Computers in Biology and Medicine 190 (2025) 110030

Contents lists available at ScienceDirect

Computers in Biology and Medicine

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

A domain adaptation model for carotid ultrasound: Image harmonization, noise reduction, and impact on cardiovascular risk markers

Mohd Usama [a],*, Emma Nyman [b], Ulf Näslund [b], Christer Grönlund [a]

[a] Department of Diagnostics and Intervention, Biomedical Engineering and Radiation Physics, Umea University, Umea, Sweden
[b] Department of Public Health and Clinical Medicine, Umea University, Umea, Sweden

ARTICLE INFO

Keywords:
Deep learning
Generative Adversarial Network
Medical image analysis
Cardiovascular disease assessment
Domain adaptation
Image harmonization
Noise reduction
Carotid ultrasound images

ABSTRACT

Deep learning has been used extensively for medical image analysis applications, assuming the training and test data adhere to the same probability distributions. However, a common challenge arises when dealing with medical images generated by different systems or even the same system with varying parameter settings. Such images often contain diverse textures and noise patterns, violating the assumption. Consequently, models trained on data from one machine or setting usually struggle to perform effectively on data from another. To address this issue in ultrasound images, we proposed a Generative Adversarial Network (GAN) based model in this paper. We formulated image harmonization and denoising tasks as an image-to-image translation task, wherein we adapt the texture pattern and reduced noise in Carotid ultrasound images while keeping the image content (the anatomy) unchanged. The performance was evaluated using feature distribution and pixel-space similarity metrics. In addition, blood-to-tissue contrast and influence on computed risk markers (Grey scale median, GSM) were evaluated. The results showed that domain adaptation was achieved in both tasks (histogram correlation 0.920 (0.043) and 0.844 (0.062)), as compared to no adaptation (0.890 (0.077) and 0.707 (0.098)), and that the anatomy of the images was retained (structure similarity index measure e.g. the arterial wall 0.71 (0.09) and 0.80 (0.08)). In addition, the image noise level (contrast) did not change in the image harmonization task (-34.1 (3.8) vs -35.2 (4.1) dB) but was improved in the noise reduction task (-23.5 (3.2) vs -46.7 (18.1) dB). To validate the performance of the proposed model, we compare its results with CycleGAN, the current state-of-the-art model. Our model outperformed CycleGAN in both tasks. Finally, the risk marker GSM was significantly changed in the noise reduction but not in the image harmonization task. We conclude that domain translation models are powerful tools for improving ultrasound image while retaining the underlying anatomy, but downstream calculations of risk markers may be affected.

1. Introduction

Deep learning has been used extensively for medical image analysis applications, assuming the training and test data adhere to the same probability distributions. However, a common challenge arises when dealing with medical images generated by different systems or even the same system with varying parameter settings. Such images often contain diverse textures and noise patterns, violating the assumption. Consequently, models trained on data from one machine or setting usually struggle to perform effectively on data from another.

Deep learning for analysis of ultrasound images has attracted much attention for multiple clinical applications, such as carotid artery plaque detection and segmentation [1,2], atherosclerosis risk assessment [3,4], breast cancer detection and segmentation [5], etc. However, images from different ultrasound systems often vary in many aspects due to different imaging parameter settings, algorithm design, hardware components, etc. Therefore, the texture pattern of images acquired by different ultrasound systems may not belong to the same feature population distribution. For example, images taken from different ultrasound machines or different settings usually differ in texture pattern and noise. In addition, the images' texture and noise can vary depending on the composition and distribution of tissues of the imaged subject. This may cause generalizability problems when using pre-trained models for new images from different systems or on images from patients with varying noise distributions. The current state-of-the-art to account for such image feature variations in carotid ultrasound imaging, is to apply simple standardization of intensities prior to computation of clinical risk markers [6,7].

* Corresponding author.
E-mail addresses: mohd.usama@umu.se (M. Usama), emma.nyman@umu.se (E. Nyman), ulf.naslund@umu.se (U. Näslund), christer.gronlund@umu.se (C. Grönlund).

https://doi.org/10.1016/j.compbiomed.2025.110030
Received 7 August 2024; Received in revised form 10 January 2025; Accepted 12 March 2025
Available online 3 April 2025

Several studies have leveraged Generative Adversarial Networks (GANs) [8] in ultrasound imaging to address generalizability, for domain adaptation and denoising individually. These networks, known for their ability to generate realistic images, have shown promise in various areas of medical imaging, such as denoising [9], domain adaptation [10], disease translation [11], synthetic image generation [12], and segmentation [5]. Several researchers have introduced different methods and techniques in ultrasound image processing to enhance these networks effectiveness and efficiency [13]. They have explored the potential of GANs, which has resulted in a surge of publications focusing on quality adaptation and denoising. Zhou et al. [14] proposed a two-stage GAN for translating low-quality ultrasound images into high-quality images, which are usually degraded due to the limited size of portable ultrasound devices. Huang et al. [15] proposed a stability-enhanced CycleGAN-based model to harmonize ultrasound images from various medical centres through domain transformation. Xia et al. [16] proposed a Multilevel Structure-Preserved Generative Adversarial Network (MSP-GAN) for domain adaptation of ultrasound images while maintaining image intravascular structures. Athreya et al. [17] proposed a CycleGAN-based model with perceptual loss for quality enhancement of ultrasound images. Noise reduction in ultrasound images plays another crucial role in enhancing the quality of the images. Zhang et al. [18] proposed a GAN-based model to adapt the noise pattern in both optical coherence tomography and ultrasound images. Wavelet-based GAN [19] have also been proposed to suppress speckle noise while preserving image features, leading to high-quality image reconstruction. Additionally, several studies [20–22] have shown that the use of GANs in ultrasound image processing yields promising improvements in image quality.

Taken together, previous work highlight that image harmonization and denoising are important challenges in carotid ultrasound images, and separate GAN models can address these tasks individually [9]. However, this requires building different models for each task which is counteracting on the issue of generalization (the model only works for one specific task which is similar to a model that is only works for images from a specific imaging system). In addition, previous work has not studied how these models influence down-stream computation of image-based risk markers using the resulting adapted texture pattern. This is crucial to understand if models are to be used in clinical research.

This paper proposes a new GAN-based model that can address image harmonization and image denoising tasks with the same architecture. The proposed model contains one generator and two discriminators, and model training is regularized with content, noise and adversarial losses. The generator aims to take a source image and generate an image with features of the target domain. The first discriminator with content loss ensures that the generated image's content (anatomy) is preserved, and the second discriminator with noise loss enforces the noise features in the image.

The main contributions of this paper are:

- We propose a GAN-based model to address the problem of denoising and domain adaptation of ultrasound images through an unpaired image-to-image translation task. The model is trained to transfer target domain image features into the source domain image without changing the underlying anatomical content of the image.
- We introduce the noise loss, which employs the Wasserstein distance to measure the dissimilarity between the stylistic features derived from the generator's early layer and the style of target domain images.
- We fine-tune the generator architecture from Jun-Yan Zhu et al. [23] by adding two extra blocks (one convolution block and one deconvolution block) to manage the feature map extraction from the first three layers.
- We evaluate the performance of the domain adaptation task and preservation of anatomical content. In addition, we evaluate the domain adaptation's impact on clinical risk markers. To our knowledge, this is the first study to do so.

The rest of the paper is organized as follows. Section 2 explains the method. Section 3 describes the results, and Section 4 discusses the findings. The last section concludes the paper.

2. Methods

This paper presents an unpaired image-to-image translation model such that an input image with one feature distribution can be converted to an output image with another feature distribution while preserving the underlying anatomical content of the images. We target the problems of image feature harmonization and noise reduction. Generally, paired images from the corresponding domains (different systems, different noise levels) are often unavailable. Therefore, we chose to solve this as an unpaired image-to-image translation task. In this work, we consider the underlying anatomical information of the images as content and noise pattern as a style and formulate the task as a style transfer task.

2.1. GAN-based model

To solve the aforementioned problem, we propose a GAN-based model consisting of one generator and two discriminators. We formulate the task as an unpaired image-to-image translation task and incorporate a noise loss. This loss is determined using the Wasserstein distance and the style of feature maps derived from the early layers of the generator and the target domain images (see Fig. 1).

Our approach is inspired by methods [18,24], demonstrating the independent processing of feature representations learned by neural networks. Our features transfer approach combines a domain-to-domain translation concept, which builds a mapping between two domains for adaptation instead of mapping between two individual images, unlike [24], with a noise loss employing the Wasserstein distance and the stylistic features derived from the generator instead of the discriminator, unlike [18].

We define the training source dataset as S and the target dataset as T. Our main objective is to learn a mapping from S to T so that the feature distribution of image $x \epsilon S$ will be translated to x' images such that anatomical content is preserved while the style of image feature distribution is similar to that of T. The generator's job is to translate the source images that map into the target domain. The D_n distinguishes the texture pattern and noise between domains S and T. The D_c distinguishes the anatomical content of images. $\mathcal{L}_{an}$ and $\mathcal{L}_{ac}$ are used for adversarial training. L_c is used to preserve the anatomical contents. L_n is used to transfer texture patterns and reduce the noise, in experiments 1 and 2, respectively. Finally, adversarial losses are combined with content and noise losses to train the model. The computation of all loss functions is given below.

2.1.1. Adversarial losses

The $\mathcal{L}_{an}$ and $\mathcal{L}_{ac}$ are used to train the D_n and D_c, respectively. The D_n job distinguishes between the target image $y \epsilon T$ and the translated image $G(x)$. The D_c distinguishes between the translated image $G(x)$ and the real image $x \epsilon S$. The Computation of the losses are as follows:

$$\mathcal{L}_{an}(G, D_n) = \mathbb{E}_{y \sim T}[\log D_n(y)] + \mathbb{E}_{x \sim S}[\log(1 - D_n(G(x)))] \quad (1)$$

$$\mathcal{L}_{ac}(G, D_c) = \mathbb{E}_{x \sim S}[\log D_c(x)] + \mathbb{E}_{x \sim S}[\log(1 - D_c(G(x)))] \quad (2)$$

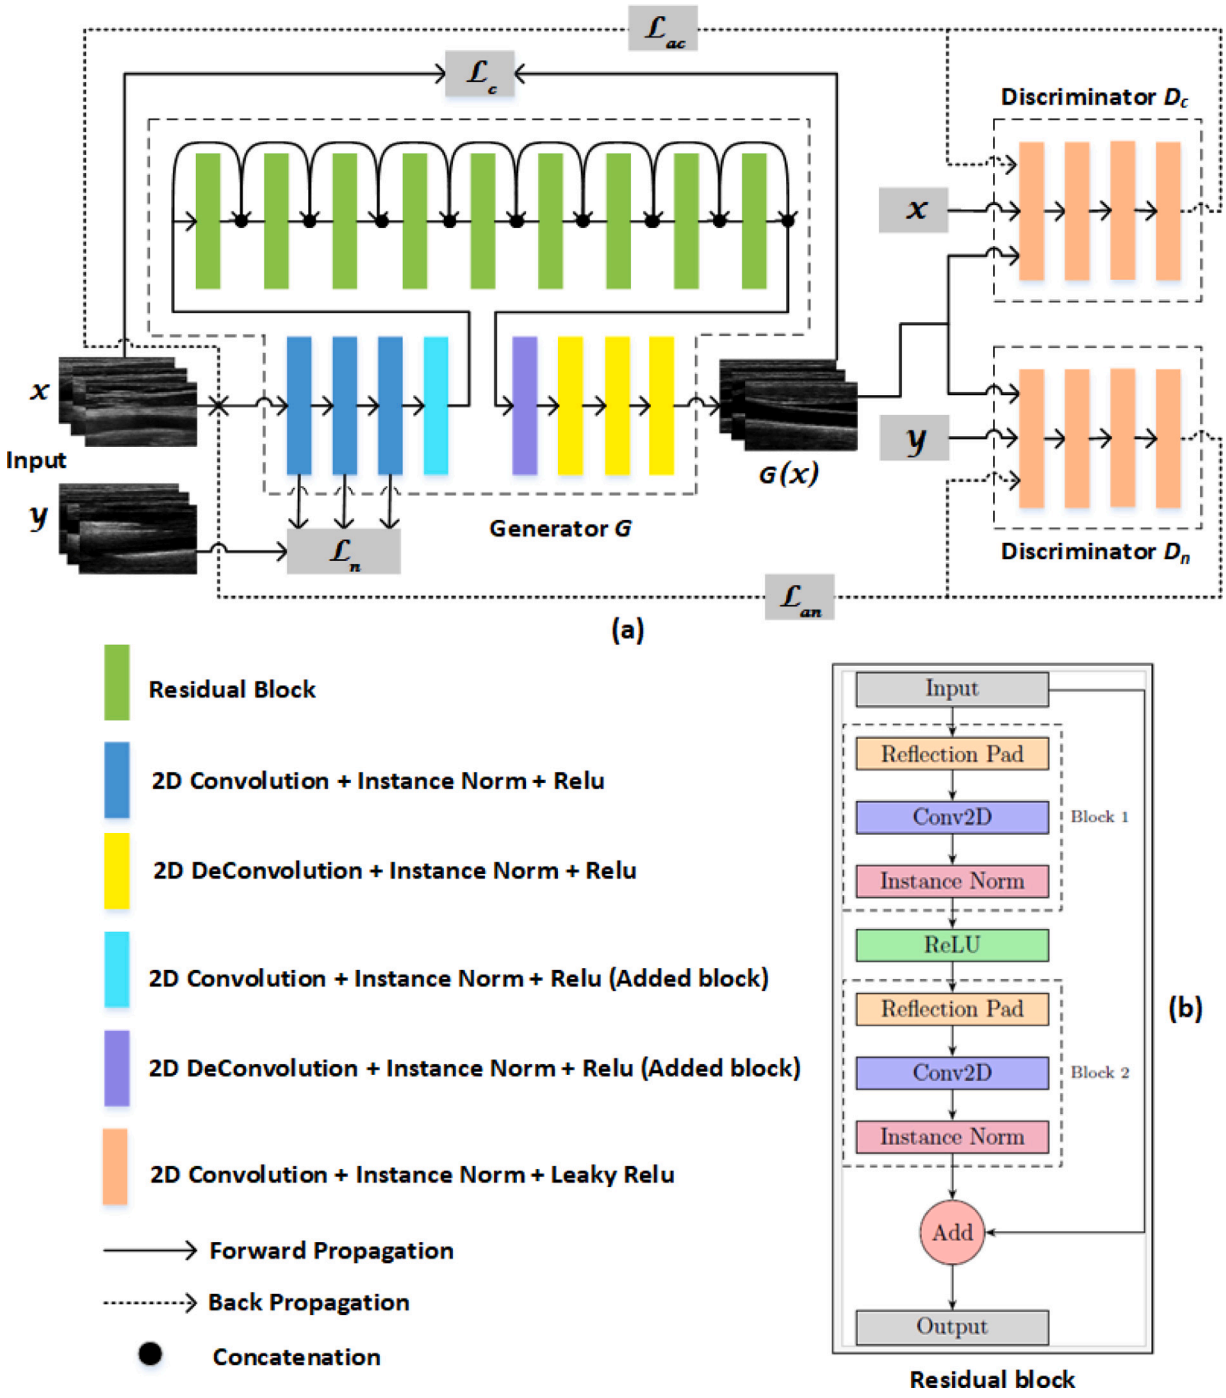

Fig. 1. (a) Proposed model. G denotes the generator network and two discriminators, D_c and D_n, trained by adversarial losses $\mathcal{L}_{ac}$ and $\mathcal{L}_{an}$, respectively. $x \sim S$ denotes the image set from the source domain. $y \sim T$ represents the image set from the target domain. $G(x)$ denotes the image set generated by G. L_c indicates the content loss calculated from the feature difference of the source and $G(x)$ image. L_n represents the noise loss calculated from the feature difference between the target image and the early layers of the generator. (b) a residual block.

2.1.2. Content loss

To preserve the anatomical contents of the image, the content loss L_c is computed from the last layer of generator G as follows:

$$\mathcal{L}_c(G) = \mathbb{E}_{x\sim S}\left[\left|F(G(x_i))^l - F(x_i)^l\right|\right] \tag{3}$$

where $F(\cdot)^l$ represents the feature maps from the last layer.

2.1.3. Noise loss

We formulate noise as style [24]. The style of an image was computed mathematically by Gram Matrix [25] G_r as follows:

$$G_r(y)^{(G,l)}_{(i,j)} = \text{vec}\left[F(y)^{(G,l)}_i\right] \otimes \text{vec}\left[F(y)^{(G,l)}_j\right] \tag{4}$$

where $G_r(y)^{(G,l)}$ is the Gram matrix of lth layer from generator's feature maps. The subscript i and j denotes the ith and jth maps of the $G_r(y)^{(G,l)}$, $vec[\cdot]$ represents the vectorization operation, and $\otimes$ represents the inner product.

L_n is computed as follows:

$$\mathcal{L}_n(G) = \mathbb{E}_{(x,y)\sim(S,T)}\left[\sum_{l=1}^{3} W\left(G_r(G(x)), G_r(y)\right)\right] \quad (5)$$

where W is the Wasserstein distance [26], we trained our model with a feature map of various early layers such as 1, 2, and 3 and decided to keep 3 for final results based on subjective visualization results decided by a US imaging expert and biomedical scientist, Christer Gronlund, who has over 20 years of experience in clinical ultrasound. Thus, we used only the first three layers of the feature map corresponding to the noise and texture pattern. The l denotes the layers in the generator.

2.1.4. Objective function

Full objective function is calculated as follows:

$$\arg\min_{G}\max_{D_n,D_c} \mathcal{L}(G, D_n, D_c) = \mathcal{L}_{an}(G, D_n) + \mathcal{L}_{ac}(G, D_c) + \lambda_1 \cdot \mathcal{L}_c(G) + \lambda_2 \cdot \mathcal{L}_n(G) \quad (6)$$

where λ_1 and λ_2 are the hyperparameters, our aim is to optimize the above equation.

2.2. Implementation details

2.2.1. Network architecture

We adopted the generator and discriminator networks architecture from Jun-Yan Zhu et al. [23]. The generator is a Resnet-15 network that we modified by adding one extra convolution block and one deconvolution block to manage the feature map extraction from the first three layers. The generator network contains nine residual blocks (including a reflection padding operation, an instance normalization layer, a 2-D convolution layer, a ReLu activation function, an instance normalization layer, a 2-D convolution layer, and a plus operation sequentially), four convolutional blocks (including a ReLu activation function, an instance normalization layer, and a 2-D convolution layer) and four deconvolutional blocks (including a ReLu activation function, an instance normalization layer, and a 2-D deconvolution layer). The structure of both discriminators is the same. It contains four convolutional blocks (including a ReLu activation function, an instance normalization layer, and a 2-D convolution layer) and a fully connected layer.

2.2.2. Training details

We used Adam optimizer [27] with batch size 1 to train the model to 200 epochs. The learning rate was kept at 0.0002 for the first 100 epochs, then linearly decreased to zero over the following 100 epochs. The λ_1 and λ_2 values in Eq. (6) were set to 10 for all experiments. The number of filters in the generator and discriminator's first convolutional layer was 64. Model training was done using two NVIDIA GeForce RTX 2080TI GPUs. These training details were applied to both architectures: the proposed model and CycleGAN.

2.2.3. Datasets

We retrospectively included carotid ultrasound 2D images from the longitudinal VIPVIZA study cohort of Umea University [28]. The study population are healthy subjects (50:50 men and women, 40–60 years old at base-line examination), with subclinical atherosclerosis. Expert operators acquired images in a longitudinal 2D B-mode of the carotid arteries. In the baseline and 3-year follow-up of the VIPVIZA study, a Cardio Health Station (CHS) and a linear 7 MHz transducer (Panasonic Healthcare Corporation of North America, Newark, NJ, USA) ultrasound machine ("US System A") were used. In the 6-year follow-up of the study, a GE Vivid IQ with the 9 MHz transducer (GE HealthCare, Chicago, IL, USA) ultrasound machine ("US System B") was used.

Experiment 1 – Image harmonization:. In this experiment, we included carotid plaque images from the two different systems (System A and System B). The task for the domain adaptation was to adapt the system A image features to resemble those of system B images while retaining the anatomical content. To robustly evaluate the effectiveness of the proposed model, we used images from different participants across the two domains. This approach ensures that the model's performance is not influenced by participant-specific features and instead demonstrates its generalizability and adaptability across distinct datasets. While images from the same subjects were available for domain a and b, they did not have the same 2D projection of the plaques, and therefore no pairvise approach e.g. pix-2-pix could be used. Instead the unparied D2D approach was chosen. Also, Using images from the same participants on both devices could potentially simplify the problem, as the model might rely on participant-specific patterns rather than truly adapting to domain differences. Furthermore, matching participants across domains could inadvertently compromise the challenge inherent to the task, undermining the objective of evaluating the model's robustness in a realistic and diverse scenario. We consider that this design choice aligns with the objectives of domain adaptation and ensures a fair and meaningful assessment of the proposed model's capabilities. For system A, the plaque acquisition mode was used, producing a B-mode image with no harmonic imaging and spatial angular compounding of three projections — here defined as "domain A". Images were reconstructed offline from the .cine format, and 493 consecutive images from different subjects were included from system A, and 546 for system B. For system B, standard B-mode acquisition mode was used, including harmonic imaging, and images were stored in dicom screen capture format — here defined as "domain B". Images had 8-bit intensity depth, and all images were resampled to 400 × 400 pixels, corresponding to a spatial resolution on the order of 0.1 mm/px.

Experiment 2 – Noise reduction:. In this experiment, we included carotid images of the intima-media complex (without plaques) from both systems A and B. For system A, the CHS IMT mode, which uses harmonic imaging for intima-media images and measurement, was used. For system B, conventional B-mode and harmonic imaging were used, and images were acquired to measure the intima-media thickness. Out of these images, noisy and clear images were manually classified by an expert operator (biomedical scientist and sonographer with +10 years of clinical ultrasound) - noisy and clear images were defined as "domain C" and "domain D", respectively. Noisy images were included based on the criteria of visual noise in the lumen of the images (a haze of bright pixels in the region of the blood), and clear images as its counterpart — with no noise in the lumen (dark pixels in the region of the blood). See Fig. 2 for examples. The task for the domain adaptation was to adapt the noisy image features to those of the clear images while retaining the anatomical content, Images had 8-bit intensity depth, and all images were resampled to 400 × 400 pixels, corresponding to a spatial resolution on the order of 0.1 mm/px. To evaluate the model, 15 test images from "domains C" were selected by an expert operator. While the test set size might appear small, these images were chosen to be representative of typical noise patterns encountered in clinical practice, providing a robust evaluation scenario. Importantly, the feature distribution of the test images is consistent with that of the training images and underwent the same preprocessing pipeline with mean(SD) values *BD 0.142 (0.043), HC 0.869 (0.037)*, and *BD 0.357 (0.152), HC 0.707 (0.098)* for training and test set, respectively.

2.3. Evaluation

The performance of the translation was evaluated from four perspectives: (1) their feature distribution similarity, (2) pixel-space similarity, (3) down-stream estimation of an image-based cardiovascular risk marker, and (4) (reverberation) noise level. For comparison, we also trained the default cycleGAN model.

Table 1
Description of the images used in the experiments. Here, GSM, CHS, GE IQ, and VIPVIZA stand for grey scale median, cardio health station, general electric image quality, and västerbotten intervention program visualization of atherosclerosis, respectively.

Characteristics	Experiment 1		Experiment 2	
	Domain A	Domain B	Domain C	Domain D
Anatomical content	Carotid plaques	Carotid plaques	Carotid intima-media	Carotid intima-media
Projection	Longitudinal	Longitudinal	Longitudinal	Longitudinal
US system	CHS (sys A)	GE IQ (sys B)	Mixed	Mixed
Source	VIPVIZA 3y FU	VIPVIZA 6y FU	Mixed	Mixed
Noise level	Mixed	Mixed	High	Low
N train	443	546	225	215
N test	50	–	15	–
Plaque GSM	44.6 (22.4)	48.6 (20.3)	–	–
Ultrasound measurements				
cIMT, mm	0.65 (0.9)	0.71 (0.8)	–	–
Min–max, mm	0.3–1.2	0.4–1.1	–	–
Plaque present	Yes	Yes	No	No
Clinical risk factors				
Low density lipo protein, (LDL)	3.51 (1.01)	3.83 (1.06)	–	–
High-density lipoprotein (HDL)	1.30 (0.39)	1.25 (0.39)	–	–
Systolic blood pressure, mmHg	127 (18)	133 (17)	–	–
Diastolic blood pressure, mmHg	83 (12)	83 (11)	–	–
Age, years	50.6 (7.6)	59.2 (8.2)	–	–
Sex, (female:men), %	50:50	55:45	–	–
Body-mass-index (BMI)	29 (5.6)	28 (5.0)	–	–
Framingham risk score (FRS)	10.4 (8.2)	10.7 (12.4)	–	–

The similarity in feature distributions was computed both without any translation and with translation, using the Bhattacharya distance (BD) and Histogram correlation (HC) [29] metrics. Mathematically, they are given by:

$$BD(h_1, h_2) = \sqrt{1 - \frac{1}{\sqrt{\bar{h}_1 \bar{h}_2 N^2}} \sum_I \sqrt{h_1(I) h_2(I)}} \tag{7}$$

$$HC(h_1, h_2) = \frac{\sum_I (h_1(I) - \bar{h}_1)(h_2(I) - \bar{h}_2)}{\sqrt{\sum_I (h_1(I) - \bar{h}_1)^2 \sum_I (h_2(I) - \bar{h}_2)^2}} \tag{8}$$

where h_1 and h_2 are the normalized histograms of the images. $\bar{h}_1$ and $\bar{h}_2$ represents the mean value of the histograms h_1 and h_2, respectively. N is the total number of bins in the histogram. I represent the background region pixels. High HC and low BD between images represent high feature distribution similarity.

The pixel-space similarity was computed using the structure similarity index measure (SSIM) between the input image and the corresponding translated image. The purpose was to evaluate the preservation of anatomical content. The SSIM metric is based on different regional pixel statistics in the image, and a high value means strong similarity. The SSIM was computed for the whole image and ROIs were selected by manually segmenting the plaque/intima-media tissues, the lumen, and the adventitia, by a US imaging expert and biomedical scientist, Christer Gronlund, who has over 20 years of experience in clinical ultrasound. These regions were then averaged to define the individual tissues of interest (e.g. Figs. 5 and 7).

To assess the impact of the translation on image-based risk markers, we computed the grey scale median (GSM), intima media grey scale median (IM-GSM), and plaque vulnerability re-classification rate. The GSM is a risk marker that is related to plaque vulnerability [30], where a low value is related to symptomatic (vulnerable) plaques and a high value to asymptomatic (stable) plaques. One threshold that has been used is GSM=25 [31]. Using this threshold, we computed the number of plaques that were re-classified, i.e. crossed this threshold due to the domain adaptation.

$$GSM = \text{median} \times \left(190 \cdot \left(I_{\text{ROI\_plaque}} - \text{mean}(I_{\text{ROI\_Lumen}})\right)\right) / \max(I_{\text{ROI\_Adventitia}}) \tag{9}$$

Finally, noise level was estimated using the lumen and wall ROIs as

$$\text{Contrast} = -20 \log_{10} \left(\frac{\text{mean}(I_{\text{ROI\_Lumen}})}{\text{mean}(I_{\text{ROI\_Adventitia}})} \right) \tag{10}$$

$$IM - GSM = \text{median} \times \left(190 \cdot \left(f(I_{\text{whole\_image}}) - \text{mean}(I_{\text{ROI\_Lumen}})\right)\right) / \max(I_{\text{ROI\_Adventitia}}) \tag{11}$$

where f indicates the contrast applied to $I_{\text{whole\_image}}$ for segmentation of intima media before the median calculation.

If there is low noise, the lumen should appear dark, and the $I_{\text{ROI\_lumen}}$ should have low values, causing the Contrast to be strongly negative in the dB scale.

3. Results

3.1. Experiment 1 – Image harmonization

Fig. 3 shows examples of two domain A images (cropped) and their corresponding translations using cycleGAN and the proposed GAN. It can be seen that the anatomical content of the input images is retained in the translated versions, but the fine-detailed texture pattern is modified. Evaluation of feature distribution similarity showed that both translation models improved the similarity, but the proposed model had the highest performance (Table 2 and Fig. 4). The proposed model achieved a lower BD and higher HC compared to CycleGAN and no translation, demonstrating that our model outperformed CycleGAN.

Fig. 5 shows an example of an input image, generated translation and the corresponding pixel-space similarity comparison using SSIM. Evaluation of pixel-space similarity showed that SSIM had higher values (higher similarity) for the proposed model than CycleGAN in all tissues except the lumen (Table 3 and Fig. 4).

The difference in plaque GSM between input and output images was 7.6 (6.5) vs. 15.0 (13.4) for the proposed and cycleGAN models, respectively (Table 4). In the training data, the difference in GSM plaque between domains B and A was 4.0 ($p < 0.006$, independent samples t-test). The contrast between the arterial wall (adventitia) and lumen ROIs was −34.1 (3.8) dB, −35.2 (4.1) dB and −35.8 (4.3) dB for the original image, cycleGAN, and proposed model translations (Table 4).

3.2. Experiment 2 – Noise reduction

Fig. 6 shows examples of two noisy domain C images (cropped) and their corresponding translations using cycleGAN and the proposed GAN, where noise in the lumen region was reduced.

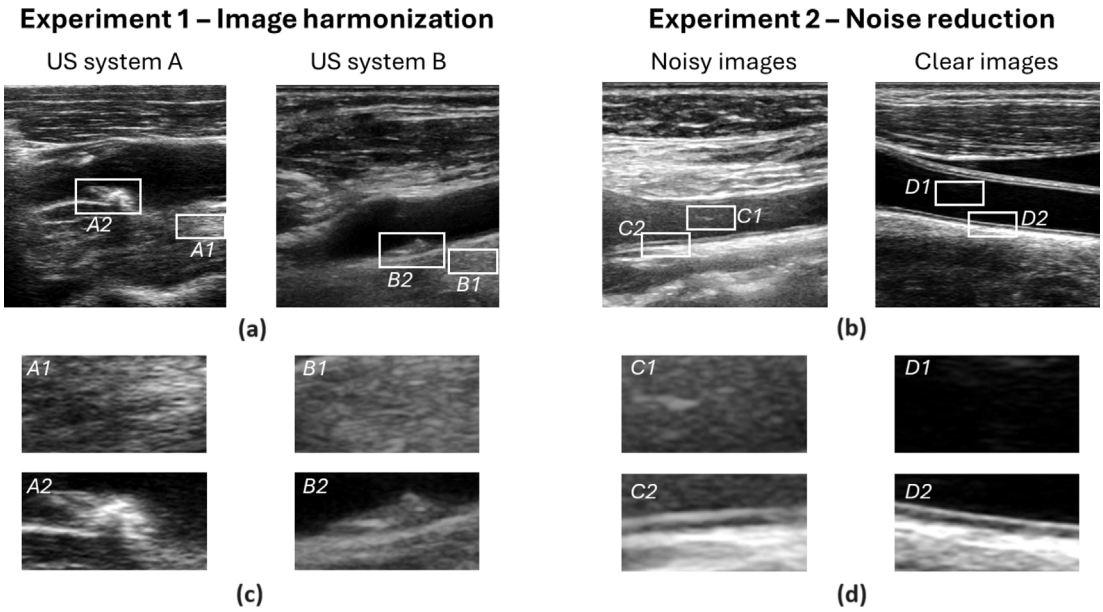

Fig. 2. Examples of carotid ultrasound images used in the training of the domain adaptation task for the two experiments. (a) show examples for experiment 1 (image harmonization) that we define as domains A and B, respectively. (b) show examples for experiment 2 (noise reduction) that we define as domains C and D. In the lower panel, (c) and (d), regions of interests (ROIs) from the upper images are shown to zoom in on a homogeneous tissue segment (A1 and B1), a carotid plaques (A2 and B2), Lumen/blood (C1 and D1), and arterial walls (C2 and D2).

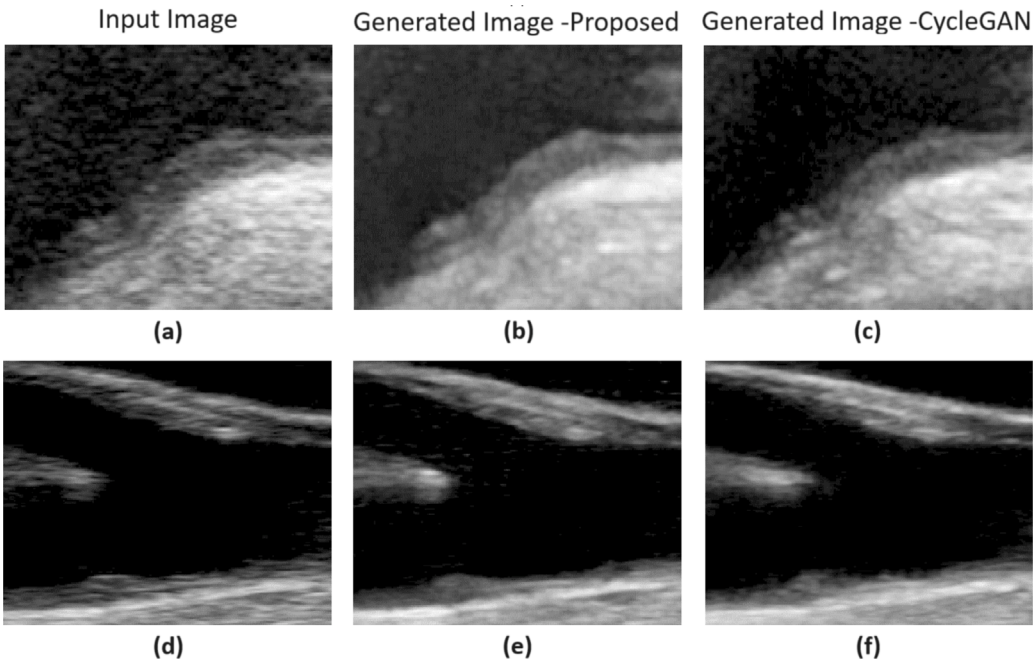

Fig. 3. Example of results for experiment 1 (Image harmonization). (a) and (d) shows input images (Domain A) and (b) and (e) shows the proposed model's generated images. (c) and (f) show the corresponding generated images of the cycleGAN model.

The feature similarity distribution metrics between input and translated images showed that both models improved the BD and HC similarities as compared to no translation and that the proposed model was the best (Table 5, Fig. 4). Fig. 7 shows example results for noise suppression in carotid ultrasound images by SSIM map and illustrates that the translation models strongly modify the lumen ROI while retaining the overall anatomy. The pixel-space similarity estimations (SSIM) in different tissues also showed that the lumen was strongly dissimilar with SSIM values close to 0, whereas tissues had strong similarity 0.7–0.8 for both models.

Evaluation of the influence of translation on risk marker showed that the intima-media GSM changed from 51.2 (13.9) to 42.6 (18.1) and 45.8 (17.7) for the cycleGAN and proposed GAN models, respectively. The contrast between adventitia and lumen ROIs changed from −23.5 (3.2) dB to −35.7 (2.8) dB and −46.7 (18.1) dB for the input image and the cycleGAN and proposed GAN translations, respectively. Note:

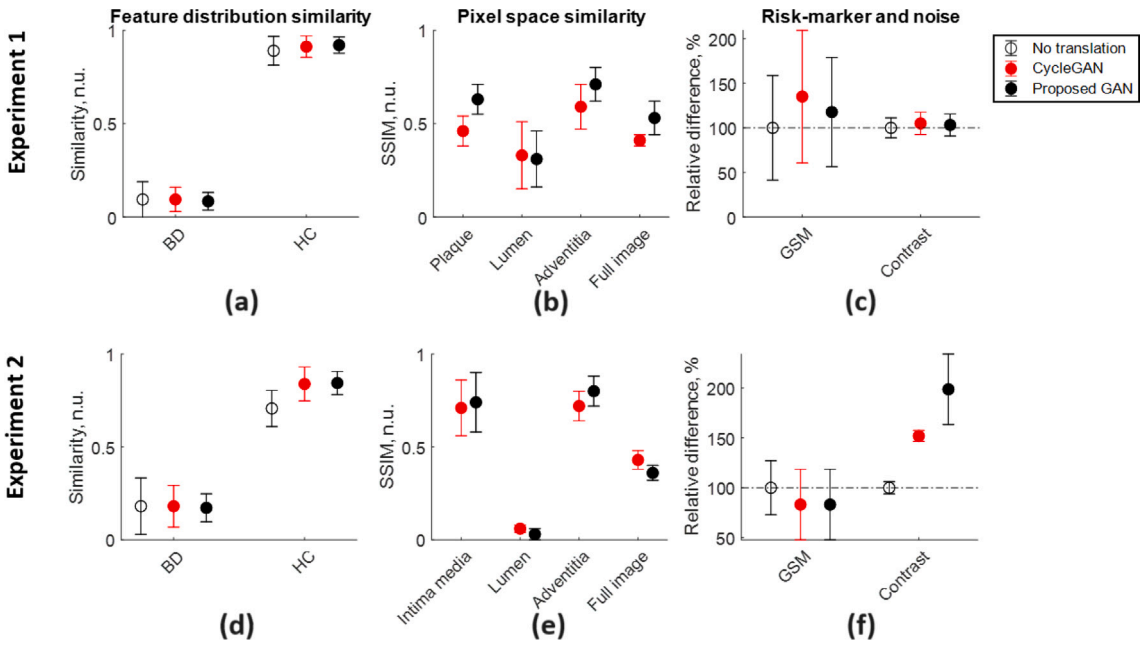

Fig. 4. Summary of results. The upper panel is for experiment 1 (image harmonization), and the lower panel is for experiment 2 (noise reduction). The evaluation was done in four perspectives: (1) feature distributions, (a) and (d), (2) pixel space (b) and (e), and (3) image-based risk marker for cardiovascular disease (GSM) and (4) noise level (Contrast), (c) and (f). Feature distributions were compared using Bhattacharyya distance (BD) and Histogram correlation (HC), and pixel-space similarity was evaluated using the structure similarity index measure (SSIM). The risk marker and contrast measures are presented as the relative change between the corresponding input and domain-adapted output images.

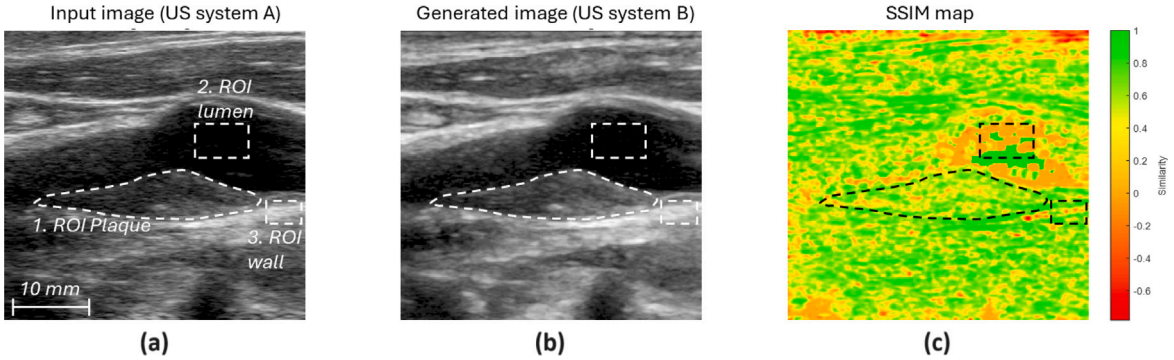

Fig. 5. Illustration of ROIs used in Experiment 1 to assess translation performance (Lumen, Plaque and Wall/Adventitia). (a) shows an example image of domain A, (b) the generated image (domain B) and (c) the corresponding structure similarity index measure (SSIM) map.

Table 2
Evaluation of feature transfer for Experiment 1 (Image Harmonization). To measure the dissimilarity and correlation between two domain images, Bhattacharyya distance (BD) and histogram correlation (HC) are used. The results include mean and standard deviation (SD) values of these metrics.

Method	Mean(SD)	
	BD	HC
No Translation A vs. A	0.044 (0.024)	0.958 (0.022)
No Translation B vs. B	0.033 (0.012)	0.968 (0.011)
No Translation A vs. B	0.120 (0.095)	0.890 (0.077)
CycleGAN A to B′	0.094 (0.065)	0.912 (0.057)
Proposed model A to B′	0.084 (0.047)	0.920 (0.043)

Table 3
Performance evaluation for pixel-space similarity for whole image and individual tissues — Experiment 1 (Image harmonization).

Model	Tissue segment			
	Plaque	Lumen	Adventitia	Whole image
Proposed GAN	0.63 (0.08)	0.31 (0.15)	0.71 (0.09)	0.53 (0.09)
CycleGAN	0.46 (0.08)	0.33 (0.18)	0.59 (0.12)	0.41 (0.03)

The proposed model had the most significant contrast, indicating the highest noise reduction.

4. Discussion

In this work, we proposed a GAN-based model for multi-purpose domain adaptation tasks and tested its performance for image harmonization between carotid ultrasound images captured with different ultrasound systems and image noise reduction in carotid ultrasound images with high noise. The main results show that (1) the model can adapt the image features between two different ultrasound systems, as well as reduce the noise in noisy images while still retaining the overall anatomy (content) in the images, and (2) the domain adaptation influenced critical risk markers for cardiovascular disease computed on adapted images.

4.1. Image harmonization

Results showed adapted feature distributions compared to those without adaptation (0.12 vs. 0.89) but not as similar to those within

Table 4
Performance evaluation — Impact on risk markers and blood-to-tissue contrast — Experiment 1 (Image harmonization). GSM stands for grey scale median. CI stands for confidence interval.

	A	B′	Diff	CI	p
	GSM				
Proposed GAN	42.7 (25.0)	50.3 (26.2)	7.6 (6.5)	[5.7 9.5]	P < 0.001 ***
CycleGAN	42.7 (25.0)	57.7 (31.8)	15.0 (13.4)	[11.0 19.0]	P < 0.001 ***
	Re-classification				
Proposed GAN	5 (10%)				
CycleGAN	12 (24%)				
	Contrast (dB)				
Proposed GAN	−34.1 (3.8)	−35.2 (4.1)		[−3.8 1.8]	P = 0.78
CycleGAN	−34.1 (3.8)	−35.8 (4.3)		[−2.4 0.0]	P = 0.73

Plaque area = 24.9 (18.0) mm^2.

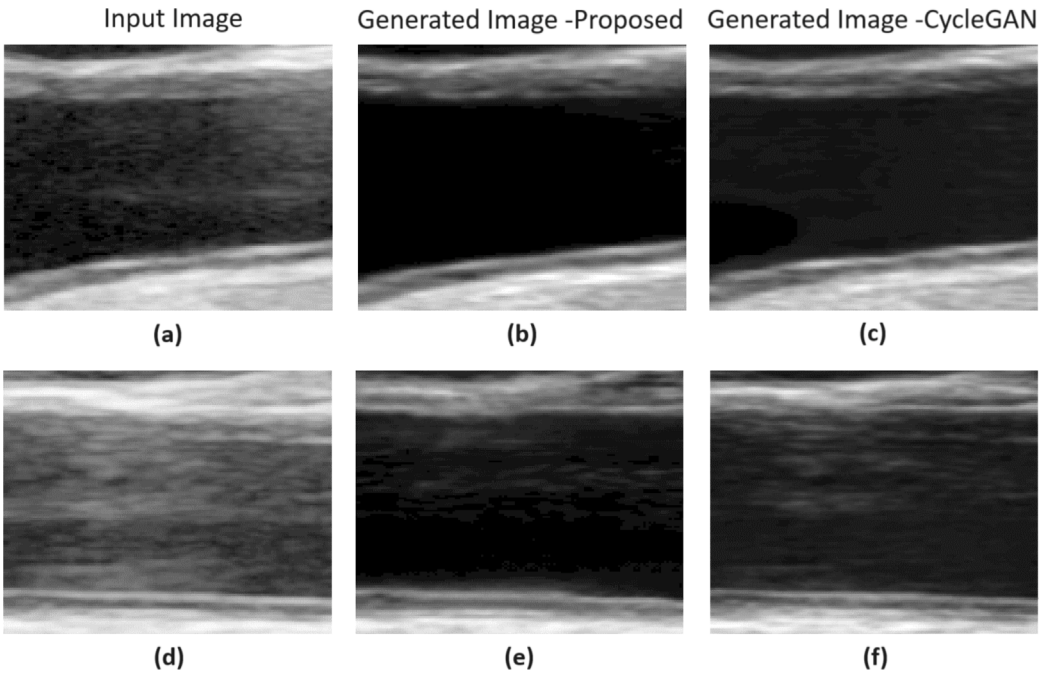

Fig. 6. Example of results for experiment 2 (Noise reduction). (a) and (d) shows input images (Domain C) and (b) and (e) shows the proposed model's generated images. (c) and (f) show the corresponding generated images of the cycleGAN model.

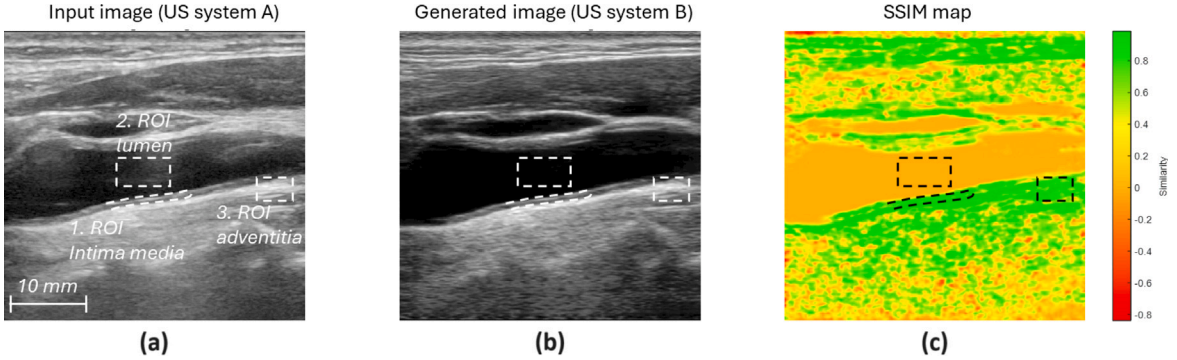

Fig. 7. Illustration of ROIs used in Experiment 2 to assess translation performance (Lumen, inner arterial wall (intima media) and outer wall (adventitia)). (a) shows an example image of domain C, (b) the generated image (domain D) and (c) the corresponding structure similarity index measure (SSIM) map.

domains (0.03 vs. 0.97) for BD and HC, respectively (Table 2). The feature similarities were higher than that of the NAGAN model [18], where they found 0.32 and 0.58 for BD and HC and 0.36 and 0.58 for the CycleGAN model. However, it should be noted that they adapted OCT and US images and, therefore, have larger differences between the image domains. The underlying content (anatomy) of the images was retained, as shown by the SSIM values (Table 3). For the whole image, SSIM was 0.53 (0.09). This is lower than those found by Ali et al. [11], translating US images of healthy and diseased arteries, with an overall SSIM of 0.78 (0.02). However, those images only included the common carotid artery, where the anatomy has less variation than the bulb. The bulb is the position where the carotid artery divides into two arteries,

Table 5
Evaluation of feature transfer for Experiment 2 (Noise Reduction). To measure the dissimilarity and correlation between two domain images, Bhattacharyya distance (BD) and histogram correlation (HC) are used. The results include mean and standard deviation (SD) values of these metrics.

Method	Mean(SD)	
	BD	HC
No Translation C vs. C	0.024 (0.008)	0.976 (0.007)
No Translation D vs. D	0.015 (0.008)	0.985 (0.007)
No Translation C vs. D	0.357 (0.152)	0.707 (0.098)
CycleGAN C to D'	0.181 (0.112)	0.839 (0.091)
Proposed model C to D'	0.172 (0.075)	0.844 (0.062)

Table 6
Performance evaluation for pixel-space similarity for whole image and individual tissues - Experiment 2 (Noise reduction).

Model	Tissue segment			
	IM	Lumen	Adventitia	Whole image
Proposed GAN	0.74 (0.16)	0.03 (0.03)	0.80 (0.08)	0.36 (0.04)
CycleGAN	0.71 (0.15)	0.06 (0.02)	0.72 (0.08)	0.43 (0.05)

the internal and external carotids. The latter was often included in the images of this study since this is where the plaques most commonly form. In addition, it should be noted that the SSIM were higher in our results for plaque and adventitia arterial wall segments SSIM were 0.63 and 0.71 as compared to the overall image — and these tissues are the ones that are used for assessing cardiovascular disease [30]. Thus, a lower preservation of the overall details of the image anatomy is likely not critical for our application.

4.2. Noise reduction

The primary source of noise in the ultrasound images is caused by diffuse reverberation which is caused by the multiple reflections of the transmitted ultrasound pulse in the tissue before returning to the transducer. The results of this work showed adapted feature distributions compared to those without adaptation (0.36 vs. 0.71) but not as similar to those within domains (0.02 vs. 0.98) for BD and HC, respectively (Table 5). The noise reduction in the image can be clearly seen in Fig. 6 and is supported by the low SSIM values in the lumen (0.03), while the arterial wall tissues are unaffected (SSIM 0.74 and 0.80), as shown in Table 6. In addition, the noise reduction was significant for both models, but the proposed model was better, as seen in the Contrast measurements changing from −23.5 dB to −46.7 dB and −35.7 dB for the proposed vs. CycleGAN models, respectively (Table 7). Recent work by other groups, e.g., by Brickson et al. [32], used a 3D CNN model on the underlying radio-frequency data of ultrasound images and showed a Contrast reduction from −18.9 (5.8) to −34.9 (8.1) dB. Thus, noise reduction performance in our work is similar and potentially better. The proposed GAN preserved the wall and adventitia tissues the best, as seen in SSIM for wall and adventitia (Table 6). It should be noted that it slightly distorted other parts of the image (see Fig. 7), e.g., muscle and vein, which is also indicated by a slightly lower overall SSIM compared to experiment 1. However, this may not be a critical problem for our application of assessing cardiovascular disease (as commented for experiment 1).

4.3. Influence on image-based computation of cardiovascular risk markers

In experiment 1, the impact of the domain adaptations on image-based risk markers showed that both models increased the GSM of the plaques, with a consequence of re-classification of 12 plaques (24%) for cycleGAN and 5 (10%) plaques for the proposed model. Thus, the proposed model preserves the computed risk marker integrity better. The SSIM was different in the lumen, adventitia, and plaque tissues, indicating that the model translated them differently (Eq. (11)). Since the GSM risk marker is computed from ROIs in the lumen, adventitia and arterial plaque, the former finding can likely explain this change in GSM. In comparison, our group previously found that the blood pressure variation over the cardiac cycle caused a 16% re-classification in the same study population [33], and 30% re-classification in another study population with asymptomatic atherosclerosis (more severe disease) [34]. This indicates that the re-classification by our model is similar to or lower than other known modulators of GSM.

In contrast, in experiment 2, the risk marker IM-GSM was used to assess the impact on clinical assessment, and it showed a slight decrease in the IM-GSM for both models but with large variability. The IM-GSM risk marker has not been as extensively studied in the literature as the GSM for plaques, and its relevance for assessing cardiovascular risk remains to be established. However, IM-GSM has been shown to be associated with the progression of atherosclerotic disease in the same study population [35], and to be associated with plaque GSM in populations with more advanced disease [36]. Since the GSM computations in experiments 1 and 2 are using the same tissue ROIs (Lumen and adventitial wall), and since the lumen was much more modified in experiment 2 compared to experiment 1, GSM estimations will be influenced differently in the two experiments. It should be noted that the increase in GSM of the plaques may be built into the models because the training target domain B had significantly higher GSM values of the plaques than source domain A (48 vs. 44, see Table 1). Furthermore, In the context of domain adaptation, the observed changes in GSM are not attributed to adaptation technique only, but rather an additive effect influenced by a combination of factors, including clinical risk factors, disease stage, vulnerability, etc. These factors collectively contribute to the overall change in GSM. Therefore, the impact of domain adaptation should be viewed as a collective result, rather than an adaptation technique only. Based on our results, we suggest that the impact on risk marker estimations, such as GSM and IM-GSM, should always be carefully evaluated when using domain adaptation models for carotid ultrasound images. This is important because different models can lead to varying changes in the computed risk markers.

4.4. The proposed model

The proposed model addressed both image harmonization and noise reduction, and results show that it performed better than the state-of-the-art model, cycleGAN, in both tasks. The proposed model was inspired by the methods [18,24], demonstrating the independent processing of feature representations learned by neural networks. In our model, we leverage the potential of Wasserstein distance to compute the noise loss using the style of feature maps driven from the early layers of the generator and the target domain images. The Wasserstein distance reduces dissimilarity by finding the most efficient way to map the distribution of one image set to another and accounting for the spatial relationships within the images by considering the geometry of the data in images [26,37]. We consider that this makes it more suitable for comparing images where the spatial relationship between pixels is crucial, such as in medical images (e.g., carotid ultrasound images). Our features transfer approach builds a domain-to-domain mapping for both image harmonization and noise reduction.

We conducted experiments with both Resnet and Unet architectures as generators and observed that Resnet outperformed Unet in our applications. Therefore, the final results are reported based on the use of Resnet. Our generator employs a ResNet-15 architecture, while the discriminator is constructed using a classical CNN. The ResNet holds an advantage over CNN due to its incorporation of residual units, enabling deeper layers that directly leverage information from shallower ones. Consequently, feature maps extracted from the ResNet-15 are likely to exhibit more refined patterns and smoother transitions compared to those from a standard CNN. In addition, we fine-tuned the generator architecture from Jun-Yan Zhu et al. [23] by adding two extra blocks

Table 7
Performance evaluation — Impact on riskmarkers and blood-to-tissue contrast — Experiment 2 (Noise reduction). GSM stands for grey scale median. CI stands for confidence interval.

Method	C	D′	Diff	CI	p
	GSM				
Proposed GAN	51.2 (13.9)	45.8 (17.7)	−5.3 (11.8)	[−11.9 1.2]	P = 0.10
CycleGAN	51.2 (13.9)	42.6 (18.1)	−8.6 (6.2)	[−12 −5.1]	P < 0.001 ***
	Contrast (dB)				
Proposed GAN	−23.5 (3.2)	−46.7 (18.1)		[−43.0 −22.4]	P < 0.001 ***
CycleGAN	−23.5 (3.2)	−35.7 (2.75)		[−14.9 −13.0]	P < 0.001 ***

(one convolution block and one deconvolution block) to manage the feature map extraction from the first three layers. We trained our model for various layers and decided to use the first three layers of the generator for stylistic feature extraction based on the training outcomes. For the proposed model, the training times were 7.85 h for experiment 1-Image harmonization and 2.32 h for experiment 2-Noise reduction while CycleGAN took 15.28 h for experiment 1-Image harmonization and 5.06 h for experiment 2-Noise reduction, approximate. These differences highlight the computational complexity, distinct architectures and learning processes inherent to each model. The proposed model demonstrates faster training times for both experiments. Code of the model is available at the GitHub repository.[1]

4.5. Limitations and future work

The proposed GAN model is designed for unsupervised image-to-image translation between two specific domains. It may not generalize well to tasks involving other domains or tasks with highly diverse datasets.

The features of carotid ultrasound images are influenced both by the specific imaging system (producing e.g. different spatial resolution and noise levels) as well as the severity of the disease (resulting in changes in the arterial wall, including plaques of various sizes and composition). While the present work and previous work by others show that adaptation between two different imaging systems can be done, the population is subclinical meaning that the amount of disease is relatively low, and it is not clear how different amounts of disease in the arteries affect the translation performance. A potential solution to this could be to, e.g. train the model with a dataset including a larger distribution of the amount of disease in the arteries.

The results of the image harmonization task are between two specific ultrasound systems and will not, by default, generalize to other US systems. For other pairs of US systems, the model needs to be retained. The noise reduction task was trained using images from both systems A and B, potentially allowing the model to learn more robust features (i.e., features not related to a specific system). Therefore, this model could potentially generalize to some degree to images of other ultrasound systems. The next step should be to train it on a larger set of images with low and high image quality and from multiple systems and use an external cohort for validation.

Moreover, in future work, the models could be trained using the risk marker data as input to the discriminator to minimize the change in risk markers that we found in this work.

The implementation of the model is a bit complex compared to traditional GAN architectures because it involves modifying the generator network to adjust harmonization features and noise levels between domains adaptively. This increased complexity can lead to higher computational overhead during training and inference.

[1] https://github.com/usama6832/Domain-Adaptation-Model-for-Carotid-Ultrasound.

5. Conclusion

This paper proposes a GAN-based domain adaptation model for image harmonization and noise reduction of carotid ultrasound images. The results show that the model can (1) adapt features of ultrasound images of one system to images from another and (2) reduce the noise in the lumen (blood), both without changing the anatomical contents of the images. The proposed model performed better than the state-of-the-art model, CycleGAN. In addition, the results showed that the downstream task of computing image-based risk markers for cardiovascular disease assessment was, in general, influenced by the domain adaptations. We conclude that domain adaptation is a powerful tool to improve image quality and harmonization of image features in carotid ultrasound images but the influence on computed risk markers should be evaluated for the specific model used.

CRediT authorship contribution statement

Mohd Usama: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Data curation, Conceptualization. **Emma Nyman:** Writing – review & editing. **Ulf Näslund:** Writing – review & editing. **Christer Grönlund:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Data curation, Conceptualization.

Ethics in publishing statement

All authors agree that:

This research presents an accurate account of the work performed, all data presented are accurate and methodologies detailed enough to permit others to replicate the work.

This manuscript represents entirely original works and or if work and/or words of others have been used, that this has been appropriately cited or quoted and permission has been obtained where necessary.

This material has not been published in whole or in part elsewhere.

The manuscript is not currently being considered for publication in another journal.

That generative AI and AI-assisted technologies have not been utilized in the writing process or if used, disclosed in the manuscript the use of AI and AI-assisted technologies and a statement will appear in the published work.

That generative AI and AI-assisted technologies have not been used to create or alter images unless specifically used as part of the research design where such use must be described in a reproducible manner in the methods section.

All authors have been personally and actively involved in substantive work leading to the manuscript and will hold themselves jointly and individually responsible for its content.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used SCISPACE in order to write some part of literature/exiting works, effectively. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

This submitted work is original, has not been published elsewhere, and is not considered for publication in any other journal. The final version of the manuscript is approved by all authors.

Acknowledgements

This work was supported by the Heart Foundation of Northern Sweden, the Kempe foundations, Sweden (JCK-3172), and Region Västerbotten, Sweden.

References

[1] Fisnik Jashari, Pranvera Ibrahimi, Gani Bajraktari, Christer Grönlund, Per Wester, Michael Y Henein, Carotid plaque echogenicity predicts cerebrovascular symptoms: a systematic review and meta-analysis, Eur. J. Neurol. 23 (7) (2016) 1241–1247.

[2] Xiao-xiao Chen, Zi-xiang Kong, Shu-fang Wei, Fei Liang, Ting Feng, Shanshan Wang, Jian-song Gao, Ultrasound lmaging-vulnerable plaque diagnostics: Automatic carotid plaque segmentation based on deep learning, J. Radiat. Res. Appl. Sci. 16 (3) (2023) 100598.

[3] Andrew Ward, Ashish Sarraju, Sukyung Chung, Jiang Li, Robert Harrington, Paul Heidenreich, Latha Palaniappan, David Scheinker, Fatima Rodriguez, Machine learning and atherosclerotic cardiovascular disease risk prediction in a multi-ethnic population, NPJ Digit. Med. 3 (1) (2020) 125.

[4] Zihao Fan, Zhi Du, Jinrong Fu, Ying Zhou, Pengyu Zhang, Chuning Shi, Yingxian Sun, Comparing the performance of machine learning and conventional models for predicting atherosclerotic cardiovascular disease in a general Chinese population, BMC Med. Inform. Decis. Mak. 23 (1) (2023) 134.

[5] Mehreen Mubashar, Hazrat Ali, Christer Grönlund, Shoaib Azmat, R2U++: a multiscale recurrent residual U-net with dense skip connections for medical image segmentation, Neural Comput. Appl. 34 (20) (2022) 17723–17739.

[6] T. Elatrozy, A. Nicolaides, T.H. Tegos, A.Z. Zarka, et al., The effect of B-mode ultrasonic image standardisation on the echodensity of symptomatic and asymptomatic carotid bifurcation plaques, Int. Angiol. 17 (3) (1998) 179.

[7] Christos P Loizou, Constantinos S Pattichis, Marios Pantziaris, T Tyllis, Andrew Nicolaides, Quality evaluation of ultrasound imaging in the carotid artery based on normalization and speckle reduction filtering, Med. Biol. Eng. Comput. 44 (2006) 414–426.

[8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio, Generative adversarial nets, Adv. Neural Inf. Process. Syst. 27 (2014).

[9] Zhihao Wang, Jian Chen, Steven C.H. Hoi, Deep learning for image super-resolution: A survey, IEEE Trans. Pattern Anal. Mach. Intell. 43 (10) (2020) 3365–3387.

[10] Hao Guan, Mingxia Liu, Domain adaptation for medical image analysis: a survey, IEEE Trans. Biomed. Eng. 69 (3) (2021) 1173–1185.

[11] Hazrat Ali, Emma Nyman, Ulf Näslund, Christer Grönlund, Translation of atherosclerotic disease features onto healthy carotid ultrasound images using domain-to-domain translation, Biomed. Signal Process. Control. 85 (2023) 104886.

[12] Tomoyuki Fujioka, Kazunori Kubota, Mio Mori, Leona Katsuta, Yuka Kikuchi, Koichiro Kimura, Mizuki Kimura, Mio Adachi, Goshi Oda, Tsuyoshi Nakagawa, et al., Virtual interpolation images of tumor development and growth on breast ultrasound image synthesis with deep convolutional generative adversarial networks, J. Ultrasound Med. 40 (1) (2021) 61–69.

[13] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, Clara I Sánchez, A survey on deep learning in medical image analysis, Med. Image Anal. 42 (2017) 60–88.

[14] Zixia Zhou, Yuanyuan Wang, Yi Guo, Yanxing Qi, Jinhua Yu, Image quality improvement of hand-held ultrasound devices with a two-stage generative adversarial network, IEEE Trans. Biomed. Eng. 67 (1) (2019) 298–311.

[15] Lihong Huang, Zixia Zhou, Yi Guo, Yuanyuan Wang, A stability-enhanced cyclegan for effective domain transformation of unpaired ultrasound images, Biomed. Signal Process. Control. 77 (2022) 103831.

[16] Menghua Xia, Hongbo Yang, Yanan Qu, Yi Guo, Guohui Zhou, Feng Zhang, Yuanyuan Wang, Multilevel structure-preserved GAN for domain adaptation in intravascular ultrasound analysis, Med. Image Anal. 82 (2022) 102614.

[17] Shreeram Athreya, Ashwath Radhachandran, Vedrana Ivezić, Vivek Sant, Corey W Arnold, William Speier, Ultrasound image enhancement using CycleGAN and perceptual loss, 2023, arXiv preprint arXiv:2312.11748.

[18] Tianyang Zhang, Jun Cheng, Huazhu Fu, Zaiwang Gu, Yuting Xiao, Kang Zhou, Shenghua Gao, Rui Zheng, Jiang Liu, Noise adaptation generative adversarial network for medical image analysis, IEEE Trans. Med. Imaging 39 (4) (2019) 1149–1159.

[19] Hee Guan Khor, Guochen Ning, Xinran Zhang, Hongen Liao, Ultrasound speckle reduction using wavelet-based generative adversarial network, IEEE J. Biomed. Heal. Inform. 26 (7) (2022) 3080–3091.

[20] Mercy N Asiedu, Alex R Benjamin, Vivek K Singh, Shuhang Wang, Kevin Wu, Anthony E Samir, Viksit S Kumar, A generative adversarial network for ultrasound signal enhancement by transforming low-voltage beamformed radio frequency data to high-voltage data, in: Medical Imaging 2022: Ultrasonic Imaging and Tomography, vol. 12038, SPIE, 2022, pp. 246–254.

[21] Silvia Seoni, Massimo Salvi, Giulia Matrone, Kristen M Meiburger, Ultrasound image beamforming optimization using a generative adversarial network, in: 2022 IEEE International Ultrasonics Symposium, IUS, IEEE, 2022, pp. 1–4.

[22] Lun Zhang, Junhua Zhang, Ultrasound image denoising using generative adversarial networks with residual dense connectivity and weighted joint loss, PeerJ Comput. Sci. 8 (2022) e873.

[23] Jun-Yan Zhu, Taesung Park, Phillip Isola, Alexei 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.

[24] Leon A. Gatys, Alexander S. Ecker, Matthias Bethge, Image style transfer using convolutional neural networks, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2414–2423.

[25] Yanghao Li, Naiyan Wang, Jiaying Liu, Xiaodi Hou, Demystifying neural style transfer, 2017, arXiv preprint arXiv:1701.01036.

[26] Martin Arjovsky, Soumith Chintala, Léon Bottou, Wasserstein generative adversarial networks, in: International Conference on Machine Learning, PMLR, 2017, pp. 214–223.

[27] Diederik P. Kingma, Jimmy Ba, Adam: A method for stochastic optimization, 2014, arXiv preprint arXiv:1412.6980.

[28] Ulf Näslund, Nawi Ng, Anna Lundgren, Eva Fhärm, Christer Grönlund, Helene Johansson, Bernt Lindahl, Bertil Lindahl, Kristina Lindvall, Stefan K Nilsson, et al., Visualization of asymptomatic atherosclerotic disease for optimum cardiovascular prevention (VIPVIZA): a pragmatic, open-label, randomised controlled trial, Lancet 393 (10167) (2019) 133–142.

[29] Anil Bhattacharyya, On a measure of divergence between two statistical populations defined by their probability distribution, Bull. Calcutta Math. Soc. 35 (1943) 99–110.

[30] AN Nicolaides, SK Kakkos, E Kyriacou, M Griffin, M Sabetai, DJ Thomas, T Tegos, G Geroulakos, N Labropoulos, CJ Doré, et al., Asymptomatic carotid stenosis and risk of stroke (ACSRS) study group. asymptomatic internal carotid artery stenosis and cerebrovascular risk stratification, J. Vasc. Surg. 52 (6) (2010) 1486–1496.

[31] Christina I Christodoulou, Constantinos S Pattichis, Marios Pantziaris, Andrew Nicolaides, Texture-based classification of atherosclerotic carotid plaques, IEEE Trans. Med. Imaging 22 (7) (2003) 902–912.

[32] Leandra L. Brickson, Dongwoon Hyun, Marko Jakovljevic, Jeremy J. Dahl, Reverberation noise suppression in ultrasound channel signals using a 3D fully convolutional neural network, IEEE Trans. Med. Imaging 40 (4) (2021) 1184–1195, http://dx.doi.org/10.1109/TMI.2021.3049307.

[33] Emma Nyman, Per Lindqvist, Ulf Näslund, Christer Grönlund, Risk marker variability in subclinical carotid plaques based on ultrasound is influenced by cardiac phase, echogenicity and size, Ultrasound Med. Biol. 44 (8) (2018) 1742–1750.

[34] Isak Stenudd, Elias Sjödin, Emma Nyman, Per Wester, Elias Johansson, Christer Grönlund, Ultrasound risk marker variability in symptomatic carotid plaque: impact on risk reclassification and association with temporal variation pattern, Int. J. Cardiovasc. Imaging 36 (2020) 1061–1068.

[35] Emma Nyman, Per Liv, Per Wester, Ulf Näslund, Christer Grönlund, Carotid wall echogenicity at baseline associates with accelerated vascular aging in a middle-aged population, Int. J. Cardiovasc. Imaging 39 (3) (2023) 575–583.

[36] Pranvera Ibrahimi, Fisnik Jashari, Elias Johansson, Christer Grönlund, Gani Bajraktari, Per Wester, Michael Y Henein, Common carotid intima–media features determine distal disease phenotype and vulnerability in asymptomatic patients, Int. J. Cardiol. 196 (2015) 22–28.

[37] Yonatan Dukler, Wuchen Li, Alex Lin, Guido Montúfar, Wasserstein of Wasserstein loss for learning generative models, in: International Conference on Machine Learning, PMLR, 2019, pp. 1716–1725.