Fused DTI/HARDI Visualization

Vesna Prčkovska, Tim H.J.M. Peeters, Markus van Almsick, Bart ter Haar Romeny, and Anna Vilanova

Abstract—High-angular resolution diffusion imaging (HARDI) is a diffusion weighted MRI technique that overcomes some of the decisive limitations of its predecessor, diffusion tensor imaging (DTI), in the areas of composite nerve fiber structure. Despite its advantages, HARDI raises several issues: complex modeling of the data, nonintuitive and computationally demanding visualization, inability to interactively explore and transform the data, etc. To overcome these drawbacks, we present a novel, multifield visualization framework that adopts the benefits of both DTI and HARDI. By applying a classification scheme based on HARDI anisotropy measures, the most suitable model per imaging voxel is automatically chosen. This classification allows simplification of the data in areas with single fiber bundle coherence. To accomplish fast and interactive visualization for both HARDI and DTI modalities, we exploit the capabilities of modern GPUs for glyph rendering and adopt DTI fiber tracking in suitable regions. The resulting framework, allows user-friendly data exploration of fused HARDI and DTI data. Many incorporated features such as sharpening, normalization, maxima enhancement and different types of color coding of the HARDI glyphs, simplify the data and enhance its features. We provide a qualitative user evaluation that shows the potentials of our visualization tools in several HARDI applications.

Index Terms—DTI, HARDI, diffusion, GPU, glyphs, multifield.

1 Introduction

Diffusion weighted magnetic resonance imaging (DW-MRI) is a noninvasive tool for measuring the local water diffusion characteristics inside a tissue, and thus, revealing orientation and anisotropy of the underlying tissue structure. One popular application of DW-MRI is the investigation of the white matter in human brain by taking several diffusion-weighted images along different directions. If the underlying structure is coherently organized in just one direction in a single imaging voxel, then the diffusion process can be approximated by a Gaussian probability density function. In diffusion tensor imaging (DTI) introduced by Basser et al. [5], one assumes that this scenario applies in every imaging voxel. However, it has been shown that white matter consists of more complex fiber structures in certain imaging voxels, and in these areas the Gaussian model is inappropriate as shown in the work of Alexander et al. [3]. To properly characterize the diffusion process in these composite structures, where one assumes that two or more fiber bundles interfere (i.e., cross, kiss, diverge), novel approaches based on High-Angular Resolution Diffusion Imaging (HARDI) were pioneered by Tuch [35].

The general idea in HARDI-based techniques is to make as few assumptions as possible concerning the underlying diffusion process, and thus, give more information about the properties of the measured tissue, specifically the orientation of multiple fiber directions. The spherical harmonics (SH) provide a convenient linear basis for representing the angular portion of the reconstructed diffusion functions, and therefore, they have been widely used in the literature for modeling the apparent diffusion coefficient (ADC) profiles by Alexander et al. [3], QBall imaging by Descoteaux et al. [11], diffusion orientation transform (DOT) by Özarslan et al. [24], etc. The benefits of the above discussed HARDI-based modeling techniques make them very attractive for biomedical research, but the drawbacks that come along with their more complex models cause a significant impediment into their employment in a clinical-research setting. There are several main reasons for that such as: overall complex modeling and computationally expensive processing of the data (even in the regions with substantial single fiber coherence where this is redundant), noninteractive visualization, and overwhelmingly complex information for the end user.

To overcome these challenges, we propose a combined DTI/HARDI visualization framework, where we exploit the SH representation of the HARDI modeled diffusion functions to classify the data in three categories: isotropic or the areas from the gray matter and the ventricles, anisotropic-Gaussian (DTI) or the areas with significant single fiber coherence, and nonGaussian (HARDI) or the areas where the white matter bundles have more complex structures. We adopt the fast and well-established visualization techniques from DTI in the anisotropic-Gaussian regions, and the more complex diffusion functions of the SH-based HARDI-modeling techniques in the nonGaussian regions, whereas the isotropic parts can be masked out since they do not contain any significant information. To acquire a fast and accurate visualization, we present a ray-casting algorithm for GPU rendering of DTI ellipsoids, and accelerate the HARDI glyphs as in the work of Peeters et al. [26]. We keep our framework general, such that any HARDI technique represented by SHs can be set as an input for the voxel classification and glyph visualization, and we demonstrate it on SH-based ADC profiles, QBalls, and DOTs. Furthermore, any of the HARDI modeling techniques that have high-order tensor (HOT) representation can be used in our framework as well, due to the direct

• *The authors are with the Biomedical Image Analysis Group, Department of Biomedical Engineering, Technische Universiteit Eindhoven, MB 5600, Eindhoven, The Netherlands. E-mail: vesna.prckovska@gmail.com, thjmpeeters@gmail.com, {B.M.terHaarRomeny, A.Vilanova}@tue.nl, m.van.almsick@arcor.de.*

Manuscript received 7 Sept. 2009; revised 16 Jan. 2010; accepted 27 June 2010; published online 29 Oct. 2010.
Recommended for acceptance by T. Möller.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-2009-09-0208.
Digital Object Identifier no. 10.1109/TVCG.2010.244.

relationship between SH and HOT shown by Özarslan and Mareci [23]. The classification is performed on SH-based anisotropy measures as in the study of Prčkovska et al. [29]. Without loss of generality, any measure (based on SH or HOT coefficients) that gives a good contrast between anisotropic-Gaussian and nonGaussian regions, can be incorporated in the framework as input for our semiautomatic threshold detection of the scalar anisotropy measure.

The main contributions of this work can be summarized as follows:

- DTI and HARDI visualization fused together for the first time in the literature. This representation of the data allows superimposing of the fibers and glyphs for validating fiber tracking algorithms.
- Significant speed up in the visualization, obtained by employing simple DTI diffusion ellipsoids or lines in the areas classified as anisotropic-Gaussian.
- Simple visual representation of the data, achieved by keeping the DT model in significant parts of the white matter where it is justified.
- Feature enhancement by different color codings for the glyphs that enhance the orientational and diffusional properties of the HARDI data, and real time deconvolution of the data that sharpens the profiles in QBall.
- User evaluation for several different HARDI research applications, that gives a first indication of the potentials of the presented framework.

The paper is organized as follows: Section 2 reviews the existing HARDI modeling techniques, explains the mathematical tools for representing HARDI data and gives short overview on the HARDI anisotropy measures for data classification. Section 3 presents the existing visualization techniques and tools for DTI and HARDI data. In Section 4, the proposed visualization framework is presented. The results with performance measurements and user evaluation follow in Section 5. Finally, we conclude the paper and address future work in Section 6.

2 Background

2.1 HARDI—From Measurements to Diffusion Propagator

DW-MRI is designed to measure the diffusion profile of water in tissue via the nuclear spin of hydrogen. During a small time interval (the so-called effective diffusion time t), the diffusion causes a displacement of water molecules. In each location of the tissue, the probability distribution of diffused water molecules with respect to the displacement vector $\mathbf{r}$ is given by the diffusion propagator $P(\mathbf{r})$. Since fiber bundles in tissue allow free diffusion of water along the fibers but tend to obstruct the diffusion perpendicular to the fibers, one assumes that the diffusion propagator $P(\mathbf{r})$ reflects the local fiber structure in a voxel.

In a homogeneous and isotropic medium, the diffusion propagator is according to Fick's second law of diffusion, given by the well-known Gaussian diffusion probability distribution $P(\mathbf{r}) \propto e^{\frac{-\mathbf{r}^2}{tD}t}$, with a scalar diffusion constant D. In a homogeneous medium of limited anisotropy, the diffusion propagator generalizes to a Gaussian diffusion probability distribution with a diffusion tensor $\mathbf{D}$ that reflects the effective diffusion constants in the different directions. However, tissue may be neither homogeneous nor isotropic. Furthermore, the limited resolution of DW-MRI causes partial volume effects, where tissue with more than one directional preference occurs in a single $P(\mathbf{r})$ measurement. Typical examples are voxels with crossing fiber bundles. These voxels do not exhibit Gaussian diffusion profiles as pointed in the work of Alexander et al. [2]. Due to the assumptions above, the Gaussian diffusion profile can only capture adequately diffusion profiles of single fiber bundles. One has to process more complex diffusion profiles to correctly interpret voxels with more complex tissue structure.

In DW-MRI, the diffusion propagator $P(\mathbf{r})$ is not measured directly. The measured signal $S(\mathbf{q})$, obtained from the nuclear spins in the presence of a diffusion sensitizing magnetic gradient is essentially the Fourier transform of the diffusion propagator with the diffusion wave vector $\mathbf{q} = (2\pi)^{-1}\gamma t\mathbf{g}$ consisting of a nuclear-specific constant γ, diffusion time t, and the gradient vector $\mathbf{g}$ of the applied magnetic field. Hence, we need to find the transformation of signal $S(\mathbf{q})$ back to diffusion kernel $P(\mathbf{r})$. The obvious solution, a discrete Fourier transformation of a densely sampled $S(\mathbf{q})$ signal, however, exceeds the time requirements of routine investigations as shown by Tuch [35]. Hence, simpler and less accurate methods are usually applied, that in general describe the signal $S(\mathbf{q})$ on a spherical shell in q-space, called q-shell.

1. *Assumption of a Mono-Exponential Signal Decay and Gaussian PDF.* The simplest approach is DTI. In DTI one assumes a Gaussian diffusion propagator, which results via Fourier transformation in a Gaussian signal $S(\mathbf{q}) = S_0\ e^{-bD(\mathbf{g})}$, where $b = 4\pi^2\mathbf{q}^2t$ is the acquisition parameter, t is the effective diffusion time, and S_0 the unweighted signal intensity. The apparent diffusivity $D(\mathbf{g})$ is modeled by a second order tensor $\mathbf{D}$, such that with gradient $\mathbf{g}$one assumes $D(\mathbf{g}) = \mathbf{g}^T \cdot \mathbf{D} \cdot \mathbf{g}$ requires very modest gradient $\mathbf{g}$ samplings on a q-shell [35] to determine its coefficients. However, as mentioned above, this approach is limited as it can resolve only one fiber bundle per voxel.
2. *Assumption of a Mono-Exponential Signal Decay* $S(\mathbf{q}) = S_0e^{-bD(\mathbf{g})}$. In HARDI, to increase the angular resolution, the apparent diffusivity $D(\mathbf{g})$ is modeled in a less crude way by SH or HOT [14], [3], [23]. This enhances the characterization of tissue properties, but also increases the number of necessary $S(\mathbf{q})$ samples, and gives incorrect results regarding the directions of the fiber crossings as shown by Özarslan et al. [24]. DOT maps the ADC into $P(\mathbf{r})$ by solving the Fourier transform using the Rayleigh expansion of a plane wave in spherical coordinates. It is worth mentioning that DOT allows multi-exponential assumption for the signal decay which allows the exploration of multiple q-shell samplings.
3. *No Assumption about Signal Decay nor the Underlying Diffusion Process.* QBall imaging [36] simply maps the measurements on a q-sphere via the Funk-Radon transform into another function on a sphere called

Orientation Distribution Function $\psi(\mathbf{u})$ (ODF). This ODF $\psi(\mathbf{u})$ is the integral in radial direction of the true diffusion propagator $P(\mathbf{r})$

$$\psi(\mathbf{u}) = \int_0^\infty P(r\mathbf{u})dr, \tag{1}$$

where $\mathbf{u} = \frac{\mathbf{r}}{|\mathbf{r}|}$. One of the biggest advantages of QBall is its analytical solution of the Funk-Radon transform [11], [17], [4] utilizing a basis of SH that significantly improves its performance.

4. Lately, *spherical deconvolution (SD) methods* proposed in the literature [34], [12] are becoming popular, especially for fiber tracking purposes due to sharper profiles of the reconstructed Fiber Orientation Distribution (FOD) F. One assumes, that the DWI signal $S(\mathbf{q})$ can be modeled by a superposition of fiber bundle signals R via the convolution $S = F*R$. The FOD that one obtains by a SD of S with R reflects the angular distribution of fiber bundles. One has to keep in mind that this approach incorporates assumptions about the single fiber response R. Recently, it has been shown that some of the bias in the fiber extraction from ODF and FOD (especially under small angles of crossings) can be improved using a tensor decomposition approach as in the work of Schultz and Seidel [31].

There are many other HARDI modeling techniques in the literature, but these exceed the scope of this paper. It is important to note, that SH provide a convenient basis for representing functions on a sphere, and that all of the diffusion weighted MRI modeling techniques render functions that reside on a sphere. Their end product can have different physical meaning (a PDF, isosurface of a PDF, ODF, FOD, etc.) but in all cases they characterize the diffusion process and/or the underlying fiber distribution. For simplicity and clarity, we will refer to all these functions as spherical distribution functions (SDF) and represent them in a SH basis in the rest of the paper.

Here, we have given a short survey on HARDI and DTI as its subset, deriving their relations and the meaning of the reconstructed diffusion patterns.

2.2 Spherical Harmonics

From the previous section, an important conclusion arises: the SDFs of any HARDI modeling technique, that characterize the local intravoxel fiber structure, can be represented as a linear combination of SH called Laplace series

$$\psi(\theta,\phi) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} a_l^m\, Y_l^m(\theta,\phi), \tag{2}$$

where a_l^m are spherical harmonic coefficients determined via a spherical harmonic transformation and Y_l^m represent the SH. The SH Y_l^m are the angular portion of the solution of the Laplace equation in spherical coordinates and are explicitly defined via associated Legendre polynomials P_l^m as

$$Y_l^m(\theta,\phi) \;=\; e^{i\,m\,\phi}\,\sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}}\,P_l^m(\cos\theta) \tag{3}$$

with $\theta \in [0,\pi], \phi \in [0,2\pi]$ and integer indices $l \geq 0$, and $|m| \leq l$.

The SHs as defined in (3) are complex valued due to the factor $e^{i\,m\,\phi}$. Complex arithmetic is not established on GPUs, so we have to rely on real-valued SHs. We adhere to the convention as in Descoteaux et al. [11] and define the real-valued SHs $\tilde{Y}_l^m$ as

$$\tilde{Y}_l^m(\theta,\phi) = \begin{cases} \sqrt{2}\,\mathrm{Re}\big(Y_l^m(\theta,\phi)\big) & \text{if} \quad m<0, \\ Y_l^0(\theta,\phi) & \text{if} \quad m=0, \\ \sqrt{2}\,\mathrm{Im}\big(Y_l^m(\theta,\phi)\big) & \text{if} \quad m>0. \end{cases} \tag{4}$$

$\mathrm{Re}(Y_l^m(\theta,\phi))$ and $\mathrm{Im}(Y_l^m(\theta,\phi))$ represent the real and imaginary part of the $Y_l^m(\theta,\phi)$.

Employing the basis functions of (4), simplifies the computation for preprocessing and rendering of the HARDI data significantly. Furthermore, note that HARDI signals are real-valued and symmetric under inversion. It is therefore, sufficient to expand a HARDI signal $\psi(\theta,\phi)$ by real-valued coefficients $\tilde{a}_l^m$ and SHs $\tilde{Y}_l^m(\theta,\phi)$ with even $l = 0, 2, 4, \ldots l_{\max}$,

$$\psi(\theta,\phi) \approx \sum_{l=0,2,\ldots}^{l_{\max}}\sum_{m=-l}^{l} \tilde{a}_l^m\;\tilde{Y}_l^m(\theta,\phi). \tag{5}$$

Since the angular sampling of the HARDI signal is finite, and due to the fact that noise becomes more dominant in higher frequencies without significant loss of the information the Laplace series can be truncated at a given $l_{\max}$. On the other hand, the higher the order $l_{\max}$, the higher the angular frequencies that are captured, thus allow more complicated and better resolved shapes for the reconstructed function $\psi(\theta,\phi)$. Hess et al. [17] has shown that currently order $l_{\max} = 4$ results in good signal-to-noise ratio (SNR) performance and least angular error for the reconstruction of the correct orientation distribution function for HARDI data acquired in a clinical setup with b-values up to $4{,}000\ \mathrm{s/mm^2}$.

2.3 HARDI Anisotropy Measures for Voxel Classification

The study of Prčkovska et al. [29] demonstrates the ability of several HARDI anisotropy measures to successfully classify HARDI data into three categories: isotropic, anisotropic-Gaussian, and nonGaussian. Most of the measures like: generalized anisotropy (GA) [25], generalized fractional anisotropy (GFA) [36], fractional multifiber index (FMI) [14], and R_0, R_2, R_i [10] are defined directly on the SH coefficients a_l^m, and give scalar output. This study only focuses on the investigation of the classification power on different HARDI anisotropy measures, and their differences. We apply these measures on any SDF from the HARDI modeled data, and thus, classify the data. There are a few important points that should be stressed from the work of Prčkovska et al. [29], when choosing among the anisotropy measures to fuse the data in our framework. The measures can significantly discriminate anisotropic-Gaussian from nonGaussian profiles even though the SDF themselves do not always discriminate the distinct fiber categories. For typical clinical applications at lower acquisition parameters (of about 60 gradient directions and $b = 1{,}000\ \mathrm{s/mm^2}$), GA and GFA applied on ADC and QBall are better classifiers of the HARDI data than SD-based methods [29].

3 Related Work

One of the most common and intuitive ways for visualizing the global structure of brain white matter is by showing reconstructed fibers with simple line rendering. Tractography is needed to understand functional coupling between cortical regions of the brain and it is important for the characterization of neurodegenerative diseases for surgical planning and for many other medical applications as pointed out by Mori and Zijl [22]. There are well-established methods for reconstructing fibers in DTI using deterministic streamline fiber tracking [38] and probabilistic methods [19]. Recently, one can observe an increase in research activity regarding deterministic and probabilistic HARDI fiber tracking algorithms [9], [28], [19]. All fiber tracking techniques exhibit several disadvantages, such as choice of initialization, sensitivity with respect to the estimated principal direction, lack of connectivity information between regions of the brain (deterministic), high computational expense (probabilistic), and most importantly the lack of thorough validation which is a precondition before fiber tracking can be used in clinical research. Therefore, it is important to obtain additional information about the local structure of the tracts in order to understand and validate the fiber tracking methods. This can be done by visualizing the local diffusion information by glyphs. Chen et al. [7] proposed a method where ellipsoid glyphs are shown on tracked fibers in order to provide local tensor information along the fibers, which show global connectivity. However, their method is limited to DTI, and thus, cannot show HARDI data.

As mentioned in the previous section, the diffusion PDF in DTI is approximated by a second order, symmetric, positive-definite tensor, and therefore, rank 2 tensor glyphs are the most useful tool for representing the full tensor information. Commonly used shapes for visualizing the DTs are cuboids, ellipsoids, and superquadrics proposed by Kindlmann [21].

HARDI techniques overcome the disadvantages of DTI by applying high-order models in the description of water molecule diffusion within a voxel. However, these methods raise the complexity of data processing and visualization significantly. Among the available software packages for visualizing and processing HARDI data are Slicer [1] with the QBall plug-in [39], and Camino [8]. The latter one has very limited visualization capabilities since its core purpose is to process HARDI data. Recently, Shattuck et al. [32] developed a set of tools for visualizing ODF models.

All of the mentioned packages apply the same concept. To obtain a geometrical SDF shape, they generate a mesh of points on a sphere and displace the vertices according to the SDF. This approach renders an interactive scene only after the geometry is calculated. Furthermore, there is always a trade-off between the visual quality and rendering speed. The interaction is reported as 10 frames/s on a single brain data slice with 225 samples on a sphere that approximates 4th order icosahedral tessellation.

Recent advances in capability and performance of GPUs have triggered several publications about the rendering of glyphs where ray casting on the GPU has been applied [16], [20], [33]. In the work of Peeters et al. [26] for the first time the capabilities of modern GPUs are exploited for fast and interactive HARDI glyph rendering. We adopt this scheme to improve the HARDI glyph rendering in the areas classified as nonGaussian.

Regarding the fiber tracking algorithms in HARDI, work by Hlawitschka et al. [18] addresses the problem of improving their speed and robustness. However, no performance information is given and the underlying field of HOT or SH representation still needs sampling on a predefined grid, which gives rise to the same disadvantages discussed above.

All of the above mentioned work uses either DTI or HARDI modeling techniques and shows only fibers or glyphs. In the first case, the visualization is fast, but in large part of the white matter unreliable. In the second case, the local structure is more accurately described by using higher order models. However, these models, due to their complexity, present overwhelming information that often do not reveal the global context. In our framework, we fuse the fiber tracts and GPU accelerated DT ellipsoids with ray casted HARDI SDFs by classifying the data using different HARDI anisotropy measures. In the areas classified as linear, we apply deterministic fiber tracking or tensor glyphs, and in regions where Gaussian profiles depart, we show the local HARDI diffusion information. In this way, we build a fast and reliable framework that allows truly interactive exploration of the DW-MRI data. Global context information is enabled by the possibility of showing the more reliable fiber tracts. However, the unreliable areas that cannot be correctly captured by the fast but limited DT model are represented by the local HARDI glyph information. This is an advantage over the current DTI (tensor glyphs or fibers) or HARDI (SDF glyphs) based visualization, due to increase of reliable information in the first case and context, and speed in the latter.

4 Combined Visualization for DTI and HARDI

In this section, we describe the components of our processing and rendering pipeline (see Fig. 2). Section 4.1 describes GPU-based ray casting for tensor glyphs. In Section 4.2, we summarize the more complicated ray casting of SH as previously published in Peeters et al. [26]. In order to improve visual perception of these glyphs, we present new feature enhancements of the glyphs in Section 4.3. In Section 4.4, we take the HARDI measures from Prčkovska et al. [29] and introduce a new semiautomatic algorithm for finding optimal thresholds for these measures. One has to keep in mind that each section below can be an individual pipeline component, but the strength of our visualization comes to full extent when the components are combined, as shown in Section 4.5.

4.1 Ray Casting Tensor Glyphs

For rendering tensor glyphs, we use GPU-based ray casting of ellipsoids. In order to do this, we render a bounding box for each glyph. These bounding boxes are simple cubes with its edges aligned with the x, y, z-axes in world space. The length of each of the edges is the maximum diameter of the glyphs in any direction. For each bounding box that we render, we pass the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ and corresponding normalized eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ computed from the diffusion tensors to the GPU as vertex attributes.

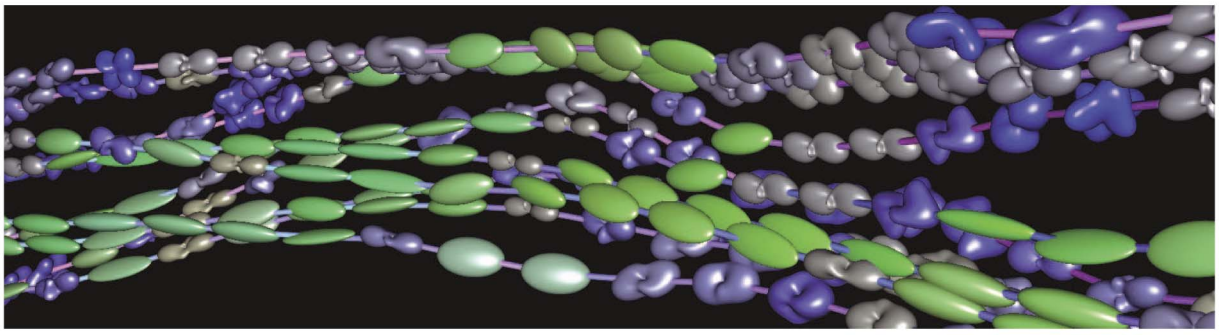

Fig. 1. Fiber tracts in the *cingulum* of a human brain with DTI (green ellipsoids) and HARDI (QBalls, blue and gray with complex shapes) glyphs representing the local diffusion.

Then, on the GPU, we use ray casting to determine the intersection of the view ray for each pixel with the ellipsoidal glyph that is being rendered.

The ellipsoidal glyphs $\mathcal{G}(Q)$ that we want to render are aligned with the eigenvectors of the tensor at a position Q. In order to solve the intersection of the view ray with $\mathcal{G}(Q)$, we transform the view ray into the local eigenvector coordinate system by multiplying the view direction and position (relative to Q) with rotation matrix $R = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)^{\mathrm{T}}$ to obtain a transformed view ray $\mathcal{V} = V + \mu\mathbf{v}$. The intersection of the view ray with the glyph can then be computed analytically by filling $V + \mu\mathbf{v}$ for $(x, y, z)^{\mathrm{T}}$ in the equation of the axis-aligned ellipsoid $\frac{x^2}{r_x^2} + \frac{y^2}{r_y^2} + \frac{z^2}{r_z^2} = 1$, where $(r_x, r_y, r_z) = (\lambda_1, \lambda_2, \lambda_3)$ are the radii of the ellipsoid. The resulting equation can be solved using the quadratic formula

$$\mu = \frac{-b \pm \sqrt{d}}{2a}, \text{ where } d = b^2 - 4ac \tag{6}$$

with $a = \sum_i r_i \mathbf{v}_i^2, b = \sum_i 2r_i(V_i - Q_i)\mathbf{v}_i$ and $c = \sum_i r_i(V_i - Q_i)^2$ for $i \in \{x, y, z\}$. For $d < 0$ there is no intersection. For $d \geq 0$, the intersection closest to V has $\mu = (-b - \sqrt{d})/2a$. Because we update the Z-buffer with the proper intersection depth, our glyph visualization integrates well with other objects in the rendered scene.

To compute the normal $\mathbf{n}$ of the ellipsoid, we compute the vector pointing from glyph center Q to intersection point I and fill it in the derivative $\nabla(\mathbf{x}^{\mathrm{T}}\mathbf{D}\mathbf{x}) = 2\mathbf{D}\mathbf{x}$, where $\mathbf{D}$ is the diffusion tensor and $\mathbf{x}$ a vector. Thus, the normal $\mathbf{n}$ is given by the gradient

$$\mathbf{n} = 2\mathbf{D}(I - Q) = 2\mathbf{R}\mathbf{M}\mathbf{R}^{\mathrm{T}}(I - Q), \tag{7}$$

where $\mathbf{M}$ is the diagonal matrix of the eigenvalues. Vector $\mathbf{n}$ can be used in standard lighting calculations. In general, the color of each individual ellipsoid is given by a direct mapping of the x, y, z components of the main eigenvector $\mathbf{e}_1$ to RGB values, but other coloring measures are also possible. In Fig. 1, the green glyphs are tensor glyphs aligned with anteroposterior direction.

4.2 Ray Casting SH Glyphs

Ray casting of HARDI glyphs is more challenging than ray casting of simple ellipsoids because there is no analytical solution for the intersection of a ray with the SDF represented by SH. The steps for ray casting of SHs are practically the same as for ray casting of ellipsoids, but the computation of the bounding box, the view ray-glyph intersection, and the normal of the glyph in the intersection are more complicated. We adhere to the work of Peeters et al. [26], and we summarize the different steps.

The computation of the bounding box for the SH glyphs is more complicated than for the ellipsoid glyphs because no maximum radius can be computed analytically. Therefore, we can use either a precomputed numerical estimation of the maximum radius, or compute an upper bound for the maximum radius [26]

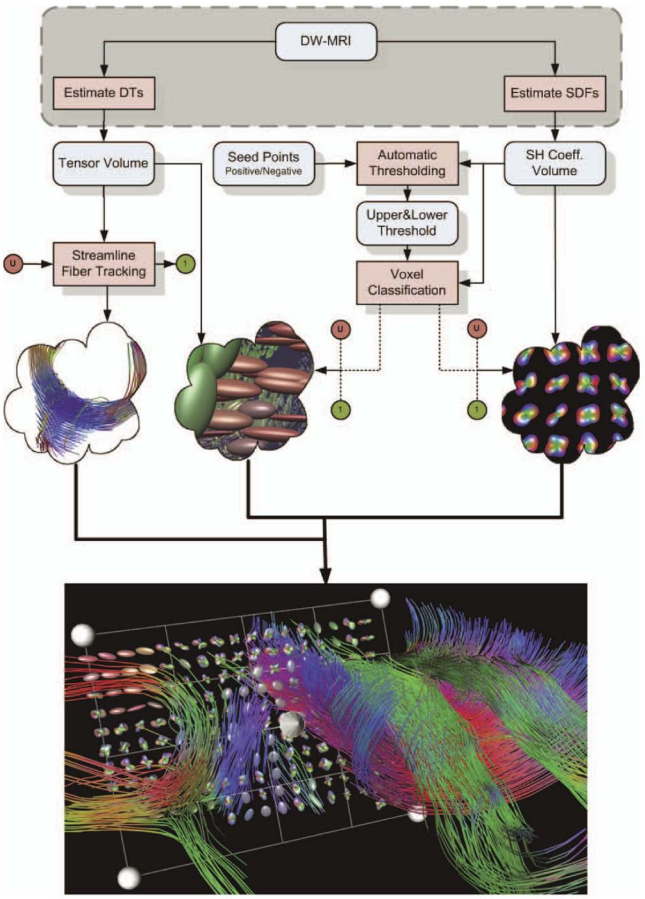

Fig. 2. Schematic diagram of the visualization pipeline. The pink rectangles represent transformations on the data and the blue rounded rectangles different data. The small red circles represent user interaction, and the small green circles the output of fiber tracking as seed points. The flow of the pipeline is captured by directed arrows where dashed lines represent optional scenarios.

$$|\psi(\theta,\phi)| \leq \sqrt{\frac{1}{2}l_{\max}+1}\sqrt{\sum_{l=0}^{l_{\max}}\left(\frac{2l+1}{4\pi}\sum_{m=-l}^{l}\left(\tilde{a}_l^m\right)^2\right)} = r_{\max}. \tag{8}$$

We use r_{max} to determine the dimensions of the bounding box on the CPU and to quickly discard pixels in the fragment shader on the GPU where the view ray does not intersect the sphere with the same origin as the current SH glyph and radius r_{max}.

In order to compute the intersection of view ray $\mathcal{V}$ with an SH glyph $\mathcal{G}(P)$, we perform linear search along the view ray to find a point on $\mathcal{V}$ inside $\mathcal{G}(P)$. Then, we further refine the exact location of the intersection using binary search. In these algorithms, the challenging part is to determine whether a point Q on $\mathcal{V}$ is inside or outside $\mathcal{G}$. For this, we implemented the computation of real-valued SHs as in (4) on the GPU. The input spherical angles of $\psi(\theta,\phi)$ are derived from the direction $(Q-P)$. We determine whether Q is inside or outside $\mathcal{G}$ from the sign of $\Phi(r,\theta,\phi) = r - \psi(\theta,\phi)$, where $r = \|(Q-P)\|$.

The glyph is an isosurface of $\Phi(r,\theta,\phi)$. Hence, the gradient of $\Phi(r,\theta,\phi)$ is a normal vector on the glyph. The gradient below is given in spherical coordinates, but the vector tuple is Cartesian [26]

$$\begin{aligned}\nabla\,\Phi(r,\theta,\phi)\Big|_{r=\psi(\theta,\phi)} &= \begin{pmatrix}\cos(\phi)\sin(\theta)\\ \sin(\theta)\sin(\phi)\\ \cos(\theta)\end{pmatrix} - \frac{1}{\psi(\theta,\phi)}\sum_{l=0}^{l_{\max}}\sum_{m=-l}^{l}\tilde{a}_l^m \\ &\quad \begin{pmatrix}\cos(\theta)\cos(\phi)\,\partial_\theta - \csc(\theta)\sin(\phi)\,\partial_\phi\\ \cos(\phi)\csc(\theta)\,\partial_\phi + \cos(\theta)\sin(\phi)\,\partial_\theta\\ -\sin(\theta)\,\partial_\theta\end{pmatrix}\tilde{Y}_l^m(\theta,\phi).\end{aligned}$$

After normalization, the gradients above are unit normal of the glyphs defined by the spherical function $\psi(\theta,\phi)$.

4.3 Feature Enhancements for SH Glyphs

For better visual perception of the spherical maxima of the SDFs (i.e., make the "peaks" of the glyphs more obvious), especially if the SDFs we want to visualize are relatively isotropic ODFs like in QBall [11], min-max normalization is most commonly used. We propose min-max normalization on the SH coefficient such that

$$\psi_{norm}(\theta,\phi) = \frac{\psi(\theta,\phi)-min}{max-min} = \sum_{l=0}^{l_{max}}\sum_{m=-l}^{l}\hat{a}_l^m\tilde{Y}_l^m(\theta,\phi), \tag{10}$$

where

$$\hat{a}_{lm} = \begin{cases}(\tilde{a}_{lm} - 2\sqrt{\pi}\,min)/(max-min) & \text{if} \quad l=m=0\\ \tilde{a}_{lm}/(max-min) & \text{if} \quad l,m\neq 0.\end{cases} \tag{11}$$

The minimum min and maximum max of each SDF per imaging voxel are precalculated from a discrete surface mesh and passed to the GPU as vertex attributes because there are no known analytical methods for finding maxima in Laplace series of a spherical function.

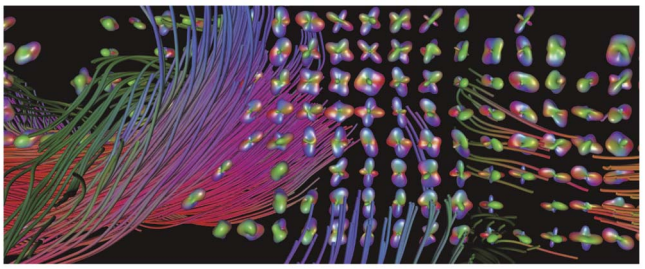

Fig. 3. Combined fiber/HARDI-glyph visualization of the *centrum semiovale* in a human brain where there is a known crossing between the *corpus callosum* and the *corona radiata*. The DOT glyphs and fibers are RGB color coded depending on their orientation.

QBall is calculated directly from the DW-MRI signal via Funk-Radon transform, as described in Section 2. The glyph profiles are therefore, relatively isotropic. An additional deconvolution sharpens these profiles as proposed by Descoteaux and Deriche [12] improves the angular resolution. We implemented sharpened fiber-ODF where the user can select areas with significant single fiber coherence, as a de-convolution kernel. The user can observe in real time the effect of the deconvolution and change the kernel. Care has to be taken here, if a constrained super-resolution regularization is not applied [34], since the deconvolution can give noisy spurious peaks. This implementation demonstrates the benefits of GPU accelerated visualization. A researcher in HARDI modeling can explore the effects of new proposed SD-based algorithms in real time. The same applies to tensor decomposition [31] techniques. The parameters can be tuned and results are immediately visible. This feature is of practical importance in this field, especially because researchers lack interactivity in the current visualization techniques.

In DW-MRI, a common standard for encoding directions is via color mapping. We, therefore, apply RGB coloring for the SDF glyphs as proposed by Tuch [35] with red indicating mediolateral, green anteroposterior, and blue superior-inferior (see Fig. 3). However, the RGB coloring enhances axes-aligned orientation, which is an arbitrary choice and diminishes the perception of maxima, especially if they are not pronounced. We propose color coding that enhances maxima (see Fig. 6), where the maxima peaks are color coded by yellow and the rest of the glyph interpolated with red. The user can interactively change the radius threshold for the yellow/red coloring.

Additionally, we can color code the whole surface of a glyph, by any of the implemented HARDI anisotropy measures, thus giving more information about the amount and type of anisotropy (see Fig. 5). This is very useful for QBalls, since min-max normalization is the most common way of enhancing the maxima due to their relatively isotropic profiles. However, this might cause considerable distortions of the glyphs coming from noise or ventricles in the brain. Adding the information about the type of anisotropy in the color coding, gives additional value in the interpretation of the glyph information.

4.4 Classification

We perform classification based on the HARDI anisotropy measures as described by Prčkovska et al. [29], where an

upper and lower threshold for crossing areas is given for HARDI anisotropy measures such as GA and GFA, which effectively classifies each voxel in one of three compartments: isotropic, nonGaussian, and anisotropic-Gaussian. One of the drawbacks of this kind of classification is that the user has to guess the "best" threshold interval that gives good contrast between the compartments. This guess is often a trial and error process, and highly dependent on the angular configuration of the nonGaussian profiles and the acquisition parameter b-value. Frequently, the order of magnitude of these measures changes, and that makes the process even more complicated. To facilitate this, we offer an algorithm for semiautomatic threshold tuning.

The user chooses a HARDI anisotropy measure $\mathbf{M}$ for classification and regions of positive (i.e., the areas where crossings are expected) Pos, and negative (i.e., linear areas) seed points Neg for the classification. Our algorithm then determines the distributions $M(Pos)$ and $M(Neg)$ of $\mathbf{M}$ in the voxels of Pos and Neg. The lower threshold is set to the minimum value of $M(Pos)$ (i.e., between isotropic and crossing). If $M(Pos)$ and $M(Neg)$ have no overlap, then the upper threshold is set to the maximum value of $M(Pos)$. If $M(Pos)$ and $M(Neg)$ (partially) overlap, then the upper threshold is set to the average of the medians of $M(Pos)$ and $M(Neg)$. In the last step, we chose the medians of $M(Pos)$ and $M(Neg)$ and not the averages because the average is prone to outliers. Then the user can give additional feedback, by tuning the selection regions, or manually refining the suggested threshold. This has proven very useful for many applications as discussed in Section 5.3.

4.5 Fused Visualization

Fig. 2 illustrates the pipeline of our fused visualization. Initially, an input of tensors and SH coefficients that represent different SDFs needs to be computed. This can be done from the same DW-MRI data, since DTI is only a subset of HARDI measurements as explained in Section 2. The pink rectangles represent transformations on the data and the blue rounded-rectangles different data. The small red circles represent user interaction (by selection seeding regions for example), and the small green circles the output of fiber tracking as seed points. The flow of the pipeline is captured by directed arrows where dashed lines represent optional scenarios. For example, for rendering the HARDI glyphs, a volume of SH coefficients are mandatory, as well as a seeding region provided by voxel classification, or the fiber tracking output, or selection by the user. On the data classified as anisotropic-Gaussian (linear), a fiber tracking algorithm can be applied. Then, this data can be visualized with simple line rendering, improved by hardware shading as in the work of Peeters et al. [27] or represented by tensor glyphs as in Section 4.1. On the other hand, the data classified as nonGaussian (e.g., crossings) can be represented by SDFs calculated from 4th order SH explained in more details in Section 4.2. However, showing only fibers or only glyphs can result in unreliable visualization in the first case or complex, and overwhelming in the second case. Even though, the user can define seeding regions and explore only parts of the data, still the complexity of the information that glyphs deliver, disables the intuition for global information. Glyphs convey only local information, and therefore, represent the data partially by fibers helps in giving global context. For that purpose, we propose to fuse the DTI and HARDI visualization techniques in a unified framework by using the voxel classification for model selection. By selecting positive and negative examples from the data (i.e., if we want to discriminate crossings from single fiber, the regions where we expect crossings are positive examples, and single fiber negative), the thresholding algorithm detects the most suitable thresholding interval that divides the data in three categories as described in Section 4.4. Afterward, the data are shown by combining tensor glyphs and HARDI glyphs (see Fig. 5). There is also a possibility to combine streamline fibers and HARDI glyphs (see Figs. 3 and 6), but care should be taken here, since the streamline fiber tracking uses as input only tensor data. Additionally, the output from fiber tracking can be an input for glyph visualization, thus showing the detailed HARDI and DTI data on the positions of the globally tracked DTI fibers (see Fig. 1). The interaction comes handy in exploring the border where the anisotropic-Gaussian profiles change into nonGaussian. By fine tuning the classification thresholds the user can see tensor glyphs disappearing and HARDI profiles appearing. This is important in investigation of structures in the brain white matter and evaluation of the amount of voxels where the DTI model is valid.

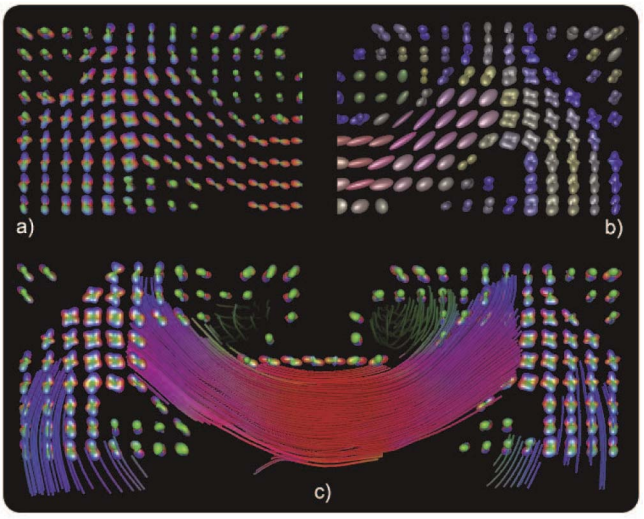

Fig. 4. Comparison of geometry-based HARDI glyph visualization with new proposed fused DTI/HATDI visualizations. The centrum semiovale is again used to show the differentiation between different types of anisotropy; a) Geometry-based QBall glyphs RGB color coded depending on their orientation; b) Fused tensor—QBall glyph visualization. The tensor glyphs can be clearly observed in the corpus callosum, and are color coded by their orientation. The QBalls appear, where corona radiata merges with corpus callosum. The color coding of the QBalls is according to GA measure for anisotropy, where yellow states high and dark blue low anisotropy; c) Fibers—QBall glyph visualization, where both fibers and QBalls are RGB color coded depending on their orientation.

Furthermore, with fibers as seeding points for the glyph visualization, the classification comes in handy in detecting uncertainties in the unreliable outcome from the fiber tracking algori), we show the traditional method of geometry glyph-based visualization. In Fig. 4b, we present the combined tensor-HARDI glyph visualization. The ellipsoids convey the information about the direction better compared to HARDI linear glyphs. This can be observed in

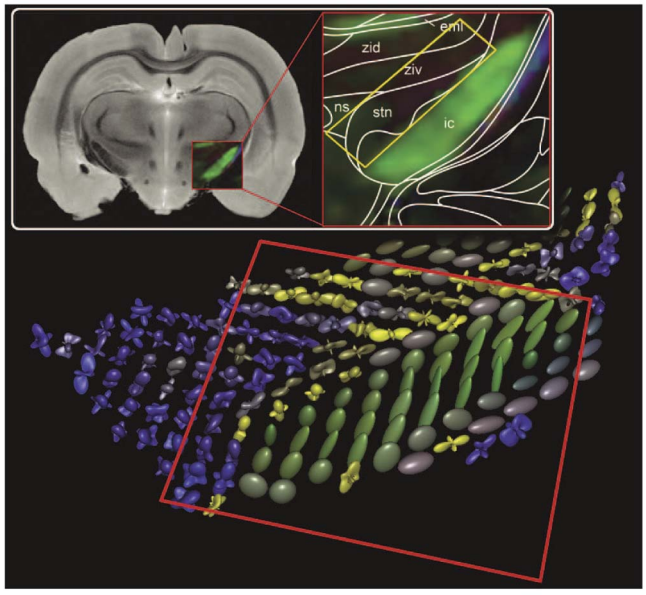

Fig. 5. Combined tensor-fODF glyph visualization of the area around the STN. The fODFs are color coded by GA measure (blue stands for low and yellow for high values). The *internal capsule* can be clearly distinguished with the tensor glyphs, and different GA categories of glyphs observed in the STN.

the area of *corpus callosum* classified as linear. The HARDI glyphs are color coded by the type of anisotropy, where yellow states high anisotropy, and dark blue low anisotropy. All the glyphs colored by yellow have high anisotropy and depending on the threshold tuning can change into linear tensor glyphs. However, it is interesting to observe the glyphs that are colored by dark blue. Even though, due to the min-max normalization, they appear as crossing, the color coding clearly shows that they belong to the gray matter or the ventricles, since crossings should not be observed above and below the corpus callosum. These glyphs can now be masked with our classification, thus removing any redundant and often misleading information. In Fig. 4c, the tensor information is furthermore simplified by showing only fibers, and the rest of the area classified as nonGaussian represented by HARDI glyphs. This kind of information is much more intuitive to grasp than only glyph visualization, especially in the case of Fig. 4a.

Furthermore, with fibers as seeding points for the glyph visualization, the classification comes in handy in detecting uncertainties in the unreliable outcome from the fiber tracking algorithms (that can be DTI deterministic as in our implementation, or generalized to any fiber tracking algorithm based on DTI or HARDI data). This allows a unique interactive way of exploring and validating the outcome of the fiber tracking algorithms.

At the end, we present the benefits of our proposed visualization in Fig. 4. In Fig. 4a we show the traditional method of geometry glyph based visualization. In Fig. 4b we present the combined tensor-HARDI glyph visualization. The ellipsoids better convey the information about the direction compared to HARDI linear glyphs. This can be observed in the area of *corpus callosum* classified as linear. The HARDI glyphs are color coded by the type of anisotropy, where yellow states high anisotropy, and dark blue low anisotropy. All the glyphs colored by yellow have high anisotropy and depending on the threshold tuning can change into linear tensor glyphs. However, it is interesting to observe the glyphs that are colored by dark blue. Even though, due to the min-max normalization, they appear as crossings, the color coding clearly shows that they belong to the grey matter or the ventricles, since crossings should not be observed above and below the corpus callosum. These glyphs can now be masked out with our classification, thus removing any redundant and often misleading information. In Fig. 4c the tensor information is furthermore simplified by showing only fibers, and the rest of the area classified as non-Gaussian represented by HARDI glyphs. This kind of information is much more intuitive to grasp than only glyph visualization, especially in the case of Fig. 4a.

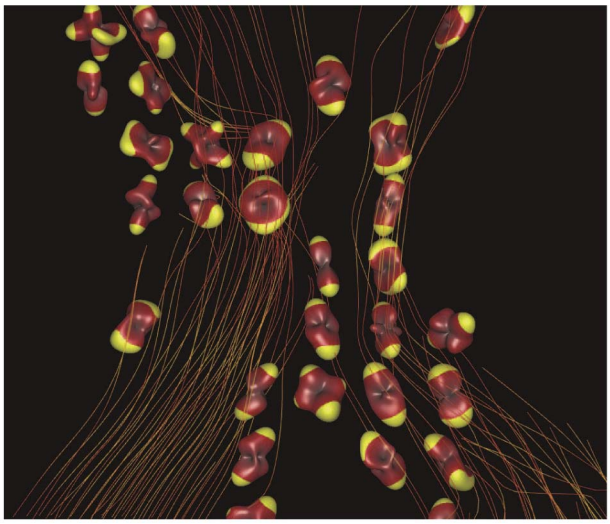

Fig. 6. Hardware phantom built for validation of 65 degree of crossing. The areas where QBall glyphs (min-max normalized) appear, non-Gaussian profiles are expected. The color coding is to enhance maxima. Our visualization shows that the phantom crossings rarely appear in an individual voxel.

5 Results

5.1 DW-MRI Data

We demonstrate the assets of our framework on three different types of DW-MRI measurements: in vivo, human data, hardware phantom data, and ex vivo rat data.

5.1.1 In Vivo Human Data

DW-MRI acquisition was performed on a healthy volunteer (25-year old female) using a twice refocused spin-echo echo-planar imaging sequence on a Siemens Allegra 3T scanner (Siemens, Erlangen, Germany). The scanning was done using FOV 208×208 mm and voxel size $2.0 \times 2.0 \times 2.0$ mm. Ten horizontal slices were positioned through the body of the corpus callosum and centrum semiovale. Data sets were acquired with 80, 106, and 132 directions, each at b-values of $1{,}000,\ 1{,}500,\ 2{,}000,\ 3{,}000\ \mathrm{s/mm^2}$.

5.1.2 Hardware Phantom Data

Hardware phantoms with crossings of 30 degrees, 50 degrees and 65 degrees as in [30] were scanned for using exactly the same DW-MRI protocol as the human subject.

TABLE 1
Rendering Performance for Fused DTI/HARDI Glyph Visualization Using GPU-Based Ray Casting

total #glyphs	GA upper threshold	%HARDI	Performance (FPS)
2601	0	0	120
2601	0.1	15%	50
2601	0.2	46%	20
2601	0.3	68%	15
2601	0.4	84%	12
2601	0.5	94%	10
2601	1	100%	10
10609	0	0%	35
10609	0.1	15%	30
10609	0.2	54%	17
10609	0.3	68%	13
10609	0.4	85%	10
10609	0.5	94%	8
10609	1	100%	8

The measurements were done on a PC with an Intel Pentium 4 3.20 GHz CPU, 2 GB RAM and a GeForce 9,600 GT graphics card with 512 MB of memory and a view port of 845×808 pixels. The step size for ray casting of SH glyphs was 0.05 and we used two refinement steps. The lower threshold for GA was set to 0 in order to render a glyph for each voxel in the selected region of interest. The upper threshold for classification using GA was varied to change the fraction of the glyphs classified as Gaussian (DTI) or nonGaussian (HARDI).

5.1.3 Ex Vivo Rat Brain

We used a 9.4T Bruker Biospec AVANCE-III scanner for acquiring the rat data. Anatomical T2-weighted images were recorded using a RARE sequence. HARDI was done using a diffusion-weighted spin-echo sequence where the voxel size was $0.2 \times 0.2 \times 0.5$ mm. We measured 132 images with different gradient directions and b-value 3,000 s/mm^2, together with an unweighted image.

5.2 Performance

The benefits of our proposed methods rely heavily on the real-time interaction that the user can have with the data in various ways. Because of this, it is very important that all the rendering can be done in real time and that the user does not have to wait for the processing of the data and that no annoying rendering delays occur while interacting with the scene by modifying the view or selecting a region of the data to be visualized. To accomplish this, we make use of the ray casting methods of Section 4. Results are shown in Table 1. In comparison, when we use a previous geometry-based method to render approximately 10 K DTI glyphs with 64 vertices per glyph, we obtain a rendering performance of 6 FPS. This rendering performance was only obtained after generating the vertices, which took 8 s, and thus, cannot be done interactively when changing seed points. Table 1 shows a performance of 35 FPS for the same amount of glyphs using ray casting, without the need to generate geometry when the user changes the region of interest.

The disadvantages of using polygonal representations of the glyphs for the rendering become even worse when visualizing HARDI data. The rendering performance for the GPU-based ray casting of SH glyphs is significantly better than the performance for the geometry-based method (see Peeters et al. [26]) and gives visually more pleasing and accurate results.

However, the biggest advantage here is again not the pure rendering performance, but the ability to change the region of interest where the glyphs are shown without the delay of conventional methods due to the need to generate geometry. For example, in our system, the user has the option to define the seed points which are used to place the glyphs by manipulating an *interactive plane* (shown in the bottom image of Fig. 2) that can be translated and rotated using the mouse, where the new glyphs are immediately visible on the new seed positions on the plane without a delay.

Although, the ray-casting based SH glyph rendering method clearly outperforms the geometry-based method, the rendering of SH glyphs is inherently slower than the GPU ellipsoid rendering, especially when rendering very large amounts of glyphs. This becomes obvious from the rows with 0 percent and 100 percent HARDI glyphs in Table 1. Furthermore, since we have voxel classification (see Fig. 2) and data each individual visualization method writes the appropriate values to the Z-buffer in order to make sure it blends well with other rendered objects in the scene, we can combine the different rendering methods without negative effects on rendering performance. Results for different percentages of HARDI and DTI glyphs are shown in Table 1. Our method optimizes rendering performance by using the more complex SH glyph rendering algorithm only in the areas where this has benefits over the ellipsoid glyph rendering. Typical percentages for HARDI/DTI profiles cannot be given since they depend greatly on scanning parameters. For example, Alexander et al. [3] classify 5 percent of the voxels as nonGaussian while Tuch [35] claims that only 25 percent of the voxels in the human white matter volume, exhibit anisotropic-Gaussian properties. They make these statements for data sets scanned using different acquisition parameters.

5.3 Evaluation

We presented our visualization framework to six independent anonymous researchers experienced in DTI and HARDI and one neurosurgeon. All of the participants in this user study did not have any previous experience in working with our framework. Making standard user evaluation for this kind of setup is a very difficult task and can be a research topic on its own. One major issue is that HARDI techniques are yet to be established, and our main target for now are researchers that develop new algorithms for processing HARDI data and explore its potentials. Performing standard user evaluation with standard tasks would be difficult in this field, especially since the researchers that participate in our study have different research targets. Also, it is impossible to compare our framework to existing HARDI visualization tools, since to the best of our knowledge there are none except a QBall plug-in for Slicer [1] as pointed in the Section 3. Most of the researchers in our user study, develop their own software for processing and viewing HARDI data and use some existing DTI oriented software packages like: MedInria [13], Brainvoyager QX [15], DTITool [37], etc., for manipulating DTI data. Despite all above mentioned difficulties, we tried to quantify the feedback from the participants in the user study, in a structured way, by asking them to state subjective ratings for different aspects of the framework.

TABLE 2
User Scores from the Evaluation, where 1 is Very Good and 5 is Very Bad

	Usability	Learning effort	Usefulness	Interaction	Comparison to existing tools	Existing tool
Expert 1	2	2	3	2	2	Mathematica
Expert 2	3	2	2	1	1	VTK based geometry visualization
Expert 3	1	2	1	1	2	In-house C++ software
Expert 4	3	2	1	2	1	BrainVoyager, TrackMark (DTI only)
Expert 5	2	2	2	1	2	MedINRIA (DTI only)
Expert 6	2	2	2	1	1	MedINRIA, Mathematica
Expert 7	3	2	4	1	X	BrainVoyager (DTI only)

We demonstrated the features of our framework, tried to answer to some of their needs that would facilitate their research, and let them use the framework on their own. The results of this user study are reported in Table 2, where score 1 states very good and 5 very bad. The aspects of consideration were the following:

- Usability: states the possibility and requirements for application of the framework in their everyday research (e.g., requirements for graphics card, compatibility with data formats, etc.).
- Learning effort: states the amount of time/effort for training, such that the framework can be applied in their research.
- Usefulness: states how useful is this framework for their research, assuming that it can be included in their work flow.
- Interaction: states how fast/interactive the users perceived the framework.
- Comparison to existing tools: states how this framework performs in comparison to the existing tools for DTI/HARDI visualization.

The qualitative feedback follows, together with subjective quantitative ratings reported in Table 2, divided according to the research expertise of the participants in our study.

5.3.1 Researchers in Local Modeling and Processing of HARDI Data (Expert 1 and Expert 2)

The work of the first researcher involves development of differential geometric models, measures, and fiber tracking methods for tensor valued DTI/HARDI data. Here, it is highly important to interactively explore the outcome from the fiber tracking techniques but still have information about the underlying local structure.

The work of the second researcher involves HOT estimation of the diffusion propagator $P(\mathbf{r})$ described earlier in Section 2. Very often in this kind of modeling of the data, exploration of the reconstructed profiles in synthetic or real data is of utmost importance. Up to now, the traditional way of rendering the HARDI data, by generating polygonal meshes forces the researcher to work with very small data sets and without any kind of interaction. For creating high-resolution images, the rendering is often put to work "over night" and the results are seen the next day without any ability for further interaction. This is tolerable for generating nice images, but extremely difficult for everyday exploratory purposes. Another drawback of this setup is the inability to select regions of interest, often forcing users to render only certain full brain slices.

Both of the users found our framework interactive (scored as 2 and 1, respectively), and intuitive to use with a small learning effort (scored as 2). They found the framework useful for their current research topics (i.e., limitation of DTI streamline tracking). In the first case, the usefulness was demonstrated with the combined fiber-glyph visualization similar to (see Fig. 1). In the second case, for the interpolation of ODFs, the framework allowed a detailed visualization, necessary for validating the accuracy of their proposed methods. The first researcher was mainly interested in the combined fiber-HARDI glyph visualization, whereas the second one in the HARDI glyph visualization. This type of framework with real-time interaction was positively accepted, with anticipation to increase their productivity if applied in everyday research (scored as useful with 3 and 2, respectively). The second researcher scored the usability of the framework as 3 due to the issues with data reading and the general problem of not having one standardized data format in DW-MRI, but this problem exceeds the scope of this paper.

5.3.2 Cognitive Neuroscience Researcher (Expert 3)

We demonstrated the capabilities of our framework to a cognitive neuroscience researcher with technical background. He found the framework very useful for his research. In his experience, the current visualization techniques in HARDI research are very poor and limited to geometry-based glyph visualization. This kind of visualization is very slow and limited in interaction, and that would be a mere obstacle once they are mature enough to be employed into clinical research. To move the HARDI techniques from pure research into application, a thorough validation must be performed. That would include exploring the outputs from different methods, and looking at the results interactively. There are few clear advantages of the combined visualization: speed and interactivity (that were scored as 1 in Table 2), observation of appearance of nonGaussian departures in real time by changing the thresholds for classification (see Fig. 4), and better visual perception with different color coding. The interaction between planes, fibers, glyphs is also very important and useful for comparing the strength of different techniques, and having more reliable information at hand. One very important feature of the framework was the interactive plane that can be moved around in real time by the user, while observing different view ports of the data (see Fig. 2). This can be a point in favor for possible usage of the tool in clinical research. Another important aspect is the intrinsic level of detail that is automatically offered by the ray casting. The user can zoom out and have a global overview of the data, but also in real time zoom in and see more details without loss of rendering quality. Sometimes, if the seeding for the glyphs is too dense, the information offered to the user can be overwhelming, but it is important that the user can still, in real time, tune the amount of the glyphs and control the complexity of the information stored in the data. This can be very useful when exploring different parts of the white matter. Therefore, the framework was scored as very useful

(scored as 1) and easy to apply in everyday research (scored as 1) with small learning effort (scored as 2).

5.3.3 Application Programmer and DTI Validation Researcher (Expert 4)

Validation of the emerging DTI and HARDI techniques is still an open issue. As the ground truth in human brain is unknown, validation is often done by software-generated synthetic data or hardware phantoms [30]. Exploring the reconstructed data is very important. Here, the expert fancied the ability to interactively move to different selections of the data, and found the fiber-HARDI glyph visualization (see Fig. 6) very important in order to validate the accuracy of the phantom construction. The framework was found to be very fast and useful (scored as 2 and 1, respectively). Also with small learning effort (scored as 2) the framework can be applied in his everyday research (scored as 3 mainly due to the issues with data conversion). A suggestion was made, for an option to quickly zoom out to see an overview of the data with context such that one has an idea of the position of the examined area, and then, zoom back again to look at the details.

5.3.4 Researchers in Imaging for Functional Neurosurgery (Expert 5 and Expert 6).

Improved localization of the target for deep brain stimulation for Parkinson's Disease, the somatomotor part of the *Subthalamic Nucleus* (STN), can reduce the cognitive and emotional side-effects of this treatment. HARDI data provides additional information that can be used to classify the different parts of the STN and to study the local fiber structure as shown by Brunenberg et al. [6]. For our user evaluation, we chose two researchers working on the imaging and functional analysis of the STN. The most important feature of our framework for these researchers is the semiautomatic classification of the data by anisotropic-Gaussian and nonGaussian profiles and different color coding by measure that emphasizes the type of anisotropy in the nonGaussian part of the data. By allowing the user to tune the thresholds, the profiles in the area around the STN could be observed in more detail as depicted in Fig. 5. By using our framework, the researchers came to several conclusions. In ODF representation of the data, for the application of STN, using higher orders of SH are important as the glyph profiles; otherwise, appear to be oblate and it is difficult to perform good discrimination of the different areas in the data. The interactive plane that allows the user to focus on certain regions of interest, was found to be very intuitive and gave the researchers a simple and interactive way to explore the data. The researchers made the same suggestion as in the previous paragraph for improving the context information. Also, integration with anatomical data was suggested. In general, the framework was found to be useful (scored 2 by both experts) and very interactive (scored 1 by both experts), and with some small learning effort (scored 2 by both experts) applied in their everyday research.

5.3.5 Neurosurgeon (Expert 7)

We also performed a user evaluation with a neurosurgeon, experienced with the DTI data (acquisition and modeling) but without technical preknowledge. He found the visualization very interactive (scored as 1) and with small learning effort (scored as 2) usable in his work (scored as 3 mainly due to data conversion issues). However, DTI and HARDI need to be thoroughly validated before they are applied in a clinical setting. Since the clinical users are used to and rely on anatomical data, it is essential for clinical usage of this framework to incorporate visualization of the anatomical data, such that a clinical user can feel more comfortable and has intuition for the current position of the view port. Therefore, without these features, he found the framework not very useful for his everyday research (scored as 4). For clinical purpose, it is important that the visualization is simple (preferably fiber visualization). However, for certain pathologies as brain infarction or multiple sclerosis, knowing the local properties of the data that can be captured with DTI and HARDI glyphs, can be of great help. Other suggestions were made on the user interface that the surgeon found too complex and nonintuitive for people without technical background. He strongly recommended only basic options for data manipulations in the user interface, and several precalculated thresholds that are proven to work best, set as default so that the framework is easier to use.

All the researchers found the framework significantly better compared to the tools that they used in their everyday research (scoring equally 1 and 2), except expert 7 who had difficulties making this comparison as discussed in the previous paragraph. However, this comparison should be taken with care, since most of the tools that were subject of this comparison do not have visualization as their main purpose, and most of them are either DTI oriented software packages or simply have geometry-based HARDI glyphs. The interaction was scored relatively high (score between 1 and 2), and all of them stated that the learning effort will be low to use it in their everyday research (all scored 2). The main problems about the usability of our framework is the data format, but that is a general problem in DW-MRI as there is not well-established standard for a data format. Most of the researchers found the framework useful for their everyday research (five experts scored 1 and 2). One of the researchers scored it as 3 due to the time it would take to be able to incorporate it in their environment. Only one of the participants in the user study scored the framework as 4. The main reason for that was that the user interface for clinical usage should be significantly simplified and many parameters precalculated such that the desirable results appear immediately. Also, integration with anatomical data is essential for this purpose. Therefore, we aim our framework only for research purposes at the moment.

6 Conclusions and Future Work

We presented a novel, interactive framework for combined DTI and HARDI visualization based on the SH representation of HARDI data. To the best of our knowledge, this is the first effort in the literature to improve the performance and control the complexity of information given to the user in the emerging field of HARDI. We exploit, the capabilities of the modern GPU to accelerate the traditionally heavy rendering of the glyph visualization. To give intuition on the data, we adopt fiber tracking techniques based on DTI data, but keep the SH-based glyph rendering of the spherical distribution functions from the HARDI data in nonGaussian areas. The classification on the data is done using measures derived

from the SH coefficients that model the data. The troublesome thresholding is facilitated by a simple algorithm for semi-automatic threshold detection. The user can interact with the data in many ways: browse through different slices, change the scale of the glyphs in real time, min-max normalization, for the ODFs applying deconvolution that sharpens the profiles, and real time classification by additional feedback from the user in fine tuning of the thresholds. We propose different color coding that enhances the maxima in the HARDI glyphs and increases the contrast between different glyph profiles. The result is a fast framework, where the user can interactively control all visualization parameters. To investigate the potentials of the proposed visualization, an informal user evaluation in different applications of DTI and HARDI was performed. In general, the users experienced the visualization as intuitive, fast, detailed, useful, and easy to apply in their every day research. However, the visualization is still too complex for use in medical practice as pointed out by expert 7. Different aspects of our framework appeared to be more useful in different domains. All of the participants in the user study fancied the visualization capabilities of our framework, over several existing tools that they use in their research. Although, this has to be taken with care due to the nature of the evaluation, we believe it shows the potentials of the presented approach.

Several suggestions from the user feedback are addressed as next steps in improving our framework: option to quickly zoom out to see an overview of the data with context, incorporation of anatomical data which is more familiar and reliable in clinical field, as well as improved and simplified user interface for the clinical users.

Future work includes other measures that can find the best separation of the data on nonGaussian and Gaussian profiles in a more robust way. Additionally, future work addresses improving the classification algorithm such that it is done in a more automatic way, or find a more robust algorithm for determining the thresholding intervals. Our current implementation of SH glyph rendering is limited to 4th order of SHs, and an extension to higher orders can be developed. Removing the overwhelming information that still remains in the nonGaussian classified regions, by levels-of-detail representation of the data, will improve the intuition of the visualization. Another important possibility for the data simplification is combining fast deterministic fiber tracking techniques in the Gaussian regions with a more sophisticated tracking techniques in the nonGaussian regions. Also, in a more general sense, the classification of the data by measure can be used as an indication of (un)certainty and can be visualized on the fiber tracts using color, transparency or changing the density of the shown fibers. Additionally, the use of texture and transparency of the glyphs can be studied in order to visualize existing information even more clear, or to show additional information. To the best of our knowledge, the proposed framework is the first attempt to give fast and intuitive insight into the complex HARDI data.

REFERENCES

[1] 3D Slicer, http://www.slicer.org, 2010.

[2] A.L. Alexander, K.M. Hasan, M. Lazar, J.S. Tsuruda, and D.L. Parker, "Analysis of Partial Volume Effects in Diffusion—Tensor MRI," *Magnetic Resonance in Medicine,* vol. 45, pp. 770-780, 2001.

[3] D.C. Alexander, G.J. Barker, and S.R. Arridge, "Detection and Modeling of Non-Gaussian Apparent Diffusion Coefficient Profiles in Human Brain Data," *Magnetic Resonance in Medicine,* vol. 48, no. 2, pp. 331-340, 2002.

[4] A.W. Anderson, "Measurement of Fiber Orientation Distributions Using High Angular Resolution Diffusion Imaging," *Magnetic Resonance in Medicine,* vol. 54, no. 5, pp. 1194-1206, 2005.

[5] P.J. Basser, J. Mattiello, and D. Lebihan, "MR Diffusion Tensor Spectroscopy and Imaging," *Biophysical J.,* vol. 66, no. 1, pp. 259-267, Jan. 1994.

[6] E. Brunenberg, V. Prčkovska, B. Platel, G. Strijkers, and B. ter Haar Romeny, "Untangling a Fiber Bundle Knot: Preliminary Results on STN Connectivity Using DTI and HARDI on Rat Brains," *Proc. First Ann. Meeting of the Benelux Int'l Soc. of Magnetic Resonance in Medicine,* p. 45, 2008.

[7] W. Chen, S. Zhang, S. Correia, and D.F. Tate, "Visualizing Diffusion Tensor Imaging Data with Merging Ellipsoids," *Proc. IEEE Pacific Visualization Symp.,* pp. 145-151, 2009.

[8] P. Cook, Y. Bai, S. Nedjati-Gilani, K. Seunarine, M. Hall, G. Parker, and D. Alexander, "Camino: Open-Source Diffusion-MRI Reconstruction and Processing," *Proc. 14th Scientific Meeting of the Int'l Soc. for Magnetic Resonance in Medicine,* p. 2759, 2006.

[9] R. Deriche and M. Descoteaux, "Splitting Tracking through Crossing Fibers: Multidirectional Q-Ball Tracking," *Proc. Int'l Symp. Biomedical Imaging (ISBI '07),* pp. 756-759, 2007.

[10] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Apparent Diffusion Coefficients from High Angular Resolution Diffusion Imaging: Estimation and Applications," *Magnetic Resonance in Medicine,* vol. 56, no. 2, pp. 395-410, 2006.

[11] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Regularized, Fast and Robust Analytical Q-Ball Imaging," *Magnetic Resonance in Medicine,* vol. 58, pp. 497-510, 2007.

[12] M. Descoteaux and R. Deriche, "Sharpening Improves Clinically Feasible Q-Ball Imaging Reconstructions," *Proc. Int'l Soc. of Magnetic Resonance in Medicine,* p. 906, 2007.

[13] P. Fillard, J. Souplet, and N. Toussaint, Medinria, http://www-sop.inria.fr/asclepios/software/medinria, 2010.

[14] L.R. Frank, "Characterization of Anisotropy in High Angular Resolution Diffusion-Weighted MRI," *Magnetic Resonance in Medicine,* vol. 47, no. 6, pp. 1083-1099, 2002.

[15] R. Goebel, "Brainvoyager QX," http://www.brainvoyager.com. Last visited on 2010.

[16] S. Gumhold, "Splatting Illuminated Ellipsoids with Depth Correction," *Proc. Vision, Modeling, and Visualization,* pp. 245-252, 2003.

[17] C. Hess, P. Mukherjee, E. Han, D. Xu, and D. Vigneron, "Q-Ball Reconstruction of Multimodal Fiber Orientations Using the Spherical Harmonic Basis," *Magnetic Resonance in Medicine,* vol. 56, pp. 104-117, 2006.

[18] M. Hlawitschka, G. Scheuermann, A. Anwander, T. Knösche, M. Tittgemeyer, and B. Hamann, "Tensor Lines in Tensor Fields of Arbitrary Order," *Proc. Int'l Conf. Advances in Visual Computing (ISVC '07),* pp. 341-350, 2007.

[19] S. Jbabdi, M.W. Woolrich, J.L.R. Andersson, and T.E.J. Behrens, "A Bayesian Framework for Global Tractography," *NeuroImage,* vol. 37, no. 1, pp. 116-129, Aug. 2007.

[20] Y. Kanamori, Z. Szego, and T. Nishita, "GPU-Based Fast Ray Casting for a Large Number of Metaballs," *Proc. Eurographics '08,* vol. 27, no. 3, pp. 351-360, 2008.

[21] G. Kindlmann, "Superquadric Tensor Glyphs," *Proc. Eurographics '04,* pp. 147-154, 2004.

[22] S. Mori and P.C. van Zijl, "Fiber Tracking: Principles and Strategies—A Technical Review," *NMR in Biomedicine,* vol. 15, nos. 7/8, pp. 468-80, 2002.

[23] E. Özarslan and T.H. Mareci, "Generalized Diffusion Tensor Imaging and Analytical Relationships between Diffusion Tensor Imaging and High Angular Resolution Diffusion Imaging," *Magnetic Resonance in Medicine,* vol. 50, no. 5, pp. 955-965, 2003.

[24] E. Özarslan, T.M. Shepherd, B.C. Vemuri, S.J. Blackband, and T.H. Mareci, "Resolution of Complex Tissue Microarchitecture Using the Diffusion Orientation Transform (DOT)," *NeuroImage,* vol. 36, no. 3, pp. 1086-1103, July 2006.

[25] E. Özarslan, B.C. Vemuri, and T.H. Mareci, "Generalized Scalar Measures for Diffusion MRI Using Trace, Variance, and Entropy," *Magnetic Resonance in Medicine,* vol. 53, no. 4, pp. 866-876, 2005.

[26] T.H.J.M. Peeters, V. Prčkovska, M. van Almsick, A. Vilanova, and B.M. ter Haar Romeny, "Fast and Sleek Glyph Rendering for Interactive HARDI Data Exploration," *Proc. IEEE Pacific Visualization Symp.*, pp. 153-160, 2009.
[27] T.H.J.M. Peeters, A. Vilanova, G.J. Strijkers, and B.M. ter Haar Romeny, "Visualization of the Fibrous Structure of the Heart," *Proc. Vision, Modeling and Visualization,* pp. 309-316, 2006.
[28] M. Perrin, Y. Cointepas, C. Poupon, B. Rieul, N. Golestani, D. Rivière, A. Constantinesco, D.L. Bihan, and J.-F. Mangin, "Fiber Tracking in Q-Ball Fields Using Regularized Particle Trajectories," *Proc. Int'l Conf. Information Processing in Medical Imaging (IPMI '05),* pp. 52-63, 2005.
[29] V. Prčkovska, A. Vilanova, C. Poupon, B.M. Haar Romeny, and M. Descoteaux, "Fast Classification Scheme for HARDI Data Simplification," *Proc. ICT Innovations '09,* pp. 345-355, Springer, 2010.
[30] W. Pullens, A. Roebroeck, and R. Goebel, "Kissing or Crossing: Validation of Fiber Tracking Using Ground Truth Hardware Phantoms," *Proc. Int'l Soc. of Magnetic Resonance in Medicine,* p. 1479, 2007.
[31] T. Schultz and H.-P. Seidel, "Estimating Crossing Fibers: A Tensor Decomposition Approach," *IEEE Trans. Visualization and Computer Graphics,* vol. 14, no. 6, pp. 1635-1642, Nov./Dec. 2008.
[32] D.W. Shattuck, M.-C. Chiang, M. Barysheva, K.L. McMahon, G.I. de Zubicaray, M. Meredith, M.J. Wright, A.W. Toga, and P.M. Thompson, "Visualization Tools for High Angular Resolution Diffusion Imaging," *Proc. Int'l Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI '08),* pp. 298-305, 2008.
[33] C. Sigg, T. Weyrich, M. Botsch, and M. Gross, "GPU-Based Ray-Casting of Quadratic Surfaces," *Proc. Eurographics '06,* pp. 59-65, 2006.
[34] J.D. Tournier, F. Calamante, and A. Connelly, "Robust Determination of the Fibre Orientation Distribution in Diffusion MRI: Non-Negativity Constrained Super-Resolved Spherical Deconvolution," *Neuroimage,* vol. 35, no. 4, pp. 1459-1472, 2007.
[35] D. Tuch, "Diffusion MRI of Complex Tissue Structure," PhD thesis, Harvard Univ., 2002.
[36] D. Tuch, "Q-Ball Imaging," *Magnetic Resonance in Medicine,* vol. 52, pp. 1358-1372, 2004.
[37] A. Vilanova, G. Berenschot, and C. van Pul, "DTI Visualization with Stream Surfaces and Evenly-Spaced Volume Seeding," *Proc. IEEE Eurographics '04,* pp. 173-182, 2004.
[38] A. Vilanova, S. Zhang, G. Kindlmann, and D. Laidlaw, "An Introduction to Visualization of Diffusion Tensor Imaging and Its Applications," *Visualization and Processing of Tensor Fields,* J. Weickert and H. Hagen, eds., pp. 121-153, Springer, 2005.
[39] D. Wassermann, M. Descoteaux, R. Deriche, and C. Westin, "QBall Plug-In for Slicer3D," http://www-sop.inria.fr/odyssee/en/software/qballslicer, 2010.

Vesna Prčkovska received the Diploma in electrical engineering (equivalent to master's degree) in 2006 from the faculty of electrical engineering, University Ss. Cyril and Methodius, R. Macedonia. She finished her PhD degree at the Biomedical Image Analysis Group of the Biomedical Engineering Department at the Eindhoven University of Technology. The title of her PhD thesis was "High Angular Resolution Diffusion Imaging, Processing, and Visualization." Her research interests include medical visualization, multifield visualization, modeling, processing, and visualization of HARDI data.

Tim H.J.M. Peeters received the MSc and PhD degrees in 2004 and 2009, respectively, from Eindhoven University of Technology. His PhD research in the Biomedical Image Analysis Group of the Department of Biomedical Engineering involved GPU-based visualization techniques for interactive exploration of diffusion-weighted MRI data.

Markus van Almsick received the Diploma (MSc degree) in physics from the Technische Universität München, Germany, in 1990. He received the PhD degree in 2007 for his research in developing context models of lines and contours. He worked for the University of Illinois, Wolfram Research Inc., and the Max-Planck Institute for Biophysics, Frankfurt, before joining the Biomedical Image Analysis Group in the Department of Biomedical Engineering at the Technische Universiteit Eindhoven. His research interests include image analysis, computer algebra, and fundamental physics.

Bart ter Haar Romeny received the MSc degree in applied physics from Delft University of Technology in 1978, and the PhD degree from Utrecht University, in biophysics, in 1983. He is a full professor at Eindhoven University of Technology, heading the Biomedical Image Analysis (BMIA) Group in the Department of Biomedical Engineering. He has more than 20 years' experience of biomedical image analysis and computer vision research. During the years from 1986 to 1989 he was the principal physicist of the Utrecht University Hospital's Radiology Department and clinical project leader of the Dutch PACS project. He was a cofounder and an associate professor at the Image Sciences Institute (ISI) of Utrecht University from 1989 to 2001. His research interests include quantitative medical image analysis, its foundations, and clinical applications. His main interests include mathematical modeling of visual perception and applying this knowledge in operational computer-aided diagnosis systems.

Anna Vilanova received the PhD degree in 2001 from the Vienna University of Technology. The title of her PhD thesis was "Visualization Techniques for Virtual Endoscopy." She is an assistant professor with the Biomedical Image Analysis Group of the Biomedical Engineering Department at the Eindhoven University of Technology. She is currently leading a research group with several PhDs in the subject of multivalued image analysis and visualization. Her research interests include medical visualization, volume visualization, multivalued visualization, and medical image analysis.