

Spatial-Angular Sparse Coding for HARDI

Evan Schwab^(✉), René Vidal, and Nicolas Charon

Center for Imaging Science, Johns Hopkins University, Baltimore, USA
eschwab3@jhu.edu

Abstract. High angular resolution diffusion imaging (HARDI) can produce better estimates of fiber orientation and richer sets of features for disease classification than diffusion tensor imaging. However, existing HARDI reconstruction algorithms require a large number of gradient directions, making the acquisition time too long to be clinically viable. State-of-the-art compressed sensing methods can reduce the number of measurements needed for accurate reconstruction by exploiting angular sparsity at each voxel, but the global sparsity level is therefore bounded below by the number of voxels. In this work, we aim to find a significantly sparser representation of HARDI by exploiting redundancies in both the spatial and angular domains jointly with a global HARDI basis. However, this leads to a massive global optimization problem over the whole brain which cannot be solved using existing sparse coding methods. We present a novel Kronecker extension to ADMM that exploits the separable spatial-angular structure of HARDI data to efficiently find a globally sparse reconstruction. We validate our method on phantom and real HARDI brain data by showing that we can achieve accurate reconstructions with a global sparsity level corresponding to less than one atom per voxel, surpassing the absolute limit of the state-of-the-art.

1 Introduction

Diffusion magnetic resonance imaging (dMRI) is a 6D neuroimaging modality that produces 3D q-space signals at every voxel of a 3D brain MRI volume and can be used to estimate the orientation and integrity of neuronal fiber bundles *in vivo*. Diffusion tensor imaging (DTI), which models the probability of water diffusion in each voxel with a 3D Gaussian distribution, requires a relatively low number of q-space signal measurements. While this makes DTI well-suited for clinical research, the diffusion tensor cannot model multiple fiber orientations crossing in a single voxel.

High angular resolution diffusion imaging (HARDI) can provide more accurate estimations of anatomical fiber networks than DTI by estimating a higher-order probability distribution function. However, HARDI reconstruction algorithms require a larger number of q-space directions, which leads to a significantly longer patient scan time in comparison to DTI. As a consequence, HARDI has not been broadly accepted as a clinically viable dMRI protocol. Therefore, reducing HARDI scan times to the rate of DTI while maintaining accurate orientation estimation is an important open question in HARDI research.

Compressed sensing (CS) is a very useful tool to solve problems of this type and hinges on finding a sparse representation of a signal with respect to some basis. If such a sparse representation exists, CS can recover nearly perfect signals with sub-Nyquist sampling, allowing the potential to reduce signal acquisition time. Many approaches [9] have applied CS to dMRI protocols like multi-shell/single-shell HARDI and DSI to sparsely reconstruct signals, estimate orientation distribution functions (ODFs), fiber orientation distribution functions (FODs) [4], and ensemble average propagators (EAPs) [7], using dictionary learning or fixed sparsifying bases such as spherical ridgelets/wavelets [15], spherical polar Fourier bases, spherical Fourier-Bessel bases, directional radial bases, higher order tensors, and many more. These approaches reduce the number of q-space measurements needed from hundreds to tens by representing a q-space signal with as few q-space basis atoms as possible. However, since these methods compute one sparse representation for each and every voxel, the sparsest representation they can possibly achieve is with a single atom per voxel. Since spatial correlations between atoms are not exploited, these methods are still likely to represent a brain volume with millions of possibly redundant parameters.

In practice, measurements from neighboring voxels share much of the same information, hence there may exist further redundancies in the spatial domain when signals are modeled sparsely over the angular domain alone. Therefore, to reduce the number of measurements further, recent work [3, 8, 10, 11] aims to apply CS to DSI/HARDI in the joint (k,q)-space. These methods apply joint (k,q) undersampling but still only apply sparse coding in the angular domain. Some works add sparse spatial regularization such as total-variation, yet with disjoint spatial and angular terms, global sparsity is still limited by the size of the data and (k,q)-CS may not be fully utilized.

In this paper, we propose to model HARDI signals using a global spatial-angular basis. However, because of the large size and complexity of HARDI data, optimizing over an entire HARDI volume is a computationally challenging problem. Our main contribution is an efficient joint spatial-angular sparse coding algorithm that exploits a separable model of the spatial and angular domains of HARDI data. With this proposed framework we aim to efficiently find a significantly sparser HARDI representation than state-of-the-art voxel-based methods can theoretically allow. In future work, joint spatial-angular sparse coding will allow us to more naturally apply (k,q)-CS with joint undersampling [8, 11] to further reduce the total number of HARDI measurements.

2 HARDI Data Representation

Angular (Voxel-Based) HARDI Representation. For each voxel v in a HARDI brain volume $\Omega \subset \mathbb{R}^3$, q-space measurements are acquired at gradient directions $\vec{g} \in \mathbb{S}^2$ and are modeled by an angular (spherical) basis, $\{\gamma_i(\vec{g})\}_{i=1}^{N_r}$, with N_r atoms such that $s_v(\vec{g}) = \sum_{i=1}^{N_r} a_i \gamma_i(\vec{g})$ where $s_v(\vec{g})$ denotes the HARDI

signal at voxel v measured at gradient direction \vec{g} . Classical HARDI reconstruction methods model the q-space signals from each voxel separately and add a regularization term \mathcal{R} to enforce desirable properties such as spatial coherence, ODF non-negativity, or sparsity. Some recent methods [1, 5, 12, 13] have considered simultaneous voxel-based reconstruction over an entire volume by solving

$$A^* = \arg \min_A \|S - \Gamma A\|_F^2 + \lambda \mathcal{R}(A), \quad (1)$$

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, $A = [a_1 \dots a_V] \in \mathbb{R}^{N_r \times V}$ is the concatenation of coefficients and $\Gamma \in \mathbb{R}^{G \times N_r}$ is the discretization of the basis γ . While these methods attempt to reduce redundancies by adding spatial regularization, signal reconstruction still only operates on angular basis Γ at every voxel.

Spatial-Angular HARDI Representation. In this work, we propose to model the HARDI signal $\mathcal{S}(v, \vec{g})$ based on a single global basis, say $\varphi(v, \vec{g})$, to explicitly reduce redundancies in both the spatial and angular domains. However, typical HARDI contains on the order of $V \approx 100^3$ voxels each with $G \approx 100$ q-space measurements for a total of $100^4 \approx 100$ million signal measurements. Since many sparse coding applications often use bases that are over-redundant, this leads to a massive matrix Φ of size greater than $100^4 \times 100^4$. Therefore efficiently optimizing over a global basis is a very difficult problem. To overcome this challenge, we introduce additional structure on the dictionary atoms by considering separable functions over Ω and \mathbb{S}^2 , namely a set of atoms of the form $\varphi(v, \vec{g})_{i,j} = (\psi_j(v)\gamma_i(\vec{g}))$, where $\psi(v)$ is a spatial basis for Ω with N_ψ atoms. The HARDI signal may then be decomposed as:

$$\mathcal{S}(v, \vec{g}) = \sum_{i=1}^{N_\gamma} \sum_{j=1}^{N_\psi} c_{i,j} \psi_j(v) \gamma_i(\vec{g}) = \sum_{k=1}^{N_\psi N_\gamma} c_k \varphi_k(v, \vec{g}), \quad (2)$$

where $c = [c_k] \in \mathbb{R}^{N_\psi N_r}$ is the vectorization of $C = [c_{i,j}] \in \mathbb{R}^{N_r \times N_\psi}$. In discretized form, our global basis φ is the separable Kronecker product matrix $\Phi \triangleq \Psi \otimes \Gamma \in \mathbb{R}^{V \times G \times N_\psi N_r}$ with $\Psi \in \mathbb{R}^{V \times N_\psi}$ and $\Gamma \in \mathbb{R}^{G \times N_r}$ such that

$$s = \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_V \end{pmatrix} = \begin{pmatrix} \Psi_{1,1}\Gamma & \Psi_{1,2}\Gamma & \cdots & \Psi_{1,N_\psi}\Gamma \\ \Psi_{2,1}\Gamma & \Psi_{2,2}\Gamma & \cdots & \Psi_{2,N_\psi}\Gamma \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_{V,1}\Gamma & \Psi_{V,2}\Gamma & \cdots & \Psi_{V,N_\psi}\Gamma \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_{N_\psi N_r} \end{pmatrix} = \Phi c. \quad (3)$$

Alternatively, in matrix form, (3) can be written compactly as $S = \Gamma C \Psi^\top$. In the special case of $\Psi = I_V$, the identity, we can see this leads to the state-of-the-art voxel-based formulation (1) with $C \equiv A$.

For HARDI, ODFs, p , can be estimated globally for all voxels with a single equation $p(v, \vec{x}; c) = \frac{1}{4\pi} + \tilde{\Phi}(v, \vec{x})c$, where $\tilde{\Phi}(v, \vec{x}) \triangleq \Psi(v) \otimes \tilde{\Gamma}(\vec{x})$ and $\tilde{\Gamma}(\vec{x})$ is the transformed angular basis into the space of ODFs for $\vec{x} \in \mathbb{S}^2$. This novel global formulation could provide a nice framework for HARDI applications like

enforcing global non-negativity, extracting global features or global fiber segmentation. Though this current spatial-angular formulation may be specific to HARDI, our framework is generalizable to any dMRI protocol such as DSI, or multi-shell methods, by choosing an appropriate Γ which represents these data in q -space.

3 Efficient Globally Sparse HARDI Reconstruction

With our proposed spatial-angular basis representation of HARDI, the goal of this paper is to accurately reconstruct an entire HARDI volume with fewer atoms than the state-of-the-art voxel-based methods can theoretically allow. To find a globally sparse representation c , we aim to solve the l_1 minimization problem:

$$c^* = \arg \min_c \frac{1}{2} \|s - \Phi c\|_2^2 + \lambda \|c\|_1, \quad (4)$$

where $\lambda > 0$ is the sparsity trade-off parameter. The Alternating Direction Method of Multipliers (ADMM) [2] is a popular method for solving (4), however, its application in the case of a large dictionary Φ remains prohibitive. To reduce computation, we first note that when using over-redundant dictionaries (i.e., $V < N_\Psi, G < N_\Gamma$), it is more efficient to apply ADMM to the dual of (4) so that we can switch from calculating $\Phi^\top \Phi$ of size $N_\Psi N_\Gamma \times N_\Psi N_\Gamma$ to $\Phi \Phi^\top$ of size $VG \times VG$. The dual of (4) is:

$$\max_\alpha -\frac{1}{2} \|\alpha\|_2^2 + \alpha^\top s \quad \text{s.t.} \quad \|\Phi^\top \alpha\|_\infty \leq \lambda, \quad (5)$$

and update equations of the Dual ADMM (DADMM) [6] are given by:

$$\alpha_{k+1} = (I + \eta \Phi \Phi^\top)^{-1} (s - \Phi (c_k - \eta \nu_k)) \quad (6)$$

$$\nu_{k+1} = P_\lambda^\infty \left(\frac{c_k}{\eta} + \Phi^\top \alpha_{k+1} \right) \quad (7)$$

$$c_{k+1} = \text{shrink}_{\lambda \eta} (c_k + \eta \Phi^\top \alpha_{k+1}), \quad (8)$$

where $P_\lambda^\infty(x)$ is an element-wise projection to $[-\lambda, \lambda]$, $\text{shrink}_\rho(x)$ is the soft-threshold operator and $\eta > 0$ is an optimization parameter. The globally sparse output vector c^* minimizes the primal problem (4). DADMM reduces inner product computations to matrices of smaller size $G \times V$ instead of $N_\Gamma \times N_\Psi$ and is therefore not dependent on the size of basis but only on the size of the data. Furthermore, soft-thresholding is now done directly on c , which reduces the number of iterations for reaching a sparsity level by building up from 0 atoms instead of descending down from the total $N_\Psi N_\Gamma$ atoms. However, this naïve formulation with large Φ still has complexity $O(GV N_\Gamma N_\Psi)$ per iteration. Even submitting Φ into memory may be an issue, as well as the expensive cost of an inverse. We address this issue by exploiting the separability of Φ to perform computations with the much smaller Ψ and Γ . Our proposed method, called Kronecker DADMM, exploits the Kronecker product in two ways:

1. Kronecker SVD. Computing $(I + \eta\Phi\Phi^\top)^{-1}$ for large Φ is challenging and even taking an SVD to reduce the inverse to a diagonal of singular values is $O((GV)^2 N_\Gamma N_\Psi)$. Instead, we can exploit the Kronecker product and compute separate SVDs of the smaller $\Psi\Psi^\top$ and $\Gamma\Gamma^\top$. Let $\Psi\Psi^\top = U_\Psi \Sigma_\Psi U_\Psi^\top$ and $\Gamma\Gamma^\top = U_\Gamma \Sigma_\Gamma U_\Gamma^\top$, then we have $(I + \eta\Phi\Phi^\top)^{-1} = (U_\Psi \otimes U_\Gamma)(I + \eta(\Sigma_\Psi \otimes \Sigma_\Gamma))^{-1}(U_\Psi \otimes U_\Gamma)^\top$, where the inverse is now simply taken over a diagonal matrix. Computing these two SVDs is now of complexity $O(V^2 N_\Psi + G^2 N_\Gamma)$. In the case of tight frames, such as Wavelets, where $\Psi\Psi^\top = I$, we can significantly reduce computations to only involve Γ . We then simplify computations for (6) by pre-multiplication, setting $\alpha' \triangleq (U_\Psi \otimes U_\Gamma)^\top \alpha$, $s' \triangleq (U_\Psi \otimes U_\Gamma)^\top s$ and $\Phi' \triangleq (U_\Psi \otimes U_\Gamma)^\top \Phi = (U_\Psi \otimes U_\Gamma)^\top (\Psi \otimes \Gamma) = (U_\Psi^\top \Psi) \otimes (U_\Gamma^\top \Gamma) \triangleq (\Psi' \otimes \Gamma')$. [6] proves that we can now replace our original variables with α' , s' , and Φ' and our primal and dual optimization problems do not change.

2. Kronecker Matrix Formulation. We can further reduce the size of our problem and avoid any computations with Φ' by using the Kronecker product matrix formulation of (3) to write $\Phi'(c_k - \eta\nu_k) = \Gamma'(C_k - \eta\mathcal{V}_k)\Psi'^\top$, where C and \mathcal{V} are the $N_\Gamma \times N_\Psi$ matrix forms of c and ν . Likewise, $\Phi'^\top \alpha' = \Gamma'^\top A' \Psi'$, where A' is the $G \times V$ matrix form of α' and we can pre-compute $S' = U_\Gamma^\top S U_\Psi$ where S is the $G \times V$ matrix form of s . To fit the $G \times V$ matrix dimensions, we simplify diagonal $(I + \eta(\Sigma_\Psi \otimes \Sigma_\Gamma))^{-1}$ to $\Sigma_\eta^{-1} \triangleq 1/(1 + \eta\Sigma)$ where $\Sigma \triangleq d_\Gamma d_\Psi^\top$ and d_Ψ and d_Γ are the diagonals of Σ_Ψ and Σ_Γ . This Kronecker matrix formulation has a significantly reduced complexity of $O(GVN_\Psi)$ per iteration compared with $O(GVN_\Gamma N_\Psi)$ for the naïve approach. Furthermore, in the case where Ψ is a tight frame transformation, such as Wavelets, for which $\Psi' = \Psi$, fast decomposition and reconstruction algorithms can be used to replace multiplication by Ψ and Ψ^\top reducing complexity to $O(GV \log_2(N_\Psi))$.

Algorithm 1. (Kron-DADMM)

Precompute: $S', \Gamma', \Psi', \Sigma$.

Initialize: $k = 0, C_0 = 0, \mathcal{V}_0 = 0$. Choose: $\eta_0, \lambda, \epsilon$.

while Duality Gap $> \epsilon$ **do**

1: $A'_{k+1} = \Sigma_{\eta_k}^{-1} \circ (S' - \Gamma'(C_k - \eta_k \mathcal{V}_k)\Psi'^\top)$;

2: $\mathcal{V}_{k+1} = P_\lambda^\infty(\frac{1}{\eta_k} C_k + \Gamma'^\top A'_{k+1} \Psi')$;

3: $C_{k+1} = \text{shrink}_{\lambda\eta_k}(C_k + \eta_k \Gamma'^\top A'_{k+1} \Psi')$;

4: $\eta_{k+1} = \min(\frac{\|S' - \Gamma' C_{k+1} \Psi'^\top\|_F^2 \|C_{k+1}\|_F^2}{2\lambda \|\Gamma' C_{k+1} \Psi'^\top\|_F^2}, \frac{\|S\|_1}{\lambda GV})$

5: $\Sigma_{\eta_{k+1}}^{-1} = 1/(1 + \eta_{k+1} \Sigma)$

6: $k = k + 1$;

end while

Return: globally sparse representation C^*

Our proposed algorithm for globally sparse HARDI reconstruction (Kron-DADMM) is presented in Algorithm 1. The symbol “ \circ ” in Step 1 denotes element-wise matrix multiplication. We follow [6] to update penalty parameter η and stop when the duality gap is sufficiently small.

4 Experiments

Data Sets. We provide experiments on the ISBI 2013 HARDI Reconstruction Challenge Phantom dataset, a $50 \times 50 \times 50$ volume consisting of 20 phantom fibers crossing intricately within an inscribed sphere, measured with $G = 64$ gradient directions (SNR = 30). Figures 1 and 2 show quantitative signal error vs. sparsity and qualitative ODF estimations at specific sparsity levels, respectively. We also experimented on a real $128 \times 128 \times 26$ HARDI brain volume with $G = 384$ from the Hippocampal Connectivity Project (FOV: 192, resolution: 1.5 mm isotropic, b-value: 1400 s/mm², TR/TE: 3500/86). Figure 3 shows sparse reconstruction using Kron-DADMM with Haar-SR compared to full voxel-based SH reconstruction. We can achieve a good reconstruction with ~ 2 atoms per voxel in about 2.5 h.

Choice of Spatial and Angular Bases. *Spatial* (Ψ): A popular choice of sparsifying basis for MRI is a wavelet basis. For our experiments we compared Haar and Daubechies wavelets to get an indication of which can more sparsely represent the spatial organization of HARDI. Importantly, we also compared to state-of-the-art voxel-wise methods by simply choosing the identity I as the spatial basis. We compare these basis choices in terms of sparsity and reconstruction error in Fig. 1 (left). *Angular* (I): The over-complete spherical ridgelet/wavelet (SR/SW) basis pair [15] has been shown to sparsely model HARDI signals/ODFs. We also compare this to the popular SH basis, though for order L the SH is a low-pass truncation and does not exude sparse signals. With order $L = 4$, SH and SR have $N_I = 15$ and $N_I = 1169$ atoms, respectively. We compare these angular choices in Fig. 1 (right). As a note, these basis choices are preliminary and future work will involve exploring more advanced basis options to increase sparsity.




Fig. 1. Phantom Data. Left: Comparison of various spatial basis choices using Kron-DADMM paired with the SR angular basis. Haar-SR achieves the lowest residual (~ 0.074) with the sparsest number of coefficients (~ 0.5). The black line is voxel-based angular reconstruction where the identity I is the chosen spatial basis. Voxel-based is unable to achieve sparsity below 1 atom/voxel without forcing some voxels to 0 atoms. Right: Comparison between Kron-OMP and Kron-DADMM. Kron-OMP takes 40 h to reach 1 atom/voxel, while Kron-DADMM takes 40 min and provides better accuracy. SR outperforms SH basis as expected.




Fig. 2. Phantom Data. Comparison of Kron-DADMM with Haar-SR (middle) and voxel-based I-SR (right) against a full voxel-based least squares reconstruction with SH (left). The global reconstruction provides a more accurate signal with less than 1 atom per voxel while the voxel-based sparse reconstruction has difficulty estimating crossing fibers at this sparsity level and is forced to model isotropic ODFs with 0 atoms.

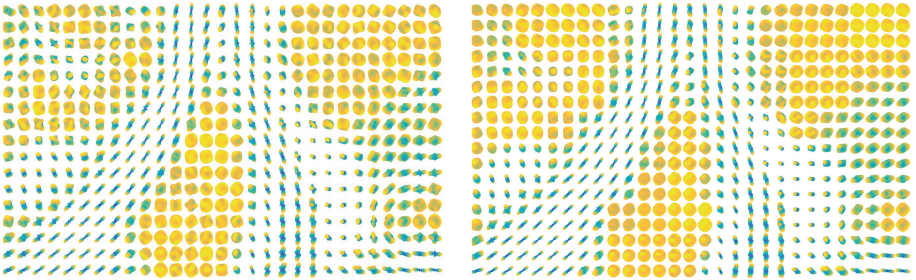


Fig. 3. Real Data. Global sparse reconstruction via Kron-DADMM with Haar-SR basis (right) compared to full voxel-based least squares reconstruction with SH basis (left). We can achieve an accurate reconstruction with only 2.23 atoms/voxel.

Comparison with a Baseline Algorithm. Orthogonal Matching Pursuit (OMP) is a widely used algorithm to approximate a solution to the l_0 problem by greedily selecting and orthogonalizing the K basis atoms that are most correlated with signal s . The Kronecker OMP (Kron-OMP) proposed in [14] exploits bases with separable structure but the method still needs to orthogonalize a $K \times K$ matrix where the sparsity level K empirically approaches the number of voxels $V \approx 100^4$. Because of this, our implementation takes on the order of 40 h to optimize over the phantom dataset compared to 40 min for Kron-DADMM. The results are presented in Fig. 1 (right).

5 Conclusion

We have presented a new efficient algorithm for globally sparse reconstruction of HARDI using a global basis representation which exploits spatial-angular

separability to significantly reduce computational complexity. Our experiments show that greater sparsity for the representations may be achieved using spatial-angular bases instead of voxel-based approaches. So far, these were conducted as a proof-of-concept with dictionaries involving very simple spatial wavelet bases such as Haar, but the versatility of the algorithm enables the use of possibly more adequate directional Wavelets (e.g. shearlets, curvelets) or dictionary learning strategies, which are both important directions for future work. Our next step is to develop spatial-angular sensing matrices to jointly subsample (k,q) -space using a form of Kronecker CS. Finally, a globally sparse representation can be utilized in many other areas of HARDI applications including fiber segmentation, global tractography, global feature extraction, and sparse disease classification.

Acknowledgements. This work funded by JHU startup funds.

References

1. Auría, A., Daducci, A., Thiran, J.-P., Wiaux, Y.: Structured sparsity for spatially coherent fibre orientation estimation in diffusion MRI. *NeuroImage* **115**, 245–255 (2015)
2. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**(1), 1–122 (2010)
3. Cheng, J., Shen, D., Basser, P.J., Yap, P.-T.: Joint 6D k - q space compressed sensing for accelerated high angular resolution diffusion MRI. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) *IPMI 2015*. LNCS, vol. 9123, pp. 782–793. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-19992-4_62](https://doi.org/10.1007/978-3-319-19992-4_62)
4. Feng, Y., Wu, Y., Rathi, Y., Westin, C.-F.: Sparse deconvolution of higher order tensor for fiber orientation distribution estimation. *Artif. Intell. Med.* **65**(3), 229–238 (2015)
5. Goh, A., Lenglet, C., Thompson, P.M., Vidal, R.: Estimating orientation distribution functions with probability density constraints and spatial regularity. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI 2009*, Part I. LNCS, vol. 5761, pp. 877–885. Springer, Heidelberg (2009)
6. Goncalves, H.R.: Accelerated sparse coding with overcomplete dictionaries for image processing applications. Ph.D. thesis (2015)
7. Gramfort, A., Poupon, C., Descoteaux, M.: Denoising and fast diffusion imaging with physically constrained sparse dictionary learning. *Med. Image Anal.* **18**(1), 36–49 (2014)
8. Mani, M., Jacob, M., Guidon, A., Magnotta, V., Zhong, J.: Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. *Magn. Reson. Med.* **73**(1), 126–138 (2015)
9. Ning, L., et al.: 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 (2015)
10. Ning, L., Setsompop, K., Michailovich, O.V., Makris, N., Shenton, M.E., Westin, C.-F., Rathi, Y.: A joint compressed-sensing and super-resolution approach for very high-resolution diffusion imaging. *NeuroImage* **125**, 386–400 (2016)

11. Sun, J., Sakhaee, E., Entezari, A., Vemuri, B.C.: Leveraging EAP-sparsity for compressed sensing of MS-HARDI in (\mathbf{k}, \mathbf{q}) -space. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) IPMI 2015. LNCS, vol. 9123, pp. 375–386. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-19992-4_29](https://doi.org/10.1007/978-3-319-19992-4_29)
12. Ouyang, Y., Chen, Y., Wu, Y.: Vectorial total variation regularisation of orientation distribution functions in diffusion weighted MRI. *Int. J. Bioinform. Res. Appl.* **10**(1), 110–127 (2014)
13. Ouyang, Y., Chen, Y., Wu, Y., Zhou, H.M.: Total variation and wavelet regularization of orientation distribution functions in diffusion MRI. *Inverse Prob. Imaging* **7**(2), 565–583 (2013)
14. Rivenson, Y., Adrian, S.: Compressed imaging with a separable sensing operator. *IEEE Signal Process. Lett.* **16**(6), 449–452 (2009)
15. Tristán-Vega, A., Westin, C.-F.: Probabilistic ODF estimation from reduced HARDI data with sparse regularization. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011. LNCS, vol. 6892, pp. 182–190. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-23629-7_23](https://doi.org/10.1007/978-3-642-23629-7_23)