

Contents lists available at ScienceDirect

MEDICAL IMAGE ANALYSIS

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Joint spatial-angular sparse coding for dMRI with separable dictionaries



Evan Schwab*, René Vidal, Nicolas Charon

Center for Imaging Science, Johns Hopkins University, Baltimore, MD, USA

ARTICLE INFO

Article history: Received 28 April 2017 Revised 1 May 2018 Accepted 7 May 2018

Keywords: Sparse coding Separable dictionaries Kronecker product Diffusion MRI

ABSTRACT

Diffusion MRI (dMRI) provides the ability to reconstruct neuronal fibers in the brain, in vivo, by measuring water diffusion along angular gradient directions in q-space. High angular resolution diffusion imaging (HARDI) can produce better estimates of fiber orientation than the popularly used diffusion tensor imaging, but the high number of samples needed to estimate diffusivity requires longer patient scan times. To accelerate dMRI, compressed sensing (CS) has been utilized by exploiting a sparse dictionary representation of the data, discovered through sparse coding. The sparser the representation, the fewer samples are needed to reconstruct a high resolution signal with limited information loss, and so an important area of research has focused on finding the sparsest possible representation of dMRI. Current reconstruction methods however, rely on an angular representation per voxel with added spatial regularization, and so, for non-zero signals, one is required to have at least one non-zero coefficient per voxel. This means that the global level of sparsity must be greater than the number of voxels. In contrast, we propose a joint spatial-angular representation of dMRI that will allow us to achieve levels of global sparsity that are below the number of voxels. A major challenge, however, is the computational complexity of solving a global sparse coding problem over large-scale dMRI. In this work, we present novel adaptations of popular sparse coding algorithms that become better suited for solving large-scale problems by exploiting spatial-angular separability. Our experiments show that our method achieves significantly sparser representations of HARDI than is possible by the state of the art.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Diffusion magnetic resonance imaging (dMRI) is a medical imaging modality used to analyze neuroanatomical biomarkers for brain diseases such as Alzheimer's. dMRI are 6D signals consisting of a set of 3D spatial MRI volumes acquired in k-space that are each weighted with a different diffusion signal measured in q-space. In each voxel of a brain dMRI, the q-space diffusion signals are reconstructed to estimate orientations and integrity of neuronal fiber tracts, in vivo. Different dMRI protocols measure qspace in different ways. For example, diffusion spectrum imaging (DSI) (Wedeen et al., 2005) measures q-space densely on a 3D grid. Alternatively, diffusion tensor imaging (DTI) (Basser et al., 1994) simplifies acquisition by modeling a Gaussian distribution on the unit *q*-sphere. High angular resolution diffusion imaging (HARDI) (Tuch, 2004) also restricts measurements to the unit sphere, but increases the angular resolution from that of DTI. Multi-Shell HARDI (MS-HARDI) (Wu and Alexander, 2007) expands its radial

* Corresponding author. E-mail address: eschwab3@jhu.edu (E. Schwab).

https://doi.org/10.1016/j.media.2018.05.002 1361-8415/© 2018 Elsevier B.V. All rights reserved. range to include multiple spheres, or shells. Since DTI collects the fewest number of measurements, it has become the most widely used clinical dMRI protocol. However, its simple tensor model is unable to capture the complex diffusion profiles in each voxel. On the other hand, protocols like HARDI, MS-HARDI, and especially DSI, collect a higher number of *q*-space measurements to estimate more accurate diffusion profiles at the expense of longer scan times, making them currently unsuitable for clinical studies.

An ongoing research goal has been to find ways to reduce acquisition times of HARDI, MS-HARDI, or DSI, while maintaining accurate estimations of diffusion. One avenue is from a hardware perspective: maintain dense signal measurement configurations while devising faster physical acquisition techniques like simultaneous multi-slice acquisition (Setsompop et al., 2012) and simultaneous image refocusing (Reese et al., 2009). The other is from a signal processing perspective: maintain accurate signal reconstructions while devising methods to exploit redundancies in the data and reduce the number of required measurements to accelerate acquisition. This paradigm is known as Compressed Sensing (CS) (Donoho et al., 2006).

CS is a class of mathematical results and algorithms that exploits sparse representations of signals, discovered through sparse

coding, to obtain extremely accurate reconstructions at sub-Nyquist rates. A classical application of CS has been to accelerate structural MRI by subsampling the spatial frequency domain, k-space (Lustig et al., 2007), known as k-space CS or k-CS. These ideas have also been previously applied to dMRI by subsampling the angular frequency domain, *q*-space, (Ning and et al., 2015) (analogously called q-CS) and more recently, to subsample both k- and q-space (Cheng et al., 2015a; Sun et al., 2015), commonly called (k, q)-space CS or (k, q)-CS, to further increase acceleration. However, because the goal of dMRI reconstruction is to estimate diffusivity profiles at each voxel, dMRI signals are traditionally represented as a set of voxel-wise q-space signals in the angular domain. Spatial regularization is an important technique used to improve these estimations over an entire dMRI volume (Goh et al., 2009), but the underlying data representation of dMRI is still angular and local to each voxel. Therefore, when applying sparse coding for dMRI, the sparsest possible global representation over an entire volume can be no less than the number of voxels since at least one dictionary atom would be required to model *q*-space signals in each voxel.

To overcome this fundamental limitation, we propose a global spatial-angular representation of dMRI that allows global sparsity levels to fall below one atom per voxel by exploiting redundancies in the spatial and angular domains, jointly with a global dictionary. A major challenge, however, of optimizing over a global dictionary is the computational complexity of solving a massive global sparse coding problem over large-scale dMRI data. Yet, by imposing that our global dictionary is separable over the spatial and angular domains we can greatly improve computational efficiency while preserving good sparsity levels for typical signals. One of our main contributions in this paper is a set of novel adaptations of popular sparse coding algorithms to solve general largescale sparse coding problems using separable dictionaries. Our experiments on phantom and real HARDI brain data show that it is possible to achieve accurate global HARDI reconstructions with a sparse representation of less than one dictionary atom per voxel, exceeding the theoretical limit of the state of the art in sparse coding. Sparse coding has many important applications like de-noising (Ouyang et al., 2013), dictionary learning (Cheng et al., 2015b) and super-resolution (Yoldemir et al., 2014), and, in particular, applying our joint spatial-angular sparse coding framework within the application of (k, q)-CS will be the subject of future work.

The remainder of this paper is organized as follows: In Section 2, we review state-of-the-art sparse coding methods for dMRI and illustrate the limitations of their performance on a phantom HARDI dataset. In Section 3, we present our joint spatial-angular dMRI representation and formalize the global spatial-angular sparse coding problem. Then, in Section 4, we develop and compare a set of novel sparse coding algorithms using separable dictionaries to efficiently solve our large-scale global optimization. Finally, in Section 5 we provide experimental results showing the performance of our method over the state-of-the-art and conclude with a discussion in Section 6.

2. State of the art

2.1. Angular (voxel-wise) reconstruction

A dMRI can be modeled as a 6D signal S(v, q), where $v \in \Omega \subset \mathbb{R}^3$ is the location of a voxel in the 3D spatial domain Ω and $q \in \mathbb{R}^3$ is a point in the so-called *q*-space.¹ A dMRI signal is measured at a discrete number of voxels, *V*, and a discrete number of *q*-space points, *G*. While dMRI signals can be viewed as a set of

G diffusion weighted images (DWIs) or volumes, the most common view-point for dMRI processing and analysis is voxel-wise, i.e.for each voxel $v \in \Omega$, we acquire a vector of *G* diffusion measurements $S(v, q_g)_{g=1}^G = s_v (q_g)_{g=1}^G = s_v$ at points q_g in 3D *q*-space. The latter interpretation is most common for modeling because a major goal of dMRI reconstruction is to estimate 3D probability distribution functions (PDFs) of fiber tract orientation at each voxel. Accordingly, the signal vector s_v is represented by a *q*-space basis, $\Gamma = [\Gamma_i(q)]_{i=1}^{N_{\Gamma}}$, with N_{Γ} atoms, such that

$$s_{\nu} = \Gamma a_{\nu}.\tag{1}$$

where a_v is the vector of angular coefficients at voxel v. (See Fig. 1). The dMRI literature has produced a wide array of dMRI reconstruction algorithms for different acquisition protocols, an artillery of q-space bases and varying models for estimating orientation distributions. The vast majority of research reconstructs q-space signals in each voxel with a q-space basis (see the dMRI challenge (Daducci et al., 2014) for a comprehensive summary and comparison of state of the art reconstruction frameworks). To enforce or exploit desirable properties of dMRI signals, many methods will add a set of constraints C on the angular coefficients such as angular smoothing (Ye, 2016), non-negativity of PDFs (Schwab et al., 2012; Wolfers et al., 2014), or orientational symmetry (Gramfort et al., 2014), solving:

$$a_{\nu}^{*} = \arg\min_{a_{\nu}} \frac{1}{2} ||\Gamma a_{\nu} - s_{\nu}||_{2}^{2} \quad \text{s.t.} \quad a_{\nu} \in \mathcal{C}.$$
⁽²⁾

The constraint of particular interest in our paper is that of enforcing sparsity on the coefficients of the reconstruction, known as *Sparse Coding*.

2.2. Angular (voxel-wise) sparse coding

Sparse coding is a reconstruction problem which seeks a sparse representation, i.e. a coefficient vector with few nonzero elements. Given a sparsifying *q*-space basis Γ for which the dMRI signal in each voxel is expected to have a sparse representation, the angular (voxel-wise) sparse coding problem can be formulated as:

$$a_{\nu}^{*} = \arg\min_{a_{\nu}} \frac{1}{2} ||\Gamma a_{\nu} - s_{\nu}||_{2}^{2} \quad \text{s.t.} \quad ||a_{\nu}||_{0} \le K_{\nu},$$
(3)

where $||a_{\nu}||_{0}$ counts the number of nonzero elements of vector a_{ν} , and K_{ν} is the sparsity level at voxel ν . This problem is known to be NP-hard, and therefore the two main methodologies to tackle (3) are to *a*) approximate a solution using greedy algorithms such as Orthogonal Matching Pursuit (OMP) (Tropp, 2004) or *b*) replace the L_{0} semi-norm by its convex relaxation, the L_{1} norm, and solve either the Basis Pursuit or LASSO problem given by:

$$a_{\nu}^{*} = \arg\min_{a_{\nu}} \frac{1}{2} ||\Gamma a_{\nu} - s_{\nu}||_{2}^{2} + \lambda ||a_{\nu}||_{1}$$
(4)

using algorithms such as Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2010) or Fast Iterative Thresholding Algorithm (FISTA) (Beck and Teboulle, 2009), where λ is the trade-off parameter between data fidelity and sparsity.

An important application of sparse coding in the dMRI community is that of CS. Angular sparse coding and *q*-CS have been widely researched for dMRI to reduce long acquisition times. Many groups have done extensive work choosing sparsifying *q*space bases (Merlet et al., 2011; Ning and et al., 2015; Aranda et al., 2015), developing dictionary learning methods (Bilgic et al., 2012; Cheng et al., 2013; 2015b; Gramfort et al., 2012; 2014; Merlet et al., 2012; 2013; Sun et al., 2013), and testing *q*-space

¹ The *q*-space is the frequency domain associated with the angular domain, while the *k*-space is the frequency domain associated with the spatial domain.



Fig. 1. Illustration of voxel-wise angular HARDI representation a_v using a sparsifying dictionary Γ .

subsampling schemes for DSI (Gramfort et al., 2014; Merlet and Deriche, 2010; Menzel et al., 2011; Merlet and Deriche, 2013; Paquette et al., 2015), MS-HARDI (Cheng et al., 2011; Rathi et al., 2011; Merlet et al., 2013; Duarte-Carvajalino et al., 2014; Daducci et al., 2015), HARDI (Michailovich and Rathi, 2010a; 2010b; Tristán-Vega and Westin, 2011; Duarte-Carvajalino et al., 2012; Alaya et al., 2016) and DTI (Landman et al., 2012) with promising results in sparsity and measurement reduction for clinical tractography (Kuhnt et al., 2013a; 2013b). However, a major limitation for this family of methods is that the sparsest possible representation of an entire dMRI dataset can be no less than the number of voxels since $||a_{\nu}||_{0} \geq 1 \forall \nu \in \Omega$. In CS applications, this induces fundamental limitations in the amount of subsampling factors that may be achievable in q-space. In practice, the spatial-angular sparsity level will be much greater than the number of voxels, to account for noise. For example the work of Michailovich and Rathi (2008, 2010b) report an average sparsity level of 6-10 atoms per voxel. The methods presented in the next section attempt to improve upon these results by exploiting spatial redundancies and reducing measurements in *k*-space.

2.3. Angular sparse coding with spatial regularization

Incorporating spatial information into voxel-wise reconstruction is a well utilized technique for increasing the accuracy of reconstruction. The following is a general formulation for including spatial regularization into the angular sparse coding problem:

$$A^{*} = \arg\min_{A} ||\Gamma A - S||_{F}^{2} + \lambda ||A||_{1} + \mathcal{R}(A),$$
(5)

where $S = [s_1 \dots s_V] \in \mathbb{R}^{G \times V}$ is the concatenation of signals $s_v \in \mathbb{R}^G$ sampled at *G* gradient directions over *V* voxels, $A = [a_1 \dots a_V] \in \mathbb{R}^{N_{\Gamma} \times V}$ is the concatenation of angular coefficients and $\mathcal{R}(A)$ is a spatial regularizer that depends on the angular representation *A*. Here $||X||_F = \sqrt{\sum_i \sum_j |X_{i,j}|^2}$ is the Frobenius norm and $||X||_1 = \sum_i \sum_j |X_{i,j}|$ is the 1-norm taken over all elements of the matrix. In particular when $\mathcal{R} = 0$, this reduces to solving (4) for each voxel. When $\lambda = 0$ and $\mathcal{R}(A) = \sum_i \sum_{j \in \mathcal{N}_i} ||a_i - a_j||^2$ (Laplacian regularization), where \mathcal{N}_i is a local spatial neighborhood of voxel *i*, this is the general non-sparse reconstruction with spatial coherence (Goh et al., 2009). Some have found incorporating both the angular sparsity constraint $\lambda ||A||_1$ and spatial coherence $\mathcal{R}(A)$ beneficial for applications such as de-noising (Ye et al., 2012; Ouyang et al., 2013; 2014; Ye, 2016) and tractography (Ye and Prince, 2017).

Spatial regularization within sparse coding is more prominently used for the application of reducing redundancies for CS. For example, (Michailovich et al., 2011; Rathi et al., 2014; Auría et al., 2015) enforce spatial smoothing for q-CS while (Ning et al., 2016; Yin et al., 2016) combine q-CS with super-resolution reconstruction of the spatial domain. To further accelerate dMRI, the recent work of Chao et al. (2017) combines CS with parallel imaging but reconstructs the signals in k-space and q-space separately in sequence. A joint (k, q)-space reconstruction is important for maintaining coherence throughout the dataset. As such, the works of Awate and DiBella (2013), Shi et al. (2015), Cheng et al. (2015a), Sun et al. (2015), McClymont et al. (2015) and

Mani et al. (2015) combine k- and q-CS by adding a data fidelity term for k-space subsampling and an additional spatial sparsity term. In total, however, while each of these works may be applied to different diffusion models and acquisition protocols testing various subsampling schemes, sparsifying transforms and dictionaries, each are based on an angular representation of dMRI data, A. In fact, they stem from the same optimization problem formulation (5) with

$$\mathcal{R}(A) = \alpha ||\mathcal{T}(\Gamma A)||_1, \tag{6}$$

where $\alpha \ge 0$ is an additional trade-off weighting parameter, and $\mathcal{T}(\cdot)$ is a sparsifying transform (or dictionary) applied to the spatial domain such as wavelets or the finite difference gradient operator, leading to the usual total variation (TV) norm. In (6), ΓA is a reconstruction of the signal *S* based on the angular representation *A*.

While adding these spatial and angular sparsity terms may exploit redundancies in both the spatial and angular domains, because they are separate disjoint terms the minimal global sparsity level will be still limited by the size of the data since $||A||_0$ should be greater than V and $||\Psi(\Gamma A)||_0$ should be greater than G. Indeed, when $||A||_0 < V$, there must exist voxels v such that $a_v = 0$, leading to a zero valued signal s_v (column of S) in that voxel. Likewise, when $||\Psi(\Gamma A)||_0 < G$, there must exist some gradient directions, q_g , such that the signal in the entire volume s_q (rows of S) equals zero. This becomes a problem because zero valued signals are not physically representative of real dMRI data. This also becomes a heuristic limitation of prior methods for appropriately choosing trade-off parameters λ and ξ that result in a physically accurate sparsity level.

In the next section we will explicitly show the limitation of sparsity on phantom HARDI data. Table 1 organizes the recent literature's usages of sparse coding and CS for dMRI and places the proposed work in context compared to the state of the art. There we use the term "Spatial + Angular" Sparse Coding to emphasize that the state of the art perform both spatial and angular sparse coding, but not jointly. As an important note, though we frame our proposed sparse coding method in Table 1 with the backdrop of CS, we do not propose or implement CS in the current manuscript.

2.4. Limitations of angular representations for sparse coding

We illustrate the limitations of sparse coding using a per-voxel angular representation on a HARDI phantom dataset with $V = 50 \times 50$ and G = 64 gradient directions (the same data is used in our experiments in Section 5). First, we solve (5) with $\mathcal{R}(A) = 0$, showing qualitative reconstruction results in Fig. 2, for various sparsity levels given by the value of λ . Our second result considers the effect of spatial regularization $\mathcal{R}(A) \neq 0$ on the amount of angular sparsity as a function of the reconstruction error in Fig. 4.

For this setting, we choose angular basis Γ to be the well performing overcomplete spherical ridglet (SR) dictionary (Michailovich and Rathi, 2008; 2010a; Tristán-Vega and Westin, 2011). Fig. 2 shows the ODF estimations (computed using the spherical wavelets (Tristán-Vega and Westin, 2011)) from the sparse signal reconstruction for various sparsity levels compared

Table 1

Summary of the state-of-the-art dMRI sparse reconstruction methods organized by domains of sparsity (spatial,angular) and CS subsampling (k, q). The literature has provided a natural extension from k-CS in MRI using spatial sparse coding to q-CS in dMRI angular sparse coding. However, for (k, q)-CS, the state of the art enforce sparsity in the spatial and angular domains separately, (called "Spatial + Angular" Sparse Coding) with a purely angular representation. In contrast, the proposed work considers a joint spatial-angular sparse coding framework is intended for the application of (k, q)-CS as illustrated by this table, the work presented in this paper is only for sparse coding.)

| | | Sparse coding | | | | |
|----|--------|-----------------------|--|--|-----------------------|--|
| | | Spatial | Angular | Spatial + Angular | Joint Spatial-Angular | |
| | k | (Lustig et al., 2007) | | | | |
| cs | q | | (Paquette et al., 2015) (Auría et al., 2015) (Ning and et al., 2015) (Deng et al., 2015b) (Daducci et al., 2015) (Aranda et al., 2015) (Alaya et al., 2016) (Ning et al., 2016) (Yin et al., 2016) | | | |
| | (k, q) | | | (Shi et al., 2015) (Cheng et al., 2015a) (Sun et al., 2015) (McClymont et al., 2015) (Mani et al., 2015) | Proposed† | |

to the ODFs estimated from the original signal, as well as closeups of a region of interest (ROI) containing ODFs with complex crossings of 2, 3 and 4 fibers. In order to compare spatial-angular sparsity levels we are interested in the average number of active dictionary atoms over all voxels, i.e. $||A||_0/V$. We use 5 different values of λ which gives us average spatial-angular sparsity levels of 0.246, 0.485, 1.11, 2.07, and 3.84 atoms per voxel. As expected, when $||A||_0/V < 1$ (see top left/middle), many voxels are forced to zero (as indicated by yellow spheres in Fig. 2). This is especially true for isotropic signals surrounding the fiber tracts. Also as expected, when $||A||_0/V \approx 1$, (see top right) many of the complex signals in the fiber crossing ROI are pushed to zero. This model requires close to $||A||_0/V = 4$ average atoms per voxel to achieve nearly accurate signal reconstruction (bottom middle). In fact, the actual number of coefficients per voxel to accurately represent typical dMRI data with angular bases is substantially higher. We illustrate this in Fig. 3 which shows the number of atoms used to represent the HARDI signals in each voxel for the reconstructions in Fig. 2. The bottom right image shows the ground truth number of fibers crossing in each voxel. This experiment demonstrates that voxels containing crossing fibers are forced to zero atoms when the average number of atoms per voxel is very small and tend to 6-12 atoms for accurate reconstruction when the sparsity level is decreased. This is consistent with the reports of Michailovich and Rathi (2008, 2010b) for the SR dictionary.

Next, we explore the effect of adding spatial regularization \mathcal{R} to the angular sparsity penalty, as a function of the reconstruction error. As a common spatial regularizer used in the literature, we consider for \mathcal{T} in (6) the finite difference (gradient) operator $\mathcal{T} = \nabla := [\partial_x, \partial_y, \partial_z]$ and the corresponding isotropic TV norm given by $||\nabla(X)||_{2,1} = ||\sqrt{|\partial_x X|^2 + |\partial_y X|^2 + |\partial_z X|^2}||_1$. This leads to the new optimization problem

$$A^{*} = \underset{A}{\arg\min} ||\Gamma A - S||_{F}^{2} + \lambda ||A||_{1} + \alpha ||\nabla(\Gamma A)||_{2,1},$$
(7)

for various λ and $\alpha \ge 0$, the relative weight of spatial regularization. This can be solved using Split-Bregman as in Michailovich et al. (2011). Fig. 4 shows the effect of nonzero α on angular sparsity compared to the case of $\alpha = 0$ ($\mathcal{R} = 0$) on a small 30×30 segment of the phantom HARDI data. As we can see, in all cases, the minimal sparsity for accurate reconstruction does not go below the limit of 5 atoms per voxel. In addition, increasing the relative weight of the TV norm spatial regularization actually results in an increase in angular sparsity for a given reconstruction error. In a sense, this is not surprising since the additional regularizer \mathcal{R} will enforce spatial smoothness of the reconstructed signal (which can be beneficial for noisy data and in compressed sensing scenarios) but cannot improve the resulting sparsity of the solution which is still represented by a set of coefficients per voxel in the angular basis Γ . As the goal of this paper is sparse coding, i.e finding sparsest possible representations of full HARDI data, in our later experimental comparisons, we will be using $\mathcal{R} = 0$ when referring to state-of-the-art reconstruction.

In the following section, we present our global spatial-angular representation of dMRI which allows for spatial regularization with the possibility to achieve accurate reconstruction at sparsity levels below the number of voxels, unachievable with an angular representation alone.

3. Joint spatial-angular dMRI representation

To overcome the sparsity limits of an angular representation, we propose to model a dMRI signal $S : \Omega \times \mathbb{R}^3 \to \mathbb{R}$ globally with a joint spatial-angular dictionary, say $\varphi(v, q)$, such that

$$S(v,q) = \sum_{k} c_k \varphi_k(v,q)$$
(8)

with a single set of global coefficients $c = [c_k]$. A global dictionary allows us to find global representations with sparsity levels below the number of voxels without forcing some voxels to have zero signal. In fact, the sparsest possible representation would be the absolute limit of 1 nonzero coefficient c_k , and so we find ourselves in a unrestricted setting for global sparse coding. To set up the spatial-angular sparse coding problem, we let $s \in \mathbb{R}^{GV}$ be the vectorization of S(v, q) where for $v = 1 \dots V$ we stack the *q*-space signals, $s_v \in \mathbb{R}^G$, and $\Phi_k \in \mathbb{R}^{GV}$ be the vectorization $\varphi_k(v, q)$ to build the global dictionary $\Phi = [\Phi_1 \dots \Phi_{N_{\Phi}}] \in \mathbb{R}^{GV \times N_{\Phi}}$, with N_{Φ} atoms. Then, to find a globally sparse *c*, we can solve the L_0 minimization problem:

$$c^* = \arg\min_{c} \frac{1}{2} ||\Phi c - s||_2^2 \quad \text{s.t.} \quad ||c||_0 \le K,$$
 (P0vec)

for a sparsity level *K* or the LASSO problem:

$$c^* = \arg\min_{c} \frac{1}{2} ||\Phi c - s||_2^2 + \lambda ||c||_1, \qquad (P1vec)$$



Fig. 2. Qualitative demonstration of state-of-the-art sparse coding limitations (5) with the spherical ridgelets (SR) dictionary for 5 different spatial-angular sparsity levels compared to the original signal (bottom right) with ROI closeups underneath. For high spatial-angular sparsity levels (top left, middle), voxels with complex signals are forced to zero (yellow spheres). Regions with crossing fibers are unable to be accurately reconstructed even when using an average of 2.07 atoms/voxel. The label I-SR refers to Identity-SR, explained in the next section. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

where $\lambda > 0$ is the sparsity trade-off parameter. However, typical dMRI contains on the order of $V \approx 100^3$ voxels each with $G \approx 100$ *q*-space measurements for a total of $100^4 = 100$ million signal measurements ($|s| \approx 10^8$). Since many sparse coding applications often utilize dictionaries that are over-redundant, this leads to a massive matrix Φ with 100^4 rows and over 100^4 columns ($|\Phi| \approx 10^{16}$).

For some datasets, even committing Φ to memory is prohibitive. Therefore solving this large-scale global dMRI sparse coding problem using traditional solvers like OMP to approximate (*POvec*) or ADMM and FISTA to solve (*P1vec*), prove intractable.

To address this challenge, we introduce additional structure on the dictionary atoms by considering separable functions over Ω



Fig. 3. Number of atoms found in each voxel corresponding to the 5 levels of spatial-angular sparsity in Fig. 2. The bottom right figure shows the ground truth number of fibers crossing in each voxel to illustrate the complexity of each angular signal in relation to how many atoms are needed to sparsely model them. Crossing fiber signals are either forced to zero for high spatial-angular sparsity levels (see: top row) or require between 3–5 atoms for single fiber signals (see: avg. sparsity 1.11 and 2.07) and 6–12 for double and triple crossing fiber signals (see: avg. sparsity 3.83). The label I-SR refers to Identity-SR, explained in the experiments Section 5.



Fig. 4. Reconstruction error vs. the average number of angular dictionary atoms per voxel using spatial regularization for the HARDI phantom data. As α , the relative weight of spatial regularization (TV) in (7), increases, the average number of angular atoms increases for a given reconstruction error. This suggests that sparser solutions for angular sparse coding can be achieved without spatial regularization, although using adequate spatial regularizers can improve the qualitative aspect of the reconstructed signal, in particular for noisy inputs.

and \mathbb{R}^3 , namely a set of atoms of the form $\{\varphi_k(v, q)\} = \{\psi_j(v) \otimes \gamma_i(q)\}$, where $\{\psi_j(v)\}$ is a spatial basis for the space of functions from $\Omega \to \mathbb{R}$ and $\{\gamma_i(q)\}$ is an angular basis for the space of functions from $\mathbb{R}^3 \to \mathbb{R}$ and \otimes is the Kronecker product. In discretized form for *V* voxels and *G* gradient directions, with $\Psi \in \mathbb{R}^{V \times N_{\Psi}}$ and $\Gamma \in \mathbb{R}^{G \times N_{\Gamma}}$, the matrix $\Phi = \Psi \otimes \Gamma$ is of the form:

$$s = \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_V \end{pmatrix} = \begin{pmatrix} \Psi_{1,1}\Gamma & \Psi_{1,2}\Gamma & \cdots & \Psi_{1,N_{\Psi}}\Gamma \\ \Psi_{2,1}\Gamma & \Psi_{2,2}\Gamma & \cdots & \Psi_{2,N_{\Psi}}\Gamma \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_{V,1}\Gamma & \Psi_{V,2}\Gamma & \cdots & \Psi_{V,N_{\Psi}}\Gamma \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_{N_{\Psi}N_{\Gamma}} \end{pmatrix} = \Phi c.$$
(9)

Fig. 5 illustrates the Kronecker structure of spatial-angular atom Φ_k . We can see that by representing a dMRI signal with this type of global spatial-angular atom, one can model an entire region of

the brain with as few as a single atom instead of angular atoms at every voxel.

A motivating model for this separable structure for dMRI is as follows: first, as is traditionally done, the signal at each voxel $v \in \Omega$ is written as a linear combination of angular basis functions { $\Gamma_i(q)$ }:

$$S(v,q) = \sum_{i=1}^{N_{\Gamma}} a_i(v) \Gamma_i(q).$$
(10)

Then, we notice that each spherical coefficient $a_i(v)$ forms a 3D volume and so can be written as a linear combination of spatial basis functions { $\Psi_i(v)$ }:

$$a_{i}(\nu) = \sum_{j=1}^{N_{\Psi}} c_{i,j} \Psi_{j}(\nu).$$
(11)



Fig. 5. Top: A separable spatial-angular dictionary composed of the Kronecker product between curvelets Ψ and spherical ridgelets Γ . A pair of spatial and angular atoms are highlighted in red and zoomed in below. Bottom: An example construction of a single spatial-angular basis atom Φ_k (right) by taking the Kronecker product of Ψ_j (left) and Γ_i (middle), i.e. $\Psi_j \otimes \Gamma_i = \Phi_k$. With this particular combination of spatial (curvelet (Candès et al., 2006)) and angular (spherical wavelet (Tristán-Vega and Westin, 2011)) atoms, we can see that it may be possible to represent an entire fiber tract with very few spatial-angular atoms. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

Sparse coding variable dimensions, where G (\approx 100) is the number of gradient directions in q-space, V (\approx 100³) is the number of voxels in the volume, N_{Γ} (\gtrsim 100) is the number of atoms of the angular dictionary Γ , and N_{Ψ} (\gtrsim 100³) is the number of atoms of the spatial dictionary Ψ .

| | Signal | | Coefficients | | Dictionaries | | |
|------------------------|---------|-------------------|---------------------------|--|---|---------------------------------|-------------------------|
| Variable Dimensions | s GV | $S \\ G \times V$ | c $N_{\Gamma}N_{\Psi}$ | $\begin{array}{c} C \\ N_{\Gamma} \times N_{\Psi} \end{array}$ | $ \Phi \\ GV \times N_{\Gamma} N_{\Psi} $ | $\Gamma \\ G \times N_{\Gamma}$ | $\Psi V 	imes N_{\Psi}$ |

Combining (10) and (11) we arrive at our proposed separable spatial-angular dictionary

$$\mathcal{S}(\nu,q) = \sum_{i=1}^{N_{\Gamma}} \sum_{j=1}^{N_{\Psi}} c_{i,j} \Psi_j(\nu) \Gamma_i(q), \qquad (12)$$

When stacking each s_v in a large vector, (12) results in the Kronecker product in (9), $s = (\Psi \otimes \Gamma)c$. Alternatively, when writing $S = [s_1, \ldots, s_V]$ as a matrix, (12) results in the equivalent matrix form:

$$S = \Gamma C \Psi^{\top}.$$
 (13)

Table 2 summaries the dimensions of the vector and matrix variables and Fig. 6 illustrates the Kronecker decompositions in the vector and matrix forms.

Decomposing signals into Kronecker (or more general multitensor) structures has been well researched to increase algorithmic efficiency by reducing computations to the smaller, separate domains. Many research groups have exploited properties of the Kronecker product, when solving problem types of the form of (*P0vec*) and (*P1vec*) for computational efficiency of larger sparse coding (Caiafa and Cichocki, 2013), dictionary learning (Hawe et al., 2013) and CS (Duarte and Baraniuk, 2012) applications. The work of Caiafa and Pestilli (2015) has applied multi-tensor sparse cod-



Fig. 6. Equivalent vector form (top) and matrix form (bottom) for the Kronecker decomposition of a signal. We propose to use the matrix form which provides a more compact representation for signals of large size and exploits the full separability of the Kronecker product, reducing matrix multiplication complexity from $O(GVN_{\Gamma}N_{\Psi})$ to $O(GVN_{\Gamma})$.

ing methods on dMRI data for the application of fiber tract data compression. In particular, a Kronecker Orthogonal Matching Pursuit (Kron-OMP) (Rivenson and S., 2009) has been utilized to solve (*P0vec*). Although Kron-OMP becomes much more efficient than the classical OMP (Tropp, 2004), the problem is not entirely separated into smaller domains, and the computationally expensive Φ matrix is still built explicitly. For large-scale problems like that of dMRI reconstruction, solving (*P0vec*) or (*P1vec*) even with a Kronecker structure dictionary remains largely intractable/expensive for memory and computation time.

In this chapter, we use the matrix form (13) which allows us to avoid the expensive uses of Φ and fully reduce computational complexity to the smaller separable basis domains of Γ and Ψ . In particular, we develop efficient algorithms to solve the completely separable spatial-angular sparse coding problems:

$$C^* = \arg\min_{C} \frac{1}{2} ||\Gamma C \Psi^{\top} - S||_F^2 \text{ s.t. } ||C||_0 \le K$$
 (P0mat)

and

$$C^* = \arg\min_{C} \frac{1}{2} ||\Gamma C \Psi^{\top} - S||_F^2 + \lambda ||C||_1.$$
 (P1mat)

This becomes a general optimization to solve large-scale sparse coding problems with separable dictionaries and can also be extended to the tensor setting.

As an important note, this matrix formulation is a generalization of the voxel-wise angular sparse coding problem (5) in the special case of $\Psi = I_V$, the $V \times V$ identity matrix, with $C \equiv A$. We use the identity as a choice for Ψ in the experiments of Section 5 when comparing the performance of purely angular sparse coding with our proposed framework.² Note that the optimization problem (*P1mat*) is on the coefficient matrix *C* of size $N_{\Gamma} \times N_{\Psi}$ in comparison to the matrix *A* of size $N_{\Gamma} \times V$ of stateof-the-art sparse coding (5). This leads to a slight increase in the dimension of the problem proportional to the redundancy factor of the spatial dictionary Ψ (i.e the ratio N_{Ψ}/V). On the other hand, our formulation only involves a single sparsity penalty in comparison to a sum of angular and spatial terms, thus reducing the number of weighting parameters to select.

4. Efficient kronecker sparse coding algorithms

In what follows we present three novel adaptations of existing sparse coding algorithms for solving large-scale sparse coding problems with a Kronecker dictionary structure. These are Kronecker extensions of OMP (Section 4.2), ADMM (Section 4.3), and FISTA (Section 4.5). We compare these to existing Kronecker sparse coding algorithms, a Kronecker OMP (Rivenson and S., 2009) (Section 4.1) as well as Kronecker Dual ADMM (Section 4.4), developed in our prior work Schwab et al. (2016) and derived in a new formulation for comparison in this paper. We compare these algorithms in terms of complexity for various types of bases in Section 4.6 and show experimental time comparisons in Section 5.

4.1. Kronecker OMP

To approximate a solution to the L_0 problem (*Povec*), Orthogonal Matching Pursuit (OMP) (Tropp, 2004) is a popular greedy algorithm that iteratively selects the atom that is most correlated with the signal, orthogonalizes it to the previously selected atoms by solving a least squares optimization, and selects the next atom that is most correlated with the resulting residual. For the case of a Kronecker structured basis, a Kronecker OMP (Kron-OMP) algorithm has been previously proposed (Rivenson and S., 2009; Caiafa and Cichocki, 2013) that reduces computations of solving the least squares subproblem in each iteration by exploiting properties of the Kronecker product. This form of Kron-OMP, however, represents the signal, coefficients, and basis atoms in vector form providing a solution to (*Povec*).

In Algorithm 1 we rewrite the Kron-OMP algorithm adapted

| Algorithm 1 Kron-OMP. |
|--|
| Choose: K, ϵ . |
| Initialize: $k = 1$, $\mathcal{I}^0 = \emptyset$, $\mathcal{J}^0 = \emptyset$, $R_0 = S$, $s = vec(S)$. |
| while $k \le K$ and error $> \epsilon$ do |
| 1: $[i^k, j^k] = \operatorname{argmax}_{[i, j]} \Gamma_i^\top R_k \Psi_j ;$ |
| 2 : $\mathcal{I}^{k} = [\mathcal{I}^{k-1}, i^{k}]; \mathcal{J}^{k} = [\mathcal{J}^{k-1}, j^{k}]; \mathcal{A}^{k} = (\mathcal{I}^{k}, \mathcal{J}^{k});$ |
| 3: $c_k = \operatorname{argmin}_{c} \frac{1}{2} (\Gamma_{\tau k} \odot \Psi_{\tau k}) c - s _2^2;$ |
| 4: $R_k = mat(s - (\Gamma_{\tau k} \odot \Psi_{\tau k})c_k);$ |
| 5: $k \leftarrow k+1$; |
| end while |
| Return: \mathcal{A}^{K} , c_{K} |
| |

to the structure of our problem, where $ve(\cdot)$ and $mat(\cdot)$ convert matrices to vectors and vice versa. The main complexity gain of Kron-OMP over the vector OMP is obtained by separating the effects of Γ and Ψ when computing the maximally correlated atoms with the residual, $|\Gamma^{\top}R\Psi|$ (See Algorithm 1 Step 1) with complexity $O(N_{\Gamma}GV + GN_{\Gamma}N_{\Psi})$ instead of computing $|\Phi^{\top}r|$ with complexity $O(N_{\Gamma}N_{\Psi}GV)$. The other gain is in solving the least squares problem (See Algorithm 1 Step 3) by exploiting properties of the Kronecker product ($A \odot B = [a_1 \otimes b_1, \ldots, a_N \otimes B_N]$) to compute a rank-1 update. However, the only real improvement on complexity is in memory since Φ can be built atom by atom from columns of Γ and Ψ instead of storing the entire matrix. The rank-1 update remains $O(k^2)$ for both vector and Kron-OMP. In the next section we present an alternative Kron-OMP algorithm that reduces complexity further by exploiting the full separability of the dictionary.

² Using $\Psi = I_V$ identity with spherical ridgelets (SR) we adopt the notation I-SR for the dictionary used in the state-of-the-art illustration Fig. 2 and Section 5.

4.2. Kronecker OMP with projected gradient descent (PGD)

In what follows, we develop a novel form of Kronecker OMP which solves the separable (*POmat*) instead of (*POvec*). This allows us to reduce computation by not building columns of Φ and not repeating individual atoms of Γ or Ψ . Instead, indices of Γ and Ψ are updated only when they each have not been chosen before, fully exploiting the separability of the dictionary. Given the previous sets of respective of indices \mathcal{I}^{k-1} and \mathcal{J}^{k-1} , we update sets by following $\mathcal{I}^k = [\mathcal{I}^{k-1} \ i^k]$ if $i^k \notin \mathcal{I}^{k-1}$ and $\mathcal{I}^k = \mathcal{I}^{k-1}$ otherwise. Likewise, $\mathcal{J}^k = [\mathcal{J}^{k-1} \ j^k]$ if $j^k \notin \mathcal{I}^{k-1}$ and $\mathcal{I}^k = \mathcal{I}^{k-1}$ otherwise. With the selected indices, the size of C_k will be $|\mathcal{I}^k| \times |\mathcal{J}^k|$ instead of $k \times k$. To find C_k , it seems natural to solve:

$$C_k = \arg\min_{C} \frac{1}{2} ||\Gamma_{\mathcal{I}^k} C \Psi_{\mathcal{J}^k}^{\top} - S||_F^2.$$
(14)

But the solution C_k will contain possible nonzero coefficients that do not coincide with the chosen selection of indices since additional indices in all combinations of pairs between \mathcal{I}^k and \mathcal{J}^k will be updated in each iteration. This is problematic for the correctness of the algorithm when choosing the next single most correlated coefficient. Therefore we must enforce that these coefficients are zero:

$$C_k = \underset{C}{\arg\min} \frac{1}{2} ||\Gamma_{\mathcal{I}^k} C \Psi_{\mathcal{J}^k}^\top - S||_F^2 \text{ s.t. } C_{i,j} = 0 \ \forall (i, j) \in \mathcal{O}^k.$$
(15)

where $\mathcal{O}^k := \overline{(\mathcal{I}^k, \mathcal{J}^k)}$. To solve this problem, we can use projected gradient descent (PGD). The gradient of $f(C) = \frac{1}{2} ||\Gamma_{\mathcal{I}^k} C \Psi_{\mathcal{J}^k}^\top - S||_F^2$ at iteration *k* is

$$\nabla f(\mathcal{C}) = \Gamma_{\mathcal{I}^{k}}^{\top} \Gamma_{\mathcal{I}^{k}} \mathcal{C} \Psi_{\mathcal{J}^{k}}^{\top} \Psi_{\mathcal{J}^{k}} - \Gamma_{\mathcal{I}^{k}}^{\top} \mathcal{S} \Psi_{\mathcal{J}^{k}}.$$
(16)

To save on computation we precompute $\mathcal{G} = \Gamma^{\top}\Gamma$, $\mathcal{P} = \Psi^{\top}\Psi$, and $\hat{S} = \Gamma^{\top}S\Psi$ and can access their entries at each iteration: $\mathcal{G}_{\underline{I}^{k},\underline{I}^{k}}, \mathcal{P}_{\mathcal{J}^{k},\mathcal{J}^{k}}, \hat{S}_{\underline{I}^{k},\mathcal{J}^{k}}$. Then setting $Z^{1} = C_{k}$ we iteratively project the update in the gradient direction to the space of feasible solutions:

$$Z^{t+1} = P_{\mathcal{O}^k}(Z^t - \epsilon \nabla f(Z^t)), \tag{17}$$

where the projection $P_{\mathcal{O}^k}$ sets all elements in \mathcal{O}^k to 0 and stepsize ϵ is estimated each iteration using a line search. Once the procedure has converged to Z^* , we set $C_k = Z^*$ and compute the residual $R_k = S - \Gamma_{\mathcal{I}^k} C_k \Psi_{\mathcal{J}^k}^{\top}$. Then, for iteration k + 1 we must find $(i^{k+1}, j^{k+1}) = \arg \max_{[i,j]} |\Gamma_i^{\top} R_k \Psi_j|$. To save significantly on computation we can instead use our precomputed \mathcal{G} and \mathcal{P} to instead find $\arg \max_{[i,j]} [\hat{R}_k]_{i,j}$, where $\hat{R}_k = |\hat{S} - \mathcal{G}_{\mathcal{I}^k} C_k \mathcal{P}_{\mathcal{J}^k}^{\top}|$ where $\mathcal{G}_{\mathcal{I}^k}, \mathcal{P}_{\mathcal{J}^k}$ respectively access the $\mathcal{I}^k, \mathcal{J}^k$ columns and all rows. Maintaining matrix forms throughout allows us to combine computing the residual and the next atoms for a large reduction in computation at each iteration k. Our proposed Kronecker OMP with projected gradient descent (Kron-OMP-PGD) is outlined in Algorithm 2.

We show a comparison of time per iteration for a small $V = 50 \times 50$, G = 64 phantom dataset in Fig. 7. The steeper time increase for Kron-OMP is due to the fact that at iteration k there is a complexity of $O(k^2 + kGV)$ that comes from Steps 3 (rank-1 update) and 4 of Algorithm 1. On the other hand, Kron-OMP-PGD has complexity involving $|\mathcal{I}^k|, |\mathcal{J}^k| \le k$ which are in practice significantly less than k. Even though a PGD sub-routine must be performed at each iteration k, we found that by incorporating Nesterov acceleration with a line search, the time per iterations k increases.

However, for dMRI data, typical sparsity levels are K = O(V). So for $V \approx 100^3$ the number of iterations as well as the time per iteration of both Kron-OMP and Kron-OMP-PGD when *k* approaches *K* becomes astronomical. Even on a relatively small 3D phantom dataset of spatial size $V = 50 \times 50$, for example, one iteration takes

Algorithm 2 Kron-OMP-PGD.

Choose: $K, \epsilon_1, \epsilon_2$. Precompute: $\hat{S} = \Gamma^{\top} S \Psi$, $\mathcal{G} = \Gamma^{\top} \Gamma$, $\mathcal{P} = \Psi^{\top} \Psi$. Initialize: k = 1, $C_0 = 1$, $\mathcal{I}^0 = \emptyset$, $\mathcal{J}^0 = \emptyset$, $\hat{R}_0 = \hat{S}$. while $k \leq K$ and error $> \epsilon_1$ do 1: $[i^k, j^k] = \operatorname{argmax}_{[i,j]}[\hat{R}_k]_{i,j};$ 2: $\mathcal{I}^k = \mathcal{I}^{k-1} \cup \{i^k\}; \quad \mathcal{J}^k = \mathcal{J}^{k-1} \cup \{j^k\}; \quad \mathcal{A}^k = (\mathcal{I}^k, \mathcal{J}^k); \quad \mathcal{O}^k = \overline{\mathcal{A}^k};$ 3: $Z^{1}_{\mathcal{I}^{k-1},\mathcal{I}^{k-1}} = C_{k-1}; n_1 = 0; t = 1;$ while error $> \epsilon_2$ do 1: $\delta = \text{linesearch}(Z^t)$; 2: $X^{t+1} = P_{\mathcal{O}^k}(Z^t - \delta(\mathcal{G}_{\mathcal{I}^k,\mathcal{I}^k}Z^t\mathcal{P}_{\mathcal{J}^k,\mathcal{J}^k} - \hat{S}_{\mathcal{I}^k,\mathcal{J}^k}));$ $\begin{array}{l} 4: \; n_{t+1} = \frac{1}{2} (1 + \sqrt{1 + 4n_t^2}); \\ 5: \; Z^{t+1} = X^{t+1} + \frac{n_t - 1}{n_{t+1}} (X^{t+1} - X^t); \end{array}$ 6: $t \leftarrow t + 1$; end while 4: $C_k = Z^*$; 5: $\hat{R}_k = |\hat{S} - \mathcal{G}_{\mathcal{I}^k} C_k \mathcal{P}_{\mathcal{I}^k}|;$ 6: $k \leftarrow k + 1$; end while Return: \mathcal{A}^{K}, C_{K} .



Fig. 7. Comparison of time per iteration for Kron-OMP and the proposed Kron-OMP-PGD. The total time to choose K = 7000 = 2.8V atoms for this $V = 50 \times 50$ slice of a phantom dataset, is 68 min for Kron-OMP and 40 min for Kron-OMP-PGD. We can see that as the number of atoms grows, the time per iteration of Kron-OMP continues to grow at a much higher rate than Kron-OMP-PGD.

on the order of a few seconds which results in over 34 hrs for these greedy algorithms to reach 1 atom/voxel atoms (K = V). In this way, greedy algorithms such as OMP are not suitable for large-scale problems that require hundreds of thousands of iterations. Instead, optimizing the LASSO problem (*P1mat*) can be accomplished with significantly less iterations, as we examine in the following section.

4.3. Kronecker ADMM

The Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2010) is a popular method for solving the LASSO problem (*P1vec*). However, its application in the case of a large dictionary Φ remains prohibitive, requiring computations involving $\Phi^{\top}s$ of order $O(GVN_{\Gamma}N_{\Psi})$. Instead, we apply ADMM to the separable LASSO problem (*P1mat*) to reduce computations by solving

$$\min_{C,Z} \frac{1}{2} ||\Gamma C \Psi^{\top} - S||_{F}^{2} + \lambda ||Z||_{1} \quad \text{s.t.} \quad C = Z.$$
(18)

The augmented Lagrangian writes:

$$\mathcal{L}_{\mu}(C, Z, \mathcal{T}) = \frac{1}{2} ||\Gamma C \Psi^{\top} - S||_{F}^{2} + \lambda ||Z||_{1} + \langle \mathcal{T}, C - Z \rangle + \frac{\mu}{2} ||C - Z||_{F}^{2},$$
(19)

and:

$$\frac{\partial \mathcal{L}_{\mu}}{\partial C} = \Gamma^{\top} (\Gamma C \Psi^{\top} - S) \Psi + \mathcal{T} + \mu (C - Z) = 0$$
(20)

$$\Rightarrow \Gamma^{\top} \Gamma C \Psi^{\top} \Psi + \mu C = \mu Z - \mathcal{T} + \Gamma^{\top} S \Psi := Q.$$
(21)

In principle, one can solve for *C* by solving a linear system of equations h(C) = Q, where $h(C) = \Gamma^{\top} \Gamma C \Psi^{\top} \Psi + \mu C$. However, solving this linear system directly is computationally challenging due to the size of the matrices involved. Therefore, to solve for *C* efficiently, we begin by taking the SVDs of Γ and Ψ . With $\Gamma = U_{\Gamma} \Sigma_{\Gamma} V_{\Gamma}^{\top}$ and $\Psi = U_{\Psi} \Sigma_{\Psi} V_{\Psi}^{\top}$, $\Gamma^{\top} \Gamma = V_{\Gamma} \Delta_{\Gamma} V_{\Gamma}^{\top}$ and $\Psi^{\top} \Psi = V_{\Psi} \Delta_{\Psi} V_{\Psi}^{\top}$, where U_{Γ}, U_{Ψ} are the matrices of eigenvectors and $\Delta_{\Gamma} = \Sigma_{\Gamma}^{\top} \Sigma_{\Gamma}, \Delta_{\Psi} = \Sigma_{\Psi}^{\top} \Sigma_{\Psi}$ are the diagonal matrices of eigenvalues for Γ and Ψ respectively. Then:

$$V_{\Gamma}\Delta_{\Gamma}V_{\Gamma}^{\top}CV_{\Psi}\Delta_{\Psi}V_{\Psi}^{\top} + \mu C = Q$$
⁽²²⁾

$$\Rightarrow \Delta_{\Gamma} \tilde{\mathcal{C}} \Delta_{\Psi} + \mu \tilde{\mathcal{C}} = \tilde{Q}$$
⁽²³⁾

where we introduced the notation $\tilde{X} = V_{\Gamma}^{\top} X V_{\Psi}$. Since Δ_{Γ} and Δ_{Ψ} are diagonal with elements δ_{Γ_i} and δ_{Ψ_j} , respectively, we can solve for \tilde{C} by:

$$\delta_{\Gamma_{i}}\tilde{C}_{i,j}\delta_{\Psi_{j}} + \mu\tilde{C}_{i,j} = \tilde{Q}_{i,j} \Rightarrow \tilde{C}_{i,j} = \frac{\tilde{Q}_{i,j}}{\delta_{\Gamma_{i}}\delta_{\Psi_{j}} + \mu}$$
(24)

To write this in matrix form we define $[\Delta_{\mu}]_{i,j} \triangleq 1/(\delta_{\Gamma_i} \delta_{\Psi_j} + \mu)$ and have $\tilde{C} = (\Delta_{\mu} \circ \tilde{Q})$ where \circ stands for element-wise matrix multiplication. Finally, we can recover $C = V_{\Gamma} \tilde{C} V_{\Psi}^{\top}$ and the complete update for *C* is:

$$C_{k+1} = V_{\Gamma} (\Delta_{\mu} \circ (V_{\Gamma}^{\top} Q_k V_{\Psi})) V_{\Psi}^{\top}$$
(25)

where $Q = \mu Z - \mathcal{T} + \Gamma^{\top} S \Psi$. When minimizing \mathcal{L}_{μ} with respect to Z, we end up with the usual proximal operator of the L_1 norm that is given by the shrinkage operator, $\operatorname{shrink}_{\kappa}(X) = \max(0, X - \kappa) - \max(0, -X - \kappa)$, applied element-wise to matrix X, giving $Z_{k+1} = \operatorname{shrink}_{\lambda/\mu}(C_{k+1} + \mathcal{T}_k)$. Similarly with respect to \mathcal{T} , we have the usual Lagrange multiplier gradient ascent update $\mathcal{T}_{k+1} = \mathcal{T}_k + C_{k+1} - Z_{k+1}$.

The formal updates for Kron-ADMM are presented in Algorithm 3. The update for *C* in (25) works well when Γ

| Algorithm 3 | Kron-ADMM | (for | undercomplete | dictionaries). |
|-------------|-----------|------|---------------|----------------|
|-------------|-----------|------|---------------|----------------|

Choose: μ , λ , ϵ . Precompute: V_{Γ} , Δ_{Γ} , V_{Ψ} , Δ_{Ψ} , Δ_{μ} . Initialize: $k = 0, Z_0 = 0, T_0 = 0, \hat{S} = \Gamma^{\top} S \Psi$. while error > ϵ do 1: $Q_k = \hat{S} + \mu Z_k - \mathcal{T}_k$; 2: $C_{k+1} = V_{\Gamma} (\Delta_{\mu} \circ (V_{\Gamma}^{\top} Q_k V_{\Psi})) V_{\Psi}^{\top}$; 3: $Z_{k+1} = \operatorname{shrink}_{\lambda/\mu} (C_{k+1} + \mathcal{T}_k)$; 4: $\mathcal{T}_{k+1} = \mathcal{T}_k + C_{k+1} - Z_{k+1}$; 5: $k \leftarrow k + 1$; end while Return: C.

and Ψ are under-complete and the eigen-decompositions of $\Gamma^{\top}\Gamma$ and $\Psi^{\top}\Psi$ are easily computable. However, dictionaries most commonly used for sparse coding and the application to CS are over-complete i.e. $G < N_{\Gamma}$ and $V < N_{\Psi}$ making these SVDs potentially expensive to compute. In the case of an over-complete Φ , for traditional vector ADMM, the matrix inversion lemma (Boyd et al., 2010) is involved in order to compute SVDs of the smaller $\Phi\Phi^{\top}$ instead of $\Phi^{\top}\Phi$. In the following proposition, we derive the equivalent result for the update of *C* in (25). **Proposition 1.** For over-complete dictionaries Γ and Ψ , update (25) is equivalent to the more compact

$$C = Q/\mu - \Gamma^{\top} U_{\Gamma} (\Delta_{\mu} \circ (U_{\Gamma}^{\top} \Gamma Q \Psi^{\top} U_{\Psi})) U_{\Psi}^{\top} \Psi/\mu.$$
⁽²⁶⁾

Proof. For over-complete dictionaries $\Gamma = U_{\Gamma}[\Sigma_{\Gamma}, \mathbf{0}]V_{\Gamma}^{\top}$ and $\Psi = U_{\Psi}[\Sigma_{\Psi}, \mathbf{0}]V_{\Psi}^{\top}$,

$$\Gamma^{\top}\Gamma = V_{\Gamma}\begin{pmatrix}\Delta_{\Gamma} & \mathbf{0}\\ \mathbf{0} & \mathbf{0}\end{pmatrix}V_{\Gamma}^{\top} \text{ and } \Psi^{\top}\Psi = V_{\Psi}\begin{pmatrix}\Delta_{\Psi} & \mathbf{0}\\ \mathbf{0} & \mathbf{0}\end{pmatrix}V_{\Psi}^{\top}.$$

For $G < i \le N_{\Gamma}, V < j \le N_{\Psi}, \ \delta_{\Gamma_{i}}, \delta_{\Psi_{j}} = 0, \text{ so } \tilde{C}_{i,j} = \frac{\tilde{Q}_{i,j}}{\delta_{\Gamma_{i}}\delta_{\Psi_{j}} + \mu} = \frac{\tilde{Q}_{i,j}}{\mu}.$
For $i \le G$ and $j \le V$, we can rewrite

$$\begin{split} \tilde{C}_{i,j} &= \frac{\tilde{Q}_{i,j}}{\delta_{\Gamma_i} \delta_{\Psi_j} + \mu} = \frac{\tilde{Q}_{i,j}}{\mu} - \frac{\delta_{\Gamma_i} \tilde{Q}_{i,j} \delta_{\Psi_j}}{\mu (\delta_{\Gamma_i} \delta_{\Psi_j} + \mu)} = \frac{\tilde{Q}_{i,j}}{\mu} - \frac{\sigma_{\Gamma_i}^2 \tilde{Q}_{i,j} \sigma_{\Psi_j}^2}{\mu (\delta_{\Gamma_i} \delta_{\Psi_j} + \mu)} \\ &= \frac{\tilde{Q}_{i,j}}{\mu} - \sigma_{\Gamma_i} \frac{\sigma_{\Gamma_i} \tilde{Q}_{i,j} \sigma_{\Psi_j}}{\mu (\delta_{\Gamma_i} \delta_{\Psi_j} + \mu)} \sigma_{\Psi_j} \\ &\Rightarrow \tilde{C} = \tilde{Q}/\mu - \Sigma_{\Gamma}^{\top} (\Delta_{\mu} \circ (\Sigma_{\Gamma} \tilde{Q} \Sigma_{\Psi}^{\top})) \Sigma_{\Psi}/\mu \\ C &= Q/\mu - V_{\Gamma} \Sigma_{\Gamma}^{\top} (\Delta_{\mu} \circ (\Sigma_{\Gamma} V_{\Gamma}^{\top} Q V_{\Psi} \Sigma_{\Psi}^{\top})) \Sigma_{\Psi} V_{\Psi}^{\top}/\mu \\ C &= Q/\mu - \Gamma^{\top} U_{\Gamma} (\Delta_{\mu} \circ (U_{\Gamma}^{\top} \Gamma Q \Psi^{\top} U_{\Psi})) U_{\Psi}^{\top} \Psi/\mu. \end{split}$$

Letting $\Gamma' = U_{\Gamma}^{\top}\Gamma$ and $\Psi' = U_{\Psi}^{\top}\Psi$, which can be precomputed, we have a final efficient update

$$C_{k+1} = Q_k/\mu - \Gamma^{\prime \top} (\Delta_\mu \circ (\Gamma' Q_k \Psi^{\prime \top})) \Psi'/\mu.$$
(27)

This allows us to compute the SVDs of $\Gamma\Gamma^{\top}$ and $\Psi\Psi^{\top}$ instead of the larger $\Gamma^{\top}\Gamma$ and $\Psi^{\top}\Psi$ and work with smaller matrices within each iteration. We present Kron-ADMM for over-complete dictionaries in Algorithm 4.

| Algorithm 4 Kron-ADMM (for overcomplete dictionaries). |
|--|
| Choose: μ , λ , ϵ . |
| Precompute: U_{Γ} , Δ_{Γ} , U_{Ψ} , Δ_{Ψ} , Γ' , Ψ' , Δ_{μ} . |
| Initialize: $k = 0, Z_0 = 0, T_0 = 0.$ |
| while error $> \epsilon$ do |
| 1: $Q_k = \Gamma^{\top} S \Psi + \mu Z_k - \mathcal{T}_k;$ |
| 2: $C_{k+1} = Q_k/\mu - \Gamma^{\prime \top} (\Delta_\mu \circ (\Gamma^{\prime} Q_k \Psi^{\prime \top})) \Psi^{\prime}/\mu;$ |
| 3: $Z_{k+1} = \text{shrink}_{\lambda/\mu} (C_{k+1} + \mathcal{T}_k);$ |
| 4: $\mathcal{T}_{k+1} = \mathcal{T}_k + C_{k+1} - Z_{k+1}$; |
| 5: $k \leftarrow k+1$; |
| end while |
| Return: C. |
| |

4.4. Kronecker dual ADMM

As an alternative to ADMM, Dual ADMM, which applies ADMM to the dual of *l*1 problem (*P1vec*), has been shown to be more efficient than ADMM for over-complete dictionaries (Goncalves, 2015) by allowing one to compute SVDs of the more affordable $\Phi\Phi^{\top}$ instead of $\Phi^{\top}\Phi$. In our previous work Schwab et al. (2016) we proposed a Kronecker Dual ADMM (Kron-DADMM) that efficiently solves the spatial-angular sparse coding problem. Below, we give an alternative derivation of this algorithm directly based on the matrix formulation of (*P1mat*). The dual of (*P1mat*) is:

$$\max_{A} -\frac{1}{2} ||A||_{F}^{2} + A^{\mathsf{T}}S \quad \text{s.t.} \quad ||\Gamma^{\mathsf{T}}A\Psi||_{\infty} \leq \lambda,$$
(28)

where $||X||_{\infty} = \max_{i,j} |X_{i,j}|$. To apply ADMM to this optimization problem, we replace $\Gamma^{\top}A\Psi$ with auxiliary variable \mathcal{V} and add the additional constraint $\Gamma^{\top}A\Psi - \mathcal{V} = 0$ to get:

$$\max_{A,\mathcal{V}} -\frac{1}{2} ||A||_F^2 + A^{\mathsf{T}} S \quad \text{s.t.} \quad ||\mathcal{V}||_{\infty} \le \lambda \quad \text{and} \quad \mathcal{V} = \Gamma^{\mathsf{T}} A \Psi.$$
(29)

Then the augmented Lagrangian is

$$\mathcal{L}_{\eta}(A, \mathcal{V}, C) = -\frac{1}{2} ||A||_{F}^{2} + A^{\top}S + \langle C, \mathcal{V} - \Gamma^{\top}A\Psi \rangle + \frac{\eta}{2} ||\mathcal{V} - \Gamma^{\top}A\Psi||_{F}^{2} + \delta_{\lambda}(\mathcal{V})$$
(30)

where

$$\delta_{\lambda}(\mathcal{V}) = \begin{cases} 0 & \text{if } ||\mathcal{V}||_{\infty} \le \lambda \\ \infty & \text{if } ||\mathcal{V}||_{\infty} > \lambda \end{cases}$$
(31)

and the Lagrange multiplier C corresponds to the primal variable C in (*P1mat*), which is our variable of interest. We then have

$$\frac{\partial \mathcal{L}_{\eta}(A, \mathcal{V}, C)}{\partial A} = -A + S - \Gamma C \Psi^{\top} - \eta \Gamma (\mathcal{V} - \Gamma^{\top} A \Psi) \Psi^{\top} = 0 \quad (32)$$

$$\Rightarrow A - \eta \Gamma \Gamma^{\top} A \Psi \Psi^{\top} = S - \Gamma (C + \eta \mathcal{V}) \Psi^{\top} := P.$$
(33)

Now with eigen-decompositions $\Gamma\Gamma^{\top} = U_{\Gamma}\Delta_{\Gamma}U_{\Gamma}^{\top}$ and $\Psi\Psi^{\top} = U_{\Psi}\Delta_{\Psi}U_{\Psi}^{\top}$ and letting $\tilde{X} = U_{\Gamma}^{\top}XU_{\Psi}$,

$$A + \eta U_{\Gamma} \Delta_{\Gamma} U_{\Gamma}^{\top} A U_{\Psi} \Delta_{\Psi} U_{\Psi}^{\top} = P$$
(34)

$$\Rightarrow \tilde{A} + \eta \Delta_{\Gamma} \tilde{A} \Delta_{\Psi} = \tilde{P}.$$
(35)

Then, \tilde{A} can be found element-wise by:

$$\tilde{A}_{i,j} + \eta \delta_{\Gamma_i} \tilde{A}_{i,j} \delta_{\Psi_j} = \tilde{P}_{i,j} \Rightarrow \tilde{A}_{i,j} = \frac{\tilde{P}_{i,j}}{1 + \eta \delta_{\Gamma_i} \delta_{\Psi_j}}.$$
(36)

Defining $[\Delta_{\eta}]_{i,j} \triangleq 1/(1 + \eta \delta_{\Gamma_i} \delta_{\Psi_j})$, the update is $\tilde{A} = \Delta_{\eta} \circ \tilde{P}$. As shown in Goncalves (2015) we can keep the update in terms of \tilde{A} instead of A since the variable we are interested in is C. We can then precompute $S' = \Gamma'^{\top} S \Psi'$, $\Gamma' = U_{\Gamma}^{\top} \Gamma$ and $\Psi' = U_{\Psi}^{\top} \Psi$. The updates of \mathcal{V} and C are as in Schwab et al. (2016) and presented in Algorithm 5, where $P_{\lambda}^{\infty}(X)$ sets all entries of matrix X that are

Algorithm 5 Kron-DADMM.

Choose: η, λ, ϵ . Precompute: $S', \Gamma', \Psi', \Delta_{\eta}$. Initialize: $k = 0, C_0 = 0, \mathcal{V}_0 = 0$. **while** Duality Gap > ϵ **do** 1: $\tilde{A}_{k+1} = \Delta_{\eta} \circ (S' - \Gamma'(C_k - \eta \mathcal{V}_k) \Psi'^{\top});$ 2: $\mathcal{V}_{k+1} = P_{\lambda}^{\infty} (\frac{1}{\eta} C_k + \Gamma'^{\top} \tilde{A}_{k+1} \Psi');$ 3: $C_{k+1} = \text{shrink}_{\lambda\eta} (C_k + \eta \Gamma'^{\top} \tilde{A}_{k+1} \Psi');$ 4: $k \leftarrow k + 1;$ **end while** Return: *C*.

greater than λ to λ .

4.5. Kronecker FISTA

The Fast Iterative Thresholding Algorithm (FISTA) (Beck and Teboulle, 2009) is another well-known method for solving LASSO. However, just as before, applying FISTA to (*P1vec*) for large-scale dMRI data is largely intractable. So here we adapt FISTA to (*P1mat*) in order to exploit the separability of our spatial-angular basis. FISTA is a proximal gradient descent

$$C_{k+1} = \operatorname{shrink}_{\lambda/L}(C_k - \nabla f(C_k)/L), \tag{37}$$

where the proximal operator is the soft-thresholding shrinkage operator associated with the l1 norm and 1/L is a chosen step size. The gradient is simply computed as:

$$\nabla f(\mathcal{C}) = \Gamma^{\top} (\Gamma \mathcal{C} \Psi^{\top}) \Psi - \Gamma^{\top} \mathcal{S} \Psi.$$
(38)

Table 3

Comparison of algorithms complexity at iteration k. For Kron-OMP-PGD, T is the number of sub-iterations of PGD.

| Algorithm | Standard | Kronecker |
|--|--|---|
| omp omp-pgd Admm DADMM FISTA | $ \begin{array}{l} k^2 + kGV + GVN_{\Gamma}N_{\Psi} \\ - \\ (GV)^2N_{\Gamma}N_{\Psi} + GV(N_{\Gamma}N_{\Psi})^2 \\ (GV)^2N_{\Gamma}N_{\Psi} \\ (N_{\Gamma}N_{\Psi})^2 + GVN_{\Gamma}N_{\Psi} \end{array} $ | $ \begin{split} & k^2 + kGV + GVN_{\Gamma} + VN_{\Gamma}N_{\Psi} \\ & TG \mathcal{I}^k \mathcal{J}^k + TGV \mathcal{J}^k + \mathcal{J}^k N_{\Gamma}N_{\Psi} \\ & (GN_{\Gamma}N_{\Psi} + GVN_{\Psi}) + GV \\ & (GN_{\Gamma}N_{\Psi} + GVN_{\Psi}) + GV \\ & (GN_{\Gamma}N_{\Psi} + GVN_{\Psi}) \end{split} $ |

To help speed convergence, we use a line search subroutine to update *L* at each iteration in addition to the usual Nesterov acceleration. By Beck and Teboulle (2009), FISTA will converge for any *L* greater than the Lipschitz constant of ∇f , which can be estimated by bounding

$$||\nabla f(C) - \nabla f(\bar{C})||_{F} = ||\Gamma^{\top}\Gamma(C - \bar{C})\Psi^{\top}\Psi||_{F} \le \lambda_{\max}^{\Gamma}\lambda_{\max}^{\Psi}||C - \bar{C}||_{F}$$
(39)

where $\lambda_{\text{max}}^{\Gamma}$ and $\lambda_{\text{max}}^{\Psi}$ are the maximum eigenvalues of $\Gamma^{\top}\Gamma$ and $\Psi^{\top}\Psi$ respectively. Therefore we initialize $L = \lambda_{\text{max}}^{\Gamma}\lambda_{\text{max}}^{\Psi}$. The Kronecker FISTA (Kron-FISTA) is presented in Algorithm 6. This natural

Algorithm 6 Kron-FISTA.

Choose: ϵ . Precompute: $\hat{S} = \Gamma^{\top} S \Psi$ Initialize: $Z_1 = C_0 = \mathbf{0}$, $n_1 = 1$, $L = \lambda_{\max}^{\Gamma} \lambda_{\max}^{\Psi}$. **while** error $> \epsilon$ **do** 1: $L = \text{linesearch}(Z_k)$; 2: $\nabla f(Z_k) = \Gamma^{\top} (\Gamma Z_k \Psi^{\top}) \Psi - \hat{S}$; 3: $C_k = \text{shrink}_{\lambda/L}(Z_k - \nabla f(Z_k)/L)$; 4: $n_{k+1} = \frac{1}{2}(1 + \sqrt{1 + 4n_k^2})$; 5: $Z_{k+1} = C_{k+1} + \frac{n_{k-1}}{n_{k+1}}(C_{k+1} - C_k)$; 6: $k \leftarrow k + 1$; **end while** Return: *C*.

Kronecker extension to FISTA has also been recently presented in Qi et al. (2016), but has not been adapted and tested on data of our scale.

4.6. Complexity analysis

To evaluate the efficiency of each algorithm and the gains of Kronecker separability compared to the original algorithms we summarize the complexity of each algorithm for general Ψ and Γ in Table 3. We notice that classical LASSO algorithms have complexity on the order of the size of the Φ matrix, including terms that multiply all four dimensions $GVN_{\Gamma}N_{\Psi}$. When applying the Kronecker LASSO algorithms, the complexity is reduced to a summation that includes only 3 of the dimensions GVN_{Ψ} , a reduction on the order of N_{Γ} (≈ 200 for some of our dictionary choices). We compare the Kronecker LASSO algorithms empirically in Section 5 to identify which is fastest for our regime. Next we address the fact that the dimensions of $\Gamma \in \mathbb{R}^{G \times N_{\Gamma}}$ and $\Psi \in \mathbb{R}^{V \times N_{\Psi}}$ will be orders of magnitude different since $G \approx 100$ and $V \approx 100^3$. We consider a few specific assumptions on the structure of spatial dictionary Ψ which can decrease the complexity and simplify computations of some of the proposed algorithms:

 Ψ *Tight Frame.* In the case that Ψ is a tight frame, *i.e.* $\Psi\Psi^{\top} =$ I, which is commonly an assumption in compressed sensing theorems, our method can still be simplified. In Kron-ADMM (overcomplete) and Kron-DADMM, we may avoid the SVD of $\Psi\Psi^{\top}$ and respective updates (23) and (35) can be simplified.

 Ψ *Fast Transform.* In the case that Ψ corresponds to a well-studied transform such as wavelets, curvelets, etc., fast transform implementations can be utilized to reduce complexity further. For the case of FISTA, for example, matrix multiplications of $\Gamma^{\top}(\Gamma Z_k \Psi^{\top})\Psi$ (See Algorithm 6 Step 2) involve fast transform reconstructions (Ψ^{\top}) of each DWI (ΓZ_k) and then deconstructions (Ψ) which we parallelize over all DWI in our implementation.

 Ψ *Orthonormal.* In the case that Ψ is orthonormal, *i.e.* $\Psi^{\top}\Psi = \Psi\Psi^{\top} = I$ then (*P1mat*) can be simplified to (5) after noticing:

$$||\Gamma C \Psi^{\top} - S||_{F}^{2} = ||\Gamma C \Psi^{\top} \Psi - S \Psi||_{F}^{2} = ||\Gamma C - \hat{S}||_{F}^{2}.$$
 (40)

This optimization can be solved using traditional methods after precomputing $\hat{S} = S\Psi$.

Ψ *Separable Tensor Product.* In the case that Ψ can be separated into a 3D tensor product $Ψ = Ψ_x \otimes Ψ_y \otimes Ψ_z$, the complexity of multiplication can be simplified by another degree, in the same vein as the decrease in complexity we gained from using $Φ = Ψ \otimes Γ$. In this case, instead of the matrix multiplication, $S = ΓCΨ^T$ can be written using n-mode products of tensors $S = C ×_x Ψ_x ×_y Ψ_y ×_z Ψ_z ×_q Γ$. Furthermore, if we consider DSI acquisition where *q*-space measurements are acquired in a grid over \mathbb{R}^3 , and assume we can represent these measurements over a separable basis over each dimension, then we can take $Γ = Γ_{q_x} \otimes Γ_{q_y} \otimes Γ_{q_z}$ and Φ becomes a 6-tensor.

5. Experiments

5.1. Data

We perform our experiments on single-shell HARDI data, though as we emphasized earlier, our framework and algorithms can be applied to any dMRI acquisition protocol with a suitable choice of the angular basis Γ . We experimented on a phantom and a real HARDI brain dataset. Specifically, we applied our methods to the ISBI 2013 HARDI Reconstruction Challenge Phantom dataset,³ a $V = 50 \times 50 \times 50$ volume consisting of 20 phantom fibers crossing intricately within an inscribed sphere, measured with G = 64 gradient directions (SNR = 30). Our initial experiments test on a 2D 50×50 slice of this data for simplification. The real HARDI brain dataset consists of a $V = 112 \times 112 \times 65$ volume with G = 127 gradient directions. We conducted experiments on the core white matter brain region of size $V = 60 \times 60 \times 30$.

5.2. Kronecker algorithm comparison

In this section we compare the computational time performance of each of the proposed Kronecker LASSO algorithms, (Kron-ADMM, Kron-DADMM, and Kron-FISTA) on a 2D 50 × 50 slice of phantom data for various values of λ using Haar-SR. For our experiment, we ran Kron-FISTA until we reached a very small mean squared error of $\epsilon = 10^{-8}$. The objective value obtained was then taken to be a rough ground truth minimum. We then tested each of Kron-ADMM, Kron-DADMM, and Kron-FISTA and recorded the time it took to reach a relative error of 10⁻⁴ from the known minimum. Fig. 8 reports the objective value of each algorithm as a function of computing time for various sparsity levels associated to choices of λ . Table 8 gives the number of iterations until completion for each method and sparsity level. For our experiments, Kron-FISTA appears to be the fastest algorithm in all cases, followed by Kron-DADMM. The superior performance of DADMM over ADMM is consistent with the findings of Goncalves (2015). With these results, we henceforth use Kron-FISTA for subsequent experiments.

Table 4

Number of iterations to completion for Kron-ADMM, Kron-DADMM, Kron-FISTA on a 2D 50 \times 50 phantom HARDI data using Haar-SR. Kron-FISTA converges in the fewest number of iterations. For computation time, see Fig. 8.

| Atoms/Voxel | 0.09 | $0.24 \\ 1.4^{-1}$ | 0.60 | 1.72 | 3.67 | 6.75 |
|-------------|------------------|--------------------|-------------------|-------------------|-------------------|-------------------|
| λ | 1.4 ¹ | | 1.4 ⁻³ | 1.4 ⁻⁵ | 1.4 ⁻⁷ | 1.4 ⁻⁹ |
| Kron-ADMM | 797 | 1462 | 2096 | 3660 | 4365 | 4667 |
| Kron-DADMM | 357 | 597 | 1060 | 1722 | 1928 | 1953 |
| Kron-FISTA | 161 | 219 | 288 | 346 | 584 | 611 |

5.3. Choice of spatial-angular dictionaries

The experiments in this section are conducted using fixed spatial and angular dictionaries. For the choice of the angular dictionary Γ , we consider the over-complete Spherical Ridgelet (SR) basis (Tristán-Vega and Westin, 2011), which has been shown to sparsely model HARDI signals. The corresponding dictionary in the space of ODFs is the set of spherical wavelets (SW) (see Fig. 5 for an example of one spherical wavelet atom). With order L = 2 and 4, the SR dictionary contains $N_{\Gamma}=210$ and $N_{\Gamma}=1169$ atoms, respectively. We used both amounts of atoms for the small 2D 50×50 phantom dataset and found roughly identical results suggesting that a basis of order L = 2 contains enough atoms if the number of gradients is below 210. This reduces computation significantly. In comparison, the spherical harmonic (SH) basis has been shown in prior work Tristán-Vega and Westin (2011) to not exude sparse signals and so we do not repeat this comparison in the current work.

For the choice of spatial dictionary Ψ , the spatial wavelet transform is a popular sparsifying basis for natural images and structural MRI volumes. In our previous work Schwab et al. (2016) we compared the performance of Daubechies wavelets and Haar wavelets and concluded that Daubechies wavelets resulted in a smoothing of the boundaries between isotropic and anisotropic regions which was not indicative of the more abrupt boundaries that real HARDI data exhibits. Haar wavelets outperformed Daubechies wavelets in terms of reconstruction error arguably due to the fact that HARDI data exhibits more rigid boundaries and piece-wise consistencies, a spatial feature that has motivated the use of totalvariation penalties in many other reconstruction methods. For this reason, we do not consider Daubechies wavelets in this work.

In addition to Haar wavelets, we consider the spatial curvelets dictionary (Candès et al., 2006) (featured as the spatial atom in Fig. 5) which, in addition to variations in position and scale, offers directional variations which may be useful for sparsely modeling the naturally directional HARDI fiber tracts regions. An important criteria for choosing our spatial basis is that they be tight frames as this choice has important theoretical implications for compressed sensing and offers computational advantages (as discussed in Section 4.6). They additionally have fast transform implementations which also reduce computational complexity. Finally, to compare our formulation to state-of-the-art voxel-wise angular sparse coding, we can simply choose Ψ to be the $V \times V$ identity I_V . For ease of notation, we use a spatial-angular $\Psi - \Gamma$ labeling: Haar-SR, Curve-SR, I-SR for Haar wavelets, curvelets, and the identity, respectively, for the spatial domain with SR for the angular domain.

5.4. Sparsity results

In this section we compare the performance of our spatialangular sparse coding method to the state-of-the-art angular sparse coding by analyzing reconstruction accuracy using very few nonzero coefficients. The first experiment is tested on the 50×50

³ http://www.hardi.epfl.ch/static/events/2013_ISBI/



Fig. 8. Comparison of time for completion of Kron-ADMM, Kron-DADMM, and Kron-FISTA on a 2D 50 × 50 phantom HARDI data using Haar-SR for various sparsity levels. Kron-FISTA consistently reaches the minimum objective in the least amount of time. For number of iterations and lambda values, see Table 4.



Fig. 9. Quantitative results of residual error vs. spatial-angular sparsity levels for I-SR, Haar-SR, and Curve-SR on 2D phantom data for various values of λ . Curve-SR out performs Haar-SR for low sparsity levels while I-SR has very high relative reconstruction error. The reconstruction of I-SR data points are displayed in Fig. 2 and Haar-SR/Curve-SR in Fig. 10. Our finding of I-SR requiring 6–8 atoms per voxel for accurate reconstruction is consistent with previous findings (Michailovich and Rathi, 2008; 2010b).

phantom data slice. We ran Kron-FISTA for various values of λ for Haar-SR, Curve-SR and I-SR. In Fig. 9 we show the results of residual reconstruction error $\frac{1}{GV}||S^* - S_{orig}||_F$ vs. spatial-angular sparsity levels in terms of the average number of atoms per voxel $(||C^*||_0/V)$. The ideal reconstruction will have a very low average number of atoms per voxel with low residual error, which happens in the lower left-hand corner of our plot. We can see that in this range, Curve-SR outperforms Haar-SR while I-SR is unable to perform at this level. Reconstruction of I-SR for various sparsity levels are visualized in Fig. 2. In comparison, Fig. 10 displays the sparse reconstruction of Haar-SR and Curve-SR with an average of 0.25 atoms/voxel. Notice that Curve-SR leads to a somehow smoother and more accurate reconstruction than the expectedly boxy reconstruction of Haar-SR at this very high sparsity level. Still, in both cases, the proposed joint spatial-angular sparse coding can reconstruct accurate signals with much fewer number of atoms than angular sparse coding, which as seen again from Fig. 2 can be achieved with an average of around 4 atoms per voxel. More strikingly, in cases of high signal complexity for crossing fibers, the sparse code requires on the order of 6–12 atoms per voxel (see Fig. 3).

We repeat this same analysis on real HARDI data. First, as was investigated for the phantom data in Section 2, we analyze the effect of adding TV regularization to the angular sparse coding with different weighting α in (7), with $\alpha = 0$ being equivalent to the purely angular I-SR model. The algorithm used to solve the resulting optimization problem is the Split Bregman procedure outlined in Michailovich et al. (2011). Consistent with the phantom experiment of Fig. 4, we observe from Fig. 11 that adding spatial regularization again increases the total number of atoms for a given reconstruction error compared to the $\alpha = 0$ case. While the reconstructed signals may have a qualitatively spatial regular-



Fig. 10. Results of the proposed spatial-angular sparse coding using Kron-FISTA for Haar-SR and Curve-SR using an average of ~ 0.25 atoms/voxel compared to original signal. Curve-SR outperforms Haar-SR in this regime due to its additional directionality. We can see a drastically better reconstruction compared to the state-of-the-art at the same sparsity level in the top left of Fig. 2. This clearly shows that we can achieve accurate reconstruction with less than 1 atom/voxel.

ity which may result in a qualitatively better output. In comparison, the joint model we propose achieves better sparsity levels for comparable reconstruction errors. Note also that both algorithms displayed very similar performances in terms of running time for that particular experiment. We finally validate the approach in the case of a full 3D HARDI volume for a very sparse number of atoms. Fig. 12 presents the reconstruction error vs. sparsity results for the state-of-the-art framework (in which we set $\alpha = 0$ as the results of Fig. 11 suggest) versus the joint Haar-SR and Curve-SR. The plot shows again that curvelets outperforms Haar for high sparsity levels in the range of



Fig. 11. Comparison of the spatial-angular sparsity levels achieved by our proposed joint method with Haar-SR dictionaries and state-of-the-art method of Michailovich et al. (2011) with different spatial regularization parameters α , both applied on a 2D slice of a real HARDI scan. Recall $\alpha = 0$ corresponds to purely angular sparse coding, I-SR. Adding TV regularization results in an increase of the number of atoms for a given reconstruction error. The joint method achieves better sparsity levels than using separate sparsity penalties.



Fig. 12. Comparison of the spatial-angular sparsity level achieved by Haar-SR and Curve-SR with respect to the state-of-the-art I-SR on the entire 3D real HARDI volume. The curvelets provide a good reconstruction error with the sparsest number of atoms, in the range of 0.5–2 atoms/voxel. The state-of-the-art error is much larger in this sparsity range and only comparable in the predicted range of 6–8 atoms/voxel, consistent with the previously reported (Michailovich and Rathi, 2008; 2010b) for I-SR.

0.5–2 atoms/voxel. As expected and consistent with our phantom data experiment, the state-of-the-art I-SR has comparable reconstruction error in the range of 6–8 atoms/voxel. Fig. 13 shows the quality of reconstruction of I-SR, Haar-SR, and Curve-SR compared to the original signal using an average of ~1 atom/voxel. Haar-SR presents boxy regions while Curve-SR maintains a smoother reconstruction with a preservation of smaller detailed fiber tract regions. In contrast, the state-of-the-art I-SR is unable to model intricate fiber regions and is forced to set most voxels to zero atoms. All in all, we can see that using our proposed method, we can achieve much higher sparsity levels than the state-of-the-art, and accurate reconstructions using less than 1 atom/voxel. In terms of efficiency, Kron-FISTA was completed on the real HARDI data of size $V = 60 \times 60 \times 30$, G = 127 in 1.5 h for our sparsity level of interest using the fast 3D wavelet transform implemented in MATLAB.

6. Discussion and conclusion

In this work, we have demonstrated that by using a joint spatial-angular dictionary, we can obtain accurate HARDI reconstruction with spatial-angular sparsity levels of less than 1 atom per voxel, surpassing the limitations of state-of-the-art angular representations. This provides a new general reconstruction framework to achieve sparser dMRI representations than previously possible with optimal choices of spatial and angular dictionaries. In particular, we have shown promising sparsity results for HARDI from the combination of curvelet (spatial) and spherical ridgelet (angular) dictionaries, but other spatial and angular dictionaries may be chosen for other dMRI protocols like DSI or MS-HARDI. In future work, we aim to further optimize sparsity levels by learning a joint spatial-angular dictionary directly from dMRI data.

Furthermore, to efficiently solve this large-scale global sparse coding problem, we have proposed three novel extensions of popular sparse coding algorithms for the Kronecker dictionary setting. All strategies improve upon previously proposed algorithms by explicitly exploiting the separability of the dictionary and each may be beneficial depending on the problem regime and size of data. For our large-scale HARDI data, Kron-FISTA was the leader in speed. In future work, we will investigate other efficient active set methods such as the recent ORacle Guided Elastic Net (ORGEN) (You et al., 2016).

In addition to sparse coding, our spatial angular representation may have novel applications in other areas of dMRI processing such as denoising, feature extraction, global ODF non-



Fig. 13. Results of proposed spatial-angular sparse coding on real HARDI brain data using Kron-FISTA for I-SR, Haar-SR and Curve-SR using an average of ~0.5 atoms/voxel compared to original signal. Curve-SR outperforms Haar-SR in this sparsity range due to its directionality. The state-of-the-art I-SR is unable to compete at this sparsity level.

negativity, fiber tract segmentation, and tractography. However, our main application for spatial-angular sparse coding framework is the promising improvements of acquisition acceleration of dMRI through CS. One natural future extension of this work will be to incorporate our joint spatial-angular sparsifying dictionaries within a unified (k, q)-CS framework to subsample signal measurements both in k- and q-space. With the adequate design of joint sensing schemes, CS recovery results such as Candès et al. (2011) predict that the minimum number of samples needed for stable and accurate reconstruction is directly linked to the sparsity of the signal in the chosen dictionary, which argues in favor of the joint representation we have proposed.

Acknowledgements

This work is supported by Johns Hopkins University startup funds.

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.
- 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.
- Auría, 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., DiBella, E., 2013. Compressed sensing HARDI via rotation-invariant concise dictionaries, flexible K-space undersampling, and multiscale spatial regularity. In: IEEE International Symposium on Biomedical Imaging, pp. 9–12.
- Basser, P., Mattiello, J., LeBihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson. 103 (3), 247–254.
- Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM J. Imaging Sci. 2 (1), 183–202.
- Bilgic, B., Setsompop, K., Cohen-Adad, J., Yendiki, A., Wald, L.L., Adalsteinsson, E., 2012. Accelerated diffusion spectrum imaging with compressed sensing using adaptive dictionaries. Magn. Reson. Med. 68 (6), 1747–1754.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., 2010. Distributed optimization and statistical learning via the alternating direction method of multipliers. Found. Trends Mach. Learn. 3 (1), 1–122.
- Caiafa, C.F., Cichocki, A., 2013. Computing sparse representations of multidimensional signals using kronecker bases. Neural Comput. 25 (1), 186–220.
- Caiafa, C. F., Pestilli, F., 2015. Sparse multiway decomposition for analysis and modeling of diffusion imaging and tractography. arXiv:1505.07170.
- Candès, E., Demanet, L., Donoho, D., Ying, L., 2006. Fast discrete curvelet transforms. Multiscale Model. Simul. 5 (3), 861–899.
- Candès, E., Eldar, Y.C., Needell, D., Randall, P., 2011. Compressed sensing with coherent and redundant dictionaries. Appl. Comput. Harmon. Anal 31 (1), 59–73.
- Chao, T.-C., Chiou, J.-y.G., Maier, S.E., Madore, B., 2017. Fast diffusion imaging with high angular resolution. Magn. Reson. Med. 7, 696–706.
- Cheng, J., Jiang, T., Deriche, R., Shen, D., Yap, P.-T., 2013. Regularized spherical polar Fourier diffusion MRI with optimal dictionary learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 639–646.
- Cheng, J., Merlet, S., Caruyer, E., Ghosh, A., Jiang, T., Deriche, R., 2011. Compressive sensing ensemble average propagator estimation via 11 spherical polar fourier imaging. MICCAI Workshop on Computational Diffusion MRI.
- Cheng, J., Shen, D., Basser, P.J., Yap, P.T., 2015. Joint 6D kq space compressed sensing for accelerated high angular resolution diffusion MRI. In: Information Processing in Medical Imaging. Springer, pp. 782–793.
- Cheng, J., Shen, D., Yap, P.-T., Basser, P.J., 2015. Tensorial spherical polar Fourier diffusion MRI with optimal dictionary learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 174–182.
- Daducci, A., Canales-Rodri, E.J., Descoteaux, M., Garyfallidis, E., Gur, Y., Lin, Y.-C., Mani, M., Merlet, S., Paquette, M., Ramirez-Manzanares, A., et al., 2014. Quantitative comparison of reconstruction methods for intra-voxel fiber recovery from diffusion MRI. IEEE Trans. Med. Imaging 33 (2), 384–399.
- 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.
- Donoho, D.L., Elad, M., Temlyakov, V.N., 2006. Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Trans. Inf. Theory 52 (1), 6–18.
- Duarte, M.F., Baraniuk, R.G., 2012. Kronecker compressive sensing. IEEE Trans. Image Process. 21 (2), 494–504.
- Duarte-Carvajalino, J.M., Lenglet, C., Ugurbil, K., Moeller, S., Carin, L., Sapiro, G., 2012. A framework for multi-task Bayesian compressive sensing of DW-MRI. In: MIC-CAI Workshop on Computational Diffusion MRI, pp. 1–13.
- Duarte-Carvajalino, J.M., Lenglet, C., Xu, J., Yacoub, E., Ugurbil, K., Moeller, S., Carin, L., Sapiro, G., 2014. Estimation of the CSA-ODF using Bayesian compressed sensing of multi-shell HARDI. Magn. Reson. Med. 72 (5), 1471–1485.
- Goh, A., Lenglet, C., Thompson, P., Vidal, R., 2009. Estimating orientation distribution functions with probability density constraints and spatial regularity. In: Medical Image Computing and Computer Assisted Intervention, 5761, pp. 877–885.
- Goncalves, H.R., 2015. Accelerated Sparse Coding with Overcomplete Dictionaries for Image Processing Applications.
- Gramfort, A., Poupon, C., Descoteaux, M., 2012. Sparse DSI: Learning DSI structure for denoising and fast imaging. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 288–296.
- Gramfort, A., Poupon, C., Descoteaux, M., 2014. Denoising and fast diffusion imaging with physically constrained sparse dictionary learning. Med. Image Anal. 18 (1), 36–49.
- Hawe, S., Seibert, M., Kleinsteuber, M., 2013. Separable dictionary learning. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 438–445.
- Kuhnt, D., Bauer, M.H., Egger, J., Richter, M., Kapur, T., Sommer, J., Merhof, D., Nimsky, C., 2013. Fiber tractography based on diffusion tensor imaging compared with high-angular-resolution diffusion imaging with compressed sensing: initial experience. Neurosurgery 72 (0 1), 165.
- Kuhnt, D., Bauer, M.H.A., Sommer, J., Merhof, D., Nimsky, C., 2013. Optic radiation fiber tractography in glioma patients based on high angular resolution diffusion imaging with compressed sensing compared with diffusion tensor imaging - initial experience. PLoS ONE 8 (7), 1–7.
- Landman, B.A., Bogovic, J.A., Wan, H., ElShahaby, F.E.Z., Bazin, P.-L., Prince, J.L., 2012. Resolution of crossing fibers with constrained compressed sensing using diffusion tensor MRI. NeuroImage 59 (3), 2175–2186.
- 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., Jacob, M., Guidon, A., Magnotta, V., Zhong, J., 2015. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. Magn. Reson. Med. 73 (1), 126–138.

- McClymont, D., Teh, I., Whittington, H.J., Grau, V., Schneider, J.E., 2015. Prospective acceleration of diffusion tensor imaging with compressed sensing using adaptive dictionaries. Magn. Reson. Med..
- Menzel, M., Tan, E., Khare, K., Sperl, J., King, K., Tao, X., Hardy, C., Marinelli, L., 2011. Accelerated diffusion spectrum imaging in the human brain using compressed sensing. Magn. Reson. Med. 66 (5), 1226–1233.
- Merlet, S., Caruyer, E., Deriche, R., 2012. Parametric dictionary learning for modeling EAP and ODF in diffusion MRI. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 10–17.
- Merlet, S., Caruyer, E., Ghosh, A., Deriche, R., 2013. A computational diffusion MRI and parametric dictionary learning framework for modeling the diffusion signal and its features. Med. Image Anal. 17 (7), 830–843.
- Merlet, S., Cheng, J., Ghosh, A., Deriche, R., 2011. Spherical polar Fourier EAP and ODF reconstruction via compressed sensing in diffusion MRI. In: IEEE International Symposium on Biomedical Imaging. IEEE, pp. 365–371.
- Merlet, S., Deriche, R., 2010. Compressed sensing for accelerated EAP recovery in diffusion MRI. In: MICCAI, pp. Page–14.
- Merlet, S., 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., Rathi, Y., 2008. On approximation of orientation distributions by means of spherical ridgelets. In: IEEE International Symposium on Biomedical Imaging. IEEE, pp. 939–942.
- Michailovich, O., Rathi, Y., 2010. Fast and accurate reconstruction of HARDI data using compressed sensing. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 607–614.
- Michailovich, O., Rathi, Y., 2010. On approximation of orientation distributions by means of spherical ridgelets. IEEE Trans. Image Process. 19 (2), 461–477.
- Michailovich, O., Rathi, Y., Dolui, S., 2011. Spatially regularized compressed sensing for high angular resolution diffusion imaging. IEEE Trans. Med. Imaging 30 (5), 1100–1115.
- Ning, L., 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.V., Makris, N., Shenton, M.E., Westin, C.-F., Rathi, Y., 2016. A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging. NeuroImage 125, 386–400.
- Ouyang, Y., Chen, Y., Wu, Y., 2014. Vectorial total variation regularisation of orientation distribution functions in diffusion weighted MRI. Int. J. Bioinform. Res. Appl. 10 (1), 110–127.
- Ouyang, Y., Chen, Y., Wu, Y., Zhou, H.M., 2013. Total variation and wavelet regularization of orientation distribution functions in diffusion MRI. Inverse Probl. Imaging 7 (2), 565–583.
- Paquette, M., Merlet, S., Gilbert, G., Deriche, R., Descoteaux, M., 2015. Comparison of sampling strategies and sparsifying transforms to improve compressed sensing diffusion spectrum imaging. Magn. Reson. Med. 73 (1), 401–416.
- Qi, N., Shi, Y., Sun, X., Yin, B., 2016. TenSR: Multi-dimensional tensor sparse representation. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 5916–5925.
- Rathi, Y., Michailovich, O., Laun, F., Setsompop, K., Grant, P.E., Westin, C.-F., 2014. Multi-shell diffusion signal recovery from sparse measurements. Med. Image Anal. 18 (7), 1143–1156.
- Rathi, Y., Michailovich, O., Setsompop, K., Bouix, S., Shenton, M., Westin, C.-F., 2011. Sparse multi-shell diffusion imaging. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 58–65.
- Reese, T.G., Benner, T., Wang, R., Feinberg, D.A., Wedeen, V.J., 2009. Halving imaging time of whole brain diffusion spectrum imaging and diffusion tractography using simultaneous image refocusing in EPI. J. Magn. Reson. Imaging 29 (3), 517–522.
- Rivenson, Y., S., A., 2009. Compressed imaging with a separable sensing operator. IEEE Signal Process. Lett. 16 (6), 449-452.
- Schwab, E., Afsari, B., Vidal, R., 2012. Estimation of non-negative ODFs using eigenvalue distribution of spherical functions. In: Medical Image Computing and Computer Assisted Intervention, 7511, pp. 322–330.
- Schwab, E., Vidal, R., Charon, N., 2016. Spatial-Angular Sparse Coding for HARDI. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 475–483.
- Setsompop, K., Cohen-Adad, J., Gagoski, B., Raij, T., Yendiki, A., Keil, B., Wedeen, V.J., Wald, L.L., 2012. Improving diffusion MRI using simultaneous multi-slice echo planar imaging. Neuroimage 63 (1), 569–580.
- 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: Information Processing in Medical Imaging. Springer, pp. 375–386.
- Sun, J., Xie, Y., Ye, W., Ho, J., Entezari, A., Blackband, S.J., Vemuri, B.C., 2013. Dictionary learning on the manifold of square root densities and application to reconstruction of diffusion propagator fields. In: International Conference on Information Processing in Medical Imaging. Springer, pp. 619–631.
- Tristán-Vega, A., Westin, C.-F., 2011. Probabilistic ODF estimation from reduced HARDI data with sparse regularization. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 182–190.
- Tropp, J., 2004. Greed is good: algorithmic results for sparse approximation. IEEE Trans. Inf. Theory 50 (10), 2231–2242.

Tuch, D., 2004. Q-ball imaging. Magn. Reson. Med. 52 (6), 1358-1372.

- Wedeen, V., Hagmann, P., Tseng, W., Reese, T., Weisskoff, R., 2005. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn. Reson. Med. 54 (6), 1377–1386.
- Wolfers, S., Schwab, E., Vidal, R., 2014. Nonnegative ODF estimation via optimal constraint selection. In: IEEE International Symposium on Biomedical Imaging, pp. 734–737.
 Wu, Y.-C., Alexander, A.L., 2007. Hybrid diffusion imaging. NeuroImage 36 (3),
- Wu, Y.-C., Alexander, A.L., 2007. Hybrid diffusion imaging. NeuroImage 36 (3), 617–629.
- Ye, C., 2016. Fiber orientation estimation using nonlocal and local information. In: Medical Image Computing and Computer Assisted Intervention. Springer, pp. 97–105.
- Ye, C., Prince, J.L., 2017. Probabilistic tractography using Lasso bootstrap. Med. Image Anal. 35, 544–553.
- Ye, W., Vemuri, B.C., Entezari, A., 2012. An over-complete dictionary based regularized reconstruction of a field of ensemble average propagators. In: IEEE International Symposium on Biomedical Imaging. Springer, pp. 940–943.
- Active reconstruction of a field of ensemble average propagators. In: IEEE International Symposium on Biomedical Imaging. Springer, pp. 940–943.
 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.
 Yoldemir, B., Bajammal, M., Abugharbieh, R., 2014. Dictionary based super-resolu-
- Yoldemir, B., Bajammal, M., Abugharbieh, R., 2014. Dictionary based super-resolution for diffusion MRI. In: MICCAI Workshop on Computational Diffusion MRI. Springer, pp. 203–213.
- You, C., Li, C.-G., Robinson, D., Vidal, R., 2016. Oracle based active set algorithm for scalable elastic net subspace clustering. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 3928–3937.