Dynamic MRI using SmooThness Regularization on Manifolds (SToRM)

Sunrita Poddar, Student Member, IEEE, Mathews Jacob, Senior Member, IEEE

Abstract—We introduce a novel algorithm to recover real time dynamic MR images from highly under-sampled k-t space measurements. The proposed scheme models the images in the dynamic dataset as points on a smooth, low dimensional manifold in high dimensional space. We propose to exploit the non-linear and non-local redundancies in the dataset by posing its recovery as a manifold smoothness regularized optimization problem. A navigator acquisition scheme is used to determine the structure of the manifold, or equivalently the associated graph Laplacian matrix. The estimated Laplacian matrix is used to recover the dataset from undersampled measurements. The utility of the proposed scheme is demonstrated by comparisons with state of the art methods in multi-slice real-time cardiac and speech imaging applications.

I. INTRODUCTION

Dynamic MR imaging plays a central role in several applications such as structural and functional imaging of the heart, lung and liver, as well as vocal tract imaging in speech. While breath-held and ECG gated imaging is the default acquisition strategy in cardiac MRI, free-breathing un-gated acquisitions can enable the imaging of patients that have difficulty holding their breath [1] (e.g. COPD, obese, and paediatric subjects). Such free running sequences, where the acquisitions are not triggered by physiological signals, can also offer higher acquisition efficiency. The main challenge with free-breathing and ungated strategies (often termed as real-time (RT) imaging), is the slow nature of MR acquisition, which severely restricts the achievable spatial and temporal resolution.

Several model-based reconstruction algorithms that recover dynamic data from undersampled measurements have been introduced to improve the spatial and temporal resolution. The popular approaches include k-t SPARSE methods [2], [3], total variation (TV) regularization [4], and low rank methods such as k-t PCA [5] or partially separable functions (PSF) [6], [7]. k-t SPARSE methods model the intensity profiles as a sparse linear combination of exponentials. Temporal TV regularization relies on the similarity of each frame with its neighbours in time. PSF and k-t PCA methods exploit the linear dependencies between the intensity profiles by modelling them as a linear combination of basis functions, which are estimated from navigator signals. The main drawback of these schemes in the context of real-time MRI is the degradation in performance with extensive inter-frame motion.

This work is supported by NSF CCF-1116067, ACS RSG-11-267-01-CCE, and ONR grant N00014-13-1-0202.

We introduce a novel RT image acquisition and reconstruction method termed SToRM (SmooThness Regularization on Manifolds), which exploits the non-linear and non-local dependencies between images in the time series. In many RT applications, each image frame in the dataset is a nonlinear function of a few physiological parameters (e.g. cardiac and respiratory phase in real-time cardiac cine). Thus the image frames can be modelled as points on a smooth and low dimensional non-linear manifold. Unlike motion resolved reconstruction strategies that bin the data to a few cardiac and respiratory phases and recover them, we propose to recover the entire dynamic dataset from the undersampled k-t data as a manifold smoothness regularized reconstruction problem. The proposed approach is inspired by the manifold regularization schemes that are widely used in machine learning applications [8], [9], [10]. This strategy requires the knowledge of the manifold structure, or equivalently the associated graph Laplacian operator. We introduce a navigator acquisition scheme to estimate the graph Laplacian matrix. We consider both ℓ_2 and ℓ_1 regularization penalties. We show that the ℓ_2 -SToRM formulation can be solved analytically in the Fourier domain in the single receiver coil setting, while it can be solved efficiently using a simple conjugate gradients algorithm in the multi-channel case. We introduce a variable splitting based algorithm to solve for the ℓ_1 -SToRM formulation. We demonstrate the utility of our method in accelerated cardiac and speech imaging. The comparisons of the proposed method with the state of the art methods show improved image quality. We expect that our proposed scheme can also be used to accelerate other MR imaging applications such as lung, bowel and liver imaging. The early version of the proposed algorithm was presented in the conference paper [11], and later extended in [12].

II. BACKGROUND

A. Acquisition scheme

We model the raw dynamic multi-channel MRI data from the i^{th} image frame \mathbf{x}_i as:

$$\mathbf{b}_{ij} = \underbrace{\mathbf{S}_i \mathbf{F} \mathbf{C}_j}_{\mathbf{A}_{ij}} \mathbf{x}_i + \boldsymbol{\eta}_{ij}, \quad j = 1, .., \mathbf{N}_{\text{coils}}$$
(1)

where C_j is the receive sensitivity of the j^{th} coil, S_i is the sampling pattern for the i^{th} frame and η_{ij} is the noise. **F** is the discrete Fourier transform matrix. The above can be simplified and re-written as:

$$\mathbf{B} = \mathcal{A}(\mathbf{X}) + \boldsymbol{\eta} \tag{2}$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_k]$ is the Casorati matrix obtained by stacking the vectorized images as columns.

S.Poddar and M.Jacob are with the Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, Iowa, 52242, USA. e-mail: sunrita-poddar, mathews-jacob@uiowa.edu.

We model the images in the dynamic dataset as points on a smooth low-dimensional image manifold i.e. $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k \in \mathcal{M} \subseteq \mathbb{R}^N$. Here \mathcal{M} is a smooth *m*-dimensional manifold $(m \ll N)$ and N specifies the number of pixels in an image. The regularized recovery of continuous multi-dimensional functions of a manifold has received considerable attention in the context of machine learning [9]. The problem is formulated as:

$$\hat{f} = \arg\min_{f} \mathcal{V}(f) + \lambda \int_{\mathcal{M}} \|\nabla_{\mathcal{M}} f\|^2 d\mathbf{x}$$
(3)

where f is the continuous function, \mathcal{V} is the desired loss function and $\nabla_{\mathcal{M}} f$ is the derivative of f on \mathcal{M} . The second term contains the roughness prior on the manifold which can also be expressed as:

$$\int_{\mathcal{M}} \|\nabla_{\mathcal{M}} f\|^2 d\mathbf{x} = \langle \nabla_{\mathcal{M}} f, \nabla_{\mathcal{M}} f \rangle = \langle f, \triangle_{\mathcal{M}} f \rangle$$
$$= \int_{\mathcal{M}} f \ \triangle_{\mathcal{M}} f \ d\mathbf{x} \qquad (4)$$

where $\triangle_{\mathcal{M}}$ is the Laplace-Beltrami operator on the manifold. When one is only interested in recovering discrete function values specified by $\mathbf{f} = f_1, f_2, \dots, f_k$ at points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$, the common practice is to approximate the problem as [9]:

$$\hat{\mathbf{f}} = \arg\min_{\mathbf{f}} \, \mathcal{V}(\mathbf{f}) + \lambda \, \sum_{i} \sum_{j} w_{ij} \, \|f_i - f_j\|^2 \qquad (5)$$

where the weights w_{ij} are specified by:

$$w_{ij} = \mathbf{e}^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma^2}} \tag{6}$$

Note that the weights decay with distance. Specifically, w_{ij} will assume a high value if \mathbf{x}_i and \mathbf{x}_j are similar to each other, and a small value if they are different. The penalty term can also be expressed as:

$$\sum_{i} \sum_{j} w_{ij} \|f_i - f_j\|^2 = 2 \operatorname{Tr} \left(\mathbf{f} \mathbf{L} \mathbf{f}^H \right)$$
(7)

where Tr denotes the trace operator and \mathbf{L} is the graph Laplacian operator. The \mathbf{L} matrix is related to the weight matrix \mathbf{W} (with entries defined by (6)) as:

$$\mathbf{L} = \mathbf{D} - \mathbf{W} \tag{8}$$

where **D** is a diagonal matrix with entries $\mathbf{D}(i, i) = \sum_{j} w_{ij}$. For example, in a three node graph, the Laplacian is:

$$\mathbf{L} = \begin{bmatrix} w_{12} + w_{13} & -w_{12} & -w_{13} \\ -w_{12} & w_{12} + w_{23} & -w_{23} \\ -w_{13} & -w_{23} & w_{13} + w_{23} \end{bmatrix}$$
(9)

Note the similarity between the discrete approximation (7) and (4). When the manifold is uniformly sampled, the discrete graph Laplacian operator converges to the Laplace Beltrami operator on the manifold in the limit (as the distance between samples tend to zero) [13]. When $\mathcal{M} = \mathbb{R}^m$, then L is exactly the finite difference discretization of the continuous Laplacian operator on a regular lattice (up to a constant factor)[10]:

$$\Delta_{\mathcal{M}} f(\mathbf{r}) = \sum_{i=1}^{m} \frac{f(\mathbf{r} + \mathbf{e}_i) + f(\mathbf{r} - \mathbf{e}_i) - 2f(\mathbf{r})}{\delta^2} = -\frac{[\mathbf{L}\mathbf{f}]_{(\mathbf{r})}}{\delta^2}$$
(10)

where $\mathbf{e}_1, \ldots \mathbf{e}_m$ form an orthogonal basis for \mathbb{R}^m with $\|\mathbf{e}_i\| = \delta$.

III. PROPOSED SCHEME

We model the images in the dynamic dataset as points on a smooth low-dimensional manifold parameterized by a few variables. For example, the images in a free-breathing and ungated cardiac MRI dataset are non-linear functions of their cardiac and respiratory phases. The proposed framework is also general enough to be applied to other dynamic imaging applications like imaging of the vocal tract in speech, where there is no concept of phases equivalent to cardiac and respiratory phases in cardiac imaging. We propose to recover the dynamic dataset from its undersampled measurements (1) by exploiting the manifold structure of the data. Motivated by (5), we pose the recovery as:

$$\{\mathbf{X}^*\} = \arg\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + \lambda \sum_i \sum_j \left(\sqrt{w_{ij}} \|\mathbf{x}_i - \mathbf{x}_j\|_p\right)^p$$
(11)

where we use the ℓ_p $(p \le 2)$ norm of the image differences in the regularizer. We will consider the special cases p = 2 and p = 1 in the later subsections. The above optimization problem promotes solutions where each image frame is similar in the lpnorm sense to its neighbours on the manifold and the degree of similarity is determined by the weights w_{ij} .

In classical manifold embedding applications, the weights are derived from the images themselves. This approach is not practical in our setting since we only have a few measurements available from each frame. We hence propose an acquisition strategy using navigators to estimate the weights to be used in (11). This approach is similar to [6], [7].

A. Estimation of Manifold Structure from Navigators

Consider that each of the k images is observed by the same $M \times N$ matrix Ψ (M < N). This mapping is a stable embedding if the distance between any two points \mathbf{x}_i and \mathbf{x}_j is preserved after the mapping $\Psi \mathbf{x}_i$. Wakin et al [14] have shown that a random orthoprojector Ψ provides a stable embedding of the manifold. Specifically, for some $0 < \epsilon \leq \frac{1}{3}$ and a sufficient number of measurements M, the following holds with high probability for every pair \mathbf{x}_i and \mathbf{x}_j :

$$(1-\epsilon)\|\mathbf{x}_i - \mathbf{x}_j\| \le \|\mathbf{\Psi}\mathbf{x}_i - \mathbf{\Psi}\mathbf{x}_j\| \le (1+\epsilon)\|\mathbf{x}_i - \mathbf{x}_j\|$$
(12)

The number of measurements M required to yield stable embedding is independent of the ambient dimension N and is almost linearly proportional to the dimension of the manifold m. The required number of measurements also depends on the characteristics of the manifold which are captured by its condition number and volume [15].

Motivated by the above result, we propose to sample the same k-space locations in every temporal frame. We term the common measurements as navigator acquisitions, which are often used in many dynamic MRI applications for calibration [6], [7]. We define the measurement operator A_{ij} corresponding to the i^{th} frame and the j^{th} coil as (see (1)):

$$\mathbf{b}_{i,j} = \underbrace{\begin{bmatrix} \mathbf{\Phi} \\ \mathbf{B}_i \end{bmatrix} \mathbf{F} \mathbf{C}_j}_{\mathbf{A}_{ij}} \mathbf{x}_i + \boldsymbol{\eta}_{ij}$$
(13)

The first operator Φ samples the same k-space locations every frame, regardless of the frame number *i*; the corresponding samples (termed navigator signals) enable the estimation of the neighbours of each frame. The second operator \mathbf{B}_i which samples different k-space locations every frame aids the image recovery algorithm by sampling the neighbours of a particular image frame at complementary k-space locations. We propose to estimate the inter-image distances as:

$$d_{ij}^{2} = \sum_{l=1}^{N_{\text{coils}}} \|\mathbf{z}_{il} - \mathbf{z}_{jl}\|^{2}$$
(14)

where $\mathbf{z}_{i,l}$ are the navigator signals given by:

$$\mathbf{z}_{i,l} = \mathbf{\Phi} \mathbf{F} \ \mathbf{C}_l \ \mathbf{x}_i + \boldsymbol{\eta}_{il}, \ l = 1, .., \mathbf{N}_{\text{coils}}$$
(15)

We compute the weights as:

$$w_{ij} = \begin{cases} \mathbf{e}^{-\frac{d_{ij}^2}{\sigma^2}} & \text{, if } \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ are neighbours.} \\ 0 & \text{, otherwise.} \end{cases}$$
(16)

We set the neighbourhood of each frame to be a fixed number of nearest neighbours. For example, in order to retain the 5 nearest neighbours for each frame, the i^{th} and the j^{th} frames are considered to be neighbours if the i^{th} frame is among the 5 frames most similar to the j^{th} frame or the j^{th} frame is among the 5 frames most similar to the i^{th} frame.

B. Special case: ℓ_2 smoothness prior

When p = 2, the recovery using (11) simplifies to:

$$\mathbf{X}^* = \arg\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + 2\lambda \operatorname{Tr}(\mathbf{X}\mathbf{L}\mathbf{X}^H), \quad (17)$$

where the Laplacian matrix \mathbf{L} is obtained from the weights using (8). We refer to this implementation as ℓ_2 -SToRM. We can view (17) as an analysis formulation since the regularizer is based on the analysis of \mathbf{X} (specified by \mathbf{XQ}), where $\mathbf{L} = \mathbf{QQ}^H$. The problem (17) can be rewritten as:

$$\{\mathbf{X}^*\} = \arg\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + 2 \ \lambda \|\mathbf{X}\mathbf{Q}\|^2.$$
(18)

The $k \times k(k-1)/2$ matrix **Q** specifies a gradient operator. For example, in a 4 node graph, the matrix **Q** is specified by:

$$\mathbf{Q}^{T} = \begin{bmatrix} \sqrt{w_{12}} & -\sqrt{w_{12}} & 0 & 0\\ 0 & \sqrt{w_{23}} & -\sqrt{w_{23}} & 0\\ 0 & 0 & \sqrt{w_{34}} & -\sqrt{w_{34}}\\ \sqrt{w_{13}} & 0 & -\sqrt{w_{13}} & 0\\ 0 & \sqrt{w_{24}} & 0 & -\sqrt{w_{24}}\\ \sqrt{w_{14}} & 0 & 0 & -\sqrt{w_{14}} \end{bmatrix}$$
(19)

Note that this approach is very similar to Tikhonov temporal regularization, when the sparse matrix \mathbf{Q} is the temporal finite difference operator. The proposed scheme uses an operator that computes differences between the neighbours on the manifold,

rather than the temporal neighbours. Since the neighbours on the manifold are expected to be more similar than the ones in time, we expect to obtain better recovery.

We will now show that this formulation is also equivalent to a synthesis formulation by a simple change of variables. In addition to providing additional insights, this offers an approach to represent the data efficiently, while working with large datasets. The Laplacian matrix has a singular value decomposition specified by:

$$\mathbf{L} = \mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^H \tag{20}$$

The eigen vectors of the Laplacian matrix denoted by \mathbf{v}_i are ideally suited to represent smooth signals on the manifold. A simple special case worth discussing is when the graph has r disjoint clusters. In this case, spectral graph theory shows that \mathbf{L} will have r zero singular values. The corresponding r singular vectors \mathbf{V}_0 with an appropriate rotation matrix \mathbf{R} will yield a set of sparse temporal basis functions:

$$\mathbf{E}_0 = \mathbf{R} \mathbf{V}_0 \tag{21}$$

Each of the basis functions in \mathbf{E}_0 will assume a value of zero for frames that are not in a particular cluster, and a constant value for all the frames in the cluster. This property is exploited in spectral clustering. If the images in the cluster are the same, these temporal basis functions are sufficient to represent the signal. Note that this representation is strikingly different from principle component analysis used in k-t PCA or PSF methods [6], [7]. Unlike the global subspace model used in these methods, the proposed approach captures the geometry of the data on the manifold, enabled by the nonlinear mapping (16). By minimizing the cross talk between images in distinct cardiac/respiratory phases, it is expected to reduce temporal blurring.

In the general setting, one would need more basis functions to account for the variability of images within clusters/on the manifold. Substituting \mathbf{L} in the regularization penalty term in (17), we obtain:

$$\operatorname{Tr}\left(\underbrace{(\mathbf{X}\mathbf{V})}_{\mathbf{U}}\mathbf{\Sigma}(\mathbf{X}\mathbf{V})^{H}\right) = \sum_{i=1}^{k} \sigma_{i} \|\mathbf{u}_{i}\|^{2}, \qquad (22)$$

where $\mathbf{u}_i = \mathbf{X}\mathbf{v}_i$ is the projection of \mathbf{X} onto the i^{th} singular vector \mathbf{v}_i and σ_i is the i^{th} singular value of \mathbf{L} . Substituting for \mathbf{X} in terms of \mathbf{U} in (17), we obtain the equivalent synthesis formulation:

$$\mathbf{U}^* = \arg\min_{\mathbf{U}} \|\mathcal{A}(\mathbf{U}\mathbf{V}^H) - \mathbf{B}\|_F^2 + 2\lambda \sum_{i=1}^k \sigma_i \|\mathbf{u}_i\|^2 \quad (23)$$

Note that the above formulation is very similar to the k-t PCA or PSF [6] algorithms that are now widely used in dynamic MRI. The columns of U correspond to representative images, while the columns of V are the corresponding temporal basis functions.

C. Special case: ℓ_1 smoothness prior

We consider the ℓ_1 norm of the differences between neighbouring images on the manifold:

$$\{\mathbf{X}^*\} = \arg\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + 2\lambda \|\mathbf{X}\mathbf{Q}\|_1$$
(24)

We term this implementation ℓ_1 -SToRM. Note that the above approach simplifies to the popular temporal total variation formulation when:

$$w_{ij} = \begin{cases} 1, & \text{if } j = i+1, i-1. \\ 0, & \text{otherwise.} \end{cases}$$
(25)

We expect our method to achieve better reconstruction than temporal TV since it enforces the differences between the closest neighbours of a frame on the manifold to be sparse. These frames might not be the frames that are close to it temporally, especially in case of high motion between frames.

Considering that \mathbf{Q} has a singular value decomposition:

$$\mathbf{Q} = \mathbf{V} \mathbf{\Sigma}^{\frac{1}{2}} \mathbf{T}$$
(26)

we can also find the equivalent synthesis formulation for the ℓ_1 problem by a change of variable $\mathbf{X} = \mathbf{U}\mathbf{V}^H$:

$$\mathbf{U}^* = \arg\min_{\mathbf{U}} \|\mathcal{A}(\mathbf{U}\mathbf{V}^H) - \mathbf{B}\|_F^2 + 2\lambda \|\mathbf{U}\boldsymbol{\Sigma}^{\frac{1}{2}}\mathbf{T}\|_1, \quad (27)$$

Note that this approach has similarities to ℓ_1 regularized PSF regularization schemes [7], except that the ℓ_1 norm of $\mathbf{U}\boldsymbol{\Sigma}^{\frac{1}{2}}\mathbf{T}$ is penalized rather than that of U.

IV. IMPLEMENTATION

We consider separately the solutions for the 2 cases described in the previous section: p = 2 and p = 1.

A. ℓ_2 smoothness prior

In the single coil case, problem (17) has an analytical solution in the Fourier domain. We rewrite (17) in this special case as:

$$\hat{\mathbf{X}}^* = \arg\min_{\hat{\mathbf{X}}} \sum_{i} \|\mathbf{S}_i \hat{\mathbf{x}}_i - \mathbf{b}_i\|_F^2 + 2\lambda \operatorname{Tr}(\hat{\mathbf{X}} \mathbf{L} \hat{\mathbf{X}}^H) \quad (28)$$

where the columns of $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_k]$ are the Fourier coefficients of the images given by: $\hat{\mathbf{x}}_i = \mathbf{F}\mathbf{x}_i$. The key observation is that the above expression can be decoupled into several independent subproblems, each involving the recovery of a row of $\hat{\mathbf{X}}$. Let $\hat{\mathbf{X}}_{(j,\cdot)}$ denote the j^{th} row of $\hat{\mathbf{X}}$ and $\mathbf{B}_{(j,\cdot)}$ denote the vector of measurements corresponding to this row. Then, we can solve for $\hat{\mathbf{X}}_{(j,\cdot)}$ analytically as:

$$\hat{\mathbf{X}}_{(j,\cdot)} = (\mathbf{D}_j^H \mathbf{D}_j + 2\lambda \mathbf{L})^{-1} \mathbf{D}_j^H \mathbf{B}_{(j,\cdot)}$$
(29)

where \mathbf{D}_j is the sampling matrix corresponding to the j^{th} row. The solutions for the different rows of $\hat{\mathbf{X}}$ can be computed in parallel. This analytic approach can give us a significant speed-up over solving for the whole matrix $\hat{\mathbf{X}}$ using iterative algorithms such as conjugate gradient.

In the multi-channel setting, it is possible to solve for each coil using the above method and combine them using a sumof-squares strategy. Since this approach is suboptimal, we propose to directly solve (17) using the conjugate gradient algorithm (accounting for the coil sensitivities) to obtain a more accurate solution. The gradient of the cost function in (17) can be computed as: $2\mathcal{A}^H\mathcal{A}(\mathbf{X}) + 4\mathbf{X}\mathbf{L}$. The computation of $\mathcal{A}^H\mathcal{A}(\mathbf{X})$ can be broken down into blocks (each containing a few temporal frames of \mathbf{X}) and the blocks can be processed in parallel in order to reduce computational complexity.

B. ℓ_1 smoothness formulation

We rely on a variable splitting strategy using an auxiliary variable \mathbf{Z} to solve (24):

$$\{\mathbf{X}^*, \mathbf{Z}^*\} = \arg\min_{\mathbf{X}, \mathbf{Z}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + 2\lambda \|\mathbf{Z}\|_{\ell_1} + \beta \|\mathbf{X}\mathbf{Q} - \mathbf{Z}\|_F^2 \quad (30)$$

We solve the above problem by alternating between minimization with respect to the 2 variables:

$$\mathbf{X}^{(n)} = \arg\min_{\mathbf{X}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_{F}^{2} + \beta \|\mathbf{X}\mathbf{Q} - \mathbf{Z}^{(n-1)}\|_{F}^{2}$$
(31)
$$\mathbf{Z}^{(n)} = \arg\min_{\mathbf{A}} \|\mathbf{X}^{(n)}\mathbf{Q} - \mathbf{Z}\|_{F}^{2} + 2) \|\mathbf{Z}\|_{L^{\infty}}$$
(32)

We use a homotopy continuation strategy on the parameter β , where β is initialized to a very small value and then increased gradually to a very large value till the algorithm converges. As in the ℓ_2 case, (31) can be solved analytically in the Fourier domain for single coil data. For multi-coil data, we use the conjugate gradient algorithm. (32) can be solved using shrinkage. The matrix **Z** is large and storing it explicitly will result in huge memory demands. We observe that the evaluation of (31) only requires \mathbf{ZQ}^T , which is considerably smaller in dimension than **Z**. We perform in-place computation of the variable \mathbf{ZQ}^T and store it instead of **Z** to reduce the memory demand of the algorithm.

C. Acquisition Scheme

The acquisition scheme used follows from the discussion in section III-A. We used a set of uniformly spaced radial navigator acquisitions (corresponding to Φ), separated by $180^{\circ}/N_l$ degrees where N_l is the number of navigator lines per frame. The remaining k-space samples (corresponding to B_i) were acquired using a golden angle radial k-space trajectory, where each line was separated by an angle of 111.25° from the previous line. Thus, B_i varies from frame to frame. The acquisition and reconstruction pipeline is illustrated in Fig 1, where we consider the single coil setup for simplicity.

D. Datasets

We use a numerical cardiac phantom and a retrospectively undersampled speech dataset for quantitative comparisons. We also consider the recovery of prospectively undersampled realtime cardiac MRI data.

1) PINCAT phantom: A short axis view of the PINCAT phantom [16] heart with matrix size 128×128 and 500 frames was used for numerical simulations. The dataset has around 26 cardiac cycles and 5 respiration cycles.



Fig. 1: Summary of the proposed data acquisition and reconstruction scheme for the single coil case. The blue radial lines denote the navigators that sample the same k-space locations in every frame. The weight matrix is estimated from the k-space data acquired using these navigator lines as described in (16). The final images are recovered from the entire measurements by solving (11).

2) Speech Imaging: We use the MR dataset titled 'F1' in the USC-TIMIT database [17] to demonstrate our method. The raw k-space data for the images in the database was acquired using a spiral trajectory and this data was gridded to reconstruct the images. The reconstructed images have been made available in the dataset as a movie in the coilcombined form with matrix size 68×68 and frame-rate 23.18 frames/s. This corresponds to a temporal resolution of around 43 ms. The Fourier data corresponding to the first 6000 image frames was retrospectively undersampled using 9 golden angle radial lines and 1 spiral navigator per frame and used for our experiments.

3) Cardiac Imaging: A prospectively undersampled freebreathing ungated radial dataset was acquired using a SSFP sequence on a Siemens 3T TIM Trio scanner with a 18 channel cardiac array from a healthy volunteer who was asked to breathe normally. The scan parameters were TR/TE = 4.2/2.2ms, number of slices = 5, slice thickness = 5 mm, FOV = 300 mm, spatial resolution = 1.17 mm. A temporal resolution of 42 ms was achieved by sampling 10 lines of k-space per frame, out of which 4 were navigator lines. 10000 radial lines of kspace were acquired per slice which resulted in an acquisition time of around 42 s per slice. For 5 slices this resulted in a total acquisition time of around 3.5 mins.

The raw k-space data was interpolated to a Cartesian grid and a SVD based coil-compression technique was used in order to create 4 virtual coil elements from the initial 18. We reconstructed low temporal resolution images for the original coils by binning k-space data from a large number of frames. We then performed an SVD on these images and retained only the 4 most significant singular vectors. The data from the original coils was coil-combined to form virtual coil data using the singular vectors obtained. This was done in order to reduce the computational complexity of the reconstruction procedure. The coil sensitivity maps were estimated from this compressed data using the method by Walsh et al [18]. To



Fig. 2: Illustration of the weight matrix and the ability of the scheme to enable implicit motion resolved recovery. (**a**,**b**) Two frames from the PINCAT dataset. (**c**) Weight matrix computed from the fully sampled k-space data. The green and blue lines show the rows corresponding to the frames in (a) and (b) respectively. The neighbours of these frames can be obtained using the weight matrix. (**d**) Temporal intensity profile corresponding to the cut shown by the red dotted line in (a). Frames (a) and (b) and a few of their neighbours are marked.

reduce computational complexity, the coil sensitivity maps were assumed to be constant over time.

4) Comparison between breath-held and free-breathing cardiac acquisitions: In order to compare the image quality obtained using our method to that obtained by a breath-held protocol, we acquired 2 cardiac datasets:

- A prospectively undersampled free-breathing ungated radial dataset.
- A fully-sampled breath-held ECG-gated radial dataset.

The first dataset was acquired using a SSFP sequence on a Siemens 3T TIM Trio scanner with a 5 channel cardiac array from a healthy volunteer who was asked to breathe normally. A TRUFI frequency scout was performed prior to data acquisition to prevent banding artifacts due to the presence of field in-homogeneity. The scan parameters were TR/TE = 3.2/1.62 ms, number of slices = 5, slice thickness = 5 mm, FOV = 300 mm, spatial resolution = 1.17 mm. A temporal resolution of 41.6 ms was achieved by sampling 13 lines of k-space per frame, out of which 4 were navigator lines. 13000 radial lines of k-space were acquired per slice which resulted in an acquisition time of around 42 s per slice. For 5 slices this resulted in a total acquisition time of around 3.3 mins.

The fully-sampled ECG-gated breath-held dataset was acquired by a SSFP sequence on the same subject immediately after the free-breathing scan. The sampling trajectory was uniform radial and the scan parameters were: TR/TE = 3.4/1.72



Fig. 3: Effect of different navigator trajectories on weight matrix estimation. (a) Percentage error in the weight matrix estimation (computed using ℓ_2 norm), using different navigator trajectories. Spiral and radial trajectories are chosen such that the time taken to acquire 1 spiral shot is the same as that for 1 radial line. (b) The 2nd, 3rd and 4th eigen vectors of the Laplacian matrix estimated from (1) fully sampled k-space, shown in blue (2) 1 radial spoke, shown in green (3) 1 spiral readout, shown in pink. We observe that these vectors capture the respiratory motion, the 2nd harmonic of the respiratory motion, and the cardiac motion modulated by the respiratory frequency respectively.

ms, number of slices = 5, slice thickness = 5 mm, number of channels = 5, FOV = 300 mm, spatial resolution = 1.17mm, number of cardiac phases = 18, radial views per cardiac phase = 253. Each slice required a breath-hold of around 16 s followed by a resting period of around 25 s. For 5 slices this resulted in a total acquisition time of around 3 mins.

Pre-interpolation to a Cartesian grid, coil sensitivity estimation and coil compression were performed using the acquired k-space data as described in the previous section. 3 virtual coils were created in this case.

E. State of the art methods used for comparison

The in vivo data reconstructed using ℓ_2 and ℓ_1 -SToRM was compared to the reconstructions by 3 other methods: kt-LR [19], temporal TV and PSF. The kt-LR and temporal TV methods do not require the acquisition of navigators. Thus, we did not include navigator lines in our sampling pattern for the speech data, for generating the results for these 2 methods. However, we could not do the same for the cardiac datasets since they were prospectively undersampled. For the PSF method, we used the Frobenius norm of the basis images as a regularizer. The approach followed was similar to [6], with the same weighting applied to all basis images. For all 3 competing methods, the regularization parameter giving the highest SER reconstruction was chosen in case of the speech dataset. For the cardiac dataset, since the ground-truth was not available, the regularization parameter which seemed to best preserve the features of the data was chosen. Spatial TV regularization was not used with any of the algorithms.

V. RESULTS

A. Simulations using PINCAT phantom data

We first conducted some numerical simulations on the PINCAT phantom. Two frames of the phantom dataset are shown in Fig 2.(a) and Fig 2.(b).

1) Weight matrix estimate from fully sampled data: We computed the weight matrices from the fully sampled k-space data, corresponding to different σ values. These matrices were thresholded to retain only the 5 nearest neighbours for each frame. The k-space data was then under-sampled (10 lines per frame sampled on a pseudo golden angle trajectory). Images were reconstructed from this under-sampled data using ℓ_2 -SToRM with the weight matrices corresponding to different σ values. The σ value giving the highest SER reconstruction was chosen to form the optimal weight matrix. This matrix is shown in Fig 2.(c). The temporal intensity profile of the original dataset (along the cut given by the red dotted line in Fig 2.(a)) is shown in Fig 2.(d). The frames in Fig 2.(a) and Fig 2.(b) and a few of their neighbours (obtained from the weight matrix) are marked along the profile. We observe that the frames in Fig 2.(a) and Fig 2.(b) are very similar to their neighbours estimated by the weight matrix.

2) Effect of navigator trajectory on weight matrix estimation: The effect of different navigator schemes on weight estimation is studied in Fig 3. The weights estimated from different trajectories were compared quantitatively to the groundtruth weights obtained from the fully sampled data (Section V-A1). The normalized ℓ_2 norm of the weight estimation error was used as the error metric. The optimal σ parameter varies from trajectory to trajectory, depending on the number of kspace points. We chose the best σ value in each case to obtain fair comparisons. We did not threshold the weight matrices for this experiment. We considered spiral and radial navigators with the same readout duration (TR = 4.3 ms). The percent errors in weight estimation (computed using ℓ_2 norm) are plotted in Fig 3.(a). We observe that 1 spiral shot (4.3 ms) is almost as accurate in estimating the weights as 5 radial lines (21.5 ms). The percent errors incurred in the two cases are 7.34% and 6.99% respectively. The 2^{nd} , 3^{rd} and 4^{th} eigenvectors of the L matrix estimated from: (1) the fully sampled data, (2) 1 radial line, and (3) 1 spiral shot are shown in Fig 3.(b). The experiments show that the eigen vectors in all three cases are quite similar. We also observe that the 2nd eigen-



Fig. 4: Effect of weight matrices estimated using different navigator trajectories on reconstruction. (a) Signal to error ratio of the reconstructions with the Laplacian matrix estimated from different navigator trajectories. The k-space samples used to reconstruct the images are the same for all cases (10 golden angle radial lines per frame). Only the navigator trajectory used to compute the weight matrix are varied. (b) A reconstructed frame is shown for a few of the trajectories reported in (a).

vector captures the respiratory motion of the data (5 respiratory cycles can be seen). The $3^{\rm rd}$ eigen-vector shows the $2^{\rm nd}$ harmonic of the respiratory motion. The dominant frequency of this eigen-vector is double that of the dominant frequency of the respiratory motion. The $4^{\rm th}$ eigen vector captures the cardiac motion modulated by the respiratory frequency (26 cardiac cycles can be seen).

3) Effect of weight matrix on image reconstruction:

The effect of different weight matrices (computed using the navigator trajectories described in Section V-A2) on image reconstruction quality is studied in Fig 4.(a). The phantom data was under-sampled in k-space using a golden angle trajectory with 10 lines per frame and this data was used for all reconstructions. The navigator data was used only for weight computation. The weight matrices were thresholded to retain only the 5 nearest neighbours for each frame. In Fig 4.(b), we show a single image frame from the time series, as reconstructed using different weight matrices. The weights computed using a 1 radial line navigator produced reconstructed images of comparable quality (31.87 dB) to the case of ground-truth weights (32.25 dB). The single shot spiral navigator trajectory, which takes the same acquisition time as 1 radial line, performed slightly better (32.18 dB) than the single radial line case. Estimation of weights using only the centre kspace signal gave very poor reconstructions (20.57 dB). Using a 3×3 patch around center k-space as the navigator signal (instead of the centre only) improved the results considerably (30.33 dB), though the error images show more artifacts than when using radial or spiral trajectories.

We clarify that for the above experiment we used the navigator data only for estimating the weights and not for reconstruction. However, the navigator data was used for reconstruction in all the subsequent in-vivo experiments on the speech and cardiac data. For the experiment in Sec V-A2, we were studying the relative merits of different sampling schemes on the weight computation. The analysis was extended in the above experiment, where we studied the effect of those computed weights on image reconstruction. If we included the navigator signals for the reconstruction step, then the quality of our reconstructed images would be dependent on: (1) The accuracy of the computed weights (2) The incoherence of the sampling patterns used for each trajectory. Since we were only studying effect (1), we used the same samples for reconstruction in each case.

B. Experiments on in vivo data

In the in vivo experiments, the parameter σ used for the calculation of the weight matrix W was automatically computed using the strategy described in [20]. For this purpose, we computed the weight matrix for a range of σ values and evaluated $l(\sigma) = \sum_{i=1}^{k} \sum_{j=1}^{k} W_{ij}(\sigma)$ for each weight matrix obtained. A log-log plot of $l(\sigma)$ revealed 2 constant asymptotes at $\sigma \to 0$ and $\sigma \to \infty$, smoothly connected by an approximately straight line. The approximate σ value at the middle of this linear portion was selected to form our weight matrix. The weight matrix was thresholded to retain only the 4-6 nearest neighbours for each frame. For the multi-slice cardiac datasets, we had to compute the weight matrix separately for each slice. This is because our acquisition scheme was 2D, i.e. the k-space samples from a particular slice were completely acquired before moving on to the next slice. The regularization parameter λ was chosen emperically. All reconstructions were done on a desktop computer (Intel Xeon E5-1620 CPU, 3.6 GHz, 32 GB RAM). The algorithms were implemented in MATLAB, and may be further optimized to produce lower reconstruction times.

1) Retrospective undersampling experiments on speech dataset: The speech dataset was reconstructed from undersampled k-space data using different techniques, as shown in Fig 5. The first row (a) shows the ground-truth images, while each subsequent row corresponds to datasets reconstructed by different methods. The techniques used for reconstruction from under-sampled k-space data along with their reconstruction times are: (b) kt-LR (4.8 hrs) (c) temporal TV (21 mins)



Fig. 5: Reconstruction of the speech dataset. (a) Ground-truth images. The subsequent rows correspond to reconstructions from under-sampled k-space data using (b) kt-LR, (c) temporal TV, (d) PSF, (e) ℓ_2 -SToRM, and (f) ℓ_1 -SToRM. The data used for (b) and (c) had a golden angle radial trajectory without navigators. The data used for (d), (e) and (f) had a spiral navigator. The arrows point out artefacts in the images reconstructed by the competing methods, which are not present in the images reconstructed by SToRM.

(d) PSF [6] (3 mins) (e) ℓ_2 -SToRM (7 mins) (f) ℓ_1 -SToRM (32 mins). For the kt-LR and temporal TV reconstructions, 10 golden angle radial lines of k-space were used per frame and no navigator lines were included. For the PSF and SToRM methods, k-space was undersampled using 9 golden angle radial lines and 1 spiral navigator per frame. SToRM produces reconstructions with higher SER than the other methods. Though the ℓ_1 and ℓ_2 -SToRM reconstructions have comparable SER, it can be seen than the ℓ_1 formulation reduces blurring and preserves borders better. The competing techniques have more artifacts compared to the proposed methods, as pointed out in the figure. The ability to recover high quality images from under-sampled data indicates that our method can be used to improve the temporal resolution and also acquire multiple slices in a shorter scan.

2) Recovery of prospectively undersampled RT cardiac dataset: The multi-slice free-breathing highly undersampled cardiac dataset described in Section IV-D3 was reconstructed using different methods, as illustrated in Fig 6. The techniques used for reconstruction in the different rows along with their reconstruction times are: (a) kt-LR (7.5 hrs) (b) temporal



Fig. 6: Reconstruction of the free-breathing cardiac dataset. Selected image frames and temporal intensity profiles along a vertical cut given by the red dotted line in (a) are shown. The images were reconstructed from under-sampled k-space data using (a) kt-LR, (b) temporal TV, (c) PSF, (d) ℓ_2 -SToRM, and (e) ℓ_1 -SToRM. The arrows point out artefacts in the images reconstructed by the competing methods, which are not present in the images reconstructed by SToRM.

TV (4.7 hrs) (c) PSF (4 mins) (d) ℓ_2 -SToRM (24 mins) (e) ℓ_1 -SToRM (4.9 hrs). The temporal intensity profile along a vertical cut of the image frames (given by the red dotted line in Fig 6.(a)) is also shown for each method. The comparisons are only qualitative since the ground truth dataset was not available. We observe that SToRM reduces streaking artifacts and spatial blurring, compared to other state of the art methods. Specifically, we observe that the myocardial borders are well captured, while details such as the papillary muscles are better defined. We also note that while the image frames of the ℓ_1 and ℓ_2 -SToRM reconstructions look similar, the temporal intensity profiles of the ℓ_1 formulation appear sharper.

3) Comparison between free-breathing and breath-held cardiac reconstructions: The quality of the reconstructed freebreathing and breath-held cardiac datasets described in Section IV-D4 are compared in Fig 7. The breath-held dataset was reconstructed using CG-SENSE [21], while the free breathing dataset was reconstructed using ℓ_2 -SToRM. We show the data corresponding to 2 out of the 5 reconstructed slices. The figure shows results from a particular slice of the breath-held dataset; it was difficult to find perfect matches between the breathheld and the free-breathing acquisitions. Fig 7 shows: (a) 3 cardiac phases from the breath-held cine reconstruction and temporal intensity profile along the yellow dotted line. (b) 3 frames from a single cardiac cycle of the free-breathing dataset



Fig. 7: Comparison between proposed free-breathing (FB) reconstruction and breath-held (BH) reconstruction. The BH dataset was reconstructed using CG-SENSE. The FB dataset was recovered using ℓ_2 -SToRM. Two matching slices from both datasets are shown. The rows represent different slices. (a) Images in different cardiac phases from the BH dataset. The voxel profiles along the yellow dotted line are also shown. (b) Image frames from a particular cardiac cycle of the FB dataset. The voxel profiles for a few cardiac cycles of the FB dataset are also shown (along the same cut as the BH dataset).

and temporal intensity profile along a vertical cut (same cut as the breath-held dataset). Note that the breath-held dataset has a few cardiac phases averaged over many cardiac cycles, while the free-breathing dataset consists of several cardiac cycles. Images from the cardiac cycle of the free-breathing reconstructions which best matched the breath-held images are shown here. We observe that the dataset reconstructed using SToRM is of comparable quality to the breath-held cine datasets.

VI. DISCUSSION

The proposed dynamic imaging scheme estimates the proximity of the images on the manifold using navigator signals, followed by a manifold aware recovery of the images from highly undersampled measurements. The reconstructed image quality was observed to be superior to that achieved by other state-of-the-art ungated reconstruction methods. Moreover, the experiment on the speech dataset demonstrated that SToRM can recover images in case of repeating frames, irrespective of whether the repetitions are periodic. In fact, the method does not distinguish between periodic and aperiodic changes. The quality of our reconstructed images is quite dependent on the degrees of freedom of the underlying physiological process. If the degrees of freedom is low, then every frame will have a sufficient number of neighbours very similar to it with high probability (provided that our acquisition time is long enough). If the degrees of freedom is high, then many frames may not have any other frames very similar to it, and the recovered frames will be of poor quality. However, in such situations, other model-based reconstruction schemes should also perform poorly due to lack of redundancy in the data.

While the original stable embedding theory deals with random ortho-projectors [14], our empirical comparisons in section V show that the radial k-space sampling scheme can estimate the neighbourhood of each image frame quite accurately. Moreover, our experiments also show that approximate estimates of the weight matrix (using one radial line of k-space) are often sufficient to ensure good recovery of images. Our experiments also reveal that spiral navigators are more efficient than radial navigators. We used the radial acquisition scheme for ease of implementation on the scanner. We will investigate the utility of spiral navigators in the future, which may translate to improved temporal resolution or reconstruction quality.

The proposed scheme has a few free parameters: (1) σ (2) the number of neighbours (3) λ . The optimal σ value is dependent on the k-space trajectory as well as the number of points. However, we observed that the reconstruction quality is not very sensitive to the exact value of σ . Specifically, changing σ by a factor of 10 does not significantly affect the reconstruction quality. The number of neighbours is a data dependent parameter determined by the degree of redundancy in the dataset. If a sufficient number of similar frames is available for each frame, then a small increase in the number of neighbours will not affect the image quality. However, if the number of neighbours is made very high, then all the neighbours of a particular frame will not be very similar to it, and the resulting reconstructed image will have motion blur. If the number of neighbours is made very low, then we will have aliasing artefacts. Similarly, the regularization parameter λ is also data dependent.

We show that ℓ_2 -SToRM has similarities to the k-t PCA

and PSF methods, with the exception that the temporal basis functions are the singular vectors of the Laplacian matrix rather than that of the covariance matrix. These basis functions promote smoother solutions on the manifold, enabling the exploitation of the non-linear dependencies between images. ℓ_1 -SToRM is similar to the temporal TV scheme, with the exception that the standard finite difference matrix is replaced by an adaptive finite difference operator; this enables the exploitation of non-local dependencies between images in the dataset. The ℓ_2 -SToRM scheme has similarities to the recent work [22]. Specifically, their solution is a clever approximation of our analytic solution in the ℓ_2 setting for the single channel case. Our approach also has conceptual similarities to [23], where the cardiac and respiratory phase information is recovered from the singular vectors of the graph Laplacian. This approach has been inspired by dimensionality reduction methods such as ISOMAP and LLE [24], [25] that are used to embed the data on a manifold to a lower dimensional subspace. [23] identifies the cardiac and respiratory phases from the dimensionality reduced data, followed by explicit motion-resolved binned reconstructions similar to [26]. In contrast, SToRM performs an implicit motion-resolved recovery of the entire RT dataset. In addition, SToRM does not need the explicit identification of individual phases, which is difficult in applications with both cardiac and respiratory motion and require additional pre-processing steps [23], [27]. The estimation of the cardiac and respiratory phases using band-pass filtering as in [27], [26] may be challenging in cases with irregular respiratory motion and arrhythmia. In addition, many applications like speech imaging have no concept of phase equivalent to cardiac and respiratory phases in cardiac imaging. SToRM extends readily to such applications. The proposed scheme also has conceptual similarities to recent kernel PCA based approaches, introduced to exploit nonlinear similarities between image patches. Specifically, [28] learns the basis functions using linear PCA on non-linearly transformed patches from low-resolution images. They then iterate between projecting each non-linearly transformed patch from the high-resolution images to this subspace, and solving for pre-images that satisfy data-consistency. This approach may be seen as a synthesis formulation of ℓ_2 -SToRM, when re-engineered for image patches.

VII. CONCLUSION

We introduced a novel acquisition scheme and reconstruction algorithm for real-time dynamic MR imaging, termed SToRM. The central assumption is that the images in the dynamic dataset are points on a smooth, low dimensional manifold embedded in high dimensional space. We formulated the recovery of the dataset from highly under-sampled measurements as a manifold smoothness regularized optimization problem. The neighbours of each image on the manifold were estimated from the navigator acquisition. SToRM was demonstrated to be useful in accelerating free breathing cardiac imaging and speech imaging, without compromising on image quality and slice coverage. This approach improves the spatio-temporal resolution, while ensuring patient comfort and reducing the total scan time. It can be easily extended to other dynamic imaging applications like liver, bowel and lung imaging.

VIII. ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their comments, as a result of which, the quality of the paper has improved significantly. We also thank Dr Krishna S. Nayak from the University of Southern California, whose insight helped us greatly during the cardiac data acquisition process.

REFERENCES

- S. B. Gay, C. L. Sistrom, C. A. Holder, and P. M. Suratt, "Breath-Holding Capability of Adults: Implications for Spiral Computed Tomography, Fast-Acquisition Magnetic Resonance Imaging, and Angiography," *Investigative Radiology*, vol. 29, 1994.
- [2] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye, "k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Magnetic Resonance in Medicine*, vol. 61, no. 1, pp. 103–116, 2009.
- [3] M. Lustig, J. Santos, D. Donoho, and J. Pauly, "k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *International Society on Magnetic Resonance in Medicine 2006.*
- [4] 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.
- [5] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke, and W. Y. Kim, "k-t PCA: Temporally constrained k-t BLAST reconstruction using principal component analysis," *Magnetic resonance in medicine*, vol. 62, no. 3, pp. 706–716, 2009.
- [6] A. G. Christodoulou, C. Brinegar, J. P. Haldar, H. Zhang, Y.-J. L. Wu, L. M. Foley, T. K. Hitchens, Q. Ye, C. Ho, and Z.-P. Liang, "Highresolution cardiac mri using partially separable functions and weighted spatial smoothness regularization," in *Annual International Conference* of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, 2010, pp. 871–874.
- [7] B. Zhao, J. P. Haldar, and Z.-P. Liang, "PSF Model-Based Reconstruction with Sparsity Constraint: Algorithm and Application to Real-Time Cardiac MRI," *Annual International Conference of the IEEE Engineering in Medicine and Biology Society.*, vol. 2010, pp. 3390– 3393, 2010.
- [8] L. J. van der Maaten, E. O. Postma, and H. J. van den Herik, "Dimensionality reduction: A comparative review," *Journal of Machine Learning Research*, vol. 10, no. 1-41, pp. 66–71, 2009.
- [9] M. Belkin, "Problems of learning on manifolds," Ph.D. dissertation, 2003.
- [10] A. J. Smola and R. Kondor, "Kernels and regularization on graphs," in Learning theory and kernel machines. Springer, 2003, pp. 144–158.
- [11] S. Poddar, S. G. Lingala, and M. Jacob, "Joint recovery of under sampled signals on a manifold: Application to free breathing cardiac MRI," in 2014 IEEE International Conference on Acoustics, speech and signal processing (ICASSP), 2014, pp. 6904–6908.
- [12] S. Poddar and M. Jacob, "Low rank recovery with manifold smoothness prior: theory and application to accelerated dynamic MRI," in *International Symposium on Biomedical Imaging (ISBI)* 2015.
- [13] A. Singer, "From graph to manifold Laplacian: The convergence rate," *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 128– 134, 2006.
- [14] A. Eftekhari and M. B. Wakin, "New analysis of manifold embeddings and signal recovery from compressive measurements," *Applied and Computational Harmonic Analysis*, vol. 39, no. 1, pp. 67–109, 2015.
- [15] P. Niyogi, S. Smale, and S. Weinberger, "Finding the homology of submanifolds with high confidence from random samples," *Discrete & Computational Geometry*, vol. 39, no. 1-3, pp. 419–441, 2008.
- [16] B. Sharif and Y. Bresler, "Physiologically improved NCAT phantom (PINCAT) enables in-silico study of the effects of beat-to-beat variability on cardiac MR," in *Proceedings of the Annual Meeting of ISMRM*, *Berlin*, vol. 3418, 2007.
- [17] S. Narayanan, A. Toutios, V. Ramanarayanan, A. Lammert, J. Kim, S. Lee, K. Nayak, Y.-C. Kim, Y. Zhu, L. Goldstein *et al.*, "Realtime magnetic resonance imaging and electromagnetic articulography database for speech production research (TC)," *The Journal of the Acoustical Society of America*, vol. 136, no. 3, pp. 1307–1311, 2014.

- [18] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magnetic Resonance in Medicine*, vol. 43, no. 5, pp. 682–690, 2000.
- [19] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic mri exploiting sparsity and low-rank structure: kt slr," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 5, pp. 1042–1054, 2011.
- [20] A. Singer, R. Erban, I. G. Kevrekidis, and R. R. Coifman, "Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps," *Proceedings of the National Academy of Sciences*, vol. 106, no. 38, pp. 16090–16095, 2009.
- [21] K. P. Pruessmann, M. Weiger, P. Brnert, and P. Boesiger, "Advances in sensitivity encoding with arbitrary k-space trajectories," *Magnetic Resonance in Medicine*, vol. 46, no. 4, pp. 638–651, 2001. [Online]. Available: http://dx.doi.org/10.1002/mrm.1241
- [22] K. K. Bhatia, J. Caballero, A. N. Price, J. V. Hajnal, and D. Rueckert, "Fast reconstruction of highly-undersampled dynamic MRI using random sampling and manifold interpolation," in *International Society on Magnetic Resonance in Medicine* 2015.
- [23] M. Usman, D. Atkinson, C. Kolbitsch, T. Schaeffter, and C. Prieto, "Manifold learning based ECG-free free-breathing cardiac CINE MRI," *Journal of Magnetic Resonance Imaging*, vol. 41, no. 6, pp. 1521–1527, 2015.
- [24] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [25] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [26] 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*, 2015.
- [27] L. Feng, D. Sodickson, and R. Otazo, "A Robust and Automatic Cardiac and Respiratory Motion Detection Framework for Self-Navigated Radial MRI," in *International Society on Magnetic Resonance in Medicine* 2014.
- [28] U. Nakarmi, Y. Wang, J. Lyu, and L. Ying, "Dynamic magnetic resonance imaging using compressed sensing with self-learned nonlinear dictionary (NL-D)," in 2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI), April 2015, pp. 331–334.