DIFFUSION BRIDGES FOR 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
Muhammad Usama Mirza
August 2024

DIFFUSION BRIDGES FOR MRI RECONSTRUCTION
By Muhammad Usama Mirza
August 2024

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

Sevinç Figen Öktem Seven

Approved for the Graduate School of Engineering and Science:

Orhan Arıkan
Director of the Graduate School

ABSTRACT

DIFFUSION BRIDGES FOR MRI RECONSTRUCTION

Muhammad Usama Mirza
M.S. in Electrical and Electronics Engineering
Advisor: Tolga Çukur
August 2024

Magnetic Resonance Imaging (MRI) reconstruction typically involves a dealiasing process to transform undersampled data into fully-sampled data. However, conventional diffusion priors perform a denoising transformation, starting from a state of Gaussian noise and ending with fully-sampled data. Since the aliasing artifacts associated with many k-space sampling patterns have spatial structures that differ significantly from white Gaussian noise, this denoising process can lead to reconstruction errors due to suboptimal suppression of artifacts. To overcome this limitation, we introduce the first Fourier-constrained diffusion bridge (FDB) for MRI reconstruction. Unlike task-agnostic diffusion priors, the FDB is specifically designed to perform a dealiasing transformation, starting with undersampled data and ending with fully-sampled data. The starting point is created using a novel stochastic degradation operator that removes a randomly selected, progressively increasing set of spatial frequencies. Unlike diffusion priors that start from a heavily degraded state, the FDB uses a moderately undersampled starting point to enhance the reverse diffusion sampling process. Furthermore, unlike existing diffusion bridges that degrade data based on a weighted average of the start and end points, the FDB uses a binary removal of k-space points, aligning more closely with the nature of accelerated MRI acquisitions. To further enhance image quality, the FDB employs a novel sampling algorithm based on a learned correction term, enabling soft dealiasing by continuously refining estimates of the recovered k-space data during reverse diffusion steps. Tests on brain MRI show that the FDB outperforms competing non-diffusion priors by 4.8 dB PSNR and 11.8% SSIM, and diffusion priors by 4.7 dB PSNR and 6.6% SSIM.

Keywords: diffusion, bridge, generative, deep learning, MRI, reconstruction.

ÖZET

MRG REKONSTRÜKSİYONU İÇİN DİFÜZYON KÖPRÜLERİ

Muhammad Usama Mirza
Elektrik ve Elektronik Mühendisliği, Yüksek Lisans
Tez Danışmanı: Tolga Çukur
Ağustos 2024

Manyetik Rezonans Görüntüleme (MRG) yeniden yapılandırması tipik olarak, az örneklenmiş verileri tam örneklenmiş verilere dönüştürmek için bir çözümleme sürecini içerir. Bununla birlikte, geleneksel yayılma öncelikleri, Gauss gürültüsü durumundan başlayıp tam örneklenmiş verilerle biten, gürültü giderici bir dönüşüm gerçekleştirir. Birçok k-uzay örnekleme deseniyle ilişkili örtüşme yapay yapıları, beyaz Gauss gürültüsünden önemli ölçüde farklı olan uzamsal yapılara sahip olduğundan, bu gürültü giderme işlemi, yapay yapıların optimum düzeyde bastırılması nedeniyle yeniden yapılandırma hatalarına yol açabilir. Bu sınırlamanın üstesinden gelmek için, MRG yeniden yapılandırması için ilk Fourier kısıtlı difüzyon köprüsünü (FDB) sunuyoruz. Görevden bağımsız yayılma önceliklerinden farklı olarak FDB, az örneklenmiş verilerle başlayıp tam örneklenmiş verilerle biten bir çözümleme dönüşümü gerçekleştirmek için özel olarak tasarlanmıştır. Başlangıç noktası, rastgele seçilmiş, giderek artan uzaysal frekanslar kümesini ortadan kaldıran yeni bir stokastik bozunma operatörü kullanılarak yaratılır. Ağır derecede bozulmuş bir durumdan başlayan difüzyon önceliklerinin aksine, FDB, ters difüzyon örnekleme sürecini geliştirmek için orta derecede yetersiz örneklenmiş bir başlangıç noktası kullanır. Ayrıca, verileri başlangıç ve bitiş noktalarının ağırlıklı ortalamasına dayalı olarak bozan mevcut difüzyon köprülerinden farklı olarak FDB, hızlandırılmış MRG edinimlerinin doğasına daha yakın bir şekilde uyum sağlayarak k-uzayı noktalarının ikili olarak kaldırılmasını kullanır. Görüntü kalitesini daha da artırmak için FDB, öğrenilmiş bir düzeltme terimini temel alan yeni bir örnekleme algoritması kullanır ve ters difüzyon adımları sırasında geri kazanılan k-uzayı verilerinin tahminlerini sürekli olarak iyileştirerek yumuşak düzeltmeyi mümkün kılar. Beyin MRG testleri, FDB'nin rakip difüzyon olmayan öncelikleri 4,8 dB PSNR ve 11,8% SSIM ile difüzyon önceliklerini ise 4,7 dB PSNR ve 6,6% SSIM ile geride bıraktığını

göstermektedir.

Anahtar sözcükler: difüzyon, köprü, üretken, derin öğrenme, MRG, yeniden yapılanma.

Acknowledgement

I would like to express my sincere gratitude to my advisor, Professor Tolga Çukur, for his invaluable guidance, support, and mentorship throughout this research. His expertise and insights have been instrumental in shaping this work. I am deeply grateful for his patience, encouragement, and unwavering belief in my abilities.

I would also like to thank the members of my thesis jury for their insightful comments, constructive feedback, and valuable suggestions. Their expertise and perspectives have significantly improved the quality of this research.

I am indebted to my lab mates for their friendship, collaboration, and support. Their contributions to the research process have been invaluable. I am particularly grateful for their willingness to share their knowledge and expertise, and for their constant encouragement and motivation.

Finally, I would like to thank my friends and family for their unwavering support and encouragement throughout my academic journey. Their love, belief, and patience have been instrumental in helping me overcome challenges and achieve my goals. I am especially grateful for their understanding and support during the demanding times of this research.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

MRI is a diagnostic powerhouse with exceptional soft-tissue contrast that suffers from long scan times. Acceleration via undersampled acquisitions helps lower operational costs and susceptibility to patient motion, albeit an ill-posed inverse problem must be solved to reconstruct images with the aid of image priors that improve problem conditioning [1]. Given their high sensitivity, image priors based on deep learning have become pervasive in MRI reconstruction over the years [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]. A prevalent framework in this domain employs task-specific priors that are directly trained to perform reconstruction, i.e., the priors capture a dealiasing transformation that maps undersampled to fully-sampled data [16, 17, 18, 19, 20, 21]. The dealiasing transformation is typically mediated by a conditional model that receives a least-squares reconstruction of undersampled data as input, and that aims to output an image that would be consistent with fully-sampled data [22, 23, 24, 25, 26]. While task-specific priors have been pervasive in the MRI reconstruction literature, they typically yield low representational diversity that compromises image fidelity [27, 28, 29].

An alternative framework that promises improved image fidelity trains task-agnostic priors to perform non-reconstruction tasks such as image generation [30]. These priors employ generative models to learn the distribution of fully-sampled data divorced from the imaging operator, (i.e., decoupled from the influence of coil

sensitivities and sampling patterns) [31, 32, 33]. Of particular promise, common diffusion priors capture a denoising transformation that maps Gaussian noise as the asymptotic start-point to fully-sampled data as the end-point on the diffusion process [34, 35]. This denoising transformation is mediated via a network-based recovery operator that receives a noisy image sample at a given time step to output a less noisy image sample. To reconstruct an image using a trained prior, sampling is initiated on a random noise image, and denoising projections via the recovery operator of the diffusion prior and data-consistency projections via the imaging operator of accelerated MRI scans are executed iteratively [36, 37, 38, 39, 40, 41]. Note that, depending on the undersampling pattern, aliasing artifacts can carry spatial structure that diverge significantly from spatially-incoherent noise [1]. This can render denoising projections task-irrelevant and hence ineffective in suppressing aliasing artifacts, compromising image quality [42, 43, 44].

To improve performance in inverse problems by increasing task relevance, several recent computer vision studies have considered cold/soft diffusion priors [45, 46] and diffusion bridges [47, 48, 49, 50]. Commonly, these methods train a prior to progressively map images degraded due to a known imaging operator (e.g., a large blur kernel) onto clean images reflecting ground truth. To derive intermediate samples during forward diffusion, cold diffusion priors granularize degradations from the imaging operator into multiple steps (e.g., a set of compact blur kernels), whereas soft diffusion priors incorporate additional degradations from Gaussian noise [51]. Yet, cold/soft diffusion priors characteristically employ an asymptotic start-point of severely degraded images similar to task-agnostic diffusion priors. This can introduce suboptimality in learning of the recovery operator due to large variability in the distribution of image samples across the diffusion process [52]. Furthermore, cold/soft diffusion priors can perform suboptimally when prompted to recover images from moderately degraded acquisitions due to mismatch between degradation levels in training versus testing [45, 46]. Indeed, a recent study on single-coil MRI reconstruction suggests that cold diffusion priors may not offer benefits in image quality over task-agnostic diffusion priors [53].

On the other hand, diffusion bridges can more effectively improve task relevance by employing a finite start-point of moderately degraded images, delimiting variability in the distribution of images samples across the diffusion process, and potential differences between training and test sets [47]. Yet, existing diffusion bridges perform forward diffusion via convex-constrained degradation operators, which express intermediate image samples in the diffusion process as a weighted linear average of degraded and clean images [49, 50, 54]. For MRI reconstruction tasks, weighted averaging of undersampled and fully-sampled data would cause all intermediate samples to have non-zero intensity across the entire k-space, causing intermediate samples to diverge significantly from the distribution of undersampled MRI acquisitions that exercise binary selection of k-space points [1]. Note that this would force the recovery operator to learn gradual amplification of relatively high-frequency k-space points as opposed to dealiasing transformations, which can restrict the performance of conventional diffusion bridges in MRI reconstruction.

Here, we introduce a novel Fourier-constrained diffusion bridge (FDB) [55] for improved performance in MRI reconstruction (Fig. 4.1). To enhance task relevance of the image prior, FDB learns a dealiasing transformation to map between undersampled and fully-sampled data. Unlike task-agnostic or cold/soft diffusion priors based on an asymptotic start-point of severely degraded images, FDB uses a finite start-point based on a moderate level of degradation (i.e., undersampling). For this purpose, FDB leverages a novel stochastic degradation operator to remove a random set of spatial frequencies per forward diffusion step, and schedules the removal with a peripheral-to-central k-space order from the end- to the start-point. Note that, unlike previous diffusion bridges based on convex-constrained degradation operators, FDB employs a Fourier-constrained operator that respects the binary k-space point selection in accelerated MRI acquisitions. To further improve reconstruction performance during inference, we introduce a novel sampling algorithm for FDB based on a learned correction term. Unlike conventional algorithms for generalized diffusion processes, the proposed algorithm enables soft dealiasing across reverse diffusion steps to continually update estimates of recovered k-space points. Comprehensive demonstrations on brain MRI datasets

indicate the superior performance of FDB against previous non-diffusion and diffusion priors. Code for FDB is available at `https://github.com/icon-lab/FDB`.

Contributions

FDB represents a pioneering diffusion bridge specifically designed for accelerated MRI reconstruction. Unlike existing methods, FDB employs a novel stochastic degradation operator that introduces random spatial-frequency removal, effectively bridging the gap between fully-sampled and moderately undersampled data. To further enhance reconstruction quality, FDB incorporates a novel sampling algorithm equipped with a learned correction term. This correction term is strategically scheduled across reverse diffusion steps, optimizing the reconstruction process through analytical derivations.

Outline

The rest of this thesis is organized as follows. Chapter 2 provides relevant background on MRI Reconstruction techniques. Chapter 3 describes the use of deep learning methods for MRI Reconstruction. Chapter 4 introduces the theory of the Fourier Constrained Diffusion Bridge. Chapter 5 lists the datasets and the competing methods used to test the proposed method, while chapter 6 depicts the results of the experiments. Chapters 7 and 8 discuss and conclude the work presented in this thesis.

Chapter 2

Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a non-invasive imaging technique that provides detailed visualizations of the body's internal organs and tissues without using ionizing radiation, making it a safer alternative to X-rays [56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70]. MRI is based on the principles of nuclear magnetic resonance, specifically the interaction between strong magnetic fields, radiofrequency (RF) pulses, and the hydrogen nuclei (protons) present in water molecules within the body [71, 72, 73, 74].

2.1 Fundamentals of MRI

The process begins by placing the patient in a strong, uniform magnetic field B_0, typically ranging from 1.5 to 7 Tesla [56, 75]. This magnetic field aligns the magnetic moments of the hydrogen nuclei along the direction of B_0. The net magnetization vector $\mathbf{M}$ can be described as:

$$\mathbf{M} = M_0 \hat{z}$$

where M_0 is the equilibrium magnetization and $\hat{z}$ is the direction of the main

magnetic field [57].

An RF pulse is then applied at the Larmor frequency ω_0, which is specific to the type of nucleus being imaged:

$$\omega_0 = \gamma B_0$$

Here, γ is the gyromagnetic ratio for hydrogen [71]. This RF pulse tips the magnetization vector $\mathbf{M}$ away from the $\hat{z}$-axis into the transverse plane.

Once the RF pulse is turned off, the magnetization vector begins to relax back to its equilibrium state, emitting RF signals during the process. The transverse component of the magnetization $\mathbf{M}_\perp(t)$ induces a time-varying voltage $V(t)$ in the receiver coil:

$$V(t) \propto \frac{d\mathbf{M}_\perp(t)}{dt}$$

This signal is known as the Free Induction Decay (FID) [72]. The measured FID signals are not directly in image space but rather in a frequency domain known as k-space, where each point $\mathbf{k}(t)$ corresponds to a spatial frequency of the object being imaged.

The relationship between the measured k-space data $S(\mathbf{k})$ and the spatial image $I(\mathbf{r})$ is governed by the Fourier transform [76]:

$$I(\mathbf{r}) = \int_{\mathbf{k}} S(\mathbf{k}) e^{i2\pi \mathbf{k}\cdot\mathbf{r}} d\mathbf{k}$$

where $\mathbf{k} = (k_x, k_y, k_z)$ is the k-space vector, and $\mathbf{r} = (x, y, z)$ represents the spatial coordinates in the image domain.

In practice, the MRI machine samples the k-space in a series of lines or spirals. The sampled k-space data is then inverse Fourier transformed to reconstruct the

spatial image:

$$I(x, y) = \mathcal{F}^{-1}\{S(k_x, k_y)\}$$

where $\mathcal{F}^{-1}$ denotes the inverse Fourier transform.

The quality of the reconstructed image $I(x, y)$ depends on the density and extent of the k-space sampling. Higher resolution and better signal-to-noise ratio (SNR) are achieved by more extensive k-space coverage. MRI excels in visualizing soft tissues, organs, and blood vessels with exceptional detail, making it an invaluable tool for diagnosing a broad spectrum of medical conditions, from neurological disorders to musculoskeletal injuries [77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87].

2.2 Prolonged Scan Times

One of the primary limitations of MRI is its extended scan times, which can vary significantly based on the body part being imaged and the desired spatial resolution. The total scan time T_{scan} is influenced by several key factors, including the repetition time T_R, the number of phase-encoding steps N_y, the number of signal averages N_{avg}, and the number of slices N_z. These parameters play crucial roles in determining the length of an MRI scan [56]:

- **Repetition Time (T_R)**: This is the time between successive pulse sequences applied to the same slice. It determines how long the system waits before re-exciting the same volume of tissue to acquire another signal. Longer T_R values allow more complete relaxation of the protons, which can improve image contrast but also increase scan time.
- **Number of Phase-Encoding Steps (N_y)**: Phase encoding is one of the steps used to localize signals in the spatial domain. The number of phase-encoding steps corresponds to the resolution of the image along the phase-encoding direction (usually the vertical direction in the image). Higher

N_y increases image resolution but also extends the scan time since each phase-encoding step requires a separate excitation.

- **Number of Signal Averages ($N_{\mathbf{avg}}$)**: Also known as the number of excitations (NEX), this refers to the number of times the same signal is acquired and averaged. Averaging improves the signal-to-noise ratio (SNR) but increases the scan time proportionally to the number of averages.
- **Number of Slices (N_z)**: In three-dimensional MRI or when acquiring multiple slices in a two-dimensional scan, N_z represents the number of slices through the body part being imaged. Increasing the number of slices provides more comprehensive coverage but also increases the scan time since each slice requires a separate acquisition.

The total scan time can be approximated by the formula:

$$T_{\text{scan}} \approx T_R \times N_y \times N_{\text{avg}} \times N_z$$

For example, in a typical brain MRI, T_R might be in the range of 500 milliseconds to several seconds, N_y could be several hundred steps for high-resolution imaging, N_{avg} is often set to 1 or 2, and N_z might range from 20 to 50 slices. This results in a scan time of around 30 minutes. For larger areas like the entire spine, where N_z is significantly higher, the scan time could extend to an hour or more.

These prolonged scan times introduce several challenges:

1. **Patient Comfort and Stress**: The extended duration can cause significant discomfort, especially in patients who must remain still for the entire scan. This can be particularly challenging for children or patients with claustrophobia. Motion during the scan can lead to artifacts, necessitating retakes and further increasing the total scan time.
2. **Emergency Situations**: In cases of acute medical conditions such as stroke or trauma, rapid diagnosis is essential. The long scan times inherent

to MRI can delay critical treatment, as the time required to acquire high-resolution images may be too long for urgent decision-making. For example, in stroke management, where every minute counts, delays in obtaining MRI results can adversely affect patient outcomes.

3. **Reduced Patient Throughput**: The length of MRI scans also limits the number of patients that can be scanned in a given time period. This reduced throughput can lead to longer wait times for MRI appointments, delaying diagnosis and treatment. Additionally, increased scan times can contribute to higher operational costs for healthcare providers, as fewer patients can be accommodated within standard operating hours.

Efforts to reduce MRI scan times focus on optimizing imaging protocols and developing advanced techniques such as parallel imaging, compressed sensing, and machine learning-based reconstruction [88, 89, 90, 91, 92, 93, 94, 95]. These methods aim to decrease the number of phase-encoding steps N_y, the number of slices N_z, or reduce the number of averages N_{avg} required, all while maintaining image quality. By doing so, the total scan time T_{scan} can be minimized.

Reducing scan times poses significant challenges, as acquiring less data typically leads to a reduction in image quality. This is due to the fact that the detailed information needed to construct a clear and accurate image is compromised when fewer data points are collected. Mathematically, this trade-off can be expressed by the following relationship:

$$T_{\text{scan}} \approx T_R \times \left(\frac{N_y}{R}\right) \times N_{\text{avg}} \times N_z$$

Here, R is the acceleration factor introduced by advanced techniques like parallel imaging, which aims to reduce the effective number of phase-encoding steps N_y. The goal is to find an optimal balance that provides sufficient diagnostic information while minimizing the scan duration [96, 97, 98].

2.3 Techniques for Accelerating MRI Acquisition

2.3.1 Parallel Imaging

Parallel imaging techniques, such as SENSE (Sensitivity Encoding) [99] and GRAPPA (Generalized Auto-Calibrating Partially Parallel Acquisition) [100], employ multiple receiver coils to gather data simultaneously. Each coil has a distinct sensitivity profile, allowing the MRI system to capture different parts of the k-space in parallel. This reduces the number of phase-encoding steps N_y needed, thereby shortening the scan time. The image is then reconstructed by combining the data from these coils:

$$S(\mathbf{k}) = \sum_{j=1}^{N_c} C_j(\mathbf{r}) \cdot I_j(\mathbf{r})$$

where $S(\mathbf{k})$ represents the k-space data, $C_j(\mathbf{r})$ is the sensitivity of the j-th coil, and $I_j(\mathbf{r})$ is the image data acquired by that coil. Despite its effectiveness, parallel imaging can suffer from artifacts if the coil sensitivities are not accurately known or if the signal-to-noise ratio (SNR) is low.

2.3.2 Compressed Sensing

Compressed sensing leverages the inherent sparsity of MRI data in a specific transform domain, such as the wavelet domain [101, 102, 103, 104, 105, 106]. This method assumes that the image can be represented with a limited number of non-zero coefficients. By acquiring fewer data points than traditionally required, compressed sensing aims to reconstruct the full image using the sparse representation:

$$\min_{I(\mathbf{r})} \|I(\mathbf{r})\|_1 \text{ subject to } \|S(\mathbf{k}) - \mathcal{F}(I(\mathbf{r}))\|_2^2 \leq \epsilon$$

where $\|I(\mathbf{r})\|_1$ is the ℓ_1-norm promoting sparsity, $\mathcal{F}$ is the Fourier transform operator, and ϵ is a small error tolerance. Although compressed sensing can significantly reduce scan times, its success depends on the validity of the sparsity assumption, which may not always hold for all types of images [1]. Additionally, the computational demands of solving the optimization problem can be substantial.

2.3.3 Deep Learning-Based Reconstruction

In recent years, deep learning has emerged as a powerful tool in the field of MRI [107, 108, 40, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134, 135, 136]. It is a subset of machine learning that uses artificial neural networks with multiple layers to learn complex patterns from data. In the context of MRI reconstruction, deep learning models are trained on large datasets of paired undersampled and fully-sampled images. These models learn to identify and extract relevant features from the undersampled data, allowing them to accurately reconstruct the missing information. By leveraging the power of deep learning, these models can overcome the limitations of traditional reconstruction techniques and produce significantly improved results. By training deep neural networks on large datasets of paired undersampled and fully-sampled images, these models can learn complex relationships between the two, enabling them to reconstruct high-quality images from undersampled data.

Chapter 3

Deep Learning Methods for MRI Reconstruction

Reconstruction aims to recover a high-quality MR image x of an anatomy given undersampled k-space acquisition y:

$$Ax + \varepsilon = y, \tag{3.1}$$

where A is the imaging operator, ε is measurement noise. $Ax{=}M\mathcal{F}(Cx)$ comprises the effects of the sampling pattern M, Fourier encoding $\mathcal{F}$, and coil sensitivities C. Reconstruction is an ill-posed inverse problem, so a regularization term is typically employed to obtain a faithful solution $\hat{x}$ [1]:

$$\hat{x} = \min_{x} \|Ax - y\|^2 + G(x, y), \tag{3.2}$$

where $G(x, y)$ is the regularizer that incorporates prior information on the distribution of MR images.

3.1 Task-Specific Priors

Deep learning frameworks have empowered leaps in image quality over traditional methods for accelerated MRI reconstruction. A mainstream approach relies on

conditional models that are trained to map undersampled acquisitions as input onto fully-sampled acquisitions as output [137, 138, 139, 140]. The resultant task-specific priors capture a dealiasing transformation particular to the prescribed imaging operator [22, 141]. Task-specific priors are typically expressed as a conditional regularizer $G(x|y)$ that captures a dealiasing transformation:

$$\hat{x} = \min_{x} \|Ax - y\|^2 + G(x|y), \tag{3.3}$$

Image samples are generated using the recovery operator to predict $x = G_\theta(y)$. The training objective for task specific priors is then given as:

$$L_{TSP} = \mathbb{E}_{x,y}[\|G_\theta(y) - x\|^2], \tag{3.4}$$

Despite their prominence, many task-specific priors produce a deterministic output limiting their representational capacity [142, 143, 144], and they might be rendered ineffective under domain shifts in imaging operators due to on-the-fly changes during scan sessions to meet clinical demand [30, 145].

3.2 Task-Agnostic Diffusion Priors

To improve generalization, an alternative approach employs task-agnostic priors based on generative models that transform random latent variables onto high-quality MR images decoupled from the imaging operator [31, 146, 147]. Thus, task-agnostic priors are instead expressed as a marginal regularizer $G(x)$ independent of the imaging operator to improve generalization:

$$\hat{x} = \min_{x} \|Ax - y\|^2 + G(x), \tag{3.5}$$

To reconstruct undersampled data, the image sets spanned by the prior and the imaging operator can be consolidated via alternating projections [148, 149]. This consolidation naturally rests on the representational capacity of the prior [40]. Diffusion priors have recently been adopted given their superior representation, enabled via a stochastic process that gradually transforms between Gaussian noise and fully-sampled data [150, 34, 35, 151].

Common task-agnostic diffusion priors degrade a clean image x_0 over T steps in the forward direction to obtain an isotropic noise sample x_T. The sample x_t at step $t \in (0\ T]$ is [150]:

$$x_t = \sqrt{1-\beta_t}x_{t-1} + \sqrt{\beta_t}z, \tag{3.6}$$

where β_t is noise variance, $z \sim \mathcal{N}(0, \mathrm{I})$ follows a Gaussian distribution, I is an identity matrix. In the reverse direction, image samples are generated via repeated denoising. One approach uses recovery operator to predict $\tilde{x}_0 = G_\theta(\hat{x}_t, t)$, and draw a denoised sample $\dot{x}_{t-1}$ from the normal distribution $q(x_{t-1}|x_t, \tilde{x}_0)$ [152]. Training objective for G_θ operationalizing the task-agnostic diffusion prior is then given as:

$$L_{TADP} = \mathbb{E}_{t,x_{(0,t)}}[\|G_\theta(x_t, t) - x_0\|^2], \tag{3.7}$$

where $\mathbb{E}$ is expectation, $t \sim U[1, T]$ follows a uniform distribution, $x_0, x_t \sim q(x_0, x_t)$ are generated via the forward process. For MRI reconstruction, starting with a random noise sample $\hat{x}_T \sim \mathcal{N}(0, \mathrm{I})$, reverse diffusion steps [152] can be interleaved with data-fidelity projections [38]:

$$\dot{x}_{t-1} = \frac{\sqrt{\gamma_t}(1-\bar{\gamma}_{t-1})}{1-\bar{\gamma}_t}\hat{x}_t + \frac{\beta_t\sqrt{\bar{\gamma}_{t-1}}}{1-\bar{\gamma}_t}\tilde{x}_0 + \sqrt{\frac{1-\bar{\gamma}_{t-1}}{1-\bar{\gamma}_t}\beta_t}z, \tag{3.8}$$

$$\hat{x}_{t-1} = \dot{x}_{t-1} + A^H(y - A\dot{x}_{t-1}), \tag{3.9}$$

where $\gamma_t = 1 - \beta_t$, $\bar{\gamma}_t = \prod_{\tau=[0,1,..,t]} \gamma_t$, and A^H is the Hermitian adjoint of the imaging operator A.

While promising results have been reported in MRI reconstruction [36, 37, 38], task-agnostic diffusion priors sample the asymptotic start-point of the diffusion process from a white Gaussian noise distribution, and perform progressive denoising transformations to reconstruct images [150]. These task-irrelevant denoising transformations can diverge significantly from the desired transformation in MRI reconstruction from undersampled to fully-sampled data, occasionally compromising efficiency and performance [43, 44].

To address this limitation, here we introduce the first diffusion bridge for MRI reconstruction in the literature to our knowledge. The proposed FDB prior leverages a stochastic degradation operator that randomly discards spatial-frequency

components from fully-sampled data, starting in peripheral k-space and gradually progressing towards central k-space. Unlike task-agnostic diffusion priors that take Gaussian noise as their asymptotic start-point and then perform denoising transformations, FDB takes moderately undersampled data as its finite start-point and then performs progressive dealiasing transformations. These unique technical attributes enable FDB to capture task-relevant information in its diffusion prior for performant MRI reconstruction.

3.3 Cold/Soft Diffusion Priors

Improvements to task-agnostic diffusion priors have recently been sought to increase flexibility in generative modeling. Computer vision studies have proposed cold diffusion priors built on deterministic degradation operators such as blurring or downsampling in lieu of noise addition [45, 51], and soft diffusion priors that use hybrid operators that combine known deterministic degradations with Gaussian noise addition [46]. A recent study on single-coil MRI reconstruction has adopted a cold diffusion prior based on k-space undersampling [53].

A degradation operator D is defined to obtain a noisy image from the clean image:

$$D(x_0, t) = x_t \tag{3.10}$$

A restoration operator is defined to reverse the degradation operator and recover the clean image:

$$G(x_t, t) \approx x_0 \tag{3.11}$$

The learning objective for the restoration operator is then defined as:

$$L_{CD} = \mathbb{E}_x[\|G_\theta(D(x, t), t) - x_0\|^2] \tag{3.12}$$

Despite benefits in flexibility and sampling speed, comparable image quality has been generally reported against task-agnostic diffusion priors [45, 53]. Note that the deterministic degradations in cold/soft diffusion priors are often assumed to

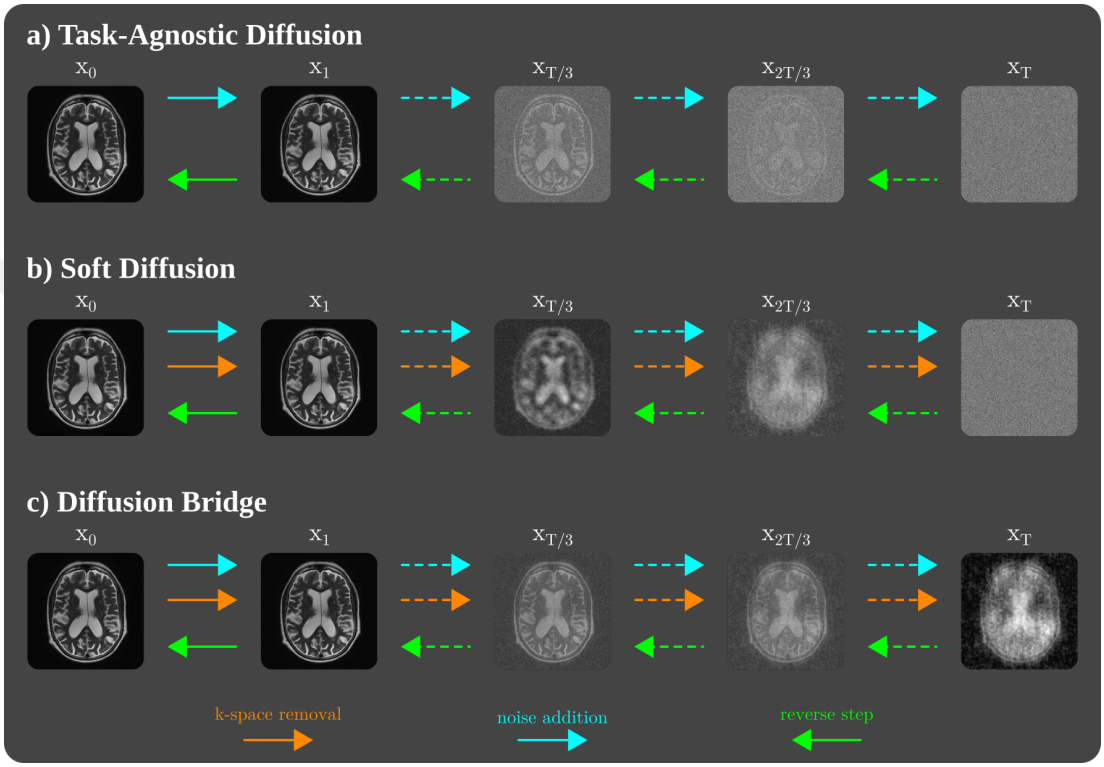

Figure 3.1: Illustration of diffusion modeling frameworks in the context of building priors for MRI reconstruction. **a)** Task-agnostic diffusion priors use a forward degradation operator that adds random noise onto a clean image (x_0), reaching an asymptotic start-point of isotropic Gaussian noise (x_T). A recovery operator then performs denoising transformations to map x_T back onto x_0. **b)** Soft diffusion priors use a degradation operator that simultaneously adds random noise and applies task-related degradations (i.e., undersampling in the case of MRI reconstruction). As aggressive noise scheduling is used towards T, an asymptotic start-point of Gaussian noise is reached similar to the case of task-agnostic diffusion priors. The recovery operator has to learn a joint denoising-dealiasing transformation. **c)** Conventional diffusion bridges use a degradation operator that adds random noise onto a weighted linear average of start- and end-points (i.e., as a weighted average of fully-sampled and undersampled data in the case of MRI reconstruction). While the start-point is undersampled data, intermediate image samples near the mid-point of the diffusion process are corrupted by high levels of noise. The recover operators has to learn a joint denoising-dealiasing transformation.

remain fixed between training-test sets, reducing the prior's stochasticity and limiting generalization [45]. They also employ asymptotic start-points at high degradation levels (e.g., over 100-fold k-space undersampling) [53], which can limit performance by elevating difficulty in reverse diffusion sampling [52]. Unlike cold/soft diffusion priors, FDB employs a stochastic degradation operator based on random frequency removal to enhance representational diversity, and a finite start-point at a moderate degradation level to facilitate image recovery during reverse diffusion sampling. In comparison to soft diffusion priors that still rely on task-irrelevant Gaussian noise addition in their forward process, FDB solely leverages task-relevant undersampling degradations for enhanced performance in MRI reconstruction.

3.4 Diffusion Bridges

An emerging alternative in generative modeling is diffusion bridges that aim to learn a transformation between two arbitrary distributions [47]. Several computer vision studies have recently adopted diffusion bridges to solve inverse problems by mapping between degraded images due to the imaging operator taken as start-point and clean images taken as end-point [48, 49]. Primarily devised for deterministic degradations, existing diffusion bridges derive intermediate image samples across the diffusion process as a weighted linear average of start- and end-points (i.e., a convex combination of degraded and clean images) [49, 50, 153], and add varying levels of Gaussian noise on the combination to introduce stochasticity that is essential for diffusion modeling [47]. This forward-process design elicits several limitations in common diffusion bridges. First, the use of Gaussian noise in the diffusion process forces the bridge to learn an auxiliary denoising transformation, which lowers the task relevance of the image prior. Second, linear averaging of undersampled and fully-sampled data produces intermediate samples with non-zero signal intensity across k-space, which is discrepant with the nature of accelerated MRI acquisitions that remove k-space points via binary selection. To avoid use of denoising transformations, FDB induces stochasticity in the forward diffusion process by removing a random set of frequencies across

k-space, and this set is also randomly varied across the training set. Instead of convex-constrained degradations in common diffusion bridges that result in linear averaging, FDB leverages a Fourier-domain constraint by removing a monotonically growing set of frequencies across forward steps, and this constraint closely matches the practical implementation of accelerated MRI acquisitions.

Chapter 4

Fourier-Constrained Diffusion Bridges

The proposed FDB prior constructs a novel diffusion process with moderately undersampled data taken as its finite start-point, while fully-sampled data constitute the end-point. In the forward direction, a stochastic degradation operator in Fourier domain removes a random, non-overlapping set of frequency components in each forward step (Fig. 4.1). In the reverse direction, a recovery operator in image domain operationalized as a neural network imputes a distinct set of frequency components in each reverse step. We derive the training objective of this recovery operator based on the interpretation that FDB's forward steps are equivalent to a generalized diffusion process with linear degradations in image domain [46]. For the first time in literature, we introduce a continually-corrected sampling algorithm that incorporates a learned correction term to mediate soft dealiasing of frequency components across reverse steps, and we present an analytical derivation of the optimal scheduling for the weighting of this correction term. Finally, we outline the inference procedures for MRI reconstruction on test data.

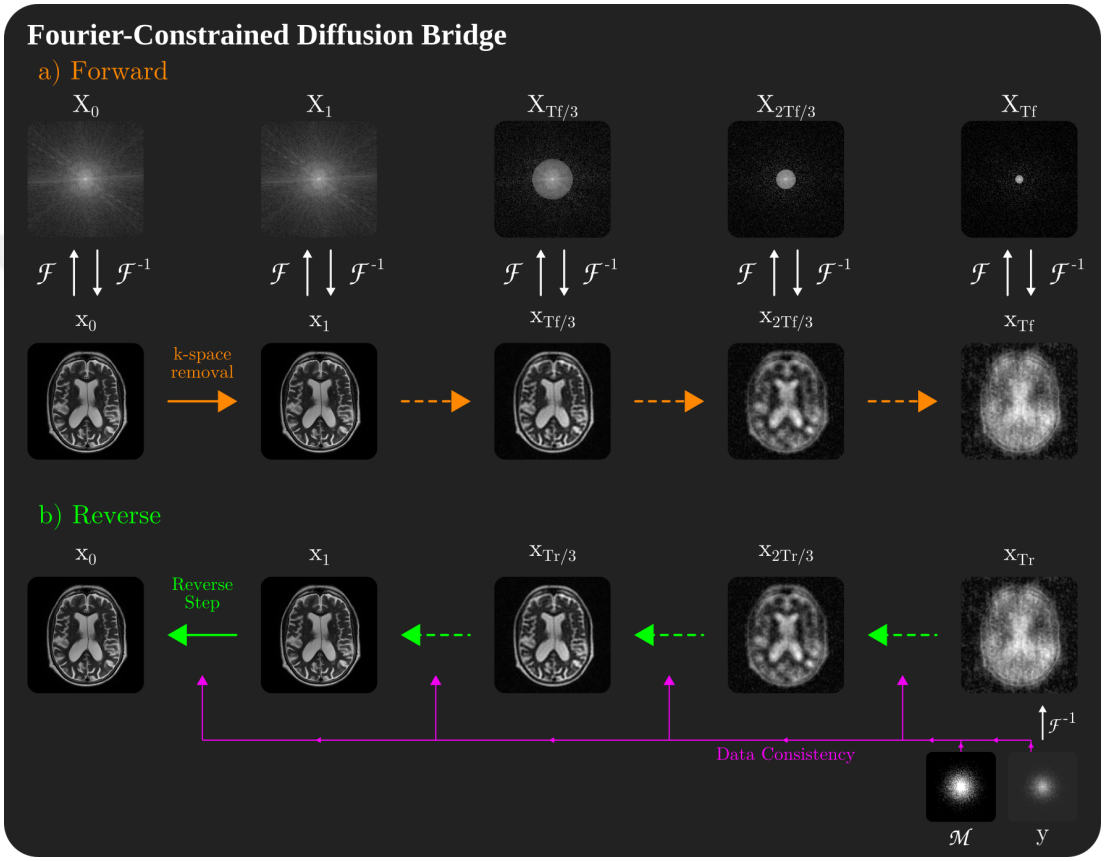

Figure 4.1: **a)** The proposed Fourier-constrained diffusion bridge (FDB) prior. FDB leverages a stochastic degradation operator based on frequency removal. This operator maps a clean image ($x_0 = \mathcal{F}^{-1}\{X_0\}$ where X_0 denotes the Fourier transform of x_0) onto a finite start-point corresponding to a least-squares reconstruction of R'-fold undersampled data ($x_{T_f} = \mathcal{F}^{-1}\{X_{T_f}\}$ where X_{T_f} denotes the Fourier transform of x_{T_f}). Note that, unlike conventional diffusion bridges, FDB does not rely on a weighted averaging of start- and end-points, and it does not use noise addition in forward steps. The recovery operator of FDB performs purely task-relevant dealiasing transformations. **b)** To reconstruct an R-fold undersampled acquisition y, sampling is initiated with the least-squares solution $x_{T_r} = A^H y$, where A^H denotes the Hermitian adjoint of the imaging operator A, and $T_r = \lfloor T_f \frac{(R-1)R'}{(R'-1)R} \rfloor$. Reverse steps with FDB for frequency imputations are interleaved with data-consistency projections.

4.1 Diffusion Process

FDB introduces a novel diffusion process based on binary removal/recovery of spatial-frequency components in MR images. Since the design of this process is key to the performance improvements enabled by FDB, here we highlight the motivations for adopting the unique design elements, preceding their detailed description. Differing from Gaussian-noise-based processes in task-agnostic diffusion priors previously considered for MRI reconstruction, FDB maps between a start-point X_{T_f} at time step T_f taken as coil-combined data degraded at level R', and an end-point $X_0 = \mathcal{F}(S^H \mathcal{F}^{-1}(y_{full}))$ taken as coil-combined Fourier-domain data derived from fully-sampled acquisitions. This start-point selection enables FDB to learn a task-relevant dealiasing transformation to improve reconstruction performance, as opposed to the denoising transformations offered by task-agnostic diffusion priors.

FDB shares a similar inspiration to recent cold/soft diffusion priors and diffusion bridges proposed to increase the level of task-relevant information embodied within the diffusion process. Yet, there are critical technical differences between FDB and these methods. Unlike cold/soft diffusion priors that utilize a start-point at very high degradation levels analogous to task-agnostic diffusion priors, FDB utilizes a reserved R' value corresponding to moderate k-space undersampling. The moderate degradation level at the start-point serves to lower difficulty in reverse diffusion mappings and thereby improve learning of the recovery operator.

Existing diffusion bridges commonly express intermediate image samples across the diffusion process as a linear average of start- and end-points. When the start-point is undersampled data, this would cause the intermediate samples to have non-zero intensity for all frequency components in k-space. In turn, the recovery operator would learn to gradually amplify missing high-frequency components, which diverges from the dealiasing transformations necessary for MRI reconstruction. In contrast, FDB exercises binary removal of random spatial-frequency components in the forward direction, such that a dealiasing transformation is

learned at each reverse step. Note that this degradation operator closely mimics the practical implementation of k-space undersampling in accelerated MRI.

In the remainder of this section, the mathmetical formulation of FDB's degradation and recovery operators are detailed. In the forward direction, the degradation operator is implemented in Fourier domain due to ease of frequency removal via multiplication with binary matrices. Yet, in the reverse direction, the recovery operator is implemented in image domain to avoid potential biases due to divergent energy levels at low versus high frequency components in k-space [1].

Degradation operator. For a two-dimensional Cartesian k-space grid with N^2 frequency components (where N is the number of grid points along each dimension), let $k \in [1\ N^2]$ denote linear component index, $r(k) \in [0\ r_{\max}]$, and $\phi(k) \in [0\ 2\pi)$ denote the k-space radius and angle of the kth component. In a forward step, the degradation operator is cast as a diagonal frequency-removal matrix $\mathbf{\Lambda}_t \in \mathbb{R}^{N^2 \times N^2}$. At step t, $\mathbf{\Lambda}_t(k,k)|_{k \in S_t} = 0$ for n components in the randomly selected set S_t, and $\mathbf{\Lambda}_t(k,k)|_{k \notin S_t} = 1$. To improve the prior via monotonically increasing degradation across time [48], we observe that components selected at t should not overlap with previously selected components at $\{t-1, .., 1\}$. To respect the energy distribution of MRI data, we observe that the selection should follow a peripheral-to-central k-space order as typical in accelerated MRI [1]. We enforce these selection guidelines via a spatio-temporal point process over k-space and t:

$$S_t = \Big\{(k_1, .., k_n) : \quad k_i \sim U[1, N^2] \ \text{ for } i = (1, .., n), \\ k_i \notin \bigcup\nolimits_{\tau=1}^{t-1} S_\tau, \ \ r(k_i) > \bar{r}_t \Big\}, \tag{4.1}$$

where U is uniform distribution, r_{max} is the maximum k-space radius, and $\bar{r}_t$ is a radius threshold scheduled as:

$$\bar{r}_t = r_{\max} - r_{\max}(1 - \sqrt{R'})(t/T_f). \tag{4.2}$$

Monotonically lowering $\bar{r}_t$ to promote peripheral-to-central ordering towards the start-point, this scheduling ensures a degradation level of R' at $t = T_f$ assuming that the number of components removed per step is $n = \lfloor N^2(R'-1)/(R'T_f) \rfloor$.

Once random removal masks are determined according to the point process in Eq. 4.1, the relationship between X_t and X_0 in Fourier domain is given as:

$$X_t = \left(\prod_{\tau=1}^{t} \mathbf{\Lambda}_\tau \right) X_0 = \bar{\mathbf{\Lambda}}_t X_0, \tag{4.3}$$

where $\bar{\mathbf{\Lambda}}_t \in \mathbb{R}^{N^2 \times N^2}$ is the cumulative frequency removal matrix at t (i.e., the product of removal matrices for {t,...,1}), and $X_t, X_0 \in \mathbb{C}^{N^2 \times 1}$ are column-vectorized forms of respective complex Fourier-domain data.

Recovery operator. To build the recovery operator, we express the relationship in Eq. 4.3 in image domain as:

$$x_t = \mathcal{F}^{-1}\big(\bar{\mathbf{\Lambda}}_t X_0\big) = \mathcal{F}^{-1}\big(\mathrm{diag}(\bar{\mathbf{\Lambda}}_t) \otimes X_0\big), \tag{4.4}$$

$$= \mathcal{F}^{-1}\big(\mathrm{diag}(\bar{\mathbf{\Lambda}}_t)\big) \circledast \mathcal{F}^{-1}\big(X_0\big), \tag{4.5}$$

$$= \bar{\kappa}_t \circledast x_0. \tag{4.6}$$

In Eq. 4.4, x_t=$\mathcal{F}^{-1}(X_t)$ denotes the image sample at t, diag$(\cdot)$ extracts the diagonal entries of $\bar{\mathbf{\Lambda}}_t$ as a column vector, and $\otimes$ is the Kronecker product. In Eq. 4.5, $\circledast$ denotes circular convolution in image domain. In Eq. 4.6, $\bar{\kappa}_t \in \mathbb{C}^{N^2 \times 1}$ is the image-domain kernel that captures the effects of frequency removal with $\bar{\mathbf{\Lambda}}_t$, and x_0=$\mathcal{F}^{-1}(X_0)$ denotes the image sample at 0. Based on linear systems theory, we note that the convolution operation in Eq. 4.6 can also be implemented via a matrix multiplication in image domain [154]:

$$x_t = D_t x_0 \tag{4.7}$$

where $D_t \in \mathbb{C}^{N^2 \times N^2}$ is a doubly block circulant matrix (i.e., a special case of Toeplitz matrices). To construct D_t using the matrix form of $\bar{\kappa}_t$, i.e., $K \in \mathbb{C}^{N \times N}$, a set of circulant matrices $b_1, ..., b_N \in \mathbb{C}^{N \times N}$ are created from individual rows of K:

$$b_i = \begin{bmatrix} K_{i,1} & K_{i,N} & \dots & K_{i,2} \\ K_{i,2} & K_{i,1} & \dots & K_{i,3} \\ \vdots & \vdots & \ddots & \vdots \\ K_{i,N} & K_{i,N-1} & \dots & K_{i,1} \end{bmatrix}, \tag{4.8}$$

where $i \in [1\ N]$ denotes row index, and $K_{i,:}$ is the ith row of K. The initial set of matrices are then organized in block circulant form to obtain D_t:

$$D_t = \begin{bmatrix} b_1 & b_N & \dots & b_2 \\ b_2 & b_1 & \dots & b_3 \\ \vdots & \vdots & \ddots & \vdots \\ b_N & b_{N-1} & \dots & b_1 \end{bmatrix}. \tag{4.9}$$

Based on the equivalence between Eqs. 4.6 and 4.9, we then observe that FDB can be viewed as a generalized diffusion process with linear degradations [46]:

$$x_t = D_t x_0 + \sigma_t z, \tag{4.10}$$

such that $\sigma_t = 0$ (i.e., no noise included). This interpretation enables us to express the training objective of FDB's recovery operator as follows:

$$L_{\text{FDB}} = \mathbb{E}_{t, x_{(0,t)}}[\|D_t(G_\theta(x_t, t) - x_0)\|^2], \tag{4.11}$$

where $t \sim U(1, T_f)$. Since $\bar{\mathbf{\Lambda}}_t$ is diagonal with maximum eigenvalue 1, and $\mathcal{F}^{-1}$, $\mathcal{F}$ are orthonormal transformations:

$$\|D_t\| = \|\mathcal{F}^{-1}\bar{\mathbf{\Lambda}}_t\mathcal{F}\| \leq \|\mathcal{F}^{-1}\|\|\bar{\mathbf{\Lambda}}_t\|\|\mathcal{F}\| = 1. \tag{4.12}$$

Further noting that $\|D_t(G_\theta(\cdot) - x_0)\| \leq \|D_t\|\|(G_\theta(\cdot) - x_0)\|$, FDB can be trained via a simplified upper-bound loss [45]:

$$L_{\text{FDB}-ub} = \mathbb{E}_{t, x_{(0,t)}}[\|(G_\theta(x_t, t) - x_0)\|^2]. \tag{4.13}$$

4.2 Sampling Algorithm

Based on the theory of diffusion modeling, the discrete forward process in Eq. 4.10 can be seen as the solution of a continuous-time stochastic differential equation (SDE) [155]:

$$\mathrm{d}x_t = (\dot{D}_t\alpha)\mathrm{d}t \tag{4.14}$$

where α denotes a single image sampled from $q(x_0)$. During reverse sampling, this SDE can be solved numerically via the Euler-Maruyama method where continuous derivatives are approximated as finite differences [151]:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha, \tag{4.15}$$

where $D_0 = \mathrm{I}$. In general, α will be unknown during reverse sampling, albeit an estimate of x_0 can be obtained via the recovery operator as a substitute: $\alpha = \{\tilde{x}_0(t) := G_\theta(x_t, t)\}$ [46]. Note that adopting the standard sampling procedure in Eq. 4.15 for FDB yields incremental imputation of missing frequency components between consecutive time steps. Here, we observe that this imputation performs hard dealiasing by only recovering n frequencies selected in the set S_t, without altering the value of previously recovered frequencies in the aggregated set $S_{pre} = \bigcup_{\tau=T_f}^{t+1} S_\tau$. Assuming that a frequency component k is imputed at step t_k, the recovered value is based on an imperfect estimate $\tilde{x}_0(t_k) := G_\theta(x_{t_k}, t_k)$, so it carries residual errors. Yet, based on Eq. 4.15, the value of k will not get updated at any other time step during $t < t_k$ or $t > t_k$, so these residual errors cannot be corrected. In turn, the updates computed based on Eq. 4.15 cannot leverage the benefits of iterated Langevin sampling since they deviate from solving the true SDE that would be obtained when $\alpha = x_0$.

Continually-corrected sampling. To address this fundamental limitation, here we propose a novel continually-corrected sampling algorithm for FDB by incorporating a correction term that enables soft dealiasing during reverse steps. In particular, we continue to check for residual errors on the previously imputed set of frequencies at each time step, and update their values to lower errors:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha + w_t D_t(\alpha - x_t), \tag{4.16}$$

where $w_t \in [0\ 1]$ is a weighting parameter, and the correction term effectively examines the difference $(\tilde{x}_0(t) - x_t)$ to revise the values of frequency components imputed in previous steps $\{T_f, .., t\}$. In theory, assuming that $\tilde{x}_0(t)$ is a good approximation to x_0, the correction term does not change the continuous-time SDE that is solved during reverse diffusion. This can be observed by setting

Algorithm 1: MRI reconstruction with FDB

Input:
y: MRI acquisition with acceleration rate R
A: Imaging operator
$G_\theta(x_t, t)$: Recovery operator
T_f: Number of diffusion steps
R': Degradation level at start-point
$\bar{w}_t$: Correction weight at time step t
Output:
$\hat{x}_0$: Reconstructed image

$T_r = \lfloor T_f \frac{(R-1)R'}{(R'-1)R} \rfloor$ ▷ number of reconstruction steps
$\hat{x}_{T_r} = A^H y$ ▷ zero-filled reconstruction
$\bar{w}_{\{1,..,T_r\}} = \text{resample}(w_{\{1,..,T_f\}}, T_r/T_f)$
for $t = T_r, \ldots, 1$ **do**
  $\tilde{x}_0(t) = G_\theta(\hat{x}_t, t)$ ▷ estimate 'clean' image
  if $t \neq 1$ **then**
    Sample D_{t-1} by drawing S_{t-1}, given D_t and S_t
  else
    $D_{t-1} = \text{I}$
  $\dot{x}_{t-1} = \hat{x}_t + (D_{t-1} - D_t)\tilde{x}_0(t) + \bar{w}_t D_t(\tilde{x}_0(t) - x_t)$
  $\hat{x}_{t-1} = \dot{x}_{t-1} + A^H(y - A\dot{x}_{t-1})$
return $\hat{x}_0$

$\alpha = x_0$ for the correction term in Eq. 4.16 and noting that $x_t = D_t x_t$:

$$\begin{aligned} w_t D_t(\alpha - x_t) &\approx w_t D_t(x_0 - x_t), && (4.17)\\ &= w_t(D_t x_0 - D_t x_t), && (4.18)\\ &= w_t(x_t - x_t) = 0. && (4.19) \end{aligned}$$

In practice, the correction term can alter the original SDE to a degree depending on both the accuracy of the approximation $\tilde{x}_0(t) \approx x_0$, and the value of w_t. To understand the nature of alterations, we can reorganize the terms in Eq. 4.16, again using $x_t = D_t x_t$ for simplification:

$$x_{t-1} - x_t = (D_{t-1} - D_t)\alpha + w_t D_t \alpha - w_t D_t x_t, \quad (4.20)$$

$$x_{t-1} - (1 - w_t)x_t = (D_{t-1} - D_t(1 - w_t))\alpha. \quad (4.21)$$

When $w_t \to 0$, Eq. 4.20 converges to iterated Langevin sampling based on the original SDE formulation; and when $w_t \to 1$, it starts ignoring the image sample at step t to give instantaneous updates based on $\alpha = \tilde{x}_0(t)$. As we show in the

next subsection, the optimal w_t schedule for FDB is attained by prescribing low w_t near step $t = T_f$ (i.e., start-point), and high w_t near step $t = 0$ (i.e., end-point). With this scheduling, the proposed algorithm performs iterated Langevin sampling near the start-point where $\tilde{x}_0(t)$ is a poor estimate of x_0, hence ignoring the correction term. Contrarily, the algorithm transitions into performing instantaneous estimates solely based on $\tilde{x}_0(t)$ near the end-point where it starts closely approximating x_0, hence emphasizing the correction term.

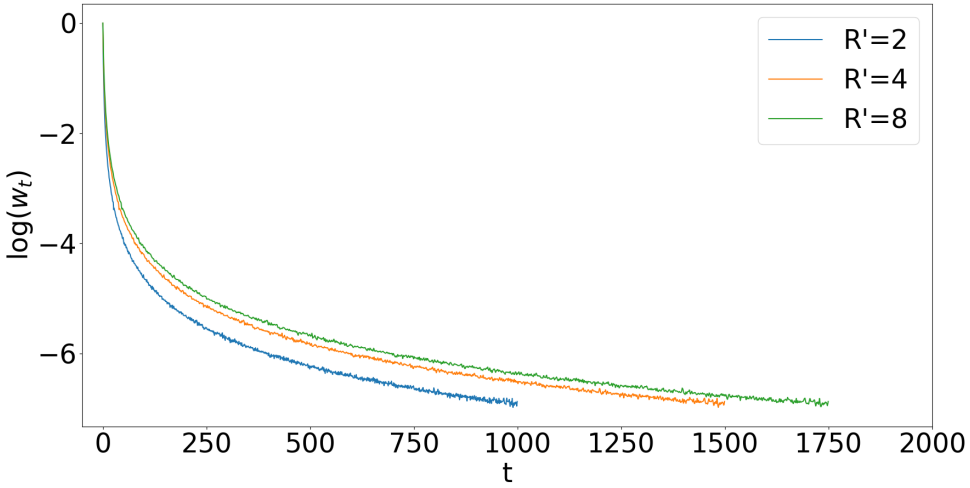

Figure 4.2: The weighting parameter w_t for the correction term in Eq. 4.16 was learned as described in Eq. 4.27. Results shown for R'=2 (T_f=1000), R'=4 (T_f=1500), and R'=8 (T_f=1750) on the IXI dataset.

Analytical derivation of the optimal w_t. Compared to earlier iterations (i.e., near T_f) where the set of recovered frequencies is compact, the proposed correction becomes more critical in later iterations as the set grows in size and the accuracy of $\tilde{x}_0(t)$ increases. Thus, we adopted a learned schedule for w_t to enhance recovery. Observing that $x_t = D_t x_t$, Eq. 4.16 can be rearranged:

$$x_{t-1} = (D_{t-1} - D_t)\alpha + D_t[w_t\alpha + (1 - w_t)x_t]. \tag{4.22}$$

The second term in Eq. 4.22 sets the values of previously recovered frequencies in x_{t-1} to a weighted average of $\alpha = \tilde{x}_0(t)$ and x_t. Assuming that $\tilde{x}_0(t)$ reasonably

approximates x_0, w_t can be learned via a regression problem in Fourier domain:

$$\min_{w_t} \mathbb{E}_{t,X_{(0,t-1,t)}} \underbrace{\left\| X_{t-1} - w_t X_0 - (1-w_t) X_t \right\|^2}_{\Gamma}. \tag{4.23}$$

Since $X_t^T X_0 = (\bar{\mathbf{\Lambda}}_t X_0)^T (\bar{\mathbf{\Lambda}}_t X_0) = \|X_t\|^2$, Γ is given as:

$$\begin{aligned}\Gamma =& \|X_{t-1}\|^2 + w_t^2 \|X_0\|^2 + (1-w_t)^2 \|X_t\|^2 - 2w_t \|X_{t-1}\|^2 \\ & - 2(1-w_t)\|X_t\|^2 + 2w_t(1-w_t)\|X_t\|^2, \qquad (4.24) \\ =& w_t^2 \left(\|X_0\|^2 - \|X_t\|^2\right) - 2w_t \left(\|X_{t-1}\|^2 - \|X_t\|^2\right) \\ & + \left(\|X_{t-1}\| - \|X_t\|^2\right). \qquad (4.25)\end{aligned}$$

Note that the partial derivative of Γ with respect to w_t is:

$$\partial\Gamma/\partial w_t = 2w_t \left(\|X_0\|^2 - \|X_t\|^2\right) - 2\left(\|X_{t-1}\|^2 - \|X_t\|^2\right). \tag{4.26}$$

Thus, the analytical solution to Eq. 4.23 can be expressed as:

$$w_t = \frac{\mathbb{E}_{t,X_{(t-1,t)}} \left(\|X_{t-1}\|^2 - \|X_t\|^2\right)}{\mathbb{E}_{t,X_{(0,t)}} \left(\|X_0\|^2 - \|X_t\|^2\right)}, \tag{4.27}$$

where the numerator and denominator reflect the expected energy differences between X_{t-1}, X_t and X_0, X_t, respectively, which can be estimated via a simple Monte-Carlo simulation over the training set. Note that w_t takes a maximum value of 1 attained at $t = 1$. Assuming that k-space energy distribution of MR images abide by an approximate inverse-power law (i.e., $\propto 1/r^\nu$), w_t should show exponential behavior [156]. Indeed, we find that w_t follows an exponential trend rising from $t = T_f$ to $t = 1$ (Fig. 4.2).

4.3 MRI reconstruction with FDB

Reconstruction is achieved by interleaving reverse diffusion steps for FDB with data-consistency projections (Alg. 1). The trained FDB prior maps R'-fold undersampled data onto fully-sampled data in T_f steps, by recovering $n \approx N^2(R'-1)/(R'T_f)$ frequency components in each step. Thus, to reconstruct an R-fold undersampled acquisition with $N^2(R-1)/R$ missing components, diffusion

sampling with the FDB prior must last for $T_r = \lfloor T_f \frac{(R-1)R'}{(R'-1)R} \rfloor$ steps. As sinusoidal time encoding is used in FDB, sampling can be extrapolated to a broader time range (i.e., $T_r > T_f$) to account for cases where $R > R'$ [152]. Yet, w_t is re-sampled across time to $\bar{w}_t$ to cover the designated range $[w_1\ w_{T_f}]$ in T_r steps. Sampling is initiated at T_r with the least-squares reconstruction of undersampled data, $\hat{x}_{T_r} = A^H y$ [42]. To maintain stochasticity, here we employ independent random instances of D_t during testing versus training. The mappings at time step t can then be expressed as:

$$\dot{x}_{t-1} = \hat{x}_t + (D_{t-1} - D_t)\tilde{x}_0(t) + \bar{w}_t D_t(\tilde{x}_0(t) - x_t), \tag{4.28}$$

$$\hat{x}_{t-1} = \dot{x}_{t-1} + A^H(y - A\dot{x}_{t-1}). \tag{4.29}$$

Finally, $\hat{x}_0$ is taken as the reconstructed image.

Chapter 5

Methods

5.1 Datasets

Experiments were performed on brain MRI data from IXI[1] and fastMRI [157]. Subjects were split into non-overlapping training, validation, test sets. In IXI, a subject-level split of (21,15,30) was used, and subjects were selected sequentially from the dataset. Coil-combined magnitude images for T_1, T_2 and PD contrasts were analyzed as single-coil acquisitions. For each contrast in a given subject, 100 axial cross-sections that contained brain tissue were selected. In fastMRI, a subject-level split of (240,60,120) was used. Subjects were selected sequentially from the dataset, albeit only subjects who data contained volumes with more than 10 cross-sections and 5 coil elements and did not suffer from poor quality due to excessive noise were selected. Multi-coil complex k-space data for T_1, T_2 and FLAIR contrasts were analyzed. Geometric coil compression was used to derive 5 virtual coils that preserved over 90% of the variance in the original data [158]. Retrospective undersampling was used based on a normal density across the two-dimensional transverse plane at acceleration rates $R = 4 - 8$ [1]. To construct the imaging operators for reconstruction, coil sensitivities were estimated using ESPIRiT on a central calibration region [159]. Volumetric k-space data were

[1] https://brain-development.org/ixi-dataset/

inverse Fourier transformed and split across the readout dimension, and each cross-section was reconstructed individually.

5.2 Competing Methods

FDB was demonstrated against eight state-of-the-art methods from the literature. Network models used two separate input and output channels to represent real and imaginary components. For each method, key hyperparameters were selected to maximize performance on the validation set. Modeling was performed via PyTorch on an Nvidia RTX 3090. Models were trained using the Adam optimizer with $(\beta_1, \beta_2) = (0.5, 0.9)$.

FDB: FDB used the architecture in [150]. Hyperparameters were set as $\eta = 10^{-4}$ learning rate, T_f=1000, E=40 epochs for IXI, E=10 for fastMRI, and R'=2.

LORAKS: A low-rank reconstruction was considered [160]. The k-space neighborhood radius and the rank of the system matrix were set as (2,6) for IXI, and (2,30) for fastMRI.

D5C5: A task-specific unrolled prior was considered [17]. Hyperparameters were $\eta = 2 \times 10^{-4}$, E=100.

PGIUN: A task-specific unrolled prior was considered [161]. Hyperparameters were $\eta = 10^{-4}$, E=100.

rGAN: A task-specific GAN prior was considered [143]. Hyperparameters were set as $\eta = 2 \times 10^{-4}$, E=100, adversarial and pixel-wise loss weights of (1,100).

DDPM: A task-agnostic diffusion prior was implemented [150]. Hyperparameters were set as $\eta = 10^{-4}$, T=1000, E=40.

Score-MRI: A score-based diffusion prior was implemented [162]. Hyperparameters were set as $\eta = 10^{-4}$, T=1000, E=40.

HFS-SDE: A score-based high frequency k-space diffusion prior was implemented [39]. Hyperparameters were set as $\eta = 10^{-4}$, T=1000, E=40 for IXI, E=10 for fastMRI.

CDiffMR: A cold diffusion prior was considered based on an asymptotic start-point derived via k-space undersampling [53]. Reconstructions were initiated via the zero-filled (ZF) reconstruction at the time step corresponding to R of the test acquisition [53]. Hyperparameters were set as $\eta = 10^{-4}$, T=1000, E=40.

Soft-Diffusion: A soft diffusion prior was considered based on an asymptotic start-point derived via k-space undersampling [46]. Reconstructions were initiated via the zero-filled (ZF) reconstruction at the time step corresponding to R of the test acquisition [46]. Hyperparameters were set as $\eta = 10^{-4}$, T=1000, E=40.

I^{2}SB: A conventional diffusion bridge with start- and end-point identical to FDB was considered [47]. Reconstructions were initiated via the ZF reconstruction of the test acquisition. Hyperparameters were $\eta = 10^{-4}$, T=200, E=40.

5.3 Analyses

For each dataset, model training was performed on data aggregated across multiple contrasts (T_1, T_2, PD in IXI and T_1, T_2, FLAIR in fastMRI). No explicit contrast information was provided during training, and data samples were randomly drawn for the aggregated training set. Task-specific models were trained for each R separately to maintain high performance. Meanwhile, task-agnostic models were trained without knowledge of R. Models were trained and tested on two-dimensional cross sections. Single-coil reconstructions were conducted on IXI, and multi-coil reconstructions were conducted on fastMRI. Reconstruction performance was assessed by quantifying peak signal-to-noise ratio (PSNR)

and structural similarity index (SSIM) between recovered and reference images. Reference images were derived via Fourier reconstruction of fully-sampled acquisitions. Significance of performance differences among competing methods were assessed via non-parametric Wilcoxon signed-rank tests.

Chapter 6

Results

6.1 Comparison Studies

To improve the task relevance of the diffusion prior it captures for MRI reconstruction, FDB leverages several key elements including a diffusion bridge that maps moderately undersampled to fully-sampled data, and stochastic degradation/recovery operators based on binary selection of k-space components. With these unique elements, FDB can offer higher fidelity diffusion priors over alternative methods that are less effective in attaining task relevance, including task-agnostic diffusion priors, cold/soft diffusion priors and conventional diffusion bridges. We reasoned that this benefit should be more prominent in within-domain settings where models are trained and tested on the same dataset and sampling density, as this scenario ensures that each trained model will retain the highest degree of task relevance when evaluated. To assess this prediction, we compared FDB against a traditional method (LORAKS), task-specific priors (D5C5, rGAN, PGIUN), task-agnostic diffusion priors (DDPM, Score-MRI), cold/soft diffusion priors (CDiffMR, Soft-Diffusion), and a diffusion bridge (I^2SB). Note that all competing methods except LORAKS were trained, so their data-driven priors naturally depend on the dataset used. Meanwhile, DDPM and DB_{blur} do not use information regarding the imaging operator and thereby the

sampling density during training. Yet, network inputs to D5C5, rGAN and PGIUN depend on the undersampling patterns used in the training set, and start-points of the diffusion process in CDiffMR, I^{2}SB and FDB are obtained by assuming a certain sampling density.

6.1.1 Within-Domain Experiments

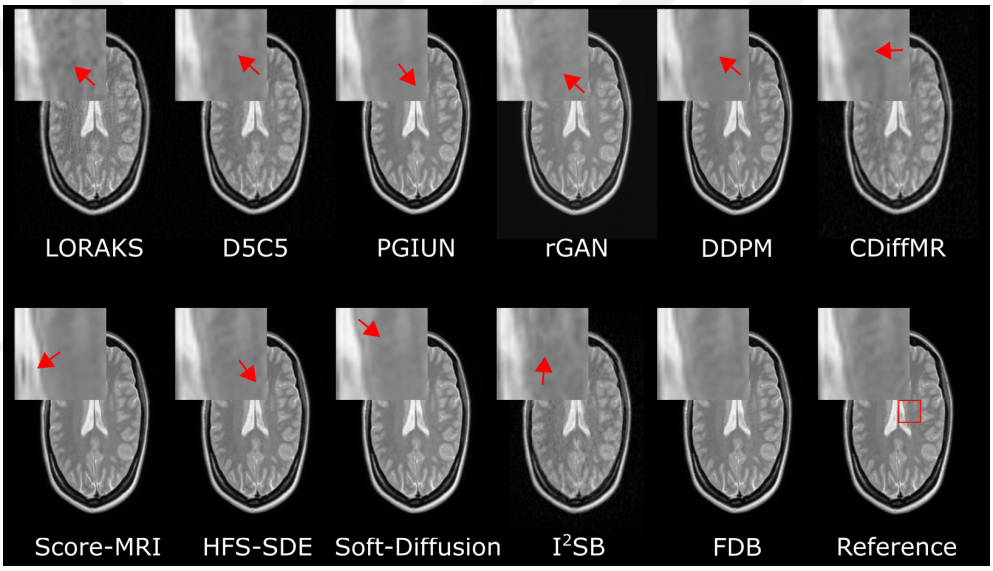

Figure 6.1: Within-domain reconstructions of representative cross-sections from a PD acquisition in IXI at $R = 4$. Images from competing methods are displayed along with the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region. In this example FDB has PSNR 44.1dB and SSIM 99.6% while HFS-SDE, which is the second best performing method has PSNR 43.6dB and SSIM 99.5%.

Performance metrics in within-domain settings are listed in Table 6.1 for IXI, and in Table 6.2 for fastMRI. Overall, FDB is the top performer in all examined tasks and datasets ($p<0.05$), except for HFS-SDE that yields similar SSIM at R=4 on IXI. On average across tasks, FDB outperforms the traditional method by 7.5dB PSNR, 24.4% SSIM, task-specific priors by 4.3dB PSNR and 12.0% SSIM, task-agnostic diffusion priors by 4.8dB and 7.6% SSIM, and cold/soft diffusion priors by 5.1dB PSNR and 8.4% SSIM, and the diffusion bridge by 6.5dB

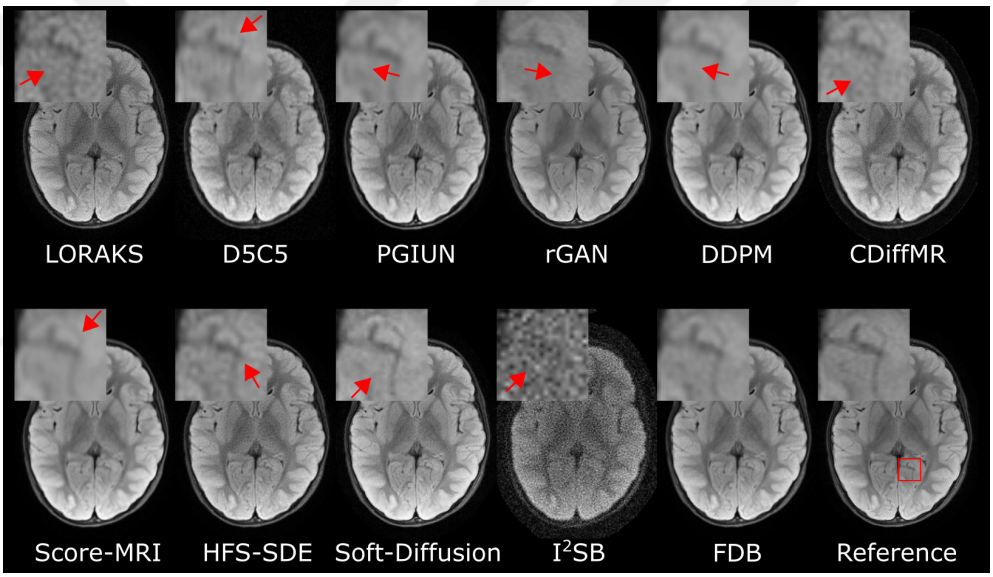

Figure 6.2: Within-domain reconstructions of representative cross-sections from a FLAIR acquisition in fastMRI at $R = 8$. Images from competing methods are displayed along with the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region. In this example FDB has PSNR 30.4dB and SSIM 96.3% while Soft-Diffusion, which is the second best performing method has PSNR 24.9dB and SSIM 93.2%.

Table 6.1: Within-domain performance on IXI at $R = 4, 6, 8$. Results averaged across T_1, T_2, PD contrasts. PSNR (dB) and SSIM (%) listed as mean±std across the test set. Boldface marks the top performing method.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	31.3±1.7	63.2±4.4	29.2±1.6	57.4±4.4	27.9±1.3	55.7±4.1
D5C5	33.8±0.6	78.3±1.6	31.4±0.6	71.5±1.8	30.0±0.6	66.7±1.9
PGIUN	38.5±2.4	96.9±0.6	36.4±2.4	97.0±0.8	34.4±2.4	96.7±0.9
rGAN	36.6±3.0	85.1±7.1	33.9±2.9	81.4±7.2	32.6±2.9	79.5±7.3
DDPM	38.0±2.5	97.6±0.7	35.7±2.5	96.7±1.0	32.9±2.6	95.4±1.4
CDiffMR	32.4±2.4	86.0±2.2	30.3±2.3	82.4±2.6	29.3±2.4	80.7±2.6
Score-MRI	37.8±2.5	97.6±0.8	35.4±2.5	96.7±1.1	32.6±2.5	95.6±1.5
HFS-SDE	43.2±2.8	**99.3±0.8**	40.1±2.8	98.7±0.9	36.3±2.8	97.7±1.1
Soft-Diffusion	38.5±2.3	98.1±0.6	36.0±2.4	97.1±0.9	33.2±2.5	95.8±1.3
I^2SB	30.3±2.6	83.0±2.8	28.8±2.6	80.6±3.0	28.2±2.6	79.9±3.1
FDB	**43.7±2.9**	**99.3±0.9**	**40.8±2.8**	**98.8±0.9**	**36.9±2.6**	**97.8±1.0**

PSNR and 9.8% SSIM. Representative reconstructions are depicted in Fig. 6.1 for IXI, and Fig. 6.2 for fastMRI. Among competing methods, LORAKS, I^2SB and to a lesser degree CDiffMR and HFS-SDE can show noise amplification; D5C5, DDPM and Score-MRI show a degree of blur due to oversmoothing; rGAN and Soft-Diffusion show visible residual noise and loss of some structural details. In contrast, FDB yields images with more accurate depiction of tissue structure along with lower artifacts and noise than competing methods. Taken together, these results suggest that the increased task-relevance of the diffusion prior captured by FDB elicits superior performance in MRI reconstruction against competing non-diffusion and diffusion methods.

6.1.2 Cross-Domain Experiments

While an image prior with increased task relevance is expected to boost performance in within-domain settings, there is risk that it limits generalization in

Table 6.2: Within-domain performance on fastMRI at $R = 4, 6, 8$. Results averaged across T_1, T_2, FLAIR contrasts.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	35.2±2.4	92.0±5.7	33.5±2.3	90.5±6.1	32.3±2.3	89.5±6.3
D5C5	35.3±1.5	90.5±2.3	33.4±1.4	87.2±2.8	32.2±1.4	85.0±3.1
PGIUN	35.7±2.3	93.7±4.5	34.3±2.2	91.8±5.1	33.7±2.2	90.9±5.5
rGAN	35.9±2.8	93.1±4.7	33.9±2.6	90.7±5.5	32.9±2.8	90.1±5.9
DDPM	36.1±2.3	93.7±4.6	34.8±2.3	92.2±5.2	34.0±2.3	91.1±5.6
CDiffMR	31.4±2.7	80.5±7.4	31.2±2.4	79.1±7.2	31.2±2.3	78.5±8.1
Score-MRI	35.6±2.2	93.4±4.0	34.3±2.2	91.8±4.7	33.5±2.2	90.5±5.0
HFS-SDE	36.3±2.1	92.9±4.1	33.9±1.8	89.0±4.7	32.2±1.7	85.1±5.2
Soft-Diffusion	36.3±2.6	91.1±6.3	34.9±2.3	89.9±6.4	34.1±2.1	89.2±6.4
I^2SB	28.4±2.1	75.7±4.5	26.9±2.1	72.6±4.8	26.2±2.1	70.9±5.0
FDB	**37.4±2.5**	**94.3±4.5**	**35.9±2.4**	**92.7±5.2**	**35.0±2.4**	**91.5±5.6**

Table 6.3: Cross-domain performance for models trained on 2D k-space sampling patterns, tested on 1D patterns for IXI at $R = 4, 6, 8$.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	24.3±0.9	42.0±2.6	22.5±0.7	37.0±2.5	21.3±0.7	34.2±2.7
D5C5	24.4±0.7	55.0±3.7	23.4± 0.6	52.6±3.8	22.7±0.6	49.8±4.0
PGIUN	32.0±2.6	95.9±1.3	29.6±2.5	91.6±2.6	27.0±2.4	89.1±2.9
rGAN	29.3±2.4	49.1±6.1	27.2±2.4	42.7±6.0	25.8±2.3	39.2±5.7
DDPM	30.1±2.8	94.3±1.9	26.7±2.7	89.6±3.0	25.1±2.5	85.5±3.6
CDiffMR	28.6±2.6	83.7±2.5	26.2±1.7	78.3±2.5	24.8±2.5	74.2±3.6
Score-MRI	29.8±2.7	94.8±1.8	26.5±2.6	90.8±2.9	25.0±2.5	87.0±3.5
HFS-SDE	33.1±3.0	96.4±1.4	28.6±3.0	91.8±2.6	25.5±2.9	85.7±4.0
Soft-Diffusion	30.5±2.6	94.5±2.0	27.8±2.6	89.1±2.7	25.2±2.6	85.6±4.0
I^2SB	27.8±2.8	85.0±3.2	26.1±2.7	81.8±3.9	25.5±2.7	81.0±4.2
FDB	**34.0±2.9**	**96.7±1.2**	**30.4±2.8**	**93.5±2.0**	**28.1±2.7**	**89.8±2.7**

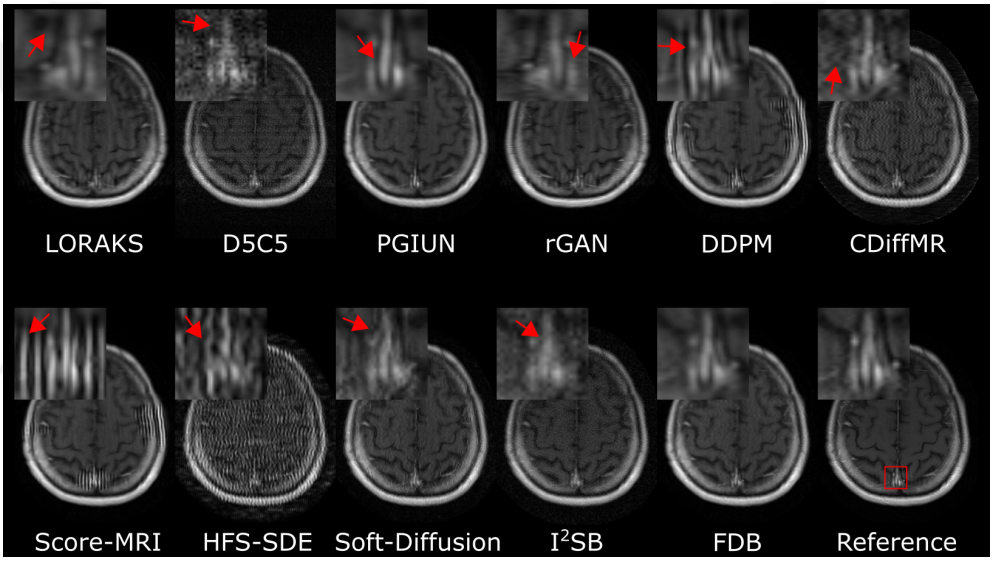

Figure 6.3: Cross-domain reconstructions of representative cross-sections. A T_1 acquisition in fastMRI at $R = 8$ is shown. Models trained on 2D k-space sampling patterns were tested on 1D patterns under matched R. Images from competing methods are displayed along with the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region. In this example FDB has PSNR 28.2dB and SSIM 96.0% while HFS-SDE, which is the second best performing method has PSNR 22.8dB and SSIM 66.5%.

cross-domain settings as a side effect, potentially yielding suboptimal performance against baselines that have lower task relevance. To ensure that FDB does not suffer from this limitation, we next compared FDB against competing methods in cross-domain settings. Specifically, shifts in sampling density and dataset between training and test sets were considered. For sampling density, shifts from 2D to 1D variable-density sampling was considered under matching acceleration rate R. Cross-domain performance metrics listed in Tables 6.3 and 6.4 indicate that FDB is the top performer in all tasks ($p<0.05$). On average across tasks, FDB outperforms the traditional method by 7.5dB PSNR, 24.4% SSIM, task-specific priors by 4.3dB PSNR and 12.0% SSIM, task-agnostic diffusion priors by 4.8dB and 7.6% SSIM, and cold/soft diffusion priors by 2.8dB PSNR and 11.2% SSIM, and the diffusion bridge by 6.5dB PSNR and 9.8% SSIM. For shifts in dataset, generalization from the multi-coil fastMRI dataset to single-coil IXI dataset was examined as listed in Table 6.5. Again, FDB achieves the highest performance among competing methods ($p<0.05$). On average across tasks, FDB outperforms the traditional method by 7.5dB PSNR, 24.4% SSIM, task-specific priors by 4.3dB PSNR and 12.0% SSIM, task-agnostic diffusion priors by 4.8dB and 7.6% SSIM, and cold/soft diffusion priors by 9.6dB PSNR and 27.5% SSIM, and the diffusion bridge by 6.5dB PSNR and 9.8% SSIM. Taken together, these results suggest that FDB captures a more reliable prior against domain shifts in sampling density and data distribution than competing methods, including diffusion priors.

Representative images from competing methods are depicted in Fig. 6.3 for domain shifts in sampling density, and Fig. 6.4 for domain shifts in dataset. Under shifts in sampling density, LORAKS, rGAN show residual aliasing artifacts; DDPM, CDiffMR, Score-MRI, HFS-SDE and to a lesser degree Soft-Diffusion show high-frequency ringing artifacts; D5C5, I^2SB shows high noise amplification. While competing methods are able to maintain relatively more reliable performance under shifts in dataset, LORAKS, D5C5, rGAN, CDiffMR, I^2SB suffer from noise amplification as particularly visible in background regions near tissue signals; DDPM, CDiffMR, HFS-SDE show a degree of blur that can lead to tissue broadening or inaccuracies in tissue signals; Soft-Diffusion can suffer from loss of

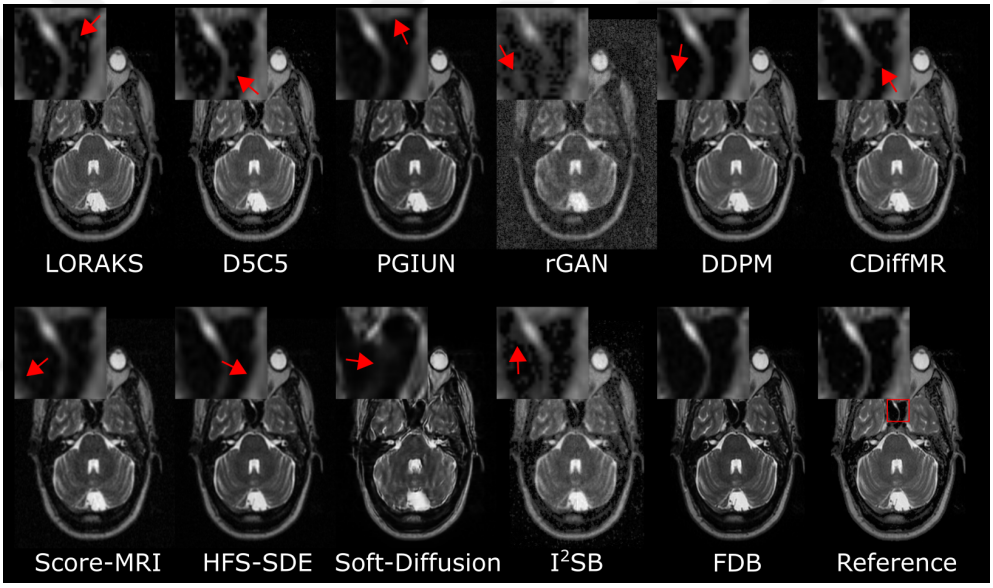

Figure 6.4: Cross-domain reconstructions of representative cross-sections. A T_2 acquisition in IXI at $R = 8$ is shown. Models trained on multi-coil fastMRI data were tested on single-coil IXI data. Images from competing methods are displayed along with the reference image computed via Fourier reconstruction of fully-sampled data. Zoom-in windows are included to highlight differences among methods, and the red box on the reference image marks the zoomed region. In this example FDB has PSNR 33.5dB and SSIM 88.9% while DDPM, which is the second best performing method has PSNR 32.2dB and SSIM 83.3%.

Table 6.4: Cross-domain performance for models trained on 2D k-space sampling patterns, tested on 1D patterns for fastMRI at $R = 4, 6, 8$.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	28.1±2.3	85.3±5.6	25.9±2.3	80.5±6.2	24.5±2.2	76.9±6.5
D5C5	26.4±1.6	68.1±5.6	25.1±1.1	64.0±8.0	24.7±9.6	53.3±4.0
PGIUN	30.1±2.0	89.4±5.5	27.9±1.7	82.3±6.0	26.4±1.7	80.6±6.2
rGAN	29.1±2.9	86.1±5.3	27.0±2.7	83.5±6.3	25.4±2.6	77.3±6.8
DDPM	29.4±2.0	88.8±5.5	27.4±1.5	84.1±5.6	25.7±1.5	80.6±6.5
CDiffMR	27.1±2.2	68.7±9.2	25.7±2.0	63.0±8.5	25.1±2.1	63.7±9.7
Score-MRI	29.0±1.8	88.1±5.0	26.7±1.6	83.5±5.7	25.2±1.5	79.7±6.2
HFS-SDE	24.3±1.2	61.1±4.9	22.6±1.1	52.6±4.9	21.6±1.1	47.9±4.9
Soft-Diffusion	29.9±1.7	84.8±5.3	28.2±1.5	79.8±5.4	26.9±1.5	75.9±5.6
I^2SB	24.3±4.5	68.6±8.8	22.4±4.9	65.1±9.3	21.3±5.2	63.3±9.5
FDB	**30.6±2.2**	**89.5±6.0**	**29.0±2.0**	**85.9±6.9**	**27.8±2.0**	**82.7±7.3**

Table 6.5: Cross-domain performance for models trained on multi-coil fastMRI data, and tested on single-coil IXI data at $R = 4, 6, 8$.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LORAKS	31.3±1.7	63.2±4.4	29.2±1.6	57.4±4.4	27.9±1.3	55.7±4.1
D5C5	34.4±0.7	80.3±1.6	31.9±0.7	73.5 ± 1.8	30.3±0.6	68.1±2.0
PGIUN	34.2±2.7	92.9±1.6	32.8±2.7	92.2±1.6	31.9±2.8	92.1±1.8
rGAN	19.9±3.5	25.1±7.7	19.4±3.5	24.2±6.9	19.2±3.1	23.5±6.7
DDPM	34.8±2.7	93.6±1.4	33.3±2.6	92.4±1.6	32.2±2.6	92.0±1.8
CDiffMR	30.6±2.7	82.7±3.0	28.9±2.6	79.5±3.3	28.2±2.6	78.2±3.3
Score-MRI	31.8±2.7	86.0±2.6	30.2±2.7	83.4±2.9	29.4±2.8	82.3±3.1
HFS-SDE	34.8±2.6	93.6±1.4	33.3±2.6	92.4±1.6	32.1±2.5	92.0±1.7
Soft-Diffusion	23.7±2.2	60.5±5.4	19.7±2.6	53.3±5.4	16.7±2.7	45.3±5.5
I^2SB	31.2±2.5	84.6±2.5	29.7±2.5	82.5±2.7	28.9±2.6	81.2±2.8
FDB	**36.0±2.7**	**95.4±1.0**	**34.1±2.6**	**93.8±1.2**	**32.7±2.7**	**93.0±1.3**

detailed features and inaccurate depiction of anatomy. In comparison, FDB produces more accurate depiction of tissue structure, and relatively low artifacts and noise than baselines. Note that task-specific priors that experience a training-test mismatch between the requirements of the reconstruction task due to change in sampling density or data distribution yield notable residual artifacts and noise amplification. Meanwhile, task-agnostic diffusion priors are relatively more resilient against shifts in dataset, they can suffer from suboptimal performance under challenging reconstruction tasks due to 1D sampling density. On the other hand, FDB maintains a favorable trade-off between task-relevance and generalizability. Taken together, our results suggest that the increased task-relevance of FDB enables significant increases in within-domain reconstruction performance, without compromising its superiority against baselines in cross-domain settings.

6.2 Ablation Studies

6.2.1 Effect of R'

A group of ablation studies were conducted to assess the influence of major method components and parameters to FDB's reconstruction performance. Note that FDB captures a delimited transformation between an end-point of fully-sampled data and a start-point of moderately undersampled data, rendering the degradation level R' at the start-point critical. First, we examined the influence of R' on reconstruction performance. Table 6.6 lists performance for FDB variants trained using R'=2-16, In general, employing lower R' values yields consistently higher performance across R prescribed for test acquisitions. This finding suggests that it is not ideal to strictly match R' and R, and thereby helps explain the superior performance of FDB against diffusion priors with asymptotic start-points based on severe degradation levels (e.g., DDPM and CDiffMR).

Table 6.6: Performance of FDB priors on IXI. $R' = 2 - 16$ is the acceleration rate equivalent to the degradation at the start-point of the diffusion process during training, and $R = 4, 6, 8$ is the acceleration rate during testing.

	$R = 4$		$R = 6$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
$R' = 2$	**43.7±2.9**	**99.3±0.9**	**40.8±2.8**	**98.8±0.9**	**36.9±2.6**	**97.8±1.0**
$R' = 4$	42.5±2.9	99.1±0.7	38.6±2.3	98.0±0.5	35.0±2.7	96.7±1.0
$R' = 8$	42.6±2.8	99.2±0.5	38.3±2.2	98.0±0.5	34.6±2.5	96.5±1.0
$R' = 16$	42.4±3.1	99.1±0.6	38.2±1.9	98.0±0.5	34.6±2.6	96.6±1.0

6.2.2 Effect of Method Components

Next, we examined the importance of key method components related to the degradation operator and sampling algorithm in FDB by forming several variants. A 'w/o Fourier constraint' variant removed the Fourier constraint in the diffusion bridge to instead express intermediate image samples as a weighted average of start- and end-points (i.e., a weighted average of fully-sampled and undersampled data) as in conventional diffusion bridges [49]. A 'w/o radius threshold' variant removed the radius threshold in the stochastic degradation operator so as to promote uniform selection of removed frequency components. A 'w/o CCS' variant ablated the corrected sampling algorithm to instead use iterated Langevin sampling as in Eq. 4.15. Lastly, a 'w/o data consistency' variant was formed by ablating data-consistency projections during inference. Table 6.7 lists performance for variants. We find that FDB generally yields superior performance against variants, except for 'w/o radius threshold' that yields on par or slightly higher SSIM. These results indicate that key method components that distinguish FDB from conventional diffusion bridges contributed significantly to MRI reconstructions.

Table 6.7: Performance of FDB variants devised to assess the importance of method components on IXI at $R = 4, 8$. Variants replacing Fourier-constrained degradations with weighted averaging (w/o Fourier const.), ablating the radial threshold from the forward degradation operator for uniform selection across k-space (w/o rad. thres.), ablating the continually corrected sampling algorithm (w/o CCS), and ablating the data-consistency projections during inference (w/o data cons.) were considered.

	$R = 4$		$R = 8$	
	PSNR	SSIM	PSNR	SSIM
FDB	**43.7±2.9**	**99.3±0.9**	**36.9±2.6**	97.8±1.0
w/o Fourier const.	17.0±2.8	50.9±5.2	17.1±2.6	51.4±4.9
w/o rad. thres.	43.0±2.7	**99.3±1.0**	36.8±2.5	**98.1±1.0**
w/o CCS	39.0±2.6	96.5±0.8	34.2±2.6	94.0±1.3
w/o data cons.	30.5±1.6	93.3±1.7	28.0±1.5	89.6±2.6

6.2.3 Effect of Scheduling Components

We also examined the importance of key parameters in FDB related to the schedules for the number of frequency components removed with the stochastic degradation operator at each time step (n), and for the weight of the correction term in the CCS algorithm (w_t). For n, the linear schedule in FDB was compared against logarithmic and cosine schedules across time steps. For w_t, the learned schedule in FDB was compared against linear and exponential schedules. Table 6.8 lists performance for FDB variants. We find that FDB yields superior performance against variants in all tasks. These findings indicate that the proposed parameter schedules contribute significantly to FDB's reconstruction performance.

Table 6.8: Performance of FDB variants devised to assess the importance of proposed schedules for the number of frequency components (n) and weighting of the correction term (w_t). For n, the proposed linear scheduling was compared against logarithmic and cosine scheduling. For w_t, the proposed learned scheduling was compared against linear and exponential scheduling.

		$R = 4$		$R = 8$	
		PSNR	SSIM	PSNR	SSIM
n	Linear (FDB)	**43.7±2.9**	**99.3±0.9**	**36.9±2.6**	**97.8±1.0**
	Log	40.7±3.0	98.7±2.9	35.5±2.8	97.3±0.9
	Cosine	43.5±2.4	99.1±0.2	36.5±2.2	97.2±0.6
w_t	Learned (FDB)	**43.7±2.9**	**99.3±0.9**	**36.9±2.6**	**97.8±1.0**
	Linear	24.4±2.4	85.1±3.1	20.8±2.4	80.2±2.7
	Exponential	25.4±1.1	88.5±2.0	24.1±1.2	88.3±1.5

Chapter 7

Discussion

FDB was demonstrated against several state-of-the-art deep-learning baselines including task-specific non-diffusion priors, task-agnostic diffusion priors, cold/soft diffusion priors and common diffusion bridges. In general, FDB offers enhanced image quality against baselines consistently across the examined acceleration rates, tissue contrasts, and datasets. In the context of diffusion-based methods, improved performance over task-agnostic diffusion priors suggest that utilization of task-relevant degradations based on undersampling serves a critical role in obtaining high-fidelity image priors. Improved performance over cold/soft diffusion priors suggests that employing a diffusion process with a start-point based on moderate degradation levels is important in improving accuracy of reverse diffusion sampling. Improved performance over common diffusion bridges suggests that utilizing a degradation operator that respects the physics of k-space undersampling yields an enhanced image prior for MRI.

Paralleling its performance in within-domain settings, here we find that FDB remains the top-performer among competing methods even a fair level of domain shifts in k-space sampling density and dataset. Admittedly, the performance benefits for FDB over diffusion baselines are relatively more modest under cross-domain versus within-domain settings. This result suggests that the design elements in FDB that improve task relevance by injecting information about

k-space undersampling can partly influence generalization. Contrarily, FDB's performance benefits over task-specific baselines are notably higher under cross-domain versus within-domain settings. This finding indicates that FDB still maintains the competitive generalization abilities commonly associated with diffusion models. That said, it remains important future work to assess performance under more drastic domain shifts (e.g., a change in anatomy between training-test sets). Recent studies have proposed adaptation of image priors to individual subjects to boost generalization to atypical anatomy [31, 147]. During inference, adaptation methods optimize network parameters that form the prior so as to improve subject-specific reconstruction performance. While promising results have been reported for adapted priors, employing adaptation on diffusion priors with thousands of sampling steps substantially elevates computational burden [40]. Although FDB improves efficiency by initiating sampling with a least-squares reconstruction, prior adaptation can be still challenging.

Conventional inference in diffusion models employs Langevin sampling across thousands of steps to generate images, resulting in characteristically slow inference. Unlike diffusion priors with asymptotic normality assumptions that initiate sampling with a random noise image, FDB initiates sampling with the least-squares reconstruction of undersampled acquisitions, recently shown to improve efficiency [42]. Yet, FDB has longer inference times compared to other generative models such as adversarial priors with one-step image sampling. Several acceleration strategies can improve efficiency in FDB. A training-phase approach is to integrate adversarial learning to enable reverse diffusion over large step sizes [40]. Alternatively, inference-phase approaches can be adopted to perform interleaved sampling followed by a refinement procedure [38], or via distillation of the trained priors [119]. Future studies are warranted to assess the influence of these acceleration approaches on FDB's performance.

Several technical limitations can be addressed to help enhance the performance and utility of FDB. Here, we observed that the continually-corrected sampling algorithm that alleviates residual errors in recovered k-space components significantly improves performance in reverse diffusion steps. This improvement suggests that estimation accuracy for $\tilde{x}_0(t)$ is relatively poor near the start-point of

the diffusion process, which not only hampers accuracy of recovered k-space components but can also lower computational efficiency. An alternative approach to improve sampling fidelity could be to adopt recursive estimation of $\tilde{x}_0(t)$ at each time step [153], which could shorten the number of steps required for approaching convergence. To gradually impute missing k-space components, FDB employs an image-domain recovery operator that first predicts the clean image at 0 given the image sample at t, and then degrades it back at a lower level to obtain the image sample at $t-1$. This procedure can cause the values of acquired k-space components to deviate away from their originally measured values in undersampled data. For this reason, FDB used additional data-consistency projections between consecutive reverse diffusion steps, which increase computational load. Improved efficiency might be attained by adopting recent data-consistent sampling algorithms that avoid the need for data-consistency projections altogether [48]. Here we demonstrated that FDB's diffusion process based on undersampling degradations helps improve image quality in MRI reconstruction. Yet, this task-specialization might hamper performance in non-reconstruction tasks such as denoising or deblurring. If a single FDB model will be deployed for multiple distinct tasks, a straightforward approach to mitigate performance losses could be to use training sets with mixed samples derived from the different degradation operators associated with each task. A task identifier could also be employed as input to the model along with the time step to attain task-specific modulation of feature maps across the recovery network [163].

Another key area for improvement is related to the learning strategies adopted in FDB. Here, image priors were trained on MRI acquisitions pooled across multiple contrasts, but later used to reconstruct single-contrast data. Performance might be improved by employing contrast-specific modulation of feature maps for elevated specificity in the prior [163], by training separate contrast-specific priors for reconstruction [22], or by performing joint reconstruction across multiple contrasts [164, 143]. Moreover, here image priors were trained on datasets with fully-sampled acquisitions, which might limit utility in cases where acquisition of fully-sampled data is challenging [165, 166]. The reliance of FDB on Nyquist-sampled

acquisitions can be lowered by adopting self-supervised [167, 168, 169, 170], cycle-consistent [171] or low-rank assisted learning [41].

Other areas for improvement include the network architecture and frequency removal operator in FDB. Here, a UNet architecture common to diffusion modeling studies was used [150]. Inverse problems in medical imaging have recently been reported to benefit from utilization of transformer backbones to capture long-range context [32, 144, 172]. Hybrid transformer-convolutional architectures might be utilized in FDB to achieve a favorable trade-off between spatial precision and contextual sensitivity [173]. Here, FDB assumed degradations based on spatial-frequency removal with peripheral-to-central ordering. This results in a variable-density sampling pattern for remaining frequency components at the start-point of the diffusion process, and FDB priors were later tested on variable-density undersampled acquisitions. The diffusion process in FDB can be adopted to implement other k-space sampling densities by modifying the scheduling of the k-space radius threshold over time steps accordingly. While our results suggest that FDB is reasonably reliable against domain shifts in sampling density, performance improvements can be sought by ensuring a match between training and test sets. It remains future work to assess the ideal architecture and frequency-removal scheduling for FDB-based MRI reconstruction.

Chapter 8

Conclusion

In this study, we introduced a novel diffusion bridge, FDB, for accelerated MRI reconstruction. FDB employs a stochastic degradation operator during forward diffusion, which performs random spatial-frequency removal to directly map between fully-sampled and undersampled data. For image reconstruction, FDB begins with the least-squares reconstruction of the undersampled acquisition and utilizes a novel reverse diffusion algorithm to progressively impute missing frequencies.

Demonstrations on brain MRI show that FDB consistently outperforms state-of-the-art reconstruction methods based on non-diffusion and diffusion priors for both within-domain and cross-domain cases. FDB's ability to maintain high reconstruction quality across these varied conditions highlights its better immunity to domain shifts and superior generalization, making it a more robust and reliable solution for MRI reconstruction across diverse clinical scenarios.

Given its demonstrated effectiveness in MRI reconstruction, FDB's principles could be adapted to other imaging tasks that involve undersampled or degraded data. For instance, in computed tomography (CT), FDB could be applied to reconstruct images from limited projection data, potentially reducing radiation exposure to patients. In positron emission tomography (PET), FDB might help

improve image quality by mitigating the effects of noise and attenuation. Additionally, FDB could be explored for applications in optical imaging, such as super-resolution microscopy, where enhancing image resolution from limited data is a critical challenge. By leveraging FDB's ability to handle undersampled data and its robustness to domain shifts, researchers can explore its potential to advance various imaging modalities and improve diagnostic capabilities.

Bibliography

[1] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.

[2] D. Liang, J. Cheng, Z. Ke, and L. Ying, "Deep magnetic resonance image reconstruction: Inverse problems meet neural networks," *IEEE Signal Process. Mag.*, vol. 37, no. 1, pp. 141–151, 2020.

[3] S. U. H. Dar, Ş. Öztürk, M. Özbey, K. K. Oguz, and T. Çukur, "Parallel-stream fusion of scan-specific and scan-general priors for learning deep mri reconstruction in low-data regimes," *Computers in Biology and Medicine*, vol. 167, p. 107610, 2023.

[4] 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 Sig. Process. Mag.*, vol. 37, no. 1, pp. 117–127, 2020.

[5] K. Hammernik, T. Küstner, B. Yaman, Z. Huang, D. Rueckert, F. Knoll, and M. Akçakaya, "Physics-driven deep learning for computational magnetic resonance imaging," *arXiv:2203.12215*, 2022.

[6] Y. Korkmaz, M. Özbey, and T. Cukur, "Mri reconstruction with conditional adversarial transformers," in *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 62–71, Springer, 2022.

[7] F. Lam, X. Peng, and Z.-P. Liang, "High-dimensional mr spatiospectral imaging by integrating physics-based modeling and data-driven machine

learning: Current progress and future directions," *IEEE Signal Process. Mag.*, vol. 40, no. 2, pp. 101–115, 2023.

[8] 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.

[9] 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.

[10] 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.

[11] E. Ilicak, E. U. Saritas, and T. Cukur, "Automated parameter selection for accelerated mri reconstruction via low-rank modeling of local k-space neighborhoods," *Zeitschrift für Medizinische Physik*, vol. 33, no. 2, pp. 203–219, 2023.

[12] B. Bilgic and T. Cukur, "Parallel imaging and reconstruction techniques," in *Advances in Magnetic Resonance Technology and Applications*, vol. 10, pp. 139–159, Elsevier, 2023.

[13] B. Güven, A. Güngör, M. U. Bahçeci, and T. Çukur, "Deep learning reconstruction for single pixel imaging with generative adversarial networks," in *2023 IEEE International Conference on Image Processing (ICIP)*, pp. 2060–2064, IEEE, 2023.

[14] M. Acar, T. Çukur, and İ. Öksüz, "Segmentation-aware mri reconstruction," in *Machine Learning for Medical Image Reconstruction* (N. Haq, P. Johnson, A. Maier, C. Qin, T. Würfl, and J. Yoo, eds.), (Cham), pp. 53–61, Springer International Publishing, 2022.

[15] A. Güngör, B. Askin, D. A. Soydan, C. B. Top, E. U. Saritas, and T. Çukur, "Deq-mpi: A deep equilibrium reconstruction with learned consistency for magnetic particle imaging," *IEEE Transactions on Medical Imaging*, vol. 43, no. 1, pp. 321–334, 2023.

[16] 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.

[17] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert, "A Deep Cascade of Convolutional Neural Networks for MR Image Reconstruction," in *IPMI*, pp. 647–658, 2017.

[18] 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," *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3055–3071, 2017.

[19] 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.

[20] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-Based deep learning architecture for inverse problems," *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 394–405, 2019.

[21] A. Güngör, B. B. Bilecen, and T. Çukur, "Bayesian conditioned diffusion models for inverse problems," *arXiv preprint arXiv:2406.09768*, 2024.

[22] S. U. H. Dar, M. Özbey, A. B. Çatli, and T. Çukur, "A transfer-learning approach for accelerated MRI using deep neural networks," *Magn. Reson. Med.*, vol. 84, no. 2, pp. 663–685, 2020.

[23] K. Kwon, D. Kim, and H. Park, "A parallel MR imaging method using multilayer perceptron," *Med. Phys.*, vol. 44, no. 12, pp. 6209–6224, 2017.

[24] 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 Trans. Med. Imaging*, vol. 38, no. 1, pp. 167–179, 2019.

[25] Z. Ramzi, C. G R, J.-L. Starck, and P. Ciuciu, "NC-PDNet: A Density-Compensated Unrolled Network for 2D and 3D Non-Cartesian MRI Reconstruction," *IEEE Trans. Med. Imaging*, vol. 41, no. 7, pp. 1625–1638, 2022.

[26] G. Wang, T. Luo, J.-F. Nielsen, D. C. Noll, and J. A. Fessler, "B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri," *IEEE Trans. Med. Imaging*, vol. 41, no. 9, pp. 2318–2330, 2022.

[27] F. Liu, A. Samsonov, L. Chen, R. Kijowski, and L. Feng, "Santis: Sampling-augmented neural network with incoherent structure for MR image reconstruction," *Magn. Reson. Med.*, vol. 82, no. 5, pp. 1890–1904, 2019.

[28] 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," *IEEE Trans. Med. Imaging*, vol. 42, no. 1, pp. 79–90, 2023.

[29] A. Pramanik, S. Bhave, S. Sajib, S. D. Sharma, and M. Jacob, "Adapting model-based deep learning to multiple acquisition conditions: Ada-modl," *Magn. Reson. Med.*, 2023.

[30] K. C. Tezcan, C. F. Baumgartner, R. Luechinger, K. P. Pruessmann, and E. Konukoglu, "Mr image reconstruction using deep density priors," *IEEE Trans. Med. Imaging*, vol. 38, no. 7, pp. 1633–1642, 2018.

[31] D. Narnhofer, K. Hammernik, F. Knoll, and T. Pock, "Inverse gans for accelerated mri reconstruction," in *Wavelets and Sparsity XVIII*, vol. 11138, pp. 381–392, SPIE, 2019.

[32] Y. Korkmaz, S. U. Dar, M. Yurt, M. Özbey, and T. Cukur, "Unsupervised mri reconstruction via zero-shot learned adversarial transformers," *IEEE Trans. Med. Imaging*, vol. 41, no. 7, pp. 1747–1763, 2022.

[33] M. U. Mirza and T. Çukur, "Super-resolution diffusion model for accelerated mri reconstruction," in *2023 31st Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2023.

[34] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, "Robust compressed sensing mri with deep generative priors," in *Adv. Neural Inf. Process. Sys.*, vol. 34, pp. 14938–14954, 2021.

[35] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated mri," *Med. Image Anal.*, vol. 80, p. 102479, 2022.

[36] G. Luo, M. Heide, and M. Uecker, "Mri reconstruction via data driven markov chain with joint uncertainty estimation," *arXiv:2202.01479*, 2022.

[37] Y. Xie and Q. Li, "Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction," *arXiv:2203.03623*, 2022.

[38] C. Peng, P. Guo, S. K. Zhou, V. Patel, and R. Chellappa, "Towards performant and reliable undersampled mr reconstruction via diffusion model sampling," in *Med. Image Comput. Comput. Assist. Int.*, pp. 623–633, 2022.

[39] C. Cao, Z.-X. Cui, Y. Wang, S. Liu, T. Chen, H. Zheng, D. Liang, and Y. Zhu, "High-frequency space diffusion model for accelerated mri," *IEEE Trans. Med. Imaging*, vol. 43, no. 5, pp. 1853–1865, 2024.

[40] A. Güngör, S. U. Dar, Ş. Öztürk, Y. Korkmaz, H. A. Bedel, G. Elmas, M. Ozbey, and T. Çukur, "Adaptive diffusion priors for accelerated mri reconstruction," *Medical image analysis*, vol. 88, p. 102872, 2023.

[41] H. Peng, C. Jiang, J. Cheng, M. Zhang, S. Wang, D. Liang, and Q. Liu, "One-shot generative prior in hankel-k-space for parallel imaging reconstruction," *IEEE Trans. Med. Imaging*, pp. 1–1, 2023.

[42] H. Chung, B. Sim, and J. C. Ye, "Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction," in *Comput. Vis. Pattern Recognit.*, pp. 12413–12422, 2022.

[43] Z.-X. Cui, C. Cao, Y. Wang, S. Jia, J. Cheng, X. Liu, H. Zheng, D. Liang, and Y. Zhu, "Spirit-diffusion: Self-consistency driven diffusion model for accelerated mri," 2024.

[44] B. Ozturkler, M. Mardani, A. Vahdat, J. Kautz, and J. M. Pauly, "Regularization by denoising diffusion process for MRI reconstruction," in *NeurIPS Workshop Deep Learn. Inv. Prob.*, 2023.

[45] A. Bansal, E. Borgnia, H.-M. Chu, J. S. Li, H. Kazemi, F. Huang, M. Goldblum, J. Geiping, and T. Goldstein, "Cold diffusion: Inverting arbitrary image transforms without noise," in *NeurIPS*, pp. 41259–41282, 2023.

[46] G. Daras, M. Delbracio, H. Talebi, A. G. Dimakis, and P. Milanfar, "Soft diffusion: Score matching for general corruptions," *Trans. Mach. Learn. Res.*, 2023.

[47] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, "I^2SB: Image-to-Image Schrödinger Bridge," in *Int. Conf. Mach. Learn.*, pp. 22042–22062, 2023.

[48] Z. Fabian, B. Tinaz, and M. Soltanolkotabi, "Diracdiffusion: Denoising and incremental reconstruction with assured data-consistency," *arXiv:2303.14353*, 2023.

[49] M. Delbracio and P. Milanfar, "Inversion by direct iteration: An alternative to denoising diffusion for image restoration," *Tran. Mach. Learn. Res.*, 2023.

[50] H. Chung, J. Kim, and J. C. Ye, "Direct diffusion bridge using data consistency for inverse problems," *arXiv:2305.19809*, 2023.

[51] S. Lee, H. Chung, J. Kim, and J. C. Ye, "Progressive deblurring of diffusion models for coarse-to-fine image synthesis," *arXiv:2207.11192*, 2022.

[52] Y. Balaji, S. Nah, X. Huang, A. Vahdat, J. Song, K. Kreis, M. Aittala, T. Aila, S. Laine, B. Catanzaro, *et al.*, "ediffi: Text-to-image diffusion models with an ensemble of expert denoisers," *arXiv:2211.01324*, 2022.

[53] J. Huang, A. I. Aviles-Rivero, C.-B. Schönlieb, and G. Yang, "Cdiffmr: Can we replace the gaussian noise with k-space undersampling for fast mri?," in *Med. Image Comput. Comput. Assist. Int.*, pp. 3–12, 2023.

[54] M. U. Mirza, F. Arslan, and T. Çukur, "Super resolution mri via upscaling diffusion bridges," in *2024 32nd Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2024.

[55] M. U. Mirza, O. Dalmaz, H. A. Bedel, G. Elmas, Y. Korkmaz, A. Gungor, S. U. Dar, and T. Çukur, "Learning fourier-constrained diffusion bridges for mri reconstruction," *arXiv preprint arXiv:2308.01096*, 2023.

[56] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons, 2014.

[57] P. T. Callaghan, *Principles of nuclear magnetic resonance microscopy*. Clarendon press, 1993.

[58] 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.

[59] 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.

[60] K. T. Kwon, H. H. Wu, T. Shin, T. Çukur, M. Lustig, and D. G. Nishimura, "Three-dimensional magnetization-prepared imaging using a concentric cylinders trajectory," *Magnetic resonance in medicine*, vol. 71, no. 5, pp. 1700–1710, 2014.

[61] 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.

[62] 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.

[63] 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.

[64] 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.

[65] 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.

[66] 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.

[67] 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.

[68] 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.

[69] K. Keskin, U. Yılmaz, and T. Çukur, "Constrained ellipse fitting for efficient parameter mapping with phase-cycled bssfp mri," *IEEE Transactions on Medical Imaging*, vol. 41, no. 1, pp. 14–26, 2021.

[70] 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.

[71] F. Bloch, "Nuclear induction," *Physical review*, vol. 70, no. 7-8, p. 460, 1946.

[72] E. M. Purcell, H. C. Torrey, and R. V. Pound, "Resonance absorption by nuclear magnetic moments in a solid," *Physical review*, vol. 69, no. 1-2, p. 37, 1946.

[73] E. Celik, U. Keles, I. 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.

[74] M. Shahdloo, E. Çelik, and T. Çukur, "Biased competition in semantic representation during natural visual search," *Neuroimage*, vol. 216, p. 116383, 2020.

[75] 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.

[76] R. N. Bracewell, C. Cherry, J. F. Gibbons, W. W. Harman, H. Heffner, E. W. Herold, J. G. Linvill, S. Ramo, and R. A. Rohrer, "Mcgraw-hill seires in electrical and computer engineering," 2000.

[77] R. Brown, Y. Cheng, E. Haacke, M. Thompson, and R. Venkatesan, "Magnetic resonance imaging," *Physical Principles and Sequence Design*, 1999.

[78] A. Webb, *Introduction to biomedical imaging*. John Wiley & Sons, 2022.

[79] 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.

[80] 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.

[81] T. Cukur, 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.

[82] 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.

[83] 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.

[84] 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.

[85] 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.

[86] 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.

[87] 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.

[88] 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.

[89] 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.

[90] 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.

[91] S. Ilbey, C. B. Top, A. Güngör, T. Çukur, E. U. Saritas, 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 IJMPI*, vol. 3, no. 1, 2017.

[92] 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.

[93] 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.

[94] 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.

[95] 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. 565976, 2021.

[96] Z.-P. Liang and P. C. Lauterbur, *Principles of magnetic resonance imaging*. SPIE Optical Engineering Press Bellingham, 2000.

[97] 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.

[98] 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.

[99] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "Sense: sensitivity encoding for fast mri," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.

[100] 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: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.

[101] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," *Journal of the ACM (JACM)*, vol. 58, no. 3, pp. 1–37, 2011.

[102] 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.

[103] T. Çukur, "Accelerated phase-cycled ssfp imaging with compressed sensing," *IEEE transactions on medical imaging*, vol. 34, no. 1, pp. 107–115, 2014.

[104] M. Shahdloo, E. Ilicak, M. Tofighi, 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.

[105] 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.

[106] 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.

[107] S. U. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, and T. Cukur, "Image synthesis in multi-contrast mri with conditional generative adversarial networks," *IEEE transactions on medical imaging*, vol. 38, no. 10, pp. 2375–2388, 2019.

[108] 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, 2018.

[109] D. Lee, J. Yoo, S. Tak, and J. C. Ye, "Deep residual learning for accelerated MRI using magnitude and phase networks," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 9, pp. 1985–1995, 2018.

[110] M. Acar, T. Çukur, and İ. Öksüz, "Segmentation-aware mri reconstruction," in *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 53–61, Springer, 2022.

[111] F. Chen, V. Taviani, I. Malkiel, J. Y. Cheng, J. I. Tamir, J. Shaikh, S. T. Chang, C. J. Hardy, J. M. Pauly, and S. S. Vasanawala, "Variable-density single-shot fast spin-echo mri with deep learning reconstruction by using variational networks," *Radiology*, vol. 289, no. 2, pp. 366–373, 2018.

[112] 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.

[113] 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.

[114] M. Yurt, S. U. Dar, A. Erdem, E. Erdem, K. K. Oguz, and T. Çukur, "mustgan: multi-stream generative adversarial networks for mr image synthesis," *Medical image analysis*, vol. 70, p. 101944, 2021.

[115] M. Yurt, O. Dalmaz, S. Dar, M. Ozbey, B. Tınaz, K. Oguz, and T. Çukur, "Semi-supervised learning of mri synthesis without fully-sampled ground truths," *IEEE Transactions on Medical Imaging*, vol. 41, no. 12, pp. 3895–3906, 2022.

[116] H. A. Bedel, I. Sivgin, O. Dalmaz, S. U. Dar, and T. Çukur, "Bolt: Fused window transformers for fmri time series analysis," *Medical image analysis*, vol. 88, p. 102841, 2023.

[117] I. Sivgin, H. A. Bedel, S. Ozturk, and T. Çukur, "A plug-in graph neural network to boost temporal sensitivity in fmri analysis," *IEEE Journal of Biomedical and Health Informatics*, 2024.

[118] O. F. Atli, B. Kabas, F. Arslan, M. Yurt, O. Dalmaz, and T. Çukur, "I2i-mamba: Multi-modal medical image synthesis via selective state space modeling," *arXiv preprint arXiv:2405.14022*, 2024.

[119] H. A. Bedel and T. Çukur, "DreaMR: Diffusion-driven counterfactual explanation for functional MRI," *arXiv:2307.09547*, 2023.

[120] V. A. Nezhad, G. Elmas, F. Arslan, B. Kabaş, and T. Çukur, "Generalizable deep mri reconstruction with cross-site data synthesis," in *2024 32nd Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2024.

[121] B. Kabas, F. Arslan, V. A. Nezhad, and T. Çukur, "Multi-contrast mr image synthesis with a brownian diffusion model," in *2024 32nd Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2024.

[122] F. Arslan, M. B. Yılmaz, and T. Cukur, "Robust brain tumor segmentation with deep residual supervision and mixed precision training," in *2024 32nd Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, IEEE, 2024.

[123] O. Dalmaz, B. Saglam, G. Elmas, M. Mirza, and T. Çukur, "Denoising diffusion adversarial models for unconditional medical image generation," in *2023 31st Signal Processing and Communications Applications Conference (SIU)*, pp. 1–5, IEEE, 2023.

[124] O. Dalmaz, U. Mirza, G. Elmas, M. Özbey, S. U. H. Dar, and T. Çukur, "A specificity-preserving generative model for federated mri translation," in *Distributed, Collaborative, and Federated Learning, and Affordable AI and Healthcare for Resource Diverse Global Health* (S. Albarqouni, S. Bakas, S. Bano, M. J. Cardoso, B. Khanal, B. Landman, X. Li, C. Qin, I. Rekik, N. Rieke, H. Roth, D. Sheet, and D. Xu, eds.), (Cham), pp. 79–88, Springer Nature Switzerland, 2022.

[125] H. A. Bedel, I. Şıvgın, and T. Çukur, "A graphical network layer for lagged analysis of fmri data," in *2022 30th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, 2022.

[126] O. Dalmaz, B. Sağlam, K. Gönç, S. U. Dar, and T. Çukur, "Bottleneck sharing generative adversarial networks for unified multi-contrast mr image

synthesis," in *2022 30th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, 2022.

[127] M. U. Mirza, O. Dalmaz, and T. Çukur, "Skip connections for medical image synthesis with generative adversarial networks," in *2022 30th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, 2022.

[128] Y. Korkmaz, M. Yurt, S. U. H. Dar, M. Özbey, and T. Cukur, "Deep mri reconstruction with generative vision transformers," in *Machine Learning for Medical Image Reconstruction* (N. Haq, P. Johnson, A. Maier, T. Würfl, and J. Yoo, eds.), (Cham), pp. 54–64, Springer International Publishing, 2021.

[129] 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.

[130] 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*, vol. 26, no. 9, pp. 4679–4690, 2022.

[131] 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," *Journal of Neuroscience*, vol. 42, no. 35, pp. 6782–6799, 2022.

[132] F. T. Yarman Vural, S. D. Newman, T. Çukur, and I. Önal Ertugrul, "Machine learning methods for human brain imaging," 2023.

[133] Ş. Öztürk, E. Çelik, and T. Çukur, "Content-based medical image retrieval with opponent class adaptive margin loss," *Information Sciences*, vol. 637, p. 118938, 2023.

[134] Q. Zhang, T. Cukur, H. Greenspan, and G. Yang, "Generative adversarial networks in cardiovascular research," 2023.

[135] I. Aytekin, O. Dalmaz, K. Gonc, H. Ankishan, E. U. Saritas, U. Bagci, H. Celik, and T. Çukur, "Covid-19 detection from respiratory sounds with hierarchical spectrogram transformers," *IEEE Journal of Biomedical and Health Informatics*, 2023.

[136] A. Zaid Alkilani, T. Çukur, and E. U. Saritas, "Fd-net: An unsupervised deep forward-distortion model for susceptibility artifact correction in epi," *Magnetic Resonance in Medicine*, vol. 91, no. 1, pp. 280–296, 2024.

[137] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo, "Deep learning for undersampled MRI reconstruction," *Phys. Med. Biol.*, vol. 63, no. 13, p. 135007, 2018.

[138] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed sensing MRI reconstruction with cyclic loss in generative adversarial networks," *IEEE Trans. Med. Imaging*, vol. 37, no. 6, pp. 1488–1497, 2018.

[139] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, and D. Firmin, "DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction," *IEEE Trans. Med. Imaging*, vol. 37, no. 6, pp. 1310–1321, 2018.

[140] 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," *Magn. Reson. Med.*, vol. 80, no. 5, pp. 2188–2201, 2018.

[141] J. Yoon, E. Gong, I. Chatnuntawech, 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.

[142] Q. Liu, Q. Yang, H. Cheng, S. Wang, M. Zhang, and D. Liang, "Highly undersampled magnetic resonance imaging reconstruction using autoencoding priors," *Magn. Reson. Med.*, vol. 83, no. 1, pp. 322–336, 2020.

[143] 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 J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1072–1087, 2020.

[144] P. Guo, Y. Mei, J. Zhou, S. Jiang, and V. M. Patel, "Reconformer: Accelerated mri reconstruction using recurrent transformer," *arXiv:2201.09376*, 2022.

[145] B. Yaman, S. A. H. Hosseini, and M. Akcakaya, "Zero-shot physics-guided deep learning for subject-specific mri reconstruction," in *NeurIPS Workshop Deep Learn. Inv. Prob.*, 2021.

[146] G. Luo, N. Zhao, W. Jiang, E. S. Hui, and P. Cao, "Mri reconstruction using deep bayesian estimation," *Magn. Reson. Med.*, vol. 84, no. 4, pp. 2246–2261, 2020.

[147] M. Z. Darestani and R. Heckel, "Accelerated mri with un-trained neural networks," *IEEE Trans. Comput. Imaging*, vol. 7, pp. 724–733, 2021.

[148] K. P. Slavkova, J. C. DiCarlo, V. Wadhwa, S. Kumar, C. Wu, J. Virostko, T. E. Yankeelov, and J. I. Tamir, "An untrained deep learning method for reconstructing dynamic mr images from accelerated model-based data," *Magn. Reson. Med.*, vol. 89, no. 4, pp. 1617–1633, 2023.

[149] G. Elmas, S. U. Dar, Y. Korkmaz, E. Ceyani, B. Susam, M. Ozbey, S. Avestimehr, and T. Çukur, "Federated learning of generative image priors for mri reconstruction," *IEEE Trans. Med. Imaging*, vol. 42, no. 7, pp. 1996–2009, 2023.

[150] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Adv. Neural Inf. Process. Sys.*, vol. 33, pp. 6840–6851, 2020.

[151] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," in *Int. Conf. Learn. Rep.*, 2022.

[152] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *Int. Conf. Mach. Learn.*, pp. 8162–8171, 2021.

[153] F. Arslan, B. Kabas, O. Dalmaz, M. Ozbey, and T. Çukur, "Self-consistent recursive diffusion bridge for medical image translation," *arXiv:2405.06789*, 2024.

[154] A. Jain, *Fundamentals of Digital Image Processing*. Prentice Hall, 1989.

[155] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," *arXiv:2111.08005*, 2021.

[156] D. L. Ruderman and W. Bialek, "Statistics of natural images: Scaling in the woods," *Phys. Rev. Lett.*, vol. 73, pp. 814–817, 1994.

[157] F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdzalv, 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," *Rad. Artif. Intell.*, vol. 2, no. 1, p. e190007, 2020.

[158] T. Zhang, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "Coil compression for accelerated imaging with cartesian sampling," *Magn. Reson. Med.*, vol. 69, no. 2, pp. 571–582, 2013.

[159] 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," *Magn. Reson. Med.*, vol. 71, no. 3, pp. 990–1001, 2014.

[160] J. P. Haldar and J. Zhuo, "P-loraks: low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magn. Reson. Med.*, vol. 75, no. 4, pp. 1499–1514, 2016.

[161] J. Jiang, Z. He, Y. Quan, J. Wu, and J. Zheng, "Pgiun: Physics-guided implicit unrolling network for accelerated mri," *IEEE Transactions on Computational Imaging*, 2024.

[162] H. Chung and J. C. Ye, "Score-based diffusion models for accelerated mri," *Med. Image Anal.*, vol. 80, p. 102479, 2022.

[163] O. Dalmaz, U. Mirza, G. Elmas, M. Ozbey, S. U. Dar, E. Ceyani, S. Avestimehr, and T. Çukur, "One model to unite them all: Personalized federated learning of multi-contrast mri synthesis," *Med. Image Anal.*, vol. 94, p. 103121, 2022.

[164] L. Xiang, Y. Chen, W. Chang, Y. Zhan, W. Lin, Q. Wang, and D. Shen, "Deep-learning-based multi-modal fusion for fast MR reconstruction," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 7, pp. 2105–2114, 2019.

[165] 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," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3172–3191, 2020.

[166] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "Rare: Image reconstruction using deep priors learned without groundtruth," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1088–1099, 2020.

[167] Y. Korkmaz, T. Cukur, and V. M. Patel, "Self-supervised mri reconstruction with unrolled diffusion models," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 491–501, Springer, 2023.

[168] Z.-X. Cui, C. Cao, S. Liu, Q. Zhu, J. Cheng, H. Wang, Y. Zhu, and D. Liang, "Self-Score: Self-Supervised Learning on Score-Based Models for MRI Reconstruction," *arXiv:2209.00835*, 2022.

[169] A. Aali, M. Arvinte, S. Kumar, and J. I. Tamir, "Solving inverse problems with score-based generative priors learned from noisy data," *arXiv:2305.01166*, 2023.

[170] T. Xiang, M. Yurt, A. B. Syed, K. Setsompop, and A. Chaudhari, "Ddm ^{2}: Self-supervised diffusion mri denoising with generative diffusion models," *arXiv preprint arXiv:2302.03018*, 2023.

[171] M. Özbey, O. Dalmaz, S. U. Dar, H. A. Bedel, Ş. Özturk, A. Güngör, and T. Çukur, "Unsupervised medical image translation with adversarial diffusion models," *IEEE Trans. Med. Imaging*, vol. 42, no. 12, pp. 3524–3539, 2023.

[172] C.-M. Feng, Y. Yan, G. Chen, Y. Xu, Y. Hu, L. Shao, and H. Fu, "Multi-modal transformer for accelerated mr imaging," *IEEE Trans. Med. Imaging*, vol. 42, no. 10, pp. 2804–2816, 2023.

[173] O. Dalmaz, M. Yurt, and T. Çukur, "ResViT: Residual vision transformers for multi-modal medical image synthesis," *IEEE Trans. Med. Imaging*, vol. 41, no. 10, pp. 2598–2614, 2022.