



Normalization of breast MRIs using cycle-consistent generative adversarial networks

Gourav Modanwal^{a,*}, Adithya Vellal^b, Maciej A. Mazurowski^a

^a Department of Radiology, Duke University, Durham, NC, USA

^b Department of Computer Science, Duke University, Durham, NC, USA

ARTICLE INFO

Article history:

Received 25 February 2021

Accepted 29 May 2021

Keywords:

MRI intensity normalization

Medical image translation

Deep learning

CycleGAN

Vendor normalization

ABSTRACT

Objectives: Dynamic Contrast Enhanced-Magnetic Resonance Imaging (DCE-MRI) is widely used to complement ultrasound examinations and x-ray mammography for early detection and diagnosis of breast cancer. However, images generated by various MRI scanners (e.g., GE Healthcare, and Siemens) differ both in intensity and noise distribution, preventing algorithms trained on MRIs from one scanner to generalize to data from other scanners. In this work, we propose a method to solve this problem by normalizing images between various scanners.

Methods: MRI normalization is challenging because it requires normalizing intensity values and mapping noise distributions between scanners. We utilize a cycle-consistent generative adversarial network to learn a bidirectional mapping and perform normalization between MRIs produced by GE Healthcare and Siemens scanners in an unpaired setting. Initial experiments demonstrate that the traditional CycleGAN architecture struggles to preserve the anatomical structures of the breast during normalization. Thus, we propose two technical innovations in order to preserve both the shape of the breast as well as the tissue structures within the breast. First, we incorporate mutual information loss during training in order to ensure anatomical consistency. Second, we propose a modified discriminator architecture that utilizes a smaller field-of-view to ensure the preservation of finer details in the breast tissue.

Results: Quantitative and qualitative evaluations show that the second innovation consistently preserves the breast shape and tissue structures while also performing the proper intensity normalization and noise distribution mapping.

Conclusion: Our results demonstrate that the proposed model can successfully learn a bidirectional mapping and perform normalization between MRIs produced by different vendors, potentially enabling improved diagnosis and detection of breast cancer. All the data used in this study are publicly available at <https://wiki.cancerimagingarchive.net/pages/viewpage.action?pageId=70226903>.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Breast cancer is one of the leading causes of death among women around the globe [1]. Dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) is widely used to complement mammography and ultrasound when evaluating breast cancer, particularly when assessing the extent of cancer before surgery [2]. In some high-risk cases, it is also used for screening.

A significant challenge related to the use of DCE-MRI is the lack of standardized imaging protocols [3,4]. Different MRI scanners use different parameters, which previous research [5] has shown to

drastically alter image appearance, quality, as well as the radiomics analysis. When the same patient is imaged using a different scanner or even the same scanner with different scanner parameters, the produced MR images may vary significantly [6,7]. The inconsistencies present in the radio-frequency (RF) coil produce intensity variations in the underlying tissue across the scanned image [8]. Additionally, varying scanner parameters alter the noise distribution of the images. An illustration of the difference in intensity and noise distribution between images obtained from two different MRI scanner manufacturers (GE Healthcare and Siemens) is shown in Fig. 1.

The high degree of inter-scanner variation proves to be a significant obstacle to the effective usage of DCE-MRI. In the context of radiomics, where a multitude of features are extracted from im-

* Corresponding author.

E-mail address: gourav.modanwal09@gmail.com (G. Modanwal).




Fig. 1. Example of images from two scanner displaying differences in intensity and noise distribution (a) GE Healthcare (b) Siemens.

ages for further processing, the features from different modalities may turn out to be incomparable, thus rendering them useless for classification and prediction. The impact of scanner parameters on breast MRI radiomic features is demonstrated by [5]. Variability in images has been shown to have an impact on the training of deep learning as well [9]. Algorithms trained on images from one scanner may not perform well on exams at a different institution that was acquired using a different scanner [10]. Finally, the inconsistency between images from different scanners may affect the outcome of computer-aided diagnosis. The ability to translate and normalize between images acquired by different vendors with varying parameters of scanner would have tremendous positive consequences. It would enable quantitative comparison of image features across various institutions. It would also improve generalization as deep models trained on one dataset could still perform inference on new datasets generated by different scanners.

In order to address this issue, we frame the problem of normalization between images generated by different MRI scanners as an application of unpaired image-to-image translation. Most of the literature in the domain of MRI pre-processing has focused on normalizing intensities but does not account for noise patterns. To our knowledge, no one has yet proposed a method for MRI vendor normalization. This process is challenging because it requires both normalizing the intensity and learn the mapping between the noise patterns. In this work, we present a vendor normalization method that attempts to perform intensity normalization as well as noise distribution mapping between MRIs obtained from different scanners. The significant contributions of this work can be summarized as follows:

- We present a method for MRI vendor normalization that performs unpaired bidirectional normalization between DCE-MRIs produced by different scanner models.
- We investigate the challenges of the standard CycleGAN approach for normalization of medical images, primarily the difficulty in maintaining the breast shape and structures within the breast between the original image and the normalized image. Then, we propose and evaluate two technical solutions to this issue, as described below.
 - We propose the incorporation of a mutual information loss with the standard CycleGAN architecture in order to ensure that the breast shape and tissue structures within the breast is maintained.
 - We propose a modified discriminator capable of preserving the breast shape as well as the dense tissues and evaluate the effect of changing the field-of-view on the performance.

We further present and compare the performance of the proposed vendor normalization methods using both quantitative and qualitative approaches. We also highlight how the proposed work can potentially enable the synthesis of larger and richer datasets that mitigate issues related to class imbalance.

The remainder of the article is organized as follows. Section 2 describes the related work. Section 3 presents details about the dataset, and the proposed methods are detailed in Section 4. Information about training is furnished in Sections 5 and 6 presents the metrics for the evaluation of the proposed method. Section 7 reports the experimental results and discussions. Finally, Section 8 concludes with a summary.

2. Related work

Unlike other imaging modalities, MRIs span a wide, non-linear spectrum of raw intensity values. They lack uniformity and often exhibit high variance between subjects. Even within a single subject, intensity variations of 10–40% have been observed [11]. This heterogeneity makes it difficult to effectively train robust medical image analysis algorithms on MRIs [12].

In response, various statistical approaches have been proposed for MRI intensity normalization. These include histogram equalization [6,13], intensity scaling based on regions of interest [14] and landmarks [15]. However, histogram-based methods rely on discrete approximations of intensity distributions, leading to high levels of inexactness [16]. Meanwhile, obtaining a high level of accuracy with landmark-based algorithms requires obtaining multiple landmarks from various tissue types in the image. Designing algorithms to perform this landmark selection task is difficult and time-consuming [15]. Another limitation of many MRI normalization methods [17,18] is that they require auxiliary inputs such as segmentation masks. This adds an intrinsic reliance on the models that perform these preprocessing tasks. Alternatively, some techniques [19] attempt to leverage the physics of MR acquisition in order to develop intensity invariant segmentation algorithms. However, using this type of approach requires integrating explicit physics-based embedding into the segmentation algorithm, thus limiting this system's ability to generalize to other downstream tasks.

Additionally, some of the methods discussed above [6,14,15] attempt to perform intensity transformation between two fixed imaging settings. That is, they make the assumption that the intensity relationship of the tissues is constant between the target group and the reference group, which is not always true [20]. If the intensity standardization needs to be done for images coming from multiple centers, multiple transforming models need to be established. Resultantly, these methods do not have the ability to process new images that are not from an MR image group that has already been included in their training data. This severely limits its usability. Work [21] presents a model based method for harmonization between patients scanned with differences in imaging parameters.

Recently, GANs have been used in a variety of applications to the domain of medical imaging. An excellent review of GANs' recent applications to the medical domain was presented by [22–24]. Most of the previous work [25–28] has focused on using GANs for multimodal translation that in turn, improved diagnosis across several modalities (e.g., ultrasound, PET, CT, and MRI). Additionally, GANs have successfully generated synthetic images [29–34] to augment training datasets for algorithms that perform downstream tasks—diagnosis, prognosis, segmentation, and registration. Lately, GANs have also been used for normalizing MRIs across different scanners. Work by [20] proposed standardization method for brain MRIs using GANs with a weak paired data strategy with focus on intensity normalization only. Another work [30] applies CycleGAN

Table 1
Distribution of patients & typical scanner parameters across different manufacturer.

Vendor/Manufacturer	Scanner model	Acquisition matrix	Echo time	Repetition time	# Subjects
GE Healthcare	Signa HDx	340×340	2.3 ms	5 ms	43
		340×340	2.4 ms	5.5 ms	33
	Signa HDxt	384×360	2.4 ms	5.5 ms	1
Siemens	MAGNETOM Avanto	448×448	1.4 ms	4.1 ms	30
		320×320	1.4 ms	4.1 ms	17

Table 2
Details about training, test and validation dataset.

	GE Healthcare (GE)	Siemens (SE)
Train Set	5045	2776
Test Set	1563	843
Validation Set	173	93

to brain MRI normalization. However, CycleGAN has an intrinsic ambiguity with respect to geometric transformations [35]. More specifically, since the anatomical structure of the images in a set of patient data is highly variable, CycleGAN is unlikely to utilize anatomical features in order to determine the realness of an image. However, algorithms for multi-modal translation and synthesis of medical images should ensure shape consistency, as these anatomical structures are crucial information for computer-based cancer detection algorithms.

One of the research work by [35] has tried to solve the problem by adding an extra penalty based on a segmentation mask generated from the CycleGAN output. However, this requires a ground-truth annotation of the dense tissue regions in the breast, which is not available in a typical use case for breast MRI normalization. Meanwhile, [36] introduce deformable convolutional layers and novel cycle-consistency losses. Other papers [25–28] utilize CycleGAN for translation between various modalities (e.g., ultrasound PET, CT and MRI) and reports many of the same issues discussed above. In this article, we present a fully unpaired algorithm for image normalization using CycleGAN. We propose and evaluate two technical solutions in order to effectively preserve the breast shape and tissue structures within the breast MRI.

3. Dataset and pre-processing

In this study, we utilize Duke-Breast-Cancer-MRI data¹ obtained from GE Healthcare (GE) and Siemens (SE) scanners (1.5 T) in the axial plane. Our dataset consists of 124 subjects: 77 imaged with a GE Healthcare scanner, and the remaining 47 with a Siemens scanner. Details about the distribution of patients & scanner parameters across different manufacturer is presented in Table 1. Each MR volume contains more than 160 2D axial image slices. The top 1% of pixel values in the entire dataset are truncated at 255, and the remaining intensities are linearly scaled to the 0–255 pixel range. The dataset are randomly divided into train, validation, and test set respectively at the patient level. We only use slices from the middle 50% of each patient volume throughout our experiment. Details regarding the number of slices used for training, testing, and validation are given in Table 2.

¹ All images used in this study are publicly available at <https://wiki.cancerimagingarchive.net/pages/viewpage.action?pageId=70226903>

4. Methods

In this section, we present various frameworks to perform normalization between MRI images acquired by GE Healthcare (GE) and Siemens (SE) scanners.

4.1. CycleGAN

We utilize the CycleGAN [37]—a bidirectional image-to-image translation method—for the normalization between the GE and SE MRIs. It consists of two generators (G_1, G_2) and two discriminators (D_1, D_2). Each generator has a corresponding discriminator, and they are trained in an adversarial setting in which the two networks compete against each other to fool their counterparts. Fig. 2 illustrates the CycleGAN network configuration where, I_1 and I_2 are training samples from $P_{data}(GE)$ and $P_{data}(SE)$, respectively. Generator G_1 normalizes from $GE \rightarrow SE$ while G_2 normalizes from $SE \rightarrow GE$.

The discriminator network D_1 discriminates between the images generated by the generator $G_1(I_1)$ and the target image I_2 while generator G_1 tries to improve the quality of the transformed image so that it can fool the discriminator. Similarly, D_2 discriminates between images generated by $G_2(I_2)$ and the target image I_1 , while G_2 tries to transform I_2 effectively enough to fool D_2 . The above task is formulated as a min-max optimization problem.

4.1.1. Network architectures

The architecture for the generators is adapted from [38]. The generator consists of an encoder, transformer, and decoder. The encoder uses convolutional down-sampling to shrink the size of the input representation and increase the number of channels. It is followed by a transformation block which retains the size of representation using residual convolution blocks. Finally, a decoder block is used which upsamples the size of representation using deconvolution.

The discriminator network uses a classical PatchGAN [39]. It is a fully convolutional neural network that processes overlapping patches of the input image instead of the entire input image. The output of the discriminator is a matrix of binary classifications of whether each patch is real or fake. A standard PatchGAN has a field of view (FOV), or patch size, of 70×70 . Our experiments with discriminator architectures with varying FOV are detailed in Section 7.

4.1.2. Losses

The objective function contains two loss terms: adversarial loss (L_{adv}) and cyclic loss (L_{cyc}). The adversarial loss [40] ensures that the generated images belong to the data distribution of the target domain. The adversarial loss is formulated as below:

$$L_{adv}(G_1, D_1, I_1, I_2) = \mathbb{E}_{I_2 \sim P_{data}(SE)} [(D_1(I_2) - 1)^2] + \mathbb{E}_{I_1 \sim P_{data}(GE)} [(D_1(G_1(I_1)))^2] \quad (1)$$




Fig. 2. CycleGAN network configuration for breast MRI normalization with two generators G_1 : normalizes $P_{data}(GE) \rightarrow P_{data}(SE)$ and G_2 normalizes $P_{data}(SE) \rightarrow P_{data}(GE)$, and associated adversarial discriminators D_1, D_2 . I_1 and I_2 are the unpaired training samples from $P_{data}(GE)$ and $P_{data}(SE)$, respectively.

$$L_{adv}(G_2, D_2, I_2, I_1) = \mathbb{E}_{I_1 \sim P_{data}(GE)} [(D_2(I_1) - 1)^2] + \mathbb{E}_{I_2 \sim P_{data}(SE)} [(D_2(G_2(I_2)))^2] \quad (2)$$

The generator tries to minimize the above adversarial loss, and the discriminator tries to maximize it. However, the adversarial loss alone is not sufficient enough to produce good target images. The adversarial loss will enforce the transformed output to be of the appropriate domain, but will not enforce the input and output to be recognizably the same. Thus an additional cycle-consistency loss is added to the overall objective. The cycle-consistency loss ensures that the translated image looks like the input image by enforcing G_1 and G_2 to be inverses of each other i.e. $(G_2(G_1(I_1))) \approx I_1$ and $(G_1(G_2(I_2))) \approx I_2$.

$$L_{cyc}(G_1, G_2) = \mathbb{E}_{I_1 \sim P_{data}(GE)} [\|G_2(G_1(I_1)) - I_1\|_1] \quad (3)$$

$$L_{cyc}(G_2, G_1) = \mathbb{E}_{I_2 \sim P_{data}(SE)} [\|G_1(G_2(I_2)) - I_2\|_1] \quad (4)$$

The overall objective is given as below where λ_{cyc} is the weighting factor for cycle-consistency loss.

$$L(G_1, G_2, D_1, D_2) = L_{adv}(G_1, D_1, I_1, I_2) + L_{adv}(G_2, D_2, I_2, I_1) + \lambda_{cyc} * (L_{cyc}(G_1, G_2) + L_{cyc}(G_2, G_1)) \quad (5)$$

4.2. CycleGAN with mutual information

The standard CycleGAN architecture detailed above, when used for normalization between GE and SE breast MRIs, may produce results that are unable to preserve the breast shape and tissue characteristics. In order to preserve the breast shape and tissue characteristics, we propose to utilize mutual information maximization between the real images and the generated images, as shown in Fig. 3. Our rationale is that while the intensity and texture of the image may change, high mutual information will indicate that the shape of the breast and the structure of dense tissue remained the same, which is desired in our application.

In practice, estimation of mutual information in images is challenging as we only have access to samples rather than the underlying distributions [41,42]. Additionally, previous sample-based estimators are brittle and do not scale well to higher dimensions [43]. Recently, Mutual Information Neural Estimation (MINE) [44] was introduced to approximate the mutual information using observed samples even when the true distribution is unknown. Their approach also scales to higher dimensions. Hence, we adopt their




Fig. 3. An illustration of the proposed mutual information loss.

method to estimate and maximize mutual information and utilize mutual information as a loss along with adversarial and cycle-consistency loss.

The mutual information is equivalent to the Kullback-Leibler (KL) divergence between the joint distribution, $P(X, Z)$, and the product of the marginal distributions $P(X)$ and $P(Z)$, as expressed below

$$I(X, Z) = D_{KL}(P_{XZ} || P_X \otimes P_Z) \quad (6)$$

where D_{KL} is defined as,

$$D_{KL}(P || Q) := \mathbb{E}_P \left[\log \frac{\partial P}{\partial Q} \right] \quad (7)$$

It uses the Donsker-Varadhan (DV) representation [45] of KL divergence, which leads to the following definition of approximate mutual information:

$$I_\phi(X, Z) = \sup_{\theta \in \mathcal{Z}} [\mathbb{E}_{P_{XZ}} [T_\theta] - \log(\mathbb{E}_{P_X \otimes P_Z} [e^{T_\theta}])] \quad (8)$$

The approximate mutual information $I_\phi(X, Z)$ is obtained by maximizing the lower bound of the objective function shown in eq. 8. The maximization is achieved by using a neural network (T_θ) with parameters θ . The neural network (T_θ) is optimized using gradient descent to characterise a family of functions which ultimately maximizes the lower bound of the above objective.




Fig. 4. Proposed discriminator architecture.

To enforce and preserve breast shape and tissue characteristics, we propose to include mutual information as a loss in the overall objective as specified below.

$$\begin{aligned}
 L(G_1, G_2, D_1, D_2) = & L_{adv}(G_1, D_1, I_1, I_2) + L_{adv}(G_2, D_2, I_2, I_1) \\
 & + \lambda_{cyc} * (L_{cyc}(G_1, G_2) + L_{cyc}(G_2, G_1)) \\
 & - \lambda_{mut} * (L_{mut}(I_1, G_1(I_1)) + L_{mut}(I_2, G_2(I_2)))
 \end{aligned} \quad (9)$$

where λ_{mut} is the weight factor for mutual information loss.

4.3. CycleGAN with modified discriminator

We also modify our discriminator to test the effects of varying fields of view (FOV)—the size of the input pixel window that contributes to a single pixel in the output map. As suggested in [39], we focus on smaller FOV to encourage the transformation learned by the generator to maintain sharp, high-frequency detail which is required in order to adequately preserve both the overall structure of the breast and the structure of the dense tissue regions inside the breast. Preliminary results were presented in [46]. Our experiments demonstrate that a (34×34) FOV discriminator architecture (shown in Fig. 4) is better at preserving morphological features of the breast tissue in comparison to the original (70×70) FOV. Quantitative analysis is presented in Table 3. Further details about the various discriminator architectures corresponding to different FOV are presented in appendix Appendix A.

4.4. CycleGAN with modified discriminator + mutual information

We also test the proposed Mutual Information loss in conjunction with the discriminator modification highlighted above.

5. Training

We optimize the network using mean squared error (MSE) instead of cross-entropy, as suggested in [47]. As a result, training becomes more stable, and higher quality images are produced. Additionally, to prevent the model from oscillation, the discriminator is fed a history of the 50 most recently generated images rather than solely the most recently generated image. Adam optimizer with the parameters $lr = 0.002$, $\beta_1 = 0.5$, and $\beta_2 = 0.999$ is used to train the network weights.

With the addition of mutual information loss, the proposed framework has additional parameters to optimize. During experiments, we found that normalization quality is susceptible to these parameters. We optimized CycleGAN for different values of λ_{cyc} (please refer Table A.5). The model with an optimal λ_{cyc} (Std. CycleGAN) was used as a baseline for comparison. We then experimented for optimal λ_{mut} (please refer Table A.6) and compared it with baseline model. The value of λ_{cyc} and λ_{mut} used in the experiments are 5.0 and 0.5 respectively.

6. Evaluation metrics

Quantitative evaluation of the normalized images is difficult in the case of unpaired images [37] as there is no standard/universal metric for assessing accuracy [48]. Hence, evaluating the quality of synthesized images is an open and challenging problem for which metrics vary depending on the specific needs of the application. Most previously published work relies either on the visual examination of the transformed images by human subjects or some application-specific metrics. Visual evaluation of the transformed image is still the most common and intuitive method for determining the quality of the transformed images.

In this work, the evaluation of our algorithms is done in two ways. First, we perform a combination of quantitative and qualitative analyses to determine the robustness of the normalization. For the quantitative analysis, we manually annotate a breast mask for 20 images both before normalization and after normalization for GE to Siemens as well as Siemens to GE, respectively. We then compute the Dice coefficient between these annotations. A higher Dice coefficient suggests that the normalization successfully preserved breast shape, while a lower value indicates distortion in breast shape. To evaluate the preservation of dense tissue, we perform qualitative analysis through visual observation.

Secondly, we evaluate the intensity normalization by manually annotating the dense tissue (10 cases) and subsequently computing the mean intensity value before and after normalization. The expectation is that while the mean intensities of dense tissue differ significantly between GE and Siemens before the normalization, they should be similar after the normalization.

7. Results and discussion

The result of the proposed MRI normalization using CycleGAN is presented in Fig. 5. Qualitatively, it can be observed that the standard CycleGAN model is unable to preserve the shape of the breast and dense tissue. Our proposed modified discriminator framework performed the best out of all explored algorithms.

A surprising result visible in Fig. 5 is that the introduction of mutual information loss is unable to preserve the shape of the breast. After further analysis, we determine that the noise pattern in the GE images is the primary cause of this failure. Specifically, the mutual information neural estimator (MINE) network tries to maximize the mutual information by matching the shape of the breast to the noisy “halo” around the breast, and in doing so, actually increases the size of the breast. Similarly, for the Siemens to GE normalization, it maximizes the mutual information by decreasing the shape of the breast. This is illustrated in Fig. 6.

We propose a modified CycleGAN framework that involves altering the discriminator architecture in order to put more stress on features pertaining to breast tissue. We experiment with various FOV in the discriminator architecture and present the effects that these changes have on performance in Fig. 7. It can be observed that the 70×70 FOV frequently modifies the dense tissues of the breast. It also modifies the shape of the breast, which is apparent from the lower Dice coefficients (Table 3). A 1×1 FOV, i.e. PixelGAN, has no effect on spatial statistics and is thus unable to learn the mapping between the noise distributions of the two domains. Additionally, the normalized images look extremely pixelated and exhibit a checkerboard pattern. The performance of a 45×45 FOV is comparatively better than the 70×70 FOV in terms of preserving both the breast shape as well as the dense tissue structures. However, visual inspection leads us to conclude that the 34×34 FOV discriminator preserves the dense tissue better and produces




Fig. 5. Representative results of the proposed image normalization (a) GE Healthcare to Siemens (b) Siemens to GE Healthcare.




Fig. 6. Effect of mutual information loss (a) GE Healthcare to Siemens (b) Siemens to GE Healthcare. Difference between the input and normalized image is shown using a composite image, where magenta shows negative value and green shows the positive value.




Fig. 7. Experiment with the field of view (FOV) in the discriminator architecture (a) GE Healthcare to Siemens (b) Siemens to GE Healthcare. Red ellipse shows the change in breast shape and dense tissue structure as compared to input.

sharper images compared to the 45×45 FOV. The Dice coefficients confirm that the 34×34 FOV is able to preserve the shape of the breast as well. This improved performance, along with its lower number of parameters, lead us to select the 34×34 FOV discriminator architecture.

Quantitative results are presented in Table 4. It can be observed that during GE to Siemens normalization, the Dice coefficient of the breast masks is the highest for the CycleGAN framework obtained by modifying the discriminator architecture. It is also apparent from Table 4 that applying the mutual information loss to

Table 3
Dice coefficients between breast mask before and after normalization obtained on validation data.

FOV	GE to SE		SE to GE	
	Mean	Std	Mean	Std
1×1	–	–	–	–
34×34	0.9762	0.0091	0.9794	0.0070
45×45	0.9236	0.0164	0.9310	0.0177
70×70	0.9138	0.0577	0.9021	0.0443

Table 4
Quantitative results: Dice coefficients between breast mask before and after normalization on test data.

Models	GE to SE		SE to GE	
	Mean	Std	Mean	Std
Std. CycleGAN	0.8913	0.0941	0.9089	0.0471
Std. CycleGAN + MINE	0.8976	0.0510	0.8949	0.0391
Proposed Discrim	0.9801	0.0061	0.9813	0.0049
Proposed Discrim + MINE	0.9082	0.0714	0.8912	0.0706




Fig. 8. Mean intensity value distribution of dense tissues (DT) in (a) original GE and original Siemens (b) original GE and normalized Siemens (c) and original Siemens and normalized GE.

the proposed discriminator causes a reduction in the Dice coefficient value ($0.9801 \rightarrow 0.9082$) due to a decrease in the shape of the breast. However, the standard CycleGAN model and its variant with mutual information both have comparable Dice coefficients. This can be explained by both methods' inability to preserve the breast shape. Similar observations can be made for the normalization between Siemens to GE. This confirms that the modified architecture with the field of view of 34×34 results in superior performance.

In summary, from a qualitative point of view, the standard CycleGAN along with mutual information leads to the worst result (See Fig. 5) This is also reflected in the quantitative results, where it achieves almost the minimum dice coefficient score. On the other hand, the proposed modified CycleGAN framework obtained by altering the discriminator architecture is able to consistently preserve the dense tissue as well as the breast shape. These observations also align with the quantitative results on test data presented in Table 4.

To evaluate the intensity transformation, we manually annotate the dense tissue in 10 cases and then measure the mean intensity of these annotated regions both before normalization and after normalization. The result is presented in Fig. 8 where Fig. 8(a) illustrates the mean intensity distribution of the dense tissue in GE and Siemens before the normalization. It can be observed from Fig. 8(b) that the mean intensity distribution of the original GE is comparable to the normalized Siemens. A similar observation can also be made from Fig. 8(c) for the original Siemens and normalized GE. This demonstrates that the proposed method is able to successfully adjust the intensity of the image as it pertains to dense tissue. It should also be noted that along with intensity adjustment, the proposed method learns to map the noise "halo" around the breast, which is a crucial aspect of vendor normalization. The proposed vendor normalization method will thus potentially increase the robustness of downstream models that do not have access to adequate training data from multiple vendors by synthesizing larger and richer datasets, which will mitigate issues related to class imbalance.

8. Conclusions

In this article, we have shown that a fully convolutional neural network can be successfully trained to learn a bidirectional mapping and perform normalization between DCE-MRI images generated from different scanners (GE Healthcare & Siemens). In

contrast to previous works, our proposed method not only performs intensity normalization but also learns the noise distribution pattern.

Our evaluation shows that when the standard CycleGAN is applied to this task, it matches the desired intensity of images but struggles with the shape of the breast and dense tissue. This is caused by the limited constraint on the images generated by the GANs and in turn, liberty that it takes to freely generate breast images. In response to this, we propose two solutions. The first one is to incorporate mutual information into the loss function. Our rationale is that this modification will ensure that the structure of the breast is maintained between the input and the output of the generator. This first solution fails to solve the problem due to a very specific characteristic of the data, which is the noise "halo" around the breast. Incorporating mutual information into a CycleGAN is not a trivial task and we believe that the method of doing so proposed in this paper will be helpful for other similar tasks in medical imaging and beyond. The second solution to the problem of maintaining the structure of the breast that we propose in this paper is a modification to the discriminator. This solution proves to be highly successful for this task as verified by our experiments.

Our study has some limitations. One limitation of this work is that it provides the capability of translation using 2D images only. While some effort in network design and parameter optimization is certainly needed, the proposed methods naturally lend themselves to 3D MR volumes. Another limitation is that our dataset consists of only two vendors and a relatively limited number of patients. While we still believe that the dataset used in this study represents the real-life problem faced in analyses of breast MRIs, further studies are needed to show that the proposed method generalizes beyond the data presented here. Finally, while we were able to demonstrate that our method results in no or minimal changes to the dense tissue structure, additional validation of the applicability for specific applications should be a topic of future studies. For example, while for some clinical applications, no changes in the breast tissue structure are acceptable, radiomics or deep learning applications are likely to be robust to some changes of this type. Data generated using our method could additionally be used to augment training data and improve deep learning model generalization. A deep learning model trained with augmented data from the various scanner will enable model generalization to real-world datasets with moderately different characteristics.

In summary, we propose a framework for normalization of breast MRIs based on CycleGAN. We also propose a few technical innovations that overcome various challenges that we experienced while applying CycleGAN framework to our task of breast MRI normalization. While the framework has only been tested using breast MRIs, it naturally lends itself to other medical imaging tasks where no paired data is available.

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.

Acknowledgments

The authors would like to thank Mr. Mateusz Buda for the insightful discussions.

Appendix A. Architectures for discriminator

Discriminator architectures with various field of view is presented in this section. Each model uses a convolution after the last layer to produce a 1-D output of size $m \times m$. Instance Norm layer was not applied to first layer in each of the architecture. The slope for LeakyReLU was 0.2.

Table A.1
Discriminator architecture (70 × 70).

Layer	Input Channel	Output Channel	Filter Size (k)	Stride (S)	Activation
Convolution	1	64	4 × 4	2	Leaky ReLU
Convolution	64	128	4 × 4	2	Leaky ReLU
Convolution	128	256	4 × 4	2	Leaky ReLU
Convolution	256	512	4 × 4	1	Leaky ReLU
Convolution	512	1	4 × 4	1	–

Table A.2
Discriminator architecture (45 × 45).

Layer	Input Channel	Output Channel	Filter Size (k)	Stride (S)	Activation
Convolution	1	64	5 × 5	2	Leaky ReLU
Convolution	64	128	5 × 5	2	Leaky ReLU
Convolution	128	256	5 × 5	1	Leaky ReLU
Convolution	256	1	5 × 5	1	–

Table A.3
Discriminator architecture (34 × 34).

Layer	Input Channel	Output Channel	Filter Size (k)	Stride (S)	Activation
Convolution	1	64	4 × 4	2	Leaky ReLU
Convolution	64	128	4 × 4	2	Leaky ReLU
Convolution	128	256	4 × 4	1	Leaky ReLU
Convolution	256	1	4 × 4	1	–

Table A.4
Discriminator architecture (PixelGAN).

Layer	Input Channel	Output Channel	Filter Size (k)	Stride (S)	Activation
Convolution	1	64	1 × 1	1	Leaky ReLU
Convolution	64	128	1 × 1	1	Leaky ReLU
Convolution	128	1	1 × 1	1	–

Table A.5
Hyperparameter (λ_{cyc}): Dice coefficients between breast mask before and after normalization obtained on validation data.

	GE to SE		SE to GE	
	Mean	Std	Mean	Std
$\lambda_{cyc} = 10$	0.85660	0.04784	0.87012	0.03411
$\lambda_{cyc} = 7.5$	0.85760	0.07617	0.88693	0.03765
$\lambda_{cyc} = 5.0$	0.91381	0.05770	0.90209	0.04426
$\lambda_{cyc} = 2.5$	0.90243	0.08136	0.90577	0.03307

Table A.6
Hyperparameter (λ_{mut}): Dice coefficients between breast mask before and after normalization obtained on validation data.

	GE to SE		SE to GE	
	Mean	Std	Mean	Std
$\lambda_{mut} = 0.01$	0.87322	0.05853	0.90161	0.03569
$\lambda_{mut} = 0.10$	0.91552	0.03438	0.89473	0.03570
$\lambda_{mut} = 0.25$	0.91112	0.04433	0.89564	0.05675
$\lambda_{mut} = 0.50$	0.91801	0.04055	0.92319	0.03903
$\lambda_{mut} = 1.0$	0.88428	0.04350	0.91006	0.03929

References

- [1] American Cancer Society, Cancer Facts & Figures, 2021 <https://www.cancer.org/content/dam/cancer-org/research/cancer-facts-and-statistics/annual-cancer-facts-and-figures/2021/cancer-facts-and-figures-2021.pdf>.
- [2] B. Reig, L. Heacock, K.J. Geras, L. Moy, Machine learning in breast MRI, *J. Magn. Reson. Imaging* 52 (4) (2020) 998–1018.
- [3] P.B. Sachs, K. Hunt, F. Mansoubi, J. Borgstede, CT And MR protocol standardization across a large health system: providing a consistent radiologist, patient, and referring provider experience, *J. Digit. Imaging* 30 (1) (2017) 11–16.
- [4] P.S. Sharma, A.M. Saindane, Standardizing magnetic resonance imaging protocols across a large radiology enterprise: barriers and solutions, *Curr. Probl. Diagn. Radiol.* 49 (5) (2020) 312–316.
- [5] A. Saha, X. Yu, D. Sahoo, M.A. Mazurowski, Effects of MRI scanner parameters on breast cancer radiomics, *Expert Syst. Appl.* 87 (2017) 384–391.
- [6] L.G. Nyul, J.K. Udupa, Xuan Zhang, New variants of a method of MRI scale standardization, *IEEE Trans. Med. Imaging* 19 (2) (2000) 143–150, doi:10.1109/42.836373.
- [7] S. Rizzo, F. Botta, S. Raimondi, D. Origi, C. Fanciullo, A.G. Morganti, M. Bellomi, Radiomics: the facts and the challenges of image analysis, *Eur. Radiol. Exp.* 2 (1) (2018) 1–8.
- [8] S. Roy, A. Carass, P.L. Bazin, J.L. Prince, Intensity inhomogeneity correction of magnetic resonance images using patches, *Medical Imaging 2011: Image Processing*, 7962, International Society for Optics and Photonics, 2011.
- [9] E.A. AlBadawy, A. Saha, M.A. Mazurowski, Deep learning for segmentation of brain tumors: impact of cross-institutional training and testing, *Med. Phys.* 45 (3) (2018) 1150–1158.
- [10] G. Mårtensson, D. Ferreira, T. Granberg, L. Cavallin, K. Oppedal, A. Padovani, I. Rektorova, L. Bonanni, M. Pardini, M.G. Kramberger, et al., The reliability of a deep learning model in clinical out-of-distribution MRI data: a multicohort study, *Med. Image Anal.* 66 (2020) 101714.
- [11] A. Simkó, T. Löfstedt, A. Garpebring, T. Nyholm, J. Jonsson, A generalized network for MRI intensity normalization, *arXiv:1909.05484* (2019).
- [12] E.H. Pooch, P. Ballester, R.C. Barros, Can we trust deep learning based diagnosis? The impact of domain shift in chest radiograph classification, in: *International Workshop on Thoracic Image Analysis*, Springer, 2020, pp. 74–83.
- [13] X. Sun, L. Shi, Y. Luo, W. Yang, H. Li, P. Liang, K. Li, V.C. Mok, W.C. Chu, D. Wang, Histogram-based normalization technique on human brain magnetic resonance images from different acquisitions, *Biomed. Eng. Online* 14 (1) (2015) 73.
- [14] G. Collewet, M. Strzelecki, F. Mariette, Influence of MRI acquisition protocols and image intensity normalization methods on texture classification, *Magn. Reson. Imaging* 22 (1) (2004) 81–91.
- [15] A. Madabhushi, J.K. Udupa, New methods of MR image intensity standardization via generalized scale, *Med. Phys.* 33 (9) (2006) 3426–3434.
- [16] S. Roy, A. Carass, J. Prince, A compressed sensing approach for MR tissue contrast synthesis, in: G. Székely, H.K. Hahn (Eds.), *Information Processing in Medical Imaging*, Springer, Berlin, Heidelberg, 2011, pp. 371–383.
- [17] R.T. Shinohara, E.M. Sweeney, J. Goldsmith, N. Shiee, F.J. Mateen, P.A. Calabresi, S. Jarso, D.L. Pham, D.S. Reich, C.M. Crainiceanu, et al., Statistical normalization techniques for magnetic resonance imaging, *NeuroImage* 6 (2014) 9–19.
- [18] J. Zhang, A. Saha, B.J. Soher, M.A. Mazurowski, Automatic deep learning-based normalization of breast dynamic contrast-enhanced magnetic resonance images, 2018, 1807.02152.
- [19] B. Fischl, D.H. Salat, A.J. Van Der Kouwe, N. Makris, F. Ségonne, B.T. Quinn, A.M. Dale, Sequence-independent segmentation of magnetic resonance images, *NeuroImage* 23 (2004) S69–S84.

- [20] Y. Gao, Y. Liu, Y. Wang, Z. Shi, J. Yu, A universal intensity standardization method based on a many-to-one weak-paired cycle generative adversarial network for magnetic resonance images, *IEEE Trans. Med. Imaging* 38 (9) (2019) 2059–2069.
- [21] B.H. van Der Velden, M.J. van Rijssel, B. Lena, M.E. Philippons, C.E. Loo, M.A. Ragusi, S.G. Elias, E.J. Sutton, E.A. Morris, L.W. Bartels, et al., Harmonization of quantitative parenchymal enhancement in t1-weighted breast MRI, *J. Magn. Reson. Imaging* 52 (5) (2020) 1374–1382.
- [22] X. Yi, E. Walia, P. Babyn, Generative adversarial network in medical imaging: areview, *Med. Image Anal.* (2019) 101552.
- [23] M.E. Tschuchnig, G.J. Oostingh, M. Gadermayr, Generative adversarial networks in digital pathology: a survey on trends and future potential, *Patterns* 1 (6) (2020) 100089.
- [24] D. Morrison, D. Harris-Birtill, P.D. Caie, Generative deep learning in digital pathology workflows, *Am. J. Pathol.* (2021).
- [25] 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.
- [26] D. Nie, R. Trullo, J. Lian, L. Wang, C. Petitjean, S. Ruan, Q. Wang, D. Shen, Medical image synthesis with deep convolutional adversarial networks, *IEEE Trans. Biomed. Eng.* 65 (12) (2018) 2720–2730.
- [27] C.-B. Jin, H. Kim, M. Liu, W. Jung, S. Joo, E. Park, Y.S. Ahn, I.H. Han, J.I. Lee, X. Cui, Deep CT to MR synthesis using paired and unpaired data, *Sensors* 19 (10) (2019) 2361.
- [28] K. Armanious, C. Jiang, S. Abdulatif, T. Küstner, S. Gatidis, B. Yang, Unsupervised medical image translation using Cycle-MedGAN, *arXiv:1903.03374* (2019).
- [29] Y.-R. Van Eycke, A. Foucart, C. Decaestecker, Strategies to reduce the expert supervision required for deep learning-based segmentation of histopathological images, *Front. Med.* 6 (2019) 222.
- [30] S.U. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, T. Çukur, Image synthesis in multi-contrast MRI with conditional generative adversarial networks, *IEEE Trans. Med. Imaging* 38 (10) (2019) 2375–2388.
- [31] V. Sandfort, K. Yan, P.J. Pickhardt, R.M. Summers, Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in ct segmentation tasks, *Sci. Rep.* 9 (1) (2019) 1–9.
- [32] G. Shi, J. Wang, Y. Qiang, X. Yang, J. Zhao, R. Hao, W. Yang, Q. Du, N.G.-F. Kazihise, Knowledge-guided synthetic medical image adversarial augmentation for ultrasonography thyroid nodule classification, *Comput. Methods Programs Biomed.* 196 (2020) 105611.
- [33] L. Sun, J. Wang, Y. Huang, X. Ding, H. Greenspan, J. Paisley, An adversarial learning approach to medical image synthesis for lesion detection, *IEEE J. Biomed. Health Inform.* 24 (8) (2020) 23032314, doi:10.1109/jbhi.2020.2964016.
- [34] S.J. Shin, S.C. You, H. Jeon, J.W. Jung, M.H. An, R.W. Park, J. Roh, Style transfer strategy for developing a generalizable deep learning application in digital pathology, *Comput. Methods Programs Biomed.* 198 (2021) 105815.
- [35] Z. Zhang, L. Yang, 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.
- [36] C. Wang, G. Macnaught, G. Papanastasiou, T. MacGillivray, D. Newby, Unsupervised learning for cross-domain medical image synthesis using deformation invariant cycle consistency networks, in: *International Workshop on Simulation and Synthesis in Medical Imaging*, Springer, 2018, pp. 52–60.
- [37] 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 (ICCV)*, 2017.
- [38] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution, in: *European Conference on Computer Vision*, 2016.
- [39] P. Isola, J.Y. Zhu, T. Zhou, A.A. Efros, Image-to-image translation with conditional adversarial networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [40] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: *Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [41] B. Poole, S. Ozair, A. Van Den Oord, A. Alemi, G. Tucker, On variational bounds of mutual information, in: *Proceedings of Machine Learning Research*, 2019, pp. 5171–5180.
- [42] D. McAllester, K. Stratos, Formal limitations on the measurement of mutual information, in: *International Conference on Artificial Intelligence and Statistics*, PMLR, 2020, pp. 875–884.
- [43] A.M. Saxe, Y. Bansal, J. Dapello, M. Advani, A. Kolchinsky, B.D. Tracey, D.D. Cox, On the information bottleneck theory of deep learning, *J. Stat. Mech.* 2019 (12) (2019) 124020.
- [44] I. Belghazi, S. Rajeswar, A. Baratin, R.D. Hjelm, A.C. Courville, MINE: Mutual information neural estimation, *ArXiv abs/1801.04062* (2018).
- [45] M.D. Donsker, S.S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time. iv, *Commun. Pure Appl. Math.* 36 (2) (1983) 183–212.
- [46] G. Modanwal, A. Vellal, M. Buda, M.A. Mazurowski, MRI image harmonization using cycle-consistent generative adversarial network, *Medical Imaging 2020: Computer-Aided Diagnosis*, 11314, International Society for Optics and Photonics, 2020.
- [47] X. Mao, Q. Li, H. Xie, R.Y. Lau, Z. Wang, S. Paul Smolley, Least squares generative adversarial networks, in: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [48] A. Borji, Pros and cons of GAN evaluation measures, *Comput. Vis. Image Underst.* 179 (2019) 41–65.