Model-Based Super-Resolution of Diffusion MRI

Alexandra Tobisch, Peter F. Neher, Matthew C. Rowe, Klaus H. Maier-Hein, and Hui Zhang

Abstract This work introduces a model-based super-resolution reconstruction (SRR) technique for achieving high-resolution diffusion-weighted MRI. Diffusion-weighted imaging (DWI) is a key technique for investigating white matter non-invasively. However, due to hardware and imaging time constraints, the technique offers limited spatial resolution. A SRR technique was recently proposed to address this limitation. This approach is attractive because it can produce high-resolution DWI data without the need for onerously long scan time. However, the technique treats individual DWI data from different diffusion-sensitizing gradients as independent, which in fact are coupled through the common underlying tissue. The proposed technique addresses this issue by explicitly accounting for this intrinsic coupling between DWI scans from different gradients. The key technical advance is in introducing a forward model that predicts the DWI data from all the diffusion gradients by the underpinning tissue microstructure. As a proof-of-concept, we show that the proposed SRR approach provides more accurate reconstruction results than the current SRR technique with synthetic white matter phantoms.

M.C. Rowe · H. Zhang (✉)
University College London, London, UK
e-mail: matthew.rowe.09@ucl.ac.uk; gary.zhang@ucl.ac.uk

A. Tobisch
German Center for Neurodegenerative Diseases, Bonn, Germany
e-mail: alexandra.tobisch.11@alumni.ucl.ac.uk

P.F. Neher · K.H. Maier-Hein
German Cancer Research Center, Heidelberg, Germany
e-mail: p.neher@dkfz-heidelberg.de; K.Maier-Hein@dkfz-heidelberg.de

T. Schultz et al. (eds.), *Computational Diffusion MRI and Brain Connectivity*, Mathematics and Visualization, DOI 10.1007/978-3-319-02475-2_3,

1 Introduction

Diffusion-weighted magnetic resonance imaging (DWI) is a key non-invasive technology for investigating brain white matter in vivo [8]. The technique enables the mapping of white matter microstructure over the whole brain [1, 11, 14, 16] and the inference of the complex structural connectivity that white matter supports [4,5,7]. However, due to its intrinsic signal-to-noise (SNR) characteristics, DWI is currently limited to relatively low spatial resolution when compared to standard anatomical scans. A voxel in typical DWI data is at least 8 times the size of a 1 mm^3 voxel in standard T1-weighted anatomical data. The latter's millimeter and sub-millimeter resolution is essential for resolving fine details in anatomical structures, such as the subfields of the hippocampus [15]. Hence, there is increasing interest in new acquisition and computational techniques that can achieve similarly high resolution in DWI so that tissue microstructure and connectivity can be examined even for very small structures.

There are a number of possible approaches to increasing the spatial resolution of DWI data. The most straightforward one is to reduce the slice thickness while simultaneously increasing the in-plane resolution to maintain isotropic voxels. However, this approach leads to substantial reduction in SNR: Reducing each voxel dimension by half decreases the SNR by a factor of 8. The loss in SNR can be compensated by repeated acquisitions but would require a 64-fold increase in imaging time, which is feasible only for postmortem studies [6, 12]. Another approach is to use reduced field-of-view (FOV) imaging. This allows isotropic sub-millimeter resolution in vivo [15] but at a price of restricted anatomical coverage. Specialist hardwares, such as the human connectome project scanner, are now able to achieve better spatial resolution with more advanced imaging sequences and hardware specifications; but such advances are not generally accessible to the standard clinical scanners.

The present paper explores a different approach known as the super-resolution reconstruction (SRR) [3]. The technique constructs a set of isotropic high-resolution (HR) DWIs from multiple sets of anisotropic low-resolution (LR) acquisitions (Fig. 1 provides an illustration). By leveraging intelligent image reconstruction, the SRR approach requires only a modest increase in imaging time, thus has the potential to become a practical solution to realizing higher resolution in DWI. Such an algorithm was first proposed by Scherrer et al. [13]. Their technique, henceforth referred to as the current SRR, carefully addresses the mismatch in diffusion-sensitizing gradients between different DWI scans due to subject motion, and is demonstrated with in vivo human brain data. However, one limitation of this technique is that it treats the DWI data for each diffusion gradient independently to the ones from the other gradients and solves the SRR problem for each in isolation. Such treatment is problematic because the DWIs from different gradients are not independent but rather dictated, and hence strongly coupled, by the same underlying tissue microstructure. Ignoring this strong dependency may lead to overfitting and thus potentially biased reconstruction of HR DWIs.

To address this limitation of the current SRR approach, we propose a technique that adopts a model-based strategy to explicitly account for the intrinsic coupling

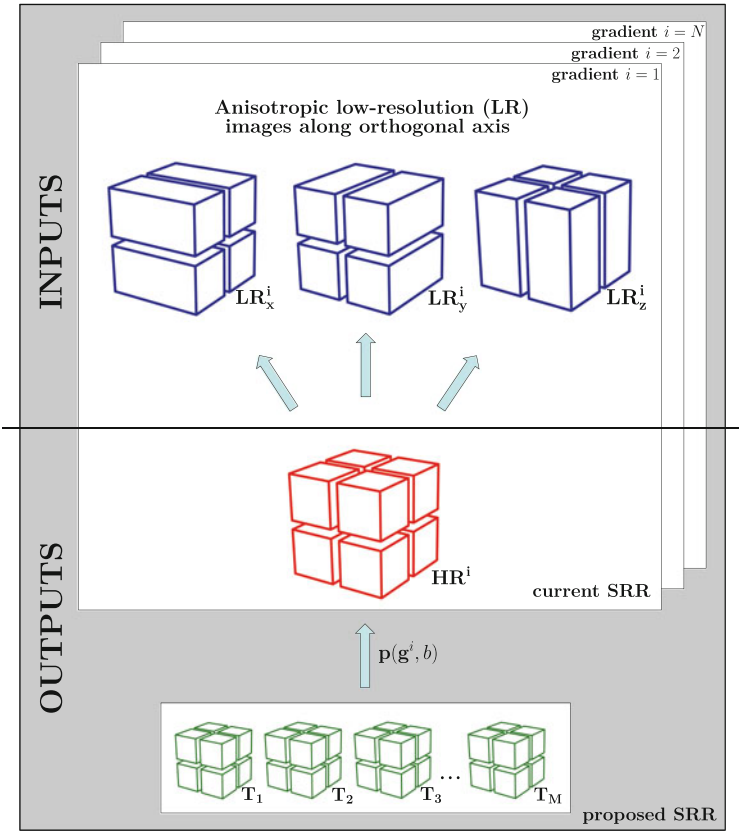

Fig. 1 Forward models of the current and the proposed SRR techniques: The imaging protocol **p** acquires $K = 3$ low-resolution diffusion-weighted images for each of the N diffusion-sensitizing gradients. For the ith gradient, the three anisotropic low-resolution diffusion-weighted images are denoted by LR^i_x, LR^i_y, and LR^i_z. The current SRR technique outputs one high-resolution diffusion-weighted image for each of the N gradients, with HR^i denoting the output for the ith gradient. In contrast, the proposed SRR technique outputs one high-resolution image for each of the M parameters of the underlying tissue model T, with T_i denoting the output for the ith tissue parameter

between the DWIs from different diffusion gradients. This leads to a new SRR formulation for DWI that treats the anisotropic LR data from different gradients as a single input and solves the SRR problem for all the gradients in one setting. As a proof of concept, the proposed technique is demonstrated with synthetic phantom data and its performance evaluated quantitatively against the current method [13]. The rest of the paper is organized as follows: Sect. 2 describes the proposed SRR framework; Sect. 3 details the design of the simulation data experiment for validating the proposed technique and reports the findings; Sect. 4 summarizes the contribution and discusses future work.

2 Model-Based Super-Resolution Reconstruction

The classic SRR problem aims to produce an isotropic HR image from a set of anisotropic LR images along each of the $K = 3$ orthogonal spatial axes [3]. Each LR image covers exactly the same FOV but is acquired in different (orthogonal) planes with high resolution in-plane but low resolution through-plane. The problem is typically solved with a maximum a posteriori (MAP) formulation which involves an image acquisition model that relates the LR images to the HR images and a spatial regularization prior to handle the ill-posedness of the problem.

In the case of DWI, the input to the SRR problem becomes a set of K anisotropic LR images for each of the N diffusion gradients. The number of diffusion gradients in a typical DWI acquisition is 30 or more. The objective is to reconstruct from these NK anisotropic LR images a set of N isotropic HR images, one for each diffusion gradient. The current SRR technique [13] solves this problem by treating the set of K LR images from each diffusion gradient as independent and determine the corresponding HR image individually with the standard algorithm. As a result, the resulting HR images for different gradients may not retain the necessary coherence required for them to faithfully reflect the underlying tissue microstructure.

To address this problem, the proposed SRR algorithm is performed simultaneously on all the NK LR images from all the diffusion gradients and utilizes an image acquisition model that relates the LR images to the HR images of the underlying tissue microstructure. The remainder of this section presents the proposed forward model, the new SRR algorithm, and the solution using the steepest gradient descent optimization.

2.1 Forward Model

The forward model in the current SRR technique [13] predicts the anisotropic LR images for each diffusion gradient from the isotropic HR images of the same gradient. The proposed forward model differs from [13] in that it aims to explicitly take into account the coupling between the images for different gradients governed by the underlying tissue. This is accomplished by predicting the HR images from all the diffusion gradients themselves from a set of HR tissue parameter maps $\mathbf{T}$, given an appropriate tissue model, and an imaging protocol $\mathbf{p}$. (See Fig. 1 for an illustration.) The tissue model prescribes the relationship between the voxel-wise DWI measurement for each gradient and the tissue microstructure at the voxel. The imaging protocol $\mathbf{p}$ specifies the diffusion weighting (the b-value) and the set of diffusion gradients $\mathbf{g}$.

We assume for simplicity that the input data has undergone the appropriate preprocessing to correct for motion and eddy-current distortion, for example

following the methods given in [13]. Then the proposed forward model takes the following simple form:

$$\tilde{\mathbf{y}}_k(\mathbf{T}, \mathbf{p}) = \mathbf{D}_k \mathbf{x}(\mathbf{T}, \mathbf{p}) + \epsilon_k \tag{1}$$

where k is the index over the K orthogonal planes, $\tilde{\mathbf{y}}_k$ is the predicted LR images for all the gradients, $\mathbf{x}(\mathbf{T}, \mathbf{p})$ is the HR images predicted by the chosen tissue model, $\mathbf{D}_k$ is the downsampling operator, and ϵ_k is the residual noise.

The proposed formulation is not limited to any particular tissue model; the choice of the tissue model depends on the application. For the present demonstration, we adopt the simple ball-and-stick model [2], which has been previously used in a different super-resolution setting for resolving partial volume effect in DWI [9]. At each voxel, the signal along a given diffusion gradient is determined by a small set of parameters: the relative fraction of the ball and stick compartments, the diffusivity, which is the same for both compartments, and the orientation of the stick.

2.2 Super-Resolution Reconstruction

The proposed algorithm determines $\hat{\mathbf{x}}$, an estimate to the unknown HR images $\mathbf{x}$, by first solving for the tissue parameter maps $\mathbf{T}$ with the MAP formulation:

$$\hat{\mathbf{T}} = \arg\max_{\mathbf{T}} p(\mathbf{T} \mid \mathbf{y}) = \arg\max_{\mathbf{T}} p(\mathbf{y} \mid \mathbf{T})\, p(\mathbf{T}) = \arg\max_{\mathbf{T}} [\ln p(\mathbf{y} \mid \mathbf{T}) + \ln p(\mathbf{T})] \tag{2}$$

The HR estimates $\hat{\mathbf{x}}$ is subsequently determined directly from $\hat{\mathbf{T}}$ using the chosen tissue model.

Assuming the Gaussian noise model with zero-mean and variance σ and statistical independence between the noise from different acquisitions, the likelihood term $p(\mathbf{y} \mid \mathbf{T})$ can be written as:

$$p(\mathbf{y} \mid \mathbf{T}, \sigma) = \prod_{k=1}^{K} p(\mathbf{y}_k \mid \mathbf{T}, \sigma) = \prod_{k=1}^{K} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\|\mathbf{y}_k - \tilde{\mathbf{y}}_k(\mathbf{T}, \mathbf{p})\|^2}{2\sigma^2}\right) \tag{3}$$

We set the same spatial regularization prior as in [13], such that $p(\mathbf{T} \mid \lambda) = \exp(-\lambda \|\mathbf{Q}\mathbf{x}(\mathbf{T}, \mathbf{p})\|^2)$ where $\mathbf{Q}$ is the 3-D discrete Laplacian. Substituting the forward model, the tissue parameter maps $\mathbf{T}$ can be determined by the following minimization:

$$\hat{\mathbf{T}} = \arg\min_{\mathbf{T}} \sum_{i=1}^{N} \sum_{k=1}^{K} \left\|\mathbf{y}_k^i - \mathbf{D}_k \mathbf{x}^i(\mathbf{T}, \mathbf{p}^i)\right\|^2 + \lambda \left\|\mathbf{Q}\mathbf{x}^i(\mathbf{T}, \mathbf{p}^i)\right\|^2 \tag{4}$$

where i is the index over the set of diffusion gradients specified by $\mathbf{p}$. The objective function differs from the one for the current SRR technique in two ways: first, the

summation over all the N diffusion gradients; second, the underlying variables are the HR tissue parameter maps $\mathbf{T}$ rather than the HR DWIs $\mathbf{x}$. Here $\mathbf{x}$ is determined from $\mathbf{T}$ by the chosen tissue model.

2.3 SRR Optimization Procedure

As in [13], we use the steepest gradient descent approach to optimize for the parameters of the tissue microstructure. The best estimate to the tissue parameter maps $\mathbf{T}$ is iteratively determined by the following update rule from differentiating Eq. (4):

$$\hat{\mathbf{T}}_t^{n+1} = \hat{\mathbf{T}}_t^n - \alpha \left[\sum_{i=1}^{N} \frac{\partial \hat{\mathbf{x}}^{i,n}(\hat{\mathbf{T}}^n, \mathbf{p}^i)}{\partial \hat{\mathbf{T}}_t^n} \left(\sum_{k=1}^{K} \mathbf{D}_k^T (\mathbf{D}_k \hat{\mathbf{x}}^{i,n}(\hat{\mathbf{T}}^n, \mathbf{p}^i) - \mathbf{y}_k^i) - \lambda \mathbf{Q}^T \mathbf{Q} \hat{\mathbf{x}}^{i,n}(\hat{\mathbf{T}}^n, \mathbf{p}^i) \right) \right] \tag{5}$$

where $\mathbf{D}_k^T$ is the transpose of $\mathbf{D}_k$, α is the step size, and $\frac{\partial \mathbf{x}^i(\mathbf{T}, \mathbf{p}^i)}{\partial \mathbf{T}_t}$ is the partial derivative of the diffusion signal predicted by the tissue model with respect to the tissue parameter $\mathbf{T}_t$. The initial estimate of the tissue parameters is obtained by fitting the tissue model to the mean of the upsampled LR images. The minimization is stopped either when the objective function, denoted as v_{obj}, does not decrease or when the difference between consecutive values of the objective function is below a threshold τ_{stop}. The following pseudo-code summarizes the proposed SRR optimization procedure:

$\widehat{\mathbf{x}}^0$ ← Compute the average of the upsampled $\mathbf{y}_k$
$\widehat{\mathbf{T}}^0$ ← Estimate tissue parameters from $\widehat{\mathbf{x}}^0$
$\widehat{\mathbf{x}}^0(\widehat{\mathbf{T}}^0, \mathbf{p})$ ← Synthesize with the tissue model

WHILE $\|v_{obj}^n - v_{obj}^{n-1}\| \geq \tau_{stop}$ || $v_{obj}^n \leq v_{obj}^{n-1}$
    FOR each tissue parameter t
        $\widehat{\mathbf{T}}_t^{n+1}$ ← Update with Eqn (5)
    ENDFOR
    $\widehat{\mathbf{x}}^{n+1}(\widehat{\mathbf{T}}^{n+1}, \mathbf{p})$ ← Synthesize with the tissue model
ENDWHILE

The parameters λ, τ_{stop} and α are set according to [13]: $\lambda = 0.001$ and $\tau_{stop} = 5 \times 10^{-5}$. The step-size α is initialized to 0.01 but automatically tuned during optimization to accelerate the convergence. If the signs of two consecutive gradients differ, α is divided by two; if the signs are the same, α is multiplied by 1.1. In addition, α is restricted to $[0.1, 10^{-6}]$.

3 Evaluation and Results

This section outlines the evaluation strategy, gives the detail of the synthetic data experiment, and reports the quantitative comparison between the proposed SRR technique to the current SRR approach.

3.1 Evaluation

We assess the proposed SRR technique with synthetic phantoms which provide ground-truth data for establishing quantitative measures of performance. In particular, we generate ground-truth DWIs from synthetic phantoms constructed with the software tool Fiberfox, part of the MITK Diffusion Imaging.[1] Fiberfox combines the advantage of synthetic DWIs with a level of realism that, so far, has only been seen in real MR acquisitions. It combines complex voxel-wise diffusion modeling, using the rich hierarchy of tissue models in [10], with the synthesis of large-scale fiber configurations.

Specifically, two fiber configurations, fanning and helix, are chosen for the present illustration (Fig. 2). Fiberfox is used to synthesize both the ground-truth isotropic HR DWIs of each phantom and the corresponding anisotropic LR DWIs. The LR acquisitions share the same in-plane voxel spacing as the HR data but double the through-plane voxel spacing. The DWI data are simulated with a standard clinical protocol: 30 isotropically-distributed diffusion gradients with $b = 1{,}000\,\mathrm{s/mm}^2$. To mimic real-world acquisitions, we perturb the LR acquisitions with Rician noise at two levels corresponding to SNR of 50 and 20.

3.2 Results

The performance of the proposed and current SRR techniques is assessed by the respective reconstruction error, which is measured as the departure from the ground-truth HR DWI data. Specifically, we compute the mean-squared error (MSE) between the ground-truth and predicted HR DWI data for each noise trial and report the average over 15 different trials. Table 1 summarizes these averaged mean-squared errors (MSE), showing that the proposed SRR consistently produces a smaller reconstruction error than the current SRR.

In particular, the relative improvement in performance of the proposed SRR over the current technique increases with decreasing SNR. For the SNR of 20, typical of in vivo data, the error of the proposed SRR is significantly smaller than that of

[1] http://www.mitk.org/DiffusionImaging

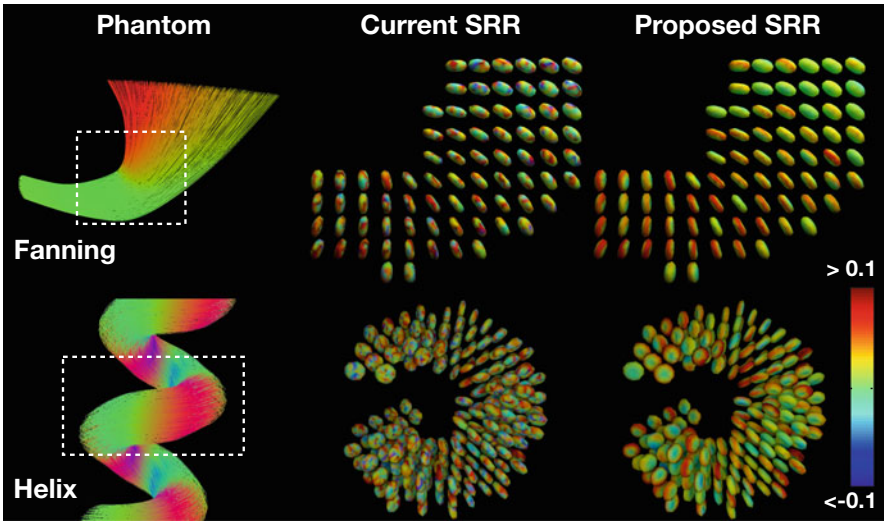

Fig. 2 Phantom data experiment: the fiber configurations of the phantoms (*left*), the reconstruction errors for the current (*middle*) and proposed (*right*) SRR techniques. The *dashed squares* mark the parts of the phantoms rendered in the middle and right columns. Each voxel-wise DWI data is rendered as a 3-D glyph colored according to the reconstruction errors of the corresponding SRR technique. For the helical phantom, the DWIs are viewed along the helical axis

Table 1 Mean-squared errors (MSE) of the reconstructed HR data with respect to the ground truth and objective function values at the convergence

		Fanning		Helix	
		SRR_{cur}	SRR_{pro}	SRR_{cur}	SRR_{pro}
MSE	SNR 50	8.65 ± 0.057	4.94 ± 0.038	12.35 ± 0.079	8.57 ± 0.072
	SNR 20	32.82 ± 0.188	8.20 ± 0.135	36.81 ± 0.326	11.88 ± 0.249
Objectives	SNR 50	4.38 ± 0.044	10.67 ± 0.065	8.02 ± 0.054	14.39 ± 0.069
	SNR 20	21.31 ± 0.254	56.80 ± 0.375	24.89 ± 0.250	60.68 ± 0.491

the current technique. This difference in performance between the two techniques can be better visualized in Fig. 2, which illustrates the voxel-wise reconstruction errors for one particular noise trial. The figure highlights additionally that the reconstruction error from the current SRR is highly irregular.

To understand the improved performance with the proposed SRR, we compare the objective function values from both techniques at the convergence in Table 1. For the proposed SRR, the objective function is given in the right hand side of Eq. (4) as the argument to arg min. For the current SRR, the objective function values reported are summed over all the diffusion gradients to be exactly comparable to the proposed SRR. The objective function in essence evaluates the squared difference between the measured (noisy) LR DWI data and the predicted LR DWI data by each

technique together with the penalty for spatial irregularity. It is interesting that the values for the current SRR are consistently lower than the proposed SRR, suggesting the former is overfitting the (noisy) data. This is consistent with the irregularity in the reconstruction error illustrated in Fig. 2.

4 Discussion

In summary, this paper proposes a new SRR formulation for DWI data that explicitly accounts for the dependency between the measurements from different diffusion gradients. This is achieved with a model-based approach that links the diffusion measurements directly to the underlying tissue model. The proof-of-concept illustration with synthetic phantom data demonstrates the importance of the proposed approach: ignoring the strong coupling between DWI data as in [13] can lead to overfitting. This is unsurprising, given the number of parameters in typical tissue models [10] are smaller than the number of diffusion gradients. The overfitting problem should be even more significant for increasingly common high-angular resolution diffusion imaging (HARDI) or multi-shell HARDI acquisitions [1, 14, 16]. Future work will consider more complex tissue models that can accommodate crossing fiber configurations and more sophisticated multi-shell HARDI acquisitions using in vivo imaging data.

References

1. Assaf, Y., Basser, P.J.: Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. NeuroImage **27**, 48–58 (2005)
2. Behrens, T.E.J., Johansen-Berg, H., Woolrich, M.W., Smith, S.M., Wheeler-Kingshott, C.A.M., Boulby, P.A., Barker, G.J., Sillery, E.L., Sheehan, K., Cicarelli, O., Thompson, A.J., Brady, J.M., Matthews, P.M.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magn. Reson. Med. **50**, 1077–1088 (2003)
3. Greenspan, H.: Super-resolution in medical imaging. Comput. J. **52**(1), 43–63 (2009)
4. Hagmann, P., Thiran, J.P., Jonasson, L., Vandergheynst, P., Clarke, S., Maeder, P., Meuli, R.: DTI mapping of human connectivity: statistical fibre tracking and virtual dissection. NeuroImage **19**, 545–554 (2003)
5. Hagmann, P., Kurant, M., Gigandet, X., Thiran, P., Wedeen, V.J., Meuli, R., Thiran, J.P.: Mapping human whole-brain structural networks with diffusion MRI. PLoS ONE **2**(7), e597 (2007)
6. McNab, J., Jbabdi, S., Deoni, S.C.L., Douaud, G., Behrens, T.E.J., Millder, K.L.: High resolution diffusion-weighted imaging in fixed human brain using diffusion-weighted steady state free precession. NeuroImage **46**, 775–785 (2009)
7. Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.M.: Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann. Neurol. **45**, 265–269 (1999)
8. Moseley, M., Cohen, Y., Kucharaczyj, M., Mintorovitch, J., Asgari, H., Wendland, M., Tsuruda, J., Norman, D.: Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system. Radiology **176**, 439–445 (1990)

9. Nedjati-Gilani, S., Parker, G.J.M., Alexander, D.C.: Regularized super-resolution for diffusion MRI. In: IEEE International Symposium on Biomedical Imaging: Macro to Nano, Paris, pp. 875–878 (2008)
10. Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C.: Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison. NeuroImage **59**(3), 2241–2254 (2012)
11. Pierpaoli, C., Jezzard, P., Basser, P.J., Barnett, A., Chiro, G.D.: Diffusion tensor MR imaging of the human brain. Radiology **201**, 637–648 (1996)
12. Roebroeck, A., Galuske, R., Formisano, E., Chiry, O., Bratzke, H., Ronen, I., Kim, D.S., Goebel, R.: High-resolution diffusion tensor imaging and tractography of the human optic chiasm at 9.4T. NeuroImage **39**(1), 157–168 (2008)
13. Scherrer, B., Gholipour, A., Warfield, S.K.: Super-resolution reconstruction to increase the spatial resolution of diffusion weighted images from orthogonal anisotropic acquisitions. Med. Image Anal. **16**(7), 1465–1476 (2012)
14. Wu, Y.C., Alexander, A.L.: Hybrid diffusion imaging. NeuroImage **36**(3), 617–629 (2007)
15. Zeineh, M.M., Holdsworth, S., Skare, S., Atlas, S.W., Bammer, R.: Ultra-high resolution diffusion tensor imaging of microscopic pathways of the medial temporal lobe. NeuroImage **62**, 2065–2082 (2012)
16. Zhang, H., Schneider, T., Wheeler-Kingshott, C., Alexander, D.C.: NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage **61**, 1000–1016 (2012)