

# Dictionary Based Super-Resolution for Diffusion MRI

Burak Yoldemir, Mohammad Bajammal, and Rafeef Abugharbieh

**Abstract** Diffusion magnetic resonance imaging (dMRI) provides unique capabilities for non-invasive mapping of fiber tracts in the brain. It however suffers from relatively low spatial resolution, often leading to partial volume effects. In this paper, we propose to use a super-resolution approach based on dictionary learning for alleviating this problem. Unlike the majority of existing super-resolution algorithms, our proposed solution does not entail acquiring multiple scans from the same subject which renders it practical in clinical settings and applicable to legacy data. Moreover, this approach can be used in conjunction with any diffusion model. Motivated by how functional connectivity (FC) reflects the underlying structural connectivity (SC), we quantitatively validate our results by investigating the consistency between SC and FC before and after super-resolving the data. Based on this scheme, we show that our method outperforms traditional interpolation strategies and the only existing single image super-resolution method for dMRI that is not dependent on a specific diffusion model. Qualitatively, we illustrate that fiber tracts and track-density maps reconstructed from super-resolved dMRI data reveal exquisite details beyond what is achievable with the original data.

## 1 Introduction

Diffusion magnetic resonance imaging (dMRI) based tractography provides a powerful non-invasive in vivo tool for localizing white matter tracts in the brain. Accurate mapping of white matter fiber tracts is important in gaining insights into the brain function since fiber tracts act as a substrate enabling communication between brain regions [8]. However, accuracy of the reconstructed fiber tracts is often hampered by the inherently low resolution of dMRI data. Currently achievable spatial dMRI resolution is around  $2 \times 2 \times 2 \text{ mm}^3$ , while the actual neuronal fiber diameter is on the order of  $1 \mu\text{m}$  [18]. A voxel can thus comprise several distinct fiber bundles with differing orientations, leading to partial volume averaging

---

B. Yoldemir (✉) • M. Bajammal • R. Abugharbieh  
Biomedical Signal and Image Computing Laboratory, The University of British Columbia,  
Vancouver, BC, Canada  
e-mail: [buraky@ece.ubc.ca](mailto:buraky@ece.ubc.ca); [bajammal@ece.ubc.ca](mailto:bajammal@ece.ubc.ca); [rafeef@ece.ubc.ca](mailto:rafeef@ece.ubc.ca)

[2]. At such locations, diffusion information typically becomes ambiguous, and tractography is often falsely terminated. Therefore, increasing the spatial resolution in dMRI holds great promise towards more accurate delineation of fiber tracts. There are practical limitations in increasing the resolution of the acquired data directly, such as reduced signal-to-noise ratio (SNR) and prolonged scanning time [13]. Such limitations motivate the search for post-processing solutions for increasing the spatial resolution, such as super-resolution techniques.

Super-resolution techniques have been previously adopted to increase the spatial detail in dMRI. In the literature, the term super-resolution is used for two distinct classes of methods which follow different paradigms. The first class of methods are based on performing multiple low-resolution acquisitions, followed by the fusion of the information in these images to generate high-resolution images. To this end, fusing images spatially shifted at sub-voxel level [16], as well as fusing multiple anisotropic images with high resolution only along one axis [17, 18] have been explored. In a fairly similar spirit, combining diffusion-weighted (DW) images acquired at two different resolutions to infer high-resolution diffusion parameters using a Bayesian model has also been proposed [20]. The inherent drawback of these approaches is the dependence on a specific acquisition protocol, limiting their usability in general settings. The second class of methods do not require multiple acquisitions, and these are typically based on examples or priors about the correspondence between low and high resolution images. Falling in this category, an approach to reconstruct diffusion tensors at a resolution higher than the underlying DW images using a single dMRI acquisition has been recently proposed [7]. Even though this method eliminates the need for multiple acquisitions, it is only geared towards estimating diffusion tensors, and cannot be easily extended to higher order diffusion models such as orientation distribution functions (ODFs). To the best of our knowledge, the only previous work that tackled the problem of super-resolving dMRI data from a single acquisition independent of the diffusion model was by Coupé et al. [4]. Specifically, the authors showed that super-resolving  $b = 0$  (non-diffusion-weighted) image using a locally adaptive patch-based strategy, and using this high-resolution  $b = 0$  image to drive the reconstruction of DW images outperforms upsampling of dMRI data using classical interpolation methods. Beyond these two classes, a new perspective to gain spatial resolution in dMRI has been proposed which is termed as super-resolution track-density imaging (TDI) [3]. This approach is fundamentally different than the aforementioned super-resolution methods in the sense that the aim is to generate high resolution track density maps through counting the number of tracts present in each element of a sub-voxel grid, rather than super-resolving the DW volumes prior to tractography.

In this paper, we employ a super-resolution approach [23] for dMRI that does not require more than a single acquisition per subject. Importantly, we apply this method on DW images before the diffusion modeling step, removing the limitation of applicability to a specific model such as diffusion tensors. The technique is based on sparse coding of DW images via dictionary learning [23]. We note that

similar methods following the sparse coding principle have been investigated before for natural images [6, 22], however the applicability of such an approach on DW images and its added value remain unknown. Given a set of training images, we start with constructing an over-complete dictionary representing the data. Another over-complete dictionary is then constructed from the downsampled versions of these training images. Notably, these two dictionaries are constructed such that the coding vectors modeling downsampled and original data as sparse linear combinations of the learned dictionary atoms are the same, hence the correspondence between low and high resolution images are automatically captured. We then exploit this correspondence between the two dictionaries to super-resolve a new input image to a higher resolution. The advantage of this method is three-fold. First, the super-resolved DW images can be used with any diffusion model as permitted by the number of gradient directions in the original dataset. Second, this method does not rely on repeated acquisitions from the same subject, allowing it to be used with legacy data and under various clinical acquisition schemes. Third, this method may still be readily applied when the imaging protocol involves multiple acquisitions, as an additional step after reconstructing a single image from multiple low resolution acquisitions [16–18].

We qualitatively validate our proposed approach by comparing the fiber tracts and track-density maps reconstructed from the original and super-resolution data. In the absence of ground truth connectivity information, in order to provide a meaningful basis for quantitative comparison, we use the consistency between intra-subject structural connectivity (SC) and functional connectivity (FC) estimates inferred from dMRI and resting state functional MRI (RS-fMRI) data, respectively. Our rationale is that FC is inherently shaped by the wiring of the brain [8, 19]. Therefore, a more accurate estimate of SC would presumably increase the SC-FC correlation. In addition, we also examine the number of fiber tracts for more insight into the observed differences in the SC-FC correlation values.

## 2 Methods

We start by presenting our assumed data acquisition model (Sect. 2.1). Given a set of acquired DW volumes, we form a training set that includes the original volumes and a set of corresponding downsampled volumes at double the voxel size. We then construct two over-complete dictionaries from the original and downsampled set of volumes (Sect. 2.2). For a previously unseen input DW volume, we obtain the super-resolution data in two steps. First, we sparsely code the volume against the dictionary learned from the downsampled volumes in the training set. We finally apply the generated sparse code to the dictionary learned from the original resolution set to obtain the super-resolution data (Sect. 2.3).

## 2.1 Acquisition Model

Let  $v_L$  be an acquired volume and  $v_H$  be the corresponding unobserved higher resolution volume. We assume that the relationship between these two volumes is modeled by [23]:

$$v_L = \mathbf{S}\mathbf{B}v_H + n \quad (1)$$

where  $\mathbf{S}$  is a downsampling operator,  $\mathbf{B}$  is a blurring operator and  $n$  is additive white Gaussian noise. We aim to invert this acquisition model to approximate the unobserved higher resolution volume through super-resolution. The maximum-likelihood solution to this problem involves the minimization of  $\|\mathbf{S}\mathbf{B}\hat{v}_H - v_L\|_2$ , where  $\hat{v}_H$  is the estimated high resolution volume. However, the inversion of  $\mathbf{S}\mathbf{B}$  is ill-posed [23], hence infinitely many maximum-likelihood solutions exist. We thus cast the problem in a dictionary learning framework instead, as explained in the following sections.

## 2.2 Dictionary Construction

We model each 3D patch in dMRI volumes as a sparse linear combination of atoms from a learned dictionary  $\mathbf{D}$ . In the proposed approach, we use two dictionaries to capture the correspondence between low and high resolution dMRI volumes. These two dictionaries are learned from the original training dataset and its downsampled version, respectively.

Let  $\mathbf{v}_O$  be the set of original training volumes concatenated across scans and  $\mathbf{v}_D$  be the corresponding set of downsampled volumes. We extract all overlapping patches in these two sets of volumes, denoted by  $\mathbf{p}_O$  and  $\mathbf{p}_D$ , respectively. Using  $\mathbf{p}_O$  and  $\mathbf{p}_D$ , we construct two over-complete dictionaries as follows:

$$\min_{\mathbf{D}_O, \mathbf{D}_D, \mathbf{y}} \sum \|\mathbf{p}_D - \mathbf{D}_D \mathbf{y}\|_2^2 + \sum \|\mathbf{p}_O - \mathbf{D}_O \mathbf{y}\|_2^2 + \psi(\mathbf{y}) \quad (2)$$

where  $\mathbf{y} = \{\mathbf{y}_{(i,j,k)}\}$  is the set of sparse coding vectors for each image location  $(i, j, k)$ , and  $\mathbf{D}_D$  and  $\mathbf{D}_O$  are the generated over-complete dictionaries of the downsampled and original volumes, respectively [23].  $\psi(\mathbf{y})$  is a regularization term which we set to be  $\psi(\mathbf{y}) = \|\mathbf{y}\|_1$ , inducing sparsity on the generated coding vector [21]. We note that the same set of coding vectors  $\mathbf{y}$  is used for both dictionaries. In other words, the learned atoms of the two dictionaries represent matched pairs. We set the number of atoms in each dictionary to 1,000 and the patch size to  $3 \times 3 \times 3$  voxels, which were empirically chosen to strike a balance between representation accuracy and overfitting.

### 2.3 Super-Resolved Volume Generation

Let  $\mathbf{p}_I$  be the set of low resolution overlapping patches obtained from a previously unseen input volume that we wish to super-resolve. We code  $\mathbf{p}_I$  with respect to  $\mathbf{D}_D$  as:

$$\min_{\mathbf{y}_I} \|\mathbf{p}_I - \mathbf{D}_D \mathbf{y}_I\|_2^2 + \psi(\mathbf{y}_I) \quad (3)$$

where  $\mathbf{y}_I$  is the set of coding vectors for  $\mathbf{p}_I$ , with  $\psi(\mathbf{y}_I)$  again being the  $l_1$  norm of  $\mathbf{y}_I$ , enforcing sparsity on the coefficients. Once the input volume is sparsely coded using  $\mathbf{D}_D$ , we generate a new set of super-resolved patches,  $\mathbf{p}_S$ , by applying the sparse coding vector  $\mathbf{y}_I$  to  $\mathbf{D}_O$  previously constructed from the training data:

$$\mathbf{p}_S = \mathbf{D}_O \mathbf{y}_I. \quad (4)$$

We note that this process results in a patch being generated for each voxel. We then reconstruct the super-resolved volume by averaging neighboring overlapping patches.

We used K-singular value decomposition (K-SVD) [1] to construct the dictionaries and orthogonal matching pursuit [15] to sparsely code the 3D patches. Theoretically,  $\mathbf{p}_O$ ,  $\mathbf{p}_D$  and  $\mathbf{p}_I$  can be extracted at once from the volumes of all gradient directions in the DW images. However, we opt to apply super-resolution for each gradient direction separately. This helps circumvent computational limitations that might arise, especially with the increasingly large number of gradient directions acquired in practice.

## 3 Materials

We validated our method on the publicly available multimodal Kirby 21 dataset.<sup>1</sup> Along with other imaging modalities, this dataset comprises dMRI and RS-fMRI scans of 21 subjects with no history of neurological disease (11 men, 10 women,  $32 \pm 9.4$  years old). We summarize the key acquisition parameters in Sects. 3.1 and 3.2. Further details on data acquisition can be found in [10]. In our experiments, we used 10 subjects for dictionary training, and 10 other subjects for testing.

---

<sup>1</sup>This dataset is available online at: <http://www.nitrc.org/projects/multimodal>.

### 3.1 *RS-fMRI Data*

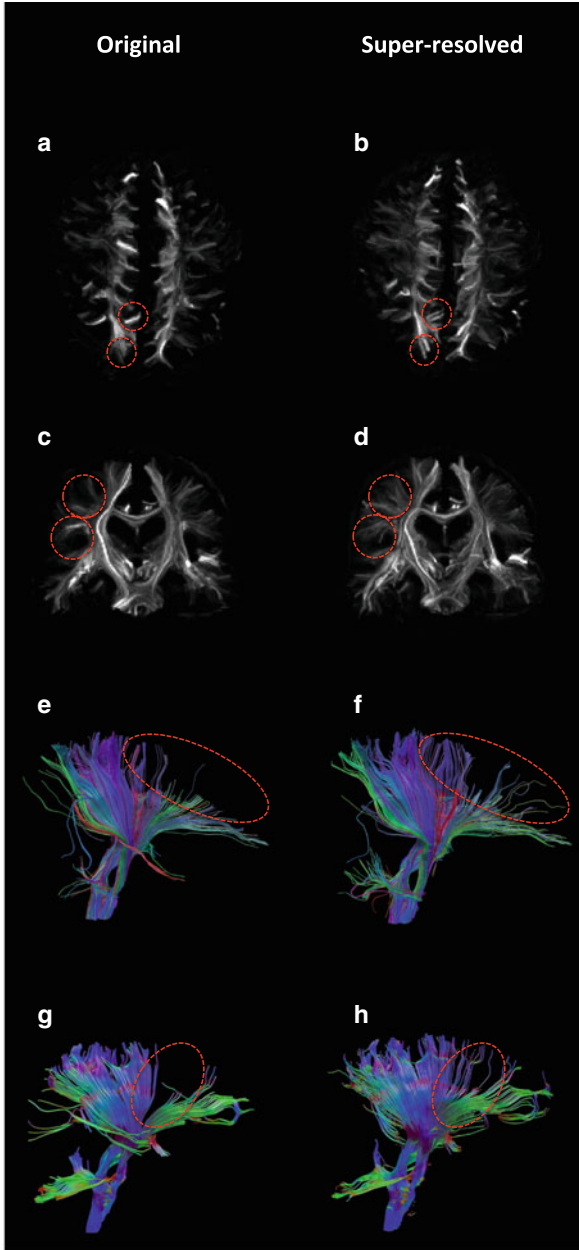
The RS-fMRI data of 7 min duration were collected with a TR of 2 s and a voxel size of 3 mm (isotropic). The data were preprocessed using in-house software written in MATLAB, and the steps followed included motion correction, bandpass filtering at 0.01 and 0.1 Hz, and removal of white matter and cerebrospinal fluid confounds. We divided the brain into 150 parcels by applying Ward clustering [12] on the voxel time courses, which were temporally concatenated across subjects. Parcel time courses were then found by averaging the voxel time courses within each parcel.

### 3.2 *dMRI Data*

The dMRI data had 32 diffusion-weighted images with a b-value of  $700 \text{ s/mm}^2$  in addition to a single  $b = 0$  image, with a voxel size of  $0.83 \times 0.83 \times 2.2 \text{ mm}^3$ . Since anisotropic voxels were previously shown to be suboptimal for fiber tractography [14], we resampled each volume to 2 mm isotropic resolution prior to any analysis. We also applied a Rician-adapted denoising filter [11] to eliminate nonstationary noise commonly observed in DW images, since our acquisition model described in Sect. 2.1 assumes Gaussian noise. We then warped our functionally derived group parcellation map to the  $b = 0$  volume of each subject using FSL [9] to facilitate the computation of fiber count.

## 4 Results and Discussion

We first present a qualitative comparison between the fiber tracts reconstructed from the original (2 mm isotropic) and super-resolved (1 mm isotropic) dMRI data. For ease of interpretation, we chose to employ deterministic streamline tractography with the diffusion tensor model, which is by far the most popular tractography approach to date. However, we highlight that our super-resolution approach can be used with any diffusion model and any tractography method. Tractography was carried out using Dipy [5], with 750,000 seed points for both the original and super-resolution data. We generated the track-density maps by calculating the total number of fiber tracts present in each voxel. Figure 1a,c and b,d show sample track-density maps with the original and super-resolved dMRI data, respectively. As observed from these figures, the track-density maps generated from the super-resolution data clearly convey more spatial information. Figure 1e,f and g,h show the corticospinal tracts extracted using a region of interest (ROI) placed on the brain stem for two representative subjects. It can be observed that fiber tracts reconstructed from the super-resolution data can capture the fan-shape configuration of the corticospinal track more fully.




**Fig. 1** Qualitative comparison between the track-density maps and fiber tracts reconstructed from the original (*left*) and super-resolved (*right*) dMRI data. Original dataset has 2 mm isotropic resolution which is super-resolved to 1 mm isotropic resolution. Each row corresponds to a different test subject. Track-density maps of super-resolved data ((**b**) and (**d**)) show markedly improved spatial detail compared to those of original data ((**a**) and (**c**)). Corticospinal tracts reconstructed from super-resolved data ((**f**) and (**h**)) can capture the fan-shape configuration more accurately than those generated from original data ((**e**) and (**g**))

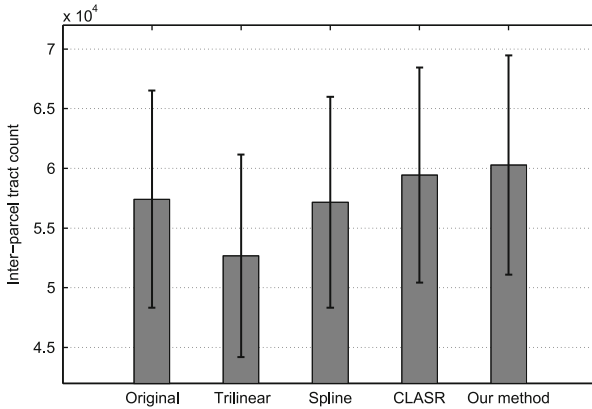
To quantify the improvement in tractography with the suggested approach, we analyzed the consistency between measures of intra-subject SC and FC. We estimated SC using the fiber counts between brain region pairs, and FC using Pearson’s correlation between parcel time courses. For each subject, SC and FC are vectors of size  $d(d - 1)/2 \times 1$  comprising the corresponding connectivity estimates between each region pair, where  $d$  is the number of brain regions. We then calculated Pearson’s correlation between intra-subject SC and FC to quantify the consistency between the two connectivity estimates. Using this correlation measure, we compared the proposed super-resolution approach with trilinear and spline interpolation in addition to an alternative super-resolution method; collaborative and locally adaptive super-resolution (CLASR) [4]. To the best of our knowledge, CLASR is the only existing single image super-resolution method for dMRI which is independent of the diffusion model employed. Figure 2 shows the SC-FC correlation for each subject tested. Taking the average SC-FC correlation across the group when using the original data as a baseline, the improvement was 5.7 % with spline interpolation, 13.6 % with CLASR, and 27.1 % with our proposed method. On the other hand, there was a 6.3 % decrease in the correlation when trilinear interpolation was used. The difference in the performance of our method and every other method tested was found to be statistically significant at  $p < 0.01$  based on the Wilcoxon signed-rank test, showing its potential for enhanced structural connectivity assessment. Our results thus suggest that low spatial resolution of dMRI data can partially account for the low SC-FC correlation, and statistically significant improvements can be achieved using super-resolved dMRI data.

To investigate why trilinear interpolation resulted in a lower SC-FC correlation compared to the original data, we calculated the number of tracts reconstructed with each method. The local intra-parcel connections were excluded since they have no effect on SC-FC correlation. Figure 3 shows the number of inter-parcel tracts averaged across the group along with the corresponding standard deviations. As observed from this figure, performing tractography on volumes upsampled with trilinear interpolation resulted in a lower number of tracts compared to the original volumes, even though the same number of seed points were used to initiate tracking for all of the methods we compared. We speculate that the reason of this phenomenon is the additional partial volume effects introduced by the blurring of the data during trilinear interpolation, which hamper the tractography quality. Spline interpolation, however, is known to cause less blurring compared to trilinear interpolation, and our results suggest that upsampling dMRI data using spline interpolation can be beneficial for tractography. The overall trend of inter-parcel track counts closely resembles to that of the SC-FC correlation, with our proposed method outperforming all other methods tested. This shows that dictionary based super-resolution is a viable post-processing solution for dMRI that can help in mapping the white matter brain architecture more accurately.





**Fig. 2** SC-FC correlation for 10 subjects with SC estimated from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method. Our method outperforms all other methods tested for eight of the subjects, and performs comparable to CLASR for two subjects (subjects 4 and 10)



**Fig. 3** Number of inter-parcel tracts reconstructed from the data at its original resolution (2 mm isotropic), and high-resolution data (1 mm isotropic) obtained using trilinear interpolation, spline interpolation, CLASR and the proposed method. Intra-parcel tracts are not included here since they do not contribute to SC-FC correlation. We emphasize that tractography is initiated with the same number of seeds (750,000) for each method

## 5 Conclusions and Future Work

Low spatial resolution is a known limitation of dMRI, which often hinders the performance of tractography significantly. We proposed the use of a simple yet very effective super-resolution technique in dMRI to capture a more accurate portrayal of the white matter architecture. This approach does not require multiple dMRI acquisitions and is applicable to legacy data. Quantitatively, we demonstrated that SC-FC consistency can be markedly increased with the use of our approach in estimating SC. We also qualitatively illustrated that the gain in spatial resolution remarkably improves the fiber tracts and track-density maps generated. Taken collectively, our results suggest that dictionary based super-resolution holds great promise in enhancing the spatial resolution in dMRI, without requiring additional scans or any modifications of the acquisition protocol.

It is important to acknowledge that the performance of the proposed method inherently depends on the training dataset, as in any machine learning method that involves training or prior information. The age span of the subjects we used in our experiments was 23–61, showing that the method can generalize to a large range of ages. However, how well abnormalities such as tumor and edema can be modeled with dictionary learning is currently unclear and warrants further research.

**Acknowledgements** The authors wish to thank Dr. Pierrick Coupé for assisting in the comparative assessment of our method with CLASR.

## References

1. Aharon, M., Elad, M., Bruckstein, A.: K-SVD: an algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Process.* **54**(11), 4311–4322 (2006)
2. Alexander, A., Hasan, K., Lazar, M., Tsuruda, J., Parker, D.: Analysis of partial volume effects in diffusion-tensor MRI. *Magn. Reson. Med.* **45**(5), 770–780 (2001)
3. Calamante, F., Tournier, J.D., Jackson, G.D., Connelly, A.: Track-density imaging (TDI): super-resolution white matter imaging using whole-brain track-density mapping. *NeuroImage* **53**(4), 1233–1243 (2010)
4. Coupé, P., Manjón, J.V., Chamberland, M., Descoteaux, M., Hiba, B.: Collaborative patch-based super-resolution for diffusion-weighted images. *NeuroImage* **83**, 245–261 (2013)
5. Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., Van Der Walt, S., Descoteaux, M., Nimmo-Smith, I.: Dipy, a library for the analysis of diffusion MRI data. *Front. Neuroinform.* **8**(8) (2014)
6. Glasner, D., Bagon, S., Irani, M.: Super-resolution from a single image. In: *Proc. IEEE 12th International Conference on Computer Vision*, pp. 349–356 (2009)
7. Gupta, V., Ayache, N., Pennec, X.: Improving DTI resolution from a single clinical acquisition: a statistical approach using spatial prior. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI, LNCS*, vol. 8151, pp. 477–484. Springer, Heidelberg (2013)
8. Honey, C., Thivierge, J.P., Sporns, O.: Can structure predict function in the human brain? *NeuroImage* **52**(3), 766–776 (2010)
9. Jenkinson, M., Beckmann, C., Behrens, T., Woolrich, M., Smith, S.: FSL. *NeuroImage* **62**(2), 782–790 (2012)

10. Landman, B., Huang, A., Gifford, A., Vikram, D., Lim, I., Farrell, J., Bogovic, J., Hua, J., Chen, M., Jarso, S., Smith, S., Joel, S., Mori, S., Pekar, J., Barker, P., Prince, J., van Zijl, P.: Multi-parametric neuroimaging reproducibility: a 3-T resource study. *NeuroImage* **54**(4), 2854–2866 (2011)
11. Manjón, J.V., Coupé, P., Buades, A., Collins, D.L., Robles, M.: New methods for MRI denoising based on sparseness and self-similarity. *Med. Image Anal.* **16**(1), 18–27 (2012)
12. Michel, V., Gramfort, A., Varoquaux, G., Eger, E., Keribin, C., Thirion, B.: A supervised clustering approach for fMRI-based inference of brain states. *Pattern Recogn.* **45**(6), 2041–2049 (2012)
13. Mori, S., Zhang, J.: Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron* **51**(5), 527–539 (2006)
14. Neher, P., Stieltjes, B., Wolf, I., Meinzer, H., Maier-Hein, K.: Analysis of tractography biases introduced by anisotropic voxels. In: *Proc. Annual Meeting ISMRM* (2013)
15. Pati, Y., Rezaifar, R., Krishnaprasad, P.S.: Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: *Proc. Asilomar Conference on Signals, Systems and Computers*, pp. 40–44 (1993)
16. Peled, S., Yeshurun, Y.: Super-resolution in MRI: application to human white matter fiber track visualization by diffusion tensor imaging. *Magn. Reson. Med.* **45**(1), 29–35 (2001)
17. Poot, D., Jeurissen, B., Bastiaensen, Y., Veraart, J., Van Hecke, W., Parizel, P., Sijbers, J.: Super-resolution for multislice diffusion tensor imaging. *Magn. Reson. Med.* **69**(1), 103–113 (2013)
18. Scherrer, B., Gholipour, A., Warfield, S.: Super-resolution in diffusion-weighted imaging. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI, LNCS*, vol. 6892, pp. 124–132. Springer, Heidelberg (2011)
19. Skudlarski, P., Jagannathan, K., Calhoun, V., Hampson, M., Skudlarska, B., Pearlson, G.: Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *NeuroImage* **43**(3), 554–561 (2008)
20. Sotiropoulos, S.N., Jbabdi, S., Andersson, J.L., Woolrich, M.W., Uğurbil, K., Behrens, T.E.J.: RubiX: combining spatial resolutions for Bayesian inference of crossing fibers in diffusion MRI. *IEEE Trans. Med. Imaging* **32**(6), 969–982 (2013)
21. Tibshirani, R.: Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B* **58**, 267–288 (1994)
22. Yang, J., Wright, J., Huang, T., Ma, Y.: Image super-resolution as sparse representation of raw image patches. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8 (2008)
23. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse-representations. In: Boissonnat, J.D., Chenin, P., Cohen, A., Gout, C., Lyche, T., Mazure, M.L., Schumaker, L. (eds.) *Curves and Surfaces, LNCS*, vol. 6920, pp. 711–730. Springer, Heidelberg (2012)