

High-Dimensional MR Reconstruction Integrating Subspace and Adaptive Generative Models

Ruiyang Zhao, *Student Member, IEEE*, Xi Peng, *Member, IEEE*, Varun A. Kelkar, Mark A. Anastasio, *Senior Member, IEEE* and Fan Lam, *Senior Member, IEEE*

Abstract—Objective: To develop a new method that integrates subspace and generative image models for high-dimensional MR image reconstruction. **Methods:** We proposed a formulation that synergizes a low-dimensional subspace model of high-dimensional images, an adaptive generative image prior serving as spatial constraints on the sequence of “contrast-weighted” images or spatial coefficients of the subspace model, and a conventional sparsity regularization. A special pretraining plus subject-specific network adaptation strategy was proposed to construct an accurate generative-network-based representation for images with varying contrasts. An iterative algorithm was introduced to jointly update the subspace coefficients and the multi-resolution latent space of the generative image model that leveraged a recently proposed intermediate layer optimization technique for network inversion. **Results:** We evaluated the utility of the proposed method for two high-dimensional imaging applications: accelerated MR parameter mapping and high-resolution MR spectroscopic imaging. Improved performance over state-of-the-art subspace-based methods was demonstrated in both cases. **Conclusion:** The proposed method provided a new way to address high-dimensional MR image reconstruction problems by incorporating an adaptive generative model as a data-driven spatial prior for constraining subspace reconstruction. **Significance:** Our work demonstrated the potential of integrating data-driven and adaptive generative priors with canonical low-dimensional modeling for high-dimensional imaging problems.

Index Terms—Generative models, high-dimensional MR imaging, subspace modeling, GAN inversion, MR parameter mapping, high-resolution MRSI.

This work was supported in part by the following grants: NSF, NSF-CBET, Grant/Award Number: 1944249. NIH, NIBIB, Grant/Award Number: 1R21EB029076A.

Ruiyang Zhao is with the Department of Electrical and Computer Engineering and Beckman Institute for Advanced Science and Technology, University of Illinois Urbana-Champaign, IL 61801, USA.

Xi Peng is with the Department of Radiology, University of Iowa, Iowa City, IA 52242, USA.

Varun A. Kelkar was with the Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, IL 61801, USA, and is currently with the Algorithmic Systems Group at Analog Devices, Inc., Boston, MA 02110, USA.

Mark A. Anastasio and Fan Lam are with the Department of Bioengineering and Beckman Institute for Advanced Science and Technology, University of Illinois Urbana-Champaign, IL 61801 USA (e-mail: fan-lam1@illinois.edu).

I. INTRODUCTION

High-dimensional imaging problems emerge in various MR imaging scenarios, including quantitative MR parameter mapping (qMR) [1], dynamic imaging [2], and MR spectroscopic imaging (MRSI) [3], [4], which introduce additional dimension(s) to encode and decode tissue properties and provide quantitative biomarkers for pathological analysis. While the introduction of additional encoding dimensions enriches the available information, it also significantly lengthens the acquisition time, a major challenge for widespread clinical applications. In recent years, reducing the number of encodings and reconstructing images from sparsely sampled data using different model constraints have been recognized as a common approach to accelerate high-dimensional imaging acquisitions.

Subspace imaging/low-rank tensor modeling is one of the most investigated approaches to enable accelerated high-dimensional MRI. While many variants have been presented, the essence is to exploit the redundancy in the imaging data due to the partial separability property that exists across various high-dimensional imaging modalities [5]–[7]. Partial separability can be induced by the spatial-temporal correlations in dynamic MRI [8]–[10], spatial-parametric correlations in qMR [11]–[13], and spatial-spectral correlations in MRSI [14], which lead to accurate low-rank representations of the multidimensional imaging data arrays. As a result, high-dimensional image reconstruction can be transformed into the recovery of a low-rank matrix/tensor which requires a significantly fewer number of measurements.

Low-rank modeling can be realized either implicitly through low-rankness encouraging regularization of high-dimensional matrices or tensors (e.g., nuclear norm as in [10], [15]–[17]), or explicitly through a matrix factorization form that directly reduces the number of unknowns during image reconstruction (e.g., [7], [14], [18]). Explicit low-rank modeling represents the signal evolution at each voxel as a linear combination of a small number of “temporal basis” with voxel-dependent “spatial coefficients” and enables effective incorporation of other prior knowledge to facilitate image recovery. In particular, special hybrid acquisition and physics-driven subspace learning strategies tailored for different applications can be designed to predetermine the basis [19]–[21]. With the predetermined basis, complementary spatial constraints can be

introduced to improve the estimation of the spatial coefficients. For example, analytical sparsifying transforms have been frequently used, e.g., [12], [22], [23]. Kernel-based methods have been proposed to leverage image features derived from complementary anatomical imaging modalities to re-represent the spatial coefficients, going beyond the traditional voxel basis paradigm [24], [25]. However, these hand-crafted “priors” did not fully exploit the prior information about the unknown images and required carefully feature/kernel selection which can be subjective and sometimes biased.

On a parallel line of pursuit, deep generative image models have provided a new and flexible way to utilize data-driven prior for image representation [26], [27], complementary to subspace modeling. State-of-the-art generative models, such as generative adversarial networks (GANs) [28], allow for automatic learning of feature-based image representations, to generate high-quality, high-resolution images and to constrain image reconstruction. For example, the unknown image can be modeled by a pretrained GAN (using images with similar contrasts and resolutions), and the reconstruction problem can be formulated as estimating a set of latent variables input to the GAN instead of the voxel values [29], [30]. Major challenges of this approach are (1) representation accuracy, i.e., even with fully sampled images as the data, the pretrained GAN may not be able to accurately represent the images; (2) optimization issue, i.e., even if the model is highly expressive, GAN inversion is solving a deep network inversion problem which is highly non-convex and sensitive to initialization.

Different models combined with various optimization strategies have been proposed to address the above mentioned issues. Asim et al. proposed to use the invertible neural networks (INNs) to mitigate the representation error [31]. As INNs use latent vectors with the same dimension as the unknown image, additional regularization can be introduced to improve reconstruction [32]. Image-adaptive GAN (IAGAN) based strategies that update both the lower-dimensional latent vector and weights [30], [33] can reduce data-specific representation error (sometimes referred to as “out of distribution” error) but may lead to overfitting (e.g., to noise and other artifacts). While untrained networks can also be used as an image model (e.g., deep image prior [34], [35]) for updating only weights during reconstruction, such strategies do not fully utilize data-driven prior available in existing datasets, and the effectiveness highly depends on the network structure. StyleGAN offers better representation power with flexible latent space adjustments to account for data distribution variations, which demonstrates great potential in accurately representing images and effectively constraining the reconstruction [36], [37].

In this work, we proposed a new method that integrates subspace model and an adaptive generative image prior for reconstructing high-dimensional MR images. Direct application of generative models as a spatial constraint for subspace imaging can be challenging. Besides the representation accuracy and network inversion issues, training generative models typically requires large-scale, high-resolution, fully sampled datasets, which can be rather difficult to obtain for applications like qMR, dynamic MRI and MRSI etc. To address these issues,

we introduced a pretraining plus subject-specific adaptation strategy for learning an accurate StyleGAN-based image representation. This adaptive prior accounted for contrasts and geometry variations between existing database and application-specific data effectively. Our method does not require end-to-end supervised training using high-dimensional MR data since the GAN can be pretrained on publicly available data, and then adapted to a subject-specific reference image to account for geometry variations. The geometry-adapted GAN representation was incorporated as a spatial constraint into a subspace-based reconstruction formalism. An alternating minimization algorithm was introduced to solve the optimization of jointly updating the spatial coefficients and GAN latent space. We demonstrated the effectiveness of our method using qMR and MRSI data as application examples.

The remaining of the paper is organized as follows: Section II provides some background information on subspace imaging and generative model based image reconstruction. Section III describes the proposed problem formulation and algorithm in details. Section IV presents results evaluating the utility of the proposed method in two application examples, accelerated parameter mapping and MRSI. Section V and VI provide some technical discussion and conclude the paper.

II. BACKGROUND

The acquisition process (after proper discretization) for many high-dimensional MRI problems can be defined as:

$$\mathbf{y} = \mathbf{A}\boldsymbol{\rho} + \mathbf{n}, \quad (1)$$

where $\boldsymbol{\rho} \in \mathbb{C}^{N \times T}$ corresponds to the discretized representation of the underlying high-dimensional image function with N being the number of voxels and T the number of “time” points to be determined, $\mathbf{y} \in \mathbb{C}^{M \times T'}$ denotes the measured data (T' may not necessarily be the same as T), $\mathbf{A}$ and $\mathbf{n}$ represent the encoding operator and measurement noise, respectively. The goal is to recover $\boldsymbol{\rho}$ from measurements $\mathbf{y}$ which can be incomplete and/or noisy.

A. Subspace Reconstruction

While direct reconstruction of the high-dimensional $\boldsymbol{\rho}$ can be rather challenging and sensitive to noise perturbation, low-dimensional subspace models offer a way to significantly reduce the dimensionality of the problem by exploiting “correlations” among different dimensions, enabling image recovery from highly sparsely sampled and/or noisy data. More specifically, if $\boldsymbol{\rho}$ is a Casorati matrix [5] with spatial dimensions cascaded along the first dimension and a temporal/parametric dimension along the second, the subspace model can be realized as a matrix factorization form: $\boldsymbol{\rho} \approx \mathbf{U}\mathbf{V}$, where $\mathbf{U} \in \mathbb{C}^{N \times R}$ and $\mathbf{V} \in \mathbb{C}^{R \times T}$ are low-rank matrices. $\mathbf{U}$ is often referred to as the spatial coefficient matrix and $\mathbf{V}$ as the temporal basis matrix. The model order/rank R , is typically much smaller than the ambient image dimensions (N and T), which means that the number of unknowns/degree-of-freedom (DOF) is significantly reduced.

With the low-rank model, the image reconstruction problem can be formulated as

$$\hat{\mathbf{U}}, \hat{\mathbf{V}} = \arg \min_{\mathbf{U}, \mathbf{V}} \|\mathbf{y} - \mathbf{A}(\mathbf{UV})\|_2^2 + \lambda R(\mathbf{U}, \mathbf{V}), \quad (2)$$

where $R(\cdot)$ is a spatial/spatiotemporal regularization that incorporates prior information such as sparsity [7], [10], piecewise smoothness [38], [39] or structured low rankness [40] etc. This regularization, often including spatial constraints, plays a critical role for the success of subspace reconstruction, as it provides a complementary prior and effectively addresses the instability/ill-conditionness issues related to subspace fitting. Furthermore, for many imaging applications, the temporal basis $\mathbf{V}$ can be pre-determined using a problem-specific hybrid data acquisition strategy [7], [9], [12], [19]–[21]. As a result, the problem in Eq. (2) can be simplified to only estimating $\mathbf{U}$, circumventing the nonconvex problem of jointly solving for $\mathbf{U}$ and $\mathbf{V}$.

B. Generative Model Constrained Image Reconstruction

Recent advances in deep generative models provide a new way to incorporate domain/data-specific prior into image reconstruction problems. With their state-of-the-art high-quality image generation performance, image representations using convolutional networks, GAN and extensions have been used in inverse problems in imaging. Specifically, a common approach is to represent the unknown image of interest $\mathbf{x}$ as $\mathbf{x} = G_{\theta}(\mathbf{w})$, where G is a generative network (can be pre-trained using available domain-specific images or untrained), θ contains the network parameters, and $\mathbf{w}$ is a set of latent variables (often sampled from i.i.d distributions) that typically has a lower dimensionality than $\mathbf{x}$. With this explicit generative model-based image function, the reconstruction problem can be formulated as

$$\hat{\mathbf{w}}, \hat{\theta} = \arg \min_{\mathbf{w}, \theta} \|\mathbf{y} - \mathbf{A}(G_{\theta}(\mathbf{w}))\|_2^2, \quad (3)$$

which seeks to recover the image by determining the latent $\mathbf{w}$ and/or θ instead of solving voxel values in the conventional paradigm. This approach assumes that the unknown image lies in the range space of G_{θ} , which may yield substantial modeling errors, depending on the network structure and capacity. This representation accuracy issue has been observed in the early works of compressed sensing using generative model (CSGM) [29], where a pretrained Deep Convolutional GAN (DCGAN) was used and the reconstruction process only updated $\mathbf{w}$. IAGAN methods [30], [33] addressed this issue by simultaneously updating $\mathbf{w}$ and the network parameters θ to fit the data in Eq. (3). Early stopping is required to avoid overfitting (e.g., to the noisy measurements), particularly when the network has a high capacity. But determining the stopping criterion is heuristic, can be problem dependent and rather challenging.

The development of advanced generative models, such as StyleGAN2, has demonstrated superior image representation power than early generations of GANs, particularly for high-resolution images [36]. This makes it a powerful tool in image

generation and representation. Many works have been proposed to utilize this representation power for image processing and reconstruction. For example, Kelkar et al. leveraged the multi-resolution latents of StyleGAN2 to constrain the reconstructed images to have similar latents with a reference image [37]. In the meantime, the high-dimensional, nonconvex network inversion problem associated with image reconstruction using GAN-based representation remains challenging and computationally expensive. There are also application specific challenges for high-dimensional MR imaging problems that are in need of formulation and algorithmic innovation.

III. PROPOSED METHOD

A. Problem Formulation

We propose to integrate a generative image model as a spatial constraint and the subspace model for high-dimensional image reconstruction via the following formulation:

$$\begin{aligned} \hat{\mathbf{U}}, \{\hat{\mathbf{w}}_t\} = \arg \min_{\mathbf{U}, \{\mathbf{w}_t\}} & \|\mathbf{y} - \mathbf{A}(\mathbf{U}\hat{\mathbf{V}})\|_2^2 \\ & + \sum_{t=1}^{N_t} \lambda_{1,t} \left\| (\mathbf{U}\hat{\mathbf{V}})_t - G_{\hat{\theta}}(\mathbf{w}_t) \right\|_F^2 + \lambda_2 R(\mathbf{U}\hat{\mathbf{V}}). \end{aligned} \quad (4)$$

The first term in Eq. (4) is the well-established subspace/low-rank model with the final reconstruction formed as $\rho = \hat{\mathbf{U}}\hat{\mathbf{V}}$ and $\hat{\mathbf{V}}$ being the predetermined “temporal” basis. The second term introduces a GAN-based image representation as a spatial regularization for images at different time points ($t = 1, 2, \dots, N_t$). The key assumption is that using a properly trained and adapted StyleGAN with fixed parameters in $\hat{\theta}$, the contrast variations in the image sequence can be accurately accounted for by updating only a set of low-dimensional latent space variables $\{\mathbf{w}_t\}$ (referred to as the latents below; Note that $\mathbf{w}_t$ denotes the entire multi-resolution latent variable for the t th image not a particular segment of the latents as described in [37]). This effectively serves as a data-adaptive reference image to constrain the spatial variations. An updating strategy to explore the multi-resolution latent structure and impose similarity for $\{\mathbf{w}_t\}$ at neighboring time points will be specified in the Algorithm section. The final term $R(\cdot)$ is a hand-crafted regularization term, e.g., the commonly used sparsity constraint [7], [12]. The subspace, sparsity and GAN prior play complementary roles to enhance the reconstruction performance.

The key challenges to integrate a GAN-based image prior are: 1) The representation accuracy of the GAN needs to be validated in application-specific context; 2) For many high-dimensional imaging applications, there may be a lack of high-resolution, high-SNR images for training GAN; and 3) Even with a sufficiently accurate representation, solving the non-convex GAN inversion problem can be quite challenging. To address these issues, we proposed a pretraining plus subject-specific adaptation strategy to construct a StyleGAN prior $G_{\hat{\theta}}(\cdot)$ adaptive to different imaging contexts, leveraging reference images readily available in a specific experiment. We also used an intermediate layer optimization (ILO) algorithm




Fig. 1. Validation of the proposed adaptive GAN-based representation: (Row 1) High-quality brain images generated by the pretrained StyleGAN2; (Row 2) Representative images with different contrasts from an independent T_2 mapping experiment (not from the database). The first image served as a reference for the GAN adaptation (see texts). The reference image can be acquired from extra acquisition or initial reconstruction; (Row 3) Images from the range space of the pretrained StyleGAN2 that were closest (in terms of ℓ_2 norm) to the images in Row 2 (relative ℓ_2 errors shown in the upper left corner), from updating only the latents of the StyleGAN2. Noticeable representation errors can be observed, especially in the error maps (Row 4); (Rows 5 and 6) The adapted GAN achieved significantly better representation accuracy. Updating both latents and network parameters yielded an image (left figure) almost the same as the reference image. With the adapted subject-specific network, updating latents only accurately accounted for contrast variations (right three columns in Rows 5&6).

to address the GAN inversion challenge [41]. More details are provided in the subsequent sections.

B. Adaptive Generative Image Model

While it is challenging to directly train a generative model that can accurately capture the geometry and contrast variations in every high-dimensional imaging problems (in many cases such large datasets are not available for training generative networks), a transfer learning plus application-specific adaptation strategy can be developed. Specifically, for a specific organ of interest, we can first train a StyleGAN2 model on large, publicly available data, e.g., the Human Connectome Project (HCP) database for brain imaging applications [42]. We chose StyleGAN2 due to its state-of-the-art performance in generating high-quality, high-resolution images, compared to older generations of DCGANs, and its superior image representation capability due to higher dimensional multi-resolution latents.

With the pretrained model, parameterized by θ_p , we proposed a subject-specific adaptation strategy to construct the

generative image representation for a specific reconstruction task. The key assumptions are 1) there exists a high-quality reference image for network geometry adaptation, and 2) the adapted network can accurately represent images with various contrasts (often acquired in high-dimensional imaging problems) by tuning only the latent space. Given a subject-specific high-resolution reference image $\mathbf{x}_p$ (e.g. from a reference T1w anatomical scan), the adaptation step can be formulated as:

$$\hat{\mathbf{w}}_p, \hat{\theta} = \arg \min_{\mathbf{w}_p, \theta} \|G_{\theta}(\mathbf{w}_p) - \mathbf{x}_p\|_F^2 + \alpha \|\theta - \theta_p\|_F^2, \quad (5)$$

where both the latents $\mathbf{w}_p$ and network parameters θ are updated to minimize the representation error of the network for $\mathbf{x}_p$. The second term is introduced to penalize large deviation of the adapted network from the pretrained one. To reduce the optimization error and fully exploit the latent-space structures, we proposed to solve Eq. (5) with a two-step algorithm, in which we first update the latents $\mathbf{w}_p$ via:

$$\hat{\mathbf{w}}_p = \arg \min_{\mathbf{w}_p} \|G_{\theta_p}(\mathbf{w}_p) - \mathbf{x}_p\|_F^2, \quad (6)$$

and then the network parameters θ via:

$$\hat{\theta} = \arg \min_{\theta} \|G_{\theta}(\hat{\mathbf{w}}_p) - \mathbf{x}_p\|_F^2 + \alpha \|\theta - \theta_p\|_F^2. \quad (7)$$

The latents $\hat{\mathbf{w}}_p$ were solved by the ILO algorithm with fixed network parameters θ_p (details provided below), and then the parameters were updated with fixed latents $\hat{\mathbf{w}}_p$. While multiple iterations can be used in general, we only alternated between Eq. (6) and Eq. (7) once as it already produced a highly accurate representation of the target image and the improvement from more iterations was marginal.

This pretraining plus adaptation strategy produces an adapted GAN prior for which updating only the low-dimensional latents can accurately account for contrast variations, as demonstrated in Fig. 1 using multicontrast brain MR images. The adaptive GAN can serve as a prior in many high-dimensional MR imaging applications, i.e., in this case as a spatial constraint to enhance the subspace reconstruction.

C. Algorithm

With the adapted $G_{\hat{\theta}}(\cdot)$, we proposed to solve the problem in Eq. (4) using an alternating minimization algorithm. Specifically, the original reconstruction problem is decoupled into two subproblems, i.e., one StyleGAN inversion problem and the other GAN-constrained spatial coefficient update problem, which can be mathematically expressed as:

Subproblem (I): Update latents $\mathbf{w}$ by solving

$$\{\hat{\mathbf{w}}_t^{i+1}\} = \arg \min_{\{\mathbf{w}_t\}} \sum_{t=1}^{N_t} \|(\hat{\mathbf{U}}^i \hat{\mathbf{V}})_t - G_{\hat{\theta}}(\mathbf{w}_t)\|_F^2; \quad (8)$$

and




Fig. 2. Illustration of the network partition in the ILO algorithm: The original StyleGAN G_{θ} can be divided into a cascade of two sub-networks, G_{θ_1} (lighter yellow) and G_{θ_2} (darker yellow), with corresponding latents $w_{1,t}$ and $w_{2,t}$. The output of G_{θ_1} , $w_{int,t}$ (originally hidden in G_{θ}), can be viewed as intermediate latents and the input to G_{θ_2} , which will be first updated in the GAN inversion process (see texts for details).

Subproblem(II): Update spatial coefficients $\mathbf{U}$ by solving

$$\begin{aligned} \hat{\mathbf{U}}^{i+1} = \arg \min_{\mathbf{U}} & \left\| \mathbf{y} - \mathbf{A}(\mathbf{U}\hat{\mathbf{V}}) \right\|_2^2 \\ & + \sum_{t=1}^{N_t} \lambda_{1,t} \left\| (\mathbf{U}\hat{\mathbf{V}})_t - G_{\theta}(\hat{\mathbf{w}}_t^{i+1}) \right\|_F^2 + \lambda_2 R(\mathbf{U}\hat{\mathbf{V}}), \end{aligned} \quad (9)$$

where i denotes the iteration number. Since the StyleGAN prior has already been adapted to the subject-specific reference image, for Subproblem (I) in Eq. (8), we minimized the image representation error with respect to the latents only. To address the optimization challenge for GAN inversion, we adapted the ILO algorithm considering the unique latent space structure in StyleGAN (similar to solving Eq. (6)). More specifically, instead of directly solving Eq. (8) w.r.t. $\{\mathbf{w}_t\}$, the loss was first minimized w.r.t. the latents in an intermediate layer and then projected back to the range space of the previous layers. Our procedure is illustrated in Fig. 2.

Mathematically, we can denote the network as a composition of $G_{\theta} = G_{\theta_2} \circ G_{\theta_1}$, with θ_1 and θ_2 containing parameters for different partitions of the network (Fig. 2). The multi-resolution latents in StyleGAN, $\mathbf{w}_t$, can then also be divided into two parts, $\mathbf{w}_{2,t}$ and $\mathbf{w}_{1,t}$ corresponding to the network partition above. With such a partition, we first minimize image representation error by updating $G_{\theta_2}(\cdot)$ w.r.t. a set of higher-resolution latents $\mathbf{w}_{2,t}$ and an intermediate layer output from the previous layers $\mathbf{w}_{int,t}$, i.e.,

$$\begin{aligned} & \{\hat{\mathbf{w}}_{int,t}^{i+1}\}, \{\hat{\mathbf{w}}_{2,t}^{i+1}\} \\ & = \arg \min_{\{\mathbf{w}_{int,t}\}, \{\mathbf{w}_{2,t}\}} \sum_{t=1}^{N_t} \left\| (\hat{\mathbf{U}}^i \hat{\mathbf{V}})_t - G_{\theta_2}(\mathbf{w}_{int,t}, \mathbf{w}_{2,t}) \right\|_F^2. \end{aligned} \quad (10)$$

Subsequently, the update for the lower-resolution latents $\mathbf{w}_{1,t}$ was obtained by fitting G_{θ_1} to the intermediate $\{\hat{\mathbf{w}}_{int,t}^{i+1}\}$

$$\{\hat{\mathbf{w}}_{1,t}^{i+1}\} = \arg \min_{\{\mathbf{w}_{1,t}\}} \sum_{t=1}^{N_t} \left\| \hat{\mathbf{w}}_{int,t}^{i+1} - G_{\theta_1}(\mathbf{w}_{1,t}) \right\|_F^2. \quad (11)$$

This step essentially “projects” the intermediate layer output $\hat{\mathbf{w}}_{int,t}$ to the range space of G_{θ_1} by representing it with the lower resolution latents $\{\hat{\mathbf{w}}_{1,t}^{i+1}\}$. The resulting $\{\hat{\mathbf{w}}_{1,t}^{i+1}\}, \{\hat{\mathbf{w}}_{2,t}^{i+1}\}$ can then be combined into an initialization for the problem in Eq. (8) to further minimize the loss. In this

work, to fully explore the multi-resolution latent structure in StyleGAN2, the network was divided into multiple partitions (with multiple intermediate latents) to break down the original challenging one step deep network inversion problem into multiple steps. Specifically, the minimization starts from a certain intermediate layer depending on the application and then backpropagated to the earlier layers (see the Experiments and Results section).

Additionally, we can assume that images sharing similar features, such as geometry and contrast, are likely to have similar latents (often demonstrated via “style mixing” experiments in the computer vision literature). Exploiting this assumption, we can mitigate potential overfitting to noise/artifacts by introducing a similarity constraint in the latent space across images with different contrasts. For example, we can introduce an l_1 ball constraint on the latents for adjacent “time” points, which can be mathematically formulated as:

$$\hat{\mathbf{w}}_t^{i+1} = \arg \min_{\mathbf{w}_t \in \hat{\mathbf{w}}_{t-1}^{i+1} \oplus B_t(l_t)} \left\| (\hat{\mathbf{U}}^i \hat{\mathbf{V}})_t - G_{\theta}(\mathbf{w}_t) \right\|_F^2. \quad (12)$$

Here $\hat{\mathbf{w}}_{t-1}^{i+1} \oplus B_t(l_t)$ denotes an l_1 ball with radius l_t centered at $\hat{\mathbf{w}}_{t-1}^{i+1}$ from an adjacent “time” point (e.g., images with adjacent contrast weightings in an acquisition). The radius was determined empirically as hyperparameters (see the Discussion and Conclusion Section). Another strategy is to enforce a portion of the latents to be the same for images with various contrasts. Selection of constraints is application dependent and will need to consider the resolution, SNR and contrast variations of the images to be reconstructed (discussed with more details in the Experiments and Results section).

For Subproblem (II), with the updated GAN representation, different regularization choices can be made for $R(\mathbf{U}\hat{\mathbf{V}})$ in Eq. (9) in an application-specific context. This is essentially the well-studied regularized subspace reconstruction, but incorporating additional, reference-adapted network constraints. The initialization of alternating minimization algorithm can be application specific, which will be discussed in the Experiments and Results section.

D. Training and Other Implementation Details

For the generative network, we used the StyleGAN2 architecture proposed in [36] with a latent space dimensionality of 512. Considering the complex-value nature of MRI data, we trained a magnitude image network and a phase image network. These two networks were trained separately due to training data availability consideration (more magnitude images available) and for more flexible training set-up and application in reconstruction. The magnitude network was trained on HCP database [42] with 120,000 T₁-weighted & 120,000 T₂-weighted images. The phase network was trained on NYU fastMRI database [43] with 70,000 phase images after coil combination [44]. The network training, subject-specific adaptation and StyleGAN2 inversion were performed on a Linux server with NVIDIA A40 GPU and implemented in PyTorch 1.12.1. We used Adam optimizer [45] with batch size 80 for training while other hyperparameters remained unchanged as described in [36]. The network training took




Fig. 3. Reconstructed multi-TE images (Left panel) by different methods from an in vivo data set and the corresponding error maps (Right panel) at AF=4. Images for different TEs are shown in different columns while results for different methods in respective rows. Results of Subspace + GAN constraint (Row 4) were produced by solving the proposed reconstruction problem with only the adaptive GAN representation as a spatial constraint (including the first two terms in Eq. (4)). The proposed method (Row 5) integrating all components outperformed the state-of-the-art Subspace+ Joint sparsity method [12] (Row 3) and achieved the lowest error, better visualized in the error images.

about 4 days. The subspace reconstruction and other methods for comparison were implemented in Matlab R2020b. Relative l_2 errors shown in percentage were used for quantitative evaluations throughout this paper.

IV. EXPERIMENTS AND RESULTS

We evaluated the proposed method in two high-dimensional MR imaging applications: quantitative MR parameter mapping and MR spectroscopic imaging. Specific GAN adaptation considering the acquisition design in each imaging scenario, as well as modification to the overall formulation will be discussed. For parameter mapping, we benchmarked the proposed reconstruction against a state-of-the-art subspace method with joint sparsity constraint, termed “Subspace+Joint sparsity” [12]. For MRSI experiments, we benchmarked the proposed method against the SPICE-based reconstruction (SPICE) [14]. We also compared the proposed method to a “Subspace+GAN” reconstruction, which is equivalent to solving the proposed formulation in Eq. (4) with only the first two terms (data consistency + GAN constraint). These comparisons aim to demonstrate both the utility of the adaptive GAN as a new spatial constraint, and the effectiveness of integrating GAN and other complementary priors. All in vivo studies were performed with local IRB approval.

A. Accelerated MR Parameter Mapping

qMR provides quantitative tissue properties beyond traditional subjective and qualitative contrast-weighted images for

diagnosis and disease characterization. One major challenge for clinical applications of qMR is the prolonged acquisition due to the need to acquire multiple images for parameter estimation. In this section, we demonstrate the effectiveness of our method for accelerating qMR (i.e., T_2 mapping) experiments.

In vivo T_2 mapping data were acquired using a multi-spin-echo sequence on a 3T scanner (Siemens Trio) using a 12-channel head coil. Specific acquisition parameters are: 16 echoes with $TE_1 = 8.8$ ms and echo spacing $\Delta TE = 8.8$ ms, $TR = 4000$ ms, slice thickness = 3 mm, matrix size = 192×192 and a FOV of 192×192 mm². The acquired fully-sampled data were retrospectively undersampled using a 1D random phase encoding pattern at different acceleration factors (AFs). The central 12 k -space lines were acquired at all TEs for subspace determination [12], [22], and the central k -space was fully sampled at the first TE with coverage dependent on AFs, e.g., 48 for AF = 4, for coil sensitivity estimation by ESPIRiT [46]. The data and reconstruction workflow for this application is further illustrated in Fig. S1 in the supplementary information.

To obtain a reference image for the proposed subject-specific GAN adaptation, we used a sum-of-squares (SoS) combination of images across all TEs from an initial subspace reconstruction. Based on our observation, the SoS image exhibits negligible aliasing even at high AFs. While additional high-resolution reference images can be acquired and used, e.g., a T_1w image typically acquired in neuroimaging experiments, our strategy might alleviate the need of additional acquisitions. The coil sensitivity maps were integrated into the




Fig. 4. Estimated T_2 maps (Row 1) from the data in Fig. 3 and the corresponding error maps (Row 2, brain region only) for different reconstruction methods (AF=4). The overall T_2 estimation errors are shown in the upper left corners of the images, which were calculated for the brain region only. As can be seen, reconstructions using joint sparsity constraint and GAN constraint yielded similar reconstruction errors. Integrating both sparsity and GAN constraints produced the best result.

forward encoding model for multichannel data and the specific reconstruction formulation becomes:

$$\begin{aligned} \hat{\mathbf{U}}, \{\hat{\mathbf{w}}_t\} = \arg \min_{\mathbf{U}, \{\mathbf{w}_t\}} \sum_{c=1}^{N_c} \left\| \mathbf{y}_c - \Omega(\mathbf{F}\mathbf{S}_c\mathbf{U}\hat{\mathbf{V}}) \right\|_2^2 \\ + \sum_{t=1}^{N_t} \lambda_{1,t} \left\| (\mathbf{U}\hat{\mathbf{V}})_t - \Phi_t \odot \mathbf{G}_{\hat{\theta}}(\mathbf{w}_t) \right\|_F^2 + \lambda_2 \|\mathbf{D}(\mathbf{U}\hat{\mathbf{V}})\|_{2,1}, \end{aligned} \quad (13)$$

where Ω , $\mathbf{F}$ and $\mathbf{S}_c$ denote the sampling operator, Fourier encoding matrix and sensitivity encoding matrix respectively. The additional regularization chosen here was a joint sparsity constraint with $\mathbf{D}$ a finite difference operator [12]. We used magnitude network only and the phase $\{\Phi_t\}$ came from a subspace reconstruction with joint sparsity constraint (Assuming the smooth phase from subspace reconstruction was good enough). The parameters $\{\lambda_{1,t}\}$ can be TE-dependent: we used 0.2 for the last five TEs and 0.1 for the other TEs in the reconstruction using subspace and adaptive GAN constraint alone. When using all three constraints, we simply chose 0.04 for all $\lambda_{1,t}$ and achieved the best reconstruction results. For the joint sparsity constraint, we used a mild choice of $\lambda_2 = 2e^{-7}$ ($\lambda_2 = 1e^{-6}$ when using joint sparsity constraint alone). We summarized the selection of regularization parameters in Table I. The algorithm converged after five iterations. More discussion on parameter selection can be found in the Discussion section.

For this application, we used the l_1 -ball similarity constraint on latent space shown in Eq. (12) for images at adjacent TEs in Subproblem (I). We are not restricting any portion of the

TABLE I

THE CHOICE OF REGULARIZATION PARAMETERS FOR ACCELERATED MR PARAMETER MAPPING

Methods	$\{\lambda_{1,t}\}$	λ_2
Subspace+Joint sparsity	0	$1e^{-6}$
Subspace+GAN	0.1, $t \in [1, 10]$ (1st to 10th TE) 0.2, $t \in [11, 15]$ (11th to 15th TE)	0
Proposed	0.04	$2e^{-7}$




Fig. 5. Reconstruction errors at different AFs for the T_2 mapping application: The proposed method shows lower errors across all AFs in both the reconstructed contrast-weighted images (left) as well as estimated T_2 maps (right).

latents in order to fully capture fine details. For initialization, we started the alternating minimization by first updating the latents (Eq. (8)). More specifically, we used a $\hat{\mathbf{U}}^0$ from a subspace reconstruction with a mild sparsity constraint, and then further updated the latents (before moving to Subproblem II) by solving the following problem each TE using k -space data directly to ensure data consistency:

$$\hat{\mathbf{w}}_t^0 = \arg \min_{\mathbf{w}_t \in \hat{\mathbf{w}}_{t-1}^0 \oplus B_t(l_t)} \sum_{c=1}^{N_c} \left\| \mathbf{y}_{c,t} - \Omega_t[\mathbf{F}\mathbf{S}_c(\Phi_t \odot \mathbf{G}_{\hat{\theta}}(\mathbf{w}_t))] \right\|_2^2. \quad (14)$$

Our StyleGAN2 network was trained with 192×192 images and 7 layers to match the reconstruction resolution of the T_2 mapping data. The ILO algorithm started from the 5th layer.

A set of multicontrast images reconstructed from data sampled with AF=4 are shown in Fig. 3. Introducing either the sparsity constraint or the adapted GAN prior improved upon the subspace reconstruction, as expected, with the GAN prior producing a slightly lower image-domain error. The proposed method integrating both the constraints yielded the best reconstruction, as shown in the error images (right panel) and the overall relative ℓ_2 errors. The T_2 estimation results corresponding to the same data (AF=4) are compared in Fig. 4, demonstrating effectiveness of the proposed method. More quantitative comparisons across different AFs are shown in Fig. 5. The proposed method consistently achieved lower errors at different AFs.

B. High-Resolution MRSI

In this section, we demonstrated the effectiveness of the proposed method on another high-dimensional imaging problem, i.e., MRSI. The goal here is to evaluate our method's utility for SNR-enhancing reconstruction from noisy, high-resolution MRSI acquisitions, benchmarking against the state-of-the-art subspace reconstruction (SPICE).

As it is rather difficult to obtain “reference” high-SNR data for high-resolution MRSI, we performed both numerical simulations (with ground truth) and in vivo experiments for evaluation. In vivo MRSI data were acquired from healthy volunteers using a fast sequence described in Ref. [47], with TR/TE = 1100/30 ms, matrix size = $64 \times 64 \times 12$, FOV = $220 \times 220 \times 64$ mm³, spectral bandwidth (BW) of 1.67kHz and 320 echoes (spectral encodings). Our sequence




Fig. 6. Simulation results for water MRSI reconstruction: (Row 1) Ground truth images simulated from interpolated water spectroscopic data ($128 \times 128 \times 12 \times 150$ spatio-spectral encodings); (Row 2) Noisy images generated by adding Gaussian noise to ground truth; (Rows 3 and 4) Reconstruction results from the SPICE subspace reconstruction and proposed method (Proposed). The proposed method produced better contrast, a higher SNR and a lower reconstruction error.

also generated a set of interleaved high-resolution, high-SNR water spectroscopic data, which is an excellent data for simulation-based evaluation, particularly for the adaptive GAN representation for different contrasts. For simulations, we first interpolated the water spectroscopic images to a grid of $128 \times 128 \times 12 \times 150$ (where 150 is the number of samples along the FID dimension) and used them as the ground truth to simulate noisy water MRSI data (Fig. 6). The noisy metabolic MRSI data from the same sequence were used to evaluate the proposed method for metabolite spatio-spectral reconstruction. In MRSI experiments, a high-resolution, anatomical scan is typically acquired and can be used as a reference for GAN adaptation. Here, a 3D T_1 w image (from an MPRAGE scan) was acquired.

The StyleGAN2 network was pretrained using the same HCP data but resized to 128×128 (as MRSI data have a lower resolution than T_2 mapping) and skull stripped. The pretrained network was first adapted to the T_1 w reference image to generate $G_{\hat{\theta}}(\cdot)$ for the proposed reconstruction. The phase network $G_{\hat{\theta}_\phi}(\cdot)$ pretrained on NYU fastMRI data was introduced to account for the complex-valued data. Considering that a) the MRSI data have a significantly larger number of “time points” than parameter mapping and b) the images at individual time points (especially later echoes) are much noisier, instead of representing images at different echoes using the adapted StyleGAN2, we proposed to use it to represent the spatial coefficient maps (for subspace reconstruction) as they can be treated as “images” with different contrasts. The

specific formulation can be written as:

$$\begin{aligned} \hat{\mathbf{U}}, \{\hat{\mathbf{w}}_r\}, \{\hat{\mathbf{w}}_{\phi,r}\} = \arg \min_{\mathbf{U}, \{\mathbf{w}_r\}, \{\mathbf{w}_{\phi,r}\}} & \left\| \mathbf{y} - \Omega[\mathbf{FB} \odot (\mathbf{U}\hat{\mathbf{V}})] \right\|_2^2 \\ & + \sum_{r=1}^R \lambda_{1,r} \left\| \mathbf{U}_r - G_{\hat{\theta}}(\mathbf{w}_r) \odot G_{\hat{\theta}_\phi}(\mathbf{w}_{\phi,r}) \right\|_F^2 + \lambda_2 \|\mathbf{D}_w(\mathbf{U}\hat{\mathbf{V}})\|_F^2, \end{aligned} \quad (15)$$

where $\mathbf{U}_r$ is the r_{th} column in $\mathbf{U}$ (spatial map of each coefficient) and R is the number of columns (aka the rank), $\mathbf{B}$ models the field inhomogeneity phase effects and $\mathbf{D}_w$ denotes an edge-weighted finite difference operator. This led to a slightly different Subproblem I than the previous application example, where the latents for magnitude network ($\{\mathbf{w}_r\}$) and phase network ($\{\mathbf{w}_{\phi,r}\}$) were determined alternatively. We performed a SPICE-based subspace reconstruction with a mild edge-preserving regularization as the initialization for $\hat{\mathbf{U}}^0$. The high-resolution portion of the latents (inputs for the last layer in StyleGAN2) were kept unchanged (after adaptation in Eq. (7)) during the GAN inversion subproblem to alleviate overfitting to noise considering the low SNR. The effects of changing latents at different resolution are further illustrated in the Discussion. We chose $\lambda_{1,r} = 1.6$ and $\lambda_2 = 0.3$ for proposed method ($\lambda_2 = 0.6$ for the SPICE reconstruction with a similar data consistency level). The ILO started from the 4th layer of the network.

Reconstruction from the simulated noisy water images supported the idea of using adaptive GAN to represent the spatial coefficient maps. The subspace for water reconstruction was obtained by applying an SVD to the “ground truth” water spectroscopic images. As shown in Fig. 6, the proposed method achieved a significantly lower reconstruction error and noticeably better “contrast recovery” than the SPICE reconstruction. Impressive SNR enhancement was achieved by both. Fig. 7 shows a set of metabolite spatio-spectral reconstruction from an in vivo MRSI acquisition. The metabolite maps from the proposed reconstruction exhibit a higher SNR and clearer contrast compared to the noisy data and subspace reconstruction (SPICE). Note that increasing the spatial regularization parameter for the SPICE reconstruction can achieve a similar level of SNR enhancement but at the expense of oversmoothing. The regularization parameters were tuned to a similar data consistency level, matching noise energy for both methods. Therefore, the proposed method provided a better trade-off in SNR and resolution by incorporating an adaptive GAN-based spatial constraint. The subspace used for the MRSI reconstruction was pre-estimated and adapted to this specific data (using the higher-SNR portion) as described in [48]–[50].

V. DISCUSSION

We presented a new adaptive generative model based spatial constraint for subspace reconstruction of high-dimensional imaging data. One unique advantage of our proposed method is that the subject-specific adaptation step allows our GAN model to be flexibly adapted to different imaging contexts, even without high-quality, application-specific training data. We validated that the subject-specific, adapted GAN (to an




Fig. 7. In vivo MRSI reconstruction: (Row 1) T_1 -weighted anatomical images (T_1w , serving as the reference for network adaptation); (Rows 2-4) Metabolite maps for the corresponding slices from noisy data (left three columns), SPICE (middle three columns) and the proposed method (right three columns), respectively. Both SPICE and the proposed method produced significant SNR enhancement for this high-resolution acquisition, revealing tissue specific metabolite variations, but better contrast and less artifacts can be observed in the proposed reconstruction (regions indicated by green arrows). The last two rows show two voxel spectra from respective methods (voxel locations marked in the T_1w images). Similar spectral quality was achieved by the SPICE and proposed methods (not surprising as both methods used the same subspace).

application-specific reference image) can be accurate representation of images from the same subject with different contrasts (a common feature in high-dimensional imaging) by updating latent variables alone. This circumvented the overfitting issue in methods updating both latents and network parameters during reconstruction, such as IAGAN. We demonstrated the effectiveness of our method in two application examples.

There are several points worth discussing. The first is the selection of regularization parameter(s) (and other hyperparameter(s)), particularly when our adaptive GAN prior is combined with other regularization functional. To alleviate possible overregularization/oversmoothing introduced by other handcrafted constraints, we chose a small λ_2 in this work for a mild additional regularization effect, and our proposed GAN constraint played a more important role. For $\{\lambda_{1,t}\}$, selection was based on a combination of discrepancy principle and visual inspection. We also observed that in our T_2 mapping experiments, similar $\{\lambda_{1,t}\}$ worked well across different AFs, and a relatively large range of $\{\lambda_{1,t}\}$ can result in similar

reconstruction results. One advantage is that $\{\lambda_{1,t}\}$ can be “time” specific. More specifically, when using adaptive GAN only to constrain the subspace reconstruction, a $2\times$ larger regularization parameter for the last five TEs (due to relatively lower SNRs) can slightly improve the T_2 mapping reconstruction. Different combinations of $\{\lambda_{1,t}\}$ may be further explored. Similarly, the determination of the radius l_t that imposes the l_1 ball constraint (Eq. (12)) can be “time” specific and further explored. The second is the initialization step. We observed that the reconstruction can be substantially affected by initialization which should be carefully chosen for different applications. For our T_2 mapping experiments, starting the iterations using $\hat{U}^0$ from a subspace reconstruction (with a mild joint sparsity constraint) may introduce potential bias. Therefore, we proposed to further update the latents using the data consistency loss described in Eq. (14) which led to a better performance. For MRSI, we found that initializing the algorithm with $\hat{U}^0$ from a initial subspace reconstruction worked well.




Fig. 8. The effects of iteration for our proposed algorithm: Reconstruction errors of the multi-TE images for a T_2 mapping data set w.r.t. the number of iteration (AF=4). The error changes substantially for the first 2-3 iterations, but minimally after 5, with and without integrating the adaptive GAN and joint sparsity constraints.

Another issue is the reconstruction speed. Currently, we used five outer iterations (Fig. 8) that achieved empirical convergence and satisfactory reconstruction results. Each iteration needs about 35-36 minutes (can be application dependent). The inference time for the 'Subspace+GAN' and proposed method is similar because Subproblem (I) which dominates the overall computing time is the same for both methods (taking about 35 minutes each iteration). Here we solved the latents TE by TE for parameter mapping, for which more efficient implementations should be considered. Other methods that better exploit parallel computation to simultaneously update the latents for all time points and/or methods that circumvent explicit network inversion will be investigated.

While our work and previous work [37] using GAN for image reconstruction have demonstrated that imposing constraints on latents can help better utilize prior information from reference images and reduce overfitting to artifacts, an important issue that presents unique research opportunities is the properties of latents (feature) space. Here, we further investigated the latent space structure through a style-mixing experiment on the adapted StyleGAN2 (after pre-training and network parameter adaptation to experimentally acquired T_1 w image). Fig. 9 demonstrated that manipulating low-resolution latents can introduce more global and significant changes in structure and contrast of the generated images than changing the high-resolution latents. This is consistent to us keeping the high-resolution latents unchanged for our MRSI reconstruction (lower-resolution data and avoiding overfitting to noise). In the meantime, it can be observed that for the current StyleGAN representation, different types of image features (e.g. contrast and geometry) are still somewhat entangled, and it is challenging to control certain semantic features by changing latents at a specific scale. Future work can focus on developing disentangled representation to enable more explicit control of different image features, which will be desirable for high-dimensional MR imaging problems, since in many situations, only parts of image features are changing across




Fig. 9. A style mixing experiment for our adapted StyleGAN2: A single image (Source A) and a set of images (Source B) were generated from the network; The rest of the images were generated by replacing a specific subset of latents in Source A by corresponding latents from Source B. Specifically, for each row, images in different columns were obtained by taking the same level of latents from different images in Source B. From the top to bottom rows, different levels of latents from Source B were used (for the same image A), from low-resolution to high-resolution subsets, respectively. As can be seen, lower-resolution latents from Source B are more related to global features such as shape and contrast, while higher-resolution latents lead to more subtle changes in fine details and less/negligible contrast and geometry variations.

specific dimensions.

Recent development of generative models may also be helpful to improve the accuracy and efficiency of image representation, and thus improve the reconstruction performance. Transformer-based GAN has been proposed [51] to better utilize the global correspondence of image features through self-attention mechanism. Korkmaz et al. [52] demonstrated that generative vision transformers were more robust to artifacts and achieved better performance than CNN-based generative models when worked as a deep image prior (without pretraining) [34]. Diffusion model is an emerging approach to generate high-quality images. Diffusion model training is more stable than GAN and can obtain better sample quality over the state-of-the-art GANs [53], [54]. While diffusion models have demonstrated its potential for image reconstruction [55], [56], many opportunities remain on how to incorporate diffusion models into high-dimensional MR imaging problems [57].

VI. CONCLUSION

We proposed a novel reconstruction method for high-dimensional MR imaging that integrated subspace modeling and an adaptive generative-network-based image prior. The proposed adaptive generative model served as an accurate representation of images from the same subject with different contrasts and an effective spatial constraint for subspace reconstruction. The complementary powers of subspace, generative image constraint, and sparsity regularization produced improved performance over state-of-the-art subspace methods in two application examples. We believe our work offers a

new perspective on high-dimensional image reconstruction via integrating learning-based spatial priors and low-dimensional modeling.

REFERENCES

- [1] D. Ma *et al.*, “Magnetic resonance fingerprinting,” *Nature*, vol. 495, pp. 187–192, 2013.
- [2] J. P. Earls *et al.*, “Cardiac MRI: recent progress and continued challenges,” *J. Magn. Reson. Imag.*, vol. 16, pp. 111–127, 2002.
- [3] S. Posse *et al.*, “MR spectroscopic imaging: principles and recent advances,” *J. Magn. Reson. Imag.*, vol. 37, pp. 1301–1325, 2013.
- [4] W. Bogner *et al.*, “Accelerated MR spectroscopic imaging—a review of current and emerging techniques,” *NMR Biomed.*, vol. 34, p. e4314, 2021.
- [5] Z.-P. Liang, “Spatiotemporal imaging with partially separable functions,” in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2007, pp. 988–991.
- [6] J. P. Haldar and Z.-P. Liang, “Spatiotemporal imaging with partially separable functions: A matrix recovery approach,” in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2010, pp. 716–719.
- [7] B. Zhao *et al.*, “Image reconstruction from highly undersampled (k,t)-space data with joint partial separability and sparsity constraints,” *IEEE Trans. Med. Imag.*, vol. 31, pp. 1809–1820, 2012.
- [8] A. S. Gupta and Z. Liang, “Dynamic imaging by temporal modeling with principal component analysis,” in *Proc. Int. Soc. Magn. Reson. Med.*, 2001, pp. 21–27.
- [9] M. Fu *et al.*, “High-resolution dynamic speech imaging with joint low-rank and sparsity constraints,” *Magn. Reson. Med.*, vol. 73, pp. 1820–1832, 2015.
- [10] S. G. Lingala *et al.*, “Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR,” *IEEE Trans. Med. Imag.*, vol. 30, pp. 1042–1054, 2011.
- [11] F. H. Petzschner *et al.*, “Fast MR parameter mapping using k-t principal component analysis,” *Magn. Reson. Med.*, vol. 66, pp. 706–716, 2011.
- [12] B. Zhao *et al.*, “Accelerated MR parameter mapping with low-rank and sparsity constraints,” *Magn. Reson. Med.*, vol. 74, pp. 489–498, 2015.
- [13] T. Zhang *et al.*, “Accelerating parameter mapping with a locally low rank constraint,” *Magn. Reson. Med.*, vol. 73, pp. 655–661, 2015.
- [14] F. Lam and Z.-P. Liang, “A subspace approach to high-resolution spectroscopic imaging,” *Magn. Reson. Med.*, vol. 71, pp. 1349–1357, 2014.
- [15] J. D. Trzasko, “Exploiting local low-rank structure in higher-dimensional MRI applications,” in *Wavelets and Sparsity XV*, D. V. D. Ville *et al.*, Eds., vol. 8858, 2013, p. 885821.
- [16] X. Qu *et al.*, “Accelerated NMR spectroscopy with low-rank reconstruction,” *Angew. Chem., Int. Ed.*, vol. 54, pp. 852–854, 2015.
- [17] T. Zhang *et al.*, “Accelerating parameter mapping with a locally low rank constraint,” *Magn. Reson. Med.*, vol. 73, pp. 655–661, 2015.
- [18] J. P. Haldar and D. Hernando, “Rank-constrained solutions to linear matrix equations using power factorization,” *IEEE Signal Process. Lett.*, vol. 16, pp. 584–587, 2009.
- [19] A. G. Christodoulou *et al.*, “Improved subspace estimation for low-rank model-based accelerated cardiac imaging,” *IEEE Trans. Biomed. Eng.*, vol. 61, pp. 2451–2457, 2014.
- [20] A. G. Christodoulou *et al.*, “Magnetic resonance multitasking for motion-resolved quantitative cardiovascular imaging,” *Nature Biomed. Eng.*, vol. 2, pp. 215–226, 2018.
- [21] F. Lam *et al.*, “High-resolution ^1H -MRSI of the brain using SPICE: data acquisition and image reconstruction,” *Magn. Reson. Med.*, vol. 76, pp. 1059–1070, 2016.
- [22] X. Peng *et al.*, “Accelerated exponential parameterization of T_2 relaxation with model-driven low rank and sparsity priors (MORASA),” *Magn. Reson. Med.*, vol. 76, pp. 1865–1878, 2016.
- [23] A. G. Christodoulou *et al.*, “High-resolution cardiovascular MRI by integrating parallel imaging with low-rank and sparse modeling,” *IEEE Trans. Biomed. Eng.*, vol. 60, pp. 3083–3092, 2013.
- [24] G. Wang and J. Qi, “PET image reconstruction using kernel method,” *IEEE Trans. Med. Imag.*, vol. 34, pp. 61–71, 2014.
- [25] Y. Li *et al.*, “Constrained MRSI reconstruction using water side information with a kernel-based method,” in *Proc. Int. Soc. Magn. Reson. Med.*, 2018, p. 0166.
- [26] A. Creswell *et al.*, “Generative adversarial networks: An overview,” *IEEE Signal Process. Mag.*, vol. 35, pp. 53–65, 2018.
- [27] M. Tschannen *et al.*, “Recent advances in autoencoder-based representation learning,” *arXiv preprint arXiv:1812.05069*, 2018.
- [28] A. Aggarwal *et al.*, “Generative adversarial network: An overview of theory and applications,” *Int. J. Inf. Manage.*, vol. 1, p. 100004, 2021.
- [29] A. Bora *et al.*, “Compressed sensing using generative models,” in *Proc. ICML*, 2017, pp. 537–546.
- [30] S. A. Hussein *et al.*, “Image-adaptive GAN based reconstruction,” in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 04, 2020, pp. 3121–3129.
- [31] M. Asim *et al.*, “Invertible generative models for inverse problems: mitigating representation error and dataset bias,” in *Proc. ICML*, 2020, pp. 399–409.
- [32] V. A. Kelkar *et al.*, “Compressible latent-space invertible networks for generative model-constrained image reconstruction,” *IEEE Trans. Comput. Imag.*, vol. 7, pp. 209–223, 2021.
- [33] S. Bhadra *et al.*, “Medical image reconstruction with image-adaptive priors learned by use of generative adversarial networks,” in *Med. Imag.: Phys. Med. Imag.*, vol. 11312. SPIE, 2020, pp. 206–213.
- [34] D. Ulyanov *et al.*, “Deep image prior,” in *Proc. IEEE CVPR*, 2018, pp. 9446–9454.
- [35] J. Yoo *et al.*, “Time-dependent deep image prior for dynamic MRI,” *IEEE Trans. Med. Imag.*, vol. 40, pp. 3337–3348, 2021.
- [36] T. Karras *et al.*, “Analyzing and improving the image quality of stylegan,” in *Proc. IEEE CVPR*, 2020, pp. 8110–8119.
- [37] V. A. Kelkar and M. Anastasio, “Prior image-constrained reconstruction using style-based generative models,” in *Proc. ICML*, 2021, pp. 5367–5377.
- [38] J. Yao *et al.*, “An efficient algorithm for dynamic MRI using low-rank and total variation regularizations,” *Med. Image Anal.*, vol. 44, pp. 14–27, 2018.
- [39] F. Shi *et al.*, “LRTV: MR image super-resolution with low-rank and total variation regularizations,” *IEEE Trans. Med. Imag.*, vol. 34, pp. 2459–2466, 2015.
- [40] M. Jacob *et al.*, “Structured low-rank algorithms: Theory, magnetic resonance applications, and links to machine learning,” *IEEE Signal Process. Mag.*, vol. 37, pp. 54–68, 2020.
- [41] G. Daras *et al.*, “Intermediate layer optimization for inverse problems using deep generative models,” in *Proc. ICML*, 2021, pp. 2421–2432.
- [42] D. C. Van Essen *et al.*, “The WU-Minn human connectome project: an overview,” *Neuroimage*, vol. 80, pp. 62–79, 2013.
- [43] F. Knoll *et al.*, “fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning,” *Radiol. Artif. Intell.*, vol. 2, p. e190007, 2020.
- [44] D. L. Parker *et al.*, “Phase reconstruction from multiple coil data using a virtual reference coil,” *Magn. Reson. Med.*, vol. 72, pp. 563–569, 2014.
- [45] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [46] M. Uecker *et al.*, “ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA,” *Magn. Reson. Med.*, vol. 71, pp. 990–1001, 2014.
- [47] Z. Wang *et al.*, “High-resolution, 3D multi-TE ^1H MRSI using fast spatio-spectral encoding and subspace imaging,” *Magn. Reson. Med.*, vol. 87, pp. 1103–1118, 2022.
- [48] F. Lam *et al.*, “Ultrafast magnetic resonance spectroscopic imaging using SPICE with learned subspaces,” *Magn. Reson. Med.*, vol. 83, pp. 377–390, 2020.
- [49] R. Guo *et al.*, “Simultaneous QSM and metabolic imaging of the brain using SPICE: Further improvements in data acquisition and processing,” *Magn. Reson. Med.*, vol. 85, pp. 970–977, 2021.
- [50] X. Peng *et al.*, “Simultaneous QSM and metabolic imaging of the brain using SPICE,” *Magn. Reson. Med.*, vol. 79, pp. 13–21, 2018.
- [51] Y. Jiang *et al.*, “TransGAN: Two pure transformers can make one strong GAN, and that can scale up,” *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 14745–14758, 2021.
- [52] Y. Korkmaz *et al.*, “Deep MRI reconstruction with generative vision transformers,” in *Proc. MLMIR workshop*, 2021, pp. 54–64.
- [53] P. Dhariwal and A. Nichol, “Diffusion models beat GANs on image synthesis,” *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 8780–8794, 2021.
- [54] J. Ho *et al.*, “Cascaded diffusion models for high fidelity image generation,” *J. Mach. Learn. Res.*, vol. 23, pp. 1–33, 2022.
- [55] S. U. Dar *et al.*, “Adaptive diffusion priors for accelerated MRI reconstruction,” *arXiv preprint arXiv:2207.05876*, 2022.
- [56] Y. Song *et al.*, “Solving inverse problems in medical imaging with score-based generative models,” in *Proc. ICLR*, 2022.
- [57] B. Levac *et al.*, “Conditional score-based reconstructions for multi-contrast mri,” *arXiv preprint arXiv:2303.14795*, 2023.