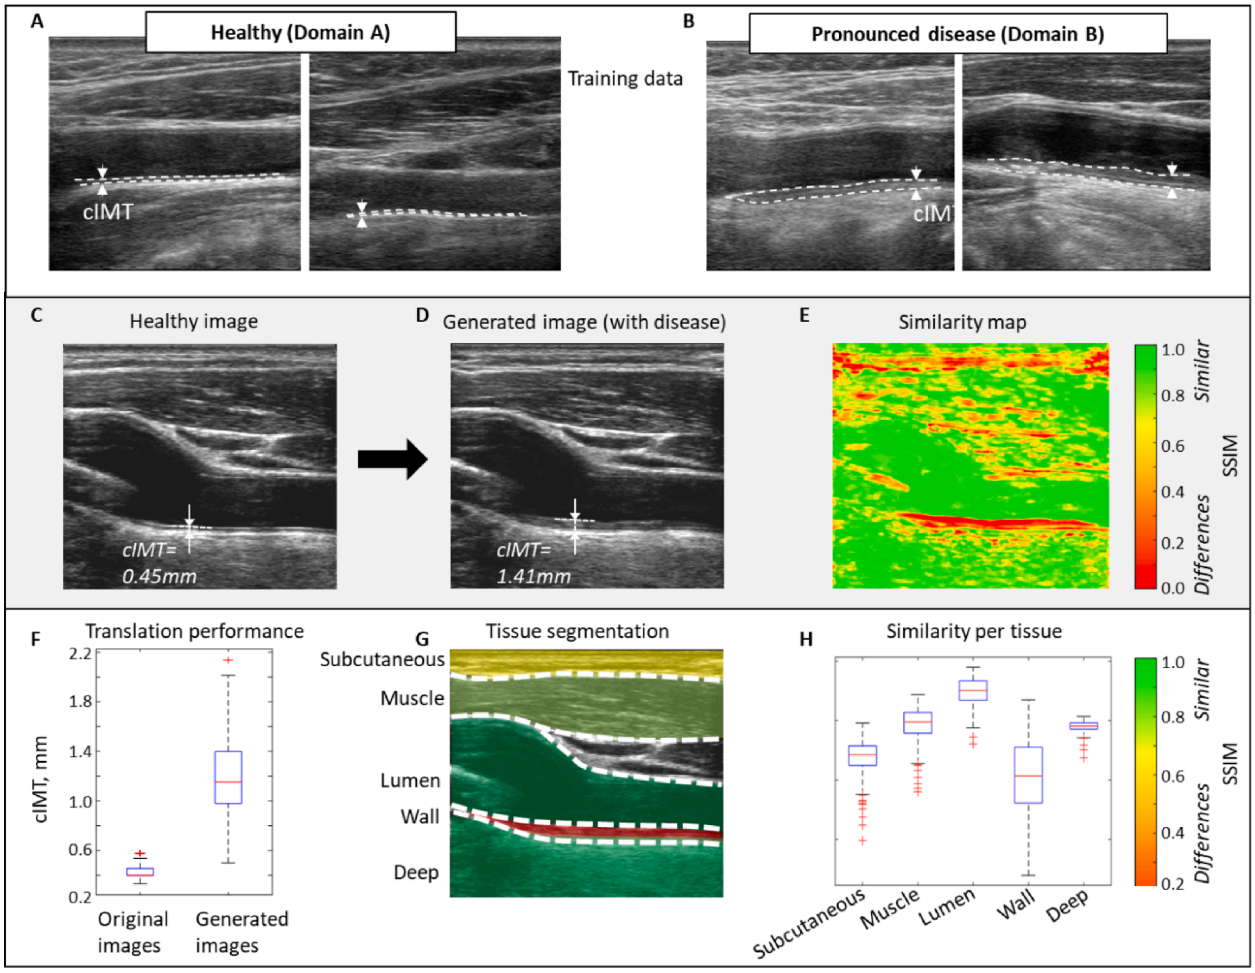

Fig. 1. **A**: Examples of training images with no atherosclerotic disease (domain A, cIMT < 0.5 mm and no plaques). **B**: Examples of training images with pronounced disease (domain B, cIMT > 1.1 mm and plaques present). **C**: Original image (healthy, thin wall, Domain A). **D**: Generated image with translated disease (thick wall) with overall anatomy retained. **E**: Similarity map (SSIM) between the two images. A low value indicates disease adaptation to this segment. **F**: Wall thickness (cIMT) of the original and generated images. **G**: Illustration of segmented regions. **H**: Similarity (SSIM) in segmented regions between original and generated image pairs (N = 460 images).

3.2. Exploration of disease translation

The translated features of the disease were explored using the SSIM between original and generated image pairs (Fig. 1E) in the full image, subcutaneous, muscle, lumen, arterial wall, and deep tissues, respectively (Fig. 1G, H). The SSIM showed a low similarity in the far wall, indicating that the model had modified this region in the generated image as compared to the original. The SSIM for the entire image was 0.78 (0.02). The far carotid wall and the subcutaneous tissue had lower SSIM as compared to the muscle, carotid artery lumen, and the deep tissue (Fig. 1H, $p < 0.001$). Fig. 2A-C shows additional examples of original and generated images with successfully translated disease.

There were a few cases where the underlying anatomy was distorted in the subcutaneous tissues of the generated images (Fig. 2D-F). Two regions with distortions could be observed: 1) the outer left and right tissues had reduced similarity for all generated images, and 2) in a few cases, the anatomy of the original image was distorted in the subcutaneous tissue segment (N = 22, ~5% of the test images).

3.3. Influence of image quality on disease translation

Motivated by the large range of SSIM of the far wall (0.25–0.70, Fig. 1H) and the observed 15 % dropout in cIMT measurements, we analysed if the image quality could account for this variation. Image quality was quantified using the ratio between wall echogenicity (average pixel intensity of the wall tissue) and the echogenicity of the lumen – Contrast tissue-to-lumen (CTL). The correlation between CTL and SSIM in the different tissues was: subcutaneous 0.08 ($p = 0.16$), muscle −0.11 ($p = 0.073$), blood −0.11 ($p = 0.61$), far wall −0.46 ($p < 0.001$), deep tissue −0.10 ($p = 0.09$), respectively. These results show that image quality only influenced the translation of disease features in the far wall tissues significantly. Thus, high noise in the lumen (low CTL) resulted in less translated disease features (high SSIM). Fig. 2G-I shows an example of how poor contrast influenced the translation.

4. Discussion

This study showed that the cycleGAN model can translate image features of pronounced atherosclerotic disease (thickened cIMT) onto images of healthy carotid arteries. The overall high similarity (SSIM) indicates that the main image features, i.e., the underlying anatomy, were retained. In contrast, the lower SSIM of the subcutaneous and far wall tissues indicate that these were adapted with the disease. The disease adaptation of the far wall was expected due to the selection of the datasets. The slight modification to the subcutaneous tissue was likely related to artifact generation by the method (see Fig. 2E-F). Noise in the lumen of the images (low CTL) was found to reduce the translation of disease features (higher SSIM). Such noise is commonly seen in carotid

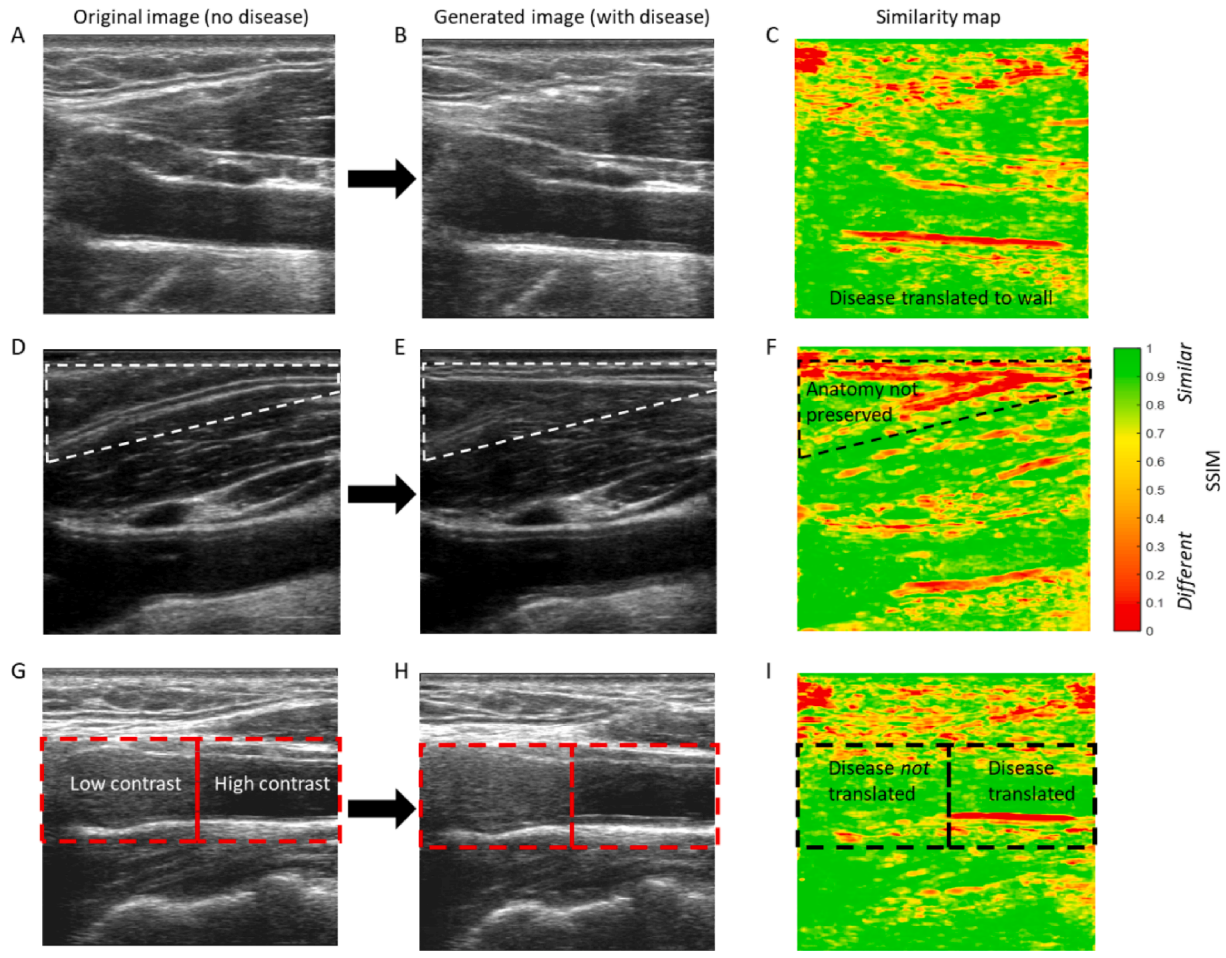

Fig. 2. **A-C**: An example with successful disease translation and retained anatomy, as indicated by the horizontal red region of the far wall in C. **D-F**: An example where the generator model has distorted the underlying anatomy of the subcutaneous tissue as indicated by red regions in the **F**. **G-I**: An example of how contrast influenced the disease translation. The original image **G** had two regions with low and high contrast (contrast tissue-to-lumen, CTL). It can be seen in the SSIM map **I** that no disease was translated to the low-contrast region.

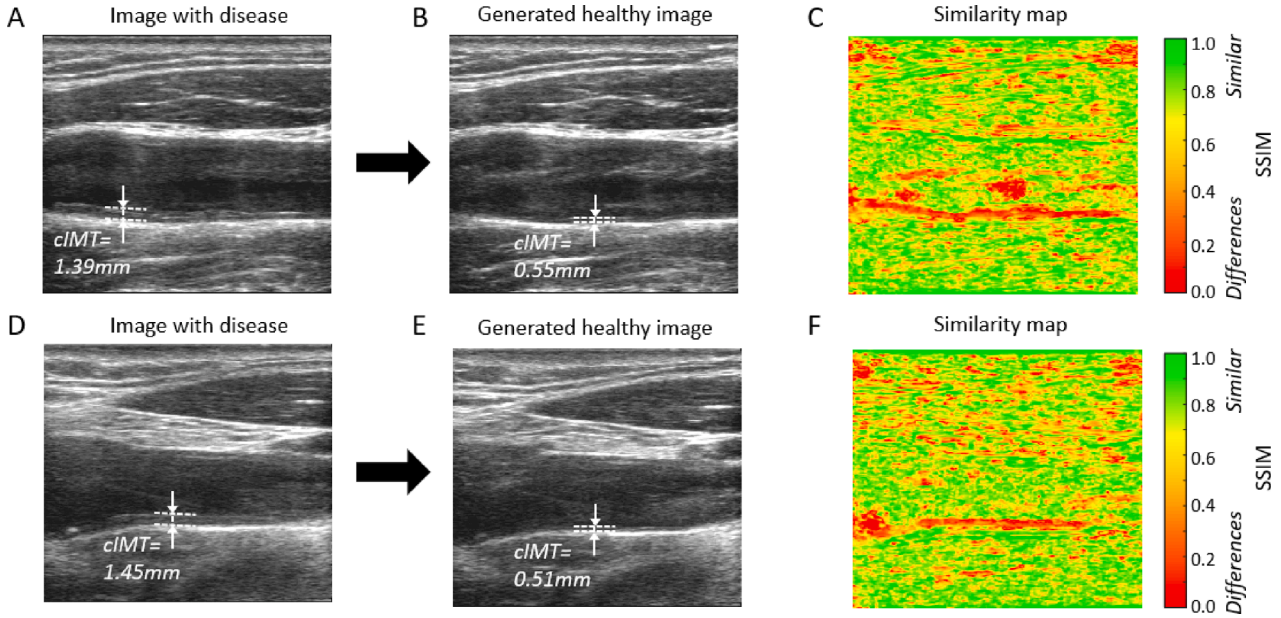

Fig. 3. **A-C** and **D-F:** Two examples of disease removal by the reverse generator.

ultrasound images but can be reduced using harmonic imaging sequences [21]. The classification by the domain experts showed that the generated images were visually perceived as real images in a majority of the cases.

The performance of the reverse generator of the cycleGAN that is trained to generate healthy images from images with pronounced disease was also assessed (Fig. 3). The cIMT was 0.56 (0.12) and 1.36 (0.29) mm for the input and generated images (Fig. 3AB, $p < 0.001$). The SSIM for the entire image was 0.64 (0.07). Thus, the performance of the reverse generator for disease removal had similar performance as the disease generation generator.

The concept of translating disease onto existing ultrasound images is new, but in other medical image modalities it has recently been demonstrated. E.g. cycleGAN models were used to generate pathological images from normal optical coherence tomography (OCT) images of the eye [22], and the Ad cycleGAN model was proposed to generate COVID-19 positive chest X-ray images from normal images [23].

The cycleGAN model was trained on relatively limited data, which could influence the performance. In addition, translation between only two stages of the disease was adapted. While the chosen model only can perform one-to-one mapping – here between healthy and diseased states, models that support the generation of different stages of disease or outcomes could be achieved by adding data constraints and allowing one-to-many mapping models e.g. [24,25].

This work presents a proof-of-concept on the generation of disease into healthy carotid ultrasound images. There are several potential applications of this concept. First, it may be used to simulate images of disease for education or for injection of disease into an existing image of a healthy subject for risk communication in a health conversation - to improve adherence to prescribed treatment [4]. Second, the current understanding of image features of atherosclerotic disease concerns mainly morphological features (e.g., cIMT), and a data-driven disease generation model may provide new insights on additional image features in different stages of the disease beyond the currently used measurements. Third, models that can learn a mapping between healthy and diseased domains, open new dimensions for personalized disease prediction modelling [26] and should be evaluated in future studies. Finally, the model could also be used to generate diseased or healthy images for data augmentation purposes, to improve the training of deep learning models on atherosclerosis and ultrasound images.

5. Conclusion

Translation of atherosclerotic disease onto existing ultrasound images of healthy carotid arteries was demonstrated. The proposed concept may be useful in education, simulations, or primary prevention to boost risk communication in a health conversation on CVD risk.

Funding

This study was supported by the Swedish Research Council (2015-04461), Region Västerbotten (RV-930168), The Heart Foundation of Northern Sweden, and the Kempe foundations (SMK-1868).

CRediT authorship contribution statement

Hazrat Ali: Conceptualization, Methodology, Software, Data curation, Writing – original draft, Writing – review & editing. **Emma Nyman:** Writing – review & editing. **Ulf Näslund:** Writing – review & editing. **Christer Grönlund:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Funding acquisition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] S. Barquera, A. Pedroza-Tobías, C. Medina, L. Hernandez-Barrera, K. Bibbins-Domingo, R. Lozano, A.E. Moran, Global Overview of the Epidemiology of Atherosclerotic Cardiovascular Disease, Arch. Med. Res. 46 (2016) 328–338, https://doi.org/10.1016/j.arcmed.2015.06.006.

[2] T.Z. Naqvi, M.S. Lee, Carotid intima-media thickness and plaque in cardiovascular risk assessment, J. Am. Coll. Cardiol. Img. 7 (2014) 1025–1038, https://doi.org/10.1016/j.jcmg.2013.11.014.

[3] A. Bengtsson, M. Norberg, N. Ng, B. Carlberg, C. Grönlund, J. Hultdin, B. Lindahl, B. Lindahl, S. Nordin, E. Nyman, P. Wennberg, P. Wester, U. Näslund, The beneficial effect over 3 years by pictorial information to patients and their physicians about subclinical atherosclerosis and cardiovascular risk: Results from the VIPVIZA randomized clinical trial, Am J Prev Cardiol 7 (2021), https://doi.org/10.1016/j.ajpc.2021.100199.

[4] U. Näslund, N. Ng, A. Lundgren, E. Fhärm, C. Grönlund, H. Johansson, B. Lindahl, B. Lindahl, K. Lindvall, S.K. Nilsson, M. Nordin, S. Nordin, E. Nyman, J. Rocklöv, D. Vanoli, L. Weinehall, P. Wennberg, P. Wester, M. Norberg, Visualization of asymptomatic atherosclerotic disease for optimum cardiovascular prevention (VIPVIZA): a pragmatic, open-label, randomised controlled trial, Lancet 393 (2019) 133–142, https://doi.org/10.1016/S0140-6736(18)32818-6.

[5] L. Bargsten, A. Schlaefer, SpeckleGAN: a generative adversarial network with an adaptive speckle layer to augment limited training data for ultrasound image processing, Int. J. Comput. Assist. Radiol. Surg. 15 (2020) 1427–1436, https://doi.org/10.1007/s11548-020-02203-1.

[6] N.J. Cronin, T. Finni, O. Seynnes, Using deep learning to generate synthetic B-mode musculoskeletal ultrasound images, Comput. Methods Programs Biomed. 196 (2020), 105583, https://doi.org/10.1016/j.cmpb.2020.105583.

[7] W. Ahmad, H. Ali, Z. Sha, S. Azmat, A new generative adversarial network for medical images super resolution, Sci. Rep. 12 (2022), https://doi.org/10.1038/s41598-022-13658-4.

[8] T. Fujioka, K. Kubota, M. Mori, L. Katsuta, Y. Kikuchi, K. Kimura, M. Kimura, M. Adachi, G. Oda, T. Nakagawa, Y. Kitazume, Virtual interpolation images of tumor development and growth on breast ultrasound image synthesis with deep convolutional generative adversarial networks, J. Ultrasound Med. 40 (2021) 61–69, https://doi.org/10.1002/jum.15376.

[9] H. Ali, J. Umander, R. Rohlén, O. Röhrle, C. Grönlund, Modelling intra-muscular contraction dynamics using in silico to in vivo domain translation, Biomed eng Online 28 (2022) 21–46, https://doi.org/10.1186/s12938-022-01016-4.

[10] J.Y. Zhu, T. Park, P. Isola, A.A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks. In Proceedings of the IEEE international conference on computer vision, 2017, pp. 2223-2232 10.1109/ICCV.2017.244.

[11] X. Huang, M.Y. Liu, S. Belongie and J. Kautz, Multimodal unsupervised image-to-image translation. In Proceedings of the European conference on computer vision (ECCV) (pp. 172-189), 2018.

[12] H.Y. Lee, H.Y. Tseng, J.B. Huang, M. Singh and M.H. Yang, Diverse image-to-image translation via disentangled representations. In Proceedings of the European conference on computer vision (ECCV), 2018, pp. 35-51.

[13] Z. Zhang, L. Yang and Y. Zheng, Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In Proceedings of the IEEE conference on computer vision and pattern Recognition, 2018, pp. 9242-9251. 10.48550/arXiv.1802.09655.

[14] X. Yi, E. Walia, P. Babyn, Generative adversarial network in medical imaging: A review, Med. Image Anal. 58 (2019), 101552, https://doi.org/10.1016/j.media.2019.101552.

[15] J. Wei, A. Suriawinata, L. Vaickus, B. Ren, X. Liu, J. Wei and S. Hassanpour, Generative image translation for data augmentation in colorectal histopathology images. Proceedings of machine learning research, 116 (2019), 10, PMCID: PMC8076951.

[16] V. Sorin, Y. Barash, E. Konen, E. Klang, Creating artificial images for radiology applications using generative adversarial networks (GANs)–A systematic review, Acad. Radiol. 27 (2020) 1175–1185, https://doi.org/10.1016/j.acra.2019.12.024.

[17] Q. Mao, H.Y. Lee, H.Y. Tseng, S. Ma and M.H. Yang, Mode seeking generative adversarial networks for diverse image synthesis. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition 2019, pp. 1429-1437, ieeecomputersociety.org/10.1109/CVPR.2019.00152.

[18] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image reconstruction, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, 770-778, ieeecomputersociety.org/10.1109/CVPR.2016.90.

[19] W. Shi, J. Caballero, F. Huszár, J. Totz, A.P. Aitken, R. Bishop, D. Rueckert, Z. Wang, Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 1874-1883, ieeecomputersociety.org/10.1109/CVPR.2016.207.

[20] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process. 13 (2004) 600–612, https://doi.org/10.1109/TIP.2003.819861.

[21] T.L. Szabo, Diagnostic ultrasound imaging: Inside-out, 1st edition, Academic Press, Elsevier, 2005.
[22] T.K. Yoo, J.Y. Choi, H.K. Kim, Feasibility study to improve deep learning in OCT diagnosis of rare retinal diseases with few-shot classification, Med Biol Eng Comput (2021) 401–415, https://doi.org/10.1007/s11517-021-02321-1.
[23] Z. Liang, J.X. Huang, S. Antani, Image translation by Ad cycleGAN for COVID-19 X-ray images: A new approach for controllable GAN, Sensors (2022) 1–15, https://doi.org/10.3390/s22249628.
[24] Z. Shen, S.K. Zhou, Y. Chen, B. Georgescu, X. Liu, T. Huang, One-to-one Mapping for Unpaired Image-to-image Translation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (pp. 1170-1179), 2020.
[25] G. Lu, Z. Zhou, Y. Song, K. Ren, Y. Yu, Guiding the one-to-one mapping in CycleGAN via optimal transport. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 33, No. 01, pp. 4432-4439), 2019.
[26] J. Corral-Acero, F. Margara, M. Marciniak, C. Rodero, F. Loncaric, Y. Feng, A. Gilbert, J.F. Fernandes, H.A. Bukhari, A. Wajdan, M.V. Martinez, The 'Digital Twin' to enable the vision of precision cardiology, Eur. Heart J. 41 (2020) 4556–4564, https://doi.org/10.1093/eurheartj/ehaa159.