

MRI Reconstruction with Side Information using Diffusion Models

1st Brett Levac
ECE

University of Texas
Austin, USA
blevac@utexas.edu

2nd Ajil Jalal
EECS

University of California
Berkeley, USA
ajiljalal@berkeley.edu

3rd Kannan Ramchandran
EECS

University of California
Berkeley, USA
kannanr@eecs.berkeley.edu

4th Jonathan I. Tamir
ECE

University of Texas
Austin, USA
jtamir@utexas.edu

Abstract—Magnetic resonance imaging (MRI) exam protocols consist of multiple contrast-weighted images of the same anatomy to emphasize different tissue properties. Due to the long acquisition times required to collect fully sampled k-space measurements, it is common to only collect a fraction of k-space for each scan and subsequently solve independent inverse problems for each image contrast. Recently, there has been a push to further accelerate MRI exams using data-driven priors, and generative models in particular, to regularize the ill-posed inverse problem of image reconstruction. These methods have shown promising improvements over classical methods. However, many of the approaches neglect the additional information present in a clinical MRI exam like the multi-contrast nature of the data and treat each scan as an independent reconstruction. In this work we show that by learning a joint Bayesian prior over multi-contrast data with a score-based generative model we are able to leverage the underlying structure between random variables related to a given imaging problem. This leads to an improvement in image reconstruction fidelity over generative models that rely only on a marginal prior over the image contrast of interest.

Index Terms—Magnetic Resonance Imaging (MRI), Multi-Contrast, Compressed sensing, Generative Models, Bayesian reconstruction

I. INTRODUCTION

Magnetic Resonance Imaging (MRI) is a medical imaging technique that has superior soft tissue contrast without the use of ionizing radiation. Unfortunately, MRI acquisition times are slow due to hardware constraints and resolution requirements. To reduce scan times, a common approach is to collect less data and subsequently solve an ill-posed inverse problem for reconstruction. A variety of advancements have been developed to address the inverse problem such as parallel imaging [1], compressed sensing [2], and more recently data-driven methods [3]–[8]. Data driven techniques have shown marked improvements over prior approaches and are closing in on clinical validation [9].

In MRI, different image weightings arise due to a variety of underlying tissue properties. These contrasts can be observed by modifying pulse sequence parameters to increase sensitivity to a desired tissue property. Common image contrasts include

Proton Density (PD), T_1 , and T_2 weightings, which are primarily based on tissue density and relaxation. Additional contrasts like diffusion weighted imaging (DWI) create images sensitive to water diffusion through tissue which can be useful in tumor identification [10]; susceptibility weighted imaging (SWI) creates contrast based on magnetic susceptibility changes in the body that can often occur in the brain due to microbleeds [11]. Each of these contrasts can assist in diagnosing and assessing different underlying health conditions.

In a typical MRI exam, data representing different contrast-weighted images are often serially collected, as illustrated in Fig. 1. Image reconstruction is typically done independently for each contrast. However, there is clearly shared information between image contrasts for the same anatomy which can be exploited to assist the reconstruction of one, or all, of the images. Past approaches have leveraged this basic property through handcrafted priors [12], [13], end-to-end deep learning [14]–[16], and generative models such as GANs [17], [18] to reduce data acquisition and enhance multi-contrast image reconstruction. Methods which utilize handcrafted priors typically rely on regularization terms based on feature similarity in a transform domain. Meanwhile, end-to-end deep learning techniques attempt to couple the reconstruction processes of different images through shared reconstruction weights. Finally, GAN based approaches can range from learning a synthesis function mapping one image contrast to another or even modified Compressed Sensing using Generative Modeling (CSGM) frameworks.

In this work we derive a joint Bayesian prior over multi-contrast data that is learned using score-based generative models [19] and can be used to solve two closely related problems: (1) Joint reconstructions, where multiple image contrasts are under-sampled; and (2) Conditional reconstructions, where one or more of the image contrasts are fully sampled. We specialize our approach to the case where two image contrasts are scanned and show that learning joint priors over contrast pairs can enhance reconstruction for both joint reconstruction and prior image constrained reconstruction without retraining for new sampling trajectories or acceleration levels.

This work was supported by NSF IFML 2019844, NSF CCF-2239687 (CAREER), NIH U24EB029240, ARO W911NF2110117, and Oracle for Research Fellowship.




Fig. 1. Example of a single brain slice with four different contrasts. Proton Density (Top Left), T1 (Top Right), T2-Flair(Bottom Left), T2 (Bottom Right).

II. BACKGROUND

A. Problem Formulation

For a desired vectorized image $x \in \mathbb{C}^N$ the corresponding MRI measurements $y \in \mathbb{C}^M$ are acquired in the Fourier domain (k-space). Letting $\mathcal{F} \in \mathbb{C}^{M \times N}$ denote a Fourier encoding matrix (2D, 3D, Cartesian, or non-Cartesian), the measured data y in the single-coil imaging case is given by

$$y = \mathcal{F}x + \epsilon, \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2 I_M)$ is zero-mean complex-valued white Gaussian noise. In the case of Cartesian sampling, we can write $\mathcal{F} = PF$ where $F \in \mathbb{C}^{N \times N}$ is the multi-dimensional DFT matrix, and $P \in \{0, 1\}^{M \times N}$ is a binary selection (diagonal) matrix.

Multi-contrast imaging entails measuring n_c images of the same anatomical structures each with different contrast weightings. In principle, these contrasts will also vary according to imaging orientation (sagittal, axial, coronal). The measurements for the i^{th} contrast are described by

$$y_i = \mathcal{F}_i x_i + \epsilon_i, \quad (2)$$

where $\mathcal{F}_i$ are independent Fourier encoding operators (potentially different sampling schemes) for each contrast, x_i is the i^{th} contrast image, and $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2 I_M)$ is an additive noise term independent between contrasts. For the remainder of this work, we will without loss of generality consider the case of two aligned contrast images ($n_c = 2$) with $\sigma_1 = \sigma_2 = \sigma$ and refer to the forward operators as $\mathcal{F}_1$ and $\mathcal{F}_2$. Under this setting, we consider two related problems: **(1)** Recover x_1 and x_2 when both sets of measurements (y_1, y_2) are under-sampled; **(2)** Recover x_2 when y_2 is under-sampled and y_1 is fully sampled giving us access to x_1 from a conventional reconstruction.

B. Optimization Approaches

Several prior works have addressed the challenges of joint and conditional reconstructions from multi-contrast data. Before the use of deep learning in medical image reconstruction, most techniques leveraged hand crafted regularization terms to couple the reconstruction process for multiple images with the assumption that there are shared structural details between images (e.g., edges, image gradients, low-rank representation) [12], [13], [20]. Ultimately these methods set out to solve an optimization problem of the following form:

$$x_1^*, x_2^* = \underset{x_1, x_2}{\operatorname{argmin}} \|\mathcal{F}_1 x_1 - y_1\|_2^2 + \|\mathcal{F}_2 x_2 - y_2\|_2^2 + \lambda R(x_1, x_2), \quad (3)$$

where $R(x_1, x_2)$ is a regularization functional that couples the image contrasts and $\lambda > 0$ is the regularization weighting. A subset of prior works justify Eq. 3 under a Bayesian framework with the goal of solving maximum a posteriori (MAP) estimation [12], [21]. Under the forward model in Eq. 2, joint image reconstruction via MAP takes the form:

$$x_1^*, x_2^* = \underset{x_1, x_2}{\operatorname{argmax}} \log p(x_1, x_2 | y_1, y_2) \quad (4)$$

$$= \underset{x_1, x_2}{\operatorname{argmax}} \log p(y_1, y_2 | x_1, x_2) + \log p(x_1, x_2). \quad (5)$$

upon including the log-likelihood model for each measurement process we obtain

$$x_1^*, x_2^* = \underset{x_1, x_2}{\operatorname{argmin}} \|\mathcal{F}_1 x_1 - y_1\|_2^2 + \|\mathcal{F}_2 x_2 - y_2\|_2^2 - \log p(x_1, x_2). \quad (6)$$

The hand-crafted priors used are approximations to the underlying statistical prior and are therefore limited in their performance. It has also been shown that solving for MAP, even with more accurate priors, can lead to biased solutions [22]. Approaches which admit optimization problems different than Eq. 3 include dictionary learning [23] and physics constrained reconstructions [24]. Joint dictionary learning builds upon single image dictionary learning [25] by creating a coupled dictionary between multi-contrast images. Knowledge about signal dynamics can also be used to constrain the multi-contrast optimization problem. This is done by creating a set of temporal basis functions from a dictionary of signal decay curves and subsequently estimating basis coefficients rather than individual multi-contrast images [26], [27].

C. Deep Learning Approaches

More recently, deep learning approaches for MRI reconstruction have been extended to handle multi-contrast data. Many of these methods use end-to-end training with shared weights between multi-contrast data channels to enhance the sharing of information during either prior image constrained reconstruction [14], [15] or joint reconstruction [16].

Another class of deep learning approaches have employed the use of generative models. These methods can be separated into **(1)** prior-image constrained reconstruction and

(2) prior-image constrained synthesis. Methods for (1) are heavily inspired by the Compressed Sensing using Generative Models (CSGM) framework proposed for more general inverse problems [28]. One such approach attempts to reconstruct high fidelity target images from low resolution target data and high resolution reference images of an alternative image contrast [17]. Later works leverage a StyleGAN architecture to reconstruct from under-sampled target data and fully sampled reference images [18]. Techniques addressing (2) aim to create a target image with access only to a reference image using conditional GANs [29], [30], CycleGANs [31], or diffusion models [32]. Our proposed method is of type (1) but can also be used for purely synthesis problems.




Fig. 2. Problem settings. (Top) Joint reconstruction with under-sampled measurements from both contrasts. (Bottom) Single image reconstruction with access to a fully sampled reference scan and under-sampled measurements from target scan.

D. Score Models

Recent work in deep learning has led to impressive progress in the area of generative modeling [33], [34], and in par-

ticular, score-based generative models [19]. These models aim to approximate the score of a distribution p , defined as $\nabla_x \log p(x)$, by training a neural network via denoising score matching with access only to i.i.d training samples from the high-dimensional target distribution [19], [35]. These models produce new samples approximately from p by initializing $x^0 \sim \mathcal{N}(0, I_N)$, and following annealed Langevin dynamics for T steps:

$$x^{t+1} = x^t + \eta^t \nabla_x \log p(x^t) + z^t, \quad (7)$$

where $z^t \sim \mathcal{N}(0, \sqrt{\eta^t} I_N)$ and η^t is the variance of the additive noise at each step. As $\eta^t \rightarrow 0, T \rightarrow \infty$, the distribution of x^T converges to p [19]. We can write similar updates for the posterior distribution, $p(x|y)$. If we know the likelihood distribution $y|x$, then using Bayes rule we can sample via:

$$x^{t+1} = x^t + \eta^t \nabla_{x^t} (\log p(y|x^t) + \log p(x^t)) + z^t. \quad (8)$$

These models have been used to solve inverse problems such as compressed sensing MRI [5]–[7], [36].

III. METHODS

A. Joint Reconstruction (Case 1)

Given under-sampled measurements for both images we can formulate our joint reconstruction algorithm as follows:

$$x_1^{t+1} = x_1^t + \eta^t \nabla_{x_1^t} \log p(x_1^t, x_2^t | y_1, y_2) + z_1^t, \quad (9)$$

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} \log p(x_1^t, x_2^t | y_1, y_2) + z_2^t. \quad (10)$$

Applying Bayes rule gives

$$x_1^{t+1} = x_1^t + \eta^t \nabla_{x_1^t} (\log p(y_1, y_2 | x_1^t, x_2^t) + \log p(x_1^t, x_2^t)) + z_1^t, \quad (11)$$

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} (\log p(y_1, y_2 | x_1^t, x_2^t) + \log p(x_1^t, x_2^t)) + z_2^t. \quad (12)$$

By substituting the approximate likelihood equations and the neural network score function approximator $(s_\theta(x_1, x_2))_i \approx \nabla_{x_i} \log p(x_i, x_j)$ we obtain

$$x_1^{t+1} = x_1^t + \eta^t \left(\frac{\mathcal{F}_1^H(\mathcal{F}_1 x_1^t - y_1)}{\gamma_t^2 + \sigma^2} + s_\theta(x_1^t, x_2^t; \beta_t)_1 \right) + z_1^t, \quad (13)$$

$$x_2^{t+1} = x_2^t + \eta^t \left(\frac{\mathcal{F}_2^H(\mathcal{F}_2 x_2^t - y_2)}{\gamma_t^2 + \sigma^2} + s_\theta(x_1^t, x_2^t; \beta_t)_2 \right) + z_2^t, \quad (14)$$

where β_t denotes the current noise level during inference and γ_t is an annealing term which is included to try and account for the mismatch between the likelihood $\nabla_x \log p(y|x)$, which we know exactly, and $\nabla_{x^t} \log p(y|x^t)$.

B. Conditional Reconstruction (Case 2)

In Case 2 our goal is to conduct posterior sampling conditioned on under-sampled k-space data (y_2) and the reconstruction from a fully-sampled prior scan (i.e., x_1). When ϵ_1 is small, we can approximate $x_1^* \approx \mathcal{F}_1^\dagger y_1$, where $\mathcal{F}_1^\dagger$ is the pseudo-inverse of the forward. Written in the Langevin update equations, this task admits the following simplified algorithm:

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} \log p(x_2^t | y_2, x_1^*) + z_2^t. \quad (15)$$

Using Bayes rule we can get the following equation:

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} (\log p(y_2, x_1^* | x_2^t) + \log p(x_2^t)) + z_2^t, \quad (16)$$

Looking at the measurement process for both images we see in Eq. 2 that given x_2 the measurements y_2 are only stochastic due to the additive noise ϵ_2 . As noise between MRI scans is independent, y_2 is conditionally independent of x_1 given x_2 . This leaves

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} (\log p(y_2 | x_2^t) + \log p(x_1^* | x_2^t) + \log p(x_2^t)) + z_2^t \quad (17)$$

$$x_2^{t+1} = x_2^t + \eta^t \nabla_{x_2^t} (\log p(y_2 | x_2^t) + \log p(x_1^*, x_2^t)) + z_2^t. \quad (18)$$

Using the analytically available derivative of the likelihood function we can further simplify the update equation as

$$x_2^{t+1} = x_2^t + \eta^t \left(\frac{\mathcal{F}_2^H (\mathcal{F}_2 x_2^t - y_2)}{\gamma_t^2 + \sigma^2} + \nabla_{x_2^t} \log p(x_1^*, x_2^t) \right) + z_2^t, \quad (19)$$

$$x_2^{t+1} = x_2^t + \eta^t \left(\frac{\mathcal{F}_2^H (\mathcal{F}_2 x_2^t - y_2)}{\gamma_t^2 + \sigma^2} + s_\theta(x_1^* + \lambda_t, x_2^t; \beta_t)_2 \right) + z_2^t, \quad (20)$$

where $\lambda_t \sim \mathcal{N}(0, \beta_t^2 I_n)$ is an additive noise term dependent the noise level at the current inference step during the annealed sampling process. The difference between Eq. 8 and our new update equations (Eq. 14, Eq. 20) is that we now need access to the joint score function $\nabla_{x_i^t} \log p(x_i^t, x_j^t)$.

C. Experiments

For proof-of-principle, we used the publicly available SKM-TEA dataset [37] which contains paired images from double-echo in steady state (DESS) gradient echo knee exams consisting of 3D multi-coil k-space data collected at two different echo times as well as reference complex-valued reconstructions based on parallel imaging. An example of paired images is show in Fig. 3, noting the weaker intensity of the second echo due to signal relaxation. Training data pairs (x_1, x_2) were normalized jointly by the 99th percentile value

of the reference magnitude image for x_1 . We trained two different score models, $s_\theta(x_2)$ and $s_\theta(x_1, x_2)$, which model the marginal and joint score functions respectively. The input sizes to the marginal and joint score networks were 2 and 4 channels to account for real and imaginary components of one or two images respectively. We used 100 slices from 150 different volumes (15,000 slices in total) to train the networks. The networks were trained for 80,000 steps with a batch size of 4. Representative random samples taken via the annealed Langevin dynamics update equations (Eq. 7) [19] are shown for both the marginal and joint score models in Fig. 4.

Test data was synthesized by Fourier transforming the reference reconstruction images and corrupting with additive Gaussian noise using $\sigma = 0.01 \cdot \max(x_1^{GT})$, and retrospectively under-sampling their Fourier coefficients with acceleration levels of $R = 3, 4, 5$ using uniform random vertical sampling patterns with fully sampled autocalibration signal (ACS) regions of k-space (examples shown in Figure 2). Although this noise level is low with respect to x_1 , x_2 has lower signal intensity leading to relatively lower SNR in simulated x_2 measurements. We note here that this is an inverse crime since test data was synthesized from single-coil images reconstructed from a 3D multi-coil acquisition [38].

In the joint reconstruction scenario, we sampled both x_1 and x_2 at the prescribed acceleration level by collecting the same 12 ACS lines and independent uniform random sampling for the remaining phase encode (PE) lines in each image contrast (Fig. 2). For the prior image conditioned reconstruction we only took under-sampled measurements from the target contrast x_2 by, again, collecting an ACS region and uniform random sampling for the remaining PE lines (Fig. 2). The test set included 100 slices from three volumes (300 slices in total) not seen in the training or validation datasets. For each test image we conducted five different reconstructions: **(1)** Zero-filled IFFT (ZF), **(2)** Sample x_2 given only the fully sampled prior image x_1 ($E[x_2 | x_1]$), **(3)** Sample x_2 using only its own under-sampled measurements y_2 ($E[x_2 | y_2]$), **(4)** Sample x_2 from under-sampled measurements of both contrasts ($E[x_1, x_2 | y_1, y_2]$), **(5)** Sample x_2 from under-sampled measurements y_2 and fully sampled prior image x_1 ($E[x_2 | y_2, x_1]$). In each case we averaged 10 posterior samples for the reconstruction. The step size η_t was tuned individually for each method using a single validation patient with 100 samples at $R = 4$.

To compare with end-to-end deep learning methods for joint reconstruction we trained an unrolled network based on MoDL [4] which shared a 4 channel denoising network between the two contrasts (MC-MoDL). We trained the network for the joint reconstruction task by retrospectively under-sampling both contrasts at an acceleration level of $R = 3$ as described above. This technique serves as a point of comparison to other unrolled end-to-end methods [16].

After all the models had been trained we conducted three different experiments: **(1)** joint reconstruction at various acceleration levels, **(2)** joint reconstruction with altered forward operator (horizontal sampling mask), **(3)** and prior image

guided reconstruction.

D. Quantitative Metrics

To evaluate reconstruction quality we used SSIM as defined in [39] (higher is better) and normalized root mean squared error (NRMSE) between ground truth and reconstructed image pairs (lower is better), where NRMSE is defined between a ground truth image x^{GT} and a reconstructed image x^* as

$$NRMSE = \frac{\|x^{GT} - x^*\|_2}{\|x^{GT}\|_2}. \quad (21)$$

Results in tables report average loss values across the entire test set for each metric.




Fig. 3. Example of an image pair in the publicly available dataset used for training and testing [37]. (Left) x_1 , (Right) x_2 . The dynamic range for both subfigures is the same.




Fig. 4. Two prior samples from marginal score model (Top), One prior sample from joint score model (Bottom).

IV. RESULTS

Numerical results for reconstructing x_2 in the joint reconstruction case are shown in Table I for each approach described above. Example reconstructions for $R = 3$ are shown in Fig. 5. The top and bottom images in the left most column

are x_2^{GT} and x_1^{GT} respectively. The remaining columns each represent a different reconstruction technique as described in Section III-C. For each reconstruction technique the top row is the reconstructed image and the bottom row contains the difference image between the reconstructed image and the ground truth image. All images are plotted on the same dynamic range and the difference images are scaled by a factor of 10. Table II shows the reconstruction metrics for when we apply the joint reconstruction techniques to data that has been sampled with a different pattern at $R = 3$. Figure 6 shows an example reconstruction for this case. Results for the prior guided image reconstruction experiment where we have access to a fully sampled prior image x_1^{GT} are shown in Table III.

V. DISCUSSION

From our results we observe in Table I that using joint priors ($E[x_1, x_2|y_1, y_2]$) via score-based generative modeling for multi-contrast image reconstruction improves reconstruction quality when compared to marginal priors ($E[x_2|y_2]$). In the same set of experiments we are able to match the performance of end-to-end techniques (MC-MoDL) in-distribution (Table I). Our technique vastly outperforms end-to-end techniques when a distribution shift occurs in the forward model as seen in Table II. This is because end-to-end techniques are trained with access to a limited distribution of forward operators at training time and thus are known to generalize poorly to alterations in the forward model at test time. This is also shown in Table III where we give MC-MoDL access to fully sampled ground truth reference images x_1^{GT} and its performance dips sharply even compared to the case where y_1 is under-sampled by a factor of 3. This, again, can be attributed to the poor generalization of end-to-end techniques in the presence of distribution shifts. Our technique does not see the forward operator at training time and thus does not experience the shift at test time. We emphasize that for both joint reconstruction ($E[x_1, x_2|y_1, y_2]$) and prior image guided reconstruction ($E[x_2|y_2, x_1]$) we need not retrain our generative prior. This robustness to forward model shift becomes increasingly important in multi-contrast imaging since the degrees of freedom with which clinicians can alter acquisition scan parameters (alter the forward model) increases rapidly and training an end-to-end model for each use case is impractical.

It should be noted that our experiments were conducted on distributions which already contained inherent noise and thus our image reconstruction metrics were inherently low even for $R = 3$. Additionally, for the specific dataset used it is not realistic to only under-sample one DESS echo. Our joint reconstruction method, as with other joint reconstruction techniques, requires the reconstruction of images after all image contrasts of interest are collected. This has the possibility to inhibit many clinical imaging sessions as most exams use intermittent image reconstructions throughout the exam to guide later scan protocol selections. This issue, however, could be mitigated by first starting with an unconditioned reconstruction and iteratively conducting heavier conditioned reconstructions as more scans are completed.




Fig. 5. Reconstruction of a slice from the test set at $R = 3$ with vertical sampling. Left-most column shows the ground-truth (GT) x_2 and x_1 . Top: reconstruction of x_2 given different side information. Bottom: Difference image with ground-truth.

TABLE I
VERTICAL SAMPLING RECONSTRUCTION METRICS

R	ZF		$E[x_2 y_2]$		$E[x_2 x_1]$		$E[x_1, x_2 y_1, y_2]$		MC-MoDL	
	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑
3	0.416	0.707	0.230	0.841	0.429	0.720	0.205	0.862	0.205	0.872
4	0.446	0.690	0.265	0.801	0.429	0.720	0.236	0.834	0.236	0.849
5	0.463	0.681	0.295	0.784	0.429	0.720	0.265	0.811	0.265	0.829




Fig. 6. Reconstruction of a slice from the test set using a horizontal sampling pattern at $R = 3$.

VI. CONCLUSION

Multi-contrast imaging is ubiquitous in clinical MRI exams due to the large variety of tissue parameters and their differing sensitivities to pathology. We have shown that by leveraging shared statistical information between images via a learned Bayesian prior we can achieve marked improvements over state-of-the-art reconstruction techniques for single contrast reconstructions and even end-to-end joint reconstruction techniques. A clear extension can be made to longitudinal imaging

TABLE II
HORIZONTAL SAMPLING RECONSTRUCTION METRICS

R	ZF		$E[x_1, x_2 y_1, y_2]$		MC-MoDL	
	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑
3	0.465	0.638	0.205	0.856	0.304	0.796

TABLE III
PRIOR IMAGE GUIDED RECONSTRUCTION METRICS

R	ZF		$E[x_2 y_2, x_1]$		MC-MoDL	
	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑	NRMSE↓	SSIM↑
3	0.416	0.707	0.184	0.878	0.234	0.861
4	0.446	0.690	0.200	0.861	0.269	0.842
5	0.463	0.681	0.211	0.850	0.297	0.828

scenarios where patients undergo scans separated by a period of time to track disease progress [40]. Our technique could be used to condition reconstruction at a present study on prior examinations. We emphasize that this work has the potential to be applied to settings where we wish to sample from $p(x|y, c)$ where c is not necessarily another image but rather additional side information (age, pre-existing conditions, etc.).

REFERENCES

- [1] Klaas P. Pruessmann, Markus Weiger, Markus B. Scheidegger, and Peter Boesiger, "Sense: Sensitivity encoding for fast mri," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.

- [2] Michael Lustig, David Donoho, and John M Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [3] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll, "Learning a variational network for reconstruction of accelerated mri data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.
- [4] Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob, "Modl: Model-based deep learning architecture for inverse problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
- [5] Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jonathan I Tamir, "Robust compressed sensing mri with deep generative priors," *Advances in Neural Information Processing Systems*, 2021.
- [6] Hyungjin Chung and Jong Chul Ye, "Score-based diffusion models for accelerated mri," *Medical Image Analysis*, vol. 80, pp. 102479, 2022.
- [7] Guanxiong Luo, Moritz Blumenthal, Martin Heide, and Martin Uecker, "Bayesian mri reconstruction with joint uncertainty estimation using diffusion models," *Magnetic Resonance in Medicine*, pp. 1–17, 2023.
- [8] Matthew Bendel, Rizwan Ahmad, and Philip Schniter, "A regularized conditional gan for posterior sampling in inverse problems," 2022.
- [9] Patricia M. Johnson, Dana J. Lin, Jure Zbontar, C. Lawrence Zitnick, Anuroop Sriram, Matthew Muckley, James S. Babb, Mitchell Kline, Gina Ciavarrà, Erin Alaia, and et al., "Deep learning reconstruction enables prospectively accelerated clinical knee mri," *Radiology*, vol. 307, no. 2, 2023.
- [10] Vinit Baliyan, Chandan J Das, Raju Sharma, and Arun Kumar Gupta, "Diffusion weighted imaging: Technique and applications," *World Journal of Radiology*, vol. 8, no. 9, pp. 785, 2016.
- [11] Sven Haller, E.Mark Haacke, Majda M. Thurnher, and Frederik Barkhof, "Susceptibility-weighted imaging: Technical essentials and clinical neurologic applications," *Radiology*, vol. 299, no. 1, pp. 3–26, 2021, PMID: 33620291.
- [12] Berkin Bilgic, Vivek K Goyal, and Elfar Adalsteinsson, "Multi-contrast reconstruction with bayesian compressed sensing," *Magnetic Resonance in Medicine*, vol. 66, no. 6, pp. 1601–1615, June 2011.
- [13] Matthias J. Ehrhardt and Marta M. Betcke, "Multicontrast mri reconstruction with structure-guided total variation," *SIAM Journal on Imaging Sciences*, vol. 9, no. 3, pp. 1084–1106, 2016.
- [14] Bo Zhou and S. Kevin Zhou, "Dudornet: Learning a dual-domain recurrent network for fast mri reconstruction with deep t1 prior," 2020.
- [15] Lei Xiang, Yong Chen, Weitang Chang, Yiqiang Zhan, Weili Lin, Qian Wang, and Dinggang Shen, "Deep-learning-based multi-modal fusion for fast mr reconstruction," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 7, pp. 2105–2114, 2019.
- [16] Daniel Polak, Stephen Cauley, Berkin Bilgic, Enhao Gong, Peter Bachert, Elfar Adalsteinsson, and Kavin Setsompop, "Joint multi-contrast variational network reconstruction (jvn) with application to rapid 2d and 3d imaging," *Magnetic Resonance in Medicine*, vol. 84, no. 3, pp. 1456–1469, 2020.
- [17] Salman U.H. Dar, Mahmut Yurt, Mohammad Shahdloo, Muhammed Emrullah Ildiz, Berk Tınaz, and Tolga Çukur, "Prior-guided image reconstruction for accelerated multi-contrast mri via generative adversarial networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1072–1087, 2020.
- [18] Varun A Kelkar and Mark A Anastasio, "Prior image-constrained reconstruction using style-based generative models," *arXiv preprint arXiv:2102.12525*, 2021.
- [19] Yang Song and Stefano Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in Neural Information Processing Systems*, pp. 11918–11930, 2019.
- [20] Aurélien Bustin, Gastão Lima da Cruz, Olivier Jaubert, Karina Lopez, René M. Botnar, and Claudia Prieto, "High-dimensionality undersampled patch-based reconstruction (hd-prost) for accelerated multi-contrast mri," *Magnetic Resonance in Medicine*, vol. 81, no. 6, pp. 3705–3719, 2019.
- [21] Matthias J Ehrhardt, Kris Thielemans, Luis Pizarro, David Atkinson, Sébastien Ourselin, Brian F Hutton, and Simon R Arridge, "Joint reconstruction of pet-mri by exploiting structural similarity," *Inverse Problems*, vol. 31, no. 1, pp. 015001, dec 2014.
- [22] Ajil Jalal, Sushrut Karmalkar, Alexandros G. Dimakis, and Eric Price, "Instance-optimal compressed sensing via posterior sampling," 2021.
- [23] Pingfan Song, Lior Weizman, João F. C. Mota, Yonina C. Eldar, and Miguel R. D. Rodrigues, "Coupled dictionary learning for multi-contrast mri reconstruction," *IEEE Transactions on Medical Imaging*, vol. 39, no. 3, pp. 621–633, 2020.
- [24] Jonathan I. Tamir, Frank Ong, Suma Anand, Ekin Karasan, Ke Wang, and Michael Lustig, "Computational mri with physics-based constraints: Application to multicontrast and quantitative imaging," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 94–104, 2020.
- [25] Saiprasad Ravishankar and Yoram Bresler, "Mr image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Transactions on Medical Imaging*, vol. 30, no. 5, pp. 1028–1041, 2011.
- [26] Jonathan I. Tamir, Martin Uecker, Weitian Chen, Peng Lai, Marcus T. Alley, Shreyas S. Vasanawala, and Michael Lustig, "T2 shuffling: Sharp, multicontrast, volumetric fast spin-echo imaging," *Magnetic Resonance in Medicine*, vol. 77, no. 1, pp. 180–195, 2017.
- [27] Anthony G. Christodoulou, Jaime L. Shaw, Christopher Nguyen, Qi Yang, Yibin Xie, Nan Wang, and Debiao Li, "Magnetic resonance multitasking for motion-resolved quantitative cardiovascular imaging," *Nature Biomedical Engineering*, vol. 2, no. 4, pp. 215–226, Apr. 2018.
- [28] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis, "Compressed sensing using generative models," in *Proceedings of the 34th International Conference on Machine Learning*, Doina Precup and Yee Whye Teh, Eds. 06–11 Aug 2017, vol. 70 of *Proceedings of Machine Learning Research*, pp. 537–546, PMLR.
- [29] Salman UH. Dar, Mahmut Yurt, Levent Karacan, Aykut Erdem, Erkut Erdem, and Tolga Çukur, "Image synthesis in multi-contrast mri with conditional generative adversarial networks," *IEEE Transactions on Medical Imaging*, vol. 38, no. 10, pp. 2375–2388, 2019.
- [30] Hongwei Li, Johannes C. Paetzold, Anjany Sekuboyina, Florian Kofler, Jianguo Zhang, Jan S. Kirschke, Benedikt Wiestler, and Bjoern Menze, "Diamondgan: Unified multi-modal generative adversarial networks for mri sequences synthesis," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan, Eds., Cham, 2019, Springer International Publishing.
- [31] Karim Armanious, Chenming Jiang, Sherif Abdulatif, Thomas Küstner, Sergios Gatidis, and Bin Yang, "Unsupervised medical image translation using cycle-medgan," in *2019 27th European signal processing conference (EUSIPCO)*, IEEE, 2019, pp. 1–5.
- [32] Muzaffer Özbey, Onat Dalmaz, Salman UH Dar, Hasan A Bedel, Şaban Öztürk, Alper Güngör, and Tolga Çukur, "Unsupervised medical image translation with adversarial diffusion models," 2022.
- [33] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [34] Diederik P Kingma and Max Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [35] Pascal Vincent, "A connection between score matching and denoising autoencoders," *Neural Computation*, vol. 23, no. 7, pp. 1661–1674, 2011.
- [36] Alper Güngör, Salman UH Dar, Şaban Öztürk, Yılmaz Korkmaz, Gökberk Elmas, Muzaffer Özbey, and Tolga Çukur, "Adaptive diffusion priors for accelerated mri reconstruction," 2022.
- [37] Arjun D Desai, Andrew M Schmidt, Elka B Rubin, Christopher M Sandino, Marianne S Black, Valentina Mazzoli, Kathryn J Stevens, Robert Boutin, Christopher Ré, Garry E Gold, Brian A Hargreaves, and Akshay S Chaudhari, "Sk-m-tea: A dataset for accelerated mri reconstruction with dense image labels for quantitative clinical evaluation," 2022.
- [38] Efrat Shimron, Jonathan I. Tamir, Ke Wang, and Michael Lustig, "Implicit data crimes: Machine learning bias arising from misuse of public data," *Proceedings of the National Academy of Sciences*, vol. 119, no. 13, pp. e2117203119, 2022.
- [39] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [40] Uzma Zahid, Emily P Hedges, Mihail Dimitrov, Robin M Murray, Gareth J Barker, and Matthew J Kempton, "Impact of physiological factors on longitudinal structural mri measures of the brain," *Psychiatry Research: Neuroimaging*, vol. 321, pp. 111446, 2022.