HARDI Denoising: Variational Regularization of the Spherical Apparent Diffusion Coefficient *sADC*⋆

Yunho Kim[1], Paul M. Thompson[2], Arthur W. Toga[2], Luminita Vese[1], and Liang Zhan[2]

[1] Mathematics Department, UCLA, Los Angeles, CA, USA
[2] LONI & Neurology Department, UCLA School of Medicine, Los Angeles, CA, USA

Abstract. We denoise HARDI (High Angular Resolution Diffusion Imaging) data arising in medical imaging. Diffusion imaging is a relatively new and powerful method to measure the 3D profile of water diffusion at each point. This can be used to reconstruct fiber directions and pathways in the living brain, providing detailed maps of fiber integrity and connectivity. HARDI is a powerful new extension of diffusion imaging, which goes beyond the diffusion tensor imaging (DTI) model: mathematically, intensity data is given at every voxel and at any direction on the sphere. However, HARDI data is usually highly contaminated with noise, depending on the b-value which is a tuning parameter preselected to collect the data. Larger b-values help to collect more accurate information in terms of measuring diffusivity, but more noise is generated by many factors as well. So large b-values are preferred, if we can satisfactorily reduce the noise without losing the data structure. We propose a variational method to denoise HARDI data by denoising the spherical Apparent Diffusion Coefficient (sADC), a field of radial functions derived from the data. We use vectorial total variation regularization, an L^1 data fidelity term and the logarithmic barrier function in the minimization. We present experiments of denoising synthetic and real HARDI data.

1 Introduction to the HARDI Data

Currently, HARDI data is used to map cerebral connectivity through fiber tractography in the brain. HARDI is a type of diffusion MRI, which was introduced in the mid-1980s by Le Bihan et al. [20,21,22] and Merboldt et al. [28]. It is based on the idea that the MR signal, which forms the basis of MRI, is attenuated when water diffuses out of a voxel, and the degree of attenuation can be used to measure the rate of water diffusion in any arbitrary 3D direction via the Stejskal-Tanner equation [33]. Water diffusion occurs preferentially in directions that are aligned with axonal fiber pathways, and is hindered in orthogonal directions by the myelin sheaths that coat the axons. Because of this diffusion anisotropy, initial approaches to assess fiber directions modeled the 3D diffusion

⋆ Funded by the NIH Roadmap for Medical Research Grant U54 RR021813 (CCB).

J.L. Prince, D.L. Pham, and K.J. Myers (Eds.): IPMI 2009, LNCS 5636, pp. 515–527, 2009.

profile at each point as a single tensor (Beaulieu et al. [4]), in which the principal eigenvector of the diffusion tensor can be used to recover the dominant fiber pathway at that voxel. The diffusion tensor model (Basser et al. [3]) describes the anisotropic nature of water diffusion in tissues (inside a typical 1-3mm sized voxel) by estimating, from a set of K diffusion-sensitized images, the 3x3 covariance matrix of a Gaussian distribution (Beaulieu et al. [4]). Each voxel's signal intensity in the k-th image is decreased, by water diffusion, according to the Stejskal-Tanner equation [33]: $S_k = S_0 \exp[-bg_k^T D g_k]$, where S_0 is the non-diffusion weighted signal intensity, D is the 3x3 diffusion tensor, g_k is the direction of the diffusion gradient and b is Le Bihan's factor with information on the pulse sequence, gradient strength, and physical constants.

Unfortunately, although it is widely used, the diffusion tensor model breaks down for voxels in which fiber pathways cross or mix together, and these are ubiquitous in the brain which is highly interconnected. More advanced image acquisition techniques, such as HARDI (Tuch et al. [39,40]), diffusion spectrum imaging [43], and q-ball imaging (Tuch et al. [41]), have been introduced in the past 5 years - these types of data recover the local microstructure of water diffusion more accurately than standard DTI data. HARDI, DTI and other similar modalities permit non-invasive quantification of the water diffusion in living tissues. The tissue structure will affect the Brownian motion of the water molecules which will lead to an anisotropic diffusion. By imaging diffusion in an arbitrary number of directions (often 100 or more), HARDI overcomes the limited accuracy of the tensor model in resolving the highly complex fiber structure of the brain, particularly in regions with fiber crossings.

HARDI data makes it possible to compute the orientation diffusion function over a sphere of possible directions. Tuch [38,40] developed the first HARDI acquisition and processing methods, and later Frank [15] used spherical harmonic expansions for processing HARDI data sets. A very active area of research has grown up in processing the HARDI signals, leading to methods for HARDI denoising, segmentation, and registration using metrics on spherical functions (Lenglet et al. [23]). Most of these signal processing methods still model the diffusion signal as a tensor, rather than exploiting the full information in the spherical harmonic expansion. For example, Khurd et al. [42] used isometric mapping and manifold learning (eigendecomposition of the distance matrix) to directly fit a manifold to the tensors, compute its dimensionality, and distinguish groups using Hotelling's T^2 statistics. Initial image processing on the full HARDI signal has focused on fitting a discrete mixture of k distinct tensors to the signal, and later on fitting a continuous mixture model for modeling the MR signal decay and multi-fiber reconstruction (Jian et al. [17], [18]), or fitting a continuous mixture of tensors using a unit-mass distribution on the symmetric positive definite tensor manifold (Leow et al. [24]).

Initial work on the nonlinear (fluid) matching of HARDI images has taken a more non-parametric approach, and has used the Kullback-Leibler divergence to measure the discrepancy between ODF fields (Chiang et al. [8,9]), using a 3D fluid transform to minimize the discrepancy between two fields of ODFs. As

information theory can be used to measure the overlap between diffusion probability density functions, there is much promising work using metrics derived from information theory (e.g., the Fisher-Rao metric, von Mises-Fisher distribution, etc.; McGraw et al. [27]; Srivastava et al. [32]; Chiang et al. [9]). Other work has modeled the HARDI signal as high-order tensors (Barmpoutis et al. [2]) or as a stratification (mixture of manifolds with different dimensions; Haro et al. [16]).

The HARDI data is the MRI signal attenuation information after time $t > 0$ modeled by $S_t(x, \theta, \phi) = S_0(x) \exp(-b \cdot d_t(x, \theta, \phi))$, where Ω is a bounded open subset of $\mathbb{R}^3$, $x \in \Omega$, $\theta \in [0, 2\pi)$, $\phi \in [0, \pi)$. S_0 is the MRI signal that is obtained when no diffusion gradient vector is applied and this is considered to be a reference image, relative to which the diffusion-attenuated signal is measured. The function $d_t(x, \theta, \phi)$ is called the spherical Apparent Diffusion Coefficient (sADC), which measures how much the water molecules diffuse in the given direction $(\cos(\theta)\sin(\phi), \sin(\theta)\sin(\phi), \cos(\phi))$, and b is a parameter pre-selected to collect the data.

In reality, in experimental data, a higher b-value (e.g., 3000 s/mm^2) tends to lead to more noise in the obtained images [10]. Hence, we are led to consider the following simplified degradation model

$$S_t(x, \theta, \phi) = S_0(x) \exp(-b \cdot d_t(x, \theta, \phi)) + noise(x, \theta, \phi). \tag{1}$$

We may say that the baseline signal collected without any diffusion gradient applied, which is S_0, may also be contaminated by noise, but here we assume that this can be neglected, or we just consider the last noise term in (1) to encompass all types of noise. This is a reasonable approximation, because in practice, it is common to collect several non-diffusion weighted images $S_{0,i}$ whose average may be used as a reference signal S_0 (e.g. Zhan et al. [44]). If we let $\tilde{S}_t(x, \theta, \phi)$ be a denoised dataset, then we expect that for all x, ϕ, θ,

$$0 \leq \tilde{S}_t(x, \theta, \phi) \leq S_0(x). \tag{2}$$

As already mentioned, the data has to be first denoised before extracting the fibers, or before registration. Although HARDI is a relatively recent type of data acquisition, several HARDI processing methods have already been proposed: we mention a few more. In [26] and [7], curve evolution techniques are applied for the segmentation of HARDI data. Descoteaux, Deriche and collaborators, among others, have also proposed a segmentation of HARDI data [12], a regularized, fast and robust analytical solution for the Q-ball imaging reconstruction of the ODF [13], and for mapping of neuronal fiber crossings [14]. [11] deals with denoising and regularization of fields of ODFs (orientation distribution functions).

The prior work most relevant to ours is by Mc Graw et al. [25]: the noisy data $S_t(x, \theta, \phi)$ is regularized to remove noise in a functional minimization approach; a standard L^2 data fidelity term is used, combined with a weighted version of vectorial total variation regularization in space, and H^1 regularization of data at every voxel with respect to direction. The data is mapped into 2-dimensional space plane using spherical coordinates and discretized using finite elements. Denoising results for synthetic and real HARDI data are presented in [25]. In

our proposed work, we also use vectorial total variation for the regularization. However, our proposed model differs from the one in [25], since we faithfully follow the signal degradation model (1) and we denoise $d_t(x, \theta, \phi)$ instead of $S_t(x, \theta, \phi)$. Results on synthetic data used in [25] will be shown for comparison.

2 Proposed Variational Denoising Model

We propose a variational denoising method that recovers a clean $d = d_t$. The HARDI data is a collection of intensity values at uniformly pre-selected directions on the sphere, to which the electromagnetic field is applied: at each position $x \in \Omega \subset \mathbb{R}^3$, we measure values at different directions. We note briefly that the actual set of directions is typically computed using an electrostatic repulsion PDE, to optimize the sampling of a spherical signal using a finite set of observations (see Tuch et al. [41] for a discussion of spherical sampling schemes).

In the continuous setting, we obtain a function defined on a manifold $\Omega \times S^2$; Ω is the spatial domain and the sphere S^2 is the space of gradient directions. It is not easy to work with the entire domain $\Omega \times S^2$ for computational purposes. Instead, we will use a discretized version of the sphere, given by n directions uniformly chosen. We drop the subscript t from S_t and thus the function that is given has the form $S = (S_1, \ldots, S_n)$ where each $S_i : \Omega \to \mathbb{R}$ corresponds to a given direction. The data S_0 has only spatial information and it is also given.

To impose the right amount of smoothness and discontinuity on the denoised data, we will use the vectorial total variation regularization, given by $|\nabla d|(\Omega) = \int_\Omega \sqrt{\sum_{i=1}^n |\nabla d_i(x)|^2} dx$ if $d \in W^{1,1}(\Omega; \mathbb{R}^n)$ (∇d is a $n \times 3$ matrix). For $d \in W^{1,1}(\Omega; \mathbb{R}^n)$, $\nabla d = (\frac{\partial d_i}{\partial x_k})_{i=1,\ldots,n, k=1,2,3}$ in the distributional sense and $|\nabla d|$ denotes the Frobenius norm of ∇d. The total variation has been successfully introduced and used in image denoising for gray-scale images by Rudin, Osher, Fatemi [31], being a convex edge-preserving regularization. The vectorial total variation for color images has been analyzed in Blomgren-Chan [6] and PhD manuscripts of Blomgren [5], Tschumperlé [34] (see also [35], [36]).

As we have mentioned, we wish to denoise the sADC (spherical Apparent Diffusion Coefficient) d_t in (1). We drop the subscript t and denote the sADC by d. We note that we can substitute $b \cdot d$ by d in (1), since b is a constant. From the HARDI data model (1), we see that knowing the true S is equivalent to knowing the true d. We directly impose the image formation model in our data-fidelity term and the constraint $d_i \geq 0$. Also, we use an L^1 noise term [1], instead of the more standard L^2 noise term, to penalize less the unknown. We minimize the energy $G(d) = G(d_1, ..., d_n)$.

We use the logarithmic barrier method [29] to realize the constraint: the energy should contain $-\mu \sum_{i=1}^n \log(d_i(x))$ (with a sequence of parameters $\mu > 0$ decreasing to zero), which realizes the constraint $d_i > 0$ for all i. Instead, we will use $-\mu \sum_{i=1}^n \left[H(S_i(x), S_0(x)) \cdot \frac{\log(d_i(x))}{d_i(x)} \right]$, with the function H depending only on the data S and S_0. One choice of the function H is: $H(a, b) = 0$ if $a \leq b$, $H(a, b) = 1$ if $a > b$. Using this weight H, we penalize the unknown only at those points $x \in \Omega$ with $S_i(x) - S_0(x) > 0$ which violate the second constraint

in (2). Also, note that, if we would have used $-\mu\log(z)$ instead of $-\mu\log(z)/z$, with $z = d_i(x)$, then the energy would have no global minimizer since for any d such that $G(d) < \infty$, violating the constraint on a set of positive measure, then $\lim_{k\to\infty} G(k+d) = -\infty$. Another advantage of the function $-\mu\log(z)/z$ is that on one hand, unlike the logarithmic function $-\mu\log(z)$, it rather spreads uniform weights on $\{z > \epsilon\}$ for $\epsilon > 0$ when $\mu > 0$ is small. On the other hand, $-\mu\log(z)/z$ generates more repelling force from $z = 0$ than $-\mu\log(z)$. In this sense, $-\mu\log(z)/z$ realizes the constraints in (2) better than $-\mu\log(z)$.

Thus, our proposed minimization model for HARDI denoising is,

$$\inf_d G(d) = |\nabla d|(\Omega) + \lambda \int_\Omega \sum_{i=1}^n |S_i(x) - S_0(x)e^{-d_i(x)}|dx \\ - \mu \int_\Omega \sum_{i=1}^n \Big[H(S_i(x), S_0(x)) \cdot \frac{\log(d_i(x))}{d_i(x)}\Big]dx. \tag{3}$$

In practice for computations, we use and discretize the Euler-Lagrange equations of model (3) in gradient descent form with respect to $d_i(t,x)$. For $i = 1,\ldots,n$, $t > 0$, these are

$$\frac{\partial d_i}{\partial t}(t,x) = \operatorname{div}\Big(\frac{\nabla d_i(t,x)}{\sqrt{\sum_{j=1}^n |\nabla d_j(t,x)|^2}}\Big) - \lambda \cdot S_0(x)e^{-d_i(t,x)} \tag{4}$$

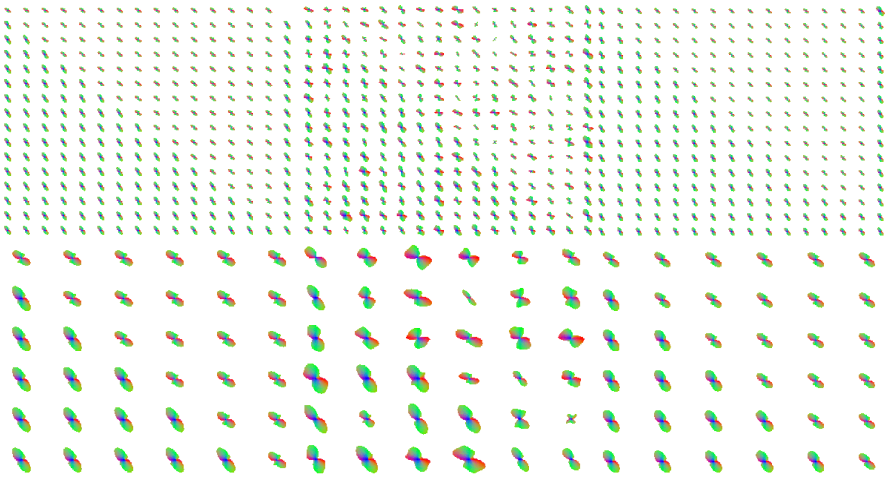

Fig. 1. ODFs. Left: noise-free synthetic data. Middle: noisy data, $M(original, noisy) = 4.2206$. Right: denoised data, $M(original, denoised) = 1.8243$. Ratio $\frac{M(original,noisy)}{M(original,denoised)} = \frac{4.2206}{1.8243} = 2.3135$, larger (better) than for the best result reported in McGraw et al. [25], $\frac{1.0409}{0.6576} = 1.5828$; $rmse(original, noisy) = 17.7079$, $rmse(original, denoised) = 7.7774$ (similar with $rmse(original, denoised) = 7.6367$ from [25]).

$$\cdot \ sign(S_i(x) - S_0(x)e^{-d_i(t,x)}) + \mu \frac{H(S_i(x), S_0(x))(1 - \log(d_i(t,x)))}{d_i(t,x)^2}$$

with boundary conditions. We use finite differences to discretize the above PDE's using an explicit scheme. The final C++ algorithm is computationally efficient.

3 Numerical Results

We recall that in practice we work with a decreasing sequence of values $\mu_k > 0$ and we find minimizers d_k^* for G, with μ substituted by μ_k. The minimizer d_k^* obtained for μ_k is the initial guess for the next minimization with $\mu_{k+1} < \mu_k$. We wish to mention that the visualization of noisy data and denoised results is done also through the ODFs (orientation distribution functions), which are obtained by a postprocessing from the HARDI signal (but the ODFs are not used in our denoising method). Note that calculating ODF of a noisy dataset means that we perform a process of smoothing the data. The ODF is typically calculated from the signal using the Fourier transform relationship between the signal and the diffusion propagator [37], [30]. Since the original noisy data usually violates the constraints we had in the model, if we want to visualize the ODF of the noisy data, then there has to be a pre-processing step to adjust those violating values.

We first show a denoising result of a synthetic 16×16 HARDI data, kindly provided by T. McGraw, for comparison with results from prior work of McGraw

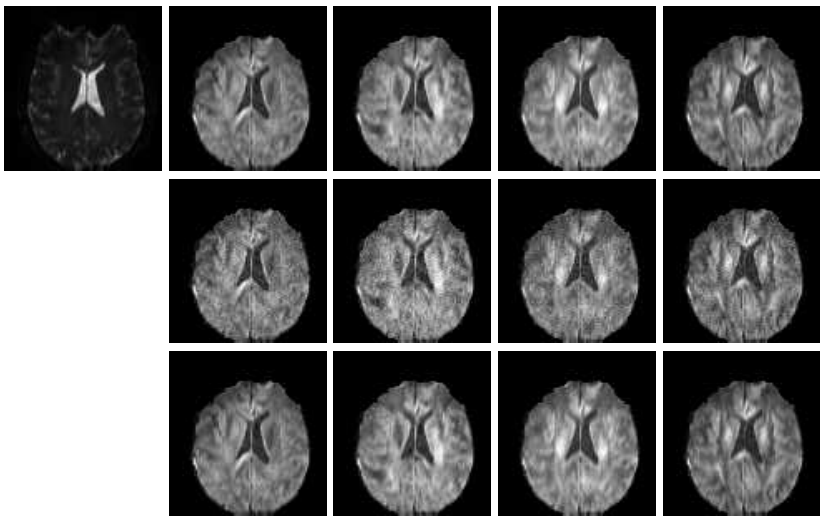

Fig. 2. Denoising experiment of real MRI data. Top: left, S_0; right, original clean slices. 2nd row: noisy slices (artificial Rician noise). 3rd row: denoised slices. $M(original, noisy) = 2.8350$, $M(original, denoised) = 1.1564$. $rmse(original, noisy) = 10.5448$, $rmse(original, denoised) = 4.7268$.

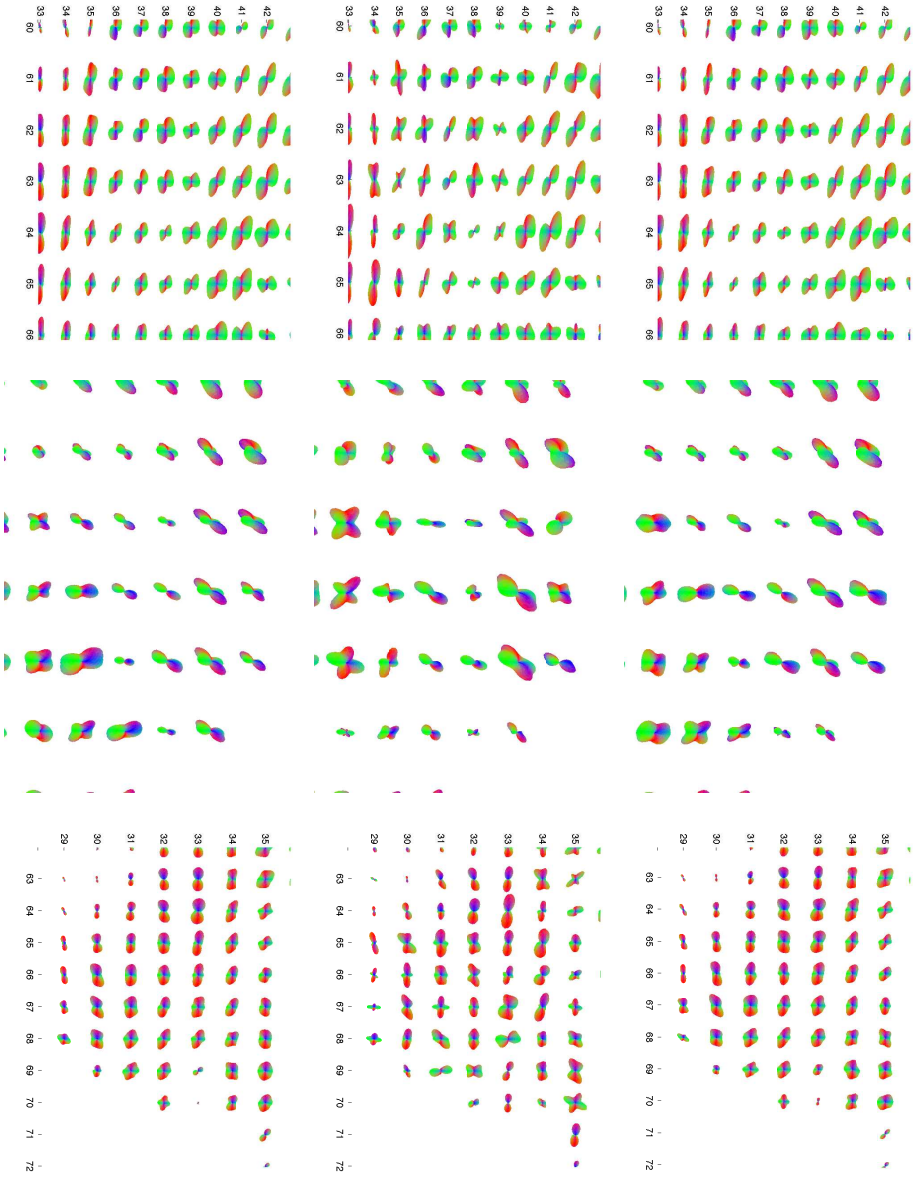

Fig. 3. ODFs of artificial denoising experiment of real MRI HARDI data from Fig. 2. Left: original clean ODFs. Middle: noisy ODFs (artificial Rician noise). Right: denoised ODFs.

et al. [25] (generated using the technique described in [30]). In Fig. 1 we show the ODF visualizations of synthetic, noise-free HARDI data and its noisy version, together with the denoised result.

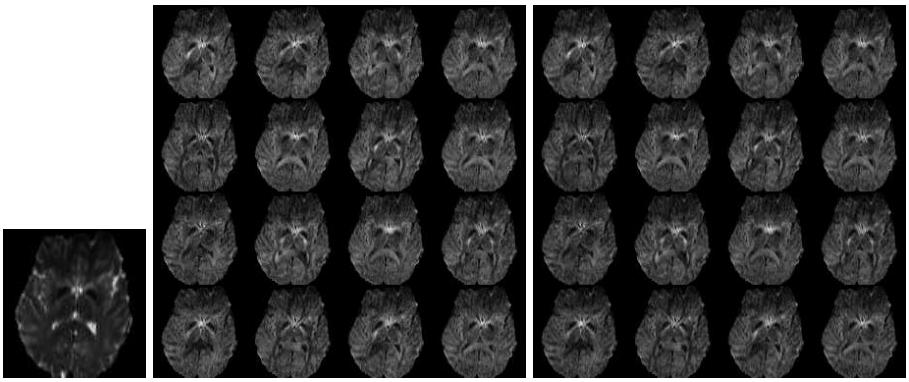

Fig. 4. 19th slice. Left: S_0. Middle: noisy data S. Right: denoised data. Only 16 randomly picked directions out of 94 are shown.

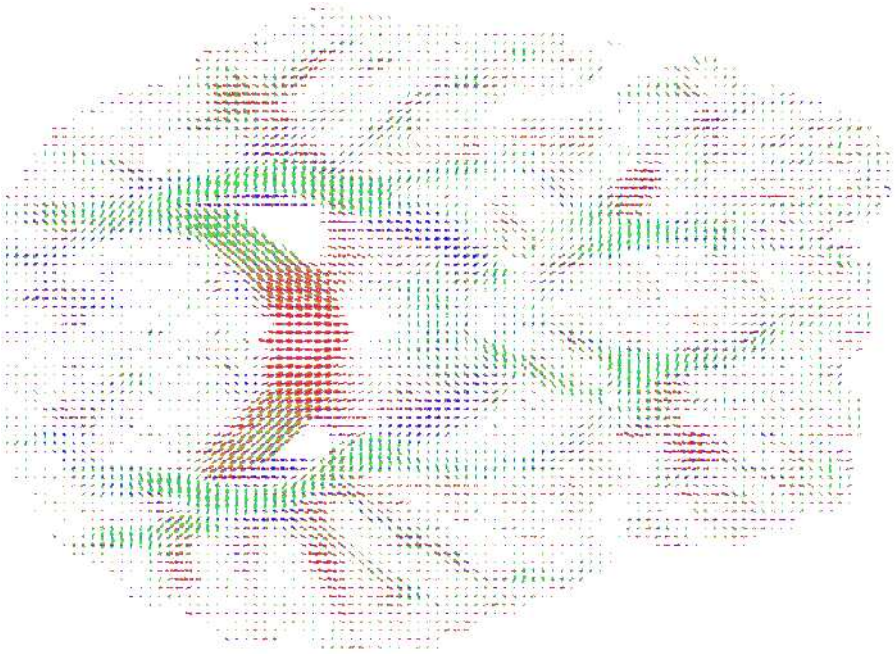

Fig. 5. ODFs: 19th slice of real clinical HARDI brain data

To assess the results'accuracy in the synthetic experiments (since we obtain the ODFs by postprocessing), we compute the mean M over all points of square root of symmetric Kullback-Leibler divergence between two probability densities $p(x), q(x)$ defined by $sKL(p,q) = \frac{1}{2}\int_\Omega \left\{p(x)\log\left(\frac{p(x)}{q(x)}\right)+q(x)\log\left(\frac{q(x)}{p(x)}\right)\right\}dx$. We

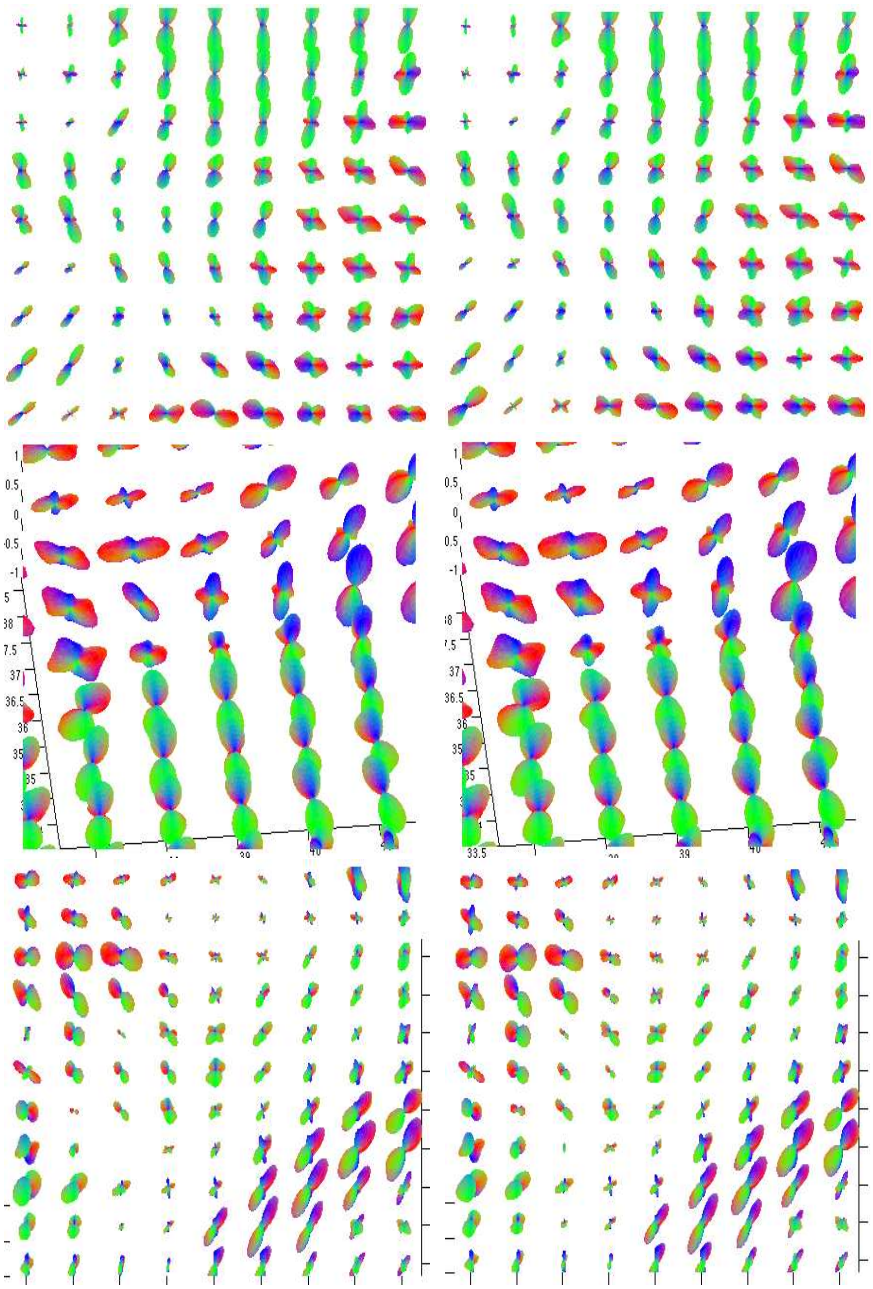

Fig. 6. ODFs of clinical noisy data (left), denoised result (right)

let q be the ODF of the noise free data and in each case, let p be the ODF of either the noisy data or the denoised data. The obtained mean distances are given in Figures 1-2. As another error measure, we also use the root mean square error in Fig. 2, which is obtained directly from our computed quantities, without ODFs.

We show next two experimental results on two real MRI HARDI data sets. We want to take only the brain region into account, thus we use a template that defined the brain region in the image. This led us to use Dirichlet boundary conditions. Since functional (3) is obviously nonconvex, there might be many local minima, which might cause visibly unsatisfying results or some numerical instability. So we need to choose an appropriate initial guess when $t = 0$. Since our minimizer d of (3) should satisfy $d_i \geq 0$, we choose the initial guess d_0 as

$$\begin{aligned}(d_0)_i(x) &= 0.005 && \text{on } \{x \in \Omega :\ H(S_i(x), S_0(x)) > 0\},\\ (d_0)_i(x) &= -\log(S_i(x)/S_0(x)) + 0.1 && \text{on } \{x \in \Omega :\ H(S_i(x), S_0(x)) = 0\}.\end{aligned}$$

We first show another artificial denoising experiment on a real MRI HARDI data set of $n = 30$ diffusion-sensitized gradient directions. Clean and noisy data are available (with artificial Rician noise), kindly provided by E. Iglesias. Slices of S_0, clean, noisy and restored data are shown in Fig. 2, with plots of clean, noisy and denoised ODFs in Fig. 3. We notice visually very good reconstruction while preserving very well the anatomic structure.

Next, we tested our model on a clinical real noisy HARDI dataset with $n = 94$ diffusion-sensitized gradient directions. Briefly, 3D structural brain MRI scans and DT-MRI scans were acquired from healthy young adults on a 4 Tesla Bruker Medspec MRI scanner using an optimized diffusion tensor sequence. Imaging parameters were: TE/TR 92.3/8250 ms, 55 x 2mm contiguous slices, FOV = 23 cm. 105 directional gradients were applied: 11 baseline images with no diffusion sensitization (i.e., T_2-weighted images) and 94 diffusion-weighted images (b-value 1159 s/mm^2) in which gradient directions were evenly distributed on the hemisphere [19]. The reconstruction matrix was 128x128, yielding a 1.8x1.8 mm^2 in-plane resolution. The total scan time was 14.5 minutes. We set S_0 to be the average of the 11 baseline images. Fig. 4 shows slices of the clinical dataset (non-diffusion weighted image S_0 and 16 directions S_i), together with the denoised results. Next we visualize ODFs of the data and we can visually compare the ODFs of the noisy data with the ODFs of the denoised data. To better see the difference between the ODFs of the noisy data and the denoised data, we take some parts of the whole brain image and magnify them especially in regions where fibers are crossing. The data and results are shown in Figures 5-6.

References

1. Alliney, S.: Digital Filters as Absolute Norm Regularizers. IEEE TSP 40(6), 1548–1562 (1992)
2. Barmpoutis, A., Jian, B., Vemuri, B.C., Shepherd, T.M.: Symmetric positive 4th order tensors & their estimation from diffusion weighted MRI. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 308–319. Springer, Heidelberg (2007)

3. Basser, P.J., Pierpaoli, C.: Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI. JMR 111, 209–219 (1996)
4. Beaulieu, C., Allen, P.S.: Water diffusion in the giant axon of the squid: Implications for diffusion-weighted MRI of the nervous system. MRM 32(5), 579–583 (1994)
5. Blomgren, P.: Total Variation Methods for Restoration of Vector Valued Images (Ph.D. thesis), UCLA CAM Report 98-30 (1998)
6. Blomgren, P., Chan, T.F.: Color TV: Total variation methods for restoration of vector-valued images. IEEE TIP 7(3), 304–309 (1998)
7. Jonasson, L., Hagmann, P., Bresson, X., Thiran, J.-P., Wedeen, V.J.: Representing Diffusion MRI in 5D for Segmentation of White Matter Tracts with a Level Set Method. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565, pp. 311–320. Springer, Heidelberg (2005)
8. Chiang, M.C., Klunder, A.D., McMahon, K., de Zubicaray, G.I., Wright, M., Toga, A.W., Thompson, P.M.: Information-theoretic analysis of brain white matter fiber orientation distribution functions. In: Karssemeijer, N., Lelieveldt, B. (eds.) IPMI 2007. LNCS, vol. 4584, pp. 172–182. Springer, Heidelberg (2007)
9. Chiang, M.C., Leow, A.D., Dutton, R.A., Barysheva, M., Rose, S., McMahon, K.L., de Zubicaray, G.I., Toga, A.W., Thompson, P.M.: Fluid Registration of Diffusion Tensor Images Using Information Theory. IEEE TMI 2008 27(4), 442–456 (2008)
10. Cihangiroglu, M., Uluğ, A.M., Firat, Z., Bayram, A., Kovanlikaya, A., Kovanlikaya, İ.: High b-value diffusion-weighted MR imaging of normal brain at 3T. European Journal of Radiology 69(3), 454–458 (2009)
11. Deputte, S., Dierckx, H., Fieremans, E., D'Asseler, Y., Achten, R., Lemahieu, I.: Postprocessing of brain white matter fiber orientation distribution functions. In: Proc. IEEE ISBI: from Nano to Macro, pp. 784–787 (2007)
12. Descoteaux, M., Deriche, R.: High Angular Resolution Diffusion MRI Segmentation Using Region-Based Statistical Surface Evolution. JMIV 33(2), 239–252 (2009)
13. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, Fast and Robust Analytical Q-Ball Imaging. MRM 58(3), 497–510 (2007)
14. Descoteaux, M., Deriche, R.: Mapping neuronal fiber crossings in the human brain. SPIE Newsroom (August 2008)
15. Frank, L.R.: Characterization of anisotropy in high angular resolution diffusion-weighted MRI. MRM 47(6), 1083–1099 (2002)
16. Haro, G., Lenglet, C., Sapiro, G., Thompson, P.M.: On the Non-Uniform Complexity of Brain Connectivity. In: IEEE ISBI: from Nano to Macro, pp. 887–890 (2008)
17. Jian, B., Vemuri, B.C.: A Unified Computational Framework for Deconvolution to Reconstruct Multiple Fibers From Diffusion Weighted MRI. IEEE TMI 26(11), 1464–1471 (2007)
18. Jian, B., Vemuri, B.C., Özarslan, E., Carney, P.R., Mareci, T.H.: A novel tensor distribution model for the diffusion-weighted MR signal. NeuroImage 37(1), 164–176 (2007)
19. Jones, D.K., Horsfield, M.A., Simmons, A.: Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. MRM 42, 515–525 (1999)
20. Le Bihan, D., Breton, E.: Imagerie de diffusion in vivo par resonance magnétique nucléaire. CRAS 301, 1109–1112 (1985)
21. Le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., Laval-Jeantet, M.: MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. Radiology 161, 401–407 (1986)

22. Le Bihan, D., Poupon, C., Amadon, A., Lethimonnier, F.: Artifacts and Pitfalls in Diffusion MRI. JMRI 24, 478–488 (2006)
23. Lenglet, C., Campbell, J.S.W., Descoteaux, M., Haro, G., Savadjiev, P., Wassermann, D., Anwander, A., Deriche, R., Pike, G.B., Sapiro, G., Siddiqi, K., Thompson, P.M.: Mathematical Methods for Diffusion MRI Processing. NeuroImage. In: Thompson, P.M., Miller, M.I., Poldrack, R., Nichols, T. (eds.) Special Issue on Mathematics in Brain Imaging, November 13 (2008)
24. Leow, A.D., Zhu, S., Zhan, L., McMahon, K., de Zubicaray, G.I., Meredith, M., Wright, M.J., Toga, A.W., Thompson, P.M.: The Tensor Distribution Function. MRM 61(1), 205–214 (2008)
25. McGraw, T., Özarslan, E., Vemuri, B.C., Chen, Y., Mareci, T.: Denoising and visualization of HARDI data. REP-2005-360, CISE, Univ. of Florida (2005)
26. McGraw, T., Vemuri, B.C., Yezierski, B., Mareci, T.: Segmentation of High Angular Resolution Diffusion MRI Modeled as a Field of von Mises-Fisher Mixtures. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3953, pp. 463–475. Springer, Heidelberg (2006)
27. McGraw, T., Vemuri, B.C., Yezierski, B., Mareci, T.: von Mises-Fisher Mixture model of the diffusion ODF. In: ISBI 2006: From Nano to Macro, pp. 65–68 (2006)
28. Merboldt, M., Hanicke, W., Frahm, J.: Self-diffusion NMR imaging using stimulated echoes. JMR 64, 479–486 (1985)
29. Nocedal, J., Wright, S.J.: Numerical Optimization. Springer Series in Operations Research. Springer, Heidelberg (1999)
30. Özarslan, E., Vemuri, B.C., Mareci, T.H.: Generalized scalar measures for diffusion MRI using trace, variance and entropy. MRM 53(4), 866–876 (2005)
31. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D 60(1-4), 259–268 (1992)
32. Srivastava, A., Jermyn, I., Joshi, S.: Riemannian Analysis of Probability Density Functions with Application in Vision. In: IEEE CVPR, pp. 1–8 (2007)
33. Stejskal, E.O., Tanner, J.E.: Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. Journal of Chemical Physics 42, 288–292 (1965)
34. Tschumperlé, D.: PDE-Based Regularization of Multivalued Images and Applications. PhD Thesis Univ. of Nice-Sophia Antipolis, France (2002)
35. Tschumperlé, D., Deriche, R.: Anisotropic Diffusion Partial Differential Equations in Multi-Channel Image Processing: Framework and Applications. In: Book chapter in Advances in Imaging and Electron Physics (AIEP). Academic Press, London (2007)
36. Tschumperlé, D., Deriche, R.: Vector-Valued Image Regularization with PDE's: A Common Framework for Different Applications. IEEE TPAMI 27(4), 506–517 (2005)
37. Tuch, D.S.: Diffusion MRI of Complex Tissue Structure. Ph. D. Thesis, Harvard-MIT Division of Health Sciences and Technology (2002)
38. Tuch, D.S., Weisskoff, R.M., Belliveau, J.W., Wedeen, V.J.: High angular resolution diffusion imaging of the human brain. In: Proc. 7th Annual Meeting of ISMRM, Philadelphia, PA, p. 321 (1999)
39. Tuch, D.S., Reese, T.G., Wiegell, M.R., Makris, N., Belliveau, J.W., Wedeen, V.J.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. MRM 48, 577–582 (2002)
40. Tuch, D.S., Reese, T.G., Wiegell, M.R., Wedeen, V.J.: Diffusion MRI of complex neural architecture. Neuron 40, 885–895 (2003)

41. Tuch, D.S.: Q-ball imaging. MRM 52, 1358–1372 (2004)
42. Verma, R., Khurd, P., Davatzikos, C.: On Analyzing Diffusion Tensor Images by Identifying Manifold Structure Using Isomaps. IEEE TMI 26(6), 772–778 (2007)
43. Wedeen, V.J., Hagmann, P., Tseng, W.Y., Reese, T.G., Weisskoff, R.M.: Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn. Reson. Med. 54(6), 1377–1386 (2005)
44. Zhan, L., Chiang, M.C., Barysheva, M., Toga, A.W., McMahon, K.L., de Zubicaray, G.I., Meredith, M., Wright, M.J., Thompson, P.M.: How Many Gradients are Sufficient in High-Angular Resolution Diffusion Imaging (HARDI)? In: MICCAI 2008, MICCAI DTI Workshop (2008)