

Accelerated Diffusion Operators for Enhancing DW-MRI

P. Rodrigues¹, R. Duits^{1,2}, B.M. ter Haar Romeny¹ and A. Vilanova¹

¹ Department of Biomedical Engineering,

² Department of Mathematics and Computer Science,
Eindhoven University of Technology,
WH 2.111, 5600 MB Eindhoven, The Netherlands

Abstract

High angular resolution diffusion imaging (HARDI) is a MRI imaging technique that is able to better capture the intra-voxel diffusion pattern compared to its simpler predecessor diffusion tensor imaging (DTI). However, HARDI in general produces very noisy diffusion patterns due to the low SNR from the scanners at high b-values. Furthermore, it still exhibits limitations in areas where the diffusion pattern is asymmetrical (bifurcations, splaying fibers, etc.). To overcome these limitations, enhancement and denoising of the data based on context information is a crucial step. In order to achieve it, convolutions are performed in the coupled spatial and angular domain. Therefore the kernels applied become also HARDI data. However, these approaches have high computational complexity of an already complex HARDI data processing. In this work, we present a framework for HARDI data enhancement and completion. The convolution operators are optimized by: pre-calculating the kernels, analysing kernels shape and utilizing look-up-tables concept. We provide an increase of speed, compared to previous brute force approaches of simpler kernels. These methods can be used as a preprocessing for tractography and lead to new ways for investigation of brain white matter.

Categories and Subject Descriptors (according to ACM CCS): Image Processing and Computer Vision [I.4.3]: Enhancement/Smoothing

1. Introduction

Diffusion Weighted imaging is a fairly new MRI acquisition Technique, first introduced by [BML94]. By measuring the directional pattern of local water diffusion, it has the capability to non-invasively allow the inspection of biological tissue such as the brain.

In Diffusion Tensor Imaging (DTI), the prominent local orientation of the fiber bundles can be estimated. In DTI the local diffusivity pattern is approximated by a 2^{nd} -order diffusion tensor (DT). Although simple and with established mathematical frameworks, these DTs fail to capture more complex fiber structures such as crossings, bifurcations and splaying configurations.

Approaches based on High Angular Resolution Diffusion Imaging (HARDI) were pioneered by Tuch [Tuc02]. In HARDI more sophisticated models are employed to reconstruct more complex fiber structures and to better capture the

intra-voxel diffusion pattern. Some of the proposed models include high-order tensors [OM03], mixture of Gaussians [Tuc02, JV07], spherical harmonic (SH) transformations [Fra02], diffusion orientation transform (DOT) [OSV*06], orientation distribution function (ODF) [DAFD07] using the Q-ball imaging [Tuc04], and the spherical deconvolution approach [TCGC04].

It is important to note that all of the diffusion weighted MRI modelling techniques model functions that reside on a sphere. For simplicity we will refer them as spherical distribution function (SDF). Whereas the physical meaning of these SDFs can be different (a probability density function (PDF), iso-surface of a PDF, ODF, FOD, etc.), in all cases they characterize the intra-voxel diffusion process, i.e. the underlying fiber distribution within a voxel. Due to the limitations in acquisition, the SDF is always antipodally symmetric and therefore can only model single fiber

tracts or symmetric fiber crossing configurations. Furthermore, HARDI produces, in general, noisy diffusion patterns due to the low SNR from the scanners at high b-values. To overcome these limitations, postprocessing of the data is crucial. As commonly done in image processing, the noise can be reduced and the data enhanced by taking into account the information in a close neighborhood (i.e. the context).

Previous research has been done on diffusion (or similar regularization) of DTI/HARDI images [FB, Flo08, HMH*06], however they do so considering the spatial and orientational domains separately. In these approaches diffusion is only performed over the spherical function per voxel (i.e. the angular part). Furthermore, by not considering the coherence in the neighborhood, these methods often fail at interesting locations where fibers cross or diverge.

In recent promising work the diffusion process is done considering the full domain, i.e. considering both spatial and orientational neighborhood information. In [ABF08] the estimated asymmetric spherical functions, called tractosemas, are able to model local complex fiber structures using inter-voxel information. Duits and Franken [DF09] proposed a framework for the cross-preserving smoothing of HARDI images by closely modelling the stochastic processes of water molecules in oriented fibrous structures. These approaches, however, increase the complexity of already complex and computationally heavy HARDI data.

In the presented work, we establish a faster framework for noise removal and enhancement of HARDI datasets. We optimize the convolution operators by: pre-calculating the kernels, analysing kernel's shape; and accelerating convolution using look-up-tables concept. Compared to previous brute force approaches, we provide a significant increase of speed, enabling a contextual processing framework of HARDI data. Thus, a basis for more accurate tractography is established, leading to new ways for the investigation of brain's white matter.

In Section 2 we start by establishing the mathematical basis on which the convolution method lives. The accelerated convolution framework is presented in Section 3. Following, in Section 4, we present experimental results, both in artificial and real HARDI data, supporting the validity and improvements of the method.

2. Background

In this section we will provide a self-contained introduction to convolution of HARDI data over the joined domain of positions and orientations (so-called $\mathbb{R}^3 \times S^2$ -convolutions). The several options concerning the kernels for these convolutions will also be addressed.

2.1. Theory

Diffusion weighted MRI modelling techniques estimate functions that reside on a sphere, the spherical distribution

functions (SDF). Therefore, a HARDI image is a function not only on positions but also on orientations:

$$\mathcal{U} : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+ : \mathcal{U}(\mathbf{y}, \tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma})) \quad (1)$$

This means that on every position $\mathbf{y} \in \mathbb{R}^3$, the probability of finding a water particle in a certain direction

$$\tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma}) = (\sin \tilde{\beta}, -\sin \tilde{\gamma} \cos \tilde{\beta}, \cos \tilde{\gamma} \cos \tilde{\beta})^T \in S^2, \quad (2)$$

i.e. a point on the sphere parameterized by $\tilde{\beta} \in [-\pi, \pi]$ and $\tilde{\gamma} \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, is given as a scalar. For visualization and processing purposes, these functions are usually discretized by nearly uniform sampling the sphere using a method such as tessellation of an icosahedron (see Figure 1).



Figure 1: Discrete samplings of the sphere corresponding to order 1, 2 and order 3 of tessellation of an icosahedron, with correspondent 12, 42 and 162 points.

To stress the coupling between orientation and positions we write $\mathbb{R}^3 \times S^2$ rather than $\mathbb{R}^3 \times S^2$. Intuitively, it is the space of fiber fragments on which the rigid motion group of 3D-rotations and translations acts by means of

$$g^{-1}(\mathbf{y}, \mathbf{n}) = (R^{-1}(\mathbf{y} - \mathbf{x}), R^{-1}\mathbf{n}), \quad (3)$$

where $\mathbf{x} \in \mathbb{R}^3$ is the spatial translation vector and R is the 3D-rotation and where $g = (\mathbf{x}, R)$ is the concatenation of the translation and the rotation. For a formal mathematical definition of $\mathbb{R}^3 \times S^2$ see [DF, DF09].

2.2. Convolutions

Obviously, an operator $U \mapsto \Phi(U)$ on an SDF should be Euclidean invariant (independent on a choice of orthonormal coordinate system), which means that the operator should commute with (3). In other words rotating and translating HARDI input $U : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}$ corresponds to rotating and translating the output $\Phi(U) : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}$. If such operators, designed for smoothing and enhancement of HARDI data, are linear then these operators can be written as a HARDI-convolution:

$$\begin{aligned} (\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) &= \int_{\mathbb{R}^3} \int_{S^2} k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}') \mathcal{U}(\mathbf{y}', \mathbf{n}') d\mathbf{y}' d\sigma(\mathbf{n}') \\ &= \int_{\mathbb{R}^3} \int_{S^2} p(R_{\mathbf{n}'}^T(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}'}^T(\mathbf{n})) \mathcal{U}(\mathbf{y}', \mathbf{n}') d\mathbf{y}' d\sigma(\mathbf{n}') \end{aligned} \quad (4)$$

where

- \mathcal{U} denotes the input HARDI dataset.
- $(\Phi(\mathcal{U}))$ denotes the output HARDI data set (obtained from the input by convolution with p)
- $k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}')$ is the full kernel in the kernel operator.

- $p(\mathbf{y}, \mathbf{n})$ is the **convolution kernel** related to $k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}')$ by means of

$$p(\mathbf{y}, \mathbf{n}) = k(\mathbf{y}, \mathbf{n}; \mathbf{0}, \mathbf{e}_z), \text{ with } \mathbf{e}_z = (0, 0, 1)^T$$

From this moment, kernels will be noted as $p(\mathbf{y}, \mathbf{n})$, i.e. the a priori probability density of finding a fiber fragment at (\mathbf{y}, \mathbf{n}) given that there is a fiber fragment at $(\mathbf{0}, \mathbf{e}_z)$.

- $R_{\mathbf{n}}$ is any rotation such that $R_{\mathbf{n}}\mathbf{e}_z = \mathbf{n}$. The choice of $R_{\mathbf{n}}$ does not matter as long as p has a symmetry w.r.t. rotations around \mathbf{e}_z (acting from the right), see [DF09, Corr.1].
- σ denotes the usual surface measure on the sphere.

As mentioned previously, convolutions can operate over different domains, obviously, with different outcomes. Consider the special cases of Equation 4:

Spatial domain Filtering can be applied for each of the directions without relating the directions between each other:

$$(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) = \int_{\mathbb{R}^3} q(\mathbf{y} - \mathbf{y}') \mathcal{U}(\mathbf{y}', \mathbf{n}) d\mathbf{y}' \quad (5)$$

This corresponds to (4) if one sets $p(\mathbf{y}, \mathbf{n}) = q(\mathbf{y})\delta_{\mathbf{e}_z}(\mathbf{n})$.

Orientalional domain Filtering can be applied to each voxel independently, i.e. considering each SDF independently from each other. This way, each voxel is smoothed locally:

$$(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) = \int_{S^2} r(R_{\mathbf{n}}^T \mathbf{n}) \mathcal{U}(\mathbf{y}, \mathbf{n}') d\sigma(\mathbf{n}') \quad (6)$$

This corresponds to (4) if one sets $p(\mathbf{y}, \mathbf{n}) = \delta_{\mathbf{0}}(\mathbf{y})r(\mathbf{n})$.

However, appropriate treatment of crossings and bifurcations requires regularization along oriented fibers (where position and orientation are coupled) and consequently our a priori fiber extension probabilities $p: \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$ should not consist of a delta-spike in position space nor in orientation space. This means we should not restrict ourselves to (5) and (6). Next we explain how to discretize full convolutions (4) on positions and orientations.

Having a discrete lattice of SDFs (the hardi image \mathcal{U}), the integral over \mathbb{R}^3 in Equation 4 becomes a summation over the lattice. Since, typically, a kernel is stronger around its center (at position \mathbf{x}), a set $N(\mathbf{x})$ can be defined containing the lattice indices neighbour of \mathbf{x} . Additionally, since the SDFs are discretized over the sphere (see Figure 1), the integral over S^2 becomes a summation over tessellation's vectors, the set S . Using these discretizations, Equation 4 becomes:

$$(\Phi(\mathcal{U}))(\mathbf{y}, \mathbf{n}) = \sum_{\mathbf{y}' \in N(\mathbf{y})} \sum_{\mathbf{n}' \in S} q_{\mathbf{y}, \mathbf{n}}(\mathbf{y}', \mathbf{n}') \mathcal{U}(\mathbf{y}', \mathbf{n}') d\mathbf{y}' d\mathbf{n}' \quad (7)$$

where $d\mathbf{y}'$ is the discrete volume measure and $d\mathbf{n}'$ the discrete surface measure, which in case of (nearly) uniform

sampling of the sphere, such as tessellations of icosahedrons, can reasonably be approximated by $\frac{4\pi}{3}$. For a slightly more accurate approach see [DF09, ch 8.2, eq. 86]. Kernel $q_{\mathbf{y}, \mathbf{n}}$, which is the rotated and translated correlation kernel (such that it is aligned with (\mathbf{y}, \mathbf{n})) associated to p as we will explain later in Section 3.

conveys the diffusion over the coupled space of positions and orientations - very much like a HARDI image. Therefore, one should note the complexity involved in these operations. Consider:

- Q : number of points in kernel's lattice
- S : number of vectors in kernel's tessellation (and SDFs of the input HARDI data)

The discretized convolution expressed in Equation 7, has the complexity of $Q \times S \times Q$ iterations, per voxel of the input HARDI. For instance, consider the convolution with a kernel discretized in a $3 \times 3 \times 3$ lattice, for 2^{nd} order of tessellation (42 directions). The discrete convolution in Equation 7, per voxel in the lattice of the input HARDI image, involves $42 \times 27 \times 42 = 47628$ iterations.

2.3. Tractosemas

In the work of Barmpoutis et al. [ABF08], a field of asymmetric spherical functions, called tractosemas, is extracted from a field of SDFs. The kernel that governs the smoothing process is defined as a function over space and orientation, i.e. over the full domain $\mathbb{R}^3 \times S^2$. The proposed kernel intuitively describes when a structure should be enhanced. It is constructed as a direct product of three parts involving von Mises and Gaussian probability distributions :

$$k(\mathbf{y}, \mathbf{n}; \mathbf{y}', \mathbf{n}') = k_{\text{dist}}(\|\mathbf{y} - \mathbf{y}'\|) \cdot k_{\text{orient}}(\mathbf{n} \cdot \mathbf{n}') \cdot k_{\text{fiber}}\left(\frac{\mathbf{n}}{\|\mathbf{y} - \mathbf{y}'\|} \cdot (-\mathbf{y} - \mathbf{y}')\right), \quad (8)$$

where the different factors are given by

$$k_{\text{dist}}(\|\mathbf{y} - \mathbf{y}'\|) = \frac{1}{(2\pi\sigma)^{\frac{3}{2}}} e^{-\frac{\|\mathbf{y} - \mathbf{y}'\|^2}{2\sigma^2}},$$

$$k_{\text{orient}}(\cos \phi) = k_{\text{fiber}}(\cos \phi) = \frac{\kappa e^{\kappa \cos(\phi)}}{4\pi \sinh(\kappa)},$$

with $\phi \in (-\pi, \pi]$ being the angle, respectively, between the vectors \mathbf{n} and \mathbf{n}' and the angle between the vectors \mathbf{n} and $(\mathbf{y} - \mathbf{y}')$. The two scale parameters σ and κ control kernel's sharpness. Figure 2 shows an example of the tractosemas kernel $p_{(\sigma, \kappa)}: \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$ given by

$$p_{(\sigma, \kappa)}(\mathbf{y}, \mathbf{n}) = \frac{1}{4\pi} k_{\text{dist}}(\|\mathbf{y}\|) k_{\text{orient}}(\mathbf{e}_z \cdot \mathbf{n}) k_{\text{fiber}}(-\|\mathbf{y}\|^{-1} \mathbf{n} \cdot \mathbf{y}). \quad (9)$$

2.4. Diffusion Kernels

Duits [DF09, DF] proposed a kernel based on solving the diffusion equation for HARDI images. The full derivation is out of the scope of this manuscript. This kernel, dubbed

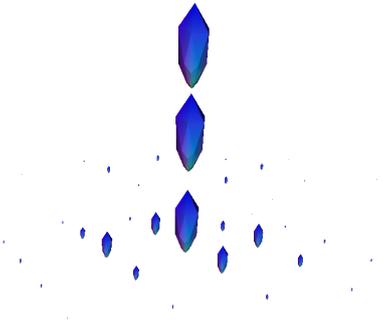


Figure 2: TODO: The tractosemas kernel (9) for $\sigma =$ and $\kappa =$. proposed in [ABF08] - convolve with delta peak, to have it oriented nicely.

from now on as the Brownian motion kernel (on the coupled space $\mathbb{R}^3 \times S^2$ of positions and orientations), satisfies the two important requirements for a diffusion kernel:

1. **left-invariant** The kernel satisfies the right symmetry constraints, [DF09, Corr.1]. Thereby rotation and translation of the input U corresponds to rotation and translation of the output $\Phi(U)$.
2. **fulfill the semigroup property** When the operator is applied iteratively, the scales can be added.
3. **diffusion equation** It closely approximates the Green's function for the diffusion equation on the coupled space $\mathbb{R}^3 \times S^2$ of positions and orientations describing Brownian motion on positions and orientations (where the angular part of a random walk prescribes the tangent vector to the trajectory), [DF09, ch 4.2, Def.5].

In this kernel, the probability function is a product of two 2D kernels on the coupled space $SE(2) \equiv \mathbb{R}^2 \times S^1$ of 2D-positions and orientations:

$$p_t^{D_{33}, D_{44}; \mathbb{R}^3 \times S^2}(x, y, z, \tilde{\mathbf{n}}(\tilde{\beta}, \tilde{\gamma})) \approx N(D_{33}, D_{44}, t) \cdot p_t^{D_{33}, D_{44}; (SE(2))}(z/2, x, \tilde{\beta}) \cdot p_t^{D_{33}, D_{44}; (SE(2))}(z/2, -y, \tilde{\gamma}), \quad (10)$$

where we recall (2) and where $N(D_{33}, D_{44}, t) \approx \frac{8}{\sqrt{2}} \sqrt{\pi t} \sqrt{D_{33}} \sqrt{D_{33} D_{44}}$ takes care of $\mathbb{L}_1(\mathbb{R}^3 \times S^2)$ -normalization. The $SE(2)$ kernel is given by:

$$p_t^{D_{33}, D_{44}; SE(2)}(x, y, \theta) \equiv \frac{1}{32\pi^2 c^4 D_{44} D_{33}} e^{-\frac{\sqrt{EN(x,y,\theta)}}{4c^2 t}} \quad (11)$$

where we use short notation for the numerator of an exponent

$$EN(x, y, \theta) = \left(\frac{\theta^2}{D_{44}} + \frac{\left(\frac{\theta y}{2} + \frac{\theta/2}{\tan(\theta/2)} x \right)^2}{D_{33}} \right)^2 + \frac{1}{D_{44} D_{33}} \left(\frac{-x\theta}{2} + \frac{\theta/2}{\tan(\theta/2)} y \right)^2$$

where one can use the estimate $\frac{\theta/2}{\tan(\theta/2)} \approx \frac{\cos(\theta/2)}{1-(\theta^2/24)}$ for $|\theta| < \frac{\pi}{10}$ to avoid numerical errors.

The diffusion parameters D_{33} and D_{44} and stopping time t allow the adaptation of the kernels to different purposes:

1. $t > 0$ determines the overall size of the kernel, i.e. how relevant is the neighbourhood;
2. $D_{33} > 0$, the diffusion along principal axis, determines how wide is the kernel;
3. $D_{44} > 0$ determines the angular diffusion, so the quotient D_{44}/D_{33} , models the bending of the fibers along which diffusion takes place.

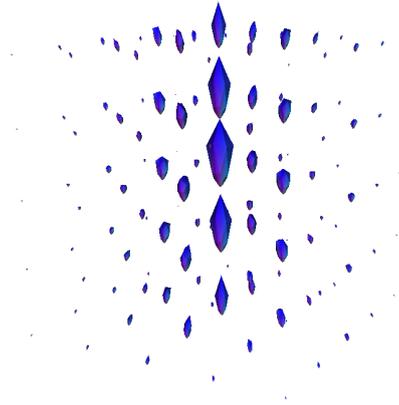


Figure 3: Brownian motion kernel proposed in [DF09].

3. Accelerated Convolution

The probability function given by Equation 8 and Equation 10, should be sampled on the sphere using the same scheme as the HARDI image: tessellation of the icosahedron, i.e. it can be seen as another HARDI image (positions and orientations). The discrete convolution of this complex data with these kernels, therefore involves a high number of operations, $(\mathbb{R}^3 \times S^2) \times (\mathbb{R}^3 \times S^2)$.

For illustration purposes, a convolution with a $3 \times 3 \times 3$ kernel, for second order of tessellation (42 directions), means 47628 iterations per voxel. Applying these operations in a real dataset and for smoother (higher) orders of tessellation (needed to avoid discretization errors) quickly escalates into a time consuming process.

How can this process be accelerated ?

3.1. Preprocessing

Pre-computing. Since the kernels are not adaptive to the data, they do not change depending on each voxel one immediate optimization is to pre-calculate and store the check convolution kernel aligned with every position and orientation.

The check convolution kernel $\check{p} : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$ is basically the correlation kernel related to the convolution kernel $p : \mathbb{R}^3 \times S^2 \rightarrow \mathbb{R}^+$:

$$\begin{aligned} p(\mathbf{y}, \mathbf{n}) &= k(\mathbf{0}, \mathbf{e}_z, \mathbf{y}, \mathbf{n}) \text{ whereas} \\ \check{p}(\mathbf{y}, \mathbf{n}) &= k(\mathbf{y}, \mathbf{n}, \mathbf{0}, \mathbf{e}_z) \end{aligned}$$

where we recall from Eq. (4) that

$$k(\mathbf{y}', \mathbf{n}', \mathbf{y}, \mathbf{n}) = p(R_{\mathbf{n}'}^T(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}}^T \mathbf{n}).$$

To align the correlation kernel with each position \mathbf{x} and orientation \mathbf{n} we define

$$q_{\mathbf{y}, \mathbf{n}}(\mathbf{y}', \mathbf{n}') = \check{p}(R_{\mathbf{n}'}^T(\mathbf{y}' - \mathbf{y}), R_{\mathbf{n}}^T \mathbf{n}'),$$

which we use in our discrete convolution scheme, Eq. (7), where we stress that

$$p(R_{\mathbf{n}'}^{-1}(\mathbf{y} - \mathbf{y}'), R_{\mathbf{n}'}^{-1} \mathbf{n}) = \check{p}(R_{\mathbf{n}}^{-1}(\mathbf{y}' - \mathbf{y}), R_{\mathbf{n}}^{-1} \mathbf{n}').$$

which explains why we must pre-compute the aligned check kernel $q_{\mathbf{y}, \mathbf{n}}$ rather than the original kernel p in Eq. (7).

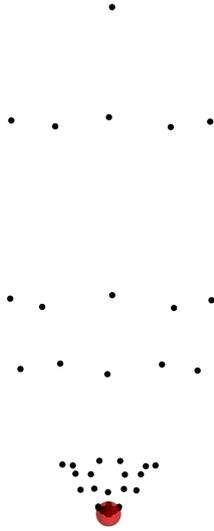


Figure 4: *TODO: figure with 1 voxel of the kernel, with the center, the \mathbf{e}_z direction, and we can see the lots of small values, far from the maxima.*

Thresholding. As we can see in Figure 4, these kernels exhibit an interesting characteristic: the probability of diffusion is larger at the locations around the starting direction \mathbf{e}_z . For all the pre-computed kernels, i.e. one per point in the tessellation, we sort the probability values per kernel voxel, and then we can then define a set $N(\mathbf{y})$, for position \mathbf{y} that contains the orientations with the largest probabilities such that:

$$N(\mathbf{y}) : \{\mathbf{n} \mid k(\mathbf{y}, \mathbf{n}) < p\} \quad (12)$$

where p is a probability threshold. Note that by sorting the values of the kernel we ensure that only the directions corresponding to the 'important' kernel values are iterated, and not any other. These tuples (value, index) are stored into a file for further use. In Figure 5, k_0 and k_1 are two simplified 2D examples of a kernel for the 2 first directions in this tessellation. The two tables the corresponding representation of the sorted and re-indexed kernels.

3.2. Look-up-table (LUT) convolution

Since the kernels are sorted and thresholded, the convolution must now take care to match the correct values per kernel direction to the corresponding HARDI image directions. Figure 5 illustrates this process. For instance, the resulting convolved HARDI image:

- for direction 0, is the outcome of the thresholded kernel k_0 with the matching values from the input HARDI image \mathbf{u} ;
- for direction 1, is the outcome of the thresholded kernel k_1 with the matching values from the input HARDI image \mathbf{u} ;
- ... and hence forth.

After convolving each direction of the input image with the corresponding kernels, we obtain the new HARDI image.

4. Results

In this section we present the experiments conducted in order to analyse the performance of the proposed optimization using a synthetic DW-MRI dataset, fibercup's hardware phantom and a real HARDI data set from a healthy brain.

In all presented experiments, to the (simulated or acquired) signal, QBalls of 4th order Spherical Harmonics (SH) were fit, and the resulting SDF was sampled on a tessellated icosahedron (3rd order, 162 points). From our experiments the choices for SH and tessellation orders are a good balance between the amount of crossing information conveyed (4th order of SH), and a relatively small amount of points on a sphere, 3th order of tessellated icosahedron).

For validation and illustration of the method, we synthesized a dataset with an underlying splaying fibers configuration, whose orientations follow the tangent of two ellipsoids centred in the bottom corners of the image. Using the multi-tensor model as in [DAFD07], we constructed a dataset with size 20×28 , with eigenvalues for each simulated tensor to be $[300, 300, 1700] \times 10^{-6} \text{mm}^2/\text{s}$, b-value of 1000 s/mm^2 and added Rician noise with realistic SNR of 15.3. Figure 6 shows this image (a) and the result of the convolution with the tractosemas kernel (b) ($\sigma = 1, \kappa = 10$ and 3 iterations).

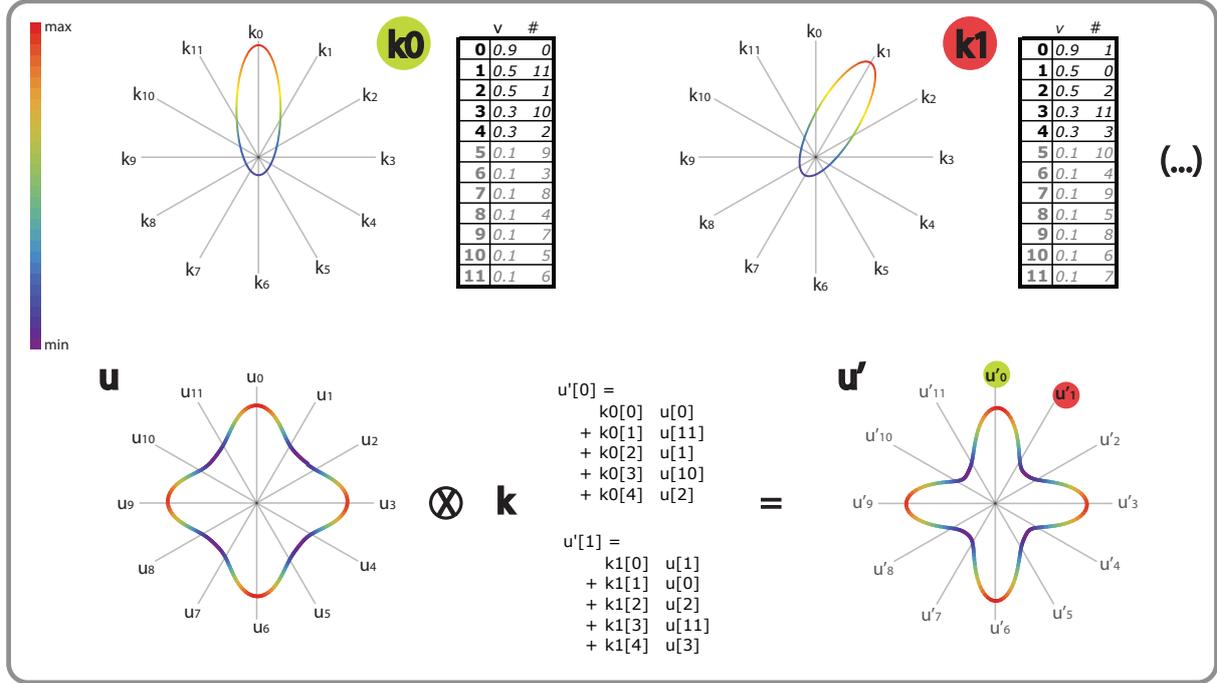


Figure 5: The optimized convolution illustrated. The pre-computed kernels, k_0 and k_1 , are sorted and the pairs value/index are stored. With a threshold $t = 0.1$, only 5 out of 12 directions are used in the convolution. In the LUT convolution, each direction in the resulting image u'_i is equal to the inner product between the corresponding kernel k_i and the matching values in the input image u .

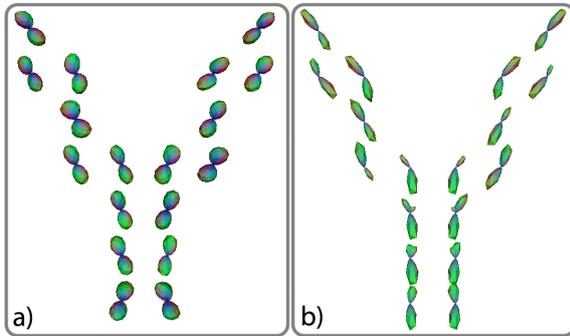


Figure 6: Synthetic splaying fibers example: a) Simulated data; b) The computed convolution with Barmpoutis' tractosemas.

We can observe the resulting asymmetric profile in the center region corresponding to the splaying fiber configuration.

The proposed toolkit was also applied to real DW-MRI datasets. For the next experiments, Duit's kernel was used with diffusion parameters $D_{33} = 0.4, D_{44} = 0.02, t = 1.4$.

From FiberCup's data [PRK*08], with b-value 1500

s/mm^2 and $3 \times 3 \times 3mm$ voxel size, we estimated QBalls as previously described. Figure 7 shows a region of interest (ROI) in the full dataset. Here two fiber bundles cross at approximately 90deg. As we can observe, the QBall model expresses a complex fiber structure in the crossing region, however due to the low b-value, few voxels actually show the 2 expected maxima. Additionally, we can also observe the perturbation caused by noise. After convolving this dataset with Duit's kernel, we obtain a regularized image where the crossing voxels are clearly enhanced, with evident maxima matching the underlying crossing bundles.

Applying the optimized convolution, again with Duit's kernels, to a healthy brain volunteer, acquired with b-value 4000 s/mm^2 , clearly illustrates the benefits of such convolution. Figure 8 shows a ROI where two major white matter structures intersect: the *corpus callosum* from the left, and the *corona radiata* from down-right. We can observe the effect of the low SNR, due to the high b-value, causing clear perturbations in the profiles, specially in the crossing voxels. After convolving this data, we obtain the expected coherency between voxels. Using the neighbourhood information allows the regularization of the data, specially in the more linear voxels, and the enhancement of the crossing voxels.

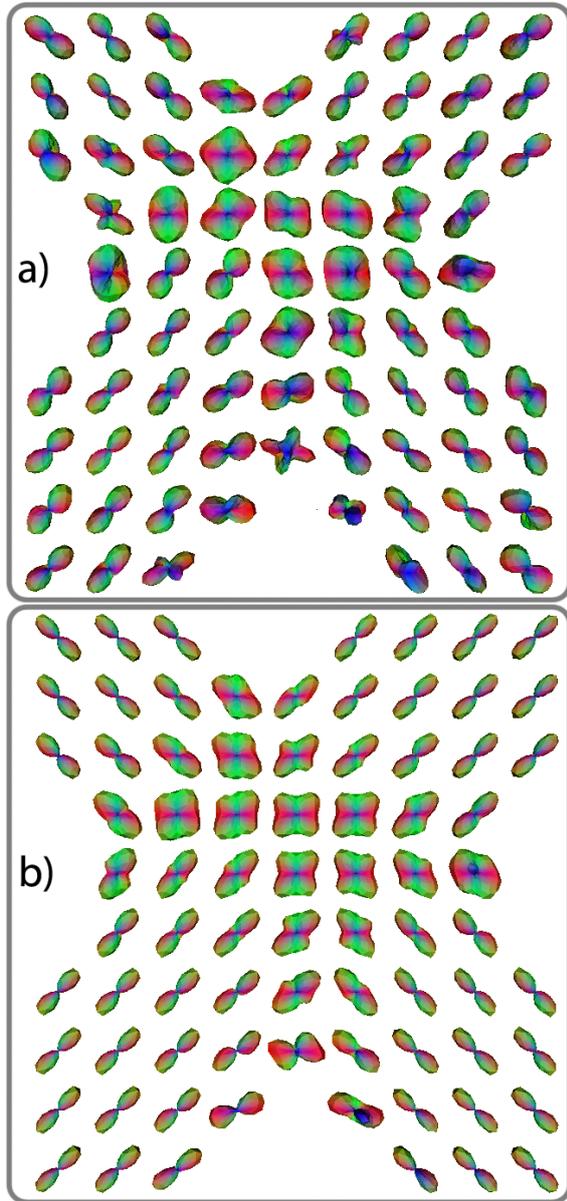


Figure 7: *fibercup, qball, bvalue 1500, min-max normalized*

4.1. Performance

In Figure 9 we present a time comparison between the different convolution methods. We show the time realizations for 4 datasets:

- **Y synthetic** - software simulated dataset [DAFD07], where the tractosemas kernel was applied
- **Fibercup** - Fibercup dataset [PRK*08], with b-value 1500 s/mm^2 and $3 \times 3 \times 3 \text{ mm}$ voxel size

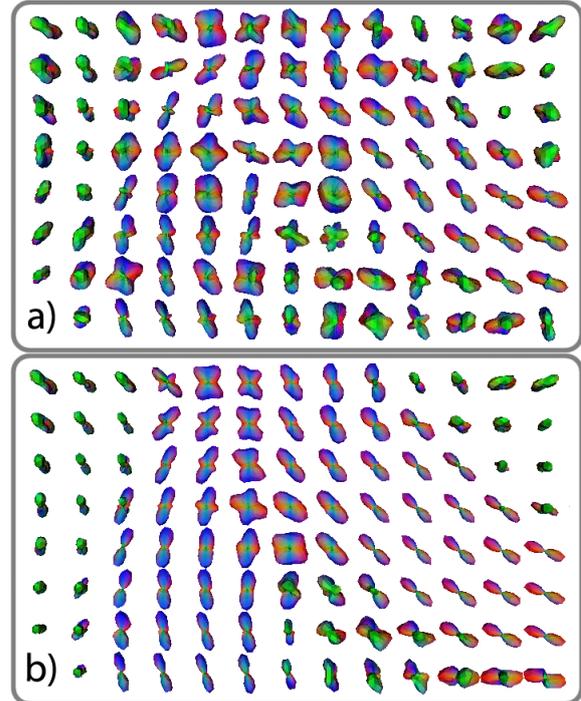


Figure 8: *132, 4000*

- **Brain slate** - one coronal slice from a healthy brain's volunteer, with b-value 4000 s/mm^2
- **Brainvisa's brain** - brain dataset [PPAM06], with b-value 700 s/mm^2

Precomputing the kernels, for 3^{rd} order tessellation, takes 47 seconds. This calculation, of course, is only needed once, per set of parameters.

Applying the proposed optimization, by thresholding the kernels at 90% ($t = 0.003$) of its L1-norm, we obtain identical results as using the full kernel, however 8 times faster. The used threshold was chosen by analysing visually the resulting output that differs minimally from the result using the full kernel. Further work will investigate the influence of the threshold on the resulting smoothed image, but our initial results shows that a great time improvement can be gained with small lost in result's accuracy.

5. Conclusions and Future Work

There are two key limitations in DW-MRI data: spherical distribution functions are symmetric, which does not always express correctly the underlying fiber structure; images can be very noisy, specially at high b-values. Processing of the data on the full domain (spatial and orientational), where contextual information plays an important role, is of utmost importance. However their complexity can be a limiting fac-

Dataset	Method	Time (min)	Ratio
Y synthetic (28x20)	full	2,7	11,7
	lut	0,23	
Fibercup (128 x 128)	full	17,7	8.13
	lut	2,1	
Brain Slate (64x10)	full	4,46	8.12
	lut	0,73	
Brain (60x50x70)	full	974,55	8,07
	lut	120,7	

Figure 9: Table a - Performance comparison between applying the convolution with **full** kernel or with optimized **lut** convolution.

tor in their use. The proposed toolkit allows the addition of these methods to the DW-MRI processing/visualization pipeline, with much improved time costs. The framework's kernel independence enables the use of different kernels, for different purposes (e.g., smoothing, enhancing, completion), but still with optimized costs.

Further work will analyse the optimal balance between optimization (i.e. which threshold to use) and results' accuracy. Further improvements can be achieved by making use of multiple processors (common in nowadays computers) since the processing an image can be atomized to voxel level.

References

- [ABF08] A. BARMPOUTIS B. C. VEMURI D. H., FORDER J. R.: Extracting tractosemas from a displacement probability field for tractography in dw-mri. In *LNCS 5241, (Springer) Proceedings of MICCAI08: Int. Conf. on Medical Image Computing and Computer Assisted Intervention* (6-10 September 2008), 9–16.
- [BML94] BASSER P. J., MATTIELLO J., LEBIHAN D.: MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 66 (1994), 259–267.
- [DAFD07] DESCOTEAUX M., ANGELINO E., FITZGIBBONS S., DERICHE R.: Regularized, fast and robust analytical q-ball imaging. *Magn. Reson. Med.* 58 (2007), 497–510.
- [DF] DUITTS R., FRANKEN E.: Left-invariant diffusions on the space of positions and orientations and their application to crossing-preserving smoothing of HARDI images. *Invited submission to International Journal of Computer Vision* 2010.
- [DF09] DUITTS R., FRANKEN E.: Left-invariant diffusions on $\mathbb{R}^3 \times S^2$ and their application to crossing-preserving smoothing on HARDI-images. *CASA report, TU/e, nr.18* (2009). Available on the web <http://www.win.tue.nl/casa/research/casareports/2009.html>.
- [FB] FLORACK L., BALMASHNOVA E.: Decomposition of high angular resolution diffusion images into a sum of self-similar polynomials on the sphere. In *Proceedings of the Eighteenth International Conference on Computer Graphics and Vision, GraphiCon'2008, Moscow, Russia, June 23–27, 2008*, Bayakovskiy Y., Moiseev E., (Eds.), Moscow State University, pp. 26–31. (invited paper).
- [Flo08] FLORACK L. M. J.: Codomain scale space and regularization for high angular resolution diffusion imaging. In *CVPR Workshop on Tensors in Image Processing and Computer Vision, Anchorage, Alaska, USA, June 24–26, 2008* (2008), Aja Fernández S., de Luis Garcia R., (Eds.), IEEE.
- [Fra02] FRANK L. R.: Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magn. Reson. Med.* 47, 6 (2002), 1083–99.
- [HMH*06] HESS C., MUKHERJEE P., HAN E., XU D., VIGNERON D.: Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn Reson Med* (2006).
- [JV07] JIAN B., VEMURI B. C.: Multi-fiber reconstruction from diffusion MRI using mixture of wisharts and sparse deconvolution. In *IPMI* (2007), pp. 384–395.
- [OM03] ÖZARSLAN E., MARECI T. H.: Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magn. Reson. Med.* 50, 5 (2003), 955–65.
- [OSV*06] ÖZARSLAN E., SHEPHERD T. M., VEMURI B. C., BLACKBAND S. J., MARECI T. H.: Resolution of complex tissue microarchitecture using the diffusion orientation transform (dot). *Neuroimage* 31, 3 (2006), 1086–103.
- [PPAM06] POUPON C., POUPON F., ALLIROL L., MANGIN J.-F.: A database dedicated to anatomo-functional study of human brain connectivity. In *12th HBM Neuroimage* (Florence, Italie, 2006), no. 646.
- [PRK*08] POUPON C., RIEUL B., KEZELE I., PERRIN M., POUPON F., MANGIN J.-F.: New diffusion phantoms dedicated to the study and validation of hardi models. *Magnetic Resonance in Medicine* 60, 6 (2008), 1276–83.
- [TCGC04] 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 (2004), 1176–85.
- [Tuc02] TUCH D. S.: *Diffusion MRI of complex tissue structure*. PhD thesis, Harvard, 2002.
- [Tuc04] TUCH D.: Q-ball imaging. *Magn. Reson. Med.* 52 (2004), 1358–1372.