

DEEP UNSUPERVISED LEARNING FOR ACCELERATED MRI RECONSTRUCTION

A THESIS SUBMITTED TO
THE GRADUATE SCHOOL OF ENGINEERING AND SCIENCE
OF BILKENT UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
MASTER OF SCIENCE
IN
ELECTRICAL AND ELECTRONICS ENGINEERING

By
Yilmaz Korkmaz
July 2022

Deep Unsupervised Learning for Accelerated MRI Reconstruction

By Yılmaz Korkmaz

July 2022

We certify that we have read this thesis and that in our opinion it is fully adequate, in scope and in quality, as a thesis for the degree of Master of Science.

Tolga Çukur(Advisor)

Emine Ülkü Sarıtaş Çukur

İlkay Öksüz

Approved for the Graduate School of Engineering and Science:

Orhan Arıkan
Director of the Graduate School

ABSTRACT

DEEP UNSUPERVISED LEARNING FOR ACCELERATED MRI RECONSTRUCTION

Yılmaz Korkmaz

M.S. in Electrical and Electronics Engineering

Advisor: Tolga Çukur

July 2022

Supervised reconstruction models are characteristically trained on matched pairs of undersampled and fully-sampled data to capture an MRI prior, along with supervision regarding the imaging operator to enforce data consistency. To reduce supervision requirements, the recent deep image prior framework instead conjoins untrained MRI priors with the imaging operator during inference. Yet, canonical convolutional architectures are suboptimal in capturing long-range relationships, and priors based on randomly initialized networks may yield suboptimal performance. To address these limitations, this thesis introduces a novel unsupervised MRI reconstruction method based on zero-Shot Learned Adversarial TransformerEs (SLATER). SLATER embodies a deep adversarial network with cross-attention transformers to map noise and latent variables onto coil-combined MR images. During pre-training, this unconditional network learns a high-quality MRI prior in an unsupervised generative modeling task. During inference, a zero-shot reconstruction is then performed by incorporating the imaging operator and optimizing the prior to maximize consistency to undersampled data. Comprehensive experiments on brain MRI datasets clearly demonstrate the superior performance of SLATER against state-of-the-art unsupervised methods.

Keywords: adversarial, transformers, MRI, unsupervised, reconstruction, zero-shot, generative.

ÖZET

DERİN DENETİMSİZ ÖĞRENME İLE HIZLANDIRILMIŞ MRG REKONSTRÜKSİYONU

Yılmaz Korkmaz

Elektrik ve Elektronik Mühendisliği, Yüksek Lisans

Tez Danışmanı: Tolga Çukur

July 2022

Denetimli rekonstrüksiyon modelleri, veri tutarlılığını sağlamak için görüntüleme operatörüyle ilgili denetimle birlikte, genelde yetersiz örneklenmiş ve tam örneklenmiş veri çiftleri üzerinde eğitilir. Denetim gereksinimlerini azaltmak için, derin görüntü modelleri rekonstrüksiyon sırasında görüntüleme operatörü ile eğitimsiz MRG modellerini birleştirir. Bunun yanı sıra, kurallı evrişimli mimariler, uzun menzilli ilişkileri yakalamada yetersizdir ve rastgele başlatılan ağlara dayanan modeller, rekonstrüksiyon sırasında performans kaybına uğrayabilir. Bu sınırlamaları ele almak için, burada transformatör ve çekişmeli üretici ağlara (SLATER) dayalı yeni bir denetimsiz MRG rekonstrüksiyon yöntemini tanıtıyoruz. SLATER, gürültüyü ve gizli değişkenleri bobin-kombine MR görüntüleri ile eşleştirmek için çapraz dikkat transformatörleri ve derin bir çekişmeli üretici ağı bünyesinde barındırır. Ön eğitim sırasında, bu koşulsuz ağ, denetimsiz bir üretici modelleme görevi ile yüksek kaliteli MR görüntüleri sentezlemeyi öğrenir. Rekonstrüksiyon sırasında, görüntüleme operatörü işleme dahil edilerek az örneklenmiş verilere tutarlılığı en üst düzeye çıkarmak için model parametreleri optimize edilir ve sıfır atışlı bir yeniden oluşturma gerçekleştirilmiş olur. Beyin MRG veri kümeleri üzerinde yapılan kapsamlı deneyler, SLATER'in son teknoloji denetimsiz yöntemlere karşı üstün performansını açıkça göstermektedir.

Anahtar sözcükler: MRG, üretici çekişmeli ağlar, transformatör, rekonstrüksiyon, sıfır-atım, denetimsiz derin ağlar.

Acknowledgement

First, I would like to thank my supervisor Prof. Tolga Cukur for his superb guidance throughout this journey. He always lighted my way when I was lost. He always had bunch of solutions to any problem I encountered. He was always understanding and helpful. I learned a lot from him and will feel his absence in my future studies for sure.

I would also like to thank my jury members, Prof. Emine Ulku Saritas for her tremendous help and very polite behaviour towards me throughout my Bachelor's and Master's studies. I would also like to thank Prof. Ilkay Oksuz for all his efforts during the examination process. I also would like to thank Prof. Yusuf Ziya Ider for his very kind and generous effort during the PhD admissions. And, special thanks to Prof. Hayrettin Koymen for his sincere guidance during my father's illness.

I want to thank a lot to all of my lab partners and cohorts. In particular, Salman for his guidance and enourmous help in the work presented in this thesis, Muzaffer for his friendship and patience in answering all the silly questions I have about the thesis, Mahmut for his friendship from the beginning of the Master's and his valuable financial advises. I also would like to thank Onat, Gokberk and Emrullah for their contributions and providing a friendly working environment.

I have to name my dearest pals for their excellent friendships. Berke was my true partner-in-crime for our eight years in Bilkent. Special thanks to Kaan for being a great companion in every part of my life for the last fifteen years, and Dilek for being the ultimate companion for him and kindest ally to me throughout Bilkent and Nanyang. I wish I know Evren, the champ, before sophomore, Bilkent would have been a lot fun. Onur was there for ME(TU) through thick and thin, he is the best.

Of course my biggest thanks goes to my family. I am deeply indebted to my father Erhan and my little (18 years old) brother Yiğit -it is now his time to make great memories in Bilkent-. All my achievements are their success, and any failure

is mine. I specially want to thank my aunt Filiz for all her love and support during my studies.

I have to express my gratefulness to Zeynep, my dearest friend and the most beautiful thing in my life for the last six years. You were always extremely supportive, encouraging, and understanding throughout this journey.

Lastly, I want to dedicate this thesis to my late mother, Ülfet. She was always proud of me and would be proud for this thesis too. I miss her every day and know we are together in an alternate universe somewhere.

Contents

1	Introduction	1
2	Fundamentals of Accelerated MRI	7
3	Theory	10
3.1	Accelerated MRI Reconstruction	10
3.2	Deep Image Prior	11
3.3	Adversarial Transformer Model	12
3.3.1	Synthesizer (G)	12
3.3.2	Mapper (M)	16
3.3.3	Discriminator (D)	18
3.4	Self- versus Cross-Attention Transformers	18
3.4.1	Self-attention transformer	19
3.4.2	Cross-attention transformer	21
3.5	Learning Procedures	22

3.5.1	Pre-training of the MRI prior	22
3.5.2	Zero-shot reconstruction	23
4	Methods	25
4.1	Network Architecture	25
4.1.1	Synthesizer	25
4.1.2	Mapper	26
4.1.3	Discriminator	28
4.2	Competing Methods	28
4.3	Datasets	31
4.4	Quantitative Assessments	32
5	Results	33
5.1	Cross-Attention Transformers	33
5.2	Model Invertability	34
5.3	Within-Domain Reconstructions	34
5.4	Across-Domain Reconstructions	35
5.5	Ablation Experiments	37
5.6	Weight Propagation	37
6	Conclusion	40

- A Supplementary Materials 72**
 - A.1 Supplementary Text 72
 - A.1.1 Positional encoding variables 72
 - A.1.2 Architectural Details 73
 - A.2 Supplementary Tables 76
 - A.3 Supplementary Figures 81

List of Figures

2.1	In parallel imaging, images gathered from independent receiver coils, which are mostly sensitive to the nearest tissues, are combined by using known coil-sensitivity maps.	9
3.1	Cross-attention maps in SLATER for a T_1 -weighted acquisition. Sample maps from the first CA sub-block are displayed by overlay onto the respective MR image across three resolutions (i.e., at network layers 4-6). Attention maps for separate latents typically show segregated spatial distribution, and tend to group tissue clusters with similar signal intensity and texture broadly distributed across the image.	18
3.2	Cross-attention maps in SLATER for a simulated phantom with varying levels of normally-distributed white noise. Sample maps from the first CA sub-block are displayed at 128x128 resolution (network layer 6). Relative to a peak signal intensity of 1, top, middle and bottom rows show results for no noise, noise variance of 0.01, and noise variance of 0.1, respectively.	23

- 4.1 **(a)** The generative model in SLATER is adapted to perform zero-shot reconstruction for accelerated MRI. To do this, the generator output is back-projected onto individual coils (C), and masked with the same sampling pattern as in the undersampled acquisition (M_Ω). **(b)** The network prior consisting of noise (n), latent variables (W) and weights (θ_G) is then optimized to maximize consistency between the reconstructed and acquired k-space data. 26
- 4.2 SLATER is based on an unconditional adversarial architecture with a synthesizer (G), a discriminator (D), and a mapper (M). G is a multi-layer architecture where image resolution is progressively increased across layers. Within intermediate layers, a convolutional upsampling block is followed by a cross-attention transformer block (see 1). The transformer block receives global and local latent variables, noise variables and positional encoding (P.E.) for latents and image features. D enables adversarial learning while it receives as input synthetic and actual MR images. It is composed of convolutional and fully-connected (FC) blocks (see 2, 3). M projects raw latent variables (Z) onto k local (w_1, \dots, w_K) variables and one global (w_g) latent variable. It is composed of self-attention and FC blocks (see 3, 4, and Supp. Fig. A.1 for details). An unsupervised generative modeling task is performed using the SLATER model to capture high-quality MRI priors. 27
- 5.1 Reconstructions of a representative T_2 -weighted acquisition at $R=4$ are shown for the Fourier method (ZF), DIP methods (GAN_{DIP} , $\text{SAGAN}_{\text{DIP}}$, $\text{SLATER}_{\text{DIP}}$) and zero-shot reconstructions ($\text{GAN}_{\text{prior}}$, SAGAN, SLATER) along with the reference image. Zoom-in display windows are added to aid visualization of performance differences. Corresponding error maps are underneath the images for each method. 35

5.2	Within-domain reconstructions of a T_1 -weighted acquisition in fastMRI at R=4. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image, and error maps in the bottom row.	36
5.3	Across-domain reconstructions of a T_2 -weighted acquisition in fastMRI at R=4. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image, and error maps in the bottom row.	38

List of Tables

5.1	Reconstruction performance of DIP and zero-shot reconstructions for T_1 - and T_2 -weighted acquisitions IXI at R=4 and 8. Performance metrics are presented as mean \pm standard deviation across test subjects.	39
5.2	Within-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in fastMRI at R=4 and 8.	39
5.3	Across-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in the IXI and fastMRI datasets at R=4. In A->B, A is the training domain and B is the test domain. . . .	39
5.4	Average inference times in sec per cross section. Run times without and with weight propagation (WP) are listed for zero-shot reconstructions. SAGAN does not perform weight optimization, so it is unaffected by WP.	39

Chapter 1

Introduction

Magnetic Resonance Imaging (MRI) is a clinical powerhouse due to its excellent soft tissue contrast. Yet, low spin polarization at mainstream field strengths limits signal-to-noise ratio and necessitates prolonged exams. Extensive MRI exams incur added economic costs and lower use efficiency, and they may not be possible in uncooperative patient populations. A dire consequence is administration of exams that are in lower quality and/or diversity, which projects poorly onto diagnostic utility. Accelerated MRI addresses this fundamental limitation by performing undersampled acquisitions, and then solving an inverse problem to reconstruct images from available k-space data [1–4].

Given their exceptional performance in inverse problems, supervised neural networks have been quickly adopted for MRI reconstruction [5]. Reconstruction refers to the task of mapping undersampled acquisitions to images that are as consistent as possible with corresponding fully-sampled acquisitions. Supervised models performing this conditional mapping are trained on matched pairs of undersampled and fully-sampled data. The training process involves multiple lines of supervision including pairing of input-output data to learn an indirect prior that reduces aliasing artifacts in MR images [6–9], and enforcement of data consistency to embed the imaging operator that reflects k-space undersampling and coil-sensitivity encoding [10–14]. Accordingly, supervised models are typically re-trained for notable changes in the data distribution (e.g, different MRI contrast)

or the imaging operator (e.g., different k-space sampling density) [15–17].

Several important approaches have been proposed in literature to reduce supervision requirements. A group of studies have reduced explicit supervision related to raw data, proposing models trained using unpaired sets of input-output data [18–20], or using only undersampled data [21–27]. That said, these models involve implicit supervision regarding the imaging operator. As such, they are trained for a specific coil-array configuration and k-space sampling density, factors assumed to be consistent across the training and test sets [21–24]. To remove other supervision aspects, a second group of studies have built unsupervised models by decoupling the MRI prior from the imaging operator. These models capture an MRI prior via generative networks that are either untrained [28–31] or trained to synthesize fully-sampled MR images [32–34]. The imaging operator is then conjoined with the MRI prior during inference on test data. This unsupervised approach excludes paired datasets for training, and promises improved generalization against deviations in the imaging operator [33,34]. Yet, previously proposed models are commonly based on convolutional neural networks (CNNs) that suffer from limited sensitivity in capturing long-range dependencies [35–37].

Contributions

In this thesis, we introduce a novel unsupervised MRI reconstruction based on zero-shot learned adversarial transformers (SLATER). SLATER decomposes the reconstruction process to decouple learning of the MRI prior from embedding of the imaging operator. During a pre-training phase, an unconditional adversarial model is used to synthesize high-quality, coil-combined MR images (Fig. 4.2). To improve capture of long-range spatial context without introducing computational burden, we propose an architecture comprising cross-attention transformer blocks between low-dimensional latent variables and high-dimensional image features. During the inference phase, the learned MRI prior is combined with the imaging operator. This is achieved by optimizing network parameters that reflect the MRI prior to enforce a data-consistency loss on undersampled test data (Fig. 4.1). To

improve inference efficiency, a weight propagation strategy is proposed where the optimized network weights are transferred across consecutive cross-sections.

The proposed method performs an unsupervised generative modeling task on coil-combined images derived from fully-sampled MRI acquisitions. Adapting the generative model to the reconstruction task without any training samples, a zero-shot reconstruction then maps undersampled data to high-quality MR images during inference. The decoupled reconstruction process and model adaptation during inference contribute to improved generalization performance for SLATER.

- For the first time in literature, we introduce an adversarial vision transformer model for MRI reconstruction.
- Our proposed model uses cross-attention transformers to capture long-range spatial dependencies without computational burden.
- Sample-specific model adaptation and cross-sectional weight propagation strategies are introduced that respectively enhance out-of-domain generalizability and inference efficiency.

Related Work

Supervised reconstruction models are current state-of-the-art in MRI with numerous successful architectures reported including basic CNNs [6, 38–40], residual CNNs [41, 42], perceptrons [17, 43], physics-guided unrolled CNN networks [7, 10, 11, 14, 44–48], recurrent CNNs [9, 49, 50], generative adversarial networks (GANs) [51–54], and variational networks [8, 13, 15, 55]. However, supervised models are trained on paired sets of undersampled and fully-sampled data, along with supervision regarding the imaging operator in the form of a data-consistency term. Compiling large sets of paired data is non-trivial [56], and supervised models often require retraining to cope with deviations in the imaging operator [17].

To improve utility of deep MRI reconstruction, a common strategy has been to target explicit supervision on raw data in order to lower reliance on large, paired

training datasets. Domain-transferred models are trained in a data-abundant source domain and then adopted for reconstruction in the target domain [16, 57]. Residual models are trained to predict residual error between ground truth and the output of a conventional MRI reconstruction [58, 59]. Both approaches permit training with relatively few samples, but training often requires paired datasets. Unpaired models are instead trained on input-output data collected from separate groups of subjects. This can be achieved with cycle-consistent models that learn bidirectional mappings between undersampled and fully-sampled data in order to enforce self-consistency of network inputs or outputs [12, 18, 60]. However, unpaired models can require substantially larger datasets for training compared to paired models [20]. Lastly, several prominent approaches were proposed to train unsupervised models in the absence of ground truth. Scan-specific models learn nonlinear interpolation kernels from an auto-calibration region in under-sampled acquisitions, and then recover missing k-space samples via interpolation during inference [61–63]. These models do not require a priori training, but their performance relies on the assumption that local dependencies between samples are largely invariant across k-space. Self-supervised models are trained via proxy loss terms substituted for the true reconstruction loss [21]. Specifically, data-consistency loss calculated on samples available in undersampled acquisitions can serve as a proxy [22, 23]. An alternative is to leverage statistical estimators for mean-squared reconstruction loss [27]. Such estimators are often analytically derived, so their use might be limited to specific types of loss functions. Recent studies have also proposed to mask a subset of k-space samples in undersampled acquisitions, where the masked subset is used to define a data-consistency loss [24, 25]. Based on this loss, self-supervised models are trained to recover masked-out samples from remaining samples. A related image-domain strategy has been proposed for multi-image MRI data [26, 64]. In this case, models can be trained to predict a masked subset of undersampled frames in dynamic MRI from remaining frames [65], or to map between differently undersampled versions of a given fully-sampled acquisition [66]. Because the indirect MRI prior to reduce aliasing artifacts is learned from undersampled data, self-supervised models might show suboptimal performance at high acceleration factors. Moreover, unsupervised methods that used fixed model weights during inference might suffer from suboptimal generalization to test data [67, 68].

A collective attribute of the aforementioned methods to limit supervision on raw data is that they are based on conditional models that map undersampled acquisitions to MR images. Thus, they involve inherent supervision regarding the imaging operator, and they are typically retrained for varying coil-array configurations and sampling densities [24]. Removing instead supervision related to the imaging operator, a fundamentally different approach decomposes the reconstruction process to decouple the MRI prior from the imaging operator. The deep image prior (DIP) approach employs an unconditional model that maps relatively low-dimensional latent variables onto images as a native MRI prior [28]. The imaging operator is embedded during inference, and the prior is adapted to minimize a data-consistency loss on undersampled test data. This decoupled approach introduces flexibility in employing the same prior for various different imaging operators. Yet, DIP methods pervasively use untrained CNN architectures with randomly initialized parameters [28–31, 69]. In turn, untrained priors might be suboptimal in capturing the distribution of MR images [32, 33], and CNN models can generally suffer from limited sensitivity toward long-range spatial interactions [35–37].

Outline

The organization for the rest of the thesis is as follows. Second chapter explains the fundamentals of accelerated MRI. Third chapter includes details of the theory behind the proposed model. Fourth chapter includes details of the proposed network architecture, competing methods and datasets used. Fifth chapter includes the quantitative results. Lastly, sixth chapter stands for the conclusion remarks. Supplementary Text, Tables and Figures can be viewed in the Appendix. The contents of this thesis reflect the work reported in the following publications:

- Y. Korkmaz, S. U. Dar, M. Yurt, M. Özbey, and T. Çukur, “Unsupervised MRI reconstruction via zero-shot learned adversarial transformers,” *arXiv:2105.08059*, 2021
- Y. Korkmaz, S. U. Dar, M. Yurt, M. Ozbey, and T. Çukur, “Zero-shot learning for unsupervised reconstruction of accelerated MRI acquisitions,” in *Proceedings of ISMRM*, 2021
- Y. Korkmaz, M. Yurt, S. U. H. Dar, M. Özbey, and T. Cukur, “Deep mri reconstruction with generative vision transformers,” in *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 54–64, Springer, 2021

Chapter 2

Fundamentals of Accelerated MRI

Magnetic Resonance Imaging (MRI) relies on the excitation and alignment of hydrogen nuclei (1H) in the human body by using magnetic field and radio-frequency pulses. While a strong external magnetic field aligns the hydrogen nuclei, a repeating sequence of radio-frequency pulses, produced by the MRI scanner, excite the protons in the transverse direction. Receiver coils in MRI scanner capture the changes in magnetic flux while these excited protons relax to realign with the magnetic field. Raw data in MRI, also known as the k-space, are these transversal components of the magnetization stored as a data matrix where every point contains a specific spatial frequency information in the target MR image. While outer rows in data matrix corresponds to the higher spatial frequency components like the details in tissues, inner rows encodes the lower spatial frequency components like contrast of the target MR image. The relationship between k-space and image domains are defined using Fourier transform.

MRI is a highly preferable imaging technique due to its exceptional soft-tissue contrast. However, long scan duration limits the usage of MRI especially in the emergency cases [73, 74]. Acceleration of MRI while keeping high image quality is an open research area. To accelerate MRI several techniques have been proposed including parallel imaging, compressed sensing, simultaneous multiple

slice imaging, and ultra fast sequences [75–116].

Among other acceleration methods, parallel imaging is a reconstruction technique where undersampled k-space data are acquired using an array of independent receiver coils where each coil is more sensitive to the specific volume of the nearest tissue (Fig 2.1). Typically, k-space acquisitions are undersampled to speed up scans which is formulated as:

$$F_p C m = y_s \quad (2.1)$$

where F_p is the partial Fourier operator, C denotes coil sensitivity maps, m is the MR image to be reconstructed and y_s is the acquired k-space data. Reconstruction of m from y_s is an underdetermined problem therefore requires prior information of m . Prior information on MR images incorporated as follows:

$$\hat{m} = \underset{m}{\operatorname{argmin}} \|y_s - F_p C m\|_2^2 + H(m) \quad (2.2)$$

where \hat{m} is the reconstructed image, and $H(m)$ is the regularization term. The regularization term can be designed based on sparsity priors [3,4], structured low-rank priors [117–120], or recently introduced deep neural network priors [121–134]. Generally, supervised deep neural network (DNN) based priors are learned during offline training sessions where network parameters are optimized to minimize the following loss function:

$$\mathcal{L}(\theta) = \|G(m_{us}^t; \theta) - m_{fs}^t\|_p \quad (2.3)$$

where m_{us}^t and m_{fs}^t are undersampled and fully sampled MR image pairs used in training respectively, $G(\cdot; \theta)$ denotes the image generator with network parameters θ , and $\|\cdot\|_p$ is the norm function where p is typically selected as 1 or 2. After offline training of the network parameters, reconstruction with prior information (Eq. 2.2) can be formulated during inference as:

$$\hat{m} = \underset{m}{\operatorname{argmin}} \|y_s - F_p C m\|_2^2 + \beta \|G(m_{us}; \hat{\theta}) - m\|_p \quad (2.4)$$

where β stands for the parameter that adjusts weighting between deep network and consistency priors, $G(m_{us}; \hat{\theta})$ stands for network output using trained parameter set $\hat{\theta}$ with undersampled input image m_{us} .

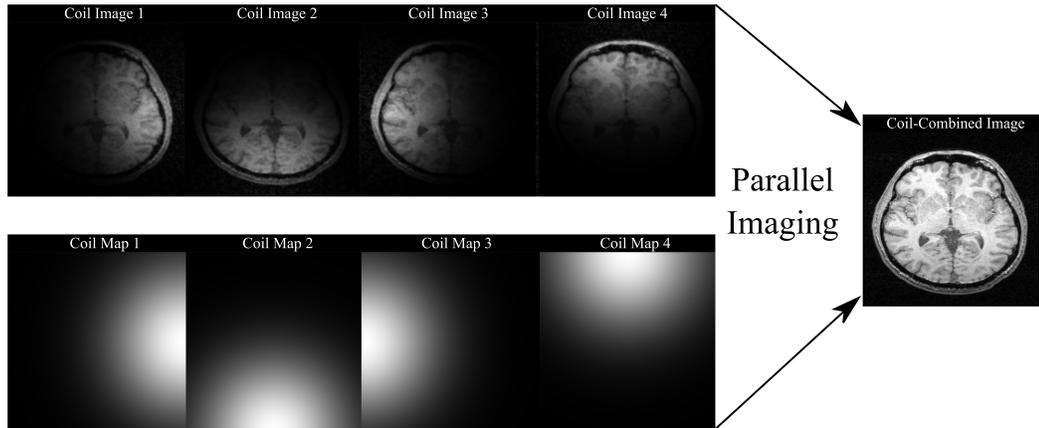


Figure 2.1: In parallel imaging, images gathered from independent receiver coils, which are mostly sensitive to the nearest tissues, are combined by using known coil-sensitivity maps.

Although the supervised deep network priors require offline training sessions, recently introduced Deep Image Prior (DIP) frameworks relies on the untrained deep models by using Convolutional Neural Networks (CNNs) as native image regularizers. Instead of the offline training sessions, DIP models optimize randomly initialized network parameters during the inference by enforcing consistency with the available corrupted images. In DIP models, optimal network parameters (θ^*) are found by solving the following minimization problem:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \|F_p C d_{\theta}(z) - y_s\|_1 \quad (2.5)$$

where θ are network parameters and $d_{\theta}(z)$ is the generated image by the deep network using latent variables z . Then, the final reconstructed image $d_{\theta^*}(z)$ is generated by using the found optimal network parameters.

Chapter 3

Theory

Inspired by the DIP framework, here we introduce a reconstruction method for accelerated MRI based on a deep generative model that maps noise and latent variables onto MR images. Unlike previous methods, this mapping is based on a style-generative adversarial network with transformer blocks. The network learns an MRI prior in pre-training phase, followed by a zero-shot inference phase where it is adapted to reconstruct undersampled acquisitions. We first overview the inverse problem formulation in accelerated MRI and DIP. We then describe the fundamental building blocks of SLATER.

3.1 Accelerated MRI Reconstruction

In accelerated MRI, undersampled k-space acquisitions are performed to speed up scans, typically with variable-density random sampling patterns:

$$F_p C m = y_s \tag{3.1}$$

F_p is the partial Fourier operator defined by the sampling pattern, C denotes coil sensitivities, m is the target MR image and y_s are the acquired multi-coil k-space samples. The target m must be computed given the available data y_s . However, the linear system in Eq.3.1 is underdetermined, so MRI reconstruction is an ill-posed inverse problem. To obtain high-quality reconstruction, the solution has

to be regularized with additional prior information on MR images:

$$\hat{m} = \underset{m}{\operatorname{argmin}} \|y_s - F_p C m\|_2^2 + H(m) \quad (3.2)$$

where \hat{m} is the reconstructed image, and $H(m)$ is the regularization term.

3.2 Deep Image Prior

The DIP framework has recently been introduced for unsupervised learning in computer vision tasks, including super resolution, inpainting and denoising [135]. DIP observes that local filtering operations in CNNs constrain the set of images that can be generated, so untrained CNNs can serve as native image regularizers. DIP performs random initialization of network inputs and weights without any pretraining. During inference, network parameters are optimized by enforcing consistency with the available corrupted image. Thus, DIP projects the corrupted image onto the space of CNN-generatable images to filter out corruptions such as blur or noise. DIP can be adopted for MRI reconstruction as follows:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \|F_p C d_\theta(z) - y_s\|_1 \quad (3.3)$$

where θ are network weights, θ^* are optimized network weights, z are latent variables, $d_\theta(z)$ is the network mapping from latents onto the reconstructed image [14, 28]. Both network weights and latents are randomly initialized, and the optimization in Eq.3.3 is performed over θ , while z is fixed. In contrast to mainstream learning-based methods, DIP inverts a random network prior to identify weights that are most consistent with the corrupted image. The reconstruction can then be expressed as:

$$\hat{m} = d_{\theta^*}(z) \quad (3.4)$$

Despite its prowess in MRI reconstruction, models with randomly initialized parameters setup a generic image prior that may not be as strongly-tuned towards the distribution of MR data as trained models [32]. Furthermore, DIP is traditionally based on CNN architectures, where local kernels introduce suboptimal sensitivity in capture of long-range spatial interactions [35].

3.3 Adversarial Transformer Model

SLATER is based on an unconditional adversarial network that receives noise and latent variables to generate MR images [71]. Here we adopt a style-generative architecture given the success of this model family in computer vision tasks [136]. We further propose to build the network layers with cross-attention transformer blocks as inspired by a recent study on natural image synthesis [137]. Self-attention transformers that compute interactions among all image pixels are prohibitive at relatively high spatial resolutions [36,138,139]. Instead, cross-attention transformers enable efficient capture of long-range context based on attentional interactions between low-dimensional latent variables and high-dimensional image features. Our model contains three sub-networks (Fig.4.2): a synthesizer that generates MR images; a mapper that prepares the set of latent variables input to the synthesizer; and a discriminator. The synthesizer aims to generate realistic images, while the discriminator aims to distinguish actual images from synthesized images. Thus, the discriminator aids the synthesizer in capturing an MRI prior that reflects the distribution of high-quality MR images.

3.3.1 Synthesizer (G)

The synthesizer contains a total of N_L layers, each comprising a convolutional upsampling block to progressively increase image resolution followed by a cross-attention transformer block. Prior style-generative models typically use a global latent variable at each layer to control high-level image features related to style [136]. In SLATER, in addition to a global latent ($w_g \in \mathbb{R}^{L_s}$, L_s : dimensionality), the synthesizer receives as input K local ($W_l \in \mathbb{R}^{K \times L_s}$) latent variables at each layer. The global latent is still used to perform a spatially-uniform affine transformation of convolutional feature maps, and so it modulates high-level image features. Meanwhile, local latents are used to perform spatially-selective modulation of feature maps via cross-attention mechanisms. Each local latent variable

focuses on a learned group of feature-map locations, so local latents serve to modulate relatively lower-level image features. For enabling cross-attention mechanism, a sinusoidal position encoding is used for spatial feature maps, whereas a learnable position encoding is assumed for local latent variables as their positions are unknown prior to training. Each cross-attention transformer block contains a serial cascade of the following sub-blocks: (cross-attention (CA^1), noise injection (NI^1), style-modulated convolution (SC), cross-attention (CA^2), noise injection (NI^2)). These sub-blocks are detailed below.

3.3.1.a Cross-Attention (CA^1)

CA derives contextual representations by mediating attentional interaction between W_l and input feature maps $X_i^0 \in \mathbb{R}^{h_1 \times h_2 \times u}$, where h_1 and h_2 denote height and width depending on the resolution at the i^{th} layer, and u is the number of feature channels. Let $X_{i,vec}^0 \in \mathbb{R}^{(h_1 \times h_2) \times u}$ be the vectorized form of X_i^0 along the spatial dimensions. Attention maps $\tilde{A}_{i,maps}^1 \in \mathbb{R}^{(h_1 \times h_2) \times K}$ that characterize the relation between W_l and $X_{i,vec}^0$ are as follows (see Fig. 3.1 and Supp. Fig. A.2 for sample maps):

$$\tilde{A}_{i,maps}^1 = \mathit{smax} \left(\frac{\tilde{q}_i^1(X_{i,vec}^0 + PE_{i,X}^1) \tilde{k}_i^1(W_l + PE_{W_l})^T}{\sqrt{u}} \right) \quad (3.5)$$

where smax is the softmax function, $\tilde{q}_i^1(\cdot) \in \mathbb{R}^{(h_1 \times h_2) \times u}$ are queries that receive $X_{i,vec}^0$ added with layer-specific position encoding variables $PE_{i,X}^1 \in \mathbb{R}^{(h_1 \times h_2) \times u}$ and perform a learnable linear projection, $\tilde{k}_i^1(\cdot) \in \mathbb{R}^{K \times u}$ are keys that receive W_l added with position encoding variables $PE_{W_l} \in \mathbb{R}^{K \times L_s}$ and perform a learnable linear projection. $PE_{i,X}^1$ are taken as learnable linear projections of pre-defined sinusoidal encoding variables (see Appendix A.1.1), while PE_{W_l} are randomly initialized learnable position encoding variables. Note that $PE_{i,X}^1$ and $A_{i,maps}^1$ are vectorized along the spatial dimensions. Attention feature maps ($\tilde{A}_{i,feat}^1 \in \mathbb{R}^{(h_1 \times h_2) \times u}$) are then obtained by multiplying $A_{i,maps}^1$ with values that are learnable linear projections ($\tilde{v}_i^1(\cdot) \in \mathbb{R}^{K \times u}$).

$$\tilde{A}_{i,feat}^1 = \tilde{A}_{i,maps}^1 \tilde{v}_i^1(W_l + PE_{W_l}) \quad (3.6)$$

Lastly, input feature maps are modulated via a single attention head that scales and shifts $X_{i,vec}^0$ with learnable linear projections of $\tilde{A}_{i,feat}^1$ to output $X_{i,vec}^1 \in \mathbb{R}^{(h_1 \times h_2) \times u}$.

$$X_{i,vec}^1 = \tilde{\gamma}_1(\tilde{A}_{i,feat}^1) \odot \left(\frac{X_{i,vec}^0 - \mu(X_{i,vec}^0)}{\sigma(X_{i,vec}^0)} \right) + \tilde{b}_1(\tilde{A}_{i,feat}^1) \quad (3.7)$$

where $\tilde{\gamma}_1(\cdot) \in \mathbb{R}^{(h_1 \times h_2) \times u}$ and $\tilde{b}_1(\cdot) \in \mathbb{R}^{(h_1 \times h_2) \times u}$ are learnable linear projections, μ denotes mean, σ denotes variance, and \odot is the Hadamard product. The mapping through CA^1 can thus be summarized as: $X_i^1 = CA^1(X_i^0)$, where $X_i^1 \in \mathbb{R}^{h_1 \times h_2 \times u}$ is the matrix form of $X_{i,vec}^1$.

3.3.1.b Noise Injection (NI^1)

To improve control over variability in fine details of feature maps, noise variables are injected onto modulated feature maps from CA^1 . Given input $X_i^1 \in \mathbb{R}^{h_1 \times h_2 \times u}$ to NI^1 , the output ($X_i^2 \in \mathbb{R}^{h_1 \times h_2 \times u}$) can be expressed as:

$$X_i^2 = \begin{bmatrix} X_i^{1,1} + \alpha_i^1 n_i^1 \\ \vdots \\ X_i^{1,u} + \alpha_i^1 n_i^1 \end{bmatrix} \quad (3.8)$$

where $X_i^{1,e} \in \mathbb{R}^{h_1 \times h_2}$ denotes e^{th} channel of X_i^1 , $n_i^1 \in \mathbb{R}^{h_1 \times h_2}$ is noise added to each e^{th} channel and α_i^1 is a learnable scalar. Note that the learnable noise variables in n_i^1 are initiated via random sampling from a standard normal distribution for each spatial location. During the course of learning, mean and standard deviation of noise variables are normalized to (0, 1) across the spatial dimensions. The final mapping through NI^1 is expressed as: $X_i^2 = NI^1(X_i^1)$.

3.3.1.c Style Modulated Convolution (SC)

Transformer blocks characteristically contain a feed-forward neural network (FFNN) sub-block following the attention sub-block to extract hidden features of attention-based contextual representations. For computational efficiency, here a convolutional FFNN is utilized to locally refine contextual representations while

modulating feature maps to control high-level style features [136]. Analogous to adaptive instance normalization (AdaIN) in style-transfer models, modulation is achieved via an affine transformation that controls the scale of feature maps [140]. Yet, we opted for a more compact implementation based on style-modulated convolution, with comparable complexity to convolution augmented with AdaIN and only requiring a trivial scaling operation over basic convolution [136]. Accordingly, feature maps are scaled by modulating convolution kernels in SC, where kernel weights are multiplied via a learnable linear projection $w_{i,s} \in \mathbb{R}^u$ of the global latent variable w_g :

$$\theta' = \begin{bmatrix} w_{i,s}^1 \theta_i^{1,1} & w_{i,s}^1 \theta_i^{1,2} & \dots & w_{i,s}^1 \theta_i^{1,u} \\ w_{i,s}^2 \theta_i^{2,1} & w_{i,s}^2 \theta_i^{2,2} & \dots & w_{i,s}^2 \theta_i^{2,u} \\ \vdots & \vdots & \ddots & \vdots \\ w_{i,s}^u \theta_i^{u,1} & w_{i,s}^u \theta_i^{u,2} & \dots & w_{i,s}^u \theta_i^{u,u} \end{bmatrix} \quad (3.9)$$

where $\theta_i^{e,h} \in \mathbb{R}^{r \times r}$ denotes the two-dimensional (2D) kernel for e^{th} input and h^{th} output channel, $w_{i,s}^e$ is the scaling coefficient for e^{th} input channel. Note that $\theta'_i \in \mathbb{R}^{r \times r \times u \times u}$ is a 4D tensor of modulated kernels, so the matrix expression in Eq. 3.9 depicts formation along the third and fourth tensor dimensions. While convolution with the modulated kernels changes the relative scaling of input feature channels for controlling style, it can also alter the overall scale of the output feature map for each channel. To restore output feature maps to unit standard deviation, the modulated kernels are normalized to unit-norm across the output channel dimension. Following the notation in Eq. 3.9, this normalization can be expressed as:

$$\theta''_i = \begin{bmatrix} \frac{\theta_i'^{1,1}}{\sqrt{\sum_c (\theta_i'^{c,1})^2}} & \frac{\theta_i'^{1,2}}{\sqrt{\sum_c (\theta_i'^{c,2})^2}} & \dots & \frac{\theta_i'^{1,u}}{\sqrt{\sum_c (\theta_i'^{c,u})^2}} \\ \frac{\theta_i'^{2,1}}{\sqrt{\sum_c (\theta_i'^{c,1})^2}} & \frac{\theta_i'^{2,2}}{\sqrt{\sum_c (\theta_i'^{c,2})^2}} & \dots & \frac{\theta_i'^{2,u}}{\sqrt{\sum_c (\theta_i'^{c,u})^2}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\theta_i'^{u,1}}{\sqrt{\sum_c (\theta_i'^{c,1})^2}} & \frac{\theta_i'^{u,2}}{\sqrt{\sum_c (\theta_i'^{c,2})^2}} & \dots & \frac{\theta_i'^{c,u}}{\sqrt{\sum_c (\theta_i'^{c,u})^2}} \end{bmatrix} \quad (3.10)$$

where $\theta_i'' \in \mathbb{R}^{r \times r \times u \times u}$ denotes de-modulated kernels and c corresponds to channel index. Finally, the output feature maps ($X_i^3 \in \mathbb{R}^{h_1 \times h_2 \times u}$) can be computed as:

$$X_i^3 = \begin{bmatrix} \sum_c X_i^{2,c} \circledast \theta_i''^{c,1} \\ \vdots \\ \sum_c X_i^{2,c} \circledast \theta_i''^{c,u} \end{bmatrix} \quad (3.11)$$

where $X_i^{2,c} \in \mathbb{R}^{h_1 \times h_2}$ are input feature maps, \circledast is convolution. The mapping through the convolution sub-block SC^1 is expressed as: $X_i^3 = SC(X_i^2)$.

3.3.1.d Cross-Attention (CA^2)

A second cross-attention sub-block is used following the SC sub-block to further boost sensitivity of the model to global context. This maps the output of the SC sub-block to attention-modulated feature maps ($X_i^4 \in \mathbb{R}^{h_1 \times h_2 \times u}$) as: $X_i^4 = CA^2(X_i^3)$.

3.3.1.e Noise-Injection (NI^2)

Finally, noise variables are injected to control variability in the fine details of feature maps following the second CA sub-block: $X_i^5 = NI^2(X_i^4)$. where $X_i^5 \in \mathbb{R}^{h_1 \times h_2 \times u}$ is the output of NI^2 .

3.3.2 Mapper (M)

The mapper projects independent and identically distributed random variables ($Z = z_1, z_2, \dots, z_K, z_g; z \in \mathbb{R}^{1 \times L_s}$) onto a refined set of local and global latent variables ($w_1, w_2, \dots, w_K, w_g; w \in \mathbb{R}^{1 \times L_s}$) expected by the synthesizer [136]. M is a multi-layered architecture comprising a first stream dedicated to the global latent and a second stream of dedicated to local latents (see Supp. Fig. A.1). The global latent is processed with a cascade of fully-connected layers [136]. Local latents are instead processed with self-attention blocks to enable interactions among these

variables. Self-attention blocks contain a cascade of self-attention (SA), fully-connected (FC^1), fully-connected (FC^2) sub-blocks as detailed below.

3.3.2.a Self-Attention (SA)

At the i^{th} self-attention block, intermediate activations for the K latent variables are concatenated as input:

$$Z_i^0 = (z_1 \oplus_1 \cdots \oplus_1 z_K)_i^0 \quad (3.12)$$

where $Z_i^0 \in \mathbb{R}^{K \times L_s}$ and \oplus_1 is the concatenation operator along the first dimension. Attention maps ($\hat{A}_{i,maps} \in \mathbb{R}^{K \times K}$) are first obtained:

$$\hat{A}_{i,maps} = \mathit{smax} \left(\frac{\hat{q}_i(Z_i^0 + PE_{Z^0}) \hat{k}_i(Z_i^0 + PE_{Z^0})^T}{\sqrt{L_s}} \right) \quad (3.13)$$

where $PE_{Z^0} \in \mathbb{R}^{K \times L_s}$ denotes learnable position encoding variables, $\hat{q}_i(\cdot) \in \mathbb{R}^{K \times L_s}$ and $\hat{k}_i(\cdot) \in \mathbb{R}^{K \times L_s}$ are queries and keys respectively. Attention feature maps ($\hat{A}_{i,feat} \in \mathbb{R}^{K \times L_s}$) are given as:

$$\hat{A}_{i,feat} = \hat{A}_{i,maps} \hat{v}_i(Z_i^0 + PE_{Z^0}) \quad (3.14)$$

where $\hat{v}_i(\cdot) \in \mathbb{R}^{K \times L_s}$ are values. $\hat{A}_{i,feat}$ is then used to scale and shift Z_i^0 :

$$Z_i^1 = \hat{\gamma}(\hat{A}_{i,feat}) \odot \left(\frac{Z_i^0 - \mu(Z_i^0)}{\sigma(Z_i^0)} \right) + \hat{b}(\hat{A}_{i,feat}) \quad (3.15)$$

where $Z_i^1 \in \mathbb{R}^{K \times L_s}$ is the output, $\hat{\gamma}(\cdot)$ and $\hat{b}(\cdot)$ are learnable linear projections. Lastly, Z_i^1 is decomposed into individual latents, $Z_i^1 \rightarrow \{z_1, \dots, z_K\}_i^1$.

3.3.2.b Fully Connected (FC)

Next, each latent is separately processed with two fully-connected sub-blocks (FC^1 and FC^2):

$$\{z_j\}_i^2 = FC^2(FC^1(\{z_j\}_i^1)) \quad (3.16)$$

where $\{z_j\}_i^2 \in \mathbb{R}^{1 \times L_s}$ is the output FC^2 for the j^{th} local latent.

3.3.3 Discriminator (D)

Adversarial models involve an interplay between the synthesizer that generates images and a separate discriminator sub-network [141]. In SLATER, D aims to accurately distinguish images generated by the synthesizer from actual MR images. A feed-forward architecture is employed here with a cascade of convolutional layers augmented with several fully-connected layers. The mapping through D can be compactly expressed as:

$$x_D = D_{\theta_D}(x) \quad (3.17)$$

where $x_D \in \mathbb{R}^1$ is the output of the discriminator, and x is either an actual MR image (x_r) or an image generated by the synthesizer $G(M(Z))$.

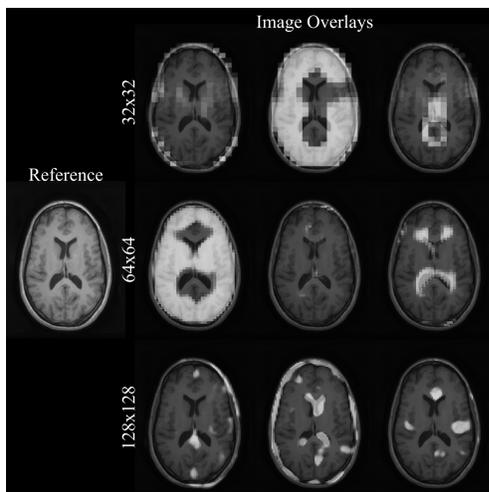


Figure 3.1: Cross-attention maps in SLATER for a T_1 -weighted acquisition. Sample maps from the first CA sub-block are displayed by overlay onto the respective MR image across three resolutions (i.e., at network layers 4-6). Attention maps for separate latents typically show segregated spatial distribution, and tend to group tissue clusters with similar signal intensity and texture broadly distributed across the image.

3.4 Self- versus Cross-Attention Transformers

The main components of vanilla transformers are a self-attention (SA) sub-block followed by a feed-forward neural network (FFNN) sub-block [36]. Self-attention

mechanisms serve to explicitly relate different positions within an image to compute contextual representations at each image pixel. Thus, SA is the primary component that learns long-range dependencies, while FFNN performs nonlinear transformations to extract hidden features of attention-based contextual representations. While vanilla transformers with fully-connected FFNN are common in natural language processing [142], computer vision studies often introduce vision-specific modifications for computational efficiency [137]. Cross-attention transformers in SLATER involve two key modifications: cross-attention sub-blocks that use a compact set of latents implicitly relating image pixels to learn contextual representations, and a convolutional FFNN to improve computational efficiency. Here, we overview self-attention and cross-attention transformers in terms of their ability to capture long-range spatial relationships and model complexity. For simplicity, remaining operations in transformers such as normalization or skip connections are ignored, and only a single feature channel and a single attention head are considered.

3.4.1 Self-attention transformer

Given an input feature map $X_{vec}^0 \in \mathbb{R}^{(h_1 \times h_2)}$, let a, b be two distant pixels whose $r \times r$ neighborhoods $\Delta a, \Delta b$ do not spatially overlap. A convolution sub-block with kernel size r would perform localized processing for each pixel in its $r \times r$ neighborhood:

$$X_{vec}^1 = Conv(X_{vec}^0) \quad (3.18)$$

where $X_{vec}^1 \in \mathbb{R}^{(h_1 \times h_2)}$. Intensities of a, b in X_{vec}^1 will thus be conditionally independent given intensities of X_{vec}^0 in $\Delta a, \Delta b$:

$$E\{X_{vec}^1[a]X_{vec}^1[b] \mid X_c\} = E\{X_{vec}^1[a] \mid X_c\}E\{X_{vec}^1[b] \mid X_c\} \quad (3.19)$$

where $E\{\cdot\}$ denotes expectation, $X_c = X_{vec}^0[\Delta a \cup \Delta b]$ and \cup is the union operator. As a result, convolutional processing does not leverage long-range relationships in feature maps.

In contrast, the SA sub-block in a transformer characterizes the relation between all spatial locations in $X_{vec}^0 \in \mathbb{R}^{(h_1 \times h_2)}$ via an attention map:

$$A_{maps} = \text{smax} (q(X_{vec}^0 + PE_X)k(X_{vec}^0 + PE_X)^T) \quad (3.20)$$

where $PE_X, q, k \in \mathbb{R}^{(h_1 \times h_2)}$ are positional encoding, query, and key, and $A_{maps} \in \mathbb{R}^{(h_1 \times h_2) \times (h_1 \times h_2)}$. Attention-modulated feature map is then derived:

$$A_{feat} = A_{maps}v(X_{vec}^0 + PE_X) \quad (3.21)$$

where $A_{feat} \in \mathbb{R}^{(h_1 \times h_2)}$, and $v \in \mathbb{R}^{(h_1 \times h_2)}$ is value. Intensities of A_{feat} at pixels a, b are:

$$\begin{aligned} A_{feat}[a] &= A_{maps}[a, 1 : (h_1 \times h_2)] v[1 : (h_1 \times h_2)] \\ A_{feat}[b] &= A_{maps}[b, 1 : (h_1 \times h_2)] v[1 : (h_1 \times h_2)] \end{aligned} \quad (3.22)$$

As seen here, both $A_{feat}[a]$ and $A_{feat}[b]$ are functions of all pixels in v and thereby X_{vec}^0 . Next, the FFNN sub-block processes A_{feat} to extract hidden representations, typically via a fully-connected architecture:

$$X_{vec}^1 = FC(A_{feat}) \quad (3.23)$$

where $X_{vec}^1 \in \mathbb{R}^{(h_1 \times h_2)}$. Note that $A_{feat}[a]$ and $A_{feat}[b]$ are readily dependent on all pixels. Intensities of pixels a, b in output feature maps are then statistically dependent even when conditioned on $X_c = X_{vec}^0[\Delta a \cup \Delta b]$:

$$\mathbb{E}\{X_{vec}^1[a]X_{vec}^1[b] \mid X_c\} \neq \mathbb{E}\{X_{vec}^1[a] \mid X_c\}\mathbb{E}\{X_{vec}^1[b] \mid X_c\} \quad (3.24)$$

Thus, vanilla transformers utilize dependencies across distant pixels to compute contextual representations.

The SA sub-block examines interactions among all possible pairs of pixels as described in Eqs. 3.20 and 3.22, so it has a computational complexity of $O((h_1 \times h_2)^2)$. Likewise, the FFNN sub-block as described in Eq. 3.23 exhaustively considers inter-pixel interactions with a complexity of $O((h_1 \times h_2)^2)$. This quadratic complexity with respect to image size limits the applicability of self-attention transformers at relatively high resolutions encountered in MRI [36].

3.4.2 Cross-attention transformer

Instead of exhaustively modeling inter-pixel interactions in high-dimensional feature maps, the CA sub-block uses a small set of K latent variables to implicitly characterize these interactions:

$$A_{maps} = \mathit{softmax}(q(X_{vec}^0 + PE_X)k(W_l + PE_{W_l})^T) \quad (3.25)$$

where $A_{maps} \in \mathbb{R}^{(h_1 \times h_2) \times K}$, $W_l \in \mathbb{R}^{K \times L_s}$: local latent variables, $PE_{W_l} \in \mathbb{R}^{K \times L_s}$: positional encoding for W_l . Attention-modulated feature map can be expressed as:

$$A_{feat} = A_{maps}v(W_l + PE_{W_l}) \quad (3.26)$$

where $A_{feat} \in \mathbb{R}^{(h_1 \times h_2)}$, and $v \in \mathbb{R}^K$ is value. In turn, intensities of pixels a, b are:

$$\begin{aligned} A_{feat}[a] &= A_{maps}[a, 1 : K] v[1 : K] \\ A_{feat}[b] &= A_{maps}[a, 1 : K] v[1 : K] \end{aligned} \quad (3.27)$$

As seen above, both $A_{feat}[a]$ and $A_{feat}[b]$ are functions of all latent variables in W_l , so they are statistically dependent.

Next, a convolutional FFNN processes A_{feat} to extract hidden representations:

$$X_{vec}^1 = \mathit{Conv}(A_{feat}) \quad (3.28)$$

where $X_{vec}^1 \in \mathbb{R}^{(h_1 \times h_2)}$. Although convolutional processing is local, the dependency introduced in the CA sub-block carries over to the output feature maps:

$$\begin{aligned} \mathbb{E}\{A_{feat}[a]A_{feat}[b] \mid X_c\} &\neq \mathbb{E}\{A_{feat}[a] \mid X_c\}\mathbb{E}\{A_{feat}[b] \mid X_c\} \\ \mathbb{E}\{X_{vec}^1[a]X_{vec}^1[b] \mid X_c\} &\neq \mathbb{E}\{X_{vec}^1[a] \mid X_c\}\mathbb{E}\{X_{vec}^1[b] \mid X_c\} \end{aligned} \quad (3.29)$$

Therefore, cross-attention transformers can model long-range dependencies to compute contextual representations.

The CA sub-block examines interactions among image pixels and local latents as described in Eqs. 3.25 and 3.27, so it has a computational complexity of $O((h_1 \times h_2) \times K)$. The convolutional FFNN considers local interactions with a complexity of $O((h_1 \times h_2) \times r^2)$. Because $K \ll (h_1 \times h_2)$ and $r^2 \ll (h_1 \times h_2)$ typically, the cross-attention transformer achieves notably lower complexity to permit use at higher spatial resolutions.

3.5 Learning Procedures

SLATER uses a two-stage strategy towards MRI reconstruction with a pre-training phase to learn the MRI prior, followed by a zero-shot reconstruction phase to embed the imaging operator. These two phases are detailed below.

3.5.1 Pre-training of the MRI prior

Since SLATER completely decouples the MRI prior from the imaging operator, pre-training assumes no prior knowledge on the imaging operator such as undersampling patterns or coil sensitivity encoding. Instead, the adversarial transformer model is trained to capture a prior on coil-combined, complex MR images, derived from fully-sampled acquisitions. Note that the synthesizer in this unconditional model maps noise and latent variables onto MR images, unlike conditional models with explicitly defined input-output relationships (i.e. undersampled versus fully-sampled data). Therefore, SLATER’s pre-training is categorized as an unsupervised generative modeling task where the distribution of MR images is learned so that new, random samples can be drawn from the distribution [143].

Adversarial models commonly involve synthesizer and discriminator sub-networks that are trained with inter-linked loss functions to improve quality of synthesized images [141]. In SLATER, the synthesizer along with the mapper that provides latent variables are trained to minimize a common adversarial loss based on non-saturating logistic function:

$$L_{G,M}(\theta_G, \theta_M) = -\mathbb{E}_{p(Z)}\{\log(f(D(G_{\theta_G}(M_{\theta_M}(Z))))))\} \quad (3.30)$$

where $\mathbb{E}_{p(\cdot)}$ is expectation with probability density p , $f(\cdot)$ is the sigmoid function, θ_G are parameters of the synthesizer, θ_M are parameters of the mapper. This particular loss is preferred in order to prevent the saturation problem in adversarial learning, where the discriminator starts outperforming the synthesizer by a significant margin and learning stops prematurely [141].

Meanwhile, the discriminator is trained to minimize an adversarial loss based

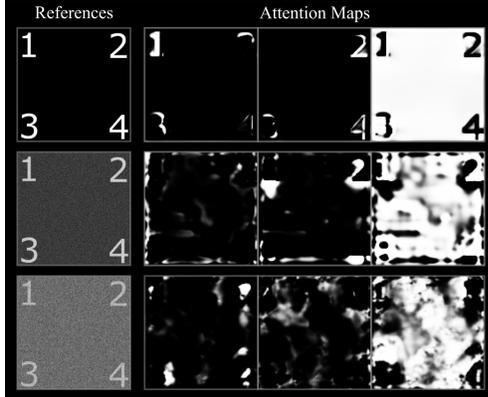


Figure 3.2: Cross-attention maps in SLATER for a simulated phantom with varying levels of normally-distributed white noise. Sample maps from the first CA sub-block are displayed at 128x128 resolution (network layer 6). Relative to a peak signal intensity of 1, top, middle and bottom rows show results for no noise, noise variance of 0.01, and noise variance of 0.1, respectively.

on non-saturating logistic function augmented with a gradient penalty term:

$$\begin{aligned}
 L_D(\theta_D) = & -\mathbb{E}_{p(Z)}\{\log(1 - f(D_{\theta_D}(G(M(Z)))))\} \\
 & -\mathbb{E}_{p(x_r)}\{\log(f(D_{\theta_D}(x_r)))\} \\
 & +\frac{\eta}{2}\mathbb{E}_{p(x_r)}\{\|\nabla D_{\theta_D}(x_r)\|^2\}
 \end{aligned} \tag{3.31}$$

where x_r denotes coil-combined, complex MR images derived from actual scans, η is the regularization parameter, and θ_D are the parameters of the discriminator. The first two terms define the adversarial loss, whereas the third term has been suggested to improve adversarial learning by enforcing limited gradients in x_r according to the learned distribution [144].

3.5.2 Zero-shot reconstruction

The learned MRI prior does not contain any information regarding the conditional mapping from undersampled to fully-sampled data. Thus, the prior is conjoined with the imaging operator during inference to reconstruct test data. To adapt the pre-trained generative model to the reconstruction task, a data-consistency loss is employed on undersampled acquisitions to optimize synthesizer parameters (Fig. 4.1). Note that zero-shot learning is an unsupervised task-adaptation

approach, where a model trained for an initial task is later transferred to a different target task without using additional training samples [145]. Analogously, SLATER adapts its adversarial model pre-trained to perform generative modeling of MR images to perform MRI reconstruction without any extra training samples. Therefore, the inference phase of SLATER is categorized as zero-shot reconstruction.

During inference, we optimize all components of the synthesizer including noise (n), latent variables (W) and weights (θ_G) to minimize the data-consistency loss. The synthesizer outputs a coil-combined, complex MR image, which is back-projected onto individual coils given sensitivity estimates, and then Fourier transformed to select available k-space coefficients according to the undersampling pattern [54]. Consistency of acquired and reconstructed k-space coefficients in the test data is then computed:

$$L_{\hat{R}}(W, n, \theta_G) = \lambda_1 \|F_p CG(W, n, \theta_G) - k_x\|_1 + \lambda_2 \|F_p CG(W, n, \theta_G) - k_x\|_2 \quad (3.32)$$

where \hat{R} : reconstructed image, k_x : acquired k-space coefficients, F_p : partial Fourier operator, G : synthesizer. Data consistency is taken as a weighted $\ell_1 - \ell_2$ -norm loss, where (λ_1, λ_2) : weightings of (ℓ_1, ℓ_2) loss components. This loss function is considered to offer a more balanced weighting of errors across k-space compared to ℓ_2 -norm that can be over-sensitive to lower spatial frequencies [24]. During pre-training, a common set of latent variables produced by the mapper are input uniformly across all CA and SC sub-blocks, whereas sub-block specific noise is included. To improve performance during zero-shot reconstruction, latent variables for each sub-block within each layer are instead segregated for independent optimization. In Eq. 3.32, $W = W_{1,\dots,N_L}$ where W_i denotes the collection of latent variables for the i^{th} layer and contains the local latents for the two CA sub-blocks ($W_{l,i}^{1,2}$) and the global latent for the SC sub-block ($w_{g,i}$); $n = n_{1,\dots,N_L}$ where n_i denotes the collection of noise components for the i^{th} layer and contains noise for the two NI sub-blocks ($n_i^{1,2}$).

Chapter 4

Methods

4.1 Network Architecture

Architecture of SLATER’s synthesizer, mapper and discriminator are described below (see also Appendix A.1.2).

4.1.1 Synthesizer

The synthesizer is a multi-layer architecture where image resolution is progressively increased. Each layer comprises a convolutional upsampling block to increase image resolution by a factor of 2 followed by a cross-attention transformer block, and a skip connection to add the upsampled input. The first layer receives a constant input randomly drawn from a standard normal distribution. As the input layer, the first layer does not contain the upsampling blocks and the first attention sub-block is replaced with an identity transformation. The last layer calculates the final synthesizer output. Attention sub-blocks are omitted in the last layer to preserve precise localization in high-resolution images with convolution operators, as contextual representations have been extracted in previous layers [36]. Two separate channels are used to output real and imaginary parts of images. The upsampling block uses transpose convolution, and upsampling and

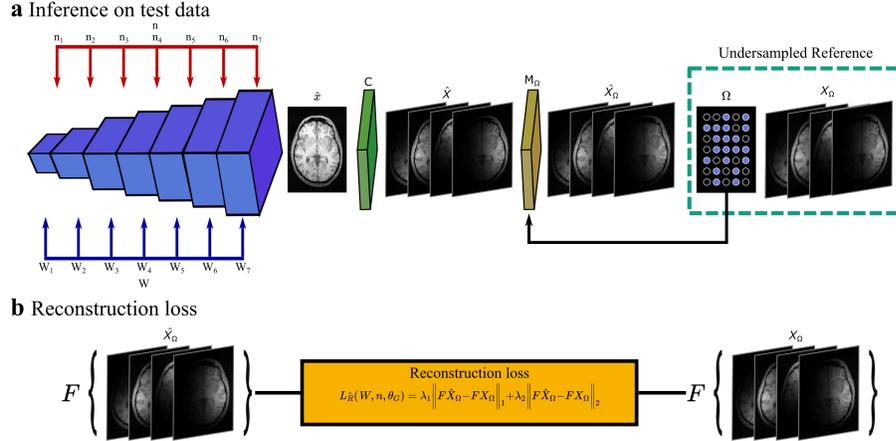


Figure 4.1: **(a)** The generative model in SLATER is adapted to perform zero-shot reconstruction for accelerated MRI. To do this, the generator output is back-projected onto individual coils (C), and masked with the same sampling pattern as in the undersampled acquisition (M_Ω). **(b)** The network prior consisting of noise (n), latent variables (W) and weights (θ_G) is then optimized to maximize consistency between the reconstructed and acquired k-space data.

modulated convolutions have a kernel size of 3×3 . The cross-attention transformer block contains a cascade of cross-attention and convolutional sub-blocks. The fully-connected sub-block in vanilla transformers is replaced with a convolutional sub-block to improve computational efficiency and permit use at high resolutions. Since the convolutional sub-block inherently focuses on local relationships, a second cross-attention block is used to reinforce long-range interactions. Note that the imaging matrix sizes differ between the IXI and fastMRI datasets analyzed here. The number of synthesizer layers was adjusted accordingly. The resolution of the final layer was set to 256×256 for IXI with a total of seven layers, whereas the resolution of the final layer was set to 512×512 for fastMRI with a total of eight layers. During unsupervised pretraining, MR images were zero-padded to the resolution of the final layer. During inference, the output of the synthesizer was cropped to match the matrix size of the MRI acquisition.

4.1.2 Mapper

The mapper comprises two streams of multi-layer architectures for processing local and global latent variables. The local stream has a total of five layers,

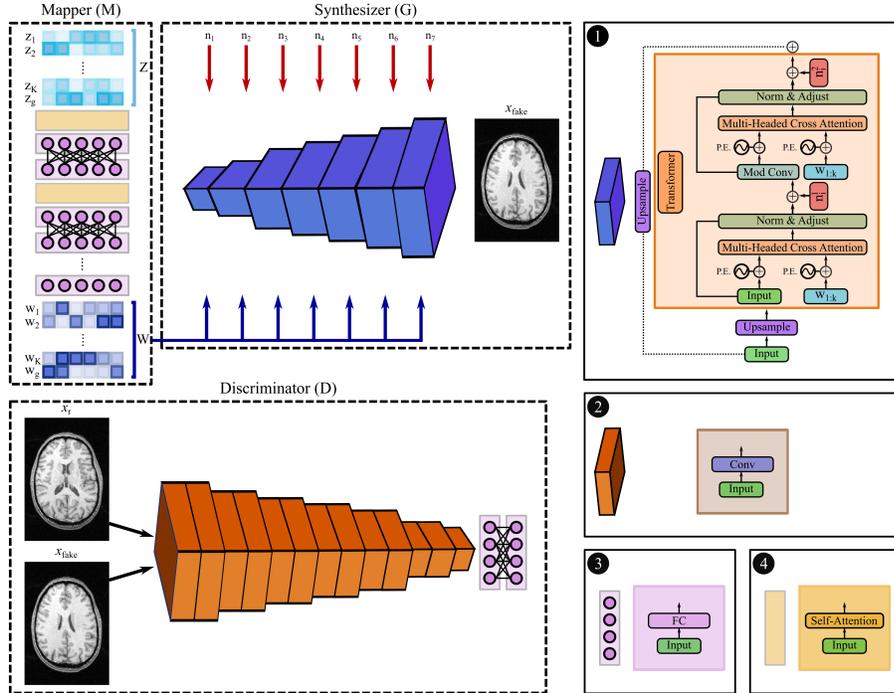


Figure 4.2: SLATER is based on an unconditional adversarial architecture with a synthesizer (G), a discriminator (D), and a mapper (M). G is a multi-layer architecture where image resolution is progressively increased across layers. Within intermediate layers, a convolutional upsampling block is followed by a cross-attention transformer block (see 1). The transformer block receives global and local latent variables, noise variables and positional encoding (P.E.) for latents and image features. D enables adversarial learning while it receives as input synthetic and actual MR images. It is composed of convolutional and fully-connected (FC) blocks (see 2, 3). M projects raw latent variables (Z) onto k local (w_1, \dots, w_K) variables and one global (w_g) latent variable. It is composed of self-attention and FC blocks (see 3, 4, and Supp. Fig. A.1 for details). An unsupervised generative modeling task is performed using the SLATER model to capture high-quality MRI priors.

with the first four containing a self-attention sub-block and the last containing a fully-connected sub-block. The global stream has a total of nine layers, each containing fully-connected sub-blocks.

4.1.3 Discriminator

The discriminator aggregates information across multiple spatial scales in a multi-layered architecture. Real and imaginary parts of synthesized and actual MR images are represented in separate channels. Each layer comprises a convolution block followed by downsampling by a factor of 2 and a skip connection to add the downsampled input. The downsampling block uses convolution with a kernel size of 3x3. The resolution of the first layer was set to 256x256 for IXI with a total of seven layers, and 512x512 for fastMRI with a total of eight layers.

4.2 Competing Methods

SLATER was comparatively demonstrated against state-of-the-art techniques based on supervised and unsupervised models, as well as a traditional method. For each technique, hyperparameter optimization was performed via cross-validation on a three-way split of subjects. Optimization was performed for number of epochs, number of inference iterations, and weights for regularization terms based on performance on the validation set. For supervised models, performance in the validation set was quantified as ℓ_2 -norm difference between the reconstructed and fully-sampled ground-truth images. For unsupervised models, validation performance was instead quantified as the ℓ_2 -norm difference between reconstructed and available k-space samples in undersampled acquisitions. A single set of hyperparameters yielding near-optimal results in all tasks were selected for each technique. Please see Supp. Fig. A.3 for sample performance curves in the validation set versus number of training epochs where training was continued up to 1500 for methods that perform generative modeling of MR images. During

inference on test data, strict data consistency was enforced to the network outputs. Codes were run on an eight-core Intel Xeon E5-2690v3 CPU for LORAKS, and in parallel on five nVidia 2080 Ti GPUs for all network models.

SLATER: SLATER was first pretrained to map random noise and latent codes onto high-quality MR images. In the IXI dataset, the model was trained to map onto single-coil magnitude images. In fastMRI dataset, the model was instead trained to map onto coil-combined complex images, with real and imaginary channel outputs. To do this, coil sensitivity maps were derived via ESPIRiT with default parameters [146]. Using these estimates, an optimal linear combination on multi-coil complex images was then performed [147]. Pre-training was performed via the Adam optimizer with $\beta_1:0.0$, $\beta_2:0.99$, $\eta = 10$ and a learning rate of 0.001 as adopted from [136]. The dimensionality of latent variables were set as $K=16$ and $L_s=32$. Network weights were randomly initialized using a standard normal distribution. Inference was performed via the RMSprop optimizer with learning rate 0.1, momentum parameter 0.9, early-stopping and learning rate schedule as adopted from [136]. Cross-validation indicated 1000 iterations as a favorable early-stopping point for maintaining near-optimal performance and computational efficiency. For IXI, 470th pretraining epoch, $\lambda_1 = 1.0, \lambda_2 = 0.0$ were selected. For fastMRI, 1280th pretraining epoch, $\lambda_1 = 0.5, \lambda_2 = 0.5$ were selected. Separate SLATER models were trained to build MRI priors for each tissue contrast within each dataset. Reconstruction was then performed via inference on individual test data conditioned by the MRI prior. The source code for SLATER can be found at: <https://github.com/icon-lab/SLATER>.

LORAKS: A traditional parallel-imaging reconstruction based on low-rank modeling of local k-space neighborhoods was performed [148] via libraries in the LORAKS V2.1 toolbox [149]. Here, an autocalibrated reconstruction was performed where the structured low-rank matrix was formed based on limited image support assumption [148]. Accordingly, the k-space neighborhood radius and the rank of the resultant matrix were selected via cross-validation as: (2,6) for IXI, and (2,30) for fastMRI.

GAN_{sup}: A fully-supervised conditional generative adversarial network

(GAN_{sup}) was trained using paired ground-truth and undersampled acquisitions. Network architecture and loss functions were adopted from [54]. Training was performed via the Adam optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$, dropout regularization rate 0.5, and a learning rate of 0.0002. Training was continued over 100 epochs, with learning rate schedule from [54]. Network weights were randomly initialized using a normal distribution with zero mean and 0.02 standard deviation. Regularization parameters for (pixel-wise, perceptual, adversarial) losses were selected as (100,100,1). A separate GAN_{sup} model was trained for each contrast within each dataset and acceleration rate.

SSDU: A self-supervised version of the conditional GAN model in GAN_{sup} was trained on undersampled data [24]. Acquired k-space samples were split into two sets of nonoverlapping points, where 60% of samples were used to estimate model weights and 40% were used to define the network loss. Analogous to Eq. 3.32, the network loss was taken as a weighted sum of ℓ_1 -, ℓ_2 -norm differences between recovered and acquired k-space samples. For both datasets, $\lambda_1 = 1.0$, $\lambda_2 = 0.0$ were selected. All other procedures were identical to GAN_{sup} . Separate SSDU models were trained for each contrast within each dataset and acceleration rate.

$\text{GAN}_{\text{prior}}$: Following [32], unsupervised pretraining on fully-sampled MRI data was performed using an unconditional GAN. Network architecture was adopted from [136] for fair comparison against SLATER. Training and inference procedures were identical to SLATER with minor modifications for enhanced performance. The synthesizer in $\text{GAN}_{\text{prior}}$ was trained to minimize the same loss as in Eq. 3.30 with an additional path length regularization parameter adopted from [136]. The discriminator was trained to minimize Eq. 3.31. A matching number of latents to SLATER were prescribed with $L_s=512$. For IXI, 60th pretraining epoch, $\lambda_1 = 1.0$, $\lambda_2 = 0.0$ were selected. For fastMRI, 180th pretraining epoch, $\lambda_1 = 1.0$, $\lambda_2 = 0.0$ were selected. Separate $\text{GAN}_{\text{prior}}$ models were trained for each contrast within each dataset.

SAGAN: Zero-shot learned reconstructions were also implemented using a self-attention GAN model. The network architecture was adopted from [136] for fair comparison. Training and inference procedures were identical to SLATER.

Optimization of network weights was not performed as it was observed to degrade reconstruction performance. Instead contrast-specific epoch selection was adopted for SAGAN since in this case it yielded enhanced performance. For T_1 and T_2 reconstructions, 798th, 399th epochs in IXI, and 967th, 2661th epoch in fastMRI, along with $\lambda_1 = 1.0, \lambda_2 = 0.0$ were selected. Separate models were trained for each contrast within each dataset.

GAN_{DIP}, SAGAN_{DIP}, SLATER_{DIP}: DIP reconstructions were performed via untrained GAN_{prior}, SAGAN and SLATER models respectively. The inference procedures were identical to pre-trained counterparts. Network weights were randomly initialized using a normal distribution with zero mean and unit standard deviation.

4.3 Datasets

Demonstrations were performed on single-coil brain MRI data from IXI (<http://brain-development.org/ixi-dataset/>) and multi-coil brain MRI data from fastMRI [150]. T_1 -weighted and T_2 -weighted acquisitions were considered. In IXI, 25 subjects were used for training, 5 for validation and 10 for testing. Parameters for T_1 -weighted scans are: repetition time (TR)=9.813 ms, echo time (TE)=4.603 ms, flip angle=8°, matrix size=256x256x150, voxel size=0.94x0.94x1.2 mm³; and those for T_2 -weighted scans are: TR=8178 ms, TE=100 ms, flip angle=90°, matrix size=256x256x130, voxel size=0.94x0.94x1.2 mm³. In fastMRI, 100 subjects were used for training, 10 for validation and 40 for testing. Data from multiple sites are included with no common protocol. For consistency, only volumes with at least 10 cross-sections and acquired with at least 5 coils were selected. To reduce computational complexity, GCC [151] was used to decrease the number of coils to 5. For both datasets, subject selection and splitting was done sequentially. Data were retrospectively undersampled using variable-density random patterns [3]. Undersampling masks were generated based on a 2D normal distribution with covariance adaptively adjusted to obtain the desired acceleration rates of R=[4, 8].

4.4 Quantitative Assessments

To assess reconstruction quality, quantitative comparisons were performed against reference images Fourier-reconstructed from fully-sampled acquisitions. Both reconstructed and reference images were normalized to a maximum of 1 prior to measurement. Peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) were calculated between the reconstructed and reference images. In Tables, summary statistics for quantitative metrics were provided as mean \pm standard deviation across test subjects. Statistical significance of differences between methods was assessed via nonparametric Wilcoxon signed-rank tests.

Chapter 5

Results

5.1 Cross-Attention Transformers

We first examined the utility of cross-attention transformers in capturing contextual representations via experiments on a simulated phantom. The phantom had numerical digits of unit intensity on four corners against a zero-intensity background (Fig. 3.2). This design creates a strong dependency among digit pixels in distant corners, and separately among background pixels. Normally-distributed white noise was added to vary the degree of contextual information present in phantom images by dampening inter-pixel correlations. DIP reconstructions at $R=4$ were then performed using SLATER. The spatial distribution of cross-attention maps in Eq. 3.5 characterize learned groupings of correlated pixels. Thus, we reasoned that the attention maps for local latent variables should span over distant albeit correlated image pixels for relatively limited noise levels, and the maps should degrade for higher noise levels due to weakened dependencies. To test this prediction, we inspected the attention maps for the phantom image in Fig. 3.2 (see Supp. Fig. A.4 for details). As expected, attention maps clearly span across digit pixels or across background pixels for limited noise. Note that attention maps learned on brain images also manifest similar grouping of correlated albeit spatially-distant pixels (see Fig. 3.1 and Supp. Fig. A.2 for

representative maps). Yet, towards substantially higher noise levels in the simulated phantom, attention maps show a less clear grouping structure as contextual information is gradually weakened. Taken together, these results suggest that SLATER can capture long-range spatial interactions among distant image pixels.

5.2 Model Invertability

During inference on test data, SLATER inverts its generative model to identify noise, latents and network weights that are most consistent with the undersampled MRI acquisition. To evaluate model invertability, we compared SLATER against CNN-based and self-attention GAN models. Each model was used in both DIP and zero-shot reconstructions. Representative results on T_1 - and T_2 -weighted acquisitions from IXI at R=4 are displayed in Fig. 5.1 and Supp. Fig. A.5 respectively. DIP reconstructions tend to suffer from visible loss of spatial resolution, and $\text{GAN}_{\text{prior}}$ and SAGAN have elevated noise and artifacts. In contrast, SLATER yields low residual errors and high visual acuity. Performance metrics are listed in Table 5.1. SLATER achieves superior performance against $\text{GAN}_{\text{prior}}$ and SAGAN in both DIP and zero-shot reconstructions ($p < 0.05$). Compared to the second-best method, SLATER yields 4.4dB higher PSNR and 7.7% higher SSIM in DIP, and 4.1dB higher PSNR and 5.9% higher SSIM in zero-shot reconstruction. Furthermore, SLATER yields 5.1dB higher PSNR and 7.0% higher SSIM over $\text{SLATER}_{\text{DIP}}$. These results indicate that the cross-attention transformer blocks in SLATER enhance model invertability compared to CNN architectures with or without self-attention and that the unsupervised pretraining stage in SLATER improves reconstruction performance.

5.3 Within-Domain Reconstructions

Next, we assessed within-domain reconstruction performance when the training and testing domains matched (e.g., T_1 reconstruction based on a T_1 -prior for SLATER). SLATER was compared against LORAKS, GAN_{sup} , SSDU, $\text{GAN}_{\text{prior}}$

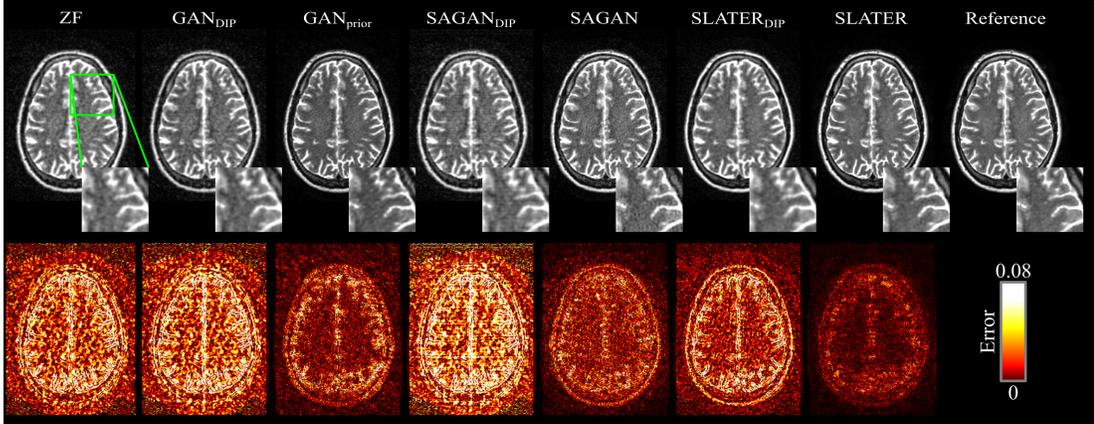


Figure 5.1: Reconstructions of a representative T_2 -weighted acquisition at $R=4$ are shown for the Fourier method (ZF), DIP methods (GAN_{DIP} , $\text{SAGAN}_{\text{DIP}}$, $\text{SLATER}_{\text{DIP}}$) and zero-shot reconstructions ($\text{GAN}_{\text{prior}}$, SAGAN, SLATER) along with the reference image. Zoom-in display windows are added to aid visualization of performance differences. Corresponding error maps are underneath the images for each method.

and SAGAN at $R=4$ and 8. Representative reconstructions are shown for IXI in Supp. Figs. A.6 and A.7, and for fastMRI in Fig. 5.2 and Supp. Fig. A.8. SLATER yields lower residual errors and higher acuity in depicting detailed tissue structure than competing methods. Quantitative assessments are listed in Supp. Table I and Table 5.2. SLATER achieves significantly enhanced reconstruction quality against all competing unsupervised methods ($p < 0.05$) offering 1.1dB higher PSNR and 1.1% higher SSIM compared to the second-best method. Furthermore, it offers 4.4dB higher PSNR and 5.5% higher SSIM compared to $\text{GAN}_{\text{prior}}$. It also yields higher performance than GAN_{sup} in all tasks ($p < 0.05$), except at $R=8$ in IXI where the two methods perform similarly and $R=8$ in fastMRI where GAN_{sup} yields higher SSIM. These results indicate that SLATER offers enhanced reconstruction quality for within-domain tasks compared to unsupervised baselines, while offering on par performance to a supervised baseline.

5.4 Across-Domain Reconstructions

We then evaluated across-domain reconstruction performance when the tissue contrast in the test domain was different than that in the domain of the trained

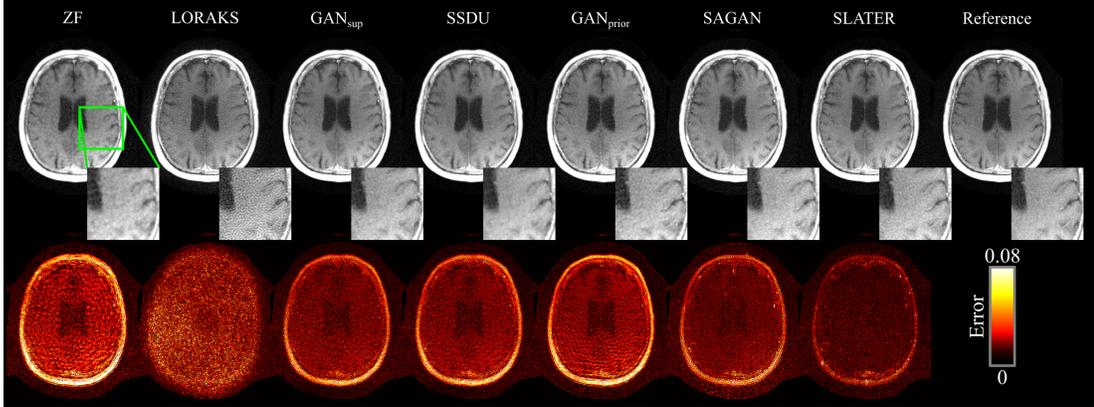


Figure 5.2: Within-domain reconstructions of a T_1 -weighted acquisition in fastMRI at $R=4$. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image, and error maps in the bottom row.

MRI prior (e.g., T_2 reconstructions based on a T_1 -prior). SLATER was compared against LORAKS, GAN_{sup} , SSDU, GAN_{prior} and SAGAN at $R=4$. Representative results are shown for IXI in Supp. Figs. A.9 and A.10, and for fastMRI in Fig. 5.3 and Supp. Fig. A.11. Quantitative assessments are listed in Table 5.3. Again, SLATER yields lower residual errors and higher acuity in depicting detailed tissue structure. It also achieves superior reconstruction quality against all competing supervised and unsupervised methods ($p < 0.05$), offering 2.5dB higher PSNR and 1.5% higher SSIM compared to the second-best method. Static models that are not adapted during inference such as GAN_{sup} and SSDU can yield suboptimal performance when the data distribution differs between the training and testing domains. Compared to within-domain reconstruction, we find that GAN_{sup} and SSDU suffer from 1.7dB PSNR, 1.7% SSIM loss on average in across-domain reconstruction. In contrast, adaptive models such as SLATER can attain more similar performance for across-domain and within-domain cases. We find that SLATER’s average performance differs less than 0.01dB PSNR and 0.11% SSIM between the two cases. Note that SLATER reconstructions based on within-domain priors generally yield on par or better performance than those based on across-domain priors, yet there are few exceptions with higher across-domain performance that might be attributed to empirical limitations in gradient-descent optimization (see Discussion).

We also assessed across-domain reconstructions when the acceleration rate of

the imaging operator differed between training and testing (i.e., trained at $R=8$ and tested at $R=4$). Performance measurements are listed in Supp. Table II. SLATER achieves superior reconstruction quality against all competing supervised and unsupervised methods ($p < 0.05$), offering 4.0dB PSNR, 1.5% SSIM improvement over the second-best method. Note that pre-training for zero-shot methods including SLATER is agnostic to the imaging operator, so these methods yield equivalent performance for within-domain and across-domain reconstructions. In contrast, GAN_{sup} and SSDU that are explicitly trained for a specific acceleration rate suffer from 3.3dB PSNR, 2.1% SSIM loss on average in across-domain reconstructions. Collectively, these results demonstrate that SLATER has improved generalization capabilities compared to static supervised and unsupervised models with fixed weights during inference, while still outperforming zero-shot reconstructions based on pure CNN architectures.

5.5 Ablation Experiments

We examined the contributions of individual parameter sets in SLATER that are optimized during inference. Variant models were built by progressively introducing optimization for latent variables, noise and network weights. Supp. Table III lists performance metrics for experiments on T_1 - and T_2 -weighted acquisitions in the IXI dataset at $R=4$. On average, the incurred performance gain in (PSNR, SSIM) is (4.1dB, 7.2%) with latent optimization, (6.5dB, 11.6%) with latent and noise optimization, and (10.8dB, 14.0%) with latent, noise and weight optimization. These results indicate the importance of each component in SLATER.

5.6 Weight Propagation

Average inference times for competing methods are listed in Table 5.4 (see Supp. Table IV for training times in the IXI dataset). Model adaptation to specific test samples in zero-shot reconstructions leads to prolonged inference. In principle, neighboring cross-sections with structural correlations should have similar

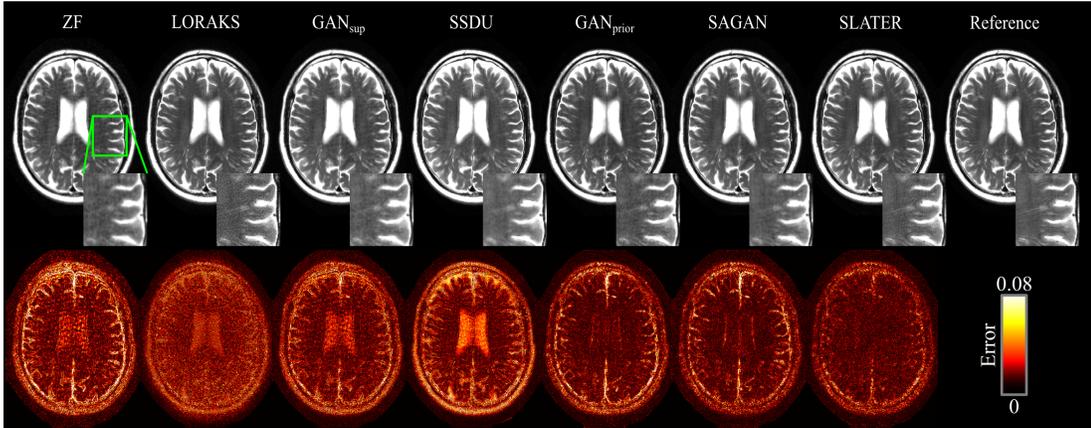


Figure 5.3: Across-domain reconstructions of a T_2 -weighted acquisition in fastMRI at $R=4$. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image, and error maps in the bottom row.

reconstructions. Thus we reasoned that propagating model parameters across consecutive cross-sections should increase efficiency by accelerating the progression onto high-quality reconstructions. Accordingly, the network weights for a given cross-section at the end of its inference optimization were stored, and then used to initialize inference optimization for the next cross-section within each subject. Latent variables and noise were not shared as they control context and fine details that could vary between cross-sections. Table 5.4 lists inference times with this weight propagation procedure (see Supp. Table V for reconstruction performance). While methods that do not perform model adaptation still provide faster inference, weight propagation substantially increases the inference efficiency for SLATER to improve its practicality.

Table 5.1: Reconstruction performance of DIP and zero-shot reconstructions for T_1 - and T_2 -weighted acquisitions IXI at R=4 and 8. Performance metrics are presented as mean \pm standard deviation across test subjects.

	GAN_{DIP}		SAGAN_{DIP}		SLATER_{DIP}		GAN_{prior}		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
T_1 , R=4	26.7 \pm 1.2	87.4 \pm 1.0	26.6 \pm 1.1	86.5 \pm 1.0	32.6 \pm 1.8	94.6 \pm 0.8	34.4 \pm 0.8	94.4 \pm 0.7	32.1 \pm 0.9	92.1 \pm 0.7	38.8 \pm 0.8	97.9 \pm 0.5
T_1 , R=8	23.5 \pm 0.9	83.0 \pm 1.1	23.4 \pm 0.9	82.5 \pm 1.2	30.3 \pm 1.7	91.6 \pm 1.4	29.3 \pm 1.2	89.7 \pm 1.4	28.6 \pm 0.9	88.3 \pm 1.2	33.2 \pm 0.9	95.2 \pm 0.9
T_2 , R=4	30.2 \pm 0.5	80.8 \pm 1.2	29.9 \pm 0.4	79.4 \pm 1.4	32.7 \pm 0.7	87.7 \pm 1.7	33.4 \pm 0.9	87.5 \pm 1.0	34.9 \pm 0.6	91.6 \pm 1.1	40.0 \pm 0.8	97.7 \pm 0.5
T_2 , R=8	27.6 \pm 0.4	76.2 \pm 1.3	27.3 \pm 0.4	74.1 \pm 1.6	29.9 \pm 0.7	84.0 \pm 1.9	31.2 \pm 0.7	85.3 \pm 1.0	30.7 \pm 0.5	86.4 \pm 1.4	34.1 \pm 0.8	94.8 \pm 0.7

Table 5.2: Within-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in fastMRI at R=4 and 8.

	LORAKS		GAN_{sup}		SSDU		GAN_{prior}		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
T_1 , R=4	33.4 \pm 2.7	82.2 \pm 7.7	37.2 \pm 2.6	94.4 \pm 5.7	37.2 \pm 3.3	94.2 \pm 7.5	32.8 \pm 2.0	92.5 \pm 5.2	36.1 \pm 2.6	94.1 \pm 5.1	37.6 \pm 3.2	93.9 \pm 9.5
T_1 , R=8	32.5 \pm 2.4	83.0 \pm 9.0	34.6 \pm 2.4	92.0 \pm 7.0	33.9 \pm 2.6	90.6 \pm 8.5	29.8 \pm 1.9	88.3 \pm 6.7	33.1 \pm 2.1	90.4 \pm 6.7	34.2 \pm 2.4	90.7 \pm 7.9
T_2 , R=4	34.3 \pm 1.0	90.8 \pm 1.6	35.4 \pm 1.2	95.5 \pm 0.5	33.0 \pm 2.9	94.6 \pm 1.3	33.5 \pm 1.1	91.5 \pm 1.8	33.5 \pm 1.3	94.1 \pm 0.8	36.3 \pm 1.2	95.5 \pm 0.7
T_2 , R=8	33.1 \pm 1.0	91.7 \pm 1.1	32.7 \pm 1.3	93.0 \pm 0.8	31.3 \pm 1.3	91.0 \pm 1.5	28.0 \pm 1.4	85.1 \pm 2.5	30.9 \pm 1.3	91.2 \pm 1.1	33.4 \pm 1.1	93.0 \pm 1.0

Table 5.3: Across-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in the IXI and fastMRI datasets at R=4. In A->B, A is the training domain and B is the test domain.

	LORAKS		GAN_{sup}		SSDU		GAN_{prior}		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
IXI $T_1 \rightarrow T_2$	35.4 \pm 0.5	92.3 \pm 1.2	36.8 \pm 0.7	95.3 \pm 0.6	34.8 \pm 1.0	90.3 \pm 1.4	36.4 \pm 0.7	93.2 \pm 0.6	32.3 \pm 0.7	86.1 \pm 1.3	39.9 \pm 0.8	97.7 \pm 0.5
IXI $T_2 \rightarrow T_1$	30.7 \pm 1.2	91.7 \pm 1.0	35.3 \pm 0.6	96.4 \pm 0.3	35.7 \pm 0.4	96.3 \pm 0.3	32.5 \pm 1.0	91.6 \pm 1.1	32.8 \pm 0.7	93.1 \pm 0.6	38.7 \pm 0.9	97.9 \pm 0.5
fastMRI $T_1 \rightarrow T_2$	34.3 \pm 1.0	90.8 \pm 1.6	33.8 \pm 1.2	94.1 \pm 0.7	34.8 \pm 1.1	95.1 \pm 0.7	33.5 \pm 1.5	94.3 \pm 0.8	34.5 \pm 1.2	94.4 \pm 0.9	36.2 \pm 1.1	94.6 \pm 1.0
fastMRI $T_2 \rightarrow T_1$	33.4 \pm 2.7	82.2 \pm 7.7	36.9 \pm 2.4	93.6 \pm 5.8	34.3 \pm 3.3	92.9 \pm 9.2	35.9 \pm 2.1	90.8 \pm 5.4	34.4 \pm 2.1	92.2 \pm 4.9	38.0 \pm 2.6	95.3 \pm 6.2

Table 5.4: Average inference times in sec per cross section. Run times without and with weight propagation (WP) are listed for zero-shot reconstructions. SAGAN does not perform weight optimization, so it is unaffected by WP.

	LORAKS	GAN_{sup}	SSDU	GAN_{prior}	SAGAN	SLATER
No WP	4.10	0.002	0.002	12.73	7.74	14.80
With WP	–	–	–	2.52	7.74	2.63

Chapter 6

Conclusion

This thesis proposes zero-shot learned adversarial transformers for unsupervised reconstruction in accelerated MRI. To our knowledge, this is the first study to introduce a transformer network for MRI reconstruction. Traditional GANs contain cascades of convolutional layers, which might limit ability to capture long-range spatial dependencies. We instead employed cross-attention transformer blocks to efficiently capture contextual image features. Note that self-attention among all feature map locations leads to excessive computational burden [36]. As such, self-attention modules following convolutional blocks have only been leveraged in layers with modest resolution to prevent quadratic complexity [35, 152–155]. In contrast, the cross-attention mechanism between low-dimensional latent variables and image features permits use at higher resolutions [137].

Our demonstrations clearly indicate the superiority of the proposed method over a fully-supervised GAN model (GAN_{sup}), a self-supervised model (SSDU), DIP based on CNN and self-attention GANs (GAN_{DIP} , $\text{SAGAN}_{\text{DIP}}$), and zero-shot reconstructions ($\text{GAN}_{\text{prior}}$, SAGAN). SLATER offers on par performance with GAN_{sup} for within-domain tasks, while it outperforms GAN_{sup} for across-domain tasks. SLATER also outperforms DIP by a substantial margin, which can be attributed to its MRI prior. Lastly, SLATER yields superior performance to $\text{GAN}_{\text{prior}}$ that indicates the importance of transformer blocks in learning a high-fidelity prior.

Recent studies aiming to reduce supervision requirements related to raw data have proposed unpaired or unsupervised learning strategies for MRI reconstruction. A successful approach is to train models on unpaired samples of under-sampled and fully-sampled acquisitions [12, 18, 60]. Another promising approach is to train models via self-supervision on undersampled data [21, 23–25]. Both approaches perform model training in conjunction with the imaging operator to map undersampled acquisitions to MR images. Trained model weights that reflect an indirect prior to reduce artifacts are then fixed during inference. Thus, reconstructions based on an across-domain prior can suffer from suboptimal generalization due to inconsistencies in the data distribution (e.g., tissue contrast) or the imaging operator (e.g., acceleration rate) between training and testing domains.

SLATER instead learns an MRI prior agnostic to the imaging operator that is later adapted to test data during inference. Prior adaptation involves an inference optimization with considerable computational burden compared to static models with fixed priors. Yet, our experiments indicate that it mitigates potential performance losses during across-domain reconstructions. Ideally, a within-domain prior would be expected to initialize the inference optimization closer to a desirable local minimum, resulting in on par or higher performance than an across-domain prior. Given complex loss surfaces for deep models, however, gradient-descent optimizers can occasionally show non-monotonic behavior approaching less preferred minima despite favorable initialization [156]. The few cases where an across-domain prior yielded relatively higher performance are likely attributed to such behavior, since we observed that the performance gap between the priors diminished with substantially prolonged optimizations in unreported experiments. That said, we opted for early stopping during inference to maintain a desirable compromise between performance and inference time while also mitigating potential risks for over-fitting [30, 135].

The inference optimization performed during zero-shot reconstruction is closely related to DIP methods [135]. Compared to DIP with randomly initialized networks, here we show that SLATER yields enhanced reconstruction quality via a pre-trained prior that is more strongly tuned towards the distribution of MR

images. Yet, it remains an important research topic to examine the convergence behavior of both untrained and pre-trained models. Several recent studies on accelerated imaging report theoretical and empirical validation of convergence for iterative optimization with traditional recovery methods [157] as well as untrained generative priors [158, 159]. Further theoretical and empirical research is warranted to investigate whether these results generalize to pre-trained priors.

Among recent efforts, the closest to the presented approach are prior-adaptation methods in [33] that uses a variational auto-encoder to learn patch priors and in [32] that uses a GAN to learn image priors. Our work differs from these efforts in the following aspects: (i) Compared to [33] that uses a patch-based implementation, we leverage an image-based implementation that can improve performance in leveraging non-local context. (ii) [33] uses a fixed patch prior during inference, whereas we perform test-sample specific adaptation of the prior. (iii) Compared to [32] that uses a CNN-based GAN, here we use cross-attention transformers to better capture long-range spatial dependencies. (iv) SLATER includes noise and latents at each synthesizer layer to better control image features.

Flow-based models have also been proposed to learn priors for inverse problems [160]. These models use a composition of fully-invertible flow steps that transform latent variables onto data samples, and they offer exact estimates of the maximum likelihood of data samples and improved immunity against representation errors due to bias. In comparison, adversarial models such as SLATER implicitly minimize discrepancy between generated and actual data distributions without explicit estimation of probability densities. Note, however, that flow-based models characteristically require high-dimensional latent spaces to retain a target level of expressiveness, which might result in a less favorable trade-off between computational burden and quality of generated samples [161]. It remains important future work to compare flow-based and adversarial methods as well as their hybrids in the context of MRI reconstruction.

Several lines of development can be pursued for the proposed technique. First,

zero-shot reconstructions can be initialized with latent estimates based on conventional parallel imaging/compressed sensing reconstructions. This can increase computational efficiency by shortening the inference procedure. Second, the reconstruction loss in SLATER can be combined with SURE-type estimates or regularization terms on network weights [27]. Here we did not explicitly regularize the weights, latents or noise to allow higher degree of consistency to acquired data. When desired, network regularization and on-line error estimates can be introduced to reduce potential for overfitting. Third, SLATER learns a coil-combined MRI prior and subsequently incorporates coil-sensitivity information during zero-shot reconstruction. This is achieved by back-projecting the synthesizer output onto individual coils. A powerful alternative is to build a synthesizer with a consistent prior across coils to generate multi-coil images [31]. Similar computational efficiency might be expected from both approaches that use a common MRI prior across coils. That said, effects of the decoupling strategy in SLATER against the channel-consistence prior on reconstruction quality remain to be investigated.

SLATER pre-training performs unsupervised generative modeling of coil-combined MR images derived from fully-sampled acquisitions. Combined with the imaging operator during inference, the learned prior is then used to perform reconstruction of undersampled acquisitions via unsupervised model adaptation. This decoupled approach bypasses the need for paired training datasets. Moreover, the MRI prior is agnostic to the imaging operator and flexibly adapted to the test domain. Thus, a learned prior can be used to reconstruct undersampled acquisitions at varying contrasts or acceleration rates. To lower reliance on fully-sampled datasets, pre-training can be instead performed on undersampled acquisitions. While the resultant prior will not entirely reflect the distribution of high-quality MR images, the model adaptation procedures in SLATER might limit potential performance losses.

Conventional supervised models perform poorly when training data are scarce, and domain shifts between the training and test sets are prominent. Recent studies have proposed supervised models of low complexity trained with limited or synthetic data to partly mitigate losses in generalization performance [162, 163].

While compact models are less prone to over-fitting, data-driven priors learned from limited or synthetic data can still elicit a biased representation of MRI data, increasing risk of inaccurate reconstruction for atypical anatomy [15,16]. Future studies are warranted to assess the utility of supervised compact versus unsupervised models in domain generalization. Supervised and unsupervised procedures can be combined to facilitate model training on larger datasets by mixing paired and unpaired data. Unsupervised adaptation of such semi-supervised models during inference might further enhance reliability against atypical anatomy.

In summary, this thesis introduced a novel unsupervised MRI reconstruction based on an unconditional deep adversarial network. SLATER leverages cross-attention transformers to improve capture of contextual image features. Benefits of SLATER over state-of-the-art supervised and unsupervised methods were demonstrated in brain MRI. SLATER can also be adopted for structural and dynamic MRI in other anatomies, or other imaging modalities [164–216]. Moreover, cross-attention transformers in SLATER may also be incorporated with other deep models designed for medical image synthesis, segmentation or analysis to improve contextual sensitivity [217–230]. Reduced supervision requirements and subject-specific adaptation render SLATER a promising candidate for high-performance accelerated MRI.

Bibliography

- [1] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, “SENSE: sensitivity encoding for fast MRI,” *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–62, 1999.
- [2] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, “Generalized autocalibrating partially parallel acquisitions (GRAPPA),” *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [3] M. Lustig, D. Donoho, and J. M. Pauly, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [4] S. Ma, W. Yin, Y. Zhang, and A. Chakraborty, “An efficient algorithm for compressed mr imaging using total variation and wavelets,” in *IEEE CVPR*, pp. 1–8, 2008.
- [5] C. M. Sandino, J. Y. Cheng, F. Chen, M. Mardani, J. M. Pauly, and S. S. Vasanawala, “Compressed sensing: From research to clinical practice with deep neural networks: Shortening scan times for Magnetic Resonance Imaging,” *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 117–127, 2020.
- [6] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang, “Accelerating magnetic resonance imaging via deep learning,” in *IEEE ISBI*, pp. 514–517, 2016.

- [7] Y. Yang, J. Sun, H. Li, and Z. Xu, “ADMM-CSNet: A deep learning approach for image compressive sensing,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 3, pp. 521–538, 2020.
- [8] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, “Learning a variational network for reconstruction of accelerated MRI data,” *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2017.
- [9] M. Mardani, E. Gong, J. Y. Cheng, S. Vasanawala, G. Zaharchuk, L. Xing, and J. M. Pauly, “Deep generative adversarial neural networks for compressive sensing MRI,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 1, pp. 167–179, 2019.
- [10] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert, “A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction,” in *Proceedings of IPMI*, pp. 647–658, 2017.
- [11] H. K. Aggarwal, M. P. Mani, and M. Jacob, “MoDL: Model-Based deep learning architecture for inverse problems,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
- [12] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, “Compressed sensing MRI reconstruction with cyclic loss in generative adversarial networks,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1488–1497, 2018.
- [13] A. Sriram, J. Zbontar, T. Murrell, A. Defazio, C. L. Zitnick, N. Yakubova, F. Knoll, and P. Johnson, “End-to-end variational networks for accelerated MRI reconstruction,” in *Proceedings of MICCAI*, pp. 64–73, 2020.
- [14] S. Biswas, H. K. Aggarwal, and M. Jacob, “Dynamic MRI using model-based deep learning and SToRM priors: MoDL-SToRM,” *Magnetic resonance in medicine*, vol. 82, no. 1, pp. 485–494, 2019.
- [15] F. Knoll, K. Hammernik, E. Kobler, T. Pock, M. P. Recht, and D. K. Sodickson, “Assessment of the generalization of learned image reconstruction and the potential for transfer learning,” *Magnetic Resonance in Medicine*, vol. 81, no. 1, pp. 116–128, 2019.

- [16] S. U. H. Dar, M. Özbey, A. B. Çatlı, and T. Çukur, “A transfer-learning approach for accelerated MRI using deep neural networks,” *Magnetic Resonance in Medicine*, vol. 84, no. 2, pp. 663–685, 2020.
- [17] B. Zhu, J. Z. Liu, B. R. Rosen, and M. S. Rosen, “Image reconstruction by domain transform manifold learning,” *Nature*, vol. 555, no. 7697, pp. 487–492, 2018.
- [18] G. Oh, B. Sim, H. Chung, L. Sunwoo, and J. C. Ye, “Unpaired deep learning for accelerated MRI using optimal transport driven cycleGAN,” *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1285–1296, 2020.
- [19] D.-i. Eun, R. Jang, W. S. Ha, H. Lee, S. C. Jung, and N. Kim, “Deep-learning-based image quality enhancement of compressed sensing magnetic resonance imaging of vessel wall: comparison of self-supervised and unsupervised approaches,” *Scientific Reports*, vol. 10, no. 1, p. 13950, 2020.
- [20] H. Chung, E. Cha, L. Sunwoo, and J. C. Ye, “Two-stage deep learning for accelerated 3D time-of-flight MRA without matched training data,” *Medical Image Analysis*, vol. 71, p. 102047, 2021.
- [21] J. I. Tamir, S. X. Yu, and M. Lustig, “Unsupervised deep basis pursuit: Learning reconstruction without ground-truth data,” in *Proceedings of ISMRM*, p. 0660, 2019.
- [22] A. Q. Wang, A. V. Dalca, and M. R. Sabuncu, “Neural network-based reconstruction in compressed sensing MRI without fully-sampled training data,” in *MLMIR*, pp. 27–37, 2020.
- [23] E. K. Cole, J. M. Pauly, S. S. Vasanawala, and F. Ong, “Unsupervised MRI reconstruction with generative adversarial networks,” *arXiv:2008.13065*, 2020.
- [24] B. Yaman, S. A. H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, and M. Akçakaya, “Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data,” *Magnetic resonance in medicine*, vol. 84, no. 6, pp. 3172–3191, 2020.

- [25] P. Huang, C. H. Li, S. K. Gaire, R. Liu, X. Zhang, X. Li, and L. Ying, “Deep MRI reconstruction without ground truth for training,” in *Proceedings of ISMRM*, p. 4668, 2019.
- [26] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, “RARE: Image reconstruction using deep priors learned without groundtruth,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1088–1099, 2020.
- [27] H. K. Aggarwal and M. Jacob, “ENSURE: Ensemble Stein’s unbiased risk estimator for unsupervised learning,” *arXiv:2010.10631*, 2020.
- [28] K. H. Jin, H. Gupta, J. Yerly, M. Stuber, and M. Unser, “Time-dependent deep image prior for dynamic MRI,” *arXiv:1910.01684*, 2019.
- [29] S. Arora, V. Roeloffs, and M. Lustig, “Untrained modified deep decoder for joint denoising and parallel imaging reconstruction,” in *Proceedings of ISMRM*, p. 3585, 2020.
- [30] Q. Zou, A. H. Ahmed, P. Nagpal, S. Kruger, and M. Jacob, “Deep generative SToRM model for dynamic imaging,” *arXiv:2101.12366*, 2021.
- [31] M. Z. Darestani and R. Heckel, “Accelerated mri with un-trained neural networks,” *IEEE Transactions on Computational Imaging*, vol. 7, pp. 724–733, 2021.
- [32] D. Narnhofer, K. Hammernik, F. Knoll, and T. Pock, “Inverse GANs for accelerated MRI reconstruction,” in *Proceedings of SPIE*, vol. 11138, pp. 381–392, 2019.
- [33] K. C. Tezcan, C. F. Baumgartner, R. Luechinger, K. P. Pruessmann, and E. Konukoglu, “MR image reconstruction using deep density priors,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 7, pp. 1633–1642, 2019.
- [34] Q. Liu, Q. Yang, H. Cheng, S. Wang, M. Zhang, and D. Liang, “Highly undersampled magnetic resonance imaging reconstruction using autoencoding priors,” *Magnetic Resonance in Medicine*, vol. 83, no. 1, pp. 322–336, 2020.

- [35] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, “Self-attention generative adversarial networks,” in *Proceedings of ICML*, pp. 7354–7363, 2019.
- [36] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, and Y. Zhou, “TransUNet: Transformers make strong encoders for medical image segmentation,” *arXiv:2102.04306*, 2021.
- [37] O. Dalmaz, M. Yurt, and T. Çukur, “ResViT: Residual vision transformers for multi-modal medical image synthesis,” *arXiv:2106.16031*, 2021.
- [38] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo, “Deep learning for undersampled MRI reconstruction,” *Physics in Medicine and Biology*, vol. 63, no. 13, p. 135007, 2018.
- [39] J. Yoon, E. Gong, I. Chatnuntaweck, B. Bilgic, J. Lee, W. Jung, J. Ko, H. Jung, K. Setsompop, G. Zaharchuk, E. Y. Kim, J. Pauly, and J. Lee, “Quantitative susceptibility mapping using deep neural network: QSMnet,” *NeuroImage*, vol. 179, pp. 199–206, 2018.
- [40] J. C. Ye, Y. Han, and E. Cha, “Deep convolutional framelets: A general deep learning framework for inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 11, no. 2, pp. 991–1048, 2018.
- [41] D. Lee, J. Yoo, S. Tak, and J. C. Ye, “Deep residual learning for accelerated MRI using magnitude and phase networks,” *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 9, pp. 1985–1995, 2018.
- [42] A. Hauptmann, S. Arridge, F. Lucka, V. Muthurangu, and J. A. Steeden, “Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning—a proof of concept in congenital heart disease,” *Magnetic Resonance in Medicine*, vol. 81, no. 2, pp. 1143–1156, 2019.
- [43] K. Kwon, D. Kim, and H. Park, “A parallel MR imaging method using multilayer perceptron,” *Medical Physics*, vol. 44, no. 12, pp. 6209–6224, 2017.
- [44] Y. Yang, J. Sun, H. Li, and Z. Xu, “Deep ADMM-Net for compressive sensing MRI,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016.

- [45] T. Eo, Y. Jun, T. Kim, J. Jang, H.-J. Lee, and D. Hwang, “KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images,” *Magnetic Resonance in Medicine*, vol. 80, no. 5, pp. 2188–2201, 2018.
- [46] J. Cheng, H. Wang, L. Ying, and D. Liang, “Model learning: Primal dual networks for fast MR imaging,” in *Proceedings of MICCAI*, pp. 21–29, 2019.
- [47] J. Adler and O. Oktem, “Learned primal-dual reconstruction,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1322–1332, 2018.
- [48] S. Wang, Z. Ke, H. Cheng, S. Jia, L. Ying, H. Zheng, and D. Liang, “DIMENSION: Dynamic MR imaging with both k-space and spatial prior knowledge obtained via multi-supervised network training,” *NMR in Biomedicine*, p. e4131, 2019.
- [49] C. Qin, J. Schlemper, J. Caballero, A. N. Price, J. V. Hajnal, and D. Rueckert, “Convolutional recurrent neural networks for dynamic MR image reconstruction,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 1, pp. 280–290, 2019.
- [50] S. A. H. Hosseini, B. Yaman, S. Moeller, M. Hong, and M. Akcakaya, “Dense recurrent neural networks for accelerated MRI: History-cognizant unrolling of optimization algorithms,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1280–1291, 2020.
- [51] Y. Chen, D. Firmin, and G. Yang, “Wavelet improved GAN for MRI reconstruction,” in *Proceedings of SPIE*, vol. 11595, p. 1159513, 2021.
- [52] S. Yu, H. Dong, G. Yang, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, D. Firmin, and Y. Guo, “DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1310–1321, 2018.
- [53] M. Mardani, H. Monajemi, V. Pappayan, S. Vasanawala, D. Donoho, and J. Pauly, “Recurrent generative adversarial networks for proximal learning and automated compressive image recovery,” *arXiv:1711.10046*, 2017.

- [54] S. U. Dar, M. Yurt, M. Shahdloo, M. E. Ildız, B. Tınaz, and T. Çukur, “Prior-guided image reconstruction for accelerated multi-contrast MRI via generative adversarial networks,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1072–1087, 2020.
- [55] 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. 84, no. 3, pp. 1456–1469, 2020.
- [56] Q. Chang, H. Qu, Y. Zhang, M. Sabuncu, C. Chen, T. Zhang, and D. N. Metaxas, “Synthetic learning: Learn from distributed asynchronized discriminator GAN without sharing medical image data,” in *IEEE CVPR*, pp. 13853–13863, 2020.
- [57] Y. Han, J. Yoo, H. H. Kim, H. J. Shin, K. Sung, and J. C. Ye, “Deep learning with domain adaptation for accelerated projection-reconstruction MR,” *Magnetic Resonance in Medicine*, vol. 80, no. 3, pp. 1189–1205, 2018.
- [58] S. U. Dar, M. Yurt, and T. Çukur, “A few-shot learning approach for accelerated MRI via fusion of data-driven and subject-driven priors,” in *Proceedings of ISMRM*, p. 1949, 2021.
- [59] Y. Arefeen, O. Beker, H. Yu, E. Adalsteinsson, and B. Bilgic, “Scan-specific, parameter-free artifact reduction in k-space (SPARK),” *arXiv:2104.01188*, 2021.
- [60] K. Lei, M. Mardani, J. M. Pauly, and S. S. Vasanawala, “Wasserstein gans for mr imaging: From paired to unpaired training,” *IEEE Transactions on Medical Imaging*, vol. 40, no. 1, pp. 105–115, 2021.
- [61] M. Akcakaya, S. Moeller, S. Weingartner, and K. Ugurbil, “Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging,” *Magnetic Resonance in Medicine*, vol. 81, no. 1, pp. 439–453, 2019.
- [62] S. A. H. Hosseini, C. Zhang, S. Weingartner, S. Moeller, M. Stuber, K. Ugurbil, and M. Akcakaya, “Accelerated coronary MRI with sRAKI:

- A database-free self-consistent neural network k-space reconstruction for arbitrary undersampling,” *PLOS ONE*, vol. 15, no. 2, pp. 1–13, 2020.
- [63] T. H. Kim, P. Garg, and J. P. Haldar, “LORAKI: Autocalibrated recurrent neural networks for autoregressive MRI reconstruction in k-space,” *arXiv:1904.09390*, 2019.
- [64] P. Huang, C. Zhang, H. Li, S. K. Gaire, R. Liu, X. Zhang, X. Li, L. Dong, and L. Ying, “Unsupervised deep learning reconstruction using the MR imaging model,” in *Proceedings of ISMRM*, p. 3617, 2020.
- [65] C. Eldeniz, W. Gan, S. Chen, J. Liu, U. S. Kamilov, and H. An, “Phase2Phase: Reconstruction of free-breathing MRI into multiple respiratory phases using deep learning without a ground truth,” in *Proceedings of ISMRM*, p. 0807, 2020.
- [66] J. Liu, C. Eldeniz, W. Sun, Yu Gan, S. Chen, H. An, and U. S. Kamilov, “RED-N2N: Image reconstruction for MRI using deep CNN priors trained without ground truth,” in *Proceedings of ISMRM*, p. 0993, 2020.
- [67] S. A. Hossein Hosseini, B. Yaman, S. Moeller, and M. Akcakaya, “High-fidelity accelerated MRI reconstruction by scan-specific fine-tuning of physics-based neural networks,” in *IEEE EMBC*, pp. 1481–1484, 2020.
- [68] H. K. Aggarwal and M. Jacob, “Model adaptation for image reconstruction using generalized Stein’s unbiased risk estimator,” *arXiv:2102.00047*, 2021.
- [69] Z. Ke, Y. Zhu, J. Cheng, L. Ying, X. Liu, H. Zheng, and D. Liang, “Assessment of the generalization of learned unsupervised deep learning method,” in *Proceedings of ISMRM*, p. 3630, 2020.
- [70] Y. Korkmaz, S. U. Dar, M. Yurt, M. Özbey, and T. Çukur, “Unsupervised MRI reconstruction via zero-shot learned adversarial transformers,” *arXiv:2105.08059*, 2021.
- [71] Y. Korkmaz, S. U. Dar, M. Yurt, M. Ozbey, and T. Çukur, “Zero-shot learning for unsupervised reconstruction of accelerated MRI acquisitions,” in *Proceedings of ISMRM*, 2021.

- [72] Y. Korkmaz, M. Yurt, S. U. H. Dar, M. Özbey, and T. Cukur, “Deep mri reconstruction with generative vision transformers,” in *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 54–64, Springer, 2021.
- [73] O. B. Demirel, T. Kilic, T. Çukur, and E. U. Saritas, “Anatomical measurements correlate with individual magnetostimulation thresholds for khz-range homogeneous magnetic fields,” *Medical physics*, vol. 47, no. 4, pp. 1836–1844, 2020.
- [74] O. Demirel, T. Kilic, T. Çukur, and E. Saritas, “Simple anatomical measures correlate with individual pns thresholds for khz-range homogeneous magnetic fields,” in *Proceedings of the 28th Annual Meeting of ISMRM*, 2020.
- [75] M. Shahdloo, E. Ilicak, M. Tofghi, E. U. Saritas, A. E. Çetin, and T. Çukur, “Projection onto epigraph sets for rapid self-tuning compressed sensing mri,” *IEEE transactions on medical imaging*, vol. 38, no. 7, pp. 1677–1689, 2018.
- [76] T. Çukur, J. M. Santos, J. M. Pauly, and D. G. Nishimura, “Variable-density parallel imaging with partially localized coil sensitivities,” *IEEE transactions on medical imaging*, vol. 29, no. 5, pp. 1173–1181, 2010.
- [77] E. Biyik, E. Ilicak, and T. Cukur, “Reconstruction by calibration over tensors for multi-coil multi-acquisition balanced ssfp imaging,” *Magnetic resonance in medicine*, vol. 79, no. 5, pp. 2542–2554, 2018.
- [78] E. Ilicak, L. K. Senel, E. Biyik, and T. Çukur, “Profile-encoding reconstruction for multiple-acquisition balanced steady-state free precession imaging,” *Magnetic resonance in medicine*, vol. 78, no. 4, pp. 1316–1329, 2017.
- [79] E. Kopanoglu, A. Güngör, T. Kilic, E. U. Saritas, K. K. Oguz, T. Çukur, and H. E. Güven, “Simultaneous use of individual and joint regularization terms in compressive sensing: Joint reconstruction of multi-channel multi-contrast mri acquisitions,” *NMR in Biomedicine*, vol. 33, no. 4, p. e4247, 2020.

- [80] L. K. Senel, T. Kilic, A. Gungor, E. Kopanoglu, H. E. Guven, E. U. Saritas, A. Koc, and T. Cukur, “Statistically segregated k-space sampling for accelerating multiple-acquisition mri,” *IEEE transactions on medical imaging*, vol. 38, no. 7, pp. 1701–1714, 2019.
- [81] E. Ilicak, S. Cetin, E. Bulut, K. K. Oguz, E. U. Saritas, G. Unal, and T. Çukur, “Targeted vessel reconstruction in non-contrast-enhanced steady-state free precession angiography,” *NMR in Biomedicine*, vol. 29, no. 5, pp. 532–544, 2016.
- [82] T. Çukur, “Accelerated phase-cycled ssfp imaging with compressed sensing,” *IEEE transactions on medical imaging*, vol. 34, no. 1, pp. 107–115, 2014.
- [83] T. Çukur, M. Lustig, E. U. Saritas, and D. G. Nishimura, “Signal compensation and compressed sensing for magnetization-prepared mr angiography,” *IEEE transactions on medical imaging*, vol. 30, no. 5, pp. 1017–1027, 2011.
- [84] D. Holland, C. Liu, X. Song, E. Mazerolle, M. Stevens, A. Sederman, L. Gladden, R. D’Arcy, C. Bowen, and S. Beyea, “Compressed sensing reconstruction improves sensitivity of variable density spiral fmri,” *Magnetic Resonance in Medicine*, vol. 70, no. 6, pp. 1634–1643, 2013.
- [85] T. Cukur, M. Lustig, and D. G. Nishimura, “Improving non-contrast-enhanced steady-state free precession angiography with compressed sensing,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 61, no. 5, pp. 1122–1131, 2009.
- [86] T. Çukur, J. H. Lee, N. K. Bangerter, B. A. Hargreaves, and D. G. Nishimura, “Non-contrast-enhanced flow-independent peripheral mr angiography with balanced ssfp,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 61, no. 6, pp. 1533–1539, 2009.
- [87] T. Çukur and D. G. Nishimura, “Multiple repetition time balanced steady-state free precession imaging,” *Magnetic Resonance in Medicine: An*

Official Journal of the International Society for Magnetic Resonance in Medicine, vol. 62, no. 1, pp. 193–204, 2009.

- [88] T. Çukur, M. Lustig, and D. G. Nishimura, “Multiple-profile homogeneous image combination: application to phase-cycled ssfp and multicoil imaging,” *Magnetic resonance in medicine*, vol. 60, no. 3, pp. 732–738, 2008.
- [89] T. Çukur, N. K. Bangerter, and D. G. Nishimura, “Enhanced spectral shaping in steady-state free precession imaging,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1216–1223, 2007.
- [90] B. Gözcü, R. K. Mahabadi, Y.-H. Li, E. Ilıcak, T. Cukur, J. Scarlett, and V. Cevher, “Learning-based compressive mri,” *IEEE transactions on medical imaging*, vol. 37, no. 6, pp. 1394–1406, 2018.
- [91] A. Güngör, E. Kopanoğlu, T. Çukur, and H. E. Güven, “Compressed multi-contrast magnetic resonance image reconstruction using augmented lagrangian method,” in *2016 24th Signal Processing and Communication Application Conference (SIU)*, pp. 1985–1988, IEEE, 2016.
- [92] A. Güngör, E. Kopanoğlu, T. Cukur, and H. E. Güven, “A synthesis-based approach to compressive multi-contrast magnetic resonance imaging,” in *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pp. 696–699, IEEE, 2017.
- [93] A. Koc, H. M. Ozaktas, B. Bartan, E. Gundogdu, and T. Cukur, “Digital computation of fractional fourier and linear canonical transforms and sparse image representation,” in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 111–117, IEEE, 2017.
- [94] P. Aggarwal and A. Gupta, “Accelerated fmri reconstruction using matrix completion with sparse recovery via split bregman,” *Neurocomputing*, vol. 216, pp. 319–330, 2016.

- [95] T. Kılıç, T. Çukur, O. Algin, and E. Ü. Saritaş, “Joint partial fourier and compressed sensing reconstruction for accelerated time-of-flight mr angiography,” in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2018.
- [96] T. Kilic, E. Ilicak, T. Çukur, and E. Saritas, “Improved spirit operator for joint reconstruction of multiple t2-weighted images,” in *Proc. 25th Annu. Meeting ISMRM*, 2017.
- [97] E. Kopanoglu, A. Güngör, T. Kilic, E. U. Saritas, T. Çukur, and H. E. Guven, “Joint reconstruction of multi-contrast images: compressive sensing reconstruction using both joint and individual regularization functions,” 2017.
- [98] M. Aksoy, C. Forman, M. Straka, T. Cukur, S. Holdsworth, S. Skare, J. Santos, J. Hornegger, and R. Bammer, “Hybrid prospective & retrospective head motion correction system to mitigate cross-calibration errors,” in *Proceedings of the 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine. Stockholm*, vol. 499, 2010.
- [99] E. Ilicak and T. Çukur, “Parameter-free profile encoding reconstruction for multiple-acquisition bssfp imaging,” in *Proc. 25th Annu. Meeting ISMRM*, 2017.
- [100] A. Güngör, E. Kopanoğlu, T. Çukur, and H. E. Güven, “Fast recovery of compressed multi-contrast magnetic resonance images,” in *Medical Imaging 2017: Image Processing*, vol. 10133, pp. 503–508, SPIE, 2017.
- [101] T. Çukur, “Spectrally selective imaging with wideband balanced steady-state free precession mri,” *Magnetic Resonance in Medicine*, vol. 75, no. 3, pp. 1132–1141, 2016.
- [102] T. Cukur, J. M. Santos, J. M. Pauly, and D. G. Nishimura, “Variable-density parallel imaging with partially localized coil sensitivities,” *IEEE transactions on medical imaging*, vol. 29, no. 5, pp. 1173–1181, 2010.
- [103] A. Koç, B. Bartan, E. Gundogdu, T. Çukur, and H. M. Ozaktas, “Sparse representation of two-and three-dimensional images with fractional fourier,

- hartley, linear canonical, and haar wavelet transforms,” *Expert Systems with Applications*, vol. 77, pp. 247–255, 2017.
- [104] U. Gamper, P. Boesiger, and S. Kozerke, “Compressed sensing in dynamic mri,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 59, no. 2, pp. 365–373, 2008.
- [105] T. Çukur, J. M. Santos, D. G. Nishimura, and J. M. Pauly, “Varying kernel-extent gridding reconstruction for undersampled variable-density spirals,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 59, no. 1, pp. 196–201, 2008.
- [106] B. Quist, B. A. Hargreaves, T. Cukur, G. R. Morrell, G. E. Gold, and N. K. Bangerter, “Simultaneous fat suppression and band reduction with large-angle multiple-acquisition balanced steady-state free precession,” *Magnetic resonance in medicine*, vol. 67, no. 4, pp. 1004–1012, 2012.
- [107] S. Ilbey, C. B. Top, A. Güngör, T. Çukur, E. Ü. Sarıtaş, and H. E. Güven, “Comparison of system-matrix-based and projection-based reconstructions for field free line magnetic particle imaging,” *International Journal on Magnetic Particle Imaging*, vol. 3, no. 1, pp. 1–8, 2017.
- [108] E. Kopanoglu, A. Güngör, T. Kilic, E. U. Saritas, T. Çukur, and H. E. Guven, “Compressive sensing reconstruction for multi-contrast data with unequal acceleration rates,” 2018.
- [109] C. F. Şenel and T. Çukur, “3 boyutlu kartezyen olmayan paralel görüntüleme deęişken görüntü alanına dayalı geriçatım,” in *2017 21st National Biomedical Engineering Meeting (BIYOMUT)*, IEEE, 2018.
- [110] C. F. Senel and T. Cukur, “Variable-fov reconstruction for 3d non-cartesian parallel imaging,” in *2017 21st National Biomedical Engineering Meeting (BIYOMUT)*, pp. i–iii, IEEE, 2017.

- [111] A. Gungor, E. Kopanoglu, T. Cukur, E. Guven, and F. T. Yarman-Vural, “Joint dictionary learning reconstruction of compressed multi-contrast magnetic resonance imaging,” in *2017 21st National Biomedical Engineering Meeting (BIYOMUT)*, pp. i–iv, IEEE, 2017.
- [112] J. P. Haldar, D. Hernando, and Z.-P. Liang, “Compressed-sensing mri with random encoding,” *IEEE transactions on Medical Imaging*, vol. 30, no. 4, pp. 893–903, 2010.
- [113] E. Ilıcak and T. Çukur, “Ruber function based reconstruction in accelerated phase-cycled bssfp acquisitions for increased detection performance,” in *2017 25th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2017.
- [114] Y. B. Can, E. Ilıcak, and T. Çukur, “Fast 3d variable-fov reconstruction for parallel imaging with localized sensitivities,” *arXiv preprint arXiv:1612.00157*, 2016.
- [115] E. Ilıcak, S. Cetin, E. U. Saritas, G. Unal, *et al.*, “Adaptive reconstruction for vessel preservation in unenhanced mr angiography,” in *2016 24th Signal Processing and Communication Application Conference (SIU)*, pp. 577–580, IEEE, 2016.
- [116] Ö. Yılmaz, E. Ü. Saritaş, and T. Çukur, “Phase-sensitive reconstruction for fat-water separation in multi-coil acquisitions,” in *2016 24th Signal Processing and Communication Application Conference (SIU)*, pp. 469–472, IEEE, 2016.
- [117] K. H. Jin, D. Lee, and J. C. Ye, “A general framework for compressed sensing and parallel mri using annihilating filter based low-rank hankel matrix,” *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 480–495, 2016.
- [118] X. Zhang, D. Guo, Y. Huang, Y. Chen, L. Wang, F. Huang, Q. Xu, and X. Qu, “Image reconstruction with low-rankness and self-consistency of k-space data in parallel mri,” *Medical Image Analysis*, vol. 63, p. 101687, 2020.

- [119] E. Ilicak, E. U. Saritas, and T. Çukur, “Automated parameter selection for accelerated mri reconstruction via low-rank modeling of local k-space neighborhoods,” *Zeitschrift für Medizinische Physik*, 2022.
- [120] J. P. Haldar, “Low-Rank Modeling of Local k-Space Neighborhoods (LO-RAKS) for Constrained MRI,” *IEEE Transactions on Medical Imaging*, vol. 33, no. 3, pp. 668–681, 2014.
- [121] M. Özbey and T. Çukur, “Multi-image reconstruction in multi-contrast mri,” in *2021 29th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2021.
- [122] İ. Kaftan, T. Çukur, *et al.*, “Synergistic reconstruction-synthesis of multi-contrast mri using transfer learning method,” in *2021 29th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2021.
- [123] S. U. H. Dar, M. Yurt, and T. Çukur, “A few-shot learning approach for accelerated mri via fusion of data-driven and subject-driven priors,” *arXiv preprint arXiv:2103.07790*, 2021.
- [124] M. Yurt, M. Özbey, S. U. H. Dar, B. Tınaz, K. K. Oğuz, and T. Çukur, “Progressively volumetrized deep generative models for data-efficient contextual learning of mr image recovery,” *arXiv preprint arXiv:2011.13913*, 2020.
- [125] M. Yurt, M. Özbey, S. U. H. Dar, B. Tınaz, K. K. Oğuz, and T. Çukur, “Progressively Volumetrized Deep Generative Models for Data-Efficient Contextual Learning of MR Image Recovery,” in *29th Annual Meeting of International Society for Magnetic Resonance Imaging (ISMRM)*, (Virtual Conference),5 2021.
- [126] M. Özbey, M. Yurt, S. U. H. Dar, and T. Çukur, “Three Dimensional MR Image Synthesis with Progressive Generative Adversarial Networks,” in *IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, (Virtual Conference),4 2020.

- [127] S. U. Dar, Ş. Öztürk, Y. Korkmaz, G. Elmas, M. Özbey, A. Güngör, and T. Çukur, “Adaptive diffusion priors for accelerated mri reconstruction,” *arXiv e-prints*, pp. arXiv–2207, 2022.
- [128] G. Elmas, S. U. Dar, Y. Korkmaz, E. Ceyani, B. Susam, M. Özbey, S. Avestimehr, and T. Çukur, “Federated learning of generative image priors for mri reconstruction,” *arXiv preprint arXiv:2202.04175*, 2022.
- [129] M. Acar, T. Çukur, and İ. Öksüz, “Self-supervised dynamic mri reconstruction,” in *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 35–44, Springer, 2021.
- [130] S. Dar, M. Yurt, M. Shahdloo, M. E. Ildız, and T. Çukur, “Joint recovery of variably accelerated multi-contrast mri acquisitions via generative adversarial networks,” in *Proceedings of the International Society of magnetic Resonance in Medicine (ISMRM)*, 2019.
- [131] S. H. Dar and T. Çukur, “Transfer learning for reconstruction of accelerated mri acquisitions via neural networks,” in *Proceedings of the 26th Scientific Meeting of ISMRM*, 2018.
- [132] M. Yurt, M. Özbey, S. U. Dar, B. Tinaz, K. K. Oguz, and T. Çukur, “Progressively volumetrized deep generative models for data-efficient contextual learning of mr image recovery,” *Medical Image Analysis*, vol. 78, p. 102429, 2022.
- [133] S. U. H. Dar, M. Özbey, A. B. Çatlı, and T. Çukur, “A transfer-learning approach for accelerated mri using deep neural networks,” *Magnetic resonance in medicine*, vol. 84, no. 2, pp. 663–685, 2020.
- [134] S. U. H. Dar, M. Yurt, M. Shahdloo, M. E. Ildız, and T. Çukur, “Synergistic reconstruction and synthesis via generative adversarial networks for accelerated multi-contrast mri,” *arXiv preprint arXiv:1805.10704*, 2018.
- [135] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” in *IEEE CVPR*, pp. 9446–9454, 2018.

- [136] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, “Analyzing and improving the image quality of StyleGAN,” in *IEEE CVPR*, pp. 8107–8116, 2020.
- [137] D. A. Hudson and C. L. Zitnick, “Generative adversarial transformers,” *arXiv:2103.01209*, 2021.
- [138] J. Yang, C. Li, P. Zhang, X. Dai, B. Xiao, L. Yuan, and J. Gao, “Focal self-attention for local-global interactions in vision transformers,” *arXiv:2107.00641*, 2021.
- [139] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” *IEEE ICCV*, 2021.
- [140] X. Huang and S. J. Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” *arXiv:1703.06868*, 2017.
- [141] 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*, pp. 2672–2680, 2014.
- [142] R. Child, S. Gray, A. Radford, and I. Sutskever, “Generating long sequences with sparse transformers,” *arXiv:1904.10509*, 2019.
- [143] C. M. Bishop, *Pattern recognition and machine learning*. Springer Verlag, 2006.
- [144] L. Mescheder, A. Geiger, and S. Nowozin, “Which training methods for GANs do actually converge?,” in *Proceedings of ICML*, vol. 80, pp. 3481–3490, 2018.
- [145] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, “Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 9, pp. 2251–2265, 2019.

- [146] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, 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.
- [147] M. Bydder, D. J. Larkman, and J. V. Hajnal, “Combination of signals from array coils using image-based estimation of coil sensitivity profiles,” *Magnetic Resonance in Medicine*, vol. 47, no. 3, pp. 539–548, 2002.
- [148] J. P. Haldar and J. Zhuo, “P-LORAKS: Low-Rank Modeling of Local k-Space Neighborhoods with Parallel Imaging Data,” *Magnetic resonance in medicine*, vol. 75, no. 4, p. 1499, 2016.
- [149] T. Kim and J. Haldar, “Loraks software version 2.0: Faster implementation and enhanced capabilities,” *University of Southern California, Los Angeles, CA, Tech. Rep. USC-SIPI-443*, 2018.
- [150] F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. De-fazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdalv, A. Romero, M. Rabbat, P. Vincent, J. Pinkerton, D. Wang, N. Yakubova, E. Owens, C. L. Zitnick, M. P. Recht, D. K. Sodickson, and Y. W. Lui, “fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning,” *Radiology: Artificial Intelligence*, vol. 2, no. 1, p. e190007, 2020.
- [151] T. Zhang, J. M. Pauly, S. S. Vasanawala, and M. Lustig, “Coil compression for accelerated imaging with Cartesian sampling.,” *Magnetic Resonance in Medicine*, vol. 69, no. 2, pp. 571–82, 2013.
- [152] Y. Wu, Y. Ma, J. Liu, J. Du, and L. Xing, “Self-attention convolutional neural network for improved MR image reconstruction,” *Information Sciences*, vol. 490, pp. 317–328, 2019.
- [153] H. Lan, , A. W. Toga, and F. Sepeshband, “SC-GAN: 3D self-attention conditional GAN with spectral normalization for multi-modal neuroimaging synthesis,” *bioRxiv:2020.06.09.143297*, 2020.

- [154] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, and D. Rueckert, “Attention U-Net: Learning where to look for the pancreas,” *arXiv:1804.03999*, 2018.
- [155] Z. Yuan, M. Jiang, Y. Wang, B. Wei, Y. Li, P. Wang, W. Menpes-Smith, Z. Niu, and G. Yang, “SARA-GAN: Self-attention and relative average discriminator based generative adversarial networks for fast compressed sensing MRI reconstruction,” *Frontiers in Neuroinformatics*, vol. 14, p. 58, 2020.
- [156] S. Ruder, “An overview of gradient descent optimization algorithms,” *arXiv:1609.04747*, 2016.
- [157] X. Zhang, H. Lu, D. Guo, L. Bao, F. Huang, Q. Xu, and X. Qu, “A guaranteed convergence analysis for the projected fast iterative soft-thresholding algorithm in parallel mri,” *Medical Image Analysis*, vol. 69, p. 101987, 2021.
- [158] W. Huang, P. Hand, R. Heckel, and V. Voroninski, “A Provably Convergent Scheme for Compressive Sensing Under Random Generative Priors,” *J Fourier Anal Appl*, vol. 27, no. 2, p. 19, 2021.
- [159] R. Heckel and M. Soltanolkotabi, “Compressive sensing with un-trained neural networks: Gradient descent finds the smoothest approximation,” *arXiv:2005.03991*, 2020.
- [160] M. Asim, M. Daniels, O. Leong, A. Ahmed, and P. Hand, “Invertible generative models for inverse problems: mitigating representation error and dataset bias,” in *Proceedings of ICML*, vol. 119, pp. 399–409, 2020.
- [161] A. Grover, M. Dhar, and S. Ermon, “Flow-gan: Bridging implicit and prescribed learning in generative models,” *arXiv:1705.08868*, 2017.
- [162] Z. Wang, C. Qian, D. Guo, H. Sun, R. Li, B. Zhao, and X. Qu, “One-dimensional deep low-rank and sparse network for accelerated mri,” *arXiv:2112.04721*, 2021.

- [163] Z. Wang, D. Guo, Z. Tu, Y. Huang, Y. Zhou, J. Wang, L. Feng, D. Lin, Y. You, T. Agback, V. Orekhov, and X. Qu, “Xcloud-modern: An artificial intelligence cloud for accelerated nmr spectroscopy,” *arXiv:2012.14830*, 2021.
- [164] E. Çelik, S. U. H. Dar, Ö. Yılmaz, Ü. Keleş, and T. Çukur, “Spatially informed voxelwise modeling for naturalistic fmri experiments,” *NeuroImage*, vol. 186, pp. 741–757, 2019.
- [165] M. Shahdloo, E. Çelik, and T. Çukur, “Biased competition in semantic representation during natural visual search,” *Neuroimage*, vol. 216, p. 116383, 2020.
- [166] Ö. Yılmaz, E. Çelik, and T. Çukur, “Informed feature regularization in voxelwise modeling for naturalistic fmri experiments,” *European Journal of Neuroscience*, vol. 52, no. 5, pp. 3394–3410, 2020.
- [167] S. Tokgoz, D. Aydogdu, B. Ilhan, Y. Sahin, N. Bariseri, B. M. Ozturkler, and T. Çukur, “Musical mirror-symmetrical movement tasks: comparison of rhythm versus melody-playing,” *NeuroReport*, vol. 31, no. 7, pp. 523–529, 2020.
- [168] T. Çukur, A. G. Huth, S. Nishimoto, and J. L. Gallant, “Functional subdomains within human ffa,” *Journal of Neuroscience*, vol. 33, no. 42, pp. 16748–16766, 2013.
- [169] T. Çukur, A. G. Huth, S. Nishimoto, and J. L. Gallant, “Functional subdomains within scene-selective cortex: parahippocampal place area, retrosplenial complex, and occipital place area,” *Journal of Neuroscience*, vol. 36, no. 40, pp. 10257–10273, 2016.
- [170] T. Cukur, S. Nishimoto, A. G. Huth, and J. L. Gallant, “Attention during natural vision warps semantic representation across the human brain,” *Nature neuroscience*, vol. 16, no. 6, pp. 763–770, 2013.
- [171] T. Zhang, J. S. Gao, T. Çukur, and J. L. Gallant, “Voxel-based state space modeling recovers task-related cognitive states in naturalistic fmri experiments,” *Frontiers in neuroscience*, vol. 14, p. 1103, 2021.

- [172] I. Kiremitçi, Ö. Yılmaz, E. Çelik, M. Shahdloo, A. G. Huth, and T. Çukur, “Attentional modulation of hierarchical speech representations in a multitalker environment,” *bioRxiv*, pp. 2020–12, 2021.
- [173] A. Ozaslan, A. Alacaoglu, O. Demirel, T. Çukur, and E. Saritas, “Fully automated gridding reconstruction for non-cartesian x-space magnetic particle imaging,” *Physics in Medicine & Biology*, vol. 64, no. 16, p. 165018, 2019.
- [174] S. Koppers, C. Haarbuerger, and D. Merhof, “Diffusion mri signal augmentation: from single shell to multi shell with deep learning,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 61–70, Springer, 2016.
- [175] E. Celik, U. Keles, İ. Kiremitçi, J. L. Gallant, and T. Cukur, “Cortical networks of dynamic scene category representation in the human brain,” *cortex*, vol. 143, pp. 127–147, 2021.
- [176] K. Keskin, U. Yılmaz, and T. Çukur, “Constrained ellipse fitting for efficient parameter mapping with phase-cycled bssfp mri,” *arXiv preprint arXiv:2106.03239*, 2021.
- [177] M. Shahdloo, E. Çelik, B. A. Urgen, J. L. Gallant, and T. Çukur, “Task-dependent warping of semantic representations during search for visual action categories,” *bioRxiv*, 2021.
- [178] B. Aşkın, A. Güngör, D. A. Soydan, C. B. Top, and T. Cukur, “A cnn based super-resolution technique for magnetic particle imaging system matrix,” in *2021 29th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2021.
- [179] T. Sanchez, B. Gözcü, R. B. van Heeswijk, A. Eftekhari, E. Ilıcak, T. Çukur, and V. Cevher, “Scalable learning-based sampling optimization for compressive dynamic mri,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8584–8588, IEEE, 2020.

- [180] E. Bıyık, K. Keskin, S. UH Dar, A. Koç, and T. Çukur, “Factorized sensitivity estimation for artifact suppression in phase-cycled bssfp mri,” *NMR in Biomedicine*, vol. 33, no. 4, p. e4228, 2020.
- [181] S. Ilbey, C. B. Top, A. Güngör, T. Çukur, E. U. Saritas, and H. E. Güven, “Fast system calibration with coded calibration scenes for magnetic particle imaging,” *IEEE transactions on medical imaging*, vol. 38, no. 9, pp. 2070–2080, 2019.
- [182] H. A. Bedel, I. Şıvgın, O. Dalmaz, S. U. H. Dar, and T. Çukur, “Bolt: Fused window transformers for fmri time series analysis,” *arXiv preprint arXiv:2205.11578*, 2022.
- [183] R. Beliy, G. Gaziv, A. Hoogi, F. Strappini, T. Golan, and M. Irani, “From voxels to pixels and back: Self-supervision in natural-image reconstruction from fmri,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [184] E. Ergun, A. Ö. Arol, E. U. Saritas, and T. Çukur, “A deblurring model for x-space mpi based on coded calibration scenes,” *International Journal on Magnetic Particle Imaging*, vol. 8, no. 1 Suppl 1, 2022.
- [185] A. Güngör, B. Askin, D. A. Soydan, C. B. Top, and T. Cukur, “Deep learned super resolution of system matrices for magnetic particle imaging,” in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 3749–3752, IEEE, 2021.
- [186] B. Aşkın, A. Güngör, D. A. Soydan, C. B. Top, and T. Çukur, “Manyetik parçacık görüntüleme için evrimsel sinir ağı tabanlı bir süper-çözünürlük tekniği,” *IEEE Xplore*, pp. 1–4, 2021.
- [187] M. Aksoy, C. Forman, M. Straka, T. Çukur, J. Hornegger, and R. Bammer, “Hybrid prospective and retrospective head motion correction to mitigate cross-calibration errors,” *Magnetic resonance in medicine*, vol. 67, no. 5, pp. 1237–1251, 2012.

- [188] T. Cukur, J. DiCarlo, N. Bangerter, B. Hargreaves, and D. Nishimura, “Flow-independent angiography of the hand with 3d balanced ssfp imaging,” in *Proceedings of the 14th Annual Meeting of ISMRM*, p. 1967, 2006.
- [189] R. Ingle, T. Cukur, and D. Nishimura, “Spectral profile design for multiple repetition time balanced ssfp,” in *Proceedings of the 18th Annual Meeting of ISMRM*, p. 75, 2010.
- [190] K. Keskin and T. Çukur, “Constrained ellipse fitting for efficient t1-t2 mapping in phase-cycled bssfp imaging,” in *Proceedings of the ISMRM 27th Annual Meeting and Exhibition. Montréal, Canada*, vol. 4533, 2019.
- [191] T. Cukur, J. Lee, N. Bangerter, B. Hargreaves, and D. Nishimura, “Comparison of phase-sensitive and alternating repetition time ssfp for flow-independent peripheral angiography,” in *Proceedings of the 15th Annual Meeting of ISMRM*, p. 178, Berlin, 2007.
- [192] J. Schlemper, G. Yang, P. Ferreira, A. Scott, L.-A. McGill, Z. Khalique, M. Gorodezky, M. Roehl, J. Keegan, D. Pennell, *et al.*, “Stochastic deep compressive sensing for the reconstruction of diffusion tensor cardiac mri,” in *International conference on medical image computing and computer-assisted intervention*, pp. 295–303, Springer, 2018.
- [193] S. G. Kafali, T. Cukur, and E. U. Saritas, “Joint non-local means reconstruction for correction of phase-induced errors in diffusion tensor imaging,” in *Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, HI*, p. 7364, 2017.
- [194] T. Cukur, B. Hargreaves, and D. Nishimura, “Reduction of the transient oscillations in alternating-tr ssfp,” in *Proceedings of the 15th Annual Meeting of ISMRM*, p. 3447, 2007.
- [195] R. R. Ingle, T. Çukur, and D. G. Nishimura, “The central signal singularity phenomenon in balanced ssfp and its application to positive-contrast imaging,” *Magnetic resonance in medicine*, vol. 67, no. 6, pp. 1673–1683, 2012.

- [196] O. Yorulmaz, O. B. Demirel, Y. Muslu, T. Çukur, E. U. Saritas, and A. E. Çetin, “A blind deconvolution technique based on projection onto convex sets for magnetic particle imaging,” *arXiv preprint arXiv:1705.07506*, 2017.
- [197] O. Yilmaz, E. U. Saritas, and T. Çukur, “Enhanced phase-sensitive ssfp reconstruction for fat-water separation in phased-array acquisitions,” *Journal of Magnetic Resonance Imaging*, vol. 44, no. 1, pp. 148–157, 2016.
- [198] V. Golkov, A. Dosovitskiy, J. I. Sperl, M. I. Menzel, M. Czisch, P. Sämann, T. Brox, and D. Cremers, “Q-space deep learning: twelve-fold shorter and model-free diffusion mri scans,” *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1344–1351, 2016.
- [199] S. G. Kafali, T. Çukur, and E. U. Saritas, “Phase-correcting non-local means filtering for diffusion-weighted imaging of the spinal cord,” *Magnetic resonance in medicine*, vol. 80, no. 3, pp. 1020–1035, 2018.
- [200] S. Ilbey, C. B. Top, T. Çukur, E. U. Saritas, and H. E. Güven, “Image reconstruction for magnetic particle imaging using an augmented lagrangian method,” in *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pp. 64–67, IEEE, 2017.
- [201] T. Çukur and D. G. Nishimura, “Fat-water separation with alternating repetition time balanced ssfp,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 60, no. 2, pp. 479–484, 2008.
- [202] N. K. Bangerter, T. Cukur, B. A. Hargreaves, B. S. Hu, J. H. Brittain, D. Park, G. E. Gold, and D. G. Nishimura, “Three-dimensional fluid-suppressed t2-prep flow-independent peripheral angiography using balanced ssfp,” *Magnetic resonance imaging*, vol. 29, no. 8, pp. 1119–1124, 2011.
- [203] T. Çukur, M. Yamada, W. R. Overall, P. Yang, and D. G. Nishimura, “Positive contrast with alternating repetition time ssfp (parts): A fast imaging technique for spio-labeled cells,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 63, no. 2, pp. 427–437, 2010.

- [204] T. Çukur, A. Shimakawa, H. Yu, B. A. Hargreaves, B. S. Hu, D. G. Nishimura, and J. H. Brittain, “Magnetization-prepared ideal bssfp: a flow-independent technique for noncontrast-enhanced peripheral angiography,” *Journal of Magnetic Resonance Imaging*, vol. 33, no. 4, pp. 931–939, 2011.
- [205] E. U. Saritas, D. Lee, T. Cukur, A. Shankaranarayanan, and D. G. Nishimura, “Hadamard slice encoding for reduced-fov diffusion-weighted imaging,” *Magnetic resonance in Medicine*, vol. 72, no. 5, pp. 1277–1290, 2014.
- [206] T. Çukur, S. Nishimoto, A. G. Huth, and J. L. Gallant, “Attention during natural vision warps semantic representation across the human brain,” *Nature neuroscience*, vol. 16, no. 6, pp. 763–770, 2013.
- [207] Ö. Yılmaz, E. Çelik, and T. Çukur, “Feature regularization in voxelwise modeling for naturalistic fmri experiments.,” *Anatomy: International Journal of Experimental & Clinical Anatomy*, vol. 14, 2020.
- [208] T. Çukur, “Voxelwise modeling for identifying cortical correlates of natural perception and attention.,” *Anatomy: International Journal of Experimental & Clinical Anatomy*, vol. 14, 2020.
- [209] K. Keskin and T. Çukur, “Elliptical signal model based simultaneous parameter estimation in bssfp imaging,” in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2018.
- [210] K. Keskin and T. Çukur, “Bssfp görüntülemeye eliptik sinyal modeline dayalı eşzamanlı parametre tahmini,” in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, IEEE, 2018.
- [211] K. Keskin and T. Çukur, “Faz döngülü bssfp görüntülemeye t1 ve t2 haritalamaya parametre tahmini yaklaşımı,” in *2017 21st National Biomedical Engineering Meeting (BIYOMUT)*, IEEE, 2018.
- [212] K. Keskin and T. Cukur, “Parameter estimation approach to t1 and t2 mapping in phase-cycled bssfp imaging,” in *2017 21st National Biomedical Engineering Meeting (BIYOMUT)*, pp. i–iv, IEEE, 2017.

- [213] S. U. H. Dar and T. Çukur, “Category-selective top-down modulation in the fusiform face area of the human brain during visual search,” in *2017 25th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2017.
- [214] S. G. Kafalı, T. Çukur, and E. Ü. Sarıtaş, “Simultaneous phase-correction and denoising for diffusion-weighted mri,” in *2016 24th Signal Processing and Communication Application Conference (SIU)*, pp. 1313–1316, IEEE, 2016.
- [215] H. K. Aggarwal, M. P. Mani, and M. Jacob, “Modl-mussels: model-based deep learning for multishot sensitivity-encoded diffusion mri,” *IEEE transactions on medical imaging*, vol. 39, no. 4, pp. 1268–1277, 2019.
- [216] T. Çukur, S. Nishimoto, A. Huth, and J. Gallant, “Category-based attention shifts tuning toward the target object category during natural visual search,” *Journal of Vision*, vol. 12, no. 9, pp. 955–955, 2012.
- [217] 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.
- [218] 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.
- [219] M. Özbey and T. Çukur, “T1-weighted contrast-enhanced synthesis for multi-contrast mri segmentation,” in *2020 28th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2020.
- [220] S. Türk, T. Cukur, C. HEPDURGUN, S. Dar, A. Karabulut, I. TAMSEL, U. Aydingoz, M. ARGIN, *et al.*, “Deep learning algorithm for sacroiliac joint evaluation: Grading with radiographs,” European Congress of Radiology-ECR 2020, 2020.
- [221] O. Dalmaz, U. Mirza, G. Elmas, M. Özbey, S. U. Dar, E. Ceyani, S. Avestimehr, and T. Çukur, “One model to unite them all: Personalized federated

- learning of multi-contrast mri synthesis,” *arXiv preprint arXiv:2207.06509*, 2022.
- [222] Ş. Öztürk and T. Çukur, “Deep clustering via center-oriented margin free-triplet loss for skin lesion detection in highly imbalanced datasets,” *IEEE Journal of Biomedical and Health Informatics*, 2022.
- [223] I. AYTEKIN, O. DALMAZ, H. ANKISHAN, E. U. SARITAS, U. BAGCI, T. ÇUKUR, and H. CELIK, “Detecting covid-19 from respiratory sound recordings with transformers,” in *Medical Imaging 2022: Computer-Aided Diagnosis*, vol. 12033, pp. 25–32, SPIE, 2022.
- [224] O. DALMAZ, B. SAGLAM, K. GÖNÇ, and T. ÇUKUR, “edagan: Encoder-decoder attention generative adversarial networks for multi-contrast mr image synthesis,” in *2022 9th International Conference on Electrical and Electronics Engineering (ICEEE)*, pp. 320–324, IEEE, 2022.
- [225] M. ÖZBEY and T. ÇUKUR, “Çoklu kontrast mrg’de çoklu görüntü geriçatımı,” *IEEE Xplore*, pp. 1–4, 2021.
- [226] Q. LYU, H. SHAN, C. STEBER, C. HELIS, C. WHITLOW, M. CHAN, and G. WANG, “Multi-contrast super-resolution mri through a progressive network,” *IEEE transactions on medical imaging*, vol. 39, no. 9, pp. 2738–2749, 2020.
- [227] İ. KAFTAN, Ö. B. GEVREK, and T. ÇUKUR, “Çoklu kontrast mrg’de tranfer öğrenme yöntemi ile birleşik geriçatım-sentez,” *IEEE Xplore*, pp. 1–4, 2021.
- [228] M. YURT, S. U. H. DAR, A. ERDEM, E. ERDEM, and T. ÇUKUR, “A multi-stream gan approach for multi-contrast mri synthesis,” in *28th annual meeting of International Society for Magnetic Resonance Imaging (ISMRM), (Virtual Conference)*, vol. 8, 2020.
- [229] M. ÖZBEY, M. YURT, S. U. H. DAR, and T. ÇUKUR, “Three dimensional mr image synthesis with progressive generative adversarial networks,” *arXiv preprint arXiv:2101.05218*, 2020.
- [230] M. YURT, S. U. H. DAR, M. ÖZBEY, B. TMAZ, K. K. OĞUZ, and T. ÇUKUR, “Semi-supervised learning of mutually accelerated mri synthesis without fully-sampled ground truths,” *arXiv preprint arXiv:2011.14347*, 2020.

Appendix A

Supplementary Materials

A.1 Supplementary Text

A.1.1 Positional encoding variables

For input feature maps $X \in \mathbb{R}^{h_1 \times h_2 \times u}$, sinusoidal position encoding variables $PE \in \mathbb{R}^{h_1 \times h_2 \times u}$ are set at location $(loc_{h_1}, loc_{h_2}, loc_u)$ as [74]:

$$PE[loc_{h_1}, loc_{h_2}, loc_u] = \begin{cases} \sin\left(\frac{loc_{h_2}}{10000^{4(\frac{loc_u}{u})}}\right) & 0 < loc_u \leq \frac{u}{4} \\ \cos\left(\frac{loc_{h_2}}{10000^{4(\frac{loc_u}{u} - \frac{1}{4})}}\right) & \frac{u}{4} \leq loc_u \leq \frac{u}{2} \\ \sin\left(\frac{loc_{h_1}}{10000^{4(\frac{loc_u}{u} - \frac{2}{4})}}\right) & \frac{u}{2} \leq loc_u \leq \frac{u}{3} \\ \cos\left(\frac{loc_{h_1}}{10000^{4(\frac{loc_u}{u} - \frac{3}{4})}}\right) & \frac{u}{3} \leq loc_u \leq u \end{cases}$$

where loc_{h_1} and loc_{h_2} lie in range $[-1, 1]$, covering complete field of view along the first two spatial dimensions, and loc_u is the channel index.

A.1.2 Architectural Details

A.1.2.a Synthesizer

- Layer 1 (4x4): Input(Constant) \rightarrow Cross-Attention Transformer Block \rightarrow Output
- Layer 2 (8x8): Input \rightarrow Upsample \rightarrow Cross-Attention Transformer Block + Upsample(Input) \rightarrow Output
- Layer 3 (16x16): Input \rightarrow Upsample \rightarrow Cross-Attention Transformer Block + Upsample(Input) \rightarrow Output
- Layer 4 (32x32): Input \rightarrow Upsample \rightarrow Cross-Attention Transformer Block + Upsample(Input) \rightarrow Output
- Layer 5 (64x64): Input \rightarrow Upsample \rightarrow Cross-Attention Transformer Block + Upsample(Input) \rightarrow Output
- Layer 6 (128x128): Input \rightarrow Upsample \rightarrow Cross-Attention Transformer Block + Upsample(Input) \rightarrow Output
- Layer 7 (256x256): Input \rightarrow Upsample \rightarrow Modulated Convolution \rightarrow Output
- Cross-attention Transformer Block: Input \rightarrow Cross-Attention + Noise \rightarrow Modulated Convolution \rightarrow Cross-Attention + Noise \rightarrow Output

A.1.2.b Mapper

A.1.2.b.1 Local Stream

- Layer 1: Input \rightarrow Self-Attention Block \rightarrow Output
- Layer 2: Input \rightarrow Self-Attention Block \rightarrow Output

- Layer 3: Input \rightarrow Self-Attention Block \rightarrow Output
- Layer 4: Input \rightarrow Self-Attention Block \rightarrow Output
- Layer 5: Input \rightarrow Fully-connected \rightarrow Output
- Self-Attention Block: Input \rightarrow Self-Attention \rightarrow Fully-connected \rightarrow Fully-connected + Input \rightarrow Output

A.1.2.b.2 Global Stream

- Layer 1: Input \rightarrow Fully-connected \rightarrow Output
- Layer 2: Input \rightarrow Fully-connected \rightarrow Output
- Layer 3: Input \rightarrow Fully-connected \rightarrow Output
- Layer 4: Input \rightarrow Fully-connected \rightarrow Output
- Layer 5: Input \rightarrow Fully-connected \rightarrow Output
- Layer 6: Input \rightarrow Fully-connected \rightarrow Output
- Layer 7: Input \rightarrow Fully-connected \rightarrow Output
- Layer 8: Input \rightarrow Fully-connected \rightarrow Output
- Layer 9: Input \rightarrow Fully-connected \rightarrow Output

A.1.2.c Discriminator

- Layer 1 (256x256): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output
- Layer 2 (128x128): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output

- Layer 3 (64x64): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output
- Layer 4 (32x32): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output
- Layer 5 (16x16): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output
- Layer 6 (8x8): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output
- Layer 7 (4x4): Input \rightarrow Convolution \rightarrow Downsample + Downsample(Input) \rightarrow Output

A.2 Supplementary Tables

Supp. Table 1: Within-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in the IXI dataset at R=4 and 8.

	LORAKS		GAN_{sup}		SSDU		GAN_{prior}		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
T_1 , R=4	30.7±1.2	91.7±1.0	37.5±0.5	97.8±0.2	37.9±0.6	97.8±0.2	34.4±0.8	94.4±0.7	32.1±0.9	92.1±0.7	38.8±0.8	97.9±0.5
T_1 , R=8	26.8±0.9	87.3±1.1	33.3±0.6	95.7±0.3	33.1±0.7	93.9±0.7	29.3±1.2	89.7±1.4	28.6±0.9	88.3±1.2	33.2±0.9	95.2±0.9
T_2 , R=4	35.4±0.5	92.3±1.2	38.7±0.8	96.8±0.3	38.9±0.7	96.3±0.4	33.4±0.9	87.5±1.0	34.9±0.6	91.6±1.1	40.0±0.8	97.7±0.5
T_2 , R=8	31.4±0.4	88.2±1.3	34.2±0.8	94.3±0.6	33.7±0.9	91.6±1.1	31.2±0.7	85.3±1.0	30.7±0.5	86.4±1.4	34.1±0.8	94.8±0.7

Supp. Table II: Across-domain reconstruction performance for T_1 - and T_2 -weighted acquisitions in the IXI and fastMRI datasets. In A->B, A and B denote the acceleration rates in training versus test domains. Because LORAKS is untrained, and GAN_{prior} , SAGAN and SLATER do not make any assumptions regarding the imaging operator during training, their across-domain reconstruction performance is equivalent to the within-domain performance for the target acceleration rate.

	LORAKS		GAN_{sup}		SSDU		GAN_{prior}		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
IXI T_1 , 8->4	30.7±1.2	91.7±1.0	32.8±0.9	96.6±0.3	33.1±1.2	94.5±0.9	34.4±0.8	94.4±0.7	32.1±0.9	92.1±0.7	38.8±0.8	97.9±0.5
IXI T_2 , 8->4	35.4±0.5	92.3±1.2	33.7±0.5	93.8±0.4	34.8±0.8	92.4±1.0	33.4±0.9	87.5±1.0	34.9±0.6	91.6±1.1	40.0±0.8	97.7±0.5
fastMRI T_1 , 8->4	33.4±2.7	82.2±7.7	34.8±2.0	93.7±5.7	35.0±2.5	92.1±7.4	32.8±2.0	92.5±5.2	36.1±2.6	94.1±5.1	37.6±3.2	93.9±9.5
fastMRI T_2 , 8->4	34.3±1.0	90.8±1.6	33.3±1.0	94.8±0.6	32.0±1.9	92.5±1.6	33.5±1.1	91.5±1.8	33.5±1.3	94.1±0.8	36.3±1.2	95.5±0.7

Supp. Table III: Reconstruction performance in ablation experiments for SLATER. Metrics are reported for T_1 - and T_2 - weighted acquisitions in the IXI dataset at $R=4$.

	None		Latent		Latent+Noise		Latent+Noise+Weight	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
T_1	26.7±1.3	87.1±1.2	32.4±1.1	94.1±0.7	34.0±1.2	96.6±0.4	38.8±0.8	97.9±0.5
T_2	30.4±0.7	80.5±1.4	32.9±0.8	88.0±0.9	36.2±0.8	94.3±0.6	40.0±0.8	97.7±0.5

Supp. Table IV: Average training time of models in min:sec format per epoch in the IXI dataset. Note that LORAKS does not perform any training.

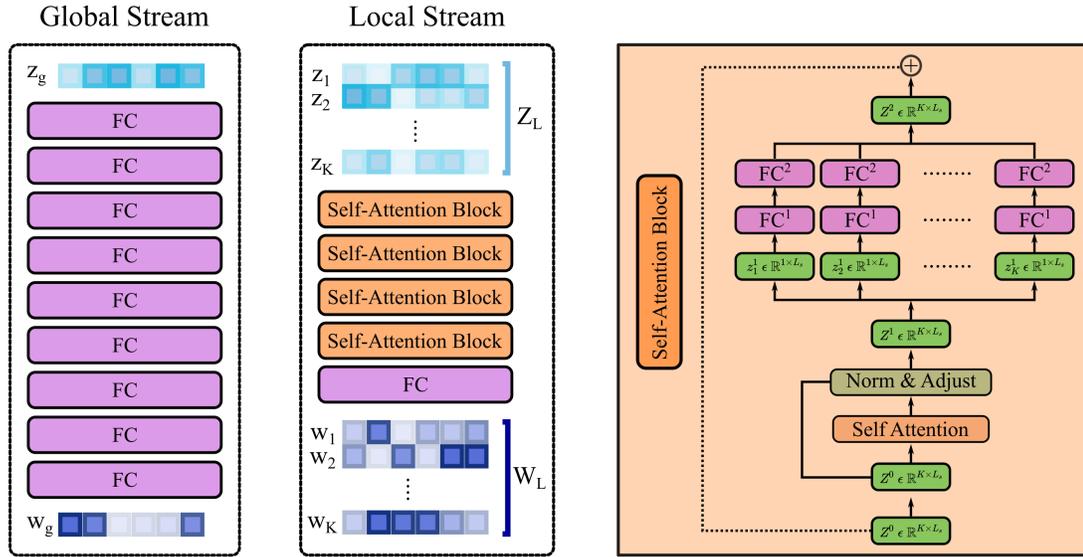
	LORAKS	GAN _{sup}	SSDU	GAN _{prior}	SAGAN	SLATER
Time (min:sec)	–	6:49	1:49	6:22	6:49	8:10

Supp. Table V: Reconstruction performance for T_1 - and T_2 -weighted acquisitions in the IXI dataset at $R=4$ and 8 based on the weight propagation procedure. Note that weight propagation only affects the performance of $\text{GAN}_{\text{prior}}$ and SLATER for which weight optimization is performed during inference.

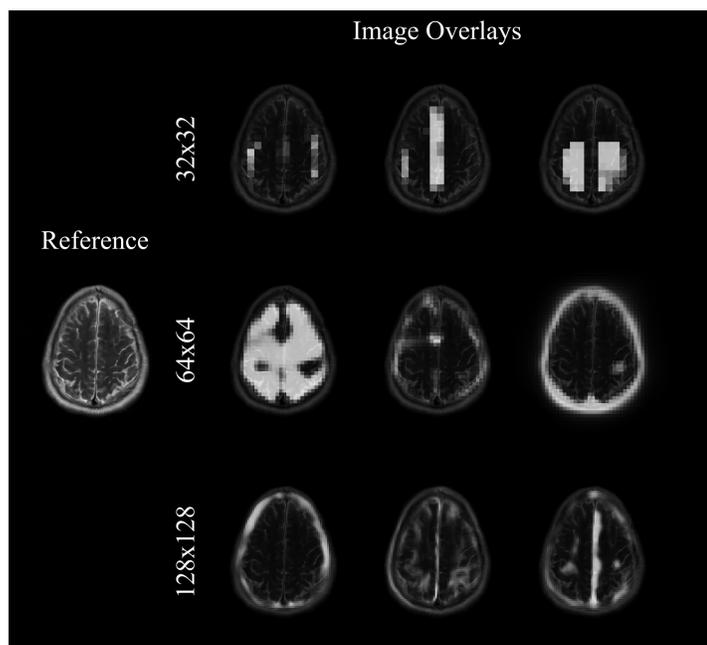
	LORAKS		GAN_{sup}		SSDU		$\text{GAN}_{\text{prior}}$		SAGAN		SLATER	
	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)	PSNR	SSIM(%)
$T_1, R=4$	30.7 ± 1.2	91.7 ± 1.0	37.5 ± 0.5	97.8 ± 0.2	37.9 ± 0.6	97.8 ± 0.2	34.15 ± 0.93	95.06 ± 0.55	32.1 ± 0.9	92.1 ± 0.7	38.63 ± 0.88	98.17 ± 0.24
$T_1, R=8$	26.8 ± 0.9	87.3 ± 1.1	33.3 ± 0.6	95.7 ± 0.3	33.1 ± 0.7	93.9 ± 0.7	29.02 ± 1.11	88.91 ± 1.41	28.6 ± 0.9	88.3 ± 1.2	33.04 ± 1.05	96.06 ± 0.52
$T_2, R=4$	35.4 ± 0.5	92.3 ± 1.2	38.7 ± 0.8	96.8 ± 0.3	38.9 ± 0.7	96.3 ± 0.4	33.04 ± 0.84	88.39 ± 1.23	34.9 ± 0.6	91.6 ± 1.1	39.80 ± 0.80	97.77 ± 0.27
$T_2, R=8$	31.4 ± 0.4	88.2 ± 1.3	34.2 ± 0.8	94.3 ± 0.6	33.7 ± 0.9	91.6 ± 1.1	30.83 ± 0.64	87.30 ± 1.17	30.7 ± 0.5	86.4 ± 1.4	33.96 ± 0.77	94.10 ± 0.46

A.3 Supplementary Figures

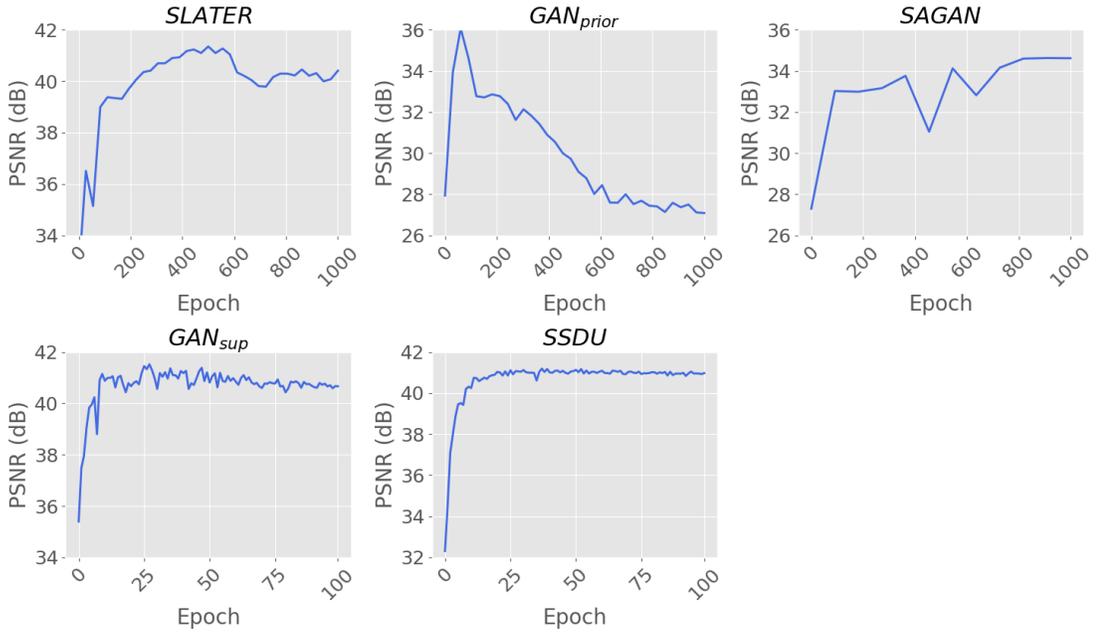
Mapper (M)



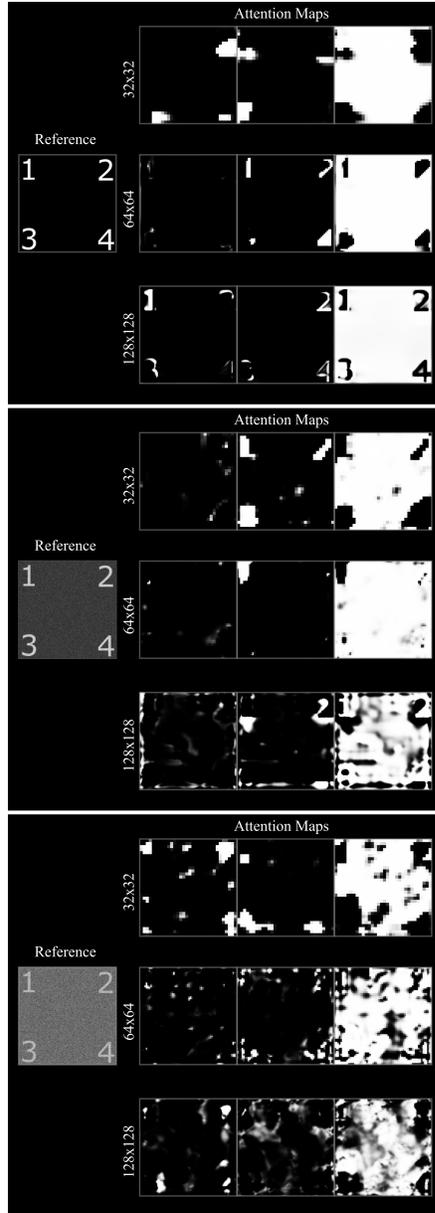
Supp. Fig. A.1: Mapper is a multi-layered architecture comprising two separate processing streams: a global stream dedicated to the global latent variable w_g , and a local stream dedicated to the local latent variables $W_l = \{w_1, w_2, \dots, w_K\}$. The global stream contains a cascade of fully-connected sub-blocks. Meanwhile, the local stream is a cascade of self-attention sub-blocks followed by a fully-connected sub-block (see rightmost panel for the architecture of the self-attention sub-block). Self-attention sub-blocks enable interactions among individual local latents.



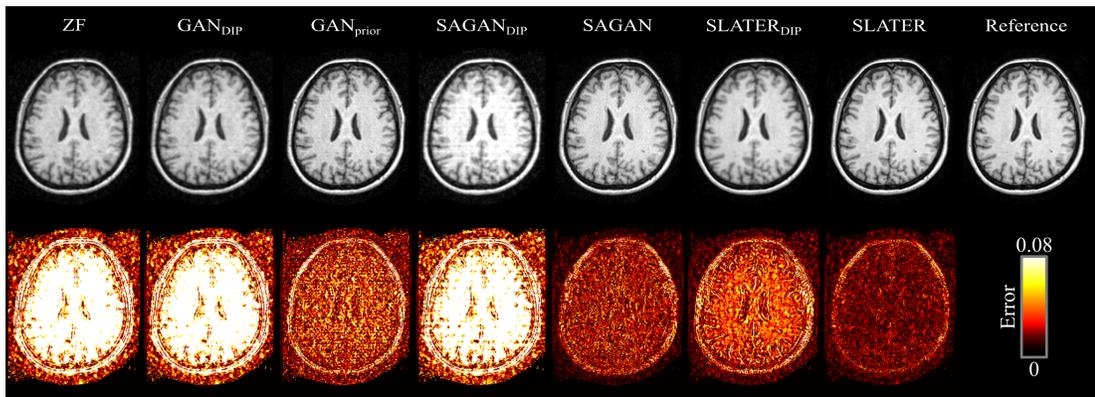
Supp. Fig. A.2: Cross-attention maps in SLATER for a T_2 -weighted acquisition. Sample attention maps from the first cross-attention transformer block are displayed across three resolutions (i.e., 32×32 , 64×64 , 128×128 at network layers 4-6). At each resolution, respective maps are displayed in overlaid format onto the MR image, and the reference MR image is also shown. Attention maps for separate latents show segregated spatial distribution. They also tend to group tissue clusters with similar signal intensity and texture, where the clusters are broadly distributed across the image and they are often spatially noncontiguous.



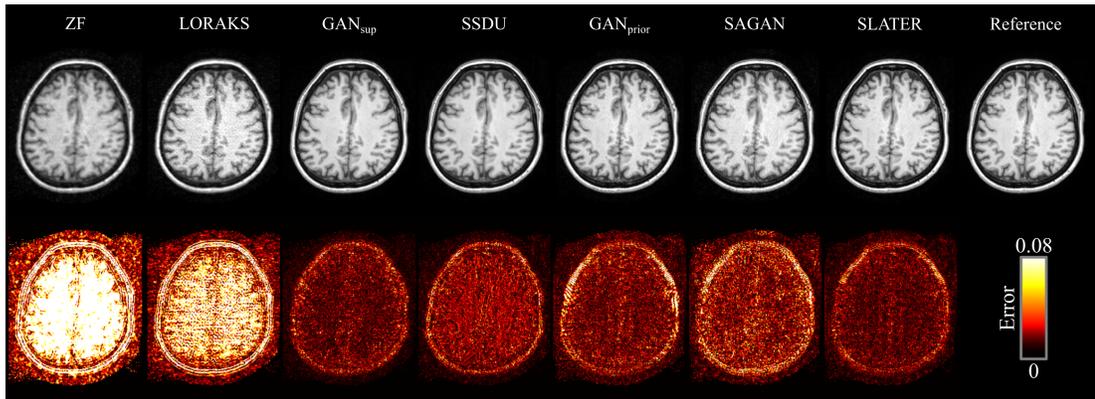
Supp. Fig. A.3: Reconstruction performance in the validation set as a function of number of training epochs. Results from supervised (GAN_{sup}) and unsupervised models (SSDU, $\text{GAN}_{\text{prior}}$, SAGAN and SLATER) are shown for T_1 -weighted acquisitions in IXI at $R=4$. For unsupervised models, hyperparameter selection in the validation set was actually performed based on the difference between recovered and acquired k-space samples in undersampled data. However, to facilitate interpretation, here performance for all methods is displayed as PSNR between reconstructed and ground-truth images.



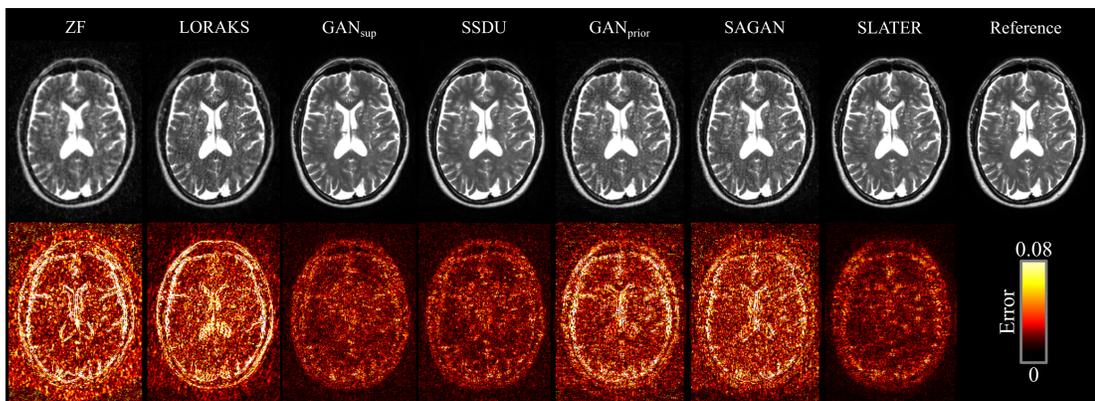
Supp. Fig. A.4: Cross-attention maps in SLATER for a simulated digit phantom with varying levels of noise. Relative to a peak signal intensity of 1, top, middle and bottom panels display sample attention maps for no noise, noise variance of 0.01, and noise variance of 0.1, respectively. Within each panel, maps from the first cross-attention sub-block are shown at three resolutions (i.e., 32x32, 64x64, 128x128 at network layers 4-6), along with the reference phantom images.



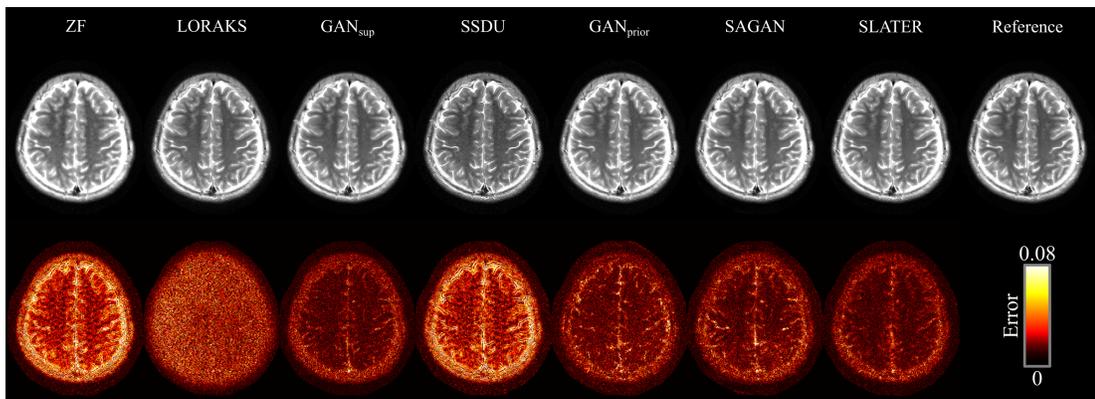
Supp. Fig. A.5: Reconstructions of a representative T₁-weighted acquisition at R=4 are shown for the Fourier method (ZF), DIP methods (GAN_{DIP}, SAGAN_{DIP}, SLATER_{DIP}) and zero-shot reconstructions (GAN_{prior}, SAGAN, SLATER) along with the reference image. Corresponding error maps are underneath the images for each method.



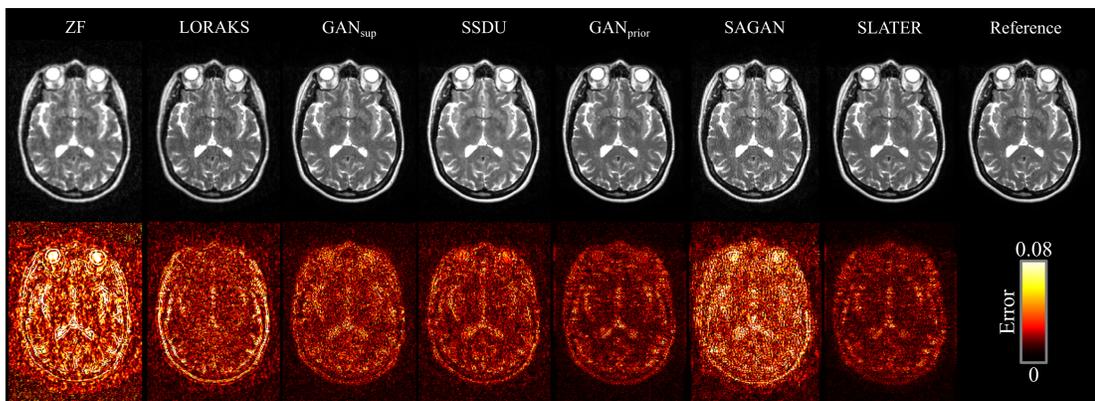
Supp. Fig. A.6: Within-domain reconstructions of a T_1 -weighted acquisition in the IXI dataset at $R=4$ are shown for the Fourier method (ZF), a traditional low-rank method (LORAKS), a supervised baseline (GAN_{sup}), unsupervised baselines (SSDU, GAN_{prior} , SAGAN) and the proposed method (SLATER) along with the reference image. Corresponding error maps are underneath the images for each method.



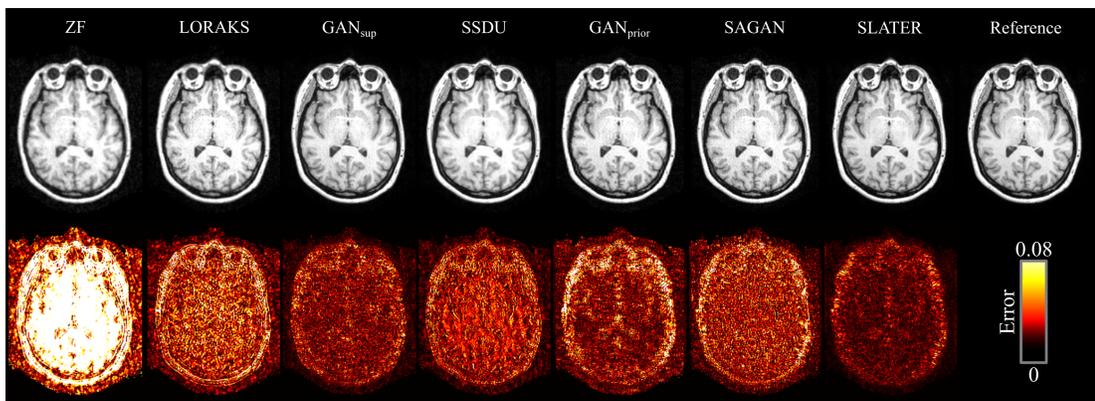
Supp. Fig. A.7: Within-domain reconstructions of a T_2 -weighted acquisition in the IXI dataset at $R=4$. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image. Corresponding error maps are underneath the images for each method.



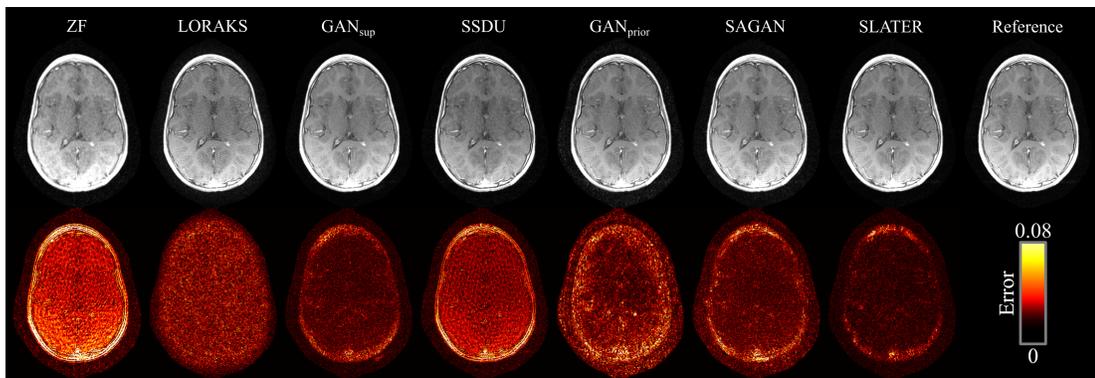
Supp. Fig. A.8: Within-domain reconstructions of a T_2 -weighted acquisition in the fastMRI dataset at $R=4$. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image. Corresponding error maps are underneath the images for each method.



Supp. Fig. A.9: Across-domain reconstructions of a T₂-weighted acquisition in the IXI dataset at R=4. Results are shown for ZF, LORAKS, GAN_{sup}, SSDU, GAN_{prior}, SAGAN and SLATER along with the reference image. Corresponding error maps are underneath the images for each method.



Supp. Fig. A.10: Across-domain reconstructions of a T₁-weighted acquisition in the IXI dataset at R=4. Results are shown for ZF, LORAKS, GAN_{sup}, SSDU, GAN_{prior}, SAGAN and SLATER along with the reference image. Corresponding error maps are underneath the images for each method.



Supp. Fig. A.11: Across-domain reconstructions of a T_1 -weighted acquisition in the fastMRI dataset at $R=4$. Results are shown for ZF, LORAKS, GAN_{sup} , SSDU, GAN_{prior} , SAGAN and SLATER along with the reference image. Corresponding error maps are underneath the images for each method.