

Prior-Guided Image Reconstruction for Accelerated Multi-Contrast MRI via Generative Adversarial Networks

Salman U.H. Dar, *Student Member, IEEE*, Mahmut Yurt, *Student Member, IEEE*, Mohammad Shahdloo, M. Emrullah Ildiz, Berk Tınaz, and Tolga Çukur*, *Senior Member, IEEE*

Abstract—Multi-contrast MRI acquisitions of an anatomy enrich the magnitude of information available for diagnosis. Yet, excessive scan times associated with additional contrasts may be a limiting factor. Two mainstream frameworks for enhanced scan efficiency are reconstruction of undersampled acquisitions and synthesis of missing acquisitions. Recently, deep learning methods have enabled significant performance improvements in both frameworks. Yet, reconstruction performance decreases towards higher acceleration factors with diminished sampling density at high-spatial-frequencies, whereas synthesis can manifest artefactual sensitivity or insensitivity to image features due to the absence of data samples from the target contrast. Here we propose a new approach for synergistic recovery of undersampled multi-contrast acquisitions based on conditional generative adversarial networks. The proposed method mitigates the limitations of pure learning-based reconstruction or synthesis by utilizing three priors: shared high-frequency prior available in the source contrast to preserve high-spatial-frequency details, low-frequency prior available in the undersampled target contrast to prevent feature leakage/loss, and perceptual prior to improve recovery of high-level features. Demonstrations on brain MRI datasets from healthy subjects and patients indicate the superior performance of the proposed method compared to pure reconstruction and synthesis methods. The proposed method can help improve the quality and scan efficiency of multi-contrast MRI exams.

Index Terms—Generative Adversarial Network (GAN), Synthesis, Reconstruction, Multi Contrast, Magnetic Resonance Imaging (MRI), Prior.

I. INTRODUCTION

MAGNETIC resonance imaging (MRI) is a preferred modality for assessment of soft tissues due to diversity of contrasts that it can provide. A typical MRI protocol comprises a set of pulse sequences that capture images of the same anatomy under different contrasts, with the aim to enhance diagnostic information. For instance, in neuroimaging protocols, T_1 -weighted images are useful for delineation of

gray and white matter, whereas T_2 -weighted images are more useful for delineation of fluids and fat. Although acquisition of multiple distinct contrasts is desirable, it may not be feasible due to scan time limitations or uncooperative patients. Thus, methods for accelerating MRI acquisitions without compromising image quality are of great interest for multiple-contrast applications.

The predominant approach for accelerated MRI relies on undersampled k-space acquisitions for scan time reduction, and on reconstruction algorithms for recovery of missing samples based on the collected evidence (i.e., acquired samples) [1]–[5]. Given the compressible nature of MR images, the state-of-the-art approach is sparse recovery [3], [4], which employs variable-density random undersampling in k-space to capture most of the energy in the MR images while ensuring low coherence of aliasing artifacts. The inverse problem of image reconstruction from sub-Nyquist sampled data is then solved via regularization from known transform domains [3], [4], learned transform domains [6] or end-to-end deep neural networks [7]–[22]. Despite the promise of deep models for image reconstruction, the evidence collected on the target MR image diminishes towards high acceleration factors due to undersampling. In turn, this significantly degrades the reconstruction performance, and causes loss in particularly high-spatial-resolution image features that may be relevant for diagnosis.

A fundamentally different approach for accelerated MRI is to perform fully-sampled acquisitions of a subset of the desired contrasts (i.e., source contrasts), and then to synthesize missing contrasts (i.e., target contrasts). This approach requires an intensity-based mapping model estimated using a collection of image pairs in both source and target contrast [23]–[50]. The model can be based on sparse linear mapping between source and target patches [32], or deep neural networks for enhanced accuracy [28], [34]–[36], [39]–[50]. Although deep models for synthesis are promising, local inaccuracies may occur in synthesized images when the source contrast is less sensitive to differences in relaxation parameters of two tissues compared to the target contrast, or vice versa. For instance, inflammation can be more clearly delineated from normal tissues in T_2 -weighted as opposed to T_1 -weighted images. In such cases, synthesized images might contain artificial pathology or fail to depict existing pathology.

Here, we propose a new approach for synergistic recovery of undersampled multi-contrast MRI acquisitions by

This work was supported in part by a European Molecular Biology Organization Installation Grant (IG 3028), by a TUBITAK 1001 Grant (118E256), by a TUBA GEBIP fellowship, and by a BAGEP fellowship awarded to T. Çukur. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan GPUs used in this study. *Asterisk indicates the corresponding author.*

S.U.H. Dar, M. Yurt, M. Shahdloo, M.E. Ildiz, and B. Tınaz are with the Department of Electrical and Electronics Engineering and the National Magnetic Resonance Research Center, Bilkent University, Bilkent, 06800 Ankara, Turkey.

T. Çukur is with the Department of Electrical and Electronics Engineering, with the National Magnetic Resonance Research Center (UMRAM), and with the Neuroscience Program, Bilkent University, Bilkent, 06800 Ankara, Turkey (e-mail: cukur@ee.bilkent.edu.tr).

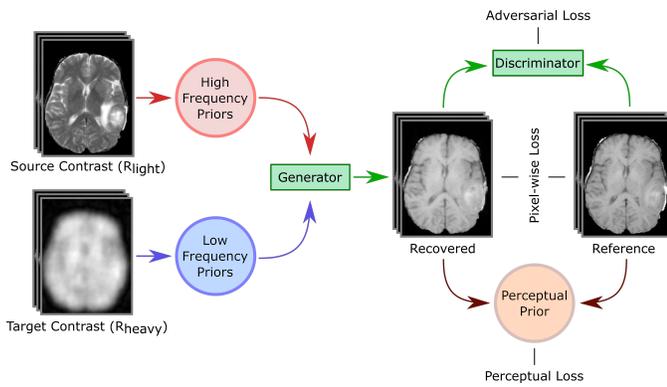


Fig. 1. The proposed rsGAN method synergistically recovers undersampled multi-contrast MRI acquisitions by complementarily using three priors: shared high-frequency priors available in fully-sampled or lightly undersampled acquisitions of one or more source contrasts to preserve high-spatial-frequency details, low-frequency priors available in highly undersampled acquisitions of one or more target contrasts to prevent feature leakage/loss, and a perceptual prior to improve recovery of high-level features. The input-to-output mapping is implemented using a conditional adversarial network with a generator and a discriminator. The generator learns to recover realistic high-quality target-contrast images by minimizing a pixel-wise, a perceptual and an adversarial loss function. The discriminator learns to discriminate between synthetic and real pairs of multi-contrast images by maximizing the adversarial loss function.

complementarily using reconstruction and synthesis models. The reconstruction branch takes as a prior the low-spatial-frequency information available in the collected evidence for the target contrast, whereas the synthesis branch takes as a prior the high-spatial-frequency information available in the fully sampled or lightly undersampled source contrast. These low-level spatial-frequency priors are complemented with a perceptual prior that improves recovery of higher-level image features [51]. The input-to-output mapping is implemented using conditional generative adversarial networks (GAN), which were recently shown to outperform traditional deep network models for image reconstruction [7]–[9], [14] and synthesis tasks [39], [40]. The proposed reconstructing-synthesizing GAN (rsGAN) contains a generator for estimating the target-contrast image given heavily undersampled target-contrast evidence and either fully sampled or lightly undersampled source-contrast image; and a discriminator to ensure that recovered images are as realistic as possible [52]. Low-spatial-frequencies are densely sampled in both target and source acquisitions, but reconstructions of the target contrast will inherently focus on low-frequency information in target acquisitions. Because the heavily undersampled target contrast misses high-frequency samples at large, the source contrast serves as the primary basis of high-frequency information. The proposed rsGAN model learns to fuse this multitude of input information in a data-driven manner.

Deep neural networks were previously proposed for recovery of multi-contrast MR acquisitions where each acquisition was accelerated at an identical rate [53]–[55]. Despite improved recovery compared to isolated reconstruction of individual contrasts, joint reconstruction may still suffer from loss of high-spatial-frequency information towards higher acceleration factors. Deep neural networks were also proposed for enhanced recovery of target-contrast images by incorporating structural information from fully sampled images of a separate

contrast [56]–[58]. Compared to [56]–[58] that employ loss terms based on mean square/absolute errors or structural similarity, rsGAN leverages an adversarial loss that demonstrated improved capture of high-spatial-frequency information [39]. A recent, independent study proposed a GAN model for super-resolution in a target contrast acquisition via the aid of fully-sampled images of a source contrast [59]. There are several technical differences between rsGAN and the model in [59]. In [59] sources have to be fully sampled, whereas rsGAN also enables light undersampling of source contrasts. For improved recovery, rsGAN further includes a perceptual prior. Lastly, the proposed rsGAN architecture can handle multi-coil complex MRI datasets, and enable reliable recovery at acceleration factors up to 50.

We demonstrated the proposed approach on several datasets: two public datasets containing normal subjects, a public dataset containing patients with high- or low-grade glioma, and a multi-coil dataset containing normal subjects. To comparatively evaluate the proposed method, following competing methods were considered: a reconstructing network (rGAN) that recovers the target-contrast image given undersampled images of the targets contrasts accelerated at identical rates, a reconstructing network (jGAN) that recovers the target-contrast image given undersampled images of the both source and target contrasts accelerated at identical rates, a synthesizing network (sGAN) that synthesizes the target-contrast image given fully sampled images of the source contrast, a joint super resolution reconstructing network (sr-sGAN) [59] that recovers the target-contrast image given undersampled images of the target contrasts and fully sampled images of the source contrast, and a variant of rsGAN deprived of the perceptual prior (rsGAN⁻). Our results indicate that rsGAN yields enhanced performance compared to the competing methods. In particular, rsGAN enables higher acceleration factors compared to rGAN and jGAN since it more reliably recovers high-spatial-frequency information. Compared to sGAN, rsGAN achieves improved reliability against artificial feature loss or leakage since it uses collected evidence from the target contrast to prevent hallucination. Compared to sr-sGAN, rsGAN achieves enhanced recovery at low to intermediate acceleration factors (up to 20x). Compared to rsGAN⁻, rsGAN improves reliability of high-level features. Overall, the proposed approach can successfully recover MR images of at acceleration factors up to 50x in the target contrasts, enabling a significant improvement in multi-contrast MRI.

Contributions

- 1) To our knowledge, this is the first GAN-based architecture that simultaneously leverages low-spatial-frequency, high-spatial-frequency and perceptual priors to accelerate multi-contrast MRI acquisitions.
- 2) The proposed approach can enable high acceleration factors up to 50x by incorporating information from both source and target contrasts.
- 3) The proposed approach can successfully recover pathologies that are either missing in the source contrast or are not clearly visible in the undersampled acquisition.

tions of the target contrast.

- 4) The proposed approach can jointly reconstruct and synthesize the target contrast even when the source contrasts are moderately undersampled.

II. THEORY AND METHODS

A. Accelerated MRI

Two mainstream approaches that can be used to accelerate MR acquisitions and enhance the diversity of acquired contrasts are reconstruction of a target contrast given randomly undersampled acquisitions of the same contrast, and synthesis of a target contrast based on fully-sampled acquisitions of a distinct source contrast. Both approaches incorporate prior information about image structure to improve the conditioning of the inverse problem of recovering images of the target contrast. However, they differ fundamentally in the type of prior information used. The problem formulations for reconstruction and synthesis are overviewed below.

a) Reconstruction: In this case, MR acquisitions are accelerated commonly via variable-density random undersampling patterns:

$$F_u m_1 = y_{1a} \quad (1)$$

where F_u is the partial Fourier operator defined at the k-space sampling locations, and m_1 is the image of the target contrast, y_{1a} are the acquired k-space data. The reconstruction task is then to recover the target image given the collected evidence (i.e., acquired data). Note that the problem in Eq. 1 is ill-posed, thus successful recovery requires additional prior information about the image. In the CS framework, this prior information reflects the sparsity of the image in a known transform domain (i.e., wavelet, TV transforms). The prior can be incorporated into the inverse problem as a regularization term:

$$\widehat{m}_1 = \arg \min_{m_1} \lambda \|F_u m_1 - y_{1a}\|_2 + R(m_1) \quad (2)$$

where the first term enforces consistency of the reconstructed and acquired data in k-space, $R(m_1)$ is the regularization term reflecting the prior, and λ controls the relative weighting of data consistency against the prior. $R(m_1)$ typically involves the ℓ_0 or ℓ_1 -norm of transform coefficients.

Recent studies have proposed neural-network methods to adaptively learn both nonlinear transform domains directly from MRI data and how to recover images from these domains. In the training stage, a large dataset of pairs of undersampled and fully-sampled acquisitions are leveraged to learn the network-based solution to the inverse problem:

$$\mathcal{L}_{rec}(\theta) = E_{m_{1t}^u, m_{1t}} \|G(m_{1t}^u; \theta) - m_{1t}\|_p \quad (3)$$

where m_{1t}^u and m_{1t} represent undersampled and fully-sampled training images, $G(m_{1t}^u; \theta)$ is the reconstructed output of the neural network based on network parameters θ , and $\|\cdot\|_p$ denotes ℓ_p -norm (where p is typically 1 or 2). Once the network parameters that minimize the objective in Eq. 3 have been learned, the following optimization problem can be cast to obtain reconstructions of undersampled acquisitions:

$$\widehat{m}_1 = \arg \min_{m_1} \lambda \|F_u m_1 - y_{1a}\|_2 + \|G(m_1^u; \theta^*) - m_1\|_2 \quad (4)$$

where m_1^u is the undersampled image, $G(m_1^u; \theta^*)$ is the reconstruction by the trained network with parameters θ^* , and \widehat{m}_1 is final recovered image. In Eq. 4, the first term again enforces consistency of reconstructed and acquired data. The second term is analogous to $R(m_1)$ in Eq. 2, and it enforces consistency of the recovered image to the network reconstruction.

b) Synthesis: In the synthesis case, fully-sampled images of the source contrast are assumed to be available. The task is then to recover target-contrast images (m_1) given source-contrast images (m_2) of the same anatomy. A learning-based procedure is used to estimate a mapping between the source and target contrast images. In the training stage, a large dataset of pairs of fully-sampled images from the source and target contrasts are used (m_{2t}, m_{1t}). In the CS-based synthesis framework, patch-based dictionaries (Φ_2, Φ_1) are formed for both source and target contrasts using m_{2t} and m_{1t} . These dictionaries are analogous to the sparsifying transform domains used in CS reconstructions. The aim is to express each patch in the source contrast images m_2 as a sparse linear combination of transform coefficients of the corresponding dictionary atoms:

$$\alpha(j) = \arg \min_{\alpha(j)} \|m_2(j) - \Phi_2 \cdot \alpha(j)\|_2 + \|\alpha(j)\|_1 \quad (5)$$

where $\alpha(j)$ is the learned combination coefficients for the j th patch, $m_2(j)$ denotes the j th patch in the source contrast, and Φ_2 denotes the dictionary formed using patches from m_{2t} . The first term ensures consistency of the synthesized patch to the true patch. The second term enforces sparsity of the vector of combination coefficients. Once the combination is learned, it can be used to synthesize target contrast images:

$$\widehat{m}_1(j) = \Phi_1 \cdot \alpha(j) \quad (6)$$

where Φ_1 denotes the dictionary formed using patches from m_{1t} , and $\widehat{m}_1(j)$ is the j th patch of the final synthesized image.

Recent studies have proposed neural-network methods to directly learn an adaptive, non-linear mapping from the source contrast to the target contrast. In the training stage, network parameters are optimized based on a loss function that reflects the error between the network output and the true target image:

$$\mathcal{L}_{synth}(\theta) = E_{m_{1t}, m_{2t}} \|G(m_{2t}; \theta) - m_{1t}\|_p \quad (7)$$

where m_{1t} and m_{2t} represent pairs of source and target images, and $G(m_{2t}; \theta)$ is the mapping from source to target contrast characterized by parameters θ . Once the network parameters that minimize the objective in Eq. 7 are learned, the network output can be directly calculated to obtain the synthesis results:

$$\widehat{m}_1 = G(m_2; \theta^*) \quad (8)$$

where \widehat{m}_1 is the prediction using the mapping $G(m_2; \theta^*)$ with parameters θ^* . Unlike the reconstruction task, here there is no evidence that has been collected about the target contrast. Therefore, no optimization procedures are needed for synthesis in the testing stage.

B. Joint Reconstruction-Synthesis via Conditional GANs

In the reconstruction task, the inverse problem solution uses undersampled acquisitions of the target contrast as evidence, and intrinsic image properties such as sparsity as prior. As the acceleration factor grows, evidence becomes scarce particularly towards high spatial frequencies that are sparsely covered by variable-density patterns. This in turn elevates the degree of aliasing artifacts; and if heavier weighting is given to the prior as a remedy, important features may be lost in the recovered images. Meanwhile, in the synthesis task, the inverse problem solution uses fully-sampled acquisitions of a distinct source contrast of the same anatomy as a prior. When the source and target contrasts exhibit similar levels of sensitivity to differences in tissue parameters, this prior can enable successful solution of the inverse problem. However, when the source and target show differential sensitivity, then features that are not supposed to be in the target may leak from the source onto the synthesized image, or features that must be present in the target may be missed.

To address the limitations of pure reconstruction or synthesis, we proposed to synergistically combine the two approaches with the aim to enhance recovery of multi-contrast MRI images. As such, the proposed approach consists of two branches: (1) A reconstruction branch that aggregates information from the target contrasts in the form of magnitude and phase images. (2) A synthesis branch that aggregates information from the source contrasts in the form of magnitude images.

Given k target contrasts and $n-k$ source contrasts, the joint recovery problem can be formulated as:

$$\begin{aligned} \hat{m}_{1,2,3,\dots,n} = \arg \min_{m_{1,2,\dots,n}} & \lambda \sum_{i=1}^k \|F_u m_i^{hu} - y_{ia}\|_2 + \\ & \lambda \sum_{j=k+1}^n \|F_u m_j^{lu} - y_{ja}\|_2 + \\ & R(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu}) \quad (9) \end{aligned}$$

where $R(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu})$ is a regularization term based on prior information, m_i^{hu} is the i th contrast that is heavily undersampled (i.e., target contrast), and m_j^{lu} is j th contrast that is lightly undersampled (i.e., source contrast), and y_{ia} denotes the acquired data for the i th contrast. We recast Eq. 9 using a neural-network based formulation:

$$\begin{aligned} \hat{m}_{1,2,3,\dots,n} = \arg \min_{m_{1,2,\dots,n}} & \lambda \sum_{i=1}^k \|F_u m_i^{hu} - y_{ia}\|_2 + \\ & \lambda \sum_{j=k+1}^n \|F_u m_j^{lu} - y_{ja}\|_2 + \\ & \sum_{l=1}^n \|G(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu}; \theta^*)[l] - m_l\|_2 \quad (10) \end{aligned}$$

Here, multiple separate channels for network output are considered since multiple contrast images can be recovered simultaneously. In Eq. 10, $G(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu}; \theta^*)[l]$ denotes the l th channel of the network output, among a total of

n channels for the entire set of contrasts. The first two terms respectively enforce the consistency of reconstructed data to acquired data in the target and source contrasts. The last term enforces consistency of the network outputs to the recovered images. Solution of Eq. 10 yields estimates of the images for each contrast separately as:

$$y_{ir}(k) = \begin{cases} \frac{F\{G(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu}; \theta^*)[i]\}(k) + \lambda y_{ia}(k)}{1 + \lambda} & \text{if } k \in \Omega \\ F\{G(m_1^{hu}, \dots, m_k^{hu}, m_{k+1}^{lu}, \dots, m_n^{lu}; \theta^*)[i]\}(k) & \text{otherwise} \end{cases}$$

$$\hat{m}_i = F^{-1}\{y_{ir}\} \quad (11)$$

where y_{ir} denotes the k -space representation of the image for the i th contrast, Ω is the set of acquired k -space samples, F is the Fourier transform operator, and F^{-1} is the inverse Fourier transform operator. The solution stated above performs two subsequent projections on the input images. The first projection takes undersampled acquisitions to generate the network predictions. The second projection enforces data consistency between data sampled that were originally acquired and those that are predicted by the network.

Based on the recent progress by generative adversarial networks in MR image synthesis and reconstruction tasks, we chose to build the joint recovery network using a conditional GAN architecture. Our network contains two subnetworks: a generator and a discriminator. The task of the generator is to learn a mapping from undersampled acquisitions onto fully-sampled acquisitions of source and target images. Both synthesis and reconstruction branches are provided to the generator part of the network as separate input channels. During the training, the network learns to adaptively fuse this information in a data-driven way. Meanwhile, the task of the discriminator is to differentiate between the images predicted by the generator and the actual images. As such, an adversarial loss function is typically used to train both subnetworks. Here for stabilized training, we used adversarial loss as in LSGAN [60]:

$$\begin{aligned} \mathcal{L}_{condAdv}(\theta_D, \theta_G) = & -E_{\mathbf{m}_t} [(D(\mathbf{m}_t; \theta_D) - 1)^2] \\ & -E_{\mathbf{m}_t^{hu}, \mathbf{m}_t^{lu}} [D(G(\mathbf{m}_t^{hu}, \mathbf{m}_t^{lu}; \theta_G))^2] \quad (12) \end{aligned}$$

where \mathbf{m}_t represents the MR images aggregated across n contrasts (m_1, m_2, \dots, m_n) in the training dataset, \mathbf{m}_t^{hu} represents the heavily undersampled acquisitions aggregated across k target contrasts (m_1, m_2, \dots, m_k), \mathbf{m}_t^{lu} represents the lightly undersampled acquisitions aggregated across $n-k$ source contrasts ($m_{k+1}, m_{k+2}, \dots, m_n$), G is the generator with parameters θ_G , D is the discriminator with parameters θ_D , and $\mathcal{L}_{Adv}(\theta_D, \theta_G)$ is the adversarial loss function. To ensure reliable recovery in each channel, a pixel-wise loss

function was incorporated to the generator:

$$\mathcal{L}_{L1}(\theta_G) = E_{\mathbf{m}_t, \mathbf{m}_t^{\text{hu}}, \mathbf{m}_t^{\text{lu}}} \left[\left\| G(\mathbf{m}_t, \mathbf{m}_t^{\text{hu}}, \mathbf{m}_t^{\text{lu}}; \theta_G) - \mathbf{m}_t \right\|_1 \right] \quad (13)$$

Recent studies on MRI reconstruction and synthesis suggest that incorporating an additional prior in the form of a perceptual loss can further enhance the image quality [39], [40]. The perceptual loss relies on high-level features extracted via networks pretrained on natural images for more general tasks. Following [51], we extracted feature maps right before the second max-pooling layer of the VGG16 model trained on the ImageNet dataset [61] for object classification. The loss function can be expressed as:

$$\mathcal{L}_{perc}(\theta_G) = E_{\mathbf{m}_t, \mathbf{m}_t^{\text{hu}}, \mathbf{m}_t^{\text{lu}}} \left[\left\| V(G(\mathbf{m}_t, \mathbf{m}_t^{\text{hu}}, \mathbf{m}_t^{\text{lu}}; \theta_G)) - V(\mathbf{m}_t) \right\|_1 \right] \quad (14)$$

where $V(\cdot)$ represents the features extracted via VGG16. The adversarial, pixel-wise and perceptual losses are finally combined to train the proposed reconstructing-synthesizing GAN (rsGAN) model:

$$\begin{aligned} \mathcal{L}_{rsGAN}(\theta_D, \theta_G) = & \lambda_p \mathcal{L}_{L1}(\theta_G) + \\ & \lambda_{perc} \mathcal{L}_{perc}(\theta_G) + \\ & \mathcal{L}_{condAdv}(\theta_D, \theta_G) \end{aligned} \quad (15)$$

where λ_p and λ_{perc} are the relative weightings of the pixel-wise and perceptual loss functions.

C. Competing Methods

To evaluate the effectiveness of rsGAN, we compared it against other GAN architectures. A GAN trained to only perform synthesis of the target-contrast images based on the respective source-contrast images. Source-contrast images were taken to be fully-sampled, high-quality images. We will refer to this network as the synthesizing GAN (sGAN). A GAN trained to only perform reconstruction of the target-contrast images based on undersampled acquisitions of all target contrasts accelerated at identical rates. We will refer to this network as the reconstructing GAN (rGAN). A GAN trained to only perform reconstruction of the target-contrast images based on undersampled acquisitions of all source and target contrasts accelerated at identical rates. We will refer to this network as the joint reconstructing GAN (jGAN). In the public datasets, rsGAN was also compared against a variant of rsGAN deprived of the perceptual prior, referred to as rsGAN⁻, and a GAN trained to only perform recovery of the target-contrast images based on fully sampled acquisitions of source contrasts and low-resolution acquisitions of target contrasts, referred to as the super-resolution synthesis GAN (sr-sGAN). Note that sGAN, jGAN, rGAN and sr-sGAN were also trained using the perceptual prior.

D. MRI Datasets

We demonstrated the proposed approach on three different public datasets and a multi-coil dataset containing multi-contrast MRI images. The public datasets MIDAS [62] and IXI (<http://brain-development.org/ixi-dataset/>) comprised images collected in healthy normals. BRATS (<https://sites.google.com/site/brain-tumor-segmentation/home/brats2015>) comprised images collected in patients with low-grade glioma (LGG) or high-grade glioma (HGG). Relevant details about each dataset are given below.

a) *MIDAS Dataset*: T_1 -weighted and T_2 -weighted images in the MIDAS dataset were considered. Data from 40 subjects were analyzed. The scan protocols were as follows:

- 1) T_1 -weighted images: 3D Gradient-Echo sequence, repetition time (TR)=14ms, echo time (TE)=7.7ms, flip angle=25°, volume size=256×176×256, voxel dimensions=1mm×1mm×1mm.
- 2) T_2 -weighted images: 2D Spin-Echo sequence, repetition time (TR)=7730ms, echo time (TE)=80ms, flip angle=180°, volume size=256×192×256, voxel dimensions=1mm×1mm×1mm.

b) *IXI Dataset*: T_1 -weighted, T_2 -weighted and PD-weighted images in the IXI dataset were considered. Data from 40 subjects were analyzed.

The scan protocols were as follows:

- 1) T_1 -weighted images: repetition time (TR)=9.813ms, echo time (TE)=4.603ms, flip angle=8°, volume size=256×256×150, voxel dimensions=0.94mm×0.94mm×1.2mm.
- 2) T_2 -weighted images: repetition time (TR)=8178ms, echo time (TE)=100ms, flip angle=90°, volume size=256×256×130, voxel dimensions=0.94mm×0.94mm×1.2mm.
- 3) PD-weighted images: Repetition time (TR)=8178ms, echo time (TE)=8ms, flip angle=90°, volume size=256×256×130, voxel dimensions=0.94mm×0.94mm×1.2mm.

c) *BRATS Dataset*: T_1 -weighted, T_2 -weighted and FLAIR images in the BRATS dataset were considered. Data from 40 Glioma patients were analyzed. Since the data were acquired in various different sites, no single scan protocol existed. In BRATS, all contrasts were already pre-registered and skull-stripped as publicly shared.

d) *Multi-Coil MR Images*: T_1 -weighted, T_2 -weighted and PD-weighted brain images from 10 subjects were acquired at Bilkent University. Images were acquired on a 3T Siemens Tim Trio scanner (maximum gradient strength of 45mT/m and slew rate of 200 T/m/s) using a 32-channel receive only coil. The scan protocols were as follows:

- 1) T_1 -weighted images: 3D MP-RAGE sequence, repetition time (TR)=2000ms, echo time (TE)=5.53ms, flip angle=20°, volume size=256×192×88, voxel dimensions=1mm×1mm×2mm, acquisition time (TA)=6:26.
- 2) T_2 -weighted images: 3D Spin-Echo sequence, repetition time (TR)=1000ms, echo time (TE)=118ms, flip angle=90°, volume size=256×192×88, voxel

dimensions=1mm×1mm×2mm, acquisition time (TA)=17:39.

- 3) PD-weighted images: 3D Spin-Echo sequence, repetition time (TR)=750ms, echo time (TE)=12ms, flip angle=90°, volume size=256×192×88, voxel dimensions=1mm×1mm×2mm, acquisition time (TA)=13:14.

In the public datasets, (25, 5, 10) subjects were used for (training, validation, testing). In the IXI dataset, one test subject was discarded due to poor registration quality. Within each subject, around 100 central cross-sections that contained brain tissues and that were relatively free of artifacts were selected. Each model was trained using a batch size of 1. This corresponds to nearly 2400-2600 iterations per epoch. In the multi-coil dataset, (7, 1, 2) subjects were used for (training, validation, testing). For each subject, around 155 central cross-sections that contained brain tissues and that were relatively free of artifacts were selected. Models were trained using a batch size of 1. This corresponds to nearly 1085 iterations per epoch.

E. Image Registration

Since the multi-contrast volumes in the MIDAS, IXI and multi-coil datasets were unregistered, these images were registered before training and testing. For the MIDAS dataset, T₂-weighted images of each subject were registered onto T₁-weighted images of the same subject using a rigid transformation. Images were registered based on mutual information loss. For the IXI dataset, T₂- and PD-weighted images of each subject were registered onto T₁-weighted images of each subject using an affine transformation. In the multi-coil dataset, T₂- and PD-weighted images of each subject were registered onto T₁-weighted images of each subject using a rigid body transformation. Images were registered based on mutual information loss. Registrations were carried out using FSL [63], [64].

F. Undersampling Patterns

For heavily undersampled acquisitions of the target contrast, we examined acceleration factors in a broad range (R= 5x, 10x, 20x, 30x, 40x, 50x; Fig. 2). For lightly undersampled acquisitions of the source contrast, we examined acceleration factors in a relatively limited range (R=1x, 2x, 3x). For rsGAN and rGAN, variable-density undersampling was used [3]. The undersampling patterns were generated using bivariate normal probability density functions. Covariance of the density functions was separately adjusted for each value of R. Fully-sampled images were Fourier transformed, and then retrospectively sampled using the generated patterns. Distinct random patterns were generated for each subject within each dataset.

For proof of concept demonstration at high acceleration rates, we simulated 2D undersampling in the transversal plane for the public datasets. For multi-coil acquisitions, 2D undersampling was performed in the coronal plane on a 192x88 grid, so a relative narrower range of accelerations were considered (R=5x, 10x, 15x, 20x, 25x, 30x).

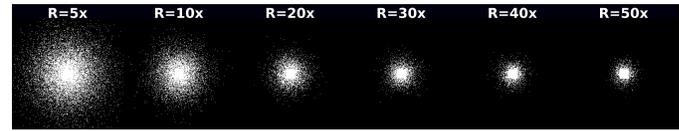


Fig. 2. Examples of undersampling patterns for acquisitions accelerated in a broad range (R= 5x, 10x, 20x, 30x, 40x, 50x)

G. Model Training Procedures

All GAN-based models were trained using an identical set of procedures. To train each conditional GAN, we adopted the generator and discriminator from [51] and [65]. The generator consisted of the following convolutional layers (Conv) connected in series: Conv (kernel-size=7, output-features=64, stride=1, activation=ReLU), Conv (kernel-size=3, output-features=64, stride=2, activation=ReLU), Conv (kernel-size=3, output-features=256, stride=2, activation=ReLU), 9x resnet blocks (kernel-size=3, output-features=256, stride=1, activation=ReLU), fractionally-strided Conv (kernel-size=3, output-features=128, stride=2, activation=ReLU), fractionally-strided Conv (kernel-size=3, output-features=64, stride=2, activation=ReLU), Conv (kernel-size=7, output-features=1, stride=1, activation=none). The discriminator consisted of the following convolutional layers (Conv) connected in series: Conv (kernel-size=4, output-features=64, stride=2, activation=leakyReLU), Conv (kernel-size=4, output-features=128, stride=2, activation=leakyReLU), Conv (kernel-size=4, output-features=256, stride=2, activation=leakyReLU), Conv (kernel-size=4, output-features=512, stride=1, activation=leakyReLU), Conv (kernel-size=4, output-features=1, stride=1, activation=none).

Generator and discriminator networks were trained for 100 epochs using the Adam optimizer [66], with decay rates for the first and second moment estimates set as 0.5 and 0.999. For the generator, the learning rate was set as 0.0002 for the initial 50 epochs and then linearly decayed to 0 during the remaining epochs. For the discriminator, the learning rate was set as 0.0001 for the first 50 epochs and then linearly decayed to 0 during the remaining epochs. Dropout regularization was used to enhance the generalizability of the network model, with a dropout rate of 0.5. Instance normalization was applied [67]. All model weights were randomly initialized based on a normally-distributed variable with 0 mean and 0.02 standard deviation.

The optimal weightings of pixel-wise loss (λ_p) and perceptual loss (λ_{perc}) terms were determined via a cross-validation procedure supplemented by visual inspection. Using the training data, separate models were obtained for λ_p in [10 150] and λ_{perc} in [10 150]. Weight selection was then performed by maximizing PSNR on the validation data. Recovered validation images were also visually inspected. When needed, selected weights were further fine-tuned to prevent low-quality recovery due to artifacts. Following these procedures, a common $\lambda_p = 100$ value was chosen for all datasets that yielded near-optimal results consistently across datasets and acceleration factors. Note that our optimum λ_p value closely matches weighting reported for conditional GAN models in the literature [39], [65]. Meanwhile, a separate λ_{perc} value was chosen for each dataset and for each acceleration factor. Relative weighting of data consistency against the prior

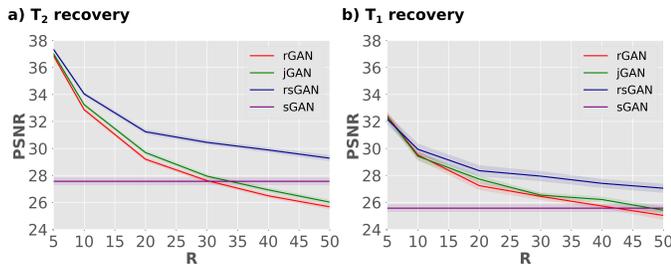


Fig. 3. The proposed rsGAN method was demonstrated for synergistic reconstruction-synthesis of T₁- and T₂-weighted images from the MIDAS dataset. The acquisition for the source contrast was fully sampled, and the acquisition for the target contrast was undersampled by R=5x, 10x, 20x, 30x, 40x, 50x. PSNR was measured between recovered and fully-sampled reference target-contrast images. (a) PSNR (mean±standard error) across the test subjects for rsGAN, rGAN, jGAN, and sGAN when T₁ is the source contrast and T₂ is the target contrast. (b) PSNR (mean±standard error) when T₂ is the source contrast and T₁ is the target contrast. The performance of sGAN remains constant across R since it does not use any evidence from the target-contrast acquisitions. As expected, the performance of rGAN, jGAN, and rsGAN gradually decreases for higher values of R where the evidence from the target contrast becomes scarce. However, rsGAN performs well even at very high acceleration factors..

(λ) was set to infinity.

Note that although the public datasets used in this study contain only coil combined magnitude images, the Fourier reconstructions of undersampled acquisitions are complex valued. Therefore, for each input contrast, two channels were designated to represent the magnitude and phase image components. For each target contrast, separate networks were trained to recover fully-sampled magnitude images. In the multi-coil dataset, first GCC [68] was used to reduce computational complexity by decreasing the number of coils from 32 to 5. For each input contrast, 5 channels were designated to represent magnitude components. In practice, the comparative performance of rsGAN models without and with phase depends on the benefits of added phase information against the disadvantages of fitting a more complex model. In the multi-coil dataset with 5 virtual coils, adding phase information for each individual contrast amounts to 5 extra input channels, considerably expanding model complexity. Since we observed that additional phase channels caused a slight decline in performance, we preferred to use rsGAN models without phase information in the multi-coil analyses of rsGAN. For each target contrast, separate networks were trained to recover fully-sampled coil-combined magnitude images. Reference coil-combined images were obtained by using coil sensitivity maps estimated via ESPiRiT [69].

To maximize model performance, a separate model was trained for each unique collection of source and target contrasts, and acceleration factors. For generalizing the rsGAN model to also handle light undersampling of the source contrast, a separate rGAN model was first trained to recover undersampled source acquisitions at each acceleration factor. In the testing phase, the reconstructed source contrast was then fed to the rsGAN model.

For multi-coil data, rsGAN was first trained to recover a coil-combined magnitude image for the target contrast from undersampled multi-coil magnitude images for source and target contrasts. Second, a coil-combined complex image for the target was obtained by adding onto the recovered magnitude

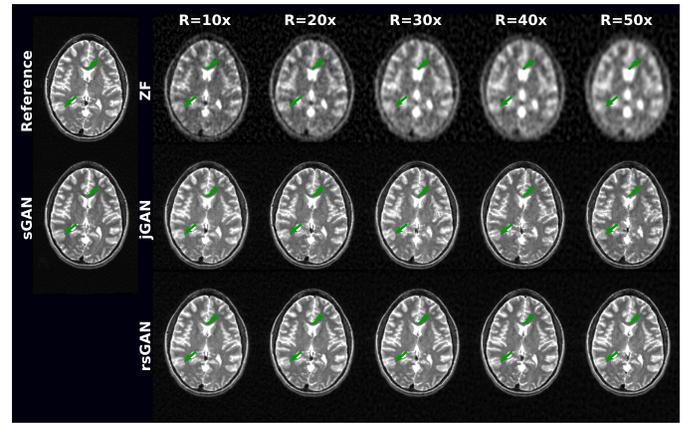


Fig. 4. T₂-weighted images in the MIDAS dataset were recovered from heavily undersampled acquisitions (R=10x, 20x, 30x, 40x, 50x). The acquisition for the source contrast (T₁-weighted) was fully sampled. Target-contrast images recovered by ZF (zero-filled Fourier reconstruction), sGAN, jGAN, and rsGAN are shown with the fully-sampled reference image. As the value of R increases the performance of jGAN degrades significantly. Meanwhile, rsGAN maintains high-quality recovered images due to use of additional information from the source contrast. Regions with enhanced recovery in rsGAN are marked with arrows.

image the phase of the coil-combined undersampled images of the target. Third, the coil-combined complex target image was back-projected onto individual coils using coil sensitivity maps. Data consistency was enforced on the resultant multi-coil complex target data, and a coil-combined complex target image was then obtained. As such, phase information in undersampled acquisitions was leveraged to enable data consistency projections.

H. Experiments

1) *Main Experiments*: To evaluate the comparative performance of the proposed approach, rsGAN, rGAN, jGAN, and sGAN were individually trained and tested on multi-contrast MRI datasets. Theoretically, as R approaches 1x, rsGAN, rGAN and jGAN should show nearly identical performance that is superior to sGAN since sGAN has no evidence collected about the target contrast. As R goes to infinity, rsGAN and sGAN should show nearly identical performance that is superior to rGAN, since no evidence from the target contrast will be available to any of the networks. In intermediate R values, we reasoned that rsGAN would outperform rGAN and jGAN in terms of reliability in recovery of high-frequency information since variable-density patterns suboptimally sample high spatial frequencies in the target contrast. We also reasoned that rsGAN would outperform sGAN especially when the source and target contrasts showed differential sensitivity to differences in tissue parameters. Based on these notions, we measured the performance of all four methods across a broad range of acceleration factors. To evaluate the effects of perceptual prior and variable-density sampling patterns, rsGAN was also compared against sr-sGAN in the public datasets. We reasoned that incorporation of the perceptual prior should enhance performance. We also reasoned that as R approaches 1, rsGAN should perform better than sr-sGAN since rsGAN contains more high-spatial frequency information from the target contrast. As R approaches infinity, rsGAN

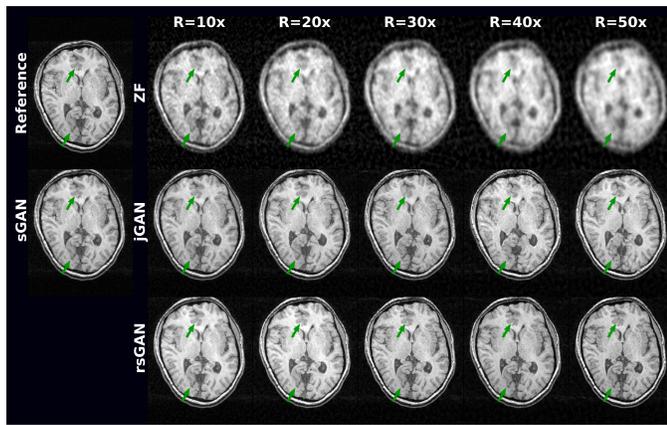


Fig. 5. T_1 -weighted images in the MIDAS dataset were recovered from heavily undersampled acquisitions ($R=10x, 20x, 30x, 40x, 50x$). The acquisition for the source contrast (T_2 -weighted) was fully sampled. Target-contrast images recovered by ZF, sGAN, jGAN, and rsGAN are shown with the fully-sampled reference image. Regions with enhanced recovery in rsGAN are marked with arrows.

should perform similar to sr-sGAN since the variable-density sampling patterns in rsGAN approach the central sampling patterns in sr-sGAN at these acceleration rates. In both MIDAS and BRATS datasets, we considered two main scenarios. First, T_1 -weighted acquisitions were taken as the source contrast ($R=1x$), and T_2 -weighted acquisitions were taken as the target contrast ($R=5x, 10x, 20x, 30x, 40x, 50x$). Second, T_2 -weighted acquisitions were taken as the source ($R=1x$), and T_1 -weighted acquisitions were taken as the target ($R=5x, 10x, 20x, 30x, 40x, 50x$).

Two distinct scenarios were examined in both IXI and multi-coil datasets. First, T_1 -weighted acquisitions were taken as the source contrast ($R=1x$), and both T_2 - and PD-weighted acquisitions were taken as the target contrasts ($R=5x, 10x, 20x, 30x, 40x, 50x$ in the IXI dataset, and $R=5x, 10x, 15x, 20x, 25x, 30x$ in the multi-coil dataset). Since T_2 - and PD-weighted acquisitions are typically performed using similar sequences, the acceleration factors for these two contrasts were always matched. Second, the source T_1 -weighted acquisitions were lightly undersampled ($R=2x, 3x$), and T_1 -, T_2 -, and PD-weighted images were jointly recovered. The overall scan time for an accelerated multi-contrast protocol depends on the distribution of R across contrasts, and individual scan times for all contrasts. To systematically examine scan efficiency, we measured recovery performance for jGAN and rsGAN with the same overall scan time. Analyses were performed on the in vivo multi-coil datasets for a fixed scan time of 250 sec, where T_1 was the source contrast and T_2 and PD were the target contrasts. For jGAN this corresponds to $R=8.9x$ across all contrasts, whereas for rsGAN this corresponds to $R_{T_1}=3x$ for the source contrast and $R=15x$ for target contrasts.

2) *Control Experiments*: Here, for more efficient model training, we preferred to focus on cross-sections that contained brain tissue. To rule out potential biases in model generalizability due to this selection, we conducted control experiments where rGAN, jGAN and rsGAN were trained on all available cross-sections in the IXI dataset without any selection (referred to as $rGAN_{All}$, $jGAN_{All}$ and $rsGAN_{All}$). These models were compared with rGAN, jGAN and rsGAN trained on the origi-

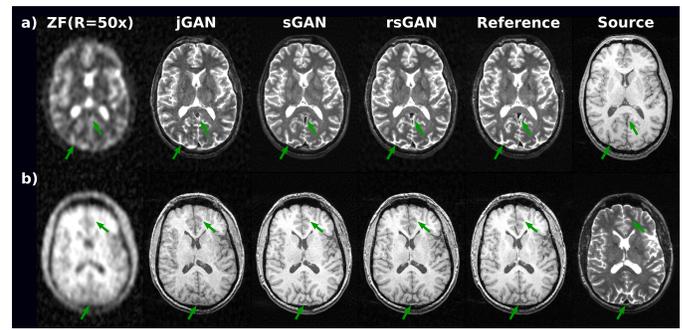


Fig. 6. Multi-contrast images in the MIDAS dataset were recovered, where the source contrast was fully sampled and the target contrast was undersampled at $R=50x$. Images were recovered using ZF, sGAN, jGAN and rsGAN. (a) Recovered T_2 -weighted images are shown along with the fully-sampled reference image and the source-contrast image. (b) Recovered T_1 -weighted images are shown along with the fully-sampled reference image and the source-contrast image. rsGAN yields visually accurate recovery of the target-contrast image compared to sGAN and jGAN. Sample regions that are better recovered by rsGAN are marked with arrows.

nally selected central cross-sections. Performance comparisons were carried out on independent test sets containing all cross-sections within subjects without any selection procedures.

In the public datasets containing coil-combined magnitude images, phase is only introduced during retrospective undersampling of k-space data, and the phase values are often small. Thus, in theory, an rsGAN model that receives as inputs only magnitude images should perform similarly to one that receives both magnitude and phase images. To test this prediction, we conducted additional analyses in the BRATS dataset where a variant of rsGAN model with only magnitude channels as inputs was trained, referred to as $rsGAN_m$. $rsGAN_m$ was then compared against rsGAN consisting of both magnitude and phase channels as inputs. T_1 was set as the source contrast and T_2 was set as the target contrast.

Here, rsGAN had different model complexity compared to the competing methods (rGAN, jGAN and sGAN). To rule out potential biases due to model complexity, we implemented additional control experiments in the BRATS dataset with rGAN, jGAN and sGAN models with matching complexity to rsGAN. Complexity was balanced across models by maintaining an identical number of input channels to the generator. In these experiments, T_1 was set as the source contrast and T_2 was set as the target contrast. Input to sGAN consisted of magnitude images of fully sampled T_1 contrast concatenated with magnitude and phase images of undersampled T_1 contrast. Input to rGAN consisted of magnitude and phase images of highly undersampled T_2 contrast concatenated with magnitude images of highly undersampled T_2 contrast. Input to jGAN consisted of magnitude and phase images of highly undersampled T_2 contrast concatenated with magnitude images of highly undersampled T_1 contrast. These models are referred to as $rGAN_{MC}$, $jGAN_{MC}$ and $sGAN_{MC}$. These models were compared with rsGAN, and regular rGAN, jGAN and sGAN.

In this study, rsGAN was mainly demonstrated for the recovery of T_1 -, T_2 -, and PD-weighted contrasts. Several diagnostic protocols also include FLAIR acquisitions. To examine the ability of rsGAN to recover FLAIR acquisitions, we also

TABLE I
QUALITY OF RECOVERED IMAGES IN THE MIDAS DATASET

	rGAN		jGAN		rsGAN($T_1=1$)	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
R=5x	36.88 ± 0.13	97.69 ± 0.10	37.03 ± 0.11	97.76 ± 0.10	37.35[†] ± 0.12	97.96[†] ± 0.10
R=10x	32.87 ± 0.14	95.18 ± 0.20	33.24 ± 0.12	95.60 ± 0.19	34.04[†] ± 0.13	96.30[†] ± 0.19
R=20x	29.20 ± 0.15	90.88 ± 0.34	29.69 ± 0.15	91.64 ± 0.33	31.23[†] ± 0.17	93.97[†] ± 0.30
R=30x	27.61 ± 0.15	87.98 ± 0.42	27.95 ± 0.14	88.82 ± 0.40	30.45[†] ± 0.16	93.25[†] ± 0.33
R=40x	26.49 ± 0.14	85.75 ± 0.47	26.93 ± 0.17	86.88 ± 0.49	29.89[†] ± 0.17	92.76[†] ± 0.35
R=50x	25.67 ± 0.17	83.77 ± 0.53	26.02 ± 0.16	84.87 ± 0.52	29.28[†] ± 0.21	92.09[†] ± 0.41
sGAN						
PSNR			SSIM (%)			
27.56 ± 0.27			90.91 ± 0.48			

PSNR and %SSIM values (mean \pm standard error) across the test subjects are listed for sGAN, rGAN, jGAN, and rsGAN. T_1 -weighted acquisitions were taken as the source contrast, and T_2 -weighted acquisitions were taken as the target contrast. The highest PSNR and SSIM values in each row are marked in bold font, and the significantly better performing values ($p < 0.05$) are marked with the '†' symbol.

trained models for recovery of FLAIR images in the BRATS dataset. T_1 was used as the source contrast, and FLAIR was used as the target contrast.

To maximize performance for individual contrasts, here the recovery of each target contrast was taken as a separate task. When multi-target-contrast images were considered, a separate rsGAN model was constructed to recover each target contrast. To assess the benefits of this strategy, we conducted additional experiments in the BRATS dataset where two distinct sets of rsGAN models were constructed. The first set consisted of the original rsGAN models (named rsGAN₁) trained to recover target contrasts individually (one model recovering T_2 from undersampled T_2 and fully-sampled T_1 acquisitions, and another model recovering FLAIR from undersampled FLAIR and fully-sampled T_1 acquisitions). The second set consisted of a unified rsGAN model (named rsGAN₂) trained to jointly

TABLE II
QUALITY OF RECOVERED IMAGES IN THE MIDAS DATASET

	rGAN		jGAN		rsGAN($T_2=1$)	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
R=5x	32.41 ± 0.31	94.11 ± 0.29	32.35 ± 0.31	94.33 ± 0.30	32.17 ± 0.39	94.53[†] ± 0.31
R=10x	29.52 ± 0.33	90.75 ± 0.40	29.42 ± 0.31	91.32 ± 0.41	29.94[†] ± 0.41	92.23[†] ± 0.45
R=20x	27.24 ± 0.31	86.15 ± 0.50	27.72 ± 0.26	87.77 ± 0.52	28.35[†] ± 0.39	90.03[†] ± 0.52
R=30x	26.44 ± 0.25	83.91 ± 0.60	26.53 ± 0.18	85.44 ± 0.55	27.94[†] ± 0.36	89.28[†] ± 0.56
R=40x	25.73 ± 0.16	82.10 ± 0.66	26.21 ± 0.21	83.98 ± 0.65	27.41[†] ± 0.31	88.66[†] ± 0.60
R=50x	25.03 ± 0.32	80.52 ± 0.66	25.41 ± 0.27	82.43 ± 0.65	27.05[†] ± 0.32	88.21[†] ± 0.66
sGAN						
PSNR			SSIM (%)			
25.57 ± 0.27			85.94 ± 0.82			

T_2 -weighted acquisitions were taken as the source contrast, and T_1 -weighted acquisitions were taken as the target contrast.

TABLE III
QUALITY OF RECOVERED IMAGES IN THE BRATS DATASET

	rGAN		jGAN		rsGAN($T_1=1$)	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
R=5x	35.54 ± 0.77	98.58 ± 0.10	36.21 ± 0.75	98.69 ± 0.07	36.15 ± 0.74	98.77 ± 0.06
R=10x	32.61 ± 0.66	97.34 ± 0.11	32.81 ± 0.64	97.49 ± 0.10	33.43[†] ± 0.64	97.73[†] ± 0.08
R=20x	30.47 ± 0.60	95.38 ± 0.18	31.17 ± 0.57	95.79 ± 0.15	31.54[†] ± 0.55	96.55[†] ± 0.12
R=30x	29.58 ± 0.52	93.91 ± 0.22	29.87 ± 0.54	94.38 ± 0.20	30.97[†] ± 0.49	95.79[†] ± 0.15
R=40x	28.84 ± 0.57	92.85 ± 0.26	28.94 ± 0.57	93.34 ± 0.23	30.00[†] ± 0.50	95.28[†] ± 0.17
R=50x	28.12 ± 0.53	91.81 ± 0.29	28.35 ± 0.50	92.39 ± 0.25	30.09[†] ± 0.39	94.88[†] ± 0.19
sGAN						
PSNR			SSIM (%)			
24.78 ± 0.48			91.69 ± 0.44			

T_1 -weighted acquisitions were taken as the source contrast, and T_2 -weighted acquisitions were taken as the target contrast.

recover T_2 and FLAIR from undersampled T_2 and FLAIR, and fully-sampled T_1 acquisitions. The two sets of models were compared in terms of average performance in recovery of T_2 and FLAIR images.

All network models and conventional reconstruction and synthesis techniques were trained and tested on the same instances of data and undersampling patterns. To quantitatively assess the quality of recovered images, the fully-sampled reference images were used. All images were first normalized to the range [0 1]. Then, peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) were calculated between the recovered and reference images. Statistical significance of differences in PSNR and SSIM between methods were assessed via a nonparametric Wilcoxon signed-rank test. In the public datasets, the statistical significance tests were performed across test subjects. In the multi-coil dataset, due to limited number of test subjects, the statistical significance tests were performed across cross-sections.

TABLE IV
QUALITY OF RECOVERED IMAGES IN THE BRATS DATASET

	rGAN		jGAN		rsGAN($T_2=1$)	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
R=5x	35.93 ± 0.57	98.71 ± 0.08	36.20 ± 0.62	98.78 ± 0.09	36.46 ± 0.60	98.86[†] ± 0.10
R=10x	32.56 ± 0.51	97.55 ± 0.14	32.63 ± 0.58	97.60 ± 0.17	33.58[†] ± 0.58	97.96[†] ± 0.16
R=20x	29.80 ± 0.51	95.82 ± 0.19	30.12 ± 0.54	96.03 ± 0.20	31.18[†] ± 0.51	96.82[†] ± 0.22
R=30x	28.52 ± 0.49	94.55 ± 0.19	28.94 ± 0.48	94.94 ± 0.20	30.16[†] ± 0.55	96.29[†] ± 0.25
R=40x	27.77 ± 0.49	93.47 ± 0.19	28.03 ± 0.52	93.96 ± 0.21	29.53[†] ± 0.61	95.86[†] ± 0.29
R=50x	27.04 ± 0.46	92.53 ± 0.21	27.26 ± 0.56	93.10 ± 0.22	29.19[†] ± 0.61	95.58[†] ± 0.30
sGAN						
PSNR			SSIM (%)			
24.78 ± 0.81			91.69 ± 0.44			

T_2 -weighted acquisitions were taken as the source contrast, and T_1 -weighted acquisitions were taken as the target contrast.

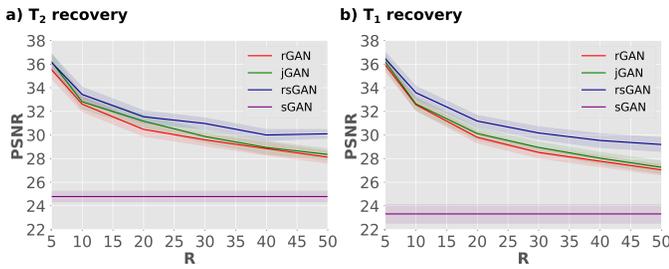


Fig. 7. The proposed rsGAN method was demonstrated for synergistic reconstruction-synthesis of T_1 - and T_2 -weighted images from the BRATS dataset. The acquisition for the source contrast was fully sampled, and the acquisition for the target contrast was undersampled by $R=5x, 10x, 20x, 30x, 40x, 50x$. PSNR was measured between recovered and fully-sampled reference target-contrast images. (a) PSNR (mean \pm standard error) across the test subjects for rsGAN, rGAN, jGAN, and sGAN when T_1 is the source contrast and T_2 is the target contrast. (b) PSNR (mean \pm standard error) when T_2 is the source contrast and T_1 is the target contrast.

III. RESULTS

A. Main Experiments

1) *Public Datasets:* We first demonstrated the proposed rsGAN method against rGAN, jGAN and sGAN on the MIDAS dataset. We considered two separate models: a model to recover T_2 -weighted images given T_1 -weighted images as source contrast, and another to recover T_1 -weighted images given T_2 -weighted images as source contrast. Tables I and II list the respective PSNR and SSIM measurements for each model, and Fig. 3 illustrates performance as a function of R . T_2 - and T_1 -weighted images in the MIDAS dataset recovered while R is varied from 10x to 50x are displayed in Figs. 4 and 5, respectively. Representative T_2 - and T_1 -weighted images recovered with ZF, jGAN, sGAN and rsGAN at $R=50x$ are shown in Fig. 6. As expected, the similarity between rsGAN and jGAN results increases towards $R=10x$, and that between rsGAN and sGAN increases towards $R=50x$. Furthermore, rsGAN recovers images of higher visual quality and acuity than both competing methods, particularly at intermediate R values. These results indicate that the incorporation of

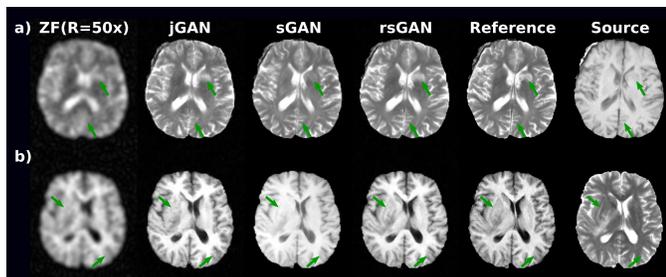


Fig. 8. Multi-contrast images in the BRATS dataset were recovered, where the source contrast was fully sampled and the target contrast was undersampled at $R=50x$. Images were recovered using ZF, sGAN, jGAN and rsGAN. (a) Recovered T_2 -weighted images along with the fully-sampled reference image and the source-contrast image. (b) Recovered T_1 -weighted images along with the fully-sampled reference image and the source-contrast image. rsGAN yields visually superior images compared to sGAN and jGAN. Note that sGAN suffers from either loss of features in the target contrast or synthesis of artefactual features. Meanwhile, jGAN suffers from excessive loss of high spatial frequency information. Sample regions that are more accurately recovered by rsGAN are marked with arrows.

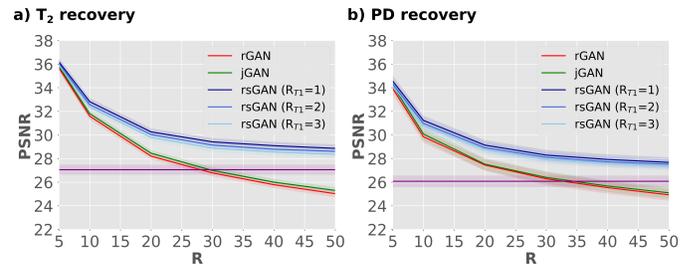


Fig. 9. The proposed rsGAN method was demonstrated for synergistic reconstruction-synthesis of T_1 -, T_2 - and PD-weighted images from the IXI dataset. The acquisition for the source contrast (T_1 -weighted) was lightly undersampled by $R_{T_1}=1x, 2x, 3x$, and the acquisitions for the target contrasts (T_2 - and PD-weighted) were heavily undersampled by $R=5x, 10x, 20x, 30x, 40x, 50x$. (a) PSNR (mean \pm standard error) across the test subjects for rsGAN, rGAN, jGAN, and sGAN when T_2 is the target contrast. (b) PSNR (mean \pm standard error) for sGAN, rGAN, jGAN, and rsGAN when PD is the target contrast. As expected, rsGAN outperforms sGAN, rGAN, and jGAN at all R . At the same time, performance of rsGAN is highly similar for distinct values of R_{T_1} .

a fully-sampled acquisitions of the source contrast enables rsGAN to more reliably recover high-frequency information compared to rGAN and jGAN, and that the use of evidence collected on the target contrast ensures that rsGAN yields more accurate recovery compared to sGAN. Next, we demonstrated the proposed method on a dataset acquired in patients with high- or low-grade gliomas. We considered two models on the BRATS dataset: a model to recover T_2 -weighted images given T_1 -weighted images, and another to recover T_1 -weighted images given T_2 -weighted images. Tables III and IV list the respective PSNR and SSIM values, and Fig. 7 illustrates model performance as a function of R . Representative T_2 - and T_1 -weighted images in the BRATS dataset recovered with ZF, jGAN, sGAN and rsGAN at $R=50x$ are shown in Fig. 8. Note that multi-contrast images can show differential sensitivity to tumor tissue, where tumors can be more easily delineated in T_2 - versus T_1 -weighted images particularly in patients with low-grade glioma. As a result, sGAN suffers from either loss of features in the target contrast or synthesis of artefactual features. Meanwhile, jGAN suffers from excessive loss of high spatial frequency information at high R . In comparison,

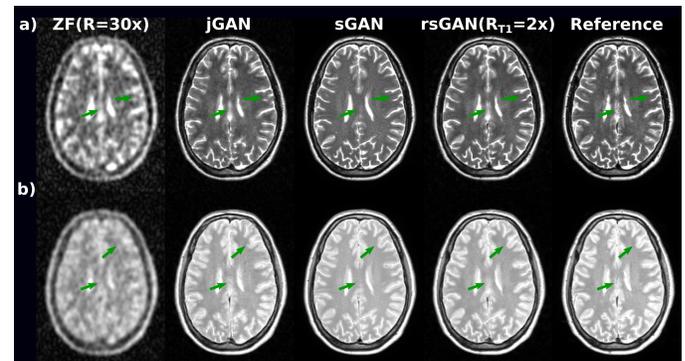


Fig. 10. Multi-contrast images in the IXI dataset were recovered, where the source contrast (T_1 -weighted) was lightly undersampled at $R_{T_1}=2x$, and the target contrasts (T_2 - and PD-weighted) were heavily undersampled at $R=30x$. Images were recovered using ZF, sGAN, jGAN and rsGAN. (a) Recovered T_2 -weighted images. (b) Recovered PD-weighted images. Sample regions where rsGAN yields sharper images and improved suppression of aliasing artifacts are marked with arrows.

rsGAN achieves higher spatial acuity while preventing feature losses and artefactual synthesis. Thus, the rsGAN method enables more reliable and accurate recovery when the source contrast is substantially less or more sensitive to differences in relaxation parameters of two tissues compared to the target contrast.

Next, we demonstrated the utility of rsGAN to recover multiple target contrasts simultaneously. The specific model tested on the IXI dataset was aimed to recover both T_2 - and PD-weighted images given T_1 -weighted images as source contrast. We examined the effect of light undersampling performed on the source contrast ($R_{T_1}=1x, 2x, 3x$) in addition to heavy undersampling on the target contrasts ($R=5x, 10x, 20x, 30x, 40x, 50x$). Tables V and VI list the PSNR and SSIM measurements for T_2 - and PD-weighted images, respectively. Fig. 9 illustrates model performance as a function of R_{T_1} and R . Representative T_2 - and PD-weighted images in the IXI dataset recovered with ZF, jGAN, sGAN and rsGAN at $R_{T_1}=2x, R=30x$ are shown in Fig. 10. The rsGAN method yields sharper images and improved suppression of aliasing artifacts compared to jGAN and sGAN, even when the source contrast acquisitions are accelerated. Across all public datasets, rsGAN achieves 1.66dB higher PSNR and 3.45% higher SSIM compared to rGAN, 1.40dB higher PSNR and 2.80% higher SSIM compared to jGAN, and 5.18dB higher PSNR and 3.83% higher SSIM compared to sGAN.

We also examined the effects of perceptual prior and variable-density sampling patterns in rsGAN. Supp. Tables I-VI list the PSNR and SSIM measurements across the recovered images in all public datasets. We find that the original rsGAN model outperforms rsGAN * on average by 0.53dB PSNR and 0.37% SSIM across the datasets. This result demonstrates the benefit of the perceptual prior for recovery performance. Comparisons among rsGAN and sr-sGAN indicate that rsGAN shows superior performance to sr-sGAN at all acceleration factors up to $R=20$ where rsGAN achieves 1.01dB higher PSNR and 0.40% higher SSIM, and the two methods perform similarly for $R>20$ where the differences are 0.16dB PSNR and 0.12% SSIM. Similar performance at very high accelerations is expected since the variable-density sampling patterns in rsGAN approach the central sampling patterns in sr-sGAN at these acceleration rates.

2) *Multi-Coil Dataset*: We then demonstrated the proposed approach on complex multi-coil dataset. The model was aimed to recover both T_2 - and PD-weighted images given T_1 -weighted images as source contrast. We examined the effect of light undersampling performed on the source contrast ($R_{T_1}=1x, 2x, 3x$) in addition to heavy undersampling on the target contrasts ($R=5x, 10x, 15x, 20x, 25x, 30x$). Tables VII and VIII list the PSNR and SSIM measurements for T_2 - and PD-weighted images, respectively. Fig. 11 illustrates model performance as a function of R_{T_1} and R . Overall, rsGAN is the leading performer. rsGAN ($R_{T_1}=1$) achieves 0.67 dB higher PSNR and 1.81% higher SSIM than rGAN, 0.58 dB higher PSNR and 1.56% higher SSIM than jGAN, and 5.61dB higher PSNR and 5.95% higher SSIM than sGAN. Even at $R_{T_1}=3$ rsGAN outperforms both rGAN and sGAN in terms of PSNR across values of $R>10x$. The only exception is at $R=5x$,

where rGAN increases T_2 recovery quality over rsGAN. This result suggests that at very low accelerations, the benefits of added prior information from the source can be outweighed by the added model complexity in rsGAN. Representative T_2 - and PD-weighted images in the multi-coil dataset recovered with ZF, rGAN, sGAN and rsGAN at $R_{T_1}=1x, R=10x$ are shown in Fig. 12. The rsGAN method yields sharper images and improved suppression of aliasing artifacts compared to rGAN and sGAN.

Next we measured recovery performance for jGAN and rsGAN with the same overall scan time of 250 sec. For jGAN this corresponds to $R=8.9x$ across all contrasts, whereas for rsGAN this corresponds to $R_{T_1}=3x$ for the source contrast and $R=15x$ for target contrasts. Across all three contrasts, rsGAN significantly outperforms jGAN in SSIM ($p<0.05$) by 0.63% while the two methods have similar PSNR. For the fixed scan time of 250 sec, we also compared the recovery performance of rsGAN explicitly for the source contrast. We observe that rsGAN outperforms jGAN by 4.69 dB in PSNR, and 6.16% in SSIM ($p<0.05$). This indicates that rsGAN is superior in recovery of the source contrasts as expected. These results showcase a scenario where rsGAN with nonuniform acceleration is preferable to jGAN with uniform acceleration.

B. Control Experiments

To rule out potential biases in model generalizability due to selection of cross-sections, we conducted control experiments in the IXI dataset where we considered two sets of models (see Methods for rGAN, jGAN, rsGAN and rGAN $_{All}$, jGAN $_{All}$, rsGAN $_{All}$). Supp. Table VII lists PSNR and SSIM measurements across the recovered images. Overall, rGAN $_{All}$ outperforms rGAN by 0.26dB in PSNR and 0.35% in SSIM, jGAN $_{All}$ outperforms jGAN by 0.30dB in PSNR and 0.36% in SSIM, and rsGAN $_{All}$ outperforms rsGAN by 0.32dB in PSNR and 0.44% in SSIM. Note that the slight performance improvement is natural since the test set contained peripheral cross-sections that were intentionally removed from the training set of rGAN, jGAN and rsGAN, but not included in the training set of rGAN $_{All}$, jGAN $_{All}$ and rsGAN $_{All}$. Second, we

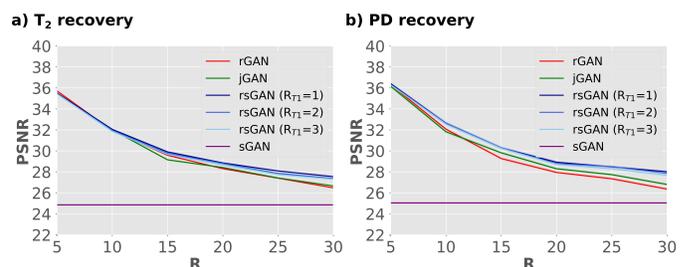


Fig. 11. The proposed rsGAN method was demonstrated for synergistic reconstruction-synthesis of T_1 -, T_2 - and PD-weighted images from the multi-coil dataset. The acquisition for the source contrast (T_1 -weighted) was lightly undersampled by $R_{T_1}=1x, 2x, 3x$, and the acquisitions for the target contrasts (T_2 - and PD-weighted) were heavily undersampled by $R=5x, 10x, 15x, 20x, 25x, 30x$. (a) PSNR (mean \pm standard error) across the test images (coronal cross-sections) for rsGAN, rGAN, jGAN, and sGAN when T_2 is the target contrast. (b) PSNR (mean \pm standard error) for sGAN, rGAN, jGAN, and rsGAN when PD is the target contrast. As expected, rsGAN outperforms sGAN, rGAN, and jGAN at high values of R . At the same time, performance of rsGAN is highly similar for distinct values of R_{T_1} .

TABLE V
QUALITY OF RECOVERED T_2 -WEIGHTED IMAGES IN THE IXI DATASET

	rGAN		jGAN		rsGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=2$)		rsGAN ($R_{T_1}=3$)	
	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)
R=5x	35.66±0.32	97.34± 0.16	35.80±0.29	97.46± 0.15	36.18±0.31[†]	97.70± 0.14[†]	36.01±0.29	97.58± 0.15	35.88±0.30	97.50± 0.15
R=10x	31.59±0.28	94.43± 0.34	31.80±0.27	94.76± 0.32	32.82±0.29[†]	95.91± 0.26[†]	32.56±0.29	95.63± 0.28	32.37±0.29	95.41± 0.29
R=20x	28.20±0.25	90.08± 0.55	28.45±0.24	90.69± 0.53	30.26±0.28[†]	93.87± 0.40[†]	30.03±0.26	93.46± 0.41	29.78±0.26	93.11± 0.42
R=30x	26.78±0.27	87.30± 0.69	26.99±0.26	88.09± 0.66	29.41±0.30[†]	93.09± 0.46[†]	29.13±0.29	92.59± 0.47	28.88±0.29	92.18± 0.50
R=40x	25.78±0.27	85.12± 0.82	26.00±0.26	85.83± 0.76	29.09±0.32[†]	92.71± 0.51[†]	28.78±0.28	92.18± 0.51	28.56±0.27	91.76± 0.52
R=50x	25.02±0.28	83.16± 0.95	25.30±0.26	84.08± 0.87	28.86±0.33[†]	92.48± 0.52[†]	28.61±0.30	91.96± 0.53	28.36±0.29	91.51± 0.55
sGAN										
					PSNR		SSIM (%)			
					27.06±0.41		90.97± 0.72			

T_1 -weighted acquisitions accelerated to various degrees (R_{T_1}) were taken as the source contrast, and T_2 - and PD-weighted acquisitions were taken as the target contrasts. PSNR and %SSIM values (mean±standard error) for T_2 -weighted images across the test subjects are listed for rGAN, jGAN, sGAN, and rsGAN. The highest PSNR and SSIM values in each row are marked in bold font, and the significantly better performing values ($p<0.05$) among rGAN, jGAN, sGAN, and rsGAN($T_1=1$) are marked with the '†' symbol.

TABLE VI
QUALITY OF RECOVERED PD-WEIGHTED IMAGES IN THE IXI DATASET

	rGAN		jGAN		rsGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=2$)		rsGAN ($R_{T_1}=3$)	
	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)
R=5x	33.93±0.48	97.46± 0.16	34.19±0.42	97.63± 0.15	34.56±0.44[†]	97.91± 0.13[†]	34.30±0.41	97.79± 0.14	34.15±0.43	97.71± 0.14
R=10x	29.84±0.49	94.62± 0.33	30.08±0.51	94.99± 0.31	31.23±0.39[†]	96.12± 0.25[†]	30.97±0.40	95.84± 0.26	30.76±0.39	95.64± 0.27
R=20x	27.45±0.48	90.35± 0.58	27.53±0.49	91.03± 0.55	29.13±0.41[†]	94.05± 0.38[†]	28.92±0.41	93.62± 0.39	28.74±0.42	93.30± 0.41
R=30x	26.29±0.54	87.68± 0.74	26.40±0.50	88.39± 0.73	28.29±0.43[†]	93.24± 0.46[†]	28.11±0.42	92.74± 0.48	27.94±0.44	92.35± 0.50
R=40x	25.53±0.49	85.42± 0.91	25.67±0.48	86.20± 0.85	27.92±0.42[†]	92.74± 0.51[†]	27.73±0.40	92.21± 0.52	27.63±0.39	91.79± 0.54
R=50x	24.92±0.48	83.70± 1.02	25.09±0.47	84.63± 0.95	27.67±0.42[†]	92.56± 0.53[†]	27.51±0.39	92.00± 0.53	27.37±0.40	91.56± 0.56
sGAN										
					PSNR		SSIM (%)			
					26.07±0.47		91.22± 0.67			

T_1 -weighted acquisitions accelerated to various degrees (R_{T_1}) were taken as the source contrast, and T_2 - and PD-weighted acquisitions were taken as the target contrasts.

TABLE VII
QUALITY OF RECOVERED T_2 -WEIGHTED IMAGES IN THE MULTI-COIL DATASET

	rGAN		jGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=2$)		rsGAN ($R_{T_1}=3$)	
	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)
R=5x	35.71±0.08[†]	97.62± 0.03	35.47±0.08	97.53± 0.03	35.51±0.08	97.61± 0.03	35.44±0.08	97.56± 0.03	35.44±0.08	97.53± 0.03
R=10x	31.93±0.07	95.27± 0.06	32.01±0.07	95.40± 0.06	32.05±0.08	95.64± 0.05[†]	31.94±0.08	95.52± 0.06	31.86±0.08	95.42± 0.06
R=15x	29.58±0.07	92.78± 0.09	29.15±0.08	91.92± 0.11	29.90±0.09[†]	93.66± 0.09[†]	29.77±0.09	93.42± 0.09	29.66±0.09	93.22± 0.10
R=20x	28.30±0.07	90.87± 0.11	28.41±0.07	91.22± 0.10	28.85±0.08[†]	92.42± 0.10[†]	28.79±0.08	92.21± 0.10	28.65±0.08	91.96± 0.11
R=25x	27.40±0.07	89.59± 0.12	27.41±0.07	89.33± 0.13	28.08±0.07[†]	91.21± 0.11[†]	27.86±0.07	90.67± 0.12	27.74±0.07	90.40± 0.11
R=30x	26.48±0.06	87.80± 0.14	26.66±0.07	87.97± 0.13	27.54±0.08[†]	90.56± 0.13[†]	27.37±0.07	90.06± 0.13	27.29±0.07	89.77± 0.13
sGAN										
					PSNR		SSIM (%)			
					28.45±0.07		87.16± 0.19			

T_1 -weighted acquisitions accelerated to various degrees (R_{T_1}) were taken as the source contrast, and T_2 - and PD-weighted acquisitions were taken as the target contrasts.

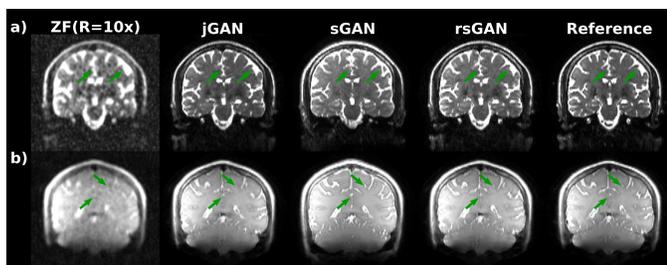


Fig. 12. Multi-contrast images in the multi-coil dataset were recovered, where the source contrast (T_1 -weighted) was fully sampled, and the target contrasts (T_2 - and PD-weighted) were heavily undersampled at $R=10x$. Images were recovered using ZF, sGAN, jGAN and rsGAN. (a) Recovered T_2 -weighted images. (b) Recovered PD-weighted images. Sample regions that are better recovered by rsGAN are marked with arrows.

of rsGAN over alternative models. Overall, rsGAN_{All} achieves 1.80dB higher PSNR and 4.27% higher SSIM than rGAN_{All}, and 1.63dB higher PSNR and 3.74% higher SSIM than jGAN_{All}. Note that rsGAN also achieves 1.48dB higher PSNR and 3.84% higher SSIM than rGAN_{All}, and 1.30dB higher PSNR and 3.31% higher SSIM than jGAN_{All}.

To examine the effects of input phase channels in recovery of coil-combined magnitude images, two sets of models were considered (rsGAN_m and rsGAN) for T_2 recovery in the BRATS dataset. Supp. Table VIII lists PSNR and SSIM measurements across the recovered images. We find that removing the phase channels from the rsGAN model decreases average PSNR and SSIM by 0.30dB and 0.05%. This difference might be attributed to the nature of phase images that typically emphasize information about tissue boundaries.

observe that results of the control experiments are consistent with the original experiments in demonstrating the superiority

Next, we conducted additional experiments to rule out any bias that might have occurred due to differences in model com-

TABLE VIII
QUALITY OF RECOVERED PD-WEIGHTED IMAGES IN THE MULTI-COIL DATASET

	rGAN		jGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=1$)		rsGAN ($R_{T_1}=2$)		rsGAN ($R_{T_1}=3$)	
	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)	PSNR	SSIM (%)
R=5x	36.31±0.11	97.55± 0.05	36.15±0.12	97.53± 0.05	36.40±0.12[†]	97.74± 0.05[†]	36.32±0.11	97.66± 0.05	36.28±0.11	97.61± 0.05
R=10x	32.06±0.12	94.88± 0.09	31.82±0.12	95.19± 0.08	32.64±0.11 [†]	95.77± 0.07[†]	32.65±0.11	95.58± 0.08	32.53±0.11	95.41± 0.08
R=15x	29.27±0.12	91.72± 0.16	29.81±0.11	92.35± 0.14	30.28±0.12 [†]	93.91± 0.10[†]	30.29±0.11	93.61± 0.11	30.23±0.11	93.41± 0.11
R=20x	27.94±0.11	89.83± 0.17	28.30±0.11	90.32± 0.17	28.92±0.12[†]	92.05± 0.16[†]	28.78±0.11	91.76± 0.16	28.71±0.11	91.51± 0.16
R=25x	27.34±0.11	88.28± 0.20	27.73±0.09	88.96± 0.18	28.48±0.11 [†]	92.01± 0.13[†]	28.50±0.11	91.62± 0.14	28.29±0.11	91.18± 0.14
R=30x	26.36±0.10	86.07± 0.25	26.80±0.10	87.52± 0.20	28.01±0.11[†]	91.38± 0.14[†]	27.85±0.11	90.86± 0.14	27.67±0.11	90.44± 0.15
sGAN										
PSNR						SSIM (%)				
25.03±0.08						88.27± 0.17				

T_1 -weighted acquisitions accelerated to various degrees (R_{T_1}) were taken as the source contrast, and T_2 - and PD-weighted acquisitions were taken as the target contrasts.

plexities among rGAN, jGAN, sGAN and rsGAN. These experiments were conducted on the BRATS dataset where T_1 was set as the source contrast and T_2 was set as the target contrast. Supp. Table IX lists PSNR and SSIM measurements across the recovered images. We find that rsGAN still outperforms rGAN_{MC}, jGAN_{MC} and sGAN_{MC} that were matched to rsGAN in model complexity. Overall rsGAN outperforms rGAN_{MC} by 1.39dB PSNR and 1.50% SSIM, jGAN_{MC} by 0.99dB PSNR and 1.19% SSIM, and sGAN_{MC} by 7.81dB PSNR and 4.99% SSIM. Furthermore, changing network complexity has minor effects in terms of model performance. Overall, performance in rGAN changes by 0.03dB PSNR and 0.02% SSIM, in jGAN changes by 0.19dB PSNR and 0.04% SSIM, and in sGAN changes by 0.56dB PSNR and 0.19% SSIM. Taken together, these experiments indicate that our results are not unduly biased by variability in model complexity.

We also evaluated the ability of rsGAN in recovering FLAIR images. rGAN, jGAN and rsGAN were compared in terms of average performance on the BRATS dataset. Supp. Table X lists PSNR and SSIM measurements across the recovered images. We find that rsGAN outperforms both rGAN and jGAN (please see Supp. Fig. 1 for representative images). In this task, rsGAN outperforms rGAN by 0.74dB PSNR and 0.85% SSIM, and jGAN by 0.62dB PSNR and 0.65% SSIM. These results suggest that the proposed rsGAN model has potential to synthesize a broader selection of contrasts.

Lastly, we compared the original rsGAN model that independently recovers all targets in a multi-target-contrast setting (rsGAN₁) against a unified rsGAN model that simultaneously recovers all target contrasts (rsGAN₂). Comparisons were performed for recovery of T_2 and FLAIR images in the BRATS dataset. Supp. Table XI lists PSNR and SSIM measurements across the recovered images. We find that rsGAN₁ yields 0.25dB higher PSNR and 0.10% higher SSIM than rsGAN₂. Note that this moderate performance drop in the unified model is expected, as rsGAN₂ has to compromise between recovery losses for the two target contrasts.

IV. DISCUSSION

A synergistic reconstruction-synthesis approach based on conditional GANs was presented for highly accelerated multi-contrast MRI. In this approach, several source- and target-contrast acquisitions accelerated to various degrees are taken as input, and high-quality images for individual contrasts are

then recovered. The proposed rsGAN method yielded superior recovery performance against state-of-the-art reconstruction and synthesis methods in three public MRI datasets and a multi-coil dataset. While rsGAN was demonstrated for multi-contrast MRI here, it may also offer improved performance in recovery of images in accelerated multi-modal datasets.

Several previous studies considered joint reconstructions of multi-contrast acquisitions to better use shared structural information among contrasts. In the CS framework, a typical scenario involves multiple acquisitions with nearly identical acceleration rates [70], [71]. Undersampled data are jointly processed, and a joint-sparsity regularization term improves recovery of shared features across contrasts. Another scenario involves the fully-sampled acquisition of a reference contrast that is then used as a structural prior for other contrasts [72]. Prior-guided reconstructions use regularization terms that enforce consistency of the magnitude and direction of image gradients across distinct contrasts. These previous approaches yield enhanced quality over independent processing of each contrast. However, hand-crafted regularization terms based on transforms such as total variation or wavelet reflect often suboptimal assumptions about structural similarity among separate contrasts. The proposed rsGAN method instead employs a data-driven approach to learn to utilize information from source contrast during recovery of target contrasts.

Few recent studies proposed a learning-based method for joint reconstruction of multi-contrast MRI data [53]–[55]. Acquisitions for separate contrasts were accelerated at identical rates. Convolutional neural network architectures were used with a subset of network weights shared across contrasts to better capture structural similarities among contrasts. While these previous methods were shown to outperform conventional CS and parallel imaging reconstructions, these are pure reconstruction approaches that can suffer from scarce sampling of high spatial frequencies at high acceleration rates. In contrast, rsGAN employs detailed structural information in a source contrast to enhance the recovery of high-frequency samples in target contrasts. Since the source acquisitions are fully-sampled or lightly undersampled, rsGAN shows improved reliability against losses in resolution. Furthermore, GANs have been shown to better learn the distribution of target images compared to conventional network architectures.

Several independent studies proposed convolutional neural networks for recovery of a target contrast by making use

of structural information from a source contrast [56]–[59]. Perhaps, the closest to our study is [59], where recovery from a low-resolution target-contrast image was enhanced by incorporating fully-sampled acquisition of a separate contrast [59]. Our study is different from [59] in the following aspects: 1) We demonstrate the proposed approach for reconstruction of multi-coil complex MR images, whereas [59] consider a model to post-process coil-combined images that were already reconstructed. 2) We demonstrate that the proposed approach can jointly reconstruct and synthesize the target contrast even when the source contrasts are undersampled. 3) We demonstrate that the proposed approach can enable high acceleration factors up to 50x by incorporating information from both source and target contrasts. 4) We incorporate an additional perceptual prior to improve recovery of high-level image features. In addition to these technical differences, we also demonstrated superiority of rsGAN over [59] at all acceleration factors up to 20 (Supp. Tables I-VI).

The synthesis framework is an alternative for recovery of images of a target contrast, where data are only available in a different source contrast. A powerful approach is to construct dictionaries from multi-resolution image patches, and to learn a mapping between the source and target dictionaries [23], [24], [27], [28], [32]. Segregation of the dictionary extraction and mapping stages might yield suboptimal performance. Network-based approaches offer a remedy to this problem by unifying the two stages [25], [34]–[36]. We recently proposed GAN-based synthesis for multi-contrast MRI that yielded enhanced performance compared to conventional methods [39]. Yet, due to lack of evidence on the target contrast, a pure synthesis approach can suffer from artificial sensitivity or insensitivity to image features. The rsGAN method, on the other hand, always collects a moderate to small amount of evidence. This helps avoid artefactual feature leakage from the source to the target contrast or loss of target-contrast features that are not apparent in the source contrast.

An important query about the proposed approach is selection of source and target contrasts. Note that rsGAN performs heavy undersampling for only target contrasts, while moderately undersampling source contrasts to preserve as much information regarding detailed tissue structure as possible. This suggests that, in a given MRI protocol, contrasts with better capture of structural details and relatively shorter scan times should be designated as source contrasts, whereas contrasts with relatively limited capture of structural details and longer scan times should be designated as target contrasts. Clinical protocols typically start with a high-resolution T_1 -weighted acquisition, so we considered T_1 as a natural source contrast in the current study. Meanwhile, remaining contrasts including T_2 , PD or FLAIR with relatively lower capture of structural details and longer scan times were designated as target contrasts. In cases where multiple candidates exist, a selection might be necessary to minimize the overall scan time. For example, diagnostic protocols for glioma typically include T_1 -weighted, postcontrast T_1 -weighted, T_2 -weighted, and FLAIR acquisitions. In this case, both T_1 and postcontrast T_1 acquisitions are possible candidates of source contrasts. Yet, given the benefits of postcontrast T_1 over T_1 in terms of

diagnostic accuracy and/or tumor segmentation, postcontrast T_1 can be selected.

Here, the experiment for a fixed scan time was based on T_1 -weighted acquisitions with a 3D MP-RAGE sequence, and T_2 - and PD-weighted acquisitions with a 3D Spin-Echo sequence. Naturally the optimal distribution of acceleration factors across contrasts might vary depending on the specific sequences prescribed for each contrast. When other sequences or different sets of sequence parameters are prescribed, the distribution of acceleration factors across contrasts can be retuned empirically to maximize reconstruction performance. To ensure optimal performance, separate networks can be trained for each MR protocol and overall scan time. But since network training is performed off-line prior to MR scans, the recovery time would not be affected.

The superiority of rsGAN over jGAN in recovering source contrasts might motivate its use in several scenarios. First, rsGAN can be used to recover a specific source contrast with higher quality within an MRI protocol. The enhanced recovery can prove useful when task-critical information is largely concentrated in this source contrast. For instance, high-quality T_1 -weighted brain images are vital for accurate segmentation of white and gray matter, or high-quality postcontrast T_1 -weighted images serve as a gold-standard tool for tumor localization. Second, rsGAN can be employed in longitudinal imaging studies [73], [74], where subjects are scanned in multiple sessions over extended periods of time. This can help reduce scan time and increase patient comfort in multi-session MRI exams. In such cases, fully-sampled source contrasts can be acquired during the initial session, and they can then aid recovery in subsequent sessions. Lastly, rsGAN might offer utility in facilitating efficient re-acquisition of problematic images in a multi-contrast exam. When a subset of acquisitions suffers from severe patient motion or artifacts, the acquisitions can be repeated at significantly higher acceleration factors to minimize undesirable increase in scan time. The resulting multi-contrast acquisitions with non-uniform undersampling across contrasts can then be recovered with rsGAN. When the source contrast is affected by motion, it could be reacquired at a moderate acceleration factor and recovered using rsGAN. Alternatively, a separate network for suppression of motion artifacts could be trained [75] and used as a pre-processing step to rsGAN.

Several technical developments are viable for improving the current implementation of the proposed method. First, the model can be generalized to simultaneously process multiple neighboring cross-sections in addition to multiple contrasts. Correlated tissue structure across cross sections might enhance recovery despite the increase in model complexity. Second, when multiple source contrasts are present, a weight sharing method can be used to enforce a shared latent representation among contrasts for improved performance. Lastly, a cycleGAN-based model [65] might be implemented to allow for learning on unpaired multi-contrast MRI datasets that are relatively more available than paired datasets.

- [38] 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.
- [39] 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. 90, no. 312, pp. 1–1, 2019.
- [40] Q. Yang, N. Li, Z. Zhao, X. Fan, E. I-Chao Chang, and Y. Xu, "MRI image-to-image translation for cross-modality image registration and segmentation," *arXiv preprint*, 2018.
- [41] B. Yu, L. Zhou, L. Wang, J. Frupp, and P. Bourgeat, "3D cGAN based cross-modality MR image synthesis for brain tumor segmentation," in *IEEE 15th International Symposium on Biomedical Imaging*. IEEE, 2018, pp. 626–630.
- [42] B. Yu, L. Zhou, L. Wang, Y. Shi, J. Frupp, 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.
- [43] A. Sharma and G. Hamarneh, "Missing MRI Pulse Sequence Synthesis using Multi-Modal Generative Adversarial Network," *arXiv preprint*, 2019.
- [44] S. Olut, Y. H. Sahin, U. Demir, and G. Unal, "Generative Adversarial Training for MRA Image Synthesis Using Multi-contrast MRI," in *Predictive Intelligence in Medicine*. Springer, Cham, 2018, pp. 147–154.
- [45] M. Yurt, S. U. Dar, A. Erdem, E. Erdem, and T. Çukur, "mustGAN: Multi-Stream Generative Adversarial Networks for MR Image Synthesis," *arXiv preprint*, 2019.
- [46] D. Lee, J. Kim, W.-J. Moon, and J. C. Ye, "CollaGAN: Collaborative GAN for Missing Image Data Imputation," *arXiv preprint*, 2019.
- [47] P. Welander, S. Karlsson, and A. Eklund, "Generative Adversarial Networks for Image-to-Image Translation on Multi-Contrast MR Images - A Comparison of CycleGAN and UNIT," *arXiv preprint*, 2018.
- [48] H. Li, J. C. Paetzold, A. Sekuboyina, F. Kofler, J. Zhang, J. S. Kirschke, and B. Menze, "DiamondGAN: Unified Multi-Modal Generative Adversarial Networks for MRI Sequences Synthesis," *arXiv preprint*, 2019.
- [49] A. Hagiwara, Y. Otsuka, M. Hori, Y. Tachibana, K. Yokoyama, S. Fujita, C. Andica, K. Kamagata, R. Irie, S. Koshino, T. Maekawa, L. Chougar, A. Wada, M. Y. Takemura, N. Hattori, and S. Aoki, "Improving the Quality of Synthetic FLAIR Images with Deep Learning Using a Conditional Generative Adversarial Network for Pixel-by-Pixel Image Translation." *AJNR. American Journal of Neuroradiology*, vol. 40, no. 2, pp. 224–230, 2019.
- [50] R. Mehta and T. Arbel, "RS-Net: Regression-Segmentation 3D CNN for Synthesis of Full Resolution Missing Brain MRI in the Presence of Tumours," *arXiv preprint*, 2018.
- [51] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Computer Vision – ECCV 2016*, 2016, pp. 694–711.
- [52] 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.
- [53] L. Sun, Z. Fan, X. Fu, Y. Huang, X. Ding, and J. Paisley, "A Deep Information Sharing Network for Multi-Contrast Compressed Sensing MRI Reconstruction," *IEEE Transactions on Image Processing*, vol. 28, no. 12, pp. 6141–6153, 2019.
- [54] D. Polak, S. Cauley, B. Bilgic, E. Gong, P. Bachert, E. Adalsteinsson, and K. Setsompop, "Joint multi-contrast variational network reconstruction (jvn) with application to rapid 2d and 3d imaging," *Magnetic Resonance in Medicine*, vol. n/a, no. n/a. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28219>
- [55] E. Gong, G. Zaharchuk, and J. Pauly, "Improving the pi+ cs reconstruction for highly undersampled multi-contrast mri using local deep network," in *25th Annual Meeting of the ISMRM*, ser. Lecture Notes in Computer Science. Honolulu: Springer, 2017, p. 5663.
- [56] 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.
- [57] A. Falvo, D. Communiello, S. Scardapane, G. Finesi, M. Scarpiniti, and A. Uncini, "A multimodal deep network for the reconstruction of t2w mr images," *arXiv preprint*, 2019.
- [58] A. Falvo, D. Communiello, S. Scardapane, M. Scarpiniti, and A. Uncini, "A multimodal dense u-net for accelerating multiple sclerosis mri," in *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2019, pp. 1–6.
- [59] K. H. Kim, W. Do, and S. 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.
- [60] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *IEEE International Conference on Computer Vision*, 2017, pp. 2813–2821.
- [61] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [62] E. Bullitt, D. Zeng, G. Gerig, S. Aylward, S. Joshi, J. K. Smith, W. Lin, and M. G. Ewend, "Vessel tortuosity and brain tumor malignancy: A blinded study," *Academic Radiology*, vol. 12, no. 10, pp. 1232–1240, 2005.
- [63] M. Jenkinson and S. Smith, "A global optimisation method for robust affine registration of brain images," *Medical Image Analysis*, vol. 5, no. 2, pp. 143–156, 2001.
- [64] M. Jenkinson, P. Bannister, M. Brady, and S. Smith, "Improved optimization for the robust and accurate linear registration and motion correction of brain images," *NeuroImage*, vol. 17, no. 2, pp. 825–841, 2002.
- [65] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE International Conference on Computer Vision*, 2017, pp. 2223–2232.
- [66] D. P. Kingma and J. L. Ba, "Adam: a Method for Stochastic Optimization," in *International Conference on Learning Representations*, 2015.
- [67] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint*, 2016.
- [68] T. Zhang, J. M. Pauly, S. S. Vasawala, and M. Lustig, "Coil compression for accelerated imaging with Cartesian sampling," *Magnetic Resonance in Medicine*, vol. 69, no. 2, pp. 571–82, 2013.
- [69] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasawala, and M. Lustig, "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA," *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.
- [70] J. Huang, C. Chen, and L. Axel, "Fast multi-contrast MRI reconstruction," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 7510 LNCS, 2012, pp. 281–288.
- [71] B. Bilgic, V. K. Goyal, and E. Adalsteinsson, "Multi-contrast reconstruction with Bayesian compressed sensing," *Magnetic Resonance in Medicine*, vol. 66, no. 6, pp. 1601–1615, 2011.
- [72] M. J. Ehrhardt and M. M. Betcke, "Multicontrast MRI Reconstruction with Structure-Guided Total Variation," *SIAM Journal on Imaging Sciences*, vol. 9, no. 3, pp. 1084–1106, 2016.
- [73] K. L. Mills and C. K. Tamnes, "Methods and considerations for longitudinal structural brain imaging analysis across development," *Developmental cognitive neuroscience*, vol. 9, pp. 172–190, 2014.
- [74] Z. Y. Shan, R. Kwiatek, R. Burnet, P. Del Fante, D. R. Staines, S. M. Marshall-Gradisnik, and L. R. Barnden, "Progressive brain changes in patients with chronic fatigue syndrome: a longitudinal mri study," *Journal of Magnetic Resonance Imaging*, vol. 44, no. 5, pp. 1301–1311, 2016.
- [75] B. A. Duffy, W. Zhang, H. Tang, L. Zhao, M. Law, A. W. Toga, and H. Kim, "Retrospective correction of motion artifact affected structural mri images using deep learning of simulated motion," 2018.