



Exploiting structural redundancy in $\mathbf{q}$ -space for improved EAP reconstruction from highly undersampled $(\mathbf{k}, \mathbf{q})$ -space in DMRI[☆]

Jiaqi Sun, Alireza Entezari, Baba C. Vemuri^{*}

Computer and Information Science and Engineering, University of Florida, Gainesville, FL, 32611, USA

ARTICLE INFO

Article history:

Received 24 February 2018

Revised 22 November 2018

Accepted 26 February 2019

Available online 13 March 2019

Keywords:

Diffusion MRI

Ensemble average propagator

$(\mathbf{k}, \mathbf{q})$ -space

Compressed sensing

Sparse coding

Parallel level sets

ABSTRACT

Accurate reconstruction of the ensemble average propagators (EAPs) from undersampled diffusion MRI (dMRI) measurements is a well-motivated, actively researched problem in the field of dMRI acquisition and analysis. A number of approaches based on compressed sensing (CS) principles have been developed for this problem, achieving a considerable acceleration in the acquisition by leveraging sparse representations of the signal. Most recent methods in literature apply undersampling techniques in the $(\mathbf{k}, \mathbf{q})$ -space for the recovery of EAP in the joint $(\mathbf{x}, \mathbf{r})$ -space. Yet, the majority of these methods follow a pipeline of first reconstructing the diffusion images in the $(\mathbf{x}, \mathbf{q})$ -space and subsequently estimating the EAPs through a 3D Fourier transform. In this work, we present a novel approach to achieve the direct reconstruction of $P(\mathbf{x}, \mathbf{r})$ from partial $(\mathbf{k}, \mathbf{q})$ -space measurements, with geometric constraints involving the parallelism of level-sets of diffusion images from proximal $\mathbf{q}$ -space points. By directly reconstructing $P(\mathbf{x}, \mathbf{r})$ from $(\mathbf{k}, \mathbf{q})$ -space data, we exploit the incoherence between the 6D sensing and reconstruction domains to the fullest, which is consistent with the CS-theory. Further, our approach aims to utilize the inherent structural similarity (parallelism) of the level-sets in the diffusion images corresponding to proximally-located $\mathbf{q}$ -space points in a CS framework to achieve further reduction in sample complexity that could facilitate faster acquisition in dMRI. We compare the proposed method to a state-of-the-art CS based EAP reconstruction method (from joint $(\mathbf{k}, \mathbf{q})$ -space) on simulated, phantom and real dMRI data demonstrating the benefits of exploiting the structural similarity in the $\mathbf{q}$ -space.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Diffusion-weighted MRI (dMRI) is an imaging technique that allows for the inference of axonal fiber connectivity in biological tissues non-invasively by sensitizing the MR signal to water diffusion. The water diffusion process is fully characterized by the ensemble average propagator (EAP), defined in the displacement $\mathbf{r}$ -space at each location $\mathbf{x}$. It is related to the dMR measurements in $(\mathbf{k}, \mathbf{q})$ -space through the 6D Fourier transform under the narrow pulse assumption (Callaghan, 1991):

$$\hat{S}(\mathbf{k}, \mathbf{q}) = \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} P(\mathbf{x}, \mathbf{r}) \exp(-2\pi j(\mathbf{x}^t \mathbf{k} + \mathbf{q}^t \mathbf{r})) d\mathbf{r} d\mathbf{x}. \quad (1)$$

In order to reconstruct the EAP with a reasonable angular accuracy, one usually needs to acquire diffusion-weighted images along a substantial number of sensitizing gradient directions, such as in multi-shell high angular resolution diffusion imaging (MS-HARDI) and diffusion spectrum imaging (DSI). For each of the

gradient direction, a full 3D acquisition in the $\mathbf{k}$ -space follows. The time incurred in this extensive data acquisition is the key problem making MS-HARDI and DSI impractical for clinical use. Thus, one particular important topic in this area is to accelerate the acquisition of diffusion MRI (especially in MS-HARDI and DSI) while maintaining accurate estimation of the diffusion process via the reconstruction of the EAPs.

Over the past decade, various techniques have been proposed to this end. From an acquisition perspective, faster MR imaging techniques such as parallel imaging can be applied to maintain dense signal measurement configurations while reducing acquisition time. On the other hand, one can strive to maintain accurate reconstruction of the EAPs by exploiting redundancies in the diffusion images to reduce the number of required measurements. In this paper, we focus our attention on the application of compressed sensing (CS) to achieve this goal, which may be employed in addition to these other forms of acceleration (Shi et al., 2015).

Compressed sensing (CS) aims to recover signals from sub-Nyquist sampled measurements, provided that the signal is compressible in some transform domain and is sampled in an incoherent manner (Donoho, 2006). Following its successful application

[☆] **Conflict of interest:** None.

^{*} Corresponding author.

E-mail address: vemuri@cise.ufl.edu (B.C. Vemuri).

to the acceleration of MR image acquisition (Lustig et al., 2007), CS has been utilized to accelerate diffusion tensor imaging (DTI) from partial $\mathbf{k}$ -space measurements (Shi et al., 2015; McClymont et al., 2016). In the context of HARDI/DSI, CS has been investigated rather extensively for the reconstruction of the signal, EAP, orientation distribution function (ODF) or fiber orientation distribution function (FOD) from partial $\mathbf{q}$ -space measurements (Michailovich et al., 2011; Landman et al., 2012; Merlet and Deriche, 2013; Bilgic et al., 2013; Aranda et al., 2015; Auria et al., 2015; Paquette et al., 2015; Ning et al., 2015; Daducci et al., 2015; Cheng et al., 2015b). These methods generally assume a model for the sparse representation of the diffusion signal and achieve signal reconstruction prior to EAP/ODF/FOD estimation. Most recent attempts among these include (Schwab et al., 2016), wherein the authors take into consideration the spatial redundancy of dMRI by defining a joint separable spatial-angular domain basis for the sparse representation of the entirety of the signal in the $(\mathbf{x}, \mathbf{q})$ -space. This joint representation model was shown to achieve better sparsity in comparison to voxel-wise angular sparsity models. Nonetheless, as pointed out by the authors, they did not apply CS jointly to the $(\mathbf{k}, \mathbf{q})$ -space. On a relevant note, CS has also been applied to super-resolution dMRI, i.e. the problem of reconstructing high-resolution diffusion images from low-resolution ones to reduce acquisition time while achieving good spatial resolution (Ning et al., 2016; Yin et al., 2016).

In the past few years, developments in the application of CS to joint $(\mathbf{k}, \mathbf{q})$ -space for the recovery of the signal/EAP have been reported. In Awate and DiBella (2013), the authors proposed an orientation-invariant dictionary to sparsely represent the diffusion signal in the $(\mathbf{x}, \mathbf{q})$ -space, which is comprised of atoms representing key types of diffusion profiles, including isotropic diffusion, single and two fiber tracts crossing at certain angles. In more recent works (Mani et al., 2015; Cheng et al., 2015a), the CS concepts were applied in a similar manner to reconstruct the $(\mathbf{x}, \mathbf{q})$ -space signal and then the EAPs. In both (Mani et al., 2015; Cheng et al., 2015a), a pre-defined basis/dictionary was used for the sparse representation of the diffusion signal, sparsity was enforced on the coefficients in the representation, and the diffusion MR signal was first estimated prior to the reconstruction of the EAP. From a CS standpoint, by sparsely representing and reconstructing $S(\mathbf{x}, \mathbf{q})$ from partial $(\mathbf{k}, \mathbf{q})$ -space samples, all three aforementioned methods effectively utilized the incoherence between the 3D Fourier dual spaces of $\mathbf{k}$ and $\mathbf{x}$, but did not exploit the incoherence between the pair $\mathbf{q}$ and $\mathbf{r}$.

1.1. Motivation and overview

In contrast to above described methods, we propose an approach to fully exploit the incoherence between the 6D Fourier pair of $(\mathbf{k}, \mathbf{q})$ and $(\mathbf{x}, \mathbf{r})$ -space. To demonstrate the advantage of harnessing such incoherence, we first present an illustrative experiment in 2-dimensions.

In this example, we aim to construct an analogy between the 2D Fourier dual space $(x, y) - (u, v)$ and the 6D Fourier dual space we are faced with in the EAP reconstruction problem $(\mathbf{x}, \mathbf{r}) - (\mathbf{k}, \mathbf{q})$ (expanded to be $(x_i, x_j, x_l, r_i, r_j, r_l) - (k_i, k_j, k_l, q_i, q_j, q_l)$). A sparse discrete signal (of size 32×32) was generated in the (x, y) -space and the problem is to reconstruct this sparse signal from its partial Fourier samples. This is analogous to the reconstruction of EAPs from partial $(\mathbf{k}, \mathbf{q})$ samples. We perform reconstruction using two methods analogous to the direct reconstruction approach proposed in this paper and the common framework (indirect method) shared in other $(\mathbf{k}, \mathbf{q})$ -space methods respectively. We randomly sample the Fourier domain using variable density sampling, and reconstruct the signal using the aforementioned two approaches from the same set of samples. In approach-1, we directly reconstruct

the spatial domain signal by solving a l_1 minimization problem enforcing the sparsity of the signal. In approach-2 (indirect method), from the (u, v) samples, we reconstruct the (x, y) space first and then recover the spatial signal by applying inverse Fourier transform in the y direction for each x . In the 6D context, this is equivalent to first reconstructing $S(\mathbf{x}, \mathbf{q})$ prior to estimating $P(\mathbf{x}, \mathbf{r})$ by applying 3D Fourier transform at each grid $(\mathbf{x})$. We vary the sampling rate and perform 10 repetitions for each sampling rate to account for the randomness in the sampling. We present the original signal in Fig. 1(a), reconstructions from an example set of 33% Fourier samples using the two methods in (b) and (c), and the average NMSE of reconstruction at various sampling rates in (d).

It is clearly evident from this toy example that in a CS-based signal recovery framework, it is critical and beneficial to leverage the incoherence between Fourier dual spaces by making sure the sensing and sparse representation/reconstruction occur in completely dual domains and directly perform reconstruction. In our context, this indicates that a direct reconstruction of $P(\mathbf{x}, \mathbf{r})$ from its Fourier dual space samples $\hat{S}(\mathbf{k}, \mathbf{q})$ will potentially enable higher undersampling rates compared to reconstructing $S(\mathbf{x}, \mathbf{q})$ as an intermediate step.

In accordance to this principle, we present a CS framework to directly reconstruct $P(\mathbf{x}, \mathbf{r})$ from jointly under-sampled $(\mathbf{k}, \mathbf{q})$ data.¹ To illustrate, since the sampling operation in MS-HARDI occurs in the $(\mathbf{k}, \mathbf{q})$ -space, only a sparse representation in the $(\mathbf{x}, \mathbf{r})$ -space enables the significant reductions in sampling rates while guaranteeing exact recovery. In this CS framework, the surfacelet transform is utilized to sparsely represent the EAP $P(\mathbf{x}, \mathbf{r})$ and a total variation (TV) penalty term was included to promote the smooth reconstruction of the entire EAP field. In lieu of this theoretically consistent technique from the CS-theory viewpoint, we will use this as our baseline technique for experimental comparisons with the proposed improvement. Validations are performed on both synthetic and real human brain data sets, and the results demonstrated the value of the method.

1.1.1. Overview of the proposed improvement

Intuitively, CS takes advantage of the redundancy in compressible signals to reduce the number of necessary samples. The redundancy manifests as sparsity in the coefficients when the signal is represented in a sparsifying basis. Thus it goes without saying that the more thoroughly the redundancy is exploited, the more time savings are to be expected. One critical question to ask in the dMRI context is then, have the above described existing CS based methods fully utilized the redundancy present in diffusion MR images? The answer is no. For each gradient direction $\mathbf{q}$, a diffusion-weighted image $S(\mathbf{x}, \mathbf{q})$ is acquired for the entire volume (over all voxels $\mathbf{x}$), which encodes the response of the water molecule to the specific sensitizing magnetic gradient at each voxel. Hence, all the diffusion-weighted images from different $\mathbf{q}$'s are measurements taken of the same subject being imaged. It is then to be expected that similar structural information will be contained in them, especially the images corresponding to nearby $\mathbf{q}$ points. In Fig. 2, we present the entire set of $\mathbf{q}$ -space sample points used in a HARDI dataset obtained from the WuMin-Human Connectome Project (HCP) (Van Essen et al., 2013) and the diffusion images acquired at 3 closely located $\mathbf{q}$ -space points. It is evident from these that a strikingly similar structural information is encoded in all 3 images, resulting in an approximately similar visual appearance in edge locations, textures and intensity gradients at corresponding spatial locations across the image. This provides unequivocal evidence that redundancy exists among the diffusion-weighted images corresponding to neighboring $\mathbf{q}$'s.

¹ Preliminary results were published in a conference paper (Sun et al., 2015).




Fig. 1. Toy example. (a) Ground-truth sparse signal, (b) direct reconstruction from 33% Fourier samples enforcing sparsity in spatial domain, (c) indirect reconstruction from the same set of samples, (d) reconstruction accuracy (averaged over 10 runs) for various sampling rates using the two methods.




Fig. 2. An example of neighboring $\mathbf{q}$ -space points and their corresponding diffusion images. 3 $\mathbf{q}$ points in a neighborhood are highlighted in (a), and their corresponding diffusion-weighted images are presented in (b) with the same coloring scheme.




Fig. 3. An overview of the proposed framework.

In the context of diffusion image reconstruction and denoising, structural similarities among images acquired at different $\mathbf{q}$ -space locations have been exploited by several researchers in the past. In Alaya et al. (2016), a geometrical approach was presented to recover diffusion images not acquired at locations in $\mathbf{q}$ -space in HARDI based on the images acquired at their neighboring $\mathbf{q}$ -space points (determined through Delaunay triangulation). In Haldar et al. (2013), the authors proposed a penalized maximal likelihood framework to jointly reconstruct a set of diffusion images from their $\mathbf{k}$ -space samples. Within this framework, the spatial smoothness of each diffusion image (for a specific $\mathbf{q}$) is enforced using a regularization functional, wherein a spatially varying line-process variable is incorporated to explicitly model the edge structure and control the power of the spatial smoothness constraint locally. In a later work of the same authors (Lam et al., 2014), a similar edge-prior was utilized in a maximum a posterior framework for the denoising of diffusion images.

In this paper, we present a CS framework for direct EAP reconstruction from highly under-sampled $(\mathbf{k}, \mathbf{q})$ data. $(\mathbf{k}, \mathbf{q})$ and $(\mathbf{x}, \mathbf{r})$ spaces are the key ingredients enabling sparse representation for $P(\mathbf{x}, \mathbf{r})$ is accomplished using surfacelet basis. The most attractive feature of surfacelet basis is the inherent directional selectivity that leads to a sparse representation in the $\mathbf{r}$ -space. The key distinction between our approach and existing approaches is that, enforcing sparsity in $P(\mathbf{x}, \mathbf{r})$ entitles us to leverage incoherent sensing, not only in $\mathbf{k}$, but also in the $\mathbf{q}$ -space simultaneously. Therefore, our approach presented here stands to benefit from practical guarantees for accurate reconstruction from partial $(\mathbf{k}, \mathbf{q})$ data. To further exploit redundancy in the $(\mathbf{k}, \mathbf{q})$ space, we utilize the structural similarity in diffusion images corresponding to proximally-located $\mathbf{q}$ points. To accomplish this, we propose an approach that emphasizes the geometric correlation between these images by considering the degree of parallelism between their level sets. We consider it as a priori knowledge and incorporate it into the 6D CS framework (described earlier) as a regularizing prior. Fig. 3 depicts a graphical self-explanatory overview of our approach. It is important to note that the level set based constraint, the key concept in this regularizing prior, is very distinct from an edge-based constraint (which is the emphasis of the previously discussed works in this context). This is because, level sets in general do not necessarily correspond to edges in images, and vice versa. To the best of our

knowledge, the parallelism of the level sets (as a prior) has never been utilized in any of the existing CS-based methods for diffusion signal/EAP/ODF reconstruction. Yet, a similar concept has been investigated in completely different applications namely, color image denoising/demosaicing and multi-modality medical image reconstruction (Ehrhardt and et al., 2014; Ehrhardt and Arridge, 2014). It has been shown in these works that the exploitation of structural similarity by enforcing the parallelism of the level sets leads to improved results in both tasks respectively.

The rest of the paper is organized as follows. In Section 2, we present a brief overview of the general CS-based framework for EAP recovery from joint- $(\mathbf{k}, \mathbf{q})$ space, followed by the details of the proposed method including the problem formulation and numerical solution. In Section 3, several aspects of our experimental design and descriptions of the datasets used for validation of our approach are first presented. We then showcase the experimental results on the various datasets and provide a discussion. Finally, in Section 4 we draw conclusions.

2. Material and methods

2.1. Compressed sensing for EAP reconstruction

In general, compressed sensing (CS) recovers the unknown $\mathbf{x}$ from partial measurements $\mathbf{y}$ by solving an underdetermined system while enforcing the sparsity of $\mathbf{x}$ in a certain transform domain along with data consistency (Donoho, 2006). The three ingredients of the CS framework necessary to guarantee accurate reconstruction are:

- **Sparsity:** The function to be reconstructed needs to be sparsely representable, possibly in some transform domain.
- **Incoherent sensing:** The data for reconstruction must be acquired in a domain incoherent (e.g., dual) to the domain in which the function is sparsely representable.
- **Nonlinear reconstruction:** The reconstruction problem involves an (convex) optimization process.

One way to formulate the problem of EAP reconstruction from partial $(\mathbf{k}, \mathbf{q})$ data appropriate in the CS settings is detailed in the following. Suppose we are interested in reconstructing the EAPs within a 3D rectangular volume Ω , containing N voxels. Let P be

a 6D matrix formed by concatenating the N EAP examples $\{P_n\}_{n=1}^N$ (each a 3D matrix) that are to be estimated. In the case where the sparsity is enforced on each individual EAP under a certain sparsifying transform Ψ , in order to reconstruct P from partial $(\mathbf{k}, \mathbf{q})$ measurements, denoted by Q , one needs to solve the following minimization problem:

$$\min_P \frac{1}{2} \|\mathcal{F}_u P - Q\|_F^2 + \mu \sum_{n=1}^N \|\Psi P_n\|_1 \quad (2)$$

where $\mathcal{F}_u$ is the partial Fourier operator in 6-dimensions, i.e. an undersampled 6-D Fourier transform evaluated at selected frequencies corresponding to the undersampling scheme in the signal space. The 6 dimensional Fourier dual spaces are spanned by $(\mathbf{k}, \mathbf{q})$ and $(\mathbf{x}, \mathbf{r})$ respectively.

In order to achieve a reconstruction of satisfactory fidelity, various regularization terms may be added to the basic CS formulation above to enforce a priori information pertaining to the desired solution. One example of such regularization terms is the total variation (TV) penalty, which is widely used in many inverse problems including MR image reconstruction from partial $\mathbf{k}$ -space samples (Lustig et al., 2007). By taking regularization terms into account, a more general formulation for EAP reconstruction in a CS framework is presented below:

$$\min_P \frac{1}{2} \|\mathcal{F}_u P - Q\|_F^2 + \mu \sum_{n=1}^N \|\Psi P_n\|_1 + \sum_{k=1}^K \gamma_k \mathcal{R}_k(P) \quad (3)$$

in which, $\mathcal{R}_k$ is the k -th regularization term and γ_k is the trade-off parameter between the data fidelity and the regularizer.

2.2. Surfacelet transform revisited

One of the key factors essential for the successful application of CS is to select the basis that provides the best sparsity for the signal of interest. In the context of dMR signal/EAP/ODF reconstruction, various basis/transforms have been utilized, such as spherical ridgelets and spherical harmonics to name a few. While spherical ridgelets and spherical harmonics are more tailored for the representation of dMR signals; surfacelets (Lu and Do, 2007), with their ability to efficiently capture directional information, are particularly well-suited for sparse EAP representation.

The surfacelet transform is implemented as a combination of a multi-scale pyramid with 3D directional filter banks (3D-DFB) (Lu and Do, 2007). The basis functions are a spatial domain representation of symmetric pyramids partitioning the frequency space. Fig. 4 depicts one example surfacelet basis in the frequency domain as well as spatial domain. We refer interested readers to Lu and Do (2007) for more technical details.

EAPs can be sparsely represented in Surfacelet basis, as depicted in Fig. 5. To demonstrate this sparsity, we constructed a field (30×30) of EAPs sampled on 3D Cartesian lattices of various sizes $(32 \times 32 \times 32$ and $64 \times 64 \times 64)$, and applied the surfacelet transform to the EAP at each voxel. We then sorted the coefficients of




Fig. 4. 3D renderings of an example surfacelet basis in frequency domain (left) and spatial domain (right) respectively.




Fig. 5. Average SNRs of EAP reconstruction from partial surfacelet transform coefficients with lattices of size 32 and 64 respectively.

the transform by absolute value and reconstructed the EAPs from only the top few coefficients. Fig. 5 shows the average SNRs of the reconstruction from a fraction of the total number of coefficients. Evidently, for both lattice sizes, with less than 3% of the coefficients, we can easily achieve an SNR of over 25, which is commonly used as a benchmark for high quality reconstructions. This experiment also showed that a denser sampling grid promotes higher sparsity. As a trade-off between sparsity and computational costs, we chose 32^3 as the size of our sampling lattice.

Following the notations introduced previously, $P_n(\mathbf{r})$ ($n = 1, \dots, N$), the EAP at the n th voxel, can be expanded in terms of surfacelet basis functions $\varphi_m^{(l)}(\cdot)$, corresponding to different scales (l) and spectral directions (m) as:

$$P_n(\mathbf{r}) = \sum_{m,l} c_{m,l} \varphi_m^{(l)}(\mathbf{r}) \quad (4)$$

Let $\mathbf{c}_n := [c_{m,l}]$ be the surfacelet coefficient vector for P_n and denote the surfacelet transform by Ψ , we can then write $\mathbf{c}_n = \Psi P_n$ and seek a sparse coefficient vector by minimizing the ℓ_1 norm of $\mathbf{c}_n$ for each voxel.

2.3. Parallel level sets

Equipped with an appropriate sparsifying transform for the representation of the EAP, what remains to guarantee an accurate reconstruction is the enforcement of the most valuable a priori information at hand. As we have illustrated, there exists considerable structural similarities within diffusion MR images corresponding to adjacent $\mathbf{q}$'s, which leads to a high degree of (similarity) parallelism between the level sets of these neighboring images. We propose to design and incorporate a regularization term to emphasize such geometric correlations by aligning the gradients (which are perpendicular to the level sets) between the corresponding images.

Consider $S(\mathbf{q}_i)$ and $S(\mathbf{q}_j)$, the diffusion-weighted image (across all voxels in Ω) for two closely located $\mathbf{q}$ points. We denote $S(\mathbf{q}_i)$ and $S(\mathbf{q}_j)$ by S^i and S^j for simplicity. The level sets of the image pair are considered parallel if the gradients ∇S^i and ∇S^j form an angle of 0 degrees at each voxel $\mathbf{x} \in \Omega$. We measure the degree of parallelism between the gradients at location $\mathbf{x}$ by,

$$\|\nabla S^i(\mathbf{x})\| \|\nabla S^j(\mathbf{x})\| - |\langle \nabla S^i(\mathbf{x}), \nabla S^j(\mathbf{x}) \rangle|. \quad (5)$$

It is a well known fact that $|\langle \nabla \mathbf{u}, \nabla \mathbf{v} \rangle| \leq \|\nabla \mathbf{u}\| \|\nabla \mathbf{v}\|$ for arbitrary vectors $\mathbf{u}, \mathbf{v}$ and equality only occurs when $\mathbf{u}$ and $\mathbf{v}$ are parallel,



Fig. 6. FA map of the selected slice and ROI.

hence the above defined term is always nonnegative and reaches minimum 0 if and only if the gradients are parallel at location $\mathbf{x}$. It measures how parallel the gradients or level sets are locally, with greater values indicating greater angles between the gradients and further distance from parallelism. Now we can define the global measure for the degree of parallelism of two diffusion-weighted images as

$$\mathcal{PL}(\nabla S^i, \nabla S^j) := \int_{\Omega} (\|\nabla S^i(\mathbf{x})\| \|\nabla S^j(\mathbf{x})\| - |\langle \nabla S^i(\mathbf{x}), \nabla S^j(\mathbf{x}) \rangle|) d\mathbf{x} \quad (6)$$

We want to guarantee that all pairs of images corresponding to neighboring $\mathbf{q}$ points have level sets that are close to or parallel to each other. Denote the entire 6D $S(\mathbf{x}, \mathbf{q})$ volume as S . One possible form of the regularization term to be incorporated into the CS framework can be defined on S as:

$$\mathcal{R}(S) = \sum_{m=1}^M \sum_{l \in \mathcal{N}(m)} \mathcal{PL}(\nabla S^m, \nabla S^l) \quad (7)$$

where $\mathcal{N}(m)$ represents the index set of all the $\mathbf{q}$ points in $\mathbf{q}_m$'s neighborhood and M is the total number of $\mathbf{q}$ sampling points. Note that the regularizer is defined on S while the EAP, P , is what we desire to reconstruct. However, S and P are related at each voxel $\mathbf{x}$ through a 3D Fourier transform, which will serve as the link that makes the incorporation of $\mathcal{R}(S)$ appropriate in the CS framework. Also, note that the gradients in the parallel-level set constraint are not computed from the raw data but from the updated estimates of S presented in step-2 of the Split Bregman technique in the next section. Details are presented in the following section.

2.4. Problem formulation and solution

In addition to the parallelism of level sets in adjacent $\mathbf{q}$ volumes, another prior we have is that the reconstructed EAP field must be smooth across the grid. To this end, a TV penalty can be included as another regularizer to enforce spatial homogeneity in the reconstruction. Finally, we formulate the EAP reconstruction problem from $(\mathbf{k}, \mathbf{q})$ data Q as the following optimization problem:

$$\begin{aligned} \min_P \quad & \frac{1}{2} \|\mathcal{F}_u P - Q\|_F^2 + \mu \sum_{n=1}^N \|\Psi P_n\|_1 + \gamma_1 \|P\|_{TV} + \gamma_2 \mathcal{R}(S) \\ \text{s. t. } \quad & S_n = \mathcal{F}_{3D} P_n, \quad n = 1, 2, \dots, N \end{aligned} \quad (8)$$

In the above formulation, S_n denotes the diffusion signal $S(\mathbf{x}, \mathbf{q})$ measured at a fixed voxel $\mathbf{x}_n$ for all M $\mathbf{q}$ points. It is related to the EAP at the voxel through a 3D Fourier transform. S is formed by stacking all the S_n 's for each voxel $n = 1, 2, \dots, N$. On the other

hand, S^m (see notation introduced in Section 2.3) denotes the diffusion signal across all N voxels for a fixed $\mathbf{q}(\mathbf{q}_m)$.

We solve the above optimization problem using a Split Bregman scheme (Goldstein and Osher, 2009). After converting it into an unconstrained problem with a Lagrange multiplier, we introduce auxiliary parameters as replacements $c_n \leftarrow \Psi P_n$, $d^m \leftarrow \nabla S^m$ for all n and m . Substituting in to Eq. (7), the split formulation of the problem becomes

$$\begin{aligned} \min_{P, S, c_n, d^m} \quad & \frac{1}{2} \|\mathcal{F}_u P - Q\|_F^2 + \mu \sum_{n=1}^N \|c_n\|_1 + \gamma_1 \|P\|_{TV} \\ & + \gamma_2 \sum_{m=1}^M \sum_{l \in \mathcal{N}(m)} \mathcal{PL}(d^m, d^l) + \frac{\lambda}{2} \sum_{n=1}^N \|\mathcal{F}_{3D} P_n - S_n\|_F^2 \\ & + \frac{\lambda_1}{2} \sum_{n=1}^N \|\Psi P_n - c_n\|_2^2 + \frac{\lambda_2}{2} \sum_{m=1}^M \|d^m - \nabla S^m\|_2^2. \end{aligned}$$

We then decompose this into subproblems using the split Bregman iterations to alternatively update P , S and the auxiliary parameters until convergence.

$$\begin{aligned} \text{Step 1: } P^{(t+1)} = \arg \min_P \quad & \frac{1}{2} \|\mathcal{F}_u P - Q\|_F^2 + \gamma_1 \|P\|_{TV} \\ & + \frac{\lambda}{2} \sum_{n=1}^N \|\mathcal{F}_{3D} P_n - S_n^{(t)}\|_F^2 \\ & + \frac{\lambda_1}{2} \sum_{n=1}^N \|\Psi P_n - c_n^{(t)} - b_n^{(t)}\|_2^2 \end{aligned}$$

$$\begin{aligned} \text{Step 2: } S^{(t+1)} = \arg \min_S \quad & \frac{\lambda}{2} \sum_{n=1}^N \|\mathcal{F}_{3D} P_n^{(t+1)} - S_n\|_F^2 \\ & + \frac{\lambda_2}{2} \sum_{m=1}^M \|d^{m(t)} - \nabla S^m - b^{m(t)}\|_2^2 \end{aligned}$$

$$\begin{aligned} \text{Step 3: } c_n^{(t+1)} = \text{shrink} \left(\Psi P_n^{(t+1)} + b_n^{(t)}, \frac{\mu}{\lambda_1} \right) \\ b_n^{(t+1)} = b_n^{(t)} + \Psi^* c_n^{(t+1)} - P_n^{(t+1)}, \quad n = 1, \dots, N \\ d^{m(t+1)} = \arg \min_{d_m} \gamma_1 \sum_{m=1}^M \sum_{l \in \mathcal{N}(m)} \mathcal{PL}(d^m, d^l) \\ + \frac{\lambda_2}{2} \sum_{m=1}^M \|d^m - \nabla S^{m(t+1)} - b^{m(t)}\|_2^2 \\ b^{m(t+1)} = b^{m(t)} + \nabla S^{m(t+1)} - d^{m(t+1)}, \quad m = 1, \dots, M \end{aligned}$$

In the above, b_n and b^m are the Bregman parameters introduced for the updates of c_n and d^m respectively, Ψ^* denotes the inverse Surfacelet transform. Most of the steps can be solved analytically, while for the d^m update in step 3, we used a limited-memory BFGS based Quasi-Newton method for a numerical solution.

3. Experiments and discussions

3.1. Experimental design

In this section, we preset several aspects of the experimental design used in this work. This involves: the choice of the method/s for comparison, the sampling scheme used in the $(\mathbf{k}, \mathbf{q})$ space for data acquisition and finally the evaluation metrics used to perform the comparisons.




Fig. 7. NMSE of EAP reconstructions, using the PLS and the baseline methods respectively, from synthetic data jointly undersampled in the $(\mathbf{k}, \mathbf{q})$ -space at various sampling rates and noise levels.

3.1.1. Comparison to the state-of-the-art

We demonstrate the power of employing the structural similarity in $\mathbf{q}$ -space for direct EAP reconstruction from highly undersampled $(\mathbf{k}, \mathbf{q})$ -space data via performance comparison of the proposed method (denoted by PLS for brevity) with the CS method used in a preliminary conference version of this work (Sun et al., 2015) (denoted by baseline). This method is the choice for baseline comparisons because, to date, it is the only method in the literature that directly reconstructs EAP from partially sensed 6D $(\mathbf{k}, \mathbf{q})$ -space without signal reconstruction as an intermediate step. Further, by directly reconstructing the EAPs from signal in its Fourier dual space, this method, when compared to other existing CS-based methods, better utilizes the CS principles and more thoroughly exploits the incoherence in the 6D Fourier dual (see Section 1 for a more detailed explanation).

3.1.2. Sampling scheme

Experiments are conducted on diffusion MR images acquired using the DSI scheme, which has been widely used as the benchmark for comparisons in reported literature (CS related reconstruction in dMRI). Since DSI sampling is designed on a Cartesian grid, in our current implementation the neighborhood of any given $\mathbf{q}$ point is chosen to consist of its 6-connected points in $\mathbf{q}$ -space. Yet, the framework is general in the sense that it is independent of the choice of neighborhood. The optimal choice of $\mathbf{q}$ -space neighborhood will be a topic of our future work.

The undersampling of the 6D space is performed in the 3D $\mathbf{k}$ and $\mathbf{q}$ -space independently. The level of undersampling is set to be equal in both spaces, for example, to achieve a sampling rate of 25% in the joint space, the sampling rate for both $\mathbf{k}$ and $\mathbf{q}$ -space is 50%. Such a choice was based on extensive experimentation on data undersampled with varying ratios between the two spaces. Two extreme cases are when the undersampling only takes place in one of the two spaces, which will henceforth be referred to as, “ $\mathbf{k}$ -only” and “ $\mathbf{q}$ -only” undersampling respectively. In the experiments to follow, in addition to the EAP reconstruction results from joint $(\mathbf{k}, \mathbf{q})$ -undersampled data, we will also showcase the results

obtained from partial data sampled with the $\mathbf{k}$ -only and $\mathbf{q}$ -only undersampling methods respectively.

In each of the 3D $\mathbf{k}$ and $\mathbf{q}$ -spaces, we select samples randomly with the sampling density scaled according to a power of the distance from the origin. This power law sampling scheme has been very widely used in related literature, including CS-based MR image reconstruction (Donoho, 2006) for the $\mathbf{k}$ -space and EAP/ODF/signal reconstruction from dMRI (Bilgic et al., 2013; Cheng et al., 2015a) for the $\mathbf{q}$ -space. It was shown to yield good CS-based recovery in both contexts.

3.1.3. Evaluation

Evaluations are carried out both quantitatively, using normalized mean squared error (NMSE) with respect to the ground-truth EAPs, and qualitatively through visual inspection. The reference ground-truth EAPs are obtained through conventional DSI reconstruction using the fully sampled data, also labeled by some in literature as “gold standard” data. In DSI reconstruction, the zero-padding in $\mathbf{q}$ -space is performed to the size of $16 \times 16 \times 16$. No Hanning filter is applied prior to reconstruction to ensure fair comparison between the ground-truth and the various reconstructions. For visualization purpose, we interpolate the ground-truth and the reconstructed EAPs (which are continuous 3D functions evaluated on a Cartesian grid) onto a sphere before performing fitting with spherical harmonics. Visualization of the spherical harmonic representations of the EAPs is done using a Java-based diffusion MRI processing software fanDTasia (<https://www.cise.ufl.edu/~abampou/lab/fanDTasia/>). For synthetic and phantom data where underlying fiber configuration is known, we also perform angular analysis based on the peaks detected in the reconstructed EAPs. The peak directions are estimated using function `peak_directions_nl` in Dipy (Garyfallidis et al., 2014) (<http://nipy.org/dipy/>), a popular open source diffusion MRI toolbox developed in Python. For the 3 parameters in the function, we kept the default value for `min_separation_angle` (25) and `xtol` ($1e-07$), and chose 0.75 as `relative_peak_threshold` to suppress subsidiary peaks.




Fig. 8. NMSE of EAP reconstructions, using the PLS and the baseline methods respectively, from synthetic data undersampled using joint-(**k**, **q**), **k**-only and **q**-only schemes at various rates, for (a) noise-free and (b) noisy data with SNR = 5.




Fig. 9. EAP reconstructions and angular analysis results for noise-free synthetic data using the PLS and the baseline methods respectively. (a) Ground-truth EAPs estimated from gold standard noise-free data, EAP reconstructions using 5% of the data using (b) the PLS method and (c) the baseline method. (d-e) Crossing angle errors within fiber crossing regions for the reconstructions displayed in (b-c) respectively. (f-g) Errors in the number of peaks detected from the reconstructions displayed in (b-c) respectively.




Fig. 10. EAP reconstructions and angular accuracy analysis results obtained using the PLS and the baseline methods respectively, from noisy synthetic data ($\text{SNR} = 5$) undersampled with various schemes. (a) Ground-truth EAPs estimated from gold standard noise-free data, EAP reconstructions from 15% jointly undersampled ($\mathbf{k}$, $\mathbf{q}$)-space data using (b) the PLS method and (c) the baseline method. (d-e) EAP reconstructions using the PLS method from 15% of the data undersampled in $\mathbf{k}$ and $\mathbf{q}$ -space only. (f-i) Crossing angle errors within fiber crossing regions for the reconstructions displayed in (b-e) respectively. (j-m) Errors in the number of peaks detected from the reconstructions displayed in (b-e) respectively.

3.2. Implementation

The proposed approach was implemented in Matlab and all computations were performed on a workstation with an Intel Core(TM) i7 CPU930 2.80 GHz x 8 processor and 24GB RAM. We have observed that the main bottleneck of the computation speed of our proposed method is in the updates of d^m (which is a part of Step 3 in the algorithm outlined in Section 2). In our current implementation, we solve the optimization problem involved in this update step using a BFGS-based Quasi-Newton method. We believe further acceleration can be achieved if an analytical solution is derived and used instead. This will be an immediate focus for our future work. The sparsity/regularization trade-off parameters μ , γ_1

and γ_2 in the objective function are tuned through a grid search within range [0.01,1].

3.3. Description of data sets

Various data sets used in our experiments are described in the following paragraphs.

3.3.1. Synthetic data

We simulated a gold standard DSI dataset using a mixture of Gaussian functions. As discussed above, the ($\mathbf{k}$, $\mathbf{q}$) measurements were then generated by applying a 3D Fourier Transform on the diffusion image simulated for each gradient direction. This simula-




Fig. 11. NMSE of EAP reconstructions, using the PLS and the baseline method respectively, from phantom data jointly undersampled in the $(\mathbf{k}, \mathbf{q})$ -space at various sampling rates and noise levels.

tion process does not take into consideration the various practical issues involved in obtaining the magnitude MR images from the complex valued $\mathbf{k}$ space acquisitions (in MR scanners), yet it is very commonly used in the literature for the study of (CS-based) MR image reconstruction from $\mathbf{k}$ -space measurements (Donoho, 2006) as well as the recovery of diffusion signal or EAP/ODF from $(\mathbf{k}, \mathbf{q})$ -space (Cheng et al., 2015a). In order to assess the proposed method in the presence of noise, we contaminated the gold standard data with various levels of Rician noise in the $(\mathbf{x}, \mathbf{q})$ -space. Noise levels are measured by $\text{SNR} = 1/\sigma$, with σ being the standard deviation of the noise.

3.3.2. ISBI HARDI challenge 2013 phantom data

We also evaluated our method on the phantom data provided at the IEEE ISBI (Intl. Symp. on Biomedical Imaging) HARDI challenge 2013 (Daducci et al., 2013). The advantage of employing this data lies in the fact that the phantom is created in a more realistic setting and the ground-truth is available which makes the quantitative assessments of the results possible. The phantom consists of a set of fiber bundles with a wide range of configurations (branching, crossing, kissing), fiber bundles radii, and fiber geometry. The diffusion MR signal is simulated in each voxel considering hindered and restricted diffusion, to account for extra-axonal and intra-axonal diffusion. Depending on the position in space, there is also an isotropic compartment, to account for the CSF contamination close to the ventricles in brain imaging. Finally, the magnitude MR signal is corrupted by Rician noise resulting in SNR of 10, 20 and 30 respectively. The generation of $(\mathbf{k}, \mathbf{q})$ -space measurements was performed in the same manner as described for the synthetic data. We selected our ROI within a slice where the various types of fiber configurations can be best observed, the FA map of which is presented in Fig. 6.

3.3.3. MGH-USC HCP human brain data

The MGH-USC HCP dataset (Setsompop et al., 2013; McNab et al., 2013; Fan et al., 2016; Keil et al., 2013; Polimeni et al., 2016; van der Kouwe et al., 2008; Fan et al., 2014; Fischl, 2012; Greve and Fischl, 2009; Andersson and Sotiropoulos, 2015; 2016;

Crawford et al., 2016) we used in our real data experiment was collected on a Siemens 3T Connectome scanner using the DSI scheme. The images were acquired at 2 mm isotropic resolution with a maximum b-value of 10,000 s/mm² and TE/TR = 77/5900 ms, resulting in a $104 \times 104 \times 55$ volume. We pick our ROI in centrum semiovale where projection, commissural and association tracts interact, and present the Fractional Anisotropy (FA) map of the selected slice with the ROI highlighted in Fig. 17(a).

3.4. Results and discussions

3.4.1. Analysis of synthetic data

We demonstrate the performance of our method on synthetic data in this section. In Fig. 7, for noise-free and noise contaminated data, we plot the NMSE of the EAP reconstructions (with respect to the ground-truth EAPs) obtained by applying the proposed method (PLS) and the baseline method on partial data undersampled jointly in the $(\mathbf{k}, \mathbf{q})$ space at various sampling rates (from 5% to 25%). Overall, the proposed method provides a decrease of the NMSE by 16%–50% with respect to the reconstructions using the baseline method at corresponding sampling rate and noise level. This improvement in the reconstruction accuracy is solely due to the exploitation of the structural similarity between diffusion images acquired at nearby $\mathbf{q}$ locations.

At SNRs higher than 15, we observed in our experiments that the reconstruction accuracy is maintained quite well by both methods (with respect to performance on noise-free data). Hence, here we present the results only for low SNRs. It can be seen from the plots that at lower sampling rates (below 20%), the effect of the added noise on the reconstruction accuracy is less pronounced for the proposed PLS method than it is for the baseline method. This is an indication that the incorporation of the parallel level set prior further assists in suppressing the noise, in addition to the TV regularization used in both approaches. Another rather promising behavior we observed is that at extremely low sampling rates (such as 5%–10%), the most significant gain is brought forth by the PLS method.




Fig. 12. NMSE of EAP reconstructions, using the PLS and the baseline methods respectively, from partial phantom data with different SNRs, undersampled using joint-($\mathbf{k}$, $\mathbf{q}$), $\mathbf{k}$ -only and $\mathbf{q}$ -only schemes at various rates. (a) SNR = 30, (b) SNR = 20 and (c) SNR = 10.

In Fig. 8, we present the comparison of the reconstruction accuracy using both approaches on partial data undersampled using joint-($\mathbf{k}$, $\mathbf{q}$), $\mathbf{k}$ -only and $\mathbf{q}$ -only schemes. The results for noise-free and noisy data (SNR = 5) are shown in sub-figure (a) and (b) respectively. It is consistent with the findings reported in our preliminary version of this paper, Sun et al. (2015), that the joint undersampling demonstrates advantage over undersampling performed in individual spaces in the context of CS-based reconstruction. For each sampling scheme, the proposed method consistently provides improvement in the reconstruction accuracy compared to baseline (Sun et al., 2015). Note that the wide range of the y-axis in the figure (necessary for the presentation of the error for baseline reconstruction using $\mathbf{q}$ -only undersampling at low sampling rates) causes the difference in the NMSE between the reconstructions from jointly undersampled data using the two methods to appear less significant than it actually is (as is already shown in Fig. 7).

To demonstrate the value of the proposed method in further reducing the amount of data needed for satisfactory EAP reconstructions, we showcase in Fig. 9, visual comparisons of the EAP reconstructions from as little as 5% of gold standard data using the proposed and baseline methods. The ground truth EAP field and EAPs reconstructed from 5% of the noise-free gold standard data using both methods are displayed in the first row. The error in the crossing angles estimated (in the fiber bundle crossing regions) and the number of peaks detected (within the entire slice)

from the two reconstructions are shown in the second and third row of the figure respectively. With the minimum amount of data given, our PLS-based approach successfully recovered most of the crossings present in the slice while achieving a significant accuracy in crossing angle estimation, with an average angular error of less than 4 degrees. On the other hand, the angular analysis results for the baseline reconstruction show a higher degree of errors within the crossing regions (brighter color in Fig. 9 (e) comparing to (d)). Further, at certain locations where only a single fiber passes through, spurious lobes were introduced more often in the EAPs reconstructed using the baseline method than the proposed. This can be clearly observed in both the EAP visualizations in sub-figure (c) and the error image for number of peaks (g).

We present the results of similar analysis for noise contaminated data (SNR = 5) in Fig. 10. Here in addition to reconstructions from jointly undersampled data, we also display the reconstructed EAPs from data undersampled in the $\mathbf{k}$ and $\mathbf{q}$ -space individually using the proposed PLS approach (labeled as PLS-konly and PLS-qonly in the figure). 15% of the gold standard data was used in all the reconstructions. When joint-($\mathbf{k}$, $\mathbf{q}$) undersampling is used, both PLS and baseline EAP reconstructions present considerable visual similarity to the ground-truth EAPs, yet further angular and peak analysis reveal the subtle differences. Comparing sub-figure (f) and (g), we see that while the 90° crossings in the center of the slice appear to be slightly better recovered in the baseline results, the




Fig. 13. EAP reconstructions and angular accuracy analysis results for the ROI in the phantom dataset with SNR = 30 using the PLS and the baseline methods respectively. (a) Ground-truth EAPs estimated from gold standard data, EAP reconstructions using 5% of the data using (b) the PLS method and (c) the baseline method. (d-e) Error in the fiber orientations estimated from the reconstructions displayed in (b-c) respectively.




Fig. 14. EAP reconstructions and angular accuracy analysis results for the ROI in the phantom dataset with SNR = 10 using the PLS and the baseline methods respectively. (a) Ground-truth EAPs estimated from gold standard data, EAP reconstructions using 10% of the data using (b) the PLS method and (c) the baseline method. (d-e) Error in the fiber orientations estimated from the reconstructions displayed in (b-c) respectively.




Fig. 15. EAP reconstructions and angular analysis results for the ROI in the phantom dataset from $(\mathbf{k}, \mathbf{q})$ -space data undersampled using different schemes. (a) Ground-truth EAP estimated from gold standard data at SNR = 20, EAP reconstructions using 20% of the data undersampled in the manner of (b) joint- $(\mathbf{k}, \mathbf{q})$, (c) $\mathbf{k}$ -only and (d) $\mathbf{q}$ -only manner. Error in the fiber orientations estimated from the reconstructions displayed in (b–d) are presented in (e–g) respectively.

angular errors are more prominent in the peripheral crossing regions where the geometric structure of the fiber bundles are more complex. It is also shown in sub-figure (k) that a larger number of false peaks were introduced in the baseline reconstruction compared to PLS. Switching the focus to the results presented in the two columns on the right, we see how performing undersampling in $\mathbf{k}$ and $\mathbf{q}$ -space individually effects the proposed CS-based EAP reconstruction. Apparently, discarding a larger portion of the $\mathbf{q}$ -space data (which takes place in $\mathbf{q}$ -only undersampling compared to the joint undersampling) substantially impacts the recovery of angular information in the EAP estimations. Nonetheless, the reconstruction appears spatially smoother to some extent compared to the $\mathbf{k}$ -only reconstruction.

3.5. Analysis of the ISBI HARDI challenge 2013 phantom data

In this section, we present the evaluation of our method on the ISBI HARDI challenge 2013 phantom data, the chosen ROI for analysis was depicted in Fig. 6 in Section 3.3.

The numerical results for this dataset are presented in Figs. 11 and 12 in a similar fashion as in the case of synthetic data. Fig. 12 comprehensively depicts the performance of the proposed and baseline technique with respect to EAP reconstruction accuracy from various percentages of the gold standard data undersampled with different schemes, while Fig. 11 is dedicated to the illustration of the difference in the two methods' performance when used in conjunction with joint- $(\mathbf{k}, \mathbf{q})$ undersampling (which is the sampling scheme of interest). Overall, it appears that the deterioration in the EAP estimation accuracy caused by the corruption of noise is more substantial for this dataset than it is for the synthetic data. Yet it is consistent across these two datasets that the proposed method better maintains its accuracy when noise level increases (shown in the figure as smaller vertical distances between the two orange dash lines compared to that between the corresponding pair of lines in blue). As the sampling rate increases, the gap between the performance of the two methods gradually diminishes, especially in the less noisy data case (SNR = 10).

We further support our numerical findings via visualizations of the reconstructed EAPs and results of fiber orientation analysis. In the visualizations, the EAP profiles are superimposed on the gray scale FA map of the ROI. The fiber orientations in the reconstruc-

tions are determined at each voxel by performing peak detection using Dipy as previously mentioned. The error in the fiber orientations (in degrees) are then computed with respect to the ground-truth fiber structure and mapped to a color according to its value.

Fig. 13 showcases the methods' ability to recover the ground-truth EAP profiles and correct underlying fiber orientations from minimum amount of the relatively clean data (5% of the data with SNR = 30), undersampled jointly in the $(\mathbf{k}, \mathbf{q})$ -space. It is evident that for this particular dataset (with a very high SNR), a sampling rate of 5% is sufficient for both methods to achieve EAP reconstructions of satisfactory quality. The majority of the fiber bundle crossings in the ROI are recovered quite accurately with both methods. Yet, in some regions where two fiber bundles cross, the subordinate direction in the EAP tends to get underestimated in the baseline method. Such phenomenon can be observed within the green rectangular box in the EAP visualizations. Further, when the two fiber bundles cross at a relatively small angle, the EAP profile could potentially get smeared in the baseline reconstruction (see blue box enclosed area). Comparing the error maps for the two reconstructed fields, we see that the orientations of the fiber bundles in these above-mentioned regions are more precisely estimated when the notion of PLS is incorporated as well. In addition, the proposed method demonstrates superiority in suppressing the introduction of spurious lobes caused by the severe undersampling of the data, one example of which is highlighted with a purple/orange box in the EAP visualizations/orientation error images.

To investigate how the two methods differ when applied to data with high noise contamination, in Fig. 14 we present visual comparisons of the EAP reconstructions from 10% jointly undersampled phantom data of which the SNR is 10. As shown in the top left sub-figure, the ground-truth EAP field estimated using the full DSI dataset appear very noisy. The two reconstructed fields, on the other hand, present more spatial smoothness primarily due to the TV regularization in both frameworks. In fact, striking similarity is demonstrated in the spherical harmonic based visualizations of the two EAP fields. Nonetheless, further insights are provided by the fiber orientation analysis. Quantitatively, the fiber orientation error within the ROI is reduced by 1.3 degrees on average with the PLS method. This can be visually observed in the error maps (bottom row) as darker color in the PLS results compared to the baseline results for corresponding voxels across the field.




Fig. 16. NMSE of EAP reconstructions, using the PLS and the baseline methods respectively, from the HCP data undersampled using joint- $(\mathbf{k}, \mathbf{q})$, $\mathbf{k}$ -only and $\mathbf{q}$ -only schemes at various rates.




Fig. 17. EAP reconstructions from jointly undersampled HCP data using the PLS and the baseline methods respectively. (a) FA map of the selected slice with ROI highlighted in a red box, (d) ground-truth EAPs, (b–c) EAP reconstructions using 5% of the data using the PLS and the baseline methods respectively. (e–f) EAP reconstructions using 15% of the data using the PLS and the baseline methods respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

EAPs reconstructed using the proposed approach from 20% of the data with $\text{SNR} = 20$, undersampled in joint- $(\mathbf{k}, \mathbf{q})$ -space, $\mathbf{k}$ -space and $\mathbf{q}$ -space only are displayed in Fig. 15. Evidently from the visualizations, joint- $(\mathbf{k}, \mathbf{q})$ undersampling provides major advantage compared to individual space undersampling. The recovered EAP profiles from jointly undersampled data are of great visual similarity to the ground-truth EAPs, and the fiber orientations

estimated from the reconstructions have an average error of less than 5 degrees. Reconstructions from $\mathbf{k}$ -only undersampled data fail to preserve some of the subsidiary components in multi-fiber voxels, while $\mathbf{q}$ -space undersampling leads to serious distortions of the EAP shapes and poor accuracy in fiber orientation estimations. Regions where this above-mentioned behavior can be best observed are highlighted with rectangular boxes.

3.6. Analysis of the MGH-USC HCP data

Comprehensive numerical evaluations of the EAP reconstruction for the MGH-USC HCP data are presented in Fig. 16. It is clearly shown here that the proposed structural correlation promoting method outperforms the baseline method at every sampling rate. A similar trend that was observed in the numerical results for both the synthetic and phantom dataset can also be seen here namely, that the incorporation of the notion of parallel level sets brings forth greater reconstruction accuracy at lower sampling rates (for the case of joint undersampling). Such an attribute is of immense value in our context as it is our ultimate goal to achieve good/acceptable reconstruction at sampling rates as low as possible.

We show visual comparisons of the EAP reconstruction performance of the proposed and baseline algorithm from jointly undersampled HCP data in Fig. 17. With 5% of the data, the EAPs can be recovered at considerably good quality with the PLS method (as shown in (b)). Marginal smearing/blurring effect can be seen in the EAP profiles but one can still easily distinguish the orientations of the crossing fiber bundles. However, this effect appears far more pronounced in the baseline reconstruction, for example in the blue box region where crossings are most noticeable within the slice in the ground-truth. In certain areas where the crossings are less predominant, the EAPs reconstructed using baseline approach can be indistinguishable from the EAPs with a single fiber direction (see orange upside-down L-shaped region). Sub-figures in the second row (e–f) show that both methods are able to achieve highly accurate estimation of the EAPs from 15% of the data. The PLS method demonstrates advantage in precisely recovering the primary direction in the EAP profiles (with respect to the ground-truth), which can be observed in the blue rectangle highlighted region. (The EAPs are colored based on the orientation of the primary lobe in the profile, hence a similar color of the profile indicates a closer approximation in the orientation). These results on HCP data quantitatively and qualitatively demonstrate the value of the proposed method in producing satisfactory EAP reconstructions from a small amount of real diffusion data which can be acquired in a much shorter time than when employing conventional methods.

4. Conclusions

In this paper, we proposed a novel CS based approach for direct EAP reconstruction from heavily undersampled $(\mathbf{k}, \mathbf{q})$ -space dMR measurements utilizing the sparsity promoting surfacelet basis and in addition imposed the structural similarity constraint between the diffusion-weighted images corresponding to proximally-located $\mathbf{q}$ points. By directly reconstructing $P(\mathbf{x}, \mathbf{r})$ from $(\mathbf{k}, \mathbf{q})$ -space data, we exploit the incoherence between the 6D sensing and reconstruction domains to the fullest, which is consistent with the CS-theory. Further, by incorporating a prior favoring parallelism of level sets at corresponding spatial locations between the diffusion images into a 6D CS framework, we take advantage of the redundancy in the $(\mathbf{x}, \mathbf{q})$ -space which in this context has never been utilized in this form (employing the notion of parallel level sets) before. We presented extensive set of experiments involving synthetic, phantom and real data, demonstrating the power of the proposed method over the baseline method, wherein the structural similarity among diffusion-weighted images were not exploited. Our experiments show that the exploitation of this inherent geometric correlation considerably enhances the reconstruction accuracy and hence leads to additional savings in dMRI acquisition time over state-of-the-art. Our future work will focus on implementation of the algorithm on clinical scanners.

Acknowledgments

This research was funded in part by the ONR grant N00014-14-1-0762 and the NSF grant IIS-1617101 to Alireza Entezari, the NSF grants IIS-1525431 and IIS-1724174 and a grant from Toshiba America Medical Systems to Baba C. Vemuri.

Data were provided [in part] by the Human Connectome Project, MGH-USC Consortium (Principal Investigators: Bruce R. Rosen, Arthur W. Toga and Van Wedeen; U01MH093765) funded by the NIH Blueprint Initiative for Neuroscience Research grant; the National Institutes of Health grant P41EB015896; and the Instrumentation grants S10RR023043, S10RR023401, S10RR019307.

References

- Alaya, I.B., Jribi, M., Ghorbel, F., Kraiem, T., 2016. A novel geometrical approach for a rapid estimation of the hardi signal in diffusion mri. In: International Conference on Image and Signal Processing. Springer, pp. 253–261.
- Andersson, J.L., Sotiropoulos, S.N., 2015. Non-parametric representation and prediction of single-and multi-shell diffusion-weighted MRI data using gaussian processes. *Neuroimage* 122, 166–176.
- Andersson, J.L., Sotiropoulos, S.N., 2016. An integrated approach to correction for off-resonance effects and subject movement in diffusion mr imaging. *Neuroimage* 125, 1063–1078.
- Aranda, R., Ramirez-Manzanares, A., Rivera, M., 2015. Sparse and adaptive diffusion dictionary (SADD) for recovering intra-voxel white matter structure. *Med. Image Anal.* 26 (1), 243–255.
- Auria, A., Daducci, A., Thiran, J.-P., Wiaux, Y., 2015. Structured sparsity for spatially coherent fibre orientation estimation in diffusion MRI. *Neuroimage* 115, 245–255.
- Awate, S.P., DiBella, E.V., 2013. Compressed sensing hardi via rotation-invariant concise dictionaries, flexible k-space undersampling, and multiscale spatial regularity. In: Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on. IEEE, pp. 9–12.
- Bilgic, B., Chatnuntawech, I., Setsompop, K., Cauley, S.F., Yendiki, A., Wald, L.L., Adalsteinsson, E., 2013. Fast dictionary-based reconstruction for diffusion spectrum imaging. *IEEE Trans. Med. Imaging* 32 (11), 2022–2033.
- Callaghan, P.T., 1991. Principles of Nuclear MR Microscopy. OUP.
- Cheng, J., Shen, D., Basser, P.J., Yap, P.-T., 2015a. Joint 6d k-q space compressed sensing for accelerated high angular resolution diffusion MRI. In: 24th Biennial International Conference on Information Processing in Medical Imaging (IPMI 2015). Springer, pp. 782–793.
- Cheng, J., Shen, D., Yap, P.-T., Basser, P.J., 2015b. Tensorial spherical polar fourier diffusion MRI with optimal dictionary learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer International Publishing, pp. 174–182.
- Crawford, K.L., Neu, S.C., Toga, A.W., 2016. The image and data archive at the laboratory of neuro imaging. *Neuroimage* 124, 1080–1083.
- Daducci, A., Canales-Rodríguez, E.J., Zhang, H., Dyrby, T.B., Alexander, D.C., Thiran, J.-P., 2015. Accelerated microstructure imaging via convex optimization (amico) from diffusion MRI data. *Neuroimage* 105, 32–44.
- Daducci, A., Caruyer, E., Descoteaux, M., Thiran, J., 2013. Hardi reconstruction challenge. IEEE International Symposium on Biomedical Imaging. <http://hardi.epfl.ch/status/events>.
- Donoho, D.L., 2006. Compressed sensing. *IEEE Trans. Inf. Theory* 52 (4), 1289–1306.
- Ehrhardt, M.J., et al., 2014. Joint reconstruction of PET-MRI by exploiting structural similarity. *Inverse Probl.* 31 (1), 015001.
- Ehrhardt, M.J., Arridge, S.R., 2014. Vector-valued image processing by parallel level sets. *IEEE Trans. Image Process.* 23 (1), 9–18.
- Fan, Q., Nummenmaa, A., Witzel, T., Zanzonico, R., Keil, B., Cauley, S., Polimeni, J.R., Tisdall, D., Van Dijk, K.R., Buckner, R.L., et al., 2014. Investigating the capability to resolve complex white matter structures with high b-value diffusion magnetic resonance imaging on the MGH-USC connectome scanner. *Brain Connect.* 4 (9), 718–726.
- Fan, Q., Witzel, T., Nummenmaa, A., Van Dijk, K.R., Van Horn, J.D., Drews, M.K., Somerville, L.H., Sheridan, M.A., Santillana, R.M., Snyder, J., et al., 2016. MGH-USC human connectome project datasets with ultra-high b-value diffusion mri. *Neuroimage* 124, 1108–1114.
- Fischl, B., 2012. Freesurfer. *Neuroimage* 62 (2), 774–781.
- Garyfallidis, E., Brett, M., Amirebian, B., Rokem, A., Van Der Walt, S., Descoteaux, M., Nimmo-Smith, I., Contributors, D., 2014. Dipy, a library for the analysis of diffusion MRI data. *Front. Neuroinform.* 8.
- Goldstein, T., Osher, S., 2009. The split Bregman method for l1-regularized problems. *SIAM J. Imaging Sci.* 2 (2), 323–343.
- Greve, D.N., Fischl, B., 2009. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* 48 (1), 63–72.
- Haldar, J.P., Wedeen, V.J., Nezamzadeh, M., Dai, G., Weiner, M.W., Schuff, N., Liang, Z.-P., 2013. Improved diffusion imaging through SNR-enhancing joint reconstruction. *Magn. Reson. Med.* 69 (1), 277–289.
- Keil, B., Blau, J.N., Biber, S., Hoecht, P., Tountcheva, V., Setsompop, K., Triantafyllou, C., Wald, L.L., 2013. A 64-channel 3t array coil for accelerated brain mri. *Magn Reson Med* 70 (1), 248–258.

- van der Kouwe, A.J., Benner, T., Salat, D.H., Fischl, B., 2008. Brain morphometry with multiecho mprage. *Neuroimage* 40 (2), 559–569.
- Lam, F., Babacan, S.D., Haldar, J.P., Weiner, M.W., Schuff, N., Liang, Z.-P., 2014. Denoising diffusion-weighted magnitude mr images using rank and edge constraints. *Magn. Reson. Med.* 71 (3), 1272–1284.
- Landman, B.A., et al., 2012. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. *Neuroimage* 59 (3), 2175–2186.
- Lu, Y.M., Do, M.N., 2007. Multidimensional directional filter banks and surfacelets. *IEEE Trans. Image Process.* 16 (4), 918–931.
- Lustig, M., Donoho, D., Pauly, J., 2007. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* 58 (6), 1182–1195.
- Mani, M., et al., 2015. Acceleration of high angular and spatial resolution diffusion imaging using CS with multichannel spiral data. *Magn. Reson. Med.* 73 (1), 126–138.
- McClymont, D., Teh, I., Whittington, H.J., Grau, V., Schneider, J.E., 2016. Prospective acceleration of diffusion tensor imaging with compressed sensing using adaptive dictionaries. *Magn. Reson. Med.* 76 (1), 248–258.
- McNab, J.A., Edlow, B.L., Witzel, T., Huang, S.Y., Bhat, H., Heberlein, K., Feiweier, T., Liu, K., Keil, B., Cohen-Adad, J., et al., 2013. The human connectome project and beyond: initial applications of 300mt/m gradients. *Neuroimage* 80, 234–245.
- Merlet, S.L., Deriche, R., 2013. Continuous diffusion signal, EAP and ODF estimation via compressive sensing in diffusion MRI. *Med. Image Anal.* 17 (5), 556–572.
- Michailovich, O., Rath, Y., Dolui, S., 2011. Spatially regularized compressed sensing for high angular resolution diffusion imaging. *IEEE Trans. Med. Imaging* 30 (5), 1100–1115.
- Ning, L., Laun, F., Gur, Y., DiBella, E.V., Deslauriers-Gauthier, S., Megherbi, T., Ghosh, A., Zucchelli, M., Menegaz, G., Fick, R., et al., 2015. Sparse reconstruction challenge for diffusion MRI: validation on a physical phantom to determine which acquisition scheme and analysis method to use? *Med. Image Anal.* 26 (1), 316–331.
- Ning, L., Setsompop, K., Michailovich, O., Makris, N., Shenton, M.E., Westin, C.-F., Rath, Y., 2016. A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging. *Neuroimage* 125, 386–400.
- Paquette, M., Merlet, S., Gilbert, G., Deriche, R., Descoteaux, M., 2015. Comparison of sampling strategies and sparsifying transforms to improve compressed sensing DSI. *Magn. Reson. Med.* 73 (1), 401–416.
- Polimeni, J.R., Bhat, H., Witzel, T., Benner, T., Feiweier, T., Inati, S.J., Renvall, V., Heberlein, K., Wald, L.L., 2016. Reducing sensitivity losses due to respiration and motion in accelerated echo planar imaging by reordering the autocalibration data acquisition. *Magn. Reson. Med.* 75 (2), 665–679.
- Schwab, E., Vidal, R., Charon, N., 2016. Spatial-angular sparse coding for HARDI. In: *MICCAI* (3), pp. 475–483.
- Setsompop, K., Kimmlingen, R., Eberlein, E., Witzel, T., Cohen-Adad, J., McNab, J.A., Keil, B., Tisdall, M.D., Hoecht, P., Dietz, P., et al., 2013. Pushing the limits of in vivo diffusion MRI for the human connectome project. *Neuroimage* 80, 220–233.
- Shi, X., Ma, X., Wu, W., Huang, F., Yuan, C., Guo, H., 2015. Parallel imaging and compressed sensing combined framework for accelerating high-resolution diffusion tensor imaging using inter-image correlation. *Magn. Reson. Med.* 73 (5), 1775–1785.
- Sun, J., Sakhaee, E., Entezari, A., Vemuri, B.C., 2015. Leveraging eap-sparsity for compressed sensing of MS-HARDI in (k, q)-space. In: *24th Biennial International Conference on Information Processing in Medical Imaging (IPMI 2015)*, pp. 375–386.
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.-M.H., et al., 2013. The WU-MINN human connectome project: an overview. *Neuroimage* 80, 62–79.
- Yin, S., You, X., Xue, W., Li, B., Zhao, Y., Jing, X.-Y., Wang, P.S., Tang, Y., 2016. A unified approach for spatial and angular super-resolution of diffusion tensor mri. In: *Chinese Conference on Pattern Recognition*. Springer, pp. 312–324.