

12 Image Synthesis in Multi-Contrast MRI with Generative Adversarial Networks

Tolga Çukur, Mahmut Yurt, Salman Ul Hassan Dar, Hyungjin Chung,
and Jong Chul Ye

12.1 Introduction

The remarkable level and diversity of soft tissue contrasts in magnetic resonance imaging (MRI) has rendered it a preferred imaging modality for diagnostic imaging. Yet, prolonged examinations and associated healthcare costs often prohibit the acquisition of comprehensive multi-contrast protocols. Even when such protocols are viable, image quality might be compromised in a subset of acquisitions owing to system imperfections and/or uncooperative patients. As a result, there is emerging interest in the synthesis of MR images that are of poor quality or completely absent from the protocol [1–5]. Learning-based methods based on deep neural networks (DNNs) are gaining immense traction in this domain by their ability to capture joint distributions of multi-contrast MR images and to identify nonlinear mappings between separate contrasts [6–10].

Two main classes of synthesis approaches have come forth for multi-contrast MRI synthesis, namely unconditional and conditional methods. In unconditional synthesis, the aim is to generate new, independent, samples from a target image distribution. The most common base neural architecture basis for unconditional MRI synthesis has been generative adversarial networks (GANs). Such networks involve a game-theoretic competition between two subnetworks: a generator that tries to generate synthetic data samples, and a discriminator that tries to distinguish the fake samples from the generator using real samples of the desired data distribution. When the generator grows sufficiently strong to deceive the discriminator, it can generate realistic data samples. Given their unparalleled ability to learn data distributions, GANs have been adopted for synthesizing high-quality MR images given simply random noise samples as input [11, 12].

In conditional synthesis, the aim is to generate samples from a target distribution that are consistent with prior information from a source distribution. For MRI synthesis, the goal is to create an image from a target contrast given input images of the same anatomy from different source contrasts. As such, the structural content of the source images are used as prior information to conditionally improve the quality of the target samples generated. As in the unconditional case, GAN models have proven their exceptional performance for conditional MR image synthesis, as they characteristically offer much better capture of structural details compared with conventional learning-based models [8–10, 13–15].

Generative adversarial network models have enabled a leap in the realism and quality of multi-contrast MRI synthesis, powering it to improve the diagnostic value of MRI examinations. In this chapter, we discuss recent GAN-based synthesis approaches and their applications in MRI synthesis. We start with a quick overview of the physics underlying tissue contrasts in MRI. We then give a review of GAN basics, followed by an overview of existing unconditional and conditional models. We highlight collaborative GANs as a representative case of unified synthesis models for multiple distinct target contrasts. Finally, we close the chapter with a summary and outlook for future research in this area.

12.2 Physics for MR Contrast

Owing to the abundance of hydrogen in the body's water, it is the most widely imaged nucleus in MRI. Single protons within the hydrogen nuclei spin around a central axis along the direction of its magnetic moment. The spin axes are incoherently aligned in the absence of external inputs, resulting in zero net magnetization. However, when the spins are inserted into a static external magnetic field (B_0), they precess around the applied field at the Larmor frequency and a net magnetic moment builds up in the direction of B_0 , i.e., the longitudinal direction. A radio-frequency (RF) pulse exerted on the spins then tips their orientation to the transverse plane, with the flip angle of the pulse determining the degree of tip. Once the RF pulse is turned off, the magnetization gradually returns to equilibrium by back-aligning itself with B_0 . This restoration results in a decay of the transverse component of magnetization, with time constant T_2 , and growth of the longitudinal component, with time constant T_1 . The transverse component of the tissue magnetization is then sensed through flux changes on dedicated receiver coils. Meanwhile, additional magnetic field gradients are applied along Cartesian axes for spatial encoding. The resulting signal recorded at the receiver coil is given as

$$s(t) = \int_x \int_y \int_z M(x, y, z) e^{-t/T_2(\mathbf{r})} \exp\left(-i\gamma \int_0^t \mathbf{G}(\tau) \cdot \mathbf{r} d\tau\right) dx dy dz \quad (12.1)$$

where $M(x, y, z)$ is the spatial distribution of the transverse magnetization, $\mathbf{r} = [x, y, z]$, $\mathbf{G} = [G_x, G_y, G_z]$, T_2 is the relaxation constant, and γ denotes the nucleus-specific gyromagnetic ratio. In the k -space formulation, the gradient terms are expressed with spatial frequency variables along each axis, k_x , k_y , and k_z . Setting $k_x(t) = (\gamma/2\pi) \int_0^t G_x(\tau) d\tau$, $k_y(t) = (\gamma/2\pi) \int_0^t G_y(\tau) d\tau$, and $k_z(t) = (\gamma/2\pi) \int_0^t G_z(\tau) d\tau$, the signal equation modifies to

$$s(t) = \int_x \int_y \int_z M(x, y, z) e^{-t/T_2(\mathbf{r})} e^{-i2\pi[k_x(t)+k_y(t)+k_z(t)]} dx dy dz. \quad (12.2)$$

Note that Eq. (12.2) corresponds to the 3D Fourier transform of the magnetization, and k -space is simply the Fourier transform of the imaged volume. The periodic execution of the RF and gradient fields determines the signal levels of various tissues in MRI, which are dependent on the tissue-specific relaxation parameters (T_1 and T_2).

Therefore, the tissue contrasts in MRI can be manipulated by varying the RF pulse strength and pulse sequence timing. A mainstream sequence in MRI is the spin-echo sequence, whose signal is expressed as:

$$S_o = K\rho(\mathbf{r})[1 - e^{-TR/T_1(\mathbf{r})}]e^{-TE/T_2(\mathbf{r})} \quad (12.3)$$

where $\rho(\mathbf{r})$ is the proton density, K is the instrumental scaling constant, TR is the repetition time of the pulse sequence, TE is the echo time, and T_1 and T_2 are relaxation parameters. In spin-echo sequences, the tissue contrast can be adjusted to weight either the T_1 or T_2 values by altering TE/TR . For instance, T_1 -weighted contrast can be obtained by choosing low TE and moderate TR , whereas T_2 -weighted contrast can be obtained by choosing moderate TE and long TR . Through multiple acquisitions with different sequence parameters, the same anatomy can be imaged under different tissue contrasts, increasing the breadth of diagnostic information captured.

12.3 Brief Review of Generative Adversarial Networks (GANs)

Generative models are of great interest thanks to their ability to model high-dimensional probability distributions. In supervised learning, pairs of input–labels are given to the model, and the model learns the mapping between the pairs. On the other hand, the purpose of generative models, or unsupervised learning methods in general, is to approximate the distribution $p(\cdot)$ over the data x . Generative models can be divided into two large categories: explicit-density models and implicit-density models. While the former try to maximize the likelihood either directly or indirectly, the latter implicitly model the distribution, leveraging the property of being able to sample from the distribution.

Concretely, GANs define a prior $p(z)$ of input-noise variables which will be used to sample the data from the modeled distribution. An input variable then passes through a generator G , which maps the input variable to the distribution of the target high-dimensional data (e.g., the image). Another component that consists of a GAN is the discriminator D . The discriminator learns to represent the probability $D(x) \in [0, 1]$ that the given data x belongs to the real data distribution. Values closer to 0 mean that the discriminator has decided that it is unlikely that the given data was sampled from the real distribution, and values closer to 1 mean that the given data was realistic. The two components, G and D are pitted against each other in a mini–max game, where D tries to maximize the probability of correctly classifying the real and fake data, while G tries to minimize the probability of D being correct by generating realistic data to fool D [16]. Often, this competitive learning strategy is described with an analogy to a police officer (D) and a counterfeiter (G), where the counterfeiter gets better and better at deceiving the police officer, while the police officer gets consistently better at catching the counterfeiter.

The actual implementation of a GAN, as depicted in Fig. 12.1, models G and D with differentiable deep neural networks, and the network parameters are updated by optimizing the following loss function:




Figure 12.1 Illustration of a generative adversarial network

$$\min_G \max_D V(D, G), \tag{12.4}$$

where

$$V(D, G) = \mathbb{E}_x[\log D(x)] + \mathbb{E}_z[\log(1 - D(G(z)))]. \tag{12.5}$$

By iterating the update steps between G and D , the objective of GAN is to reach a Nash equilibrium, which means that each component is at its optimal state in relation to the opposite component. When the training is properly performed, the generator is able to produce very realistic data that are hard to distinguish from the real data.

While one acknowledges the strong representational power of GANs, they are also notorious for being hard to optimize. Often, optimization fails and the model either stagnates in bad local minima or ends up with mode collapse [17]. Therefore, the design of GAN loss functions has been an area of great interest [18–21]. When Eq. (12.5) is decoupled into two separate components, it can be written as

$$\begin{aligned} V^{(D)} &= -\mathbb{E}_x[\log D(x)] - \mathbb{E}_z[\log(1 - D(G(z)))], \\ V^{(G)} &= \mathbb{E}_z[\log(1 - D(G(z)))]. \end{aligned} \tag{12.6}$$

Although the vanilla (i.e., basic) GAN loss described in Eq. (12.6) is mathematically sound, it does not perform particularly well in practice. A simple heuristic fix is to change the cost function for the generator to

$$V^{(G)} = -\mathbb{E}_z[\log D(G(z))]. \tag{12.7}$$

By changing the cost function to Eq. (12.7), the generator receives a stronger gradient at the early stage of training, and thus is able to escape from bad local minima:

$$\begin{aligned} V^{(D)} &= \mathbb{E}_x[D(x)] - \mathbb{E}_z[D(G(z))] \\ V^{(G)} &= -\mathbb{E}_z[D(G(z))] \end{aligned} \quad (12.8)$$

Later, the authors of the Wasserstein GAN (WGAN) [18] proposed to use the “earth mover’s” (EM) distance as the metric between the generated distribution and the target distribution, as in Eq. (12.8). In contrast to minimizing the Kullback–Leibler (KL) divergence in the vanilla GAN [16], it was shown that the WGAN loss is superior at driving the cost to a global optimum both theoretically and experimentally. This idea was further extended in WGAN-gp [19], where the authors imposed a gradient penalty on top of the WGAN loss function to match a Lipschitz constraint.

Least squares GAN (LSGAN) [20] is another widely used GAN objective, described in Eq. (12.9):

$$\begin{aligned} V^{(D)} &= \mathbb{E}_x[(D(x) - b)^2] + \mathbb{E}_z[(D(G(z)) - a)^2], \\ V^{(G)} &= \mathbb{E}_z[(D(G(z)) - c)^2]. \end{aligned} \quad (12.9)$$

The main motivation of LSGAN is to pull the distribution of the generated data to the real-data manifold; using this simple loss formulation has proved to be efficient in many applications including image-to-image translation [22, 23].

Thanks to the great advances in training strategy and the development of efficient neural architectures, GANs have revolutionized deep learning and generative modeling over the past few years. Interestingly, although the first intuitive use of GANs was to generate images out of random noise, GANs have been applied to various fields, such as text-to-image generation [24], photo blending [25], and image-to-image translation [22, 23]. Consequently, GANs are nowadays the main workhorse of unsupervised deep learning, achieving the state-of-the-art in many areas including the field of MR image synthesis.

12.4 MR Contrast Conversion using GAN

Generative adversarial models with convolutional neural network (CNN) backbones are being widely adopted for MR image synthesis with a high degree of realism and structural detail. Depending on whether prior information from source contrasts is available, two main classes of model emerge: unconditional and conditional GANs. In this section, we will overview the basics of the two types of model and discuss their existing applications.

12.4.1 Unconditional GANs

Unconditional GANs learn to generate samples from a random noise vector without any extra information regarding the target MR images [11, 12, 26–31]. The generator is trained with the aim of generating realistic-looking MR images, and the discrimi-

nator is trained to distinguish real images from synthetic images. Classical GANs are trained to minimize the following adversarial loss function:

$$\ell_{adv} = \mathbb{E}_{x_t} [\log(D(x_t))] + \mathbb{E}_z [\log(1 - D(G(z)))] \quad (12.10)$$

where z is the noise vector, x_t denotes an image from the target distribution, G denotes the generator, and D is the discriminator; G is trained to minimize ℓ_{adv} while D is trained to maximize ℓ_{adv} .

Unconditional GANs are highly suited for data augmentation, to help improve the training of learning-based methods. As such, they have found broad use in image analysis tasks including segmentation and classification. For instance, synthetic cardiac MRI samples have been shown to improve classification performance [12], and volumetric brain MRI samples have been leveraged for enhanced segmentation [27].

As discussed before, a main limitation of vanilla GANs is their training instabilities, which could compromise the realism and diversity of generated images. Enhanced divergence measures have been proposed to alleviate this limitation, including the earth mover's distance in Wasserstein GAN (WGAN) [18]; this has been demonstrated for multi-contrast brain MRI synthesis [28, 31, 32]. Another prominent approach rests on the Pearson χ^2 divergence in least squares GAN (LSGAN) [20]; this approach has been demonstrated in prostate MR image synthesis for subsequent classification [32].

Another limitation of vanilla GANs involves the difficulty of training them on full-resolution images directly. Progressively growing GANs (PCGANs) that sequentially increase the image resolution during synthesis have been successfully demonstrated for brain image synthesis in glioma patients [29] and later were adopted for segmentation [30].

12.4.2 Conditional GANs

Conditional GANs learn to synthesize images in a target domain, given input images from a separate source domain as prior information. In the case of MRI, the source and target can be images of the same anatomy under separate MR contrasts.

In cross-contrast image synthesis, the aim is to learn the nonlinear mapping among images of distinct MRI tissue contrasts. A pioneering study in this field introduced pixel-GAN (pGAN) and cycleGAN (cGAN) methods for paired and unpaired MRI synthesis, respectively [8]. In the presence of training data consisting of co-registered images of source and target contrasts, pGAN leverages the pixel-wise and perceptual losses between the synthesized and real target images in conjunction with an adversarial loss (Fig. 12.2).

The pixel-wise loss can be expressed as

$$\ell_{pix} = \mathbb{E}_{x_t, x_s} [\|G(x_s) - x_t\|_1] \quad (12.11)$$

where x_s denotes the image of the source contrast. The perceptual loss function can be expressed as:

$$\ell_{perc} = \mathbb{E}_{x_t, x_s} [\|V(G(x_s)) - V(x_t)\|_1] \quad (12.12)$$

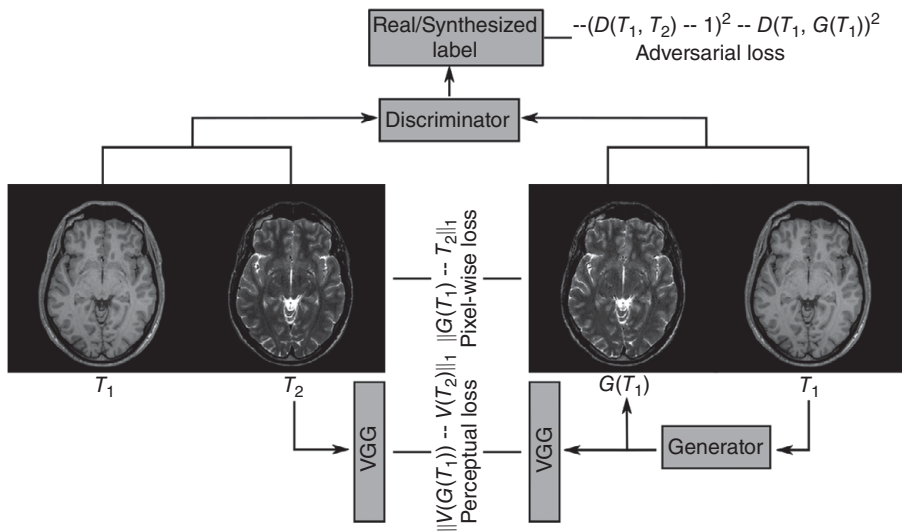


Figure 12.2 pGAN is a conditional GAN trained using co-registered images of source and target contrasts [8]. pGAN is trained to minimize an adversarial loss, a pixel-wise loss, and a perceptual loss. Note that, in the figure, T_1 and T_2 correspond respectively to x_s and x_t in Eq. (12.13). © 2019 IEEE. Reprinted, with permission, from [8]

where $V(\cdot)$ denotes a pre-trained computer vision model. The pGAN method enables a leap in accuracy of synthetic MRI over both conventional and prior deep-learning methods in the brain, as illustrated in Fig. 12.3.

For cases when paired images of source and target contrasts from the same set of subjects are unavailable, cGAN can be used to allow unpaired training with cycle-consistency loss (Fig. 12.4). Cycle-consistency enforces self-supervision on the GAN model:

$$\ell_{cc} = \mathbb{E}_{x_t, x_s} [\|G_{t \rightarrow s}(G_{s \rightarrow t}(x_s)) - x_s\|_1 + \|G_{s \rightarrow t}(G_{t \rightarrow s}(x_t)) - x_t\|_1] \quad (12.13)$$

where $G_{s \rightarrow t}$ is trained to recover an image of the target contrast given an image of the source contrast, and $G_{t \rightarrow s}$ is trained to synthesize an image of the source contrast given an image of the target contrast.

On the basis of the conditional GAN framework for multi-contrast MRI synthesis introduced by [8], later studies have either adopted it for other multi-contrast MRI synthesis applications [33, 34] or proposed additional loss terms to enforce enhanced priors [35, 36], and there have been many-to-one variants that aggregate information from multiple source contrasts [35, 37–39]. The authors of [37] incorporated multiple source contrasts by concatenating them at the input level in the form of a many-to-one GAN, and showed that this can enhance the overall recovery performance. In [38] it was demonstrated that traditional many-to-one GANs based on concatenation at the input level might be less sensitive to the unique information present within each source contrast, and this could lead to sub-optimal performance. The authors of [38] further




Figure 12.3 Representative T_1 -weighted images of a healthy subject, and T_2 -weighted images of a glioma patient recovered via pGAN: multimodal based on a traditional convolutional neural networks, and replica based on random forests are shown [8]. Compared with the basic the multimodal and replica, pGAN shows a remarkable recovery performance in both the healthy subject and the glioma patient. © 2019 IEEE. Reprinted, with permission, from [8]

showed that fusing features from multiple one-to-one GANs dedicated for each source contrast and a many-to-one GAN can lead to efficient recovery of both complementary and shared information across the source contrasts.

Several fundamental advances have been introduced to the main conditional GAN framework in recent studies. First, multi-tasking has been suggested as a means to improve the quality of MRI synthesis [40–44]. In [42] the authors performed synergistic synthesis and reconstruction by providing highly undersampled images of the target contrasts as additional priors. Similarly, [44] jointly super-resolved and synthesized the target contrast images.

Second, training instabilities were investigated particularly for 3D models whose complexity renders the learning process suboptimal for inevitably limited medical datasets. Spectral normalization and feature matching were proposed as two common methods to improve GAN training [45]. Self-attention modules were incorporated to reduce residual errors in focal, important, image regions [45]. Despite these advances, the need for massive datasets for training 3D models remains. Recently, the authors of [46] proposed the progressive decomposition of volumetric mapping into 2D mappings in a multi-planar fashion. The progressive model enforces spatial context to prevent the incoherence and artifacts commonly encountered in 2D models and can train accurate models for datasets that are several orders of magnitude smaller.

Most GAN-based models for MRI synthesis assume the availability of either paired or unpaired samples of multi-contrast MRI images from source and target domains. In practice, however, scan time limitations render it difficult to acquire fully sampled images. A recent study addresses this vital gap in the literature by introducing a semi-supervised GAN (ssGAN) model that can directly learn to synthesize MR images

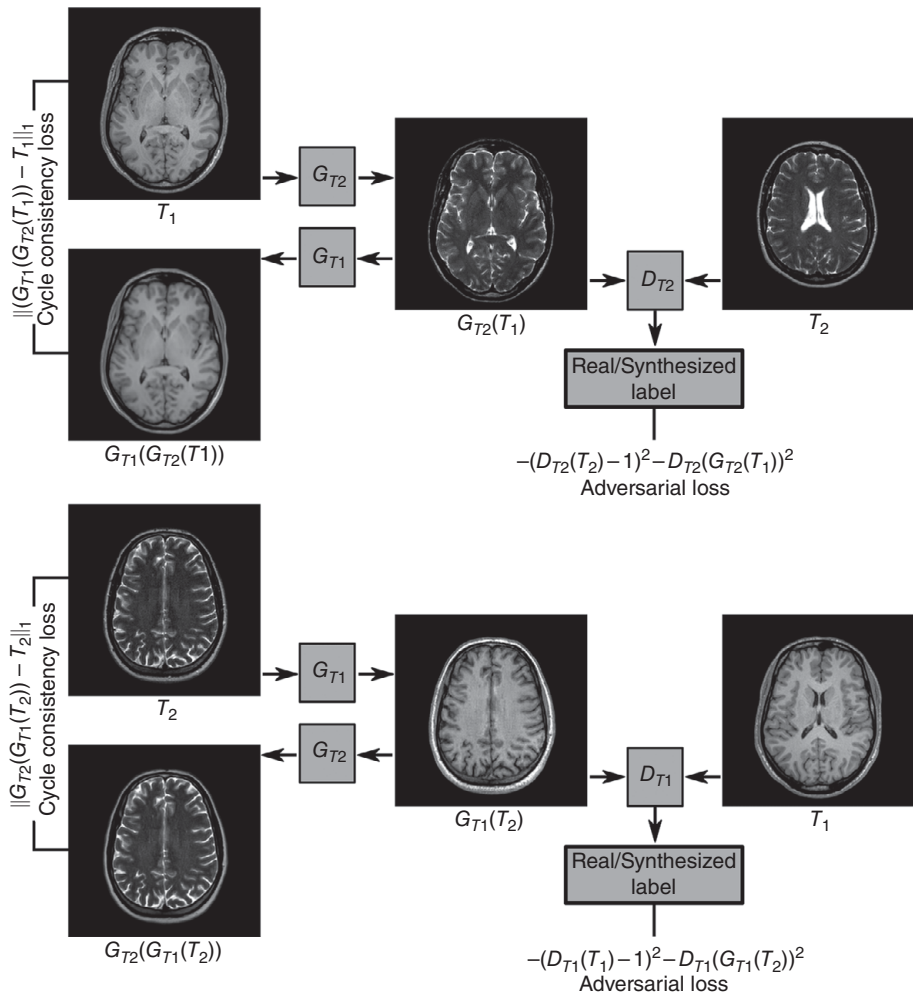


Figure 12.4 cGAN is a conditional GAN trained using unregistered images of source and target contrasts [8]. cGAN is trained to minimize an adversarial loss and a cycle-consistency loss. Note that, in the figure, T_1 and T_2 correspond respectively to x_s and x_t in Eq. (12.13). © 2019 IEEE. Reprinted, with permission, from [8]

from undersampled source and target contrasts [47]. A selective k -space loss function in ssGAN delivers a performance on a par with state-of-the-art supervised methods; it has highly accelerated datasets that are easier to collect.

12.5 Collaborative GAN for MR Contrast Conversion

Magnetic resonance contrast conversion is related to a topic that is broadly investigated in the field of statistics and machine learning: *missing data imputation*, or, more simply, *imputation*. Formally, the goal of imputation is to estimate the missing

element in the dataset using the remainder of the data and statistical modeling. Based on different modeling assumptions and theories, there are several well-known methods for imputation: regression imputation, nonnegative matrix factorization, stochastic imputation, etc. [48, 49]. However, these conventional methods have limitations when dealing with images. Unlike low-dimensional tabular data, image data exist in a high-dimensional manifold, which makes them very hard to model directly. Fortunately, owing to the recent development of deep learning, image imputation has now become relevant. One typical method to tackle image imputation in the field of deep learning is to consider the problem as image-to-image translation [9, 10, 22, 23, 50].

Among many studies, cycleGAN [23] has been the main workhorse owing to its capability to train a network even in the case of unpaired settings. While in a typical GAN there exists a single generator–discriminator set, cycleGAN exploits two sets of generator–discriminator pairs to formally address the mapping between two different image spaces,

$$\hat{x}_t = G_{i \rightarrow t}(x_i), \quad (12.14)$$

$$\hat{x}_i = G_{t \rightarrow i}(x_t), \quad (12.15)$$

where i, t are the input domain and the target domain, respectively, and $G_{i \rightarrow t}, G_{t \rightarrow i}$ are the generators that execute the mapping. However, a major shortcoming of using cycleGAN is that it cannot be used in situations where we have N different domains of interest. If we were to use cycleGAN to deal with all N domains, then we would need $N(N - 1)$ different generators, which would be prohibitive.

To consider multiple domains effectively, starGAN [50] was proposed. In starGAN, a *single* generator–discriminator set is used. Formally put, the mapping in starGAN is described as

$$\hat{x}_t = G(x_i; t), \quad (12.16)$$

where the single generator G is conditioned with t , and the subscript from Eq. (12.14) is dropped. Here, the generator knows the target domain t to be estimated by a mask vector that is concatenated with the input image. However, starGAN is also not a perfect fit for image imputation. Although it is able to exploit multiple domains with a single generator, it uses a *single* input to estimate the missing component, when it would be much more plausible to use multiple information from different domains synergistically.

12.5.1 Collaborative GAN

Accordingly, it would be beneficial to design a mapping to consider all other input images for missing data imputation. This idea was first proposed in collaborativeGAN (collaGAN) [9, 10]. A conceptual diagram of collaGAN is depicted in Fig. 12.5, where an example of synthesizing MR contrast using collaGAN is shown.

In this subsection, we formally address the collaGAN framework, handling multiple input domains to generate more feasible image imputation. For simplicity in expla-

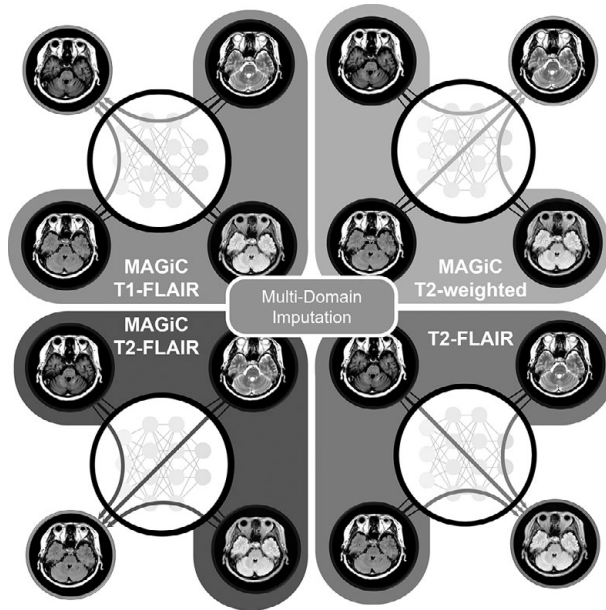


Figure 12.5 Conceptual diagram of collaGAN. © [2020] *Nature Machine Intelligence*, vol. 2, no. 1, pp. 34–42. Springer Nature. CollaGAN enables flexible multi-domain imputation within a single GAN architecture.

nation, we will assume that we have $N = 4$ domains: a, b, c , and d . When the target domain is a , and a target image x_a exists, collaGAN tries to estimate a collaborative mapping from the set of images that exist in different domains: $\{x_a\}^C = \{x_b, x_c, x_d\}$. Here, the superscript C denotes the complementary set. More properly, we have

$$\hat{x}_t = G(\{x_t\}^C; t), \quad (12.17)$$

where $t \in \{a, b, c, d\}$ is the corresponding target domain for each missing image. Notice that there are N different combinations of multiple inputs and single outputs. These combinations are randomly chosen at the training stage so that the generator learns to synergistically combine multiple images from different domains to synthesize the missing image. In the following, we review the overall flow of collaGAN and the training methodology.

Multiple Cycle-Consistency Loss

In cycleGAN, the cycle-consistency loss is used to impose a constraint that the mapping between the two domains should be the inverse of each other. In collaGAN, this cycle-consistency loss is redefined so that the inverse mapping can be achieved in *any* other domains. Specifically, let us assume that we have generated an output in domain a , denoted as \hat{x}_a , with the generator G , as shown in the middle panel of Fig. 12.6. Subsequently, we can define three different inverse mappings:

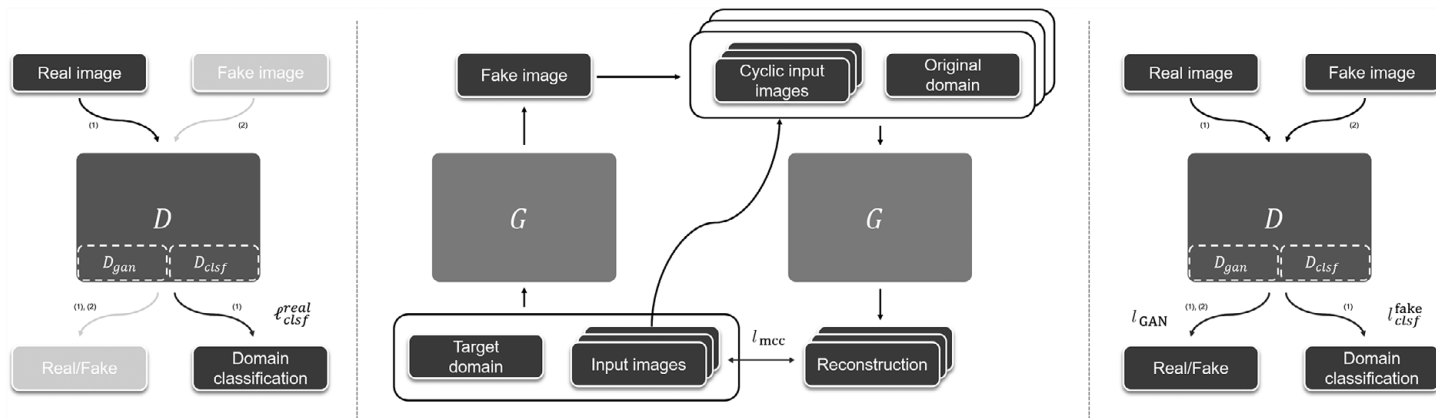


Figure 12.6 Training pipeline for collaGAN.

$$\tilde{x}_{b|a} = G(\{\hat{x}_a, x_c, x_d\}; b) \quad (12.18)$$

$$\tilde{x}_{c|a} = G(\{\hat{x}_a, x_b, x_d\}; c) \quad (12.19)$$

$$\tilde{x}_{d|a} = G(\{\hat{x}_a, x_b, x_c\}; d). \quad (12.20)$$

Then, we can calculate the *multiple cycle-consistency loss* for domain a as the following:

$$\ell_{mcc,a} = \|x_b - \tilde{x}_{b|a}\|_1 + \|x_c - \tilde{x}_{c|a}\|_1 + \|x_d - \tilde{x}_{d|a}\|_1. \quad (12.21)$$

This is illustrated in the middle panel of Fig. 12.6 indexed with ℓ_{mcc} and a bidirectional arrow. Generally, ℓ_{mcc} reads

$$\ell_{mcc,t} = \sum_{t' \neq t} \|x_{t'} - \tilde{x}_{t'|t}\|_1. \quad (12.22)$$

In the original paper [9], the authors claim to use a more sophisticated cycle-consistency loss on top of the ℓ_1 loss, coined the structural similarity index loss. The new quantity directly minimizes the structural similarity index (SSIM), which leads to a more feasible output:

$$\tilde{x}_{t'|t} = G(\{\hat{x}_t\}^C; t'). \quad (12.23)$$

Interested readers are referred to [9].

Discriminator Loss

Now that the cycle-consistency loss has been defined, we are ready to define the discriminator loss designed for multi-domain image imputation. In most GAN architectures, the role of the discriminator is to learn to tell how realistic the given image is. However, there is one more role of the discriminator in collaGAN. It also has to classify to which domain the given image belongs. A schematic diagram of the collaGAN discriminator is depicted in Fig. 12.6 (left) and (right). Within the discriminator D we see two components: D_{gan} and D_{clsf} . Here, D_{gan} refers to the adversarial loss, defined analogously to the other GAN architectures, and D_{clsf} refers to the domain classification loss. Before jointly training the generator and the discriminator, the domain classifier part of the discriminator, D_{clsf} , is trained with only real images:

$$\ell_{clsf}^{real}(D_{clsf}) = \mathbb{E}_{x_t} [-\log(D_{clsf}(t; x_t))], \quad (12.24)$$

where $(D_{clsf}(t; x_t))$ outputs a probability value of x_t belonging to the class t . Training D_{clsf} prior to jointly training G and D stabilizes training by building a strong classifier to provide a better guide to the generator. At the actual joint training stage, the generated fake images are also used to update the parameters in G with the following loss function:

$$\ell_{clsf}^{fake}(G) = \mathbb{E}_{\hat{x}_{t|t}} [-\log(D_{clsf}(t; \hat{x}_{t|t}))]. \quad (12.25)$$

The other part of the discriminator minimizes the LSGAN loss [20], known to stabilize the training process while preventing mode collapse. The parameters of D_{gan} are updated by minimizing

$$\ell_{gan}(D_{gan}) = \mathbb{E}_{x_t} [(D_{gan}(x_t) - 1)^2] + \mathbb{E}_{\tilde{x}_t|t} [(D_{gan}(\tilde{x}_t|t))^2], \quad (12.26)$$

whereas the parameters of G are updated by minimizing

$$\ell_{gan}(G) = \mathbb{E}_{\tilde{x}_t|t} [(D_{gan}(\tilde{x}_t|t) - 1)^2]. \quad (12.27)$$

12.5.2 MR Contrast Synthesis using CollaGAN

As discussed in earlier sections, different combinations of MR contrast images deliver diverse information about the patient being scanned. Among them, T_1 -weighted (T_1), T_2 -weighted (T_2), gadolinium-contrast-enhanced T_1 -weighted (T_1 Gd), T_2 fluid-attenuated inversion recovery (T_2 F) are canonical examples of MR contrasts that are widely used in clinical situations. Full acquisition of different MR contrast images would be the most beneficial, but in most cases this is hard to achieve: acquiring a complete set of MR contrast images requires a painfully long scanning time, and the protocols among different medical centers vary. Even if a complete set has been acquired, systematic and operational errors have often corrupted one of the images. Subsequently, such images cannot be used, which eventually leads to statistical errors that hinder exact analysis [48].

Recently, a synthetic MRI technique called magnetic resonance image compilation (MAGiC, GE Healthcare) [51] has grown in popularity owing to its ability to generate multiple contrast MR images using the newly developed multidynamic multiecho (MDME) scan. With MAGiC, one can generate different contrast images such as T_1 , T_2 , and T_2 -FLAIR, but it is known to generate substantial artifacts [51–53]. For example, it is known that the SNR is degraded with MAGiC FLAIR in comparison with conventional FLAIR, and the flow-related artifacts are enhanced. Hence, even after the acquisition of the MAGiC sequence, additional MR scans need to be incorporated to confirm the diagnosis, which eventually increases time and cost.

Instead, when we have erroneous acquisition among one of the contrast images, we can directly use collaGAN to impute the missing data, indicated by the question marks in Fig. 12.7. In the figure we can see that collaGAN is indeed able to impute the missing data in any given domain: MAGIC T_1 -FLAIR, T_2 , MAGIC T_2 -FLAIR, and T_2 -FLAIR. It is notable that by jointly combining the information from multiple domains, collaGAN is able to provide more accurate results compared with methods such as cycleGAN or starGAN which utilize a single-input domain.

The versatility of collaGAN readily extends to the imputation of pathological data. For example, in Fig. 12.8, we see the results of reconstruction with visible lesions. In (a), collaGAN is able to perform reconstruction even with an erroneous MAGIC T_2 -FLAIR image, showing a hyperintensity signal of the cerebrospinal fluid (CSF). Abnormal signals are also well reconstructed in (b). In (c), while the MAGIC T_2 -FLAIR image does not capture the hyperintensity signal, indicated by the arrow, collaGAN is able to reconstruct the signal. In (d), we see that collaGAN again corrects for the error induced in MAGIC T_2 -FLAIR. From the reconstruction results, it is clear that collaGAN can accurately reconstruct missing MR contrasts.




Figure 12.7 MR contrast imputation results using different methods. © [2020] *Nature Machine Intelligence*, vol. 2, no. 1, pp. 34–42. Springer Nature.

Assessing the Importance of MR Contrast

Interesting enough, this also leads to a fundamental question – can collaGAN synthesize *any* contrast? If not, what are the conditions required for collaGAN to perform well? In fact, this question was rigorously investigated in [10]. Specifically, by estimating specific contrasts one by one, it is possible to verify which contrasts can or cannot be generated through the imputation process, and ultimately which contrasts are irreplaceable. To ponder this question, the present authors performed a quantitative study by comparing the segmentation performance of the synthesized MR contrasts, replacing real contrast images with the synthesized images, one by one. If the segmentation performance did not drop even with the synthesized image, then it would mean that the contrast was replaceable. On the other hand, if there were a sufficient performance drop, that would mean that such contrasts were irreplaceable. From the results, the authors concluded that T1, T2, and T2F images are replaceable with synthetic images from collaGAN, which shows almost no difference in performance. However, a clear distinction was seen for a gadolinium contrast agent injected T_1 -weighted image. From the quantitative results, we can again verify that Gd-enhanced contrasts are crucial for segmentation performance, and they are not replaceable with collaGAN imputation.

12.6 Summary and Outlook

For both unconditional synthesis and conditional synthesis, GANs have shown to be the perfect fit to such means thanks to their ability to learn probability distributions.

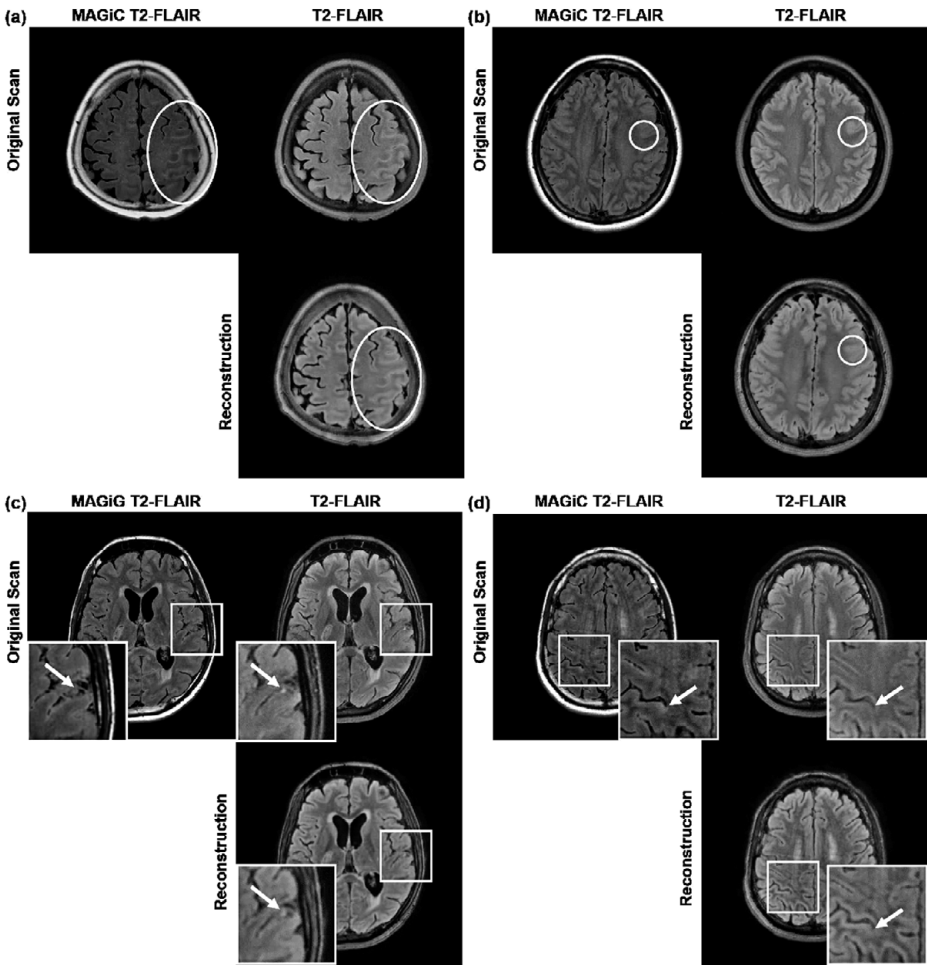


Figure 12.8 MR contrast imputation results on pathological data using collaGAN. © [2020] *Nature Machine Intelligence*, vol. 2, no. 1, pp. 34–42. Springer Nature.

For unconditional synthesis, the objective is to stochastically generate MR images of target contrast. Conditional synthesis refers to the case where the model is to learn a nonlinear mapping to the different MR tissue contrasts without altering the physiological information. Furthermore, by merging collaborative information about multiple contrast images, missing data imputation among many different domains is also effectively solved with GANs. Although promising results have been shown, developments in the area are still at an early stage. Interesting research directions have been proposed in previous work, which include application of the more advanced methods and rigorous validation in clinical settings. In effect, MRI image synthesis techniques should be able to reduce the burden of costly MR scans, benefiting both patients and hospitals.

References

- [1] S. Roy, A. Carass, and J. Prince, “A compressed sensing approach for MR tissue contrast synthesis,” in *Proc. Conference on Information Processing in Medical Imaging*, vol. 22, 2011, pp. 371–383.
- [2] —, “Magnetic resonance image example-based contrast synthesis,” *IEEE Transactions on Medical Imaging*, vol. 32, no. 12, pp. 2348–2363, 2013.
- [3] Y. Huang, L. Beltrachini, L. Shao, and A. F. Frangi, “Geometry regularized joint dictionary learning for cross-modality image synthesis in magnetic resonance imaging,” in *Proc. International Workshop on Simulation and Synthesis in Medical Imaging*. Springer, 2016, pp. 118–126.
- [4] Y. Huang, L. Shao, and A. F. Frangi, “Cross-modality image synthesis via weakly-coupled and geometry co-regularized joint dictionary learning,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 3, pp. 815 – 827, 2018.
- [5] A. Jog, A. Carass, S. Roy, D. L. Pham, and J. L. Prince, “Random forest regression for magnetic resonance image synthesis,” *Medical Image Analysis*, vol. 35, pp. 475–488, 2017.
- [6] H. Van Nguyen, K. Zhou, and R. Vemulapalli, “Cross-domain synthesis of medical images using efficient location-sensitive deep network,” in *Proc. Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 677–684.
- [7] A. Chartsias, T. Joyce, M. Valerio Giuffrida, and S. Tsaftaris, “Multimodal MR synthesis via modality-invariant latent representation,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 3, pp. 803–814, 2018.
- [8] S. U. H. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, and T. Çukur, “Image synthesis in multi-contrast MRI with conditional generative adversarial networks,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 10, pp. 2375–2388, 2019.
- [9] D. Lee, J. Kim, W.-J. Moon, and J. C. Ye, “Collagan: Collaborative GAN for missing image data imputation,” in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2487–2496.
- [10] D. Lee, W.-J. Moon, and J. C. Ye, “Assessing the importance of magnetic resonance contrasts using collaborative generative adversarial networks,” *Nature Machine Intelligence*, vol. 2, no. 1, pp. 34–42, 2020.
- [11] F. Calimeri, A. Marzullo, C. Stamile, and G. Terracina, “Biomedical data augmentation using generative adversarial neural networks,” in *Proc. Conference on Artificial Neural Networks and Machine Learning*. Springer, 2017, pp. 626–634.
- [12] L. Zhang, A. Gooya, and A. F. Frangi, “Semi-supervised assessment of incomplete LV coverage in cardiac MRI using generative adversarial nets,” in *Proc. International Workshop on Simulation and Synthesis in Medical Imaging*. Springer, 2017, pp. 61–68.
- [13] B. Yu, L. Zhou, L. Wang, J. Fripp, and P. Bourgeat, “3D cGAN based cross-modality MR image synthesis for brain tumor segmentation,” in *Proc. IEEE 15th International Symposium on Biomedical Imaging*, 2018, pp. 626–630.
- [14] F. Liu, “SUSAN: Segment unannotated image structure using adversarial network,” *Magnetic Resonance in Medicine*, vol. 81, no. 5, pp. 3330–3345, 2018.
- [15] S. Olut, Y. H. Sahin, U. Demir, and G. Unal, “Generative adversarial training for MRA image synthesis using multi-contrast MRI,” in *Proc. Conference on Predictive Intelligence in Medicine*. Springer, 2018, pp. 147–154.

- [16] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” in *Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [17] I. Goodfellow, “Nips 2016 tutorial: Generative adversarial networks,” *arXiv:1701.00160*, 2016.
- [18] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *Proc. International Conference on Machine Learning*, 2017, pp. 214–223.
- [19] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, “Improved training of wasserstein GANs,” in *Advances in Neural Information Processing Systems*, 2017, pp. 5767–5777.
- [20] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, “Least squares generative adversarial networks,” in *Proc. IEEE International Conference on Computer Vision*, 2017, pp. 2794–2802.
- [21] J. H. Lim and J. C. Ye, “Geometric GAN,” *arXiv:1705.02894*, 2017.
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1125–1134.
- [23] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proc. IEEE International Conference on Computer Vision*, 2017, pp. 2223–2232.
- [24] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas, “StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks,” in *Proc. International Conference on Computer Vision*, 2017, pp. 5907–5915.
- [25] H. Wu, S. Zheng, J. Zhang, and K. Huang, “Gp-GAN: Towards realistic high-resolution image blending,” in *Proc. 27th ACM International Conference on Multimedia*, 2019, pp. 2487–2495.
- [26] C. Bermudez, A. J. Plassard, L. T. Davis, A. T. Newton, S. M. Resnick, and B. A. Landman, “Learning implicit brain MRI manifolds with deep learning,” in *Proc. Conference on Medical Imaging 2018: Image Processing*, vol. 10574, SPIE, 2018, pp. 408–414.
- [27] A. K. Mondal, J. Dolz, and C. Desrosiers, “Few-shot 3D multi-modal medical image segmentation using generative adversarial learning,” *arXiv*, 2018.
- [28] Y. Han and J. C. Ye, “Framing U-Net via deep convolutional framelets: Application to sparse-view CT,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1418–1429, 2018.
- [29] A. Beers, J. Brown, K. Chang, J. P. Campbell, S. Ostmo, M. F. Chiang, and J. Kalpathy-Cramer, “High-resolution medical image synthesis using progressively grown generative adversarial networks,” *arXiv:1805.03144*, 2018.
- [30] C. Bowles, L. Chen, R. Guerrero, P. Bentley, R. Gunn, A. Hammers, D. A. Dickie, M. V. Hernández, J. Wardlaw, and D. Rueckert, “GAN augmentation: Augmenting training data using generative adversarial networks,” *arXiv:1810.10863*, 2018.
- [31] G. Kwon, C. Han, and D. Shik Kim, “Generation of 3D brain MRI using auto-encoding generative adversarial networks,” in *Proc. Conference on Medical Image Computing and Computer Assisted Intervention*, vol. 11766. Springer, 2019, pp. 118–126.
- [32] H. Yu and X. Zhang, “Synthesis of prostate MR images for classification using capsule network-based GAN model,” *Sensors*, vol. 20, no. 20, p. 5736, 2020.

- [33] D. Abramian and A. Eklund, "Refacing: Reconstructing anonymized facial features using GANs," in *2019 IEEE 16th International Symposium on Biomedical Imaging*, 2019, pp. 1104–1108.
- [34] X. Gu, H. Knutsson, M. Nilsson, and A. Eklund, "Generating diffusion MRI scalar maps from T1 weighted images using generative adversarial networks," in *Image Analysis*. Springer, 2019, pp. 489–498.
- [35] B. Yu, L. Zhou, L. Wang, Y. Shi, J. Fripp, and P. Bourgeat, "Ea-GANs: Edge-aware generative adversarial networks for cross-modality MR image synthesis," *IEEE Transactions on Medical Imaging*, vol. 38, no. 7, pp. 1750–1762, 2019.
- [36] X. Liu, F. Xing, J. L. Prince, A. Carass, M. Stone, G. E. Fakhri, and J. Woo, "Dual-cycle constrained bijective vae-GAN for tagged-to-cine magnetic resonance image synthesis," *arXiv:2101.05439*, 2021.
- [37] A. Sharma and G. Hamarneh, "Missing MRI pulse sequence synthesis using multi-modal generative adversarial network," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1170–1183, 2020.
- [38] M. Yurt, S. U. Dar, A. Erdem, E. Erdem, K. K. Oguz, and T. Çukur, "mustGAN: Multi-stream generative adversarial networks for MR image synthesis," *Medical Image Analysis*, vol. 70, p. 101944, 2021.
- [39] B. E. Dewey, C. Zhao, J. C. Reinhold, A. Carass, K. C. Fitzgerald, E. S. Sotirchos, S. Saidha, J. Oh, D. L. Pham, P. A. Calabresi, P. C. van Zijl, and J. L. Prince, "DeepHarmony: A deep learning approach to contrast harmonization across scanner changes," *Magnetic Resonance Imaging*, vol. 64, pp. 160–170, 2019.
- [40] L. Xiang, Y. Chen, W. Chang, Y. Zhan, W. Lin, Q. Wang, and D. Shen, "Deep-learning-based multi-modal fusion for fast MR reconstruction," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 7, pp. 2105–2114, 2019.
- [41] A. Falvo, D. Communiello, S. Scardapane, M. Scarpiniti, and A. Uncini, "A multimodal dense U-Net for accelerating multiple sclerosis MRI," in *Proc. IEEE 29th (MLSP)*. IEEE, 2019.
- [42] S. Dar, M. Yurt, M. Shahdloo, M. Ildiz, B. Tinaz, and T. Cukur, "Prior-guided image reconstruction for accelerated multi-contrast MRI via generative adversarial networks," *IEEE Journal on Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1072–1087, 2020.
- [43] W.-J. Do, S. Seo, Y. Han, J. C. Ye, S. H. Choi, and S.-H. Park, "Reconstruction of multicontrast MR images through deep learning," *Medical Physics*, vol. 47, no. 3, pp. 983–997, 2020.
- [44] K. H. Kim, W. J. Do, and S. H. Park, "Improving resolution of MR images with an adversarial network incorporating images with different contrast," *Medical Physics*, vol. 45, no. 7, pp. 3120–3131, 2018.
- [45] H. Lan, A. W. Toga, and F. Seppehrband, "SC-GAN: 3D self-attention conditional GAN with spectral normalization for multi-modal neuroimaging synthesis," *bioRxiv*, 2020.
- [46] M. Yurt, M. Özbey, S. U. H. Dar, B. Tinaz, K. K. Oğuz, and T. Çukur, "Progressively volumetrized deep generative models for data-efficient contextual learning of MR image recovery," *arXiv preprint arXiv:2011.13913*, 2020.
- [47] M. Yurt, S. U. H. Dar, B. Tinaz, M. Özbey, and T. Çukur, "Semi-supervised learning of mutually accelerated multi-contrast MRI synthesis without fully-sampled ground-truths," *arXiv:2011.14347*, 2020.

- [48] A. N. Baraldi and C. K. Enders, “An introduction to modern missing data analyses,” *Journal of School Psychology*, vol. 48, no. 1, pp. 5–37, 2010.
- [49] C. K. Enders, *Applied Missing Data Analysis*. Guilford Press, 2010.
- [50] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, “StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8789–8797.
- [51] L. N. Tanenbaum, A. J. Tsiouris, A. N. Johnson, T. P. Naidich, M. C. DeLano, E. R. Melhem, P. Quarterman, S. Parameswaran, A. Shankaranarayanan, M. Goyen *et al.*, “Synthetic MRI for clinical neuroimaging: Results of the magnetic resonance image compilation (magic) prospective, multicenter, multireader trial,” *American Journal of Neuroradiology*, vol. 38, no. 6, pp. 1103–1110, 2017.
- [52] A. Hagiwara, M. Warntjes, M. Hori, C. Andica, M. Nakazawa, K. K. Kumamaru, O. Abe, and S. Aoki, “SyMRI of the brain: Rapid quantification of relaxation rates and proton density, with synthetic MRI, automatic brain segmentation, and myelin measurement,” *Investigative Radiology*, vol. 52, no. 10, p. 647, 2017.
- [53] A. Hagiwara, M. Hori, K. Yokoyama, M. Takemura, C. Andica, T. Tabata, K. Kamagata, M. Suzuki, K. Kumamaru, M. Nakazawa *et al.*, “Synthetic MRI in the detection of multiple sclerosis plaques,” *American Journal of Neuroradiology*, vol. 38, no. 2, pp. 257–263, 2017.