

Fiber Continuity Based Spherical Deconvolution in Spherical Harmonic Domain

Marco Reisert¹ and Henrik Skibbe²

¹ Department of Diagnostic Radiology, Medical Physics, University Medical Center, Breisacher Street 60a, 79106 Freiburg, Germany

² Graduate School of Informatics, Kyoto University, Gokasho, 611-0011 Uji, Kyoto, Japan

Abstract. The reliable estimation of fiber orientation distributions from diffusion-sensitized magnetic resonance imaging is an important processing step for revealing structural connectivity in the human brain. Fiber orientation distributions are usually represented in spherical harmonic (SH) domain, but many modern spatial regularization techniques were not yet formulated in SH-domain such that they cannot benefit from the well known advantages. In this article we propose a novel SH-formulation of the contour enhancement propagator (fiber continuity) which enables to implement a variety of processing steps in SH-domain.

1 Introduction

Magnetic resonance imaging (MRI) has the potential to visualize non-invasively the fibrous structure of the human brain white matter [5]. Based on the directional dependency of water diffusivity in fibrous tissue it is possible to reveal underlying connectivity information. There are numerous methods for estimating orientation distributions on the basis of the diffusion-weighted MR-signal. We will focus on spherical deconvolution [8], which is one way to estimate the so called fiber orientation distributions (FOD) on the basis of the diffusion-weighted MR-signal. The idea is based on a model-driven deconvolution scheme to turn the diffusion weighted MR-signal into a FOD. Fiber orientation distributions are usually represented in spherical harmonic (SH) domain. However, sophisticated spatial regularizers based on $\mathbb{R}^3 \times S^2$ -diffusion [1,2,3,6] are difficult to implement in SH domain (apart from the simple isotropic ones [4]). They all rely on an equiareal discretization of the two-sphere S_2 . Indeed, there are implementations [3] that use spherical harmonics as an intermediate S_2 -interpolation scheme, but they cannot benefit from the well-known advantages of the spherical harmonic representation such as the compact and memory efficient storage, the analytic and efficient computations of S_2 -convolutions, and the closeness under rotations. In this article we overcome these restrictions by showing how the generators of diffusion and convection on $\mathbb{R}^3 \times S^2$ can be described and implemented in terms of spherical harmonics and how the spherical discretization as proposed in [3,6] can be avoided.

2 Theory

We are interested in solving or propagating a partial differential equation of the form $\partial_t \phi(\mathbf{r}, \mathbf{n}, t) = H\phi(\mathbf{r}, \mathbf{n}, t)$ where ϕ is a function (typically the FOD) on $\mathbb{R}^3 \times S^2$ and

H a linear differential operator in the orientation $\mathbf{n} \in S_2$ and position $\mathbf{r} \in \mathbb{R}^3$. The orientation \mathbf{n} corresponds the local fiber direction at location \mathbf{r} in the brain. In [3] it was shown that, if H generates directed diffusion/convection, then it is a linear/quadratic form in $\mathbf{n}^\top \nabla$ (convection and anisotropic diffusion) and linear in $\Delta = \partial_x^2 + \partial_y^2 + \partial_z^2$ and \mathcal{J}^2 , which denotes the Laplace-Beltrami operator on the two-sphere. Our goal is to understand the action of the generator H , if the field ϕ is written in terms of a spherical harmonics expansion

$$\phi(\mathbf{r}, \mathbf{n}) = \sum_{j=0}^{\infty} \mathbf{f}^j(\mathbf{r})^\top \mathbf{Y}^j(\mathbf{n}) = \sum_{j=0}^{\infty} \sum_{m=-j}^{m=j} \overline{f_m^j(\mathbf{r})} Y_m^j(\mathbf{n})$$

where the spherical tensor fields \mathbf{f}^j are obtained by the projections $\mathbf{f}^j = \frac{1}{2j+1} \langle \mathbf{Y}^j, \phi \rangle$, where \mathbf{Y}^j are semi-Schmidt normalized spherical harmonics. Hence we are interested in matrix elements $\hat{H}_{j'm'}^{j'm} = \langle Y_{m'}^{j'}, H Y_m^j \rangle$ of H in spherical harmonic representation such that the propagating equation can be written as $\partial_t f_m^j(\mathbf{r}, t) = \sum_{j',m'} \hat{H}_{j'm'}^{j'm} f_{m'}^{j'}(\mathbf{r}, t)$ where $\hat{H}_{j'm'}^{j'm}$ is a differential operator in \mathbf{r} , but purely algebraic in the orientation coordinate. The spherical Laplace-Beltrami operator is well known in this representation $\langle Y_m^j, \mathcal{J}^2 \phi \rangle = -j(j+1) f_m^j$. Also for the spatial Laplacian Δ the result is trivial $\langle Y_m^j, \Delta \phi \rangle = \Delta f_m^j$ (see [7]). But for the directed convection $\mathbf{n} \cdot \nabla$ and the directed diffusion $(\mathbf{n} \cdot \nabla)^2$ generator the results are more intricate. In our companion report [7] we give the general proof for $SE(3)$ diffusion, but state here the more simple case for $\mathbb{R}^3 \times S_2$. For the convection generator one finds

$$\langle Y_m^j, (\mathbf{n} \cdot \nabla) \phi \rangle = \sum_{\substack{j'=-1,1 \\ m=m'+q}} \frac{2j'+1}{2j+1} \langle jm|j'm', 1q \rangle \langle j0|j'0, 10 \rangle \partial_q^1 f_{m'}^{j'} \quad (1)$$

where $\partial_m^j = R_m^j(\nabla)$ is the so called spherical derivative operator, which is a j -homogeneous differential as the solid harmonic $R_m^j(\mathbf{r}) = r^j Y_m^j(\mathbf{r}/r)$ is a j -homogeneous polynomial. The $\langle jm|j'm', 1q \rangle$ are the so called Clebsch Gordan coefficients, which are just real numbers. The sum runs over the indices j', m' and q , but note that due to the selection rules of the Clebsch Gordan coefficients only very few terms remain (at most 6). On the other hand the diffusion generator takes the form

$$\langle Y_m^j, (\mathbf{n} \cdot \nabla)^2 \phi \rangle = \frac{\Delta}{3} + \frac{2}{3} \sum_{\substack{j'=-2,0,2 \\ m=m'+q}} \frac{2j'+1}{2j+1} \langle jm|j'm', 2q \rangle \langle j0|j'0, 20 \rangle \partial_q^2 f_{m'}^{j'} \quad (2)$$

Again the sum runs over j', m' and q and can contain at most 10 terms.

Spatially Regularized Spherical Deconvolution. The goal of spherical deconvolution is to find a FOD ϕ such that $\|\mathbf{F}\phi - S\|^2$ is minimized. Here S is typically the diffusion weighted measurement normalized by the b0-image. The operator \mathbf{F} denotes the spherical convolution with the so called fiber response function. As \mathbf{F} is highly non-regular a regularizer is needed. In this article we consider the contour enhancement [3] regularizer (also called fiber continuity [6]) $R_{FC}(\phi) = \lambda \int (\mathbf{n} \cdot \nabla \phi)^2 d\mathbf{x}d\mathbf{n}$ that prevents

‘arbitrary’ smoothing and preserves and emphasize the fibrous nature of the data. Instead of smoothing isotropically like done in [4] the smoothing takes place only along the current fiber direction. The idea behind is that a fiber which goes through voxel \mathbf{r} with direction \mathbf{n} is very likely to go also through voxel $\mathbf{r} + \epsilon\mathbf{n}$ for some small ϵ .

The choice of regularization strength is a crucial issue. We found that too strong regularization emphasizes discretization artifacts. On the other hand too low values lead to less stable results. We found by a simple visual inspection of a simulated crossing a value of $\lambda = 0.005$ is a good trade-off between stability and accuracy.

Optimization and Implementation. In order to find the optimum of $J_{\text{reg}} = \|\mathbf{F}\phi - S\|^2 + R_{\text{FC}}(\phi)$, one has to compute its derivative and set it to zero. This leads to $(\mathbf{F}^T \mathbf{F} - \lambda(\mathbf{n} \cdot \nabla)^2)\phi = \mathbf{F}^T S$. The second term of the equation is formulated with the proposed SH representation from equation (2) and solved by an ordinary conjugate gradients (CG) scheme. The operator \mathbf{F} is simple to implement in terms of spherical harmonics, because it is a diagonal matrix. For details about the discretized sphere (512 directions) approach we refer to our companion report [7]. For both methods the CG algorithm was iterated 100 times, which was enough for convergence. For initialization a ‘zero’ FOD was used.




Fig. 1. Two simulated crossing situations to compare the SH representation with the angular discrete version. The simulation was performed with 64 gradient directions at a b-value of $1000s/mm^2$ and with a diffusion coefficient of $10^{-3}mm^2/s$

3 Experiments

The goal of the experiment is to compare the proposed spherical harmonic representation with the discrete angular representation of the fiber orientation distributions. Therefore, we simulated the MR-signal of a crossing region with 64 gradient directions. The signal was simulated by using the standard exponential model $S_{\mathbf{n}_{\text{fib}}}(\mathbf{n}) = e^{-bD(\mathbf{n} \cdot \mathbf{n}_{\text{fib}})^2}$, that is, no diffusion perpendicular to a fiber is assumed. We have chosen $bD = 1$, which emulates a b-value of $1000\text{s}/\text{mm}^2$ and a typical diffusion coefficient for the human brain. The generated signal was distorted by Rician noise $S_{\text{noisy}} = \sqrt{(S + n_{\text{real}})^2 + n_{\text{imag}}^2}$, where n_{real} and n_{imag} are normally distributed real numbers with standard deviation σ . The signal-to-noise ratio is defined as $\text{SNR} = 1/\sigma$, that is, the SNR is calculated with respect to the $b = 0$ measurement. The crossing was created on a 24×24 voxel grid, where the tracts of the crossing are on average 5 voxels thick. To get an impression look at Figure 1. To measure the performance of the deconvolution method, the local maxima of the estimated FODs are extracted and compared to the ground truth direction (see [7] for details). A ground truth direction is said to be detected, when it is in a range of 10° from a detected direction. To measure the performance we used precision, recall and the f-score¹. The precision rate reflects how many of the directions detected by the algorithm are true fiber directions, on the other hand, the recall rate gives the percentage of true fiber directions that were detected by the algorithm. The f-score is just a measure reflecting the trade-off between recall and precision. To generate the performance measures the simulations were repeated 100 times.

Method Comparison and Discussion. In Figure 1 we visually compare the spherical harmonic implementation for a cutoff of $L = 8$ with the discrete version with 512 directions on the sphere. We consider a crossing angle of 35° and 45° . For both situations one fiber bundle direction was chosen along the underlying Cartesian coordinate axis. For the larger crossing, which is easier to resolve, we assumed a relatively low SNR of 7. For this case one can see that the discrete version is more susceptible to noise than the spherical harmonic version, which is not astonishing due to the implicit regularization by the finite SH-cutoff of $L = 8$. For the smaller crossing angle of 35° a higher SNR of 20 was assumed. In this case a SH-representation of $L = 8$ is nearly at its limits to discriminate between both directions, while the discrete version can still well distinguish. Further, one can observe that for both methods the FODs along the horizontal Cartesian axis are sharper than along the skew axis. But this effect is more prominent for the discretized version than for the spherical harmonic representation.

In Figure 2 we show quantitative results. The crossing was simulated for varying crossing angles between 30° and 90° . Additionally we varied the absolute pose α of the crossing. For $\alpha = 0^\circ$ the horizontal bundle is aligned with underlying x-axis of

¹ Let TP be the number of successfully found ground truth directions, let FP be the number of detections that are not in a range of 10 degree to a ground truth direction, and let FN the number of ground truth direction that are not detected, then precision = $\text{TP}/(\text{TP} + \text{FP})$ and recall = $\text{TP}/(\text{TP} + \text{FN})$ and f-score = $2 \text{ precision} \cdot \text{recall}/(\text{precision} + \text{recall})$.




Fig. 2. Detection accuracies in terms of precision/recall and f-score for different crossing angles and absolute angles α . The measurement was simulated at a $SNR = 20$ with a value of $bD = 1$ and 64 gradient directions.

the Cartesian grid. With growing α the whole configuration is rotated clockwise. The crossing was simulated at an SNR of 20, which is a realistic scenario. As a baseline experiment we show results of the so called Constrained Spherical Deconvolution (CSD) approach [9], where an additional positivity constraint is used to obtain more stable results (original implementation obtained by the author). Apart from the CSD approach the results obviously depend on the absolute angle of the configuration due to the underlying Cartesian voxel grid. The discrete approach has advantages for small crossing angles, but suffers from a more severe dependency on the absolute orientation of the configuration and less precision. For crossing angles above 45° the SH-based approach solves the task nearly without any error.

In Figure 3 the results are further investigated by scatter plots in the ϕ, θ -plane. Each extracted local maxima is represented as a blue dot in the ϕ, θ -plane, while the true directions are depicted by solid red lines and the tolerance used for quantitative evaluation

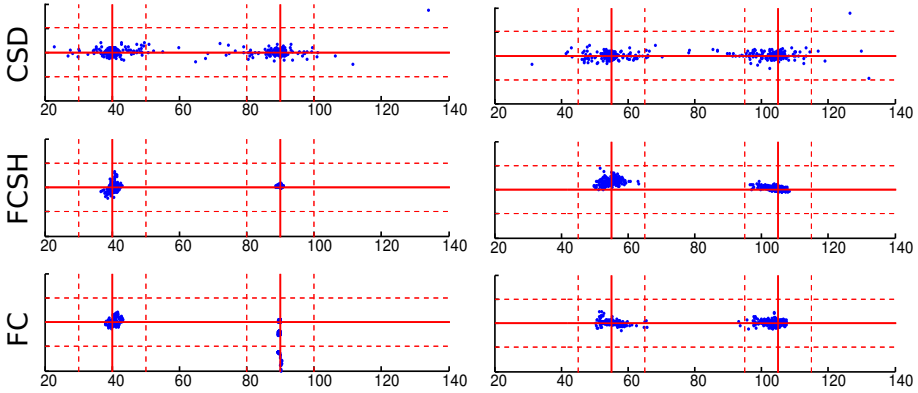


Fig. 3. Scatter plots in the ϕ - θ plane for the crossing configuration with a crossing angle of 50° . The angle ϕ is plotted along the x-axis, θ along the y-axis of the scatter plots. The intersections of the thick red lines indicate the expected ground truth directions. The dotted lines indicate the 10° detection tolerance. On the left, the results for the configuration with an absolute angle of $\alpha = 0$ are shown, on the right the results for $\alpha = 15^\circ$ are given.

by dashed red lines. With the same parameters as above a crossing of 50° was simulated and reconstructed by the three methods. The crossing was simulated twice, for an absolute angle of $\alpha = 0$, i.e. one direction is along the x-axis, and secondly for an absolute angle of $\alpha = 15^\circ$. For $\alpha = 0$ the SH-based approach is able to perfectly resolve the direction along the x-axis ($\phi = 90^\circ$), while the other direction ($\phi = 40^\circ$) is a bit more blurry. On the other hand, the discrete approach has severe problems with the $\phi = 90^\circ$ direction, which explains the lack of precision, i.e. apart from the true direction there are some additional local maxima that produce false positives. The main reason is the interplay of the 64 gradient directions and the 512 discrete directions of the FOD. The effect is reduced by an increase of measurement directions. But also the SH-based approach has problems when the number of measurement directions is too low. They are revealed for an absolute angle of $\alpha = 15^\circ$. Besides the uncertainty caused by the measurement noise one can observe a systematic bias. For example, for the direction along $\phi = 55^\circ$, the center of the distribution is shifted in θ by approximately 5° . Also for the other direction the distribution is a bit squeezed. We also found that the main reason is the low number of measurement directions. For example, for 128 gradient directions the estimated directions are unbiased. Another way to reduce the effect is to decrease the expansion cutoff of the spherical harmonic representation. For $L = 6$ and 64 gradient directions the estimates do not show a bias. To conclude the differences: both methods have problems when the number of measurements becomes too low. While the discrete approach shows scattered, multimodal distributions, the SH-based approach shows a slight systematic bias but the distributions stay unimodal.

In Figure 4 we show a real world example of the human brain. The setting of the measurement is nearly the same like in the simulations. A b-value of $1000s/mm^2$ and 61 gradient directions were used with an isotropic resolution of $2mm$. A kernel of the form $S_{\mathbf{n}_{\text{fib}}}(\mathbf{n}) = e^{-bD(\mathbf{n} \cdot \mathbf{n}_{\text{fib}})^2}$ with $bD = 1$ was used as a model for deconvolution.

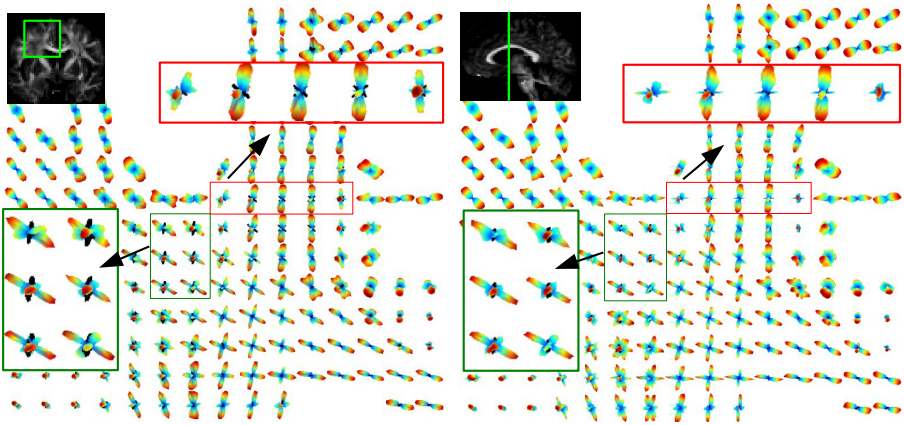


Fig. 4. A real world example of the human brain. A coronal section is shown. Left: results of the SH-based approach with $L = 8$, right: the discrete approach with 512 directions.

Figure 4 shows a coronal section in glyph representation. The green rectangle highlights a region with the largest differences. The red rectangle shows a regions where the discrete approach shows a direction which does not appear for the SH-based approach. Whether the direction is true or not is difficult to say, but the fact that it precisely points along the x-axis makes it dubious.

Memory Consumption and Running Time. The memory consumption of the SH-based and discrete approach is easy to compare. We want to consider the above real world experiment as an example. The whole volume has a size of $96 \times 96 \times 60$. Exploiting the symmetry and realness each voxels consumes $(L + 2)^2/4$ complex numbers in SH representation, resulting in $50 \cdot 8$ bytes for $L = 8$ and double precision. Overall, one volume needs $96^2 \cdot 60 \cdot 50 \cdot 8 = 220$ MB. On the other hand, in discrete representation with 512 direction needs $96^2 \cdot 60 \cdot 512 \cdot 8 = 2264$ MB, which is 10-times more compared to the SH-representation. Recall, that the conjugate gradient algorithm needs four instances of the volume in memory. The running time is more difficult to compare. The application of $(\mathbf{n} \cdot \nabla)^2$ and the spherical convolution \mathbf{H} are the important processing steps. In case of the discrete approach one has to compute for each of the 512 components six second order finite differences to get $(\mathbf{n} \cdot \nabla)^2 \phi$. For the SH-based approach one also have to compute six finite differences, but the linear combinations of them to obtain the final values are more expensive. In practice we found that one application of $(\mathbf{n} \cdot \nabla)^2$ with 256 discrete directions is comparable to an SH-based application of $(\mathbf{n} \cdot \nabla)^2$ with $L = 10$. The computation of the operator \mathbf{H} is negligible in SH-representation, while it is the bottleneck for the discrete approach. The running time can heavily differ depending on the used library. For example, the highly optimized BLAS matrix-multiplication shipped with MATLAB is 10 times faster than the standard non-optimized version. To give an example, to reconstruct the above volume with 100 CG iteration with

512 directions takes on a *Intel Xeon X7560 @ 2.27GHz* about 20 minutes with a highly optimized multiplication. For comparison, our SH-implementation for $L = 8$ takes about 10 minutes on the same machine.

4 Conclusion

We proposed a SH-formulation for a family of diffusion/convection generators that are of particular interest for the spatial regularization of FODs. The SH-formulation offers advantages in terms of running time and memory consumption and keeps approximately the performance and behavior of the original discrete implementation. While the discrete implementation tends to produce additional local maxima and spikey distributions, the distributions of the SH-implementation are more smooth, but shows a bit worse performance for smaller crossing angles.

Acknowledgements. This work was partly supported by 'Bioinformatics for Brain Sciences' under the Strategic Research Program for Brain Sciences, by the MEXT (Japan), the 'Baden-Württemberg Stiftung - Eliteprogramme for Postdocs' and by the Deutsche Forschungsgemeinschaft (DFG) Grant KI 3286/2-1. The authors would also like to thank D. Tournier for providing his original CSD implementation.

References

1. Barmpoutis, A., Vemuri, B.C., Howland, D., Forder, J.R.: Extracting tractosemas from a displacement probability field for tractography in DW-MRI. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part I. LNCS, vol. 5241, pp. 9–16. Springer, Heidelberg (2008)
2. Delputte, S., Dierckx, H., Fieremans, E., D'Asseler, Y., Achten, E., Lemahieu, I.: Postprocessing of brain white matter fiber orientation distribution functions. In: ISBI 2007, pp. 784–787 (2007)
3. Duits, R., Franken, E.: Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of hardi images. *International Journal of Computer Vision* 92, 231–264 (2011)
4. 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)
5. Jones, D.K. (ed.): *Diffusion MRI: Theory, Methods and Applications*. Oxford University Press (2010)
6. Reisert, M., Kiselev, V.: Fiber continuity: An anisotropic prior for odF estimation. *IEEE Trans. Med. Imaging* 30(6), 1274–1283 (2011)
7. Reisert, M., Skibbe, H.: Left-Invariant Diffusion on the Motion Group in terms of the Irreducible Representations of $SO(3)$. ArXiv e-prints (February 2012)
8. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *Neuroimage* 23(3), 1176–1185 (2004)
9. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35(4), 1459–1472 (2007)