

Bi-Linear Modeling of Data Manifolds for Dynamic-MRI Recovery

Gaurav N. Shetty,^{*} *Student Member, IEEE*, Konstantinos Slavakis,^{*} *Senior Member, IEEE*,
 Abhishek Bose,[†] *Student Member, IEEE*, Ukash Nakarmi,[‡] *Member, IEEE*,
 Gesualdo Scutari,[¶] *Senior Member, IEEE*, and Leslie Ying,[§] *Senior Member, IEEE*

Abstract—This paper puts forth a novel bi-linear modeling framework for data recovery via manifold-learning and sparse-approximation arguments and considers its application to dynamic magnetic-resonance imaging (dMRI). Each temporal-domain MR image is viewed as a point that lies onto or close to a smooth manifold, and landmark points are identified to describe the point cloud concisely. To facilitate computations, a dimensionality reduction module generates low-dimensional/compressed renditions of the landmark points. Recovery of high-fidelity MRI data is realized by solving a non-convex minimization task for the linear decompression operator and affine combinations of landmark points which locally approximate the latent manifold geometry. An algorithm with guaranteed convergence to stationary solutions of the non-convex minimization task is also provided. The aforementioned framework exploits the underlying spatio-temporal patterns and geometry of the acquired data without any prior training on external data or information. Extensive numerical results on simulated as well as real cardiac-cine MRI data illustrate noteworthy improvements of the advocated machine-learning framework over state-of-the-art reconstruction techniques.

Index Terms—Dynamic MRI, manifold learning, low rank, sparsity, dimensionality reduction.

I. INTRODUCTION

Magnetic-resonance imaging (MRI), a non-invasive, non-ionizing and high-fidelity visualization technology, has found widespread applications in cardiac-cine, dynamic contrast-enhanced and neuro-imaging, playing a key role in medical research and diagnosis [1]. MRI’s inherent limitations and various physiological constraints often incur slow data acquisition,

while long scanning times make MRI an expensive process, may cause patient discomfort, and hinder thus its usefulness.

Raw MRI data are observed in the k -space domain; the image (visual) data are computed by the inverse Fourier transform of the k -space ones [1]. A prominent way to speed up data acquisition is to sample the k -space or frequency domain densely enough to ensure reconstruction of a high-fidelity and artifact-free image via Fourier-transform arguments [2]. In the case of dynamic (d)MRI, where an extra-temporal dimension is added to the spatial domain, sampling is also performed along the time axis. Due to MRI’s slow scanning times, it becomes difficult for the data acquisition process to keep up with the motion of the organs or the fluid flow in the field of view (FOV) [1]. It is a usual case to not be able to reach the necessary sampling density, not only in the temporal direction but also in the k -space domain, to guarantee artifact-free reconstructed images. This “under-sampling” inflicts signal aliasing and distortion [3].

Naturally, a lot of the MRI-research effort has been focusing on developing reconstruction algorithms that improve the spatio-temporal resolution of MR images given the highly under-sampled k -space data. To this end, artificial-intelligence (AI) approaches have been very recently placed at the focal point of MRI research; examples are convolutional neural networks (CNNs) [4]–[9], deep variational networks (VNs) [10], and generative adversarial neural networks (GANs) [11], [12]. AI methods learn non-linear mappings via extensive offline training on large-scale datasets, different from or in addition to the acquired data, and use those learned non-linear mappings to map the observed low-resolution, or, undersampled data to their high-fidelity counterparts. In contrast to the *data-driven* AI methods, the present work, as well as the following prior-art schemes, assume *no* offline training on large-scale datasets and resort *solely* to the observed data.

A popular approach to exploit the underlying spatio-temporal patterns within dMRI data is compressed sensing (CS) [13]–[16]. Low-rank structures [17], [18] and total-variation-based schemes [19]–[21] have also been explored at length and found to produce promising results for slow varying dynamic data. For instance, [18] proposes the estimation, first, of a temporal basis of image time series via singular-value decomposition, prior to formulating a sparsity inducing convex-recovery task. Nevertheless, these schemes seem to be less

This work was supported in part by the NSF grants 1514056 and 1718796.

^{*}G. N. Shetty (contact author) and K. Slavakis are with the Dept. of Electrical Engineering, University at Buffalo, Buffalo, NY 14260 USA (e-mail: {gauravna,kslavaki}@buffalo.edu).

[†]A. Bose is with INRIA, Université Côte d’Azur, 2004 Route des Lucioles, B.P. 93, 06902, Sophia Antipolis CEDEX 06902, France (e-mail: abhishek.bose@inria.fr).

[‡]U. Nakarmi is with the Depts. of Electrical Engineering and Radiology, Stanford University, Stanford, CA 94305 USA (e-mail: nakarmi@stanford.edu).

[¶]G. Scutari is with the Dept. of Industrial Engineering, Purdue University, West Lafayette, IN 47907-2023 USA (e-mail: gscutari@purdue.edu).

[§]L. Ying is with the Depts. of Electrical Engineering and Biomedical Engineering, University at Buffalo, Buffalo, NY 14260 USA (e-mail: leiyang@buffalo.edu).

Copyright (c) 2019 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

effective when it comes to dMRI with extensive inter-frame motion or with a low number of temporal frames, as mentioned in [22], [23]. Acceleration capabilities of CS combined with the motion-robustness of radial imaging techniques have also found their place in MRI reconstruction and have shown promising results in cardiac and respiratory motion correction. The motion resolved strategy [24] introduces extra motion dimensions by sorting the continuously acquired k-space data into distinct motion states followed by a sparsity and total variation based CS technique for reconstruction. There has also been a growing interest in MRI-recovery by dictionary-learning (DL) schemes [25]–[29]. In [23], for example, the dMRI data are decomposed in two components: A low-rank one, that captures the temporal (video) background, and a sparse one, described via spatio-temporal patch-based DL that models the (dynamic) foreground.

Manifold-learning techniques have also been employed to recover dMRI data from highly under-sampled observations [22], [30]–[35]. In [22], a graph-Laplacian matrix is formed via the Euclidean distances between points of a data cloud and is subsequently used as a regularizer in a convex-recovery task. Improving on [22], a bandlimited modeling of the data points, coupled with an improved estimation of the Laplacian matrix, was proposed in [31] to make the reconstruction resilient to noise and patient motion, computationally inexpensive and less demanding memory-wise. The work in [32] proposes a kernel variation to the graph Laplacian estimation in [22]. A popular path followed by manifold-learning schemes is to perform dimensionality reduction of the collected high-dimensional data prior to applying a reconstruction algorithm, *e.g.*, [33], [34]. In the non-MRI context, [36] capitalizes, also, on a Euclidean-distance-based Laplacian matrix to perform dimensionality reduction prior to reconstructing data by local principal component analysis. In the previous schemes, *all* of the observed points participate in the dimensionality-reduction task, raising thus computational burdens, especially in cases where the number of data is excessively large. Methods that perform dimensionality reduction on properly chosen small-cardinality subsets of the observed data cloud have been introduced for clustering and classification, but *not* for regression tasks [37]–[39].

This paper follows the manifold-learning path and serves a two-fold objective: **i)** Present a machine-learning framework, the bi-linear modeling of data manifolds (BiLMDM), that contributes novelties to exploiting local and latent data structures via a sparsity-aware and bi-linear optimization task; and **ii)** apply BiLMDM to the dMRI-data recovery problem. In a nutshell, BiLMDM can be described as follows. Each vector of data, observed from an undersampled dMRI temporal frame, is modeled as a point onto or close to an unknown manifold, embedded in a Euclidean space. The only assumption imposed on the manifold is smoothness [40]. Landmark points are chosen to concisely describe the observed data-vector cloud, and are mapped to a lower-dimensional space, as in [34], [41], to effect compression and enable low-computational footprints in the proposed algorithmic solutions. Motivated by the smooth-manifold hypothesis, the proposed work approximates each data vector as an affine combinations of neighboring *landmark*

points (*cf.* Fig. 2). A *locally bi-linear* factorization model, novel for data representations, is then formed to model/fit the point cloud: One factor gathers the coefficients of the previous affine combinations of the landmark points, while the other one serves as the linear *decompression* operator that unfolds the points back to the image dimensions. Improving on our preliminary results [42], a highly modular bi-linear optimization task is tailored to the dMRI-data recovery problem, penalized by terms which account for sparsity along the temporal axis and other modeling assumptions. A successive-convex-approximation algorithm is proposed to guarantee convergence to a stationary solution of the previous bi-linear optimization task.

The proposed work is validated against the following state-of-the-art schemes: Partially separable sparsity aware model (PS-Sparse) [18], joint manifold learning and sparsity aware (MLS) framework [34], smoothness regularization on manifolds (SToRM) [22], low rank and adaptive sparse signal model (LASSI) [23] and extra dimensional golden angle radial sparse parallel (XD-GRASP) [24]. In spite of sharing with the previous state-of-the-art schemes the principle of identifying low-dimensional structures for high-dimensional data and exploiting this structure to reconstruct data, BiLMDM's original contributions in data approximations can be summarized as follows: **i)** BiLMDM departs from the mainstream manifold-learning approach of identifying a data-graph Laplacian matrix that penalizes an optimization (data-recovery) task as in [22], [31], [35], and uses instead the fundamental geometric principle of tangent spaces of smooth manifolds to search for data dependencies/patterns through local affine combinations of landmark points (*cf.* Fig. 2); **ii)** it offers a novel bi-linear factorization model, where one linear factor captures the local smoothness of the manifolds, via local affine combinations, while the other one accounts for the linear decompression operator which maps the low-dimensional affine approximations back to the high-dimensional raw-data/input space. The affine combinations and the decompression operator are jointly identified via a highly modular non-convex data recovery task, which may accommodate also additional prior information, *e.g.*, periodicity over the temporal axis of the dMRI data. Unlike the traditional bi-linear schemes, BiLMDM doesn't follow the classical dictionary learning approach of decomposing the image into two components as discussed earlier and model the dynamic foreground of the image as a sparse combination of adaptive dictionary atoms (which act as image basis) of MR image signals. A more detailed description of the differences of BiLMDM with the state-of-the-art dMRI recovery schemes is deferred to Sec. III. The efficacy of BiLMDM was observed on both synthetically generated and experimentally acquired cardiac cine MR data, undersampled via both Cartesian and radial trajectories. BiLMDM has consistently outperformed all the competing techniques and produced reconstructions which were sharp, free of undesirable artifacts, deformations and temporal bleeding.

The rest of the paper is organized as follows: Section II-A describes, in short, the dMRI acquisition scheme in the k-space domain. Section II-B details the BiLMDM's modeling assumptions, while Sec. III describes the algorithm to solve the proposed non-convex minimization task. The extensive




Fig. 1. (a) k-space with 1-D Cartesian sampling pattern; (b) k-space with radial sampling pattern; (c) The $N_p \times N_f \times N_{fr}$ (k,t)-space. “Navigator (pilot) data” comprise the gray-colored $\nu \times N_f \times N_{fr}$ area of the (k,t)-space ($\nu \ll N_p$).

numerical tests of Section IV showcase that BiLMDM outperforms state-of-the-art dMRI recovery schemes. The manuscript summarises the numerical results in Section V and concludes in Section VI. Finally, the appendix gathers basic mathematical facts and expressions that are essential for the implementation of the algorithm described in Sec. III.

II. BI-LINEAR MODELING OF DATA MANIFOLDS

A. dMRI data description

MRI data $\mathbf{Y} \in \mathbb{C}^{N_p \times N_f}$ ($\mathbb{C}$ denotes the set of all complex-valued numbers) are observed in *k-space* (frequency domain), which spans an area of size $N_p \times N_f$ (cf. Figs. 1a and 1b), with N_p standing for the number of phase-encoding lines and N_f for the number of frequency-encoding ones [1]. Data $\mathbf{Y}$ can be considered as the two-dimensional (discrete) Fourier transform $\mathcal{F}(\cdot)$ of the image-domain data $\mathbf{X} \in \mathbb{C}^{N_p \times N_f}$, i.e., $\mathbf{Y} = \mathcal{F}(\mathbf{X})$ [1]. Without any loss of generality, this study assumes that the “low-frequency” part of $\mathbf{Y}$ is located around the center of the $N_p \times N_f$ area. Availability of the data over the whole k-space is infeasible in practice; k-space is usually severely under-sampled [3]. There exist several strategies to sample the k-space; examples are the 1-D Cartesian (Fig. 1a) and the radial (Fig. 1b) ones, where the “white” lines in Figs. 1a and 1b denote the available/sampled data, while data in the “black” areas are not observed. A general trend among sampling strategies is to put more emphasis on low-frequency components, which carry contrast information and with high SNR, and select few high-frequency components, which comprise high-resolution image details. The 1-D Cartesian sampling pattern emulates the acquisition of k-space pixels via the 1-D Gaussian distribution, acquiring a large number of samples in the central k-space area while sampling few ones from the “high-frequency” area (cf. Fig. 1a). The radial-sampling pattern consists of radial spokes which yield dense sampling at the center of k-space, while the sampling density is decreased as the spokes move away from the center (cf. Fig. 1b).

In dMRI, an additional dimension is added to the MRI k-space to accommodate time (the axis vertical on the paper in Fig. 1c), resulting in the augmented (k,t)-space. The dMRI (k,t)-space can be viewed, in other words, as the N_{fr} -fold Cartesian product of the $(N_p \times N_f)$ -sized MRI k-space, where N_{fr} represents the number of observed MRI frames over time. In dMRI, k-space $\mathbf{Y}_j$ and image-domain $\mathbf{X}_j$ data are connected via $\mathbf{Y}_j = \mathcal{F}(\mathbf{X}_j)$, $j \in \{1, \dots, N_{fr}\}$. The (k,t) space is usually highly under-sampled. To extract reliable information from the (k,t)-space data, this work follows [18], [22], [34] and considers a small number ν ($\ll N_p$) of phase-encoding lines, coined




Fig. 2. Landmark points $\{\ell_{k_i}\}_{i=1}^3$ are affinely combined to describe $\mathbf{y}_j^{nav}$. All affine combinations of $\{\ell_{k_i}\}_{i=1}^3$ are depicted by the gray-colored plane.

“navigator (pilot) data” (the gray-colored area in Fig. 1c), to learn the intrinsic low-dimensional structure of the data.

To facilitate processing, the (k,t)-space data are vectorized. More specifically, $\text{vec}(\mathbf{Y}_j)$ stacks one column of $\mathbf{Y}_j$ below the other to yield the complex-valued $N_k \times 1$ vector $\mathbf{y}_j := \text{vec}(\mathbf{Y}_j)$ where $N_k := N_p N_f$ denotes the number of pixels in every frame. To avoid notation clutter, $\mathcal{F}$ still denotes the two-dimensional (discrete) Fourier transform even when applied to vectorized versions of image frames: $\mathcal{F}[\text{vec}(\mathbf{X}_j)] := \text{vec}[\mathcal{F}(\mathbf{X}_j)] = \text{vec}(\mathbf{Y}_j)$. All vectorized k-space frames are gathered in the $N_k \times N_{fr}$ matrix $\mathbf{Y} := [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{N_{fr}}]$ so that the vectorized image-domain data are $\mathbf{X} := \mathcal{F}^{-1}(\mathbf{Y}) := [\mathcal{F}^{-1}(\mathbf{y}_1), \mathcal{F}^{-1}(\mathbf{y}_2), \dots, \mathcal{F}^{-1}(\mathbf{y}_{N_{fr}})]$, where $\mathcal{F}^{-1}(\cdot)$ denotes the inverse two-dimensional (discrete) Fourier transform. The navigator data of the j th k-space frame (cf. Fig. 1c), $j \in \{1, 2, \dots, N_{fr}\}$, are gathered into a $\nu N_f \times 1$ vector $\mathbf{y}_j^{nav}$. All navigator data comprise the $\nu N_f \times N_{fr}$ matrix $\mathbf{Y}_{nav} := [\mathbf{y}_1^{nav}, \mathbf{y}_2^{nav}, \dots, \mathbf{y}_{N_{fr}}^{nav}]$.

B. Modeling assumptions

The high-dimensional navigator data $\{\mathbf{y}_j^{nav}\}_{j=1}^{N_{fr}}$ carry useful information about spatio-temporal dependencies in the (k,t)-space. To promote parsimonious data representations, especially in cases where N_{fr} attains large values, it is desirable to extract a subset $\{\ell_k\}_{k=1}^{N_\ell} \subset \{\mathbf{y}_j^{nav}\}_{j=1}^{N_{fr}}$ ($N_\ell \leq N_{fr}$), called *landmark points*, which provide a “concise description,” in a user-defined sense, of the data cloud $\{\mathbf{y}_j^{nav}\}_{j=1}^{N_{fr}}$. To this end, the following assumption, often met in manifold-learning approaches [43], imposes structure on $\{\mathbf{y}_j^{nav}\}_{j=1}^{N_{fr}}$.

Assumption 1. Data $\{\mathbf{y}_j^{nav}\}_{j=1}^{N_{fr}}$ lie on a smooth low-dimensional manifold $\mathcal{M}$ [40] embedded in the high-dimensional Euclidean space $\mathbb{C}^{\nu N_f}$ (cf. Fig. 2).

For example, the most well-known case of $\mathcal{M}$ is a linear subspace, which is the main hypothesis behind principal component analysis (PCA). PCA assumes that the data lie close to or onto a linear subspace in a low-dimensional space. Along the same lines, based on As. 1 and the concept of the tangent space of a smooth manifold, it is conceivable that neighbouring landmark points cooperate affinely to describe vector $\mathbf{y}_j^{nav}$ (the gray-colored area in Fig. 2 depicts *all* possible affine combinations of $\{\ell_{k_1}, \ell_{k_2}, \ell_{k_3}\}$). In other words, every temporal frame is approximated by a combination of landmark points, where these landmark points can be viewed as a representative of a particular group of frames sharing similar phase and motion characteristics. Upon defining the $\nu N_f \times N_\ell$ matrix $\mathbf{\Lambda} := [\ell_1, \ell_2, \dots, \ell_{N_\ell}]$, the above notion can be established by the existence of an $N_\ell \times 1$ vector $\mathbf{b}_j$ that renders the approximation error $\|\mathbf{y}_j^{nav} - \mathbf{\Lambda} \mathbf{b}_j\|$ small, where $\|\cdot\|$ denotes the standard Euclidean norm of space $\mathbb{C}^{\nu N_f}$.

Since affine combinations are desirable, $\mathbf{b}_j$ is constrained to satisfy $\mathbf{1}_{N_\ell}^\top \mathbf{b}_j = 1$, where $\mathbf{1}_{N_\ell}$ stands for the all-one $N_\ell \times 1$ vector and superscript $\top$ denotes vector/matrix transposition. Moreover, motivated by the low-dimensional nature of $\mathcal{M}$ (As. 1), only a few landmark points are considered to cooperate to represent $\mathbf{y}_j^{\text{nav}}$ in its closed vicinity, *i.e.*, $\mathbf{b}_j$ is sparse. The previous arguments are summarized into the following modeling hypothesis.

Assumption 2. There exist a sparse $N_\ell \times N_{\text{fr}}$ matrix $\mathbf{B}$, with $\mathbf{1}_{N_\ell}^\top \mathbf{B} = \mathbf{1}_{N_{\text{fr}}}^\top$, and a $\nu N_f \times N_{\text{fr}}$ matrix $\mathbf{E}_2$, which gathers approximation errors, such that (s.t.) $\mathbf{Y}_{\text{nav}} = \mathbf{A}\mathbf{B} + \mathbf{E}_2$.

The previous assumption holds true in the prototypical case where $\mathcal{M}$ is a linear subspace and $\mathbf{Y}_{\text{nav}}$ comprises column vectors which lie close to or onto $\mathcal{M}$. Any matrix which includes as column vectors any basis of $\mathcal{M}$ may serve as $\mathbf{A}$. Since the affine hull [44] of the columns of $\mathbf{A}$ coincides with $\mathcal{M}$, there exists surely a coefficient matrix $\mathbf{B}$ that satisfies the affine constraints $\mathbf{1}_{N_\ell}^\top \mathbf{B} = \mathbf{1}_{N_{\text{fr}}}^\top$ and $\mathbf{Y}_{\text{nav}} = \mathbf{A}\mathbf{B}$. It is not now difficult to see that As. 2 covers also the case where $\mathcal{M}$ is a union $\cup_{q=1}^Q \mathcal{M}_q$ of linear subspaces $\{\mathcal{M}_q\}_{q=1}^Q$, with $\mathbf{Y}_{\text{nav}}$ comprising columns vectors which lie close to or onto $\{\mathcal{M}_q\}_{q=1}^Q$. Hence, As. 2 lays the foundations for more general cases where data vectors lie close to or onto a union of smooth manifolds, with different dimensions. In such a way, As. 2 offers theoretical support and justification in cases where data comprise disparate groups, where intra-group dependencies can be captured by a single smooth manifold, while inter-group relations can be only viewed via the union of those manifolds, *e.g.*, the union of data drawn from a static gray-colored background and data describing the movement of a beating heart.

Although several strategies may be implemented to identify the landmark points $\mathbf{A}$, a greedy optimization methodology, introduced in [39], is adopted here. In short, at every step of the algorithm, a landmark point is selected from $\{\mathbf{y}_j^{\text{nav}}\}_{j=1}^{N_{\text{fr}}}$ that maximizes, over all *un-selected* $\{\mathbf{y}_j^{\text{nav}}\}_{j=1}^{N_{\text{fr}}}$, the minimum distance to the landmark points which have been already selected up to the previous step of the algorithm. The algorithm of [39] scores a computational complexity of order $\mathcal{O}(N_\ell N_{\text{fr}})$, which is naturally heavier than that of a procedure that selects $\{\ell_k\}_{k=1}^{N_\ell}$ randomly from $\{\mathbf{y}_j^{\text{nav}}\}_{j=1}^{N_{\text{fr}}}$.

Still, the landmark points (columns of $\mathbf{A}$) are high dimensional. To meet restrictions imposed by finite computational resources, it is desirable to reduce the dimensionality of $\mathbf{A}$. To this end, the methodology of [41], which is motivated by [43], [45], is employed. The approach comprises of two steps:

1) Given $\mathbf{A}$ and a user-defined $\lambda_W > 0$, solve

$$\begin{aligned} \min_{\mathbf{W} \in \mathbb{C}^{N_\ell \times N_\ell}} \|\mathbf{A} - \mathbf{A}\mathbf{W}\|_F^2 + \lambda_W \|\mathbf{W}\|_1 \\ \text{s.t. } \mathbf{1}_{N_\ell}^\top \mathbf{W} = \mathbf{1}_{N_\ell}^\top \text{ and } \text{diag}(\mathbf{W}) = \mathbf{0}, \end{aligned} \quad (1a)$$

where $\|\cdot\|_F$ stands for the Frobenius norm of a matrix. Since $\{\ell_k\}_{k=1}^{N_\ell}$ lie on the manifold $\mathcal{M}$, then according to As. 1 and Fig. 2, any point taken from $\{\ell_k\}_{k=1}^{N_\ell}$ may be faithfully approximated by an affine combination of the rest of the landmark points. In other words, there exists a matrix $\mathbf{W}$ s.t. $\mathbf{A} \approx \mathbf{A}\mathbf{W}$. With $\mathbf{1}_{N_\ell}^\top \mathbf{W} = \mathbf{1}_{N_\ell}^\top$ manifesting the previous

desire for affine combinations, the constraint $\text{diag}(\mathbf{W}) = \mathbf{0}$ is used to exclude the trivial solution of the identity matrix $\mathbf{I}_{N_\ell}$ for $\mathbf{W}$. Task (1a) is an affinely constrained composite convex minimization task, and, hence, the framework of [46] can be employed to solve it, due to the flexibility by which [46] deals with affine constraints when compared with state-of-the-art convex optimization techniques.

2) Once $\mathbf{W}$ is obtained from the previous step and for a user-defined integer number $d \leq \min\{N_\ell, \nu N_f\}$, solve

$$\min_{\tilde{\mathbf{A}} \in \mathbb{C}^{d \times N_\ell}} \|\tilde{\mathbf{A}} - \tilde{\mathbf{A}}\mathbf{W}\|_F^2 \text{ s.t. } \tilde{\mathbf{A}}\tilde{\mathbf{A}}^H = \mathbf{I}_d, \quad (1b)$$

where the constraint $\tilde{\mathbf{A}}\tilde{\mathbf{A}}^H = \mathbf{I}_d$ is used to exclude the trivial solution of $\tilde{\mathbf{A}} = \mathbf{0}$, and the superscript H denotes the Hermitian transpose of a matrix. The solution of the previous task is nothing but the complex conjugate transpose of the matrix which comprises the d minimal eigenvectors of $(\mathbf{I}_{N_\ell} - \mathbf{W})(\mathbf{I}_{N_\ell} - \mathbf{W})^H$.

To this point, the identified landmark points were used to capture a latent structure of the MR data cloud and the structure was further used to compress the landmark points to a lower dimension, thus, facilitating the search for affine combinations in a lower dimensional space. However, there arises a need to unfold the points back to the dimensions of the navigator data which leads us to the following model.

Assumption 3. There exist an $N_k \times d$ matrix $\mathbf{G}_3$ and an $N_k \times N_\ell$ matrix $\mathbf{E}_3$, which gathers all approximation errors, s.t. $\mathbf{A} = \mathbf{G}_3\tilde{\mathbf{A}} + \mathbf{E}_3$.

Matrix $\mathbf{G}_3$ can be viewed as the “decompression” operator which reconstructs the “full” $\mathbf{A}$ from its low-dimensional representation $\tilde{\mathbf{A}}$.

The relationship between the undersampled k-space frames and navigator frames can be established, for the case of Fig. 1c, as $\mathbf{Y}_{\text{nav}} = \mathbf{\Omega}\mathbf{Y}$, where $\mathbf{\Omega}$ is a matrix with binary entries $\{0, 1\}$ that select entries of $\mathbf{Y}$. The inverse problem of recovering $\mathbf{Y}$ from its partial $\mathbf{Y}_{\text{nav}}$ is viable under errors, *e.g.*, an approximation of $\mathbf{Y}$ can be obtained via $\mathbf{\Omega}^\dagger \mathbf{Y}_{\text{nav}}$, where the decompressor $\mathbf{\Omega}^\dagger$ stands for the Moore-Penrose pseudo-inverse of $\mathbf{\Omega}$. The following modeling hypothesis generalizes this argument.

Assumption 4. There exist an $N_k \times \nu N_f$ matrix $\mathbf{G}_1$ and an $N_k \times N_{\text{fr}}$ matrix $\mathbf{E}_1$, which gathers all approximation errors, s.t. $\mathbf{Y} = \mathbf{G}_1 \mathbf{Y}_{\text{nav}} + \mathbf{E}_1$.

A similar modeling assumption can be found in the classical principal component analysis (PCA) [47], where $\mathbf{G}_1$ serves as the “decompression” operator (usually an orthogonal matrix) that maps the “compressed” $\mathbf{Y}_{\text{nav}}$ back to the original data $\mathbf{Y}$ under the approximation error $\mathbf{E}_1$. However, PCA identifies best-linear-fit approximations to data, which might be an inappropriate modeling assumption for several types of data [43]. The following discussion departs from PCA and establishes modeling assumptions that allow non-linear data geometries.

Putting modeling assumptions 2, 3 and 4 together, it can be verified that there exist matrices $\mathbf{G}$ and $\mathbf{E}$ s.t. $\mathbf{Y} = \mathbf{G}\tilde{\mathbf{A}}\mathbf{B} + \mathbf{E}$. Upon defining $\mathbf{U} := \mathcal{F}^{-1}(\mathbf{G})$, and since $\mathbf{G}\tilde{\mathbf{A}}\mathbf{B} = \mathcal{F}(\mathbf{U})\tilde{\mathbf{A}}\mathbf{B} = \mathcal{F}(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B})$, by virtue of the linearity of $\mathcal{F}$, the following bi-linear model between $\mathbf{Y}$ and the unknowns $(\mathbf{U}, \mathbf{B})$

is established:

$$\mathbf{Y} = \mathcal{F}(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B}) + \mathbf{E}. \quad (2)$$

Bi-linearity means that if $\mathbf{U}$ (or $\mathbf{B}$) is fixed to a specific value, then $\mathbf{Y}$ is linear with respect to $\mathbf{B}$ (or $\mathbf{U}$), modulo the error $\mathbf{E}$ term. Interestingly, the linearity of $\mathcal{F}^{-1}$ suggests that the previous modeling hypothesis holds true also in the image domain: $\mathcal{F}^{-1}(\mathbf{Y}) = \mathbf{U}\tilde{\mathbf{A}}\mathbf{B} + \mathcal{F}^{-1}(\mathbf{E})$.

III. THE BI-LINEAR RECOVERY TASK AND ITS ALGORITHMIC SOLUTION

In practice, only few (k,t)-space data are known. To explicitly take account of the limited number of data, a (linear) sampling/masking operator $\mathcal{S}(\cdot)$ is introduced, where $\mathcal{S}(\mathbf{Y})$ leaves the entries of $\mathbf{Y}$ as they are at sampled or observed positions of the k-space domain while nullifying all the rest. The sampling operator is capable of mimicking any sampling strategy, such as Cartesian, radial, spiral, *etc.* It is also often in dMRI that image frames capture a periodic process, *e.g.*, heart movement, other than the static background. In other words, it is reasonable to assume that in (2), the one-dimensional Fourier transform $\mathcal{F}_t$ of the $1 \times N_{\text{fr}}$ time profile of every one of the N_k pixels, *i.e.*, every row of the matrix $\mathcal{F}_t(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B})$, is a sparse vector.

All of the previous modeling assumptions are incorporated in the following recovery task: Given the positive real-valued parameters $\lambda_1, \lambda_2, \lambda_3, C_U$, solve

$$\begin{aligned} \min_{(\mathbf{U}, \mathbf{B}, \mathbf{Z})} \quad & \underbrace{\frac{1}{2} \|\mathcal{S}(\mathbf{Y}) - \mathcal{S}\mathcal{F}(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B})\|_{\text{F}}^2}_{\text{T1}} + \underbrace{\frac{\lambda_1}{2} \|\mathbf{Z} - \mathcal{F}_t(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B})\|_{\text{F}}^2}_{\text{T2}} \\ & + \underbrace{\lambda_2 \|\mathbf{Z}\|_1}_{\text{T3}} + \underbrace{\lambda_3 \|\mathbf{B}\|_1}_{\text{T4}} \\ \text{s.to} \quad & \underbrace{\|\mathbf{U}\mathbf{e}_i\| \leq C_U, \forall i \in \{1, \dots, d\}}_{\text{C1}}, \underbrace{\mathbf{1}_{N_{\text{fr}}}^{\text{T}} \mathbf{B} = \mathbf{1}_{N_{\text{fr}}}^{\text{T}}}_{\text{C2}}; \\ & \mathbf{U} \in \mathbb{C}^{N_k \times d}, \mathbf{B} \in \mathbb{C}^{N_{\text{fr}} \times N_{\text{fr}}}, \mathbf{Z} \in \mathbb{C}^{N_k \times N_{\text{fr}}}, \end{aligned} \quad (3)$$

where $\mathbf{e}_i$ denotes the i th column of the identity matrix $\mathbf{I}_d$. Elaborating more on the recovery task (3), T1 corresponds to the data-fit term, while T2 and T3 introduce the auxiliary variable $\mathbf{Z}$, used to impose a sparsity constraint on $\mathcal{F}_t(\mathbf{U}\tilde{\mathbf{A}}\mathbf{B})$. T4 imposes a sparsity constraint on $\mathbf{B}$, following the discussion on As. 2. Bound C_U in C1 is used to prevent unbounded solutions for $\mathbf{U}$ due to the scaling ambiguity in the bi-linear term $\mathbf{U}\tilde{\mathbf{A}}\mathbf{B}$. C2 adds the affine constraint discussed in As. 2. Moreover, it is worth noticing here that the reduction of dimensionality, achieved via $\tilde{\mathbf{A}}$, reduces also the number of columns and unknowns, *i.e.*, degrees of freedom, of $\mathbf{U}$.

The proposed scheme shares with PS-Sparse [18], STORM [22], (and its other variants bandlimited-STORM [31], navigator-less STORM [35]) and MLS [34] similarities in terms of using navigator lines to learn the underlying data manifold. Instead of using the eigen vectors of the co-variance (PS-Sparse) and Laplacian (STORM and its variants) matrices as the temporal basis, BiLMDM relies on robust sparse embedding [41] to estimate a low dimensional latent structure. STORM and BiLMDM both aim at learning the structure of

Algorithm 1 Recovering the dMRI data

Input: Available are data $\mathcal{S}(\mathbf{Y})$, including the navigator $\mathbf{Y}_{\text{nav}}$ ones. Choose parameters $\lambda_1, \lambda_2, \lambda_3, C_U, \tau_U, \tau_B > 0$, as well as $\zeta \in (0, 1)$ and $\gamma_0 \in (0, 1]$.

Output: Extract the limit points $\mathbf{U}_*$ and $\mathbf{B}_*$ of sequences $(\mathbf{U}_n)_n$ and $(\mathbf{B}_n)_n$, respectively, and recover the dMRI data by the estimate $\hat{\mathbf{X}} := \mathbf{U}_* \tilde{\mathbf{A}} \mathbf{B}_*$.

- 1: Identify landmark points $\mathbf{\Lambda}$ from the columns of $\mathbf{Y}_{\text{nav}}$ according to [39] (*cf.* Sec. II-B).
- 2: Compute the “compressed” $\tilde{\mathbf{A}}$ according to (1).
- 3: Arbitrarily fix $(\mathbf{U}_0, \mathbf{B}_0, \mathbf{Z}_0)$ and set $n = 0$.
- 4: **while** $n \geq 0$ **do**
- 5: Available are $(\mathbf{U}_n, \mathbf{B}_n, \mathbf{Z}_n)$ and γ_n .
- 6: Let $\gamma_{n+1} := \gamma_n(1 - \zeta\gamma_n)$.
- 7: Obtain $\hat{\mathbf{U}}_n$ of (4a) and $\hat{\mathbf{B}}_n$ of (4b) via Alg. 2, and the (i, j) th entry of $\hat{\mathbf{Z}}_n$, $\forall (i, j)$, via the following soft-thresholding rule:

$$[\hat{\mathbf{Z}}_n]_{ij} := [\mathcal{F}_t(\mathbf{U}_n \tilde{\mathbf{A}} \mathbf{B}_n)]_{ij} \cdot \left(1 - \frac{\lambda_2/\lambda_1}{\max\{\lambda_2/\lambda_1, |[\mathcal{F}_t(\mathbf{U}_n \tilde{\mathbf{A}} \mathbf{B}_n)]_{ij}|\}} \right).$$

- 8: Update $(\mathbf{U}_{n+1}, \mathbf{B}_{n+1}, \mathbf{Z}_{n+1}) := (1 - \gamma_{n+1})(\mathbf{U}_n, \mathbf{B}_n, \mathbf{Z}_n) + \gamma_{n+1}(\hat{\mathbf{U}}_n, \hat{\mathbf{B}}_n, \hat{\mathbf{Z}}_n)$.
- 9: Set n equal to $n + 1$ and go to step 4.
- 10: **end while**

“smooth” data manifolds. Smoothness in STORM is established globally over the data cloud by penalizing the recovery task via several operator norms of a graph-Laplacian matrix. On the other hand, BiLMDM enforces smoothness locally over the manifold via neighborhoods and affine approximations (patches) of the tangent spaces (*cf.* Fig. 2), and the data manifold can be viewed as an approximation of the union of all those patches (as discussed in As. 2). PCA-inspired schemes, such as PS-Sparse, search also for local affine/linear patches to manifolds. Nevertheless, BiLMDM does not impose any orthogonality constraints on any of its factors $\mathbf{U}$ and $\mathbf{B}$, as PCA via the singular value decomposition (SVD) does. Identification of landmark points to learn an underlying manifold is a novel contribution of the proposed work, given that all of the prior-art methods rely on the entire data cloud for learning. The incorporation of the landmark points in the bi-linear factorization model $\mathbf{U}\tilde{\mathbf{A}}\mathbf{B}$ serves also as a natural encapsulation of the underlying data geometry into the optimization task, extending thus the usual way that PCA interprets its factors as an orthogonal/decorrelating basis and coefficients. Given that MLS employs the same approach of robust sparse embedding, the key feature of BiLMDM that sets itself apart from MLS, as well as other manifold-learning methods, is the bi-linear factorization model used to represent the MR data cloud. This bi-linear factorization model helps to capture the geometry of the point cloud locally by imposing a sparsity-promoting penalty on the affine combinations, thus, restricting the sharing

of data among frames that show common phase and structural characteristics; local exploitation of dependencies is promoted over any global one. BiLMDM offers the flexibility of updating (per iteration) *both* the spatial ($\mathbf{U}$) and temporal parts ($\mathbf{B}$) of the data-matrix decomposition. On the contrary, PS-Sparse and MLS follow a two-stage strategy, where first a temporal basis is sought, and after this temporal basis is fixed, a spatial basis is computed. Such a bi-linear model differs also from LASSI [23], which adopts the popular approach of viewing the data matrix as the superposition of a low-rank and a sparse component. The low-rank component describes the static image background, while the sparse component, modeled via a classical bi-linear dictionary-learning term, describes the dynamic foreground. In contrast to LASSI, BiLMDM uses bi-linear modeling to describe the dynamic foreground and the static background together. BiLMDM does not also require any pre-processing procedures, like sorting the acquired data into cardiac and respiratory phases, as in the motion resolved strategies of XD-GRASP [24], before solving the reconstruction task.

The successive-convex-approximation framework of [48] is employed to solve (3) and is presented in a concise form in steps 4–10 of Alg. 1. Convergence to a stationary solution of (3) is guaranteed [48]. Step 7 of Alg. 1 comprises convex minimization sub-tasks. More specifically, at every step of the algorithm, given $(\mathbf{U}_n, \mathbf{B}_n, \mathbf{Z}_n)$, the following estimates are required (for $\tau_U, \tau_B > 0$):

$$\begin{aligned} \hat{\mathbf{U}}_n \in \arg \min_{\mathbf{U}} \frac{1}{2} \|\mathcal{S}(\mathbf{Y}) - \mathcal{SF}(\mathbf{U} \tilde{\mathbf{A}} \mathbf{B}_n)\|_{\mathbf{F}}^2 + \frac{\tau_U}{2} \|\mathbf{U} - \mathbf{U}_n\|_{\mathbf{F}}^2 \\ + \frac{\lambda_1}{2} \|\mathbf{Z}_n - \mathcal{F}_t(\mathbf{U} \tilde{\mathbf{A}} \mathbf{B}_n)\|_{\mathbf{F}}^2 \\ \text{s.to } \|\mathbf{U} \mathbf{e}_i\| \leq C_U, \forall i \in \{1, \dots, d\}. \end{aligned} \quad (4a)$$

$$\begin{aligned} \hat{\mathbf{B}}_n \in \arg \min_{\mathbf{B}} \frac{1}{2} \|\mathcal{S}(\mathbf{Y}) - \mathcal{SF}(\mathbf{U}_n \tilde{\mathbf{A}} \mathbf{B})\|_{\mathbf{F}}^2 + \frac{\tau_B}{2} \|\mathbf{B} - \mathbf{B}_n\|_{\mathbf{F}}^2 \\ + \frac{\lambda_1}{2} \|\mathbf{Z}_n - \mathcal{F}_t(\mathbf{U}_n \tilde{\mathbf{A}} \mathbf{B})\|_{\mathbf{F}}^2 + \lambda_3 \|\mathbf{B}\|_1 \\ \text{s.to } \mathbf{1}_{N_\ell}^T \mathbf{B} = \mathbf{1}_{N_{fr}}^T. \end{aligned} \quad (4b)$$

Both tasks in (4) can be viewed as affinely constrained composite convex minimization tasks, hence allowing the use of [46], as described in Alg. 2. From a computational complexity perspective, it is worth pointing out that the proposed scheme relies on minimization sub-tasks, and computational complexities depend thus on the solver adopted for solving those sub-tasks. Here, [46] employs only first-order information (gradients) and proximal mappings, *e.g.*, soft-thresholding rules and projection mappings. The implementation of [46] for the specific tasks (4) is presented in Alg. 2, and details are deferred to the appendix section of this manuscript. Notice also that the previous minimization sub-tasks can be solved in parallel. Furthermore, problems 4 can be solved inexactly at each iteration (with increasing precision, as described in [48]), which contributes to reducing the computation cost of each iteration; we refer to [48] for more details.

IV. NUMERICAL RESULTS

The proposed framework is tested and validated on three datasets, used also in [17], [18], [30], [34], [49]: **i)** Magnetic-resonance extended cardiac-torso (MRXCAT) cine phantom [50]; **ii)** cardiac phantom generated from real MR

Algorithm 2 Computing $\hat{\mathbf{U}}_n$ of (4a) and $\hat{\mathbf{B}}_n$ of (4b)

Input: $\mathbf{D}$ is either $\mathbf{U}_n$ or $\mathbf{B}_n$ in step 7 of Alg. 1. Choose parameter K_0 , $\alpha \in [0.5, 1)$.

Output: $\hat{\mathbf{U}}_n$ and $\hat{\mathbf{B}}_n$ are set equal to the limit $\mathbf{H}_{K_0}$ of the sequence $(\mathbf{H}_k)_k$.

- 1: Compute the Lipschitz coefficient L via (13) or (20), and choose $\lambda \in (0, 2(1 - \alpha)/L)$.
- 2: $\mathbf{H}_0 := \mathbf{D}$.
- 3: Letting T be either (15) or (22), define $T_\alpha := \alpha T + (1 - \alpha) \text{Id}$, where Id denotes the identity operator.
- 4: Set $\mathbf{H}_{1/2} := T_\alpha(\mathbf{H}_0) - \lambda \nabla g_1(\mathbf{H}_0)$, where the gradient $\nabla g_1(\mathbf{H}_0)$ takes the form of either (11) or (19).
- 5: Set $\mathbf{H}_1 := \text{Prox}_{\lambda g_2}(\mathbf{H}_{1/2})$, where the proximal operator takes the form of either (14) or (21).
- 6: **while** $k \leq K_0$ **do**
- 7: $\mathbf{H}_{k+3/2} := \mathbf{H}_{k+1/2} + T(\mathbf{H}_{k+1}) - \lambda \nabla g_1(\mathbf{H}_{k+1}) - T_\alpha(\mathbf{H}_k) + \lambda \nabla g_1(\mathbf{H}_k)$.
- 8: $\mathbf{H}_{k+2} := \text{Prox}_{\lambda g_2}(\mathbf{H}_{k+3/2})$.
- 9: Set k equal to $k + 1$ and go to step 6.
- 10: **end while**

scans [18] **iii)** prospectively undersampled free breathing, real time, cardiac cine data. All experiments were conducted on a 12-core Intel(R) 2.40 GHz Linux-based system with 48GB RAM, with implementations realized in MATLAB [51]. The proposed method is compared with: PS-Sparse [18], MLS [34], STORM [22], LASSI [23] and XD-GRASP [24].

The publicly available MATLAB implementations of STORM [52], LASSI [53] and XD-GRASP [54] were employed. MATLAB code was also written to realize the algorithmic solutions of [46] to the convex-optimization problems posed in PS-Sparse and MLS. Parameters were tuned to produce the least NRMSE for each algorithm at each sampling ratio and on every dataset.

Tasks (3), (4a) and (4b) require parameter tuning to produce the desirable MR images from the scanner acquired data. These parameters were chosen empirically to minimize the reconstruction error for the datasets where the gold standard was known. A range of values for every parameter was realized, and as long as these parameters are within that range, the NRMSEs were not so sensitive to the parameters. This empirical research into the behaviour of the parameter values revealed the following: for hyper parameters penalizing the periodicity along the time series (λ_1, λ_2), higher the ratio λ_2/λ_1 , severe temporal bleeding was observed; hence, it should be low enough to avoid temporal bleeding and at the same time exploit the periodicity along the temporal axis. The parameter which penalizes the sparsity on $\mathbf{B}$, λ_3 , controls the size of the neighborhood used for approximation of the time series image. Higher the value of λ_3 , lower the size of the neighborhood, allowing thus frames with very similar characteristics to share data among themselves. However, decreasing the value excessively would result in dissimilar frames influencing each other. Parameter N_ℓ determines the number of landmark points which act as a representative of a collection of frames with similar phase and motion characteristics. The value for N_ℓ




Fig. 3. NRMSE values [cf. (5)] computed for (a) MRXCAT cardiac cine data, and (b) real cardiac cine data vs. acceleration/undersampling rates. The NRMSE plot for BiLMDM, given the non-convex nature of the recovery task, is averaged over 25 independent trials, with different initialization points for each trial. Error bars are also used to indicate the deviation from the sample means due to the random initializations of the non-convex algorithmic scheme. There are points where the error bars are too small to be clearly visible.

should be high enough to describe the data cloud concisely. For the datasets used in this work, desirable values range between 15%–20% of N_{fr} . Any increase N_ℓ beyond that didn't yield any significant changes in the image quality. Bounding constant C_U set to 1 usually yields adequate estimates for $\mathbf{U}$, although it is worth pointing out that the estimates for $\mathbf{U}$ are not sensitive to a change in C_U . The step sizes for convergence of Alg. 1 ($\gamma_0 = 0.9$, $\zeta_0 = 0.001$) and the step size for Alg. 2 ($\alpha = 0.5$) ensures desirable results. These values were reached by observing the effects of these parameters on the NRMSE values for the reconstruction of the two synthetically generated cardiac cine dataset.

Quality of reconstruction is evaluated by the normalized-root-mean-square error (NRMSE), defined as

$$\text{NRMSE} := \frac{\|\mathbf{X} - \hat{\mathbf{X}}\|_F}{\|\mathbf{X}\|_F}, \quad (5)$$

where $\mathbf{X}$ represents the fully sampled, high-fidelity and original image-domain data, while $\hat{\mathbf{X}}$ represents an estimate of $\mathbf{X}$ computed by the reconstruction schemes. In addition to the




Fig. 4. Framewise-NRMSE values for (a) MRXCAT cardiac cine data (acceleration rate: 20x), and (b) real cardiac cine data (acceleration rate: 17x) for PS-Sparse [(a) $0.054 \pm 3 \times 10^{-3}$ (b) $0.0863 \pm 2.1 \times 10^{-2}$], MLS [(a) $0.051 \pm 3.1 \times 10^{-3}$ (b) $0.0892 \pm 2 \times 10^{-2}$], StoRM [(a) $0.0844 \pm 2.6 \times 10^{-2}$ (b) $0.0583 \pm 6.6 \times 10^{-3}$], LASSI [(a) $0.0732 \pm 6 \times 10^{-3}$ (b) $0.0791 \pm 3.1 \times 10^{-2}$] XD-GRASP [(a) $0.065 \pm 2.5 \times 10^{-2}$ (b) $0.0587 \pm 2.7 \times 10^{-2}$] and BiLMDM [(a) $0.0488 \pm 1.6 \times 10^{-3}$, (b) $0.0527 \pm 6.4 \times 10^{-3}$]. The previous numerical values demonstrate the (sample mean over all N_{fr} frames) $\pm$ (standard deviation from the sample mean). Consistent with simulation results, BiLMDM exhibits less fluctuations in error across all the frames over the rest of the methods.

NRMSE, the effectiveness of the algorithms are also inspected along the edges using the high-frequency error norm (HFEN) and sharpness measures. HFEN is defined as

$$\text{HFEN} = \frac{\|\text{LoG}(\mathbf{X}) - \text{LoG}(\hat{\mathbf{X}})\|_2}{\|\text{LoG}(\mathbf{X})\|_2}, \quad (6)$$

where LoG is a rotationally symmetric Laplacian of Gaussian filter. As in [55], a filter kernel of size 15×15 with standard deviation of 1.5 pixels is employed. Two sharpness measures are used: M1 (intensity variance based) and M2 (energy of the image gradient based), as described in [56]. Better insights into the structural information are extracted using the Structural Similarity (SSIM) index [57] which investigates the local patterns in the pixel intensities after normalizing for luminance and contrast. The metrics (except for sharpness measure) can be used for assessment only when the ground truth images are present and hence are used for the two phantom datasets and not the prospectively undersampled data.




Fig. 5. Temporal cross-sections for MRXCAT cardiac cine (acceleration rate: 20x). Left to right: Gold standard (spatial frame), gold standard (temporal cross-section), PS-Sparse (0.055), MLS (0.0512), STORM (0.0844), LASSI (0.0732), XD-GRASP (0.0653) and BiLMDM ($0.0488 \pm 2.5 \times 10^{-4}$). The previous numerical values indicate the NRMSE for the complete dataset, in addition to the standard deviation (for BiLMDM only) obtained after running the non-convex algorithmic scheme for 25 independent trials. Top to bottom: Temporal cross section and error maps. The temporal location of the frames is indicated by the yellow dotted line in the gold-standard (spatial-frame) image.

TABLE I
RECONSTRUCTION OF MRXCAT PHANTOM UNDER RADIAL
SAMPLING TRAJECTORIES: NRMSE VALUES

	4x	8x	12x	16x	20x	24x
PS-Sparse	0.0538	0.0556	0.0578	0.0591	0.0621	0.073
MLS	0.049	0.0521	0.0562	0.0573	0.06	0.072
STORM	0.0491	0.0503	0.0521	0.0564	0.0595	0.0634
LASSI	0.0514	0.0548	0.0593	0.0654	0.069	0.082
XD-GRASP	0.048	0.0497	0.0523	0.058	0.064	0.0673
BiLMDM	0.0441	0.0456	0.0473	0.0489	0.0498	0.0512

The proposed framework is validated over a range of undersampling/acceleration rates, defined by $N_k N_{fr} / (\# \text{ of acquired voxels})$. 1-D Cartesian (Fig. 1a) as well as radial (Fig. 1b) sampling were applied to both the phantom datasets. To save space, only the 1-D Cartesian-sampling results are demonstrated for the MRXCAT dataset, while radial-sampling ones are shown for the phantom generated using real MR scans. Nevertheless, BiLMDM's performance against the competing reconstruction algorithms follows a similar trend also for sampling strategies not included in the manuscript due to space limitations.

A. MRXCAT phantom

The extended cardiac torso (XCAT) framework was used to generate the MRXCAT phantom [50]; a breath-hold cardiac cine data of spatial size $(N_p, N_f) = (408, 408)$ corresponding to a spatial resolution of $1.56 \times 1.56 \text{ mm}^2$ for a FOV of $400 \times 400 \text{ mm}^2$. The cardiac cine phantom generated is spread across $N_{fr} = 360$ time frames in the temporal direction, consisting of 15 cardiac cycles and 24 cardiac phases. The MRXCAT-phantom dataset is characterized by its limited inter-frame variations and periodic nature along the temporal direction, unlike the other two datasets.

On the quantitative front, Fig. 3a provides quantitative comparisons in terms of NRMSE for a range of acceleration factors; while Fig. 4a provides the fluctuations in the NRMSE across the time frames in addition to Tab. I and II which summarizes the NRMSE values under radial sampling and all the performance metrics for an acceleration factor of 20x

respectively. It is evident from Fig. 3a that the proposed BiLMDM consistently outperforms the state-of-the-art schemes over the entire range of acceleration rates. It is worth noticing here that BiLMDM scores similar NRMSE values to PS-Sparse and MLS, at specific acceleration rates, even though it uses a small subset of the navigator data (landmark points) in this case $N_\ell = 50$, $d = 4$, in contrast to PS-Sparse and MLS that utilize the whole set of navigator data. BiLMDM's performance is consistent for every frame in the time series as displayed in Fig. 4a, and exhibits the least NRMSE and least NRMSE deviation for every frame, unlike XD-GRASP, LASSI and STORM. Tab. II provides evidence that BiLMDM not only produces result closest to the ground truth, as exhibited by both the NRMSE and SSIM values, but also produces the sharpest image of all indicated by the smallest HFEN value and the largest M1 measure supported by a high M2 measure.

Shifting focus towards the qualitative aspect, Fig. 6 focuses on a dynamic region of interest (ROI) from the end-diastolic and end-systolic phases. The proposed framework exhibits reconstruction of a high-resolution image with contrast close to the gold standard, as opposed to the alias-infected reconstruction of XD-GRASP and STORM or the noisy reconstructions of PS-Sparse, MLS and LASSI. The error maps in Figs. 5 and 6 provide visual proof of the improvements in reconstruction. BiLMDM shows significant improvements at the image edges, as visualized in the error maps for both the diastolic and systolic phases and further supported by HFEN, M1 and M2 values. Moreover, PS-Sparse, MLS and LASSI exhibit faint-right artifacts (marked by red arrows in Fig. 6), unlike BiLMDM. The temporal cross sections in Fig. 5 are consistent with the spatial results in Fig. 6. Reconstructions from XD-GRASP and STORM exhibit temporal blurring, while LASSI reconstructions show presence of noisy artifacts in Fig. 5. STORM, LASSI and XD-GRASP also witness breaks in the periodicity of the cardiac phases as pointed by red arrows in Fig. 5.

The reconstruction times required by each method on an average are as follows: PS-Sparse (5 hours), MLS (5 hours), STORM (43 mins), LASSI (1.25 hours), XD-GRASP (1.5




Fig. 6. Spatial results for region of interest in (marked with a yellow rectangle in the gold standard spatial-frame image of Fig. 5) MRXCAT cardiac cine (acceleration rate: 20x). Left to right: Gold standard, PS-Sparse (0.055), MLS (0.0512), STORM (0.0844), LASSI (0.0732), XD-GRASP (0.0653) and BiLMDM ($0.0488 \pm 2.5 \times 10^{-4}$). The previous numerical values indicate the NRMSE for the complete dataset, in addition to the standard deviation (for BiLMDM only) obtained after running the non-convex algorithmic scheme for 25 independent trials. Top to bottom: Diastole phase (frame 1 of the time series), the under-sampled image followed by error maps, systole phase (frame 12 of the time series) and the under-sampled image followed by error maps.

TABLE II
QUANTITATIVE PERFORMANCE ANALYSIS FOR MRXCAT
PHANTOM (ACCELERATION RATE: 20X)

	NRMSE	SSIM	HFEN	M1	M2
PS-Sparse	0.055	0.8913	0.1548	4856.8	1.7×10^7
MLS	0.0512	0.8905	0.1558	4855.1	1.7×10^7
STORM	0.0844	0.776	0.2882	4594	1.9×10^7
LASSI	0.0732	0.7954	0.1785	4616.1	1.6×10^7
XD-GRASP	0.0653	0.8205	0.2478	4764.3	1.8×10^7
BiLMDM	0.0488	0.9218	0.1314	4958.1	1.8×10^7

hours) and BiLMDM (5.5 hours). Experiments were conducted for parameter $N_\ell = 360$ which basically skips the landmark selection process and considers all the 360 frames for the identification of the underlying low-dimensional manifold. The reconstructions took around 11.5 hours due to the cumbersome RSE procedure employed for manifold learning followed by the determination of the increased number of unknowns in $\mathbf{B}$. The experiment yielded no significant increase in NRMSE value. The reconstructions from k-space acquisitions sampled under an acceleration rate of 20x (Cartesian trajectory) exhibited an NRMSE of 0.0468 compared to the 0.0488 obtained when N_ℓ was set to 50.

B. Cardiac phantom generated from real MR scan

Besides synthetically generated MR images, real human-cardiac cine MR data were also acquired [18] with the

TABLE III
QUANTITATIVE PERFORMANCE ANALYSIS FOR CARDIAC
PHANTOM GENERATED FROM REAL MR SCANS (ACCELERATION
RATE: 17X)

	NRMSE	SSIM	HFEN	M1	M2
PS-Sparse	0.0863	0.9216	0.1375	2.6×10^{-6}	175.6
MLS	0.0892	0.9134	0.2684	2.6×10^{-6}	176.6
STORM	0.0583	0.9336	0.1375	2.6×10^{-6}	191.9
LASSI	0.0791	0.7425	0.368	2.2×10^{-6}	180.7
XD GRASP	0.0587	0.9348	0.1193	2.6×10^{-6}	194.4
BiLMDM	0.0527	0.946	0.089	2.7×10^{-6}	194.6

following parameters: spatial size $(N_p, N_f) = (200, 256)$, $\text{FOV} = 273 \times 350 \text{ mm}^2$, resulting in a spatial resolution of $1.36 \times 1.36 \text{ mm}^2$. The data were acquired during a single breath hold. Multiple time-wraps (introducing temporal variations) and quasi-periodic spatial deformation [58] (to model respiration) were then used to generate a temporal sequence with $N_{fr} = 256$ frames. The results in Fig. 3b, similar to those on the previous cine dataset, show that BiLMDM outperforms the rest of the techniques with NRMSE values ranging from 0.05 to 0.061 for $N_\ell = 35$, $d = 12$. Fig. 4b, further supports the ability of BiLMDM to create reconstructions with the least NRMSE fluctuations across the many frames. As seen in Tab. III, BiLMDM exhibits the least HFEN, M1 and M2 measure values which are testament to the fact that reconstructions obtained are sharper and are precise around the edges. The high similarity




Fig. 7. Temporal cross-sections for real cardiac cine data (acceleration rate: 17x). Left to right: Gold standard (spatial frame), gold standard (temporal cross-section), PS-Sparse (0.0863), MLS (0.0892), STORM (0.0583), LASSI (0.0791), XD-GRASP (0.0587) and BiLMDM ($0.0527 \pm 1.1 \times 10^{-4}$). The previous numerical values indicate the NRMSE for the complete dataset, in addition to the standard deviation (for the proposed scheme only) obtained after running the non-convex algorithmic scheme for 25 independent trials. Top to bottom: Temporal cross section and error maps. The temporal location of the frames is indicated by the yellow dotted line in the gold-standard (spatial-frame) image.




Fig. 8. Spatial results for region of interest marked with a yellow rectangle in the gold standard (spatial-frame) image of Fig. 7 real cardiac cine data (acceleration rate: 17x). Left to right: Gold standard, PS-Sparse (0.0863), MLS (0.0892), STORM (0.0583), LASSI (0.0791), XD-GRASP (0.0587) and BiLMDM ($0.0527 \pm 1.1 \times 10^{-4}$). The previous numerical values indicate the NRMSE for the complete dataset, in addition to the standard deviation (for BiLMDM only) obtained after running the non-convex algorithmic scheme for 25 independent trials. Top to bottom: Diastole phase (frame 1 of the time series), the under-sampled image followed by error maps, systole phase (frame 16 of the time series) and the under-sampled image followed by error maps.

index proves that the BiLMDM is capable of producing results closest to the ground truth.

With regards to Figs. 7 and 8, significant blurring and deformations can be observed for PS-Sparse, MLS and LASSI, in contrast to the sharper images produced by BiLMDM for both the end-diastolic and end-systolic phases. There are slight improvements over STORM and XD-GRASP, which are illustrated via the error maps. At the same time, it is worth noting that in the end-systolic phase, PS-Sparse, MLS and LASSI produce an enlargement (deformation) of the area marked with a red rectangle on the image. The temporal cross sections in Fig. 7 indicate significant motion blurring for PS-Sparse, MLS and LASSI. Moreover, it can be verified that a

large number of artifacts are present in the temporal cross section produced by LASSI. Significant temporal bleeding is observed in results from PS-Sparse and MLS which is consistent with the undesirable enlargement in the end-systolic phase discussed earlier. Even though STORM and XD-GRASP provide reconstructions almost as good as BiLMDM, it can be seen that there are signs of temporal blurring in the marked areas in Fig. 7 for the corresponding methods. On the contrary, BiLMDM provides a temporal cross-section very close to the gold standard. The reconstruction times required by each method are as follows: PS-Sparse (3.75 hours), MLS (3.75 hours), STORM (32 mins), LASSI (1 hour), XD-GRASP (1 hour) and BiLMDM (4 hours).




Fig. 9. Spatial results for prospectively undersampled cardiac cine data (acceleration rate: 10x). Left to right: PS-Sparse, MLS, STORM, LASSI, XD-GRASP and BiLMDM. Top to bottom: Reconstructed frame 1 of the time series, the region of interest (cardiac region marked with a yellow rectangle in the images of the topmost row), temporal cross-section of the reconstructed time series along the yellow dotted line in the images of the topmost row.

C. Prospectively Undersampled Cardiac Cine Data

The prospectively undersampled real-time free breathing cardiac cine data was acquired using a FLASH sequence from a volunteer breathing normally under the following acquisition parameters: TR/TE = 5.8/4 ms, FOV 284×350 mm, spatial resolution = 1.8 mm. A 12 channel scanner was used to continuously acquire 4500 phase encoding lines under 1-D Cartesian trajectory. These phase encoded lines were divided in groups of 15 (including 5 navigator lines) to form 300 frames, resulting in a data matrix of size $156 \times 192 \times 300$ for each channel. The above acquisition corresponds to a undersampling factor of nearly 10x. For fairness in comparison, the multi-channel data was reconstructed coil by coil for all methods and then combined with the sum-of-squares strategy.

Fig. 9 shows the reconstruction results from the proposed and state-of-the methods. The validation skips most of the performance metrics except for the sharpness measure due to the absence of a ground truth. Hence the validation is solely dependent on the visual quality as shown in Fig. 9 and the sharpness measures. The full FOV images exhibit that BiLMDM along with STORM, unlike other methods show well-defined structures and result in a sharper image. LASSI exhibits some aliasing effects, while PS-Sparse and MLS observes some blurring in the spatial frame. This is further supported by sharpness measures (M1 followed by M2): PS-Sparse (8.0×10^{-8} , 81.7), MLS (8.2×10^{-8} , 79.9), STORM (8.6×10^{-7} , 91.8), LASSI (8.2×10^{-8} , 83.6), XD-GRASP (8.2×10^{-7} , 80.6) and BiLMDM (8.8×10^{-7} , 91.7). Comparing the temporal cross sections for the reconstructed images, even though the FOV images look pretty sharp for STORM and BiLMDM, STORM suffers from more temporal blurring than BiLMDM, while LASSI and PS-Sparse show grainy artifacts in the temporal cross section. The reconstruction times required by each method are as follows: PS-Sparse (5 hours), MLS (5 hours), STORM (35 mins), LASSI (1.5 hours), XD-GRASP

(1.5 hours) and BiLMDM (5.25 hours)

D. Significance of matrices $\mathbf{U}$, $\tilde{\mathbf{A}}$ and $\mathbf{B}$

BiLMDM decomposes the MR image series into three matrices $\mathbf{U}$, $\tilde{\mathbf{A}}$ and $\mathbf{B}$, resulting in distinct identification of the “spatial” and “temporal bases” for the given MR image series. Fig. 10 exhibits the ability of matrix $\mathbf{U}$ to capture the spatial features of the image, like cardiac structures, edges, contrast, etc., from a highly undersampled k-space acquisition (leftmost image in Fig. 10). The captured spatial basis obtained by reshaping two of the $d = 4$ (parameter value as set in Sec. IV-A) columns of $\mathbf{U}$ to the image dimensions is shown in Fig. 10.

Matrix $\tilde{\mathbf{A}}$ gathers landmark points to offer a low-dimensional approximation of the high-dimensional manifold onto which each temporal frame lies, while $\mathbf{B}$ affinely combines the low-dimensional landmark points to capture any inter-frame dependencies. Hence, product $\tilde{\mathbf{A}}\mathbf{B}$ may be thought to serve as the temporal basis describing the cardiac cycles and phases of the MR image series. Fig. 11 presents the 2^{nd} and 3^{rd} row of the product $\tilde{\mathbf{A}}\mathbf{B}$, which correspond to the 2^{nd} and 3^{rd} “elements” of the temporal basis for the MRXCAT phantom. Evidently, this temporal basis is successful in capturing the 15 cardiac cycles of MRXCAT phantom and its periodic nature. It is also worth noting that the temporal basis generated from only four navigator lines in k-space (represented by the blue dashed line in Fig. 11) sufficiently approximates the temporal basis acquired from a fully sampled k-space data (ground truth curve in the Fig 11), thus overcoming the need to employ the entire dataset in the estimation of the ground-truth temporal basis.

V. DISCUSSION

The proposed BiLMDM scheme relies on data sharing among different frames by approximating each frame by an




Fig. 10. Left to right: MRXCAT phantom data acquired under 12x acceleration rate followed by columns of matrix $\mathbf{U}$ reshaped to image dimensions.




Fig. 11. Comparisons of the reconstructed temporal basis with that of the ground truth. Left to Right: 2nd and 3rd temporal basis capturing the cardiac cycles and phases.

affine combination of its neighboring frames on the manifold that describes locally the MR data cloud. This manifold and sparse-approximation based recovery model enables the reconstruction of desirable, good quality and artifact free MR images from severely undersampled k-space measurements. The quality of the reconstructed images and the corresponding performance in numbers discussed under Sec. IV is a testament to the fact that the BiLMDM outperforms the competing state-of-the-art methods. BiLMDM has consistently provided better results for both synthetically generated, retrospectively undersampled as well as experimentally acquired, prospectively undersampled cardiac cine data.

Methods like PS-Sparse and MLS have reconstructed desirable MR images in the case of the MRXCAT phantom, under cartesian sampling trajectories; however, they suffered severe motion blurring and deformations (enlargements) in the case of the cardiac cine data retrospectively sampled via radial trajectories. In contrast, STORM produced good quality images for the data using radial sampling but suffered severe spatial blurring and temporal bleeding in the cases of Cartesian sampling. However, BiLMDM reconstructions have been artifact free, without any temporal bleeding irrespective of what sampling strategy was used to acquire data. Similar results to the ones discussed here hold also for the sampling strategies not included in the manuscript, for all employed algorithms. Even in the case of the experimentally acquired data, where techniques like PS-Sparse, MLS, XD-GRASP and LASSI suffered from some missing structures and distortions, while STORM suffered temporal blurring, BiLMDM produced images which seem to best preserve the features of the cardiac structure spatially and temporally. BiLMDM seems capable of reconstructing desirable MR images irrespective of the sampling strategy into consideration, which appears to be a

dominant factor in the competing state-of-the-art methods. The extension of BiLMDM to kernel-function spaces is currently under development, and a detailed discussion of the same in addition to numerical comparisons versus various other kernel-based techniques, such as the very recent versions of STORM [31], [32], [35], is reserved for a future journal publication.

As seen previously in the numerical validations, the proposed framework currently underperforms in terms of computation times versus the competing algorithms and their *optimized* software implementations. The high computation times are a result of the inefficient implementation of the proposed algorithm. The implementation can benefit from the use of special and optimized packages along with parallel processing. The present manuscript serves as a proof of concept of BiLMDM's potential, while an optimized and time-efficient C/C++ version of the developed MATLAB code is still a work in progress and will be presented in the near future via an open-source software-distribution venue.

VI. CONCLUSIONS

This paper proposed the novel bi-linear modeling for data manifolds (BiLMDM); a new framework for data reconstruction using manifold-learning and sparse-approximation arguments. BiLMDM comprises several modules: Extracting a set of landmark points from a data cloud helps in learning the latent manifold geometry while identifying low-dimensional renditions of the landmark points facilitates efficient means for data storage and computations. Finally, a bi-linear optimization task is used to achieve data recovery. Quantitative and qualitative analyses on dynamic MRI data, described in Sec. IV, provided evidence that the proposed BiLMDM achieves improvements in dMRI image reconstruction and artifact suppression over state-of-the-art approaches such as PS-Sparse, MLS, LASSI, XD-GRASP and STORM in cardiac cine data. BiLMDM reconstruction is not limited to cine MR data but has also produced reliable reconstructions for perfusion MR data. The efficacy of BiLMDM on perfusion data and extensive comparison tests versus state-of-the-art schemes are reserved for a future publication. By introducing BiLMDM, this work paves the way for devising efficient ways to utilize fewer data points for reconstruction than state-of-the-art solutions and opens the door for further advances in geometric data-approximation methods.

APPENDIX A MATHEMATICAL PRELIMINARIES

For positive integers m, n , the space of matrices $\mathbb{C}^{m \times n}$ is equipped with the inner product $\langle \mathbf{A} | \mathbf{B} \rangle := \text{trace}(\mathbf{A}^H \mathbf{B})$, $\forall \mathbf{A}, \mathbf{B} \in \mathbb{C}^{m \times n}$, where the superscript H denotes the Hermitian transpose of a matrix. It is worth noting that the inner product is not commutative: $\langle \mathbf{B} | \mathbf{A} \rangle = \overline{\langle \mathbf{A} | \mathbf{B} \rangle}$, where the overline symbol denotes the complex conjugate of a number. The induced norm of $\mathbb{C}^{m \times n}$ by the previous inner product coincides with the Frobenius norm of a matrix: $\|\cdot\|_F = \langle \cdot | \cdot \rangle^{1/2}$. Moreover, the spectral norm $\|\mathbf{A}\|_2$ of matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is defined as $\lambda_{\max}^{1/2}(\mathbf{A}^H \mathbf{A})$, where $\lambda_{\max}(\cdot)$ stands

for the maximum eigenvalue of a symmetric matrix. In the case where $m = n$, then $\|\mathbf{A}\|_2 = \lambda_{\max}(\mathbf{A})$.

Given the positive integers m_1, m_2, n_1, n_2 and the linear mapping $\mathcal{L} : \mathbb{C}^{m_1 \times n_1} \rightarrow \mathbb{C}^{m_2 \times n_2}$, the adjoint of $\mathcal{L}$ is the linear mapping $\mathcal{L}^* : \mathbb{C}^{m_2 \times n_2} \rightarrow \mathbb{C}^{m_1 \times n_1}$ defined as $\langle \mathbf{A} | \mathcal{L}(\mathbf{B}) \rangle = \langle \mathcal{L}^*(\mathbf{A}) | \mathbf{B} \rangle$, $\forall \mathbf{A} \in \mathbb{C}^{m_2 \times n_2}$, $\forall \mathbf{B} \in \mathbb{C}^{m_1 \times n_1}$. For example, with regards to the sampling mapping $\mathcal{S}(\cdot)$ in Sec. III, its adjoint $\mathcal{S}^*(\cdot) = \mathcal{S}(\cdot)$, i.e., $\mathcal{S}(\cdot)$ is self-adjoint, and $\mathcal{S}^2(\cdot) = \mathcal{S}(\cdot)$. Moreover, for the MATLAB implementation of the Fourier transform [51], $\mathcal{F}^* = N_k \mathcal{F}^{-1}$ and $\mathcal{F}_t^* = N_{fr} \mathcal{F}_t^{-1}$. If matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ is viewed as a linear mapping $\mathbf{A} : \mathbb{C}^n \rightarrow \mathbb{C}^m$, then $\mathbf{A}^* = \mathbf{A}^H$. Mapping $\mathcal{L}(\cdot)$ is called Lipschitz continuous, with coefficient $L > 0$, if $\|\mathcal{L}(\mathbf{A}) - \mathcal{L}(\mathbf{B})\|_F \leq L \|\mathbf{A} - \mathbf{B}\|_F$, $\forall \mathbf{A}, \mathbf{B}$.

Given a convex function $g(\cdot) : \mathbb{C}^{m \times n} \rightarrow \mathbb{R}$ and a positive real number λ , the proximal mapping $\text{Prox}_{\lambda g}(\cdot) : \mathbb{C}^{m \times n} \rightarrow \mathbb{C}^{m \times n}$ is defined as $\text{Prox}_{\lambda g}(\mathbf{A}) := \arg \min_{\mathbf{B}} \lambda g(\mathbf{B}) + (1/2) \|\mathbf{A} - \mathbf{B}\|_F^2$. For example, in the case where g becomes the indicator function ι_C with respect to a closed convex set $C \subset \mathbb{C}^{m \times n}$, i.e., $\iota_C(\mathbf{A}) := 0$, if $\mathbf{A} \in C$, while $\iota_C(\mathbf{A}) := +\infty$, if $\mathbf{A} \notin C$, then $\text{Prox}_{\lambda \iota_C}$ becomes the metric projection mapping onto C : $\text{Prox}_{\lambda \iota_C}(\mathbf{A}) = \arg \min_{\mathbf{B} \in C} \|\mathbf{A} - \mathbf{B}\|_F$. Moreover, in the case where g is the ℓ_1 -norm $\|\cdot\|_1$, then the (i, j) th entry of $\text{Prox}_{\lambda \|\cdot\|_1}(\mathbf{A})$ is given by the soft-thresholding rule [59, Lemma V.I]

$$[\text{Prox}_{\lambda \|\cdot\|_1}(\mathbf{A})]_{ij} = [\mathbf{A}]_{ij} \left(1 - \frac{\lambda}{\max\{\lambda, |[\mathbf{A}]_{ij}|\}} \right). \quad (7)$$

APPENDIX B SOLVING FOR $\hat{\mathbf{U}}_n$

This refers to the convex minimization sub-task (4a) and provides important details essential to the implementation of Alg. 2. All derivations, including those in Sec. C, are performed on the basis of viewing the complex-valued $\mathbf{U}$ as $(\Re(\mathbf{U}), \Im(\mathbf{U}))$, where $\Re(\cdot)$ and $\Im(\cdot)$ stand for the real and imaginary parts, respectively, of a complex-valued matrix. Gradients are *not* considered in the complex-differentiability sense (Cauchy-Riemann conditions) [60]. Nevertheless, to save space and use compact mathematical expressions, all subsequent results, including those in Sec. C, are stated in their complex-valued form.

Upon defining the convex constraint $\mathcal{C}_i := \{\mathbf{U} \mid \|\mathbf{U} \mathbf{e}_i\| \leq C_U\}$, $\forall i \in \{1, \dots, d\}$, (4a) can be expressed as

$$\hat{\mathbf{U}}_n \in \arg \min_{\mathbf{U}} g_1(\mathbf{U}) + g_2(\mathbf{U}), \quad (8)$$

where

$$g_1(\mathbf{U}) := \frac{1}{2} \|\mathcal{S}(\mathbf{Y}) - \mathcal{S}\mathcal{F}(\mathbf{U} \check{\mathbf{A}} \mathbf{B}_n)\|_F^2 + \frac{\tau_U}{2} \|\mathbf{U} - \mathbf{U}_n\|_F^2 + \frac{\lambda_1}{2} \|\mathbf{Z}_n - \mathcal{F}_t(\mathbf{U} \check{\mathbf{A}} \mathbf{B}_n)\|_F^2, \quad (9)$$

$$g_2(\mathbf{U}) := \sum_{i=1}^d \iota_{\mathcal{C}_i}(\mathbf{U}). \quad (10)$$

It can be verified that $\forall \mathbf{U}$,

$$\begin{aligned} \nabla g_1(\mathbf{U}) &= [N_k \mathcal{F}^{-1} \mathcal{S}\mathcal{F}(\mathbf{U} \check{\mathbf{A}} \mathbf{B}_n) + \lambda_1 N_{fr} \mathbf{U} \check{\mathbf{A}} \mathbf{B}_n] \mathbf{B}_n^H \check{\mathbf{A}}^H \\ &\quad + \tau_U (\mathbf{U} - \mathbf{U}_n) \\ &\quad - [N_k \mathcal{F}^{-1} \mathcal{S}(\mathbf{Y}) + \lambda_1 N_{fr} \mathcal{F}_t^{-1}(\mathbf{Z}_n)] \mathbf{B}_n^H \check{\mathbf{A}}^H. \end{aligned} \quad (11)$$

By virtue of the fact $\|\mathcal{F}^{-1} \mathcal{S}\mathcal{F}(\mathbf{U})\|_F \leq \|\mathbf{U}\|_F$, $\forall \mathbf{U}$, it can be also verified that $\forall \mathbf{U}_1, \mathbf{U}_2$,

$$\begin{aligned} &\|\nabla g_1(\mathbf{U}_1) - \nabla g_1(\mathbf{U}_2)\|_F \\ &\leq \|(N_k + \lambda_1 N_{fr}) \check{\mathbf{A}} \mathbf{B}_n \mathbf{B}_n^H \check{\mathbf{A}}^H + \tau_U \mathbf{I}_d\|_2 \cdot \|\mathbf{U}_1 - \mathbf{U}_2\|_F, \end{aligned} \quad (12)$$

which yields the Lipschitz coefficient

$$L := (N_k + \lambda_1 N_{fr}) \cdot \lambda_{\max}(\check{\mathbf{A}} \mathbf{B}_n \mathbf{B}_n^H \check{\mathbf{A}}^H) + \tau_U. \quad (13)$$

The i th column of $\text{Prox}_{\lambda g_2}(\mathbf{U})$, $\forall \lambda > 0$ and $\forall \mathbf{U}$, is computed by the (metric) projection mapping onto $\mathcal{C}_i$:

$$\text{Prox}_{\lambda g_2}(\mathbf{U}) \mathbf{e}_i = \frac{C_U}{\max\{C_U, \|\mathbf{U} \mathbf{e}_i\|\}} \mathbf{U} \mathbf{e}_i. \quad (14)$$

Lastly, since (8) is not affinely constrained, [46] suggests

$$T(\mathbf{U}) = \mathbf{U}, \quad \forall \mathbf{U}, \quad (15)$$

for Alg. 2.

APPENDIX C SOLVING FOR $\hat{\mathbf{B}}_n$

This refers to the convex minimization sub-task (4b) and provides important details, essential to the implementation of Alg. 2. Upon defining the affine constraint $\mathcal{C}_{\text{aff}} := \{\mathbf{B} \mid \mathbf{1}_{N_\ell}^T \mathbf{B} = \mathbf{1}_{N_{fr}}^T\}$, (4b) can be expressed as

$$\hat{\mathbf{B}}_n \in \arg \min_{\mathbf{B} \in \mathcal{C}_{\text{aff}}} g_1(\mathbf{B}) + g_2(\mathbf{B}), \quad (16)$$

where

$$g_1(\mathbf{B}) = \frac{1}{2} \|\mathcal{S}(\mathbf{Y}) - \mathcal{S}\mathcal{F}(\mathbf{U}_n \check{\mathbf{A}} \mathbf{B})\|_F^2 + \frac{\tau_B}{2} \|\mathbf{B} - \mathbf{B}_n\|_F^2 + \frac{\lambda_1}{2} \|\mathbf{Z}_n - \mathcal{F}_t(\mathbf{U}_n \check{\mathbf{A}} \mathbf{B})\|_F^2, \quad (17)$$

$$g_2(\mathbf{B}) = \lambda_3 \|\mathbf{B}\|_1. \quad (18)$$

It can be verified that $\forall \mathbf{B}$,

$$\begin{aligned} \nabla g_1(\mathbf{B}) &= \check{\mathbf{A}}^H \mathbf{U}_n^H [N_k \mathcal{F}^{-1} \mathcal{S}\mathcal{F}(\mathbf{U}_n \check{\mathbf{A}} \mathbf{B}) + \lambda_1 N_{fr} \mathbf{U}_n \check{\mathbf{A}} \mathbf{B}] \\ &\quad + \tau_B (\mathbf{B} - \mathbf{B}_n) \\ &\quad - \check{\mathbf{A}}^H \mathbf{U}_n^H [N_k \mathcal{F}^{-1} \mathcal{S}(\mathbf{Y}) + \lambda_1 N_{fr} \mathcal{F}_t^{-1}(\mathbf{Z}_n)], \end{aligned} \quad (19)$$

and by steps similar to those in Sec. B, $\nabla g_1(\cdot)$ is Lipschitz continuous with coefficient

$$L = (N_k + \lambda_1 N_{fr}) \cdot \lambda_{\max}(\check{\mathbf{A}}^H \mathbf{U}_n^H \mathbf{U}_n \check{\mathbf{A}}) + \tau_B. \quad (20)$$

According to (7), $\text{Prox}_{\lambda g_2}(\cdot)$ becomes: $\forall \mathbf{B}$,

$$[\text{Prox}_{\lambda g_2}]_{ij} = [\mathbf{B}]_{ij} \left(1 - \frac{\lambda \lambda_3}{\max\{\lambda \lambda_3, |[\mathbf{B}]_{ij}|\}} \right). \quad (21)$$

Moreover, according to [46], mapping $T(\cdot)$, used in Alg. 2, is set equal to the projection mapping onto $\mathcal{C}_{\text{aff}}$: $\forall \mathbf{B}$,

$$T(\mathbf{B}) = \mathbf{B} - \frac{1}{N_\ell} \mathbf{1}_{N_\ell} (\mathbf{1}_{N_\ell}^T \mathbf{B} - \mathbf{1}_{N_{fr}}^T). \quad (22)$$

REFERENCES

- [1] Z.-P. Liang and P. C. Lauterbur, *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. IEEE Press, 2000.
- [2] D. P. Petersen and D. Middleton, "Sampling and reconstruction of wave-number-limited functions in N-dimensional Euclidean spaces," *Information and Control*, vol. 5, no. 4, pp. 279–323, 1962.

- [3] Z.-P. Liang and P. C. Lauterbur, "An efficient method for dynamic magnetic resonance imaging," *IEEE Trans. Medical Imag.*, vol. 13, no. 4, pp. 677–686, 1994.
- [4] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang, "Accelerating magnetic resonance imaging via deep learning," in *Proc. ISBI*, 2016, pp. 514–517.
- [5] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang, "Low-dose CT with a residual encoder-decoder convolutional neural network," *IEEE Trans. Medical Imag.*, vol. 36, no. 12, pp. 2524–2535, 2017.
- [6] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, p. 487, 2018.
- [7] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [8] H. K. Aggarwal, M. P. Mani, and M. Jacob, "Model based image reconstruction using deep learned priors (MODL)," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 2018, pp. 671–674.
- [9] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans. Medical Imag.*, vol. 37, no. 2, pp. 491–503, 2018.
- [10] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.
- [11] M. Mardani, E. Gong, J. Y. Cheng, J. Pauly, and L. Xing, "Recurrent generative adversarial networks for compressive imaging," in *Proc. IEEE CAMSAP*, 2017, pp. 1–5.
- [12] M. Mardani, E. Gong, J. Y. Cheng, S. Vasanaawala, G. Zaharchuk, M. Alley, N. Thakur, S. Han, W. Dally, J. M. Pauly *et al.*, "Deep generative adversarial networks for compressed sensing automates MRI," *arXiv e-print*, 2017, 1706.00051.
- [13] M. Lustig, J. M. Santos, D. L. Donoho, and J. M. Pauly, "k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *Proc. ISMRM*, vol. 2420, 2006.
- [14] H. Jung, J. C. Ye, and E. Y. Kim, "Improved k-t BLAST and k-t SENSE using FOCUS," *Physics in Medicine and Biology*, vol. 52, no. 11, pp. 3201–3226, 2007.
- [15] R. Otazo, D. Kim, L. Axel, and D. K. Sodickson, "Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI," *Magnetic Resonance in Medicine*, vol. 64, no. 3, pp. 767–776, 2010.
- [16] D. Liang, E. V. R. DiBella, R.-R. Chen, and L. Ying, "k-t ISD: Dynamic cardiac MR imaging using compressed sensing with iterative support detection," *Magnetic Resonance in Medicine*, vol. 68, no. 1, pp. 41–53, 2012.
- [17] S. G. Lingala, Y. Hu, E. V. R. DiBella, and M. Jacob, "Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR," *IEEE Trans. Medical Imag.*, vol. 30, no. 5, pp. 1042–1054, 2011.
- [18] B. Zhao, J. P. Haldar, A. G. Christodoulou, and Z.-P. Liang, "Image reconstruction from highly undersampled (k,t)-space data with joint partial separability and sparsity constraints," *IEEE Trans. Medical Imag.*, vol. 31, no. 9, pp. 1809–1820, 2012.
- [19] K. T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils: Iterative image reconstruction using a total variation constraint," *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1086–1098, 2007.
- [20] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (TGV) for MRI," *Magnetic Resonance in Medicine*, vol. 65, no. 2, pp. 480–491, 2011.
- [21] L. Feng, R. Grimm, K. T. Block, H. Chandarana, S. Kim, J. Xu, L. Axel, D. K. Sodickson, and R. Otazo, "Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI," *Magnetic Resonance in Medicine*, vol. 72, no. 3, pp. 707–717, 2014.
- [22] S. Poddar and M. Jacob, "Dynamic MRI using smoothness regularization on manifolds (SToRM)," *IEEE Trans. Medical Imag.*, vol. 35, no. 4, pp. 1106–1115, 2016.
- [23] S. Ravishanker, B. E. Moore, R. R. Nadakuditi, and J. A. Fessler, "Low-rank and adaptive sparse signal (LASSI) models for highly accelerated dynamic imaging," *IEEE Trans. Medical Imag.*, vol. 36, no. 5, pp. 1116–1128, 2017.
- [24] L. Feng, L. Axel, H. Chandarana, K. T. Block, D. K. Sodickson, and R. Otazo, "XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing," *Magnetic resonance in medicine*, vol. 75, no. 2, pp. 775–788, 2016.
- [25] S. P. Awate and E. V. R. DiBella, "Spatiotemporal dictionary learning for undersampled dynamic MRI reconstruction via joint frame-based and dictionary-based sparsity," in *Proc. ISBI*, 2012, pp. 318–321.
- [26] Y. Wang and L. Ying, "Compressed sensing dynamic cardiac cine MRI using learned spatiotemporal dictionary," *IEEE Trans. Biomed. Engineering*, vol. 61, no. 4, pp. 1109–1120, 2014.
- [27] J. Caballero, A. N. Price, D. Rueckert, and J. V. Hajnal, "Dictionary learning and time sparsity for dynamic MR data reconstruction," *IEEE Trans. Medical Imag.*, vol. 33, no. 4, pp. 979–994, 2014.
- [28] U. Nakarmi, Y. Zhou, J. Lyu, K. Slavakis, and L. Ying, "Accelerating dynamic magnetic resonance imaging by nonlinear sparse coding," in *Proc. ISBI*, 2016.
- [29] Y. Wang, N. Cao, Z. Liu, and Y. Zhang, "Real-time dynamic MRI using parallel dictionary learning and dynamic total variation," *Neurocomputing*, vol. 238, pp. 410–419, 2017.
- [30] U. Nakarmi, K. Slavakis, J. Lyu, and L. Ying, "M-MRI: A manifold-based framework to highly accelerated dynamic magnetic resonance imaging," in *Proc. ISBI*, 2017, pp. 19–22.
- [31] S. Poddar, Y. Mohsin, D. Ansah, B. Thattaliyath, R. Ashwath, and M. Jacob, "Free-breathing cardiac MRI using bandlimited manifold modelling," *arXiv preprint arXiv:1802.08909*, 2018.
- [32] S. Poddar and M. Jacob, "Recovery of noisy points on bandlimited surfaces: Kernel methods re-explained," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4024–4028.
- [33] M. Usman, D. Atkinson, C. Kolbitsch, T. Schaeffter, and C. Prieto, "Manifold learning based ECG-free free-breathing cardiac CINE MRI," *J. Magnetic Resonance Imag.*, vol. 41, no. 6, pp. 1521–1527, 2015.
- [34] U. Nakarmi, K. Slavakis, and L. Ying, "MLS: Joint manifold-learning and sparsity-aware framework for highly accelerated dynamic magnetic resonance imaging," in *Proc. ISBI*, 2018, pp. 1213–1216.
- [35] A. H. Ahmed, Y. Mohsin, R. Zhou, Y. Yang, M. Salerno, P. Nagpal, and M. Jacob, "Free-breathing and ungated cardiac cine using navigator-less spiral SToRM," *arXiv preprint arXiv:1901.05542*, 2019.
- [36] Y. Shen, P. A. Traganitis, and G. B. Giannakis, "Nonlinear dimensionality reduction on graphs," in *Proc. IEEE CAMSAP*, 2017.
- [37] J. Silva, J. Marques, and J. Lemos, "Selecting landmark points for sparse manifold learning," in *Proc. NIPS*, 2006, pp. 1241–1248.
- [38] Y. Chen, M. Crawford, and J. Ghosh, "Improved nonlinear manifold learning for land cover classification via intelligent landmark selection," in *Proc. IEEE IGARSS*, 2006, pp. 545–548.
- [39] V. De Silva and J. B. Tenenbaum, "Sparse multidimensional scaling using landmark points," Stanford University, Tech. Rep., 2004.
- [40] L. W. Tu, *An Introduction to Manifolds*. New York: Springer, 2008.
- [41] K. Slavakis, G. B. Giannakis, and G. Leus, "Robust sparse embedding and reconstruction via dictionary learning," in *Proc. CISS*, Baltimore: USA, Mar. 2013.
- [42] K. Slavakis, G. N. Shetty, A. Bose, U. Nakarmi, and L. Ying, "Bi-linear modeling of manifold-data geometry for dynamic-MRI recovery," in *Proc. IEEE CAMSAP*, 2017.
- [43] L. K. Saul and S. T. Roweis, "Think globally, fit locally: Unsupervised learning of low dimensional manifolds," *J. Machine Learning Research*, vol. 4, pp. 119–155, 2003.
- [44] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton University Press, 1970.
- [45] E. Elhamifar and R. Vidal, "Sparse manifold clustering and embedding," in *Proc. NIPS*, Granada: Spain, Dec. 2011.
- [46] K. Slavakis and I. Yamada, "Fejér-monotone hybrid steepest descent method for affinely constrained and composite convex minimization tasks," *Optimization*, 2018, DOI: 10.1080/02331934.2018.1505885.
- [47] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning*, 2nd ed. Springer, 2009.
- [48] F. Facchinei, G. Scutari, and S. Sagratella, "Parallel selective algorithms for nonconvex big data optimization," *IEEE Trans. Signal Process.*, vol. 63, no. 7, pp. 1874–1889, 2015.
- [49] E. J. Candès, C. A. Sing-Long, and J. D. Trzasko, "Unbiased risk estimates for singular value thresholding and spectral estimators," *IEEE Trans. Signal Process.*, vol. 61, no. 19, pp. 4643–4657, 2013.
- [50] L. Wissmann, C. Santelli, W. P. Segars, and S. Kozierke, "MRXCAT: Realistic numerical phantoms for cardiovascular magnetic resonance," *Journal of Cardiovascular Magnetic Resonance*, vol. 16, no. 1, p. 63, 2014.
- [51] "MATLAB (R2018b)," The Mathworks, Inc., Natick, Massachusetts, 2018.

- [52] S. Poddar and M. Jacob, “SToRM MATLAB Package,” accessed in 2018. [Online]. Available: <https://github.com/sunrita-poddar/12SToRM>
- [53] S. Ravishankar, B. E. Moore, R. R. Nadakuditi, and J. A. Fessler, “LASSI MATLAB Package,” accessed in 2018. [Online]. Available: <https://web.eecs.umich.edu/~fessler/irt/reproduce/17/ravishankar-17-lra/>
- [54] L. Feng, L. Axel, H. Chandarana, K. T. Block, D. K. Sodickson, and R. Otazo, “XD-GRASP MATLAB Package,” accessed in 2019. [Online]. Available: <https://cai2r.net/resources/software/xd-grasp-matlab-code>
- [55] S. Ravishankar and Y. 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.
- [56] M. Subbarao, T.-S. Choi, and A. Nikzad, “Focusing techniques,” *Optical Engineering*, vol. 32, no. 11, pp. 2824–2837, 1993.
- [57] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli *et al.*, “Image quality assessment: From error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [58] Y.-C. Tsai, H.-D. Lin, Y.-C. Hu, C.-L. Yu, and K.-P. Lin, “Thin-plate spline technique for medical image deformation,” *J. Medical and Biological Eng.*, vol. 20, no. 4, pp. 203–210, 2000.
- [59] A. Maleki, L. Anitori, Z. Yang, and R. G. Baraniuk, “Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP),” *IEEE Trans. Information Theory*, vol. 59, no. 7, pp. 4290–4308, 2013.
- [60] S. Lang, *Complex Analysis*, 4th ed. New York: Springer, 1999.