

DiamondGAN: Unified Multi-Modal Generative Adversarial Networks for MRI Sequences Synthesis

Hongwei Li^{1*}, Johannes C. Paetzold^{1*}, Anjany Sekuboyina¹, Florian Kofler¹,
Jianguo Zhang³, Jan S. Kirschke², Benedikt Wiestler², and Bjoern Menze¹

1. Department of Computer Science, Technical University of Munich, Germany
2. Klinikum rechts der Isar, Technical University of Munich, Germany
3. School of Computing, University of Dundee, United Kingdom
{hongwei.li, bjoern.menze@tum.de}

Abstract. Synthesizing MR imaging sequences is highly attractive for clinical practice, as often single sequences are missing or of poor quality (e.g. due to motion). Naturally, the idea arises that a target modality would benefit from multi-modal input. However, existing methods fail to scale up to non-aligned image volumes with multiple modalities, facing common drawbacks of complex multi-modal imaging sequences. We propose a novel, scalable and multi-modal approach called *DiamondGAN*. Our model is capable of performing flexible non-aligned cross-modality synthesis and data infill, when given multiple modalities or any of their arbitrary subsets, learning structured information in an end-to-end fashion. We synthesize two MRI sequences with clinical relevance (i.e., double inversion recovery (DIR) and contrast-enhanced T1 (T1-c)), reconstructed from three common MRI sequences. In addition, we perform multi-rater visual evaluation experiment and find that trained radiologists are unable to distinguish our synthetic DIR images from real ones.

1 Introduction

In clinical practice, imaging datasets often consists of high-dimensional image volumes, with multiple imaging protocols and repeated scans at multiple time points. Given the multiplicity of possible sequence parameters, protocols largely vary according to the scanner and imaging center, hindering their comparability. This often leads to repeated exams or severely limits the clinical information that can be drawn from those MRI studies. Particularly in the case of multiple sclerosis, longitudinal comparisons of MRI studies are the main reason for treatment decisions while lesion quantification tools require complete identical modalities at multiple time points. Potentially cross-modality image synthesis technique can resolve those obstacles through efficient data infilling and re-synthesis.

Recently, generative adversarial networks (GANs) have been applied in translating MRI sequences, positron emission tomography (PET) and computed tomography (CT) images. Most of them are one-to-one cross-modality synthesis

* equal contribution

approaches, for example, PET-CT [1], CT-MRI [16], MRI sequences translation [4]. Recent multi-modal synthesis method [13] has limited scalability because the input modalities are required to be uniform and spatially aligned. Although there are several multi-domain translation algorithms [9,3] in the computer vision community, these approaches design one-to-multiple domain translation but do not model the multiple-to-one domain mapping. Especially in medical images synthesis, multiple-to-one cross-modality mapping is highly relevant as proprietary information of individual and non-aligned modalities is probably synergistic.

There are three main challenges in the scenarios of multi-modal cross-modality medical image synthesis: 1) the input modalities are assumed to be *not* spatially-aligned because registration methods for aligning multiple modalities may fail, limiting the applicability of conventional regression approaches. 2) input modalities may be missing due to different clinical settings between centers, thus a traditional regression-based data infill would be restricted to the smallest uniform subset or rely on iterative data infill methods. 3) existing approaches have limited scalability, e.g. in a *Cycle-GAN* [17] setting, one would therefore have to train individual models for possible combinations of the input modalities.

Contributions 1) We propose *DiamondGAN*, a unified multi-modal generative adversarial network that learns the multiple-to-one cross-modality mapping among non-aligned modalities using only a pair of generators and discriminators, training effectively with a multi-modal cycle-consistency loss function. 2) We provide both qualitative and quantitative results on two clinically-relevant MRI sequences synthesis tasks, showing *DiamondGAN*'s superiority over baseline models. 3) We present the results of extensive visual evaluation, performed by fourteen experienced radiologists to confirm the quality of synthetic images.

2 Methodology

2.1 Multi-Modal Cross-Modality Synthesis

Given an input set of n modalities: $X = \{x_i | i = 1, \dots, n\}$ and a target modality T . Our goal is to learn a generator G that learns mappings from multiple input modalities to one target modality. We assume that 1) all the modalities, i.e., X and T , are *not* spatially-aligned because it is rather difficult to obtain *strictly* spatially-aligned images as mentioned in Section 1; 2) the input modalities can be any subset of X , denoted as X' during the training and inference stages as some modalities of a subject may be missing in clinical practice.

We enforce G to be capable of translating any subset X' into a target modality T using a condition c which indicates the presence of the input modalities, i.e., $G(X', c) \rightarrow T$. This condition handles the missing modality issue and makes it a scalable model in both the training and the inference stages. We further introduce a multi-modal cycle-consistency loss to handle the "non-aligned modalities" issues among the input and output. Fig. 1 illustrates the main idea of our proposed approach. We regularly generate the condition c and the corresponding multi-modal data X_c of all possible combinations, so that G learns to flexibly translate the arbitrary multi-modal input. As mentioned in the caption of Fig. 1, we use an *availability* condition to serve as an indicator of the input

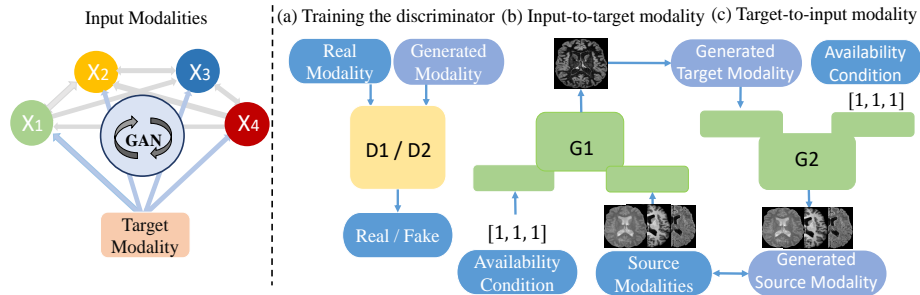


Fig. 1: Left: The high-level idea behind our *DiamondGAN*, which is capable of learning mappings between any subset of multiple input modalities (X) to a target modality in a single model. This mapping represents a diamond-shape topology. Right: Overview of *DiamondGAN*. It consists of two modules, a pair of discriminators D and a pair of generators G . (a) $D1$ and $D2$ learn to distinguish between real and synthetic images from multi-modal input and the target output respectively. (b) $G1$ takes both multi-modal input and the condition as input and generates a target modality. The condition c is a binary vector: $c = \{c_1, c_2, \dots, c_n\}$, where $c_i = 0$ or 1 indicates the corresponding input modality as available (1) or not (0). It is spatially replicated and concatenated with the input image in the feature-map level. (c) $G2$ tries to generate the original modalities from the synthetic target modality given the original availability condition.

modalities. It is spatially replicated to the image size ($1 \times H \times W$) and is a part of the two-stream network input. In the case of 3 modalities as the input, the condition $c = [1, 1, 1]$ would indicate that every input modality is given.

Multi-Modal Reconstruction Loss We aim to train G to guarantee that synthetic target modality preserves the content of its input modalities. The input modalities are assumed to be not spatially aligned or not from the same subject as mentioned above. In this situation, the traditional cycle loss [17] as well as the regression loss [7] would fail to tackle with the multi-modal and non-aligning issues. To alleviate the two problems, we extend the traditional cycle-consistency loss [17] to a multi-modal one. Specifically, we concatenate the source modalities into a multi-channel input and define a multi-channel output as the target modality. We then simultaneously train two generators $G_1 : X \rightarrow T$ and $G_2 : T \rightarrow X$ in a cycle-consistency fashion. Please note that the output target modality is with multiple channels which correspond to the input modalities. The loss function of the generator is defined as:

$$\mathcal{L}_{rec} = \mathbb{E}_{X,T,c} \{ \|X - G_2(G_1(X, c), c)\|_1 + \|T - G_1(G_2(T, c), c)\|_1 \} \quad (1)$$

Adversarial Loss To make the generated images indistinguishable from real images, we adopt an adversarial loss:

$$\begin{aligned} \mathcal{L}_{adv} = & \mathbb{E}_{X,T} \{ \log [D_1(X) \cdot D_2(T)] \} \\ & + \mathbb{E}_{X,T,c} \{ \log [(1 - D_2(G_1(X, c))) \cdot (1 - D_1(G_2(T, c)))] \} \end{aligned} \quad (2)$$

where G_1 generates a target modality $G_1(X, c)$ conditioned on the presence of input modalities X , while D_1 tries to distinguish between real input modalities and generated ones. Similarly, G_2 generates the original input modalities $G_2(T, c)$ conditioned on the presence of original input modalities X and D_2 tries to distinguish between real target modality and generated target one. The generators try to minimize this objective, while the discriminators try to maximize it.

Full Objective The objective functions to optimize D and G respectively are

$$\mathcal{L}_D = -\mathcal{L}_{adv}; \quad \mathcal{L}_G = \mathcal{L}_{adv} + \lambda_{rec}\mathcal{L}_{rec} \quad (3)$$

where λ_{rec} is the hyper-parameter that balances the reconstruction loss and adversarial loss.

2.2 Implementation

Two-Stream Network Architecture To leverage the information from both input modalities and corresponding availability conditions, we build a two-stream network architecture based on the popular encoder-decoder network [8]. It takes the multi-modal images and condition as two inputs and merges them in the feature level. This network contains stride-2 convolutions, residual blocks [6] and fractionally-strided convolutions (1/2 stride). We use 6 blocks for the input size of $N \times H \times W$, where N , H and W are the number of modalities, height and width of the images respectively. The input images and availability conditions pass through two encoders and are merged in the last feature layer before the decoder. *PatchGANs* [8] is used for the discriminator network, which classifies the patch features maps to real or fake, instead of using a fully-connected layer.

Training Details We apply two recent techniques to stabilize the training of the model. First, for \mathcal{L}_{adv} (Eq. 2), we replace the negative log likelihood objective by a least-squares loss [12]. Second, to reduce the model oscillation [5], we update the discriminators using a history of generated images rather than the ones produced by the latest generators, as proposed in [14]. Thus we put the 25 previously generated images in an image buffer. We set $\lambda_{rec} = 10$ in Equation 3 for all the experiments. We use the Adam solver [10] with a batch size of 5. All networks were trained from scratch with a learning rate of 0.0002 and for 20 epochs. When given n input modalities, for each epoch the parameters in both generator and discriminator are updated for $2^n - 1$ times given $2^n - 1$ training subsets of input modalities excluding empty set.

2.3 Expert Rating and Evaluation Protocol

Quantitative evaluation of generated images in terms of standard scores for errors and correlation remains a debatable task [2]. Additionally, the evaluation with common metrics such as PSNR and MAE [15] would not tell us to whether the algorithm captures clinically relevant small substructures. Therefore, we strive to get experts' estimates of the image quality. We design a multi-rater quality evaluation experiment. Neuro-radiologists rated the images in a browser-application.

In each trial, they were provided with two images. On the left side, one real source image of a T1 or Flair images is presented. On the other side, a paired image of the target modality is shown which is either a real image or a generated one. The displayed paired images were randomly chosen in the pool of generated images and real images. This particular setup enables the experts to identify very small inconsistencies or implausibilities between the two images immediately. For evaluation, the doctors were asked to rate the plausibility of the image on the right based on the real image on the left. To rate the plausibility, they were asked to assign a 6-star rating, where 6 stars denoted a perfectly plausible image and 1 star a completely implausible image. The images were presented in 280 trials. The sequence of trials was randomized across participants.

3 Experiments

Datasets *Dataset 1* consists of 65 scans of patients with MS lesions from a local hospital, acquired with a multi-parametric protocol, which includes co-registered Flair, T1, T2, double inversion recovery (DIR) and contrast-enhanced T1 (T1-c) after skull-stripping. The first three modalities are common modalities in most MS lesion exams. DIR is a MRI pulse sequence, which suppresses signal from the cerebrospinal fluid and the white matter, enhancing the inflammatory lesion. T1-c is a MRI sequence which requires a paramagnetic contrast agent (usually gadolinium) that reduces the T1 relaxation time and thereby increases the signal intensity. Synthesizing DIR and T1-c is of clinical relevance because it can substantially reduce medical costs. We mainly report our result on *Dataset 1*. Additional *Dataset 2* is used for demonstrating our approach can work on multiple datasets with incomplete and non-aligned modalities. It is a part of the public MICCAI-WMH dataset [11], and includes 40 subjects with two modalities (Flair and T1). 2D axial slices are used for training the network. All the slices are cropped or padded to a uniform size of 240×240 and intensity values are rescaled to $[-1, 1]$.

Reconstructing DIR and T1-c from Common Modalities We perform two image synthesis tasks on two clinically-relevant MRI sequences (DIR and T1-c), using three common modalities (i.e., Flair, T1 and T2). We separate the *Dataset 1* into a training set, a validation set and a test set, resulting in 30 scans (2015 slices for each modality) for training and 35 scans for testing (2100 slices for each modality). To obtain the optimal hyper-parameters of the model, we use 5 out of the 30 training scans as a validation set. A common approach for quantitative evaluation of medical GAN images is to calculate relative errors and signal to noise ratio between the synthetic image and the real image [15]. Table 1 shows the results of peak signal-to-noise ratio (PSNR) and mean absolute error (MAE) by comparing the synthetic images and real T1-c and DIR images. For the synthetic DIR and T1-c images, we report the highest PSNR and the lowest MAE for a combined T1+T2+Flair input to our model. In the DIR synthesis experiment, the listed scores of using multiple inputs to our GAN are comparable (MAE 0.058-0.065). Whereas, the scores for single inputs are substantially worse (MAE 0.073-0.084). For the T1-c synthesis task, we find that any combination




Fig. 2: Samples of synthetic T1-c and DIR images given the combination of T1, T2 and Flair modalities. Difference images are generated and visualized in heat maps. Our generated images preserve the tissue contrast and the anatomy. However, we find more differences in synthetic DIR images than in synthetic T1-c ones, especially around the brain boundary. This could be due to the alignment error by registration methods.

of multi-modal inputs involving the T1 modality (MAE 0.045-0.048) results in better scores compared to other input. This indicates that our model successfully extracts the relevant information, as T1-c is essentially a T1 scan with a contrast enhancing agent. For comparison, we use *CycleGAN* [17] to perform one-to-one cross-modality synthesis, the best results of *CycleGAN* were listed in Table. 1. For DIR synthesis, using Flair images as the input of *CycleGAN* achieves the highest PSNR and lowest MAE while for T1-c, using T1 as the input gets the best performance. The proposed model outperforms *CycleGAN* in both tasks. We further replace a part of the training Flair and T1 images in *Dataset 1* with images from *Dataset 2* (totally 794 images for each modality) and we find the result on same testing set is comparable to using the original *Dataset 1*.

Rating Experiment by Fourteen Experts Fourteen neuro-radiologists with median 5+ years of professional experience participated and each of them evaluates 210 synthetic images and 70 original images. The 210 synthetic images are generated enforcing 6 different input conditions in which each condition includes 35 samples. The rating results of the 14 raters are averaged and the box plots of the results are shown in Figure 3. For the synthesis of T1-c images, we found that three multi-modal combinations (i.e., $T1$, $T1+Flair$ and $T1+T2+Flair$) gave comparable results, while the ones based solely on a Flair were consistently rated as implausible. The plausibility of DIR images synthesized with $T1+T2+Flair$ input was rated in average 0.83 star higher than that with solely T1 input. This is plausible as the DIR is a complex sequence containing proprietary information, its synthesis thus benefits from multiple input

Table 1: Quantitative evaluation of our generated images compared to the real DIR and T1-c images. We evaluate using PSNR and the mean absolute error (MAE) across 2100 testing images. Results show that the generated images benefit from a multi-modal input. \uparrow indicates that the higher value corresponds to better image quality.

	DIR $PSNR\uparrow$	DIR $MAE\downarrow$	T1-c $PSNR\uparrow$	T1-c $MAE\downarrow$
<i>CycleGAN</i> [17]	17.34	0.068	20.36	0.045
<i>DiamonGAN</i> _{T1}	15.46	0.084	20.21	0.048
<i>DiamonGAN</i> _{T2}	15.99	0.073	19.34	0.054
<i>DiamonGAN</i> _{Flair}	16.16	0.078	17.15	0.068
<i>DiamonGAN</i> _{T1+T2}	17.41	0.065	20.75	0.046
<i>DiamonGAN</i> _{T2+Flair}	18.58	0.059	19.78	0.051
<i>DiamonGAN</i> _{T1+Flair}	18.02	0.062	20.40	0.047
<i>DiamonGAN</i> _{T1+T2+Flair}	18.63	0.058	20.86	0.045




Fig. 3: Box plots showing the rating scores of synthetic images and real ones for T1-c modality on the left and DIR modality on the right. The means are shown as black numbers. *DiamondGAN* achieves comparable plausibility levels for the DIR modality.

sources. For the synthetic images with $T1+T2+Flair$ input, the experts assigned an identical rating to the synthetic and original images (4.54 stars *vs* 4.7 stars).

We conduct Wilcoxon rank-sum tests on the paired rating scores of synthetic and real images from 14 raters on 6 conditions which results in 6 pairs of 14 observations. Results show that the pair of rating scores on synthetic DIR images by $T1+T2+Flair$ input and real DIR images are not significantly different (p-value = 0.1432) while all other pairs are significantly different (p-values < 0.0001). This demonstrates that trained radiologists are unable to distinguish our synthetic DIR images from real ones. Furthermore, the experts ratings for the individual conditions of synthetic images are in agreement with the metrical evaluation in Table 1. For T1-c synthesis, the PSNR and MAE scores are consistently good when T1 modality is fed to *DiamondGAN*.

4 Conclusion and Discussion

This work introduces a novel approach for medical image synthesis using multiple non-aligned modalities, along with an extensive multi-rater experiment and

statistic tests. This multi-modal approach allows us to mine the structured information inside the existing extensive MRI sequences. Pathological evaluation is the ultimate goal of this work. Our approach is evaluated by clinical partners who contribute the datasets. We compared synthetic DIR sequence with conventional FLAIR sequence in the detection task of MS lesions in a cohort study. We show that the synthetic DIR images capture modality specific diagnostic features, reaching comparable results with real DIR images in detection MS lesions. The results demonstrate that *DiamondGAN* which synthesizes DIR sequence, has the potential to reduce medical costs.

References

1. Ben-Cohen, A., et al.: Cross-modality synthesis from ct to pet using fcn and gan networks for improved automated lesion detection. *Engineering Applications of Artificial Intelligence* 78, 186–194 (2019)
2. Borji, A.: Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding* 179, 41–65 (2019)
3. Choi, Y., et al.: Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In: *CVPR*. pp. 8789–8797 (2018)
4. Dar, S.U., et al.: Image synthesis in multi-contrast mri with conditional generative adversarial networks. *IEEE transactions on medical imaging* (2019)
5. Goodfellow, I.: Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160* (2016)
6. He, K., et al.: Deep residual learning for image recognition. In: *CVPR*. pp. 770–778 (2016)
7. Isola, P., et al.: Image-to-image translation with conditional adversarial networks. In: *CVPR*. pp. 1125–1134 (2017)
8. Johnson, J., et al.: Perceptual losses for real-time style transfer and super-resolution. In: *ECCV*. pp. 694–711. Springer (2016)
9. Karras, T., et al.: A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948* (2018)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
11. Kuijf, H.J., et al.: Standardized assessment of automatic segmentation of white matter hyperintensities; results of the wmh segmentation challenge. *TMI* (2019)
12. Mao, X., et al.: Least squares generative adversarial networks. In: *CVPR*. pp. 2794–2802 (2017)
13. Olut, S., et al.: Generative adversarial training for mra image synthesis using multi-contrast mri. In: *PRIME*. pp. 147–154. Springer (2018)
14. Shrivastava, A., et al.: Learning from simulated and unsupervised images through adversarial training. In: *CVPR*. pp. 2107–2116 (2017)
15. Welander, P., et al.: Generative adversarial networks for image-to-image translation on multi-contrast mr images-a comparison of cyclegan and unit. *arXiv preprint arXiv:1806.07777* (2018)
16. Wolterink, J.M., et al.: Deep mr to ct synthesis using unpaired data. In: *SASHIMI*. pp. 14–23. Springer (2017)
17. Zhu, J.Y., et al.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *CVPR*. pp. 2223–2232 (2017)