Visualization of uncertainty in fiber tracking based on diffusion tensor imaging

Visualization of Uncertainty in Fiber Tracking Based on Diffusion Tensor Imaging

Colophon

This thesis was typeset by the author using LaTeX2$_\varepsilon$.

About the cover

The front and back covers are designed by Kim Rauwerdink.

The left image on the front cover shows white matter fibers color-coded by fractional anisotropy threshold. The middle image shows how a fixed safety margin surrounding a single fiber does not adequately cover the potential fibers created after repeated measurements. The right image shows a tumor color-coded with distances to the white matter fibers (indicated in blue).

A catalogue record is available from the Eindhoven University of Technology Library
ISBN: 978-90-386-3159-2

Printed by Offpage, The Netherlands

Visualization of Uncertainty in Fiber Tracking Based on Diffusion Tensor Imaging

Proefschrift

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op woensdag 13 juni 2012 om 16.00 uur

door

Ralph Brecheisen

geboren te Eersel

Dit proefschrift is goedgekeurd door de promotor:

prof.dr.ir. B.M. ter Haar Romeny

Copromotoren:
dr.ir. A. Vilanova
en
dr.ir. B. Platel

Contents

And bring me a hard copy of the Internet so I can do some serious surfing

- Dilbert

1 Introduction

1.1 Introduction

In this thesis we present techniques for the visualization of uncertainty in 3-dimensional reconstructions of white matter pathways in the human brain. In this chapter we will start by outlining the main research problems which these techniques attempt to address.

Recent advances in medical imaging have led to a wealth of new information about the anatomy and function of the human body. These developments are especially noticeable in medical imaging of the brain. The brain remains one of most mysterious organs of the human body and even though significant progress has been made in revealing and understanding brain function on the basis of medical imaging we are still far removed from a complete understanding. Part of the research problem lies in identifying and delineating specific functional regions in various locations of the brain. Another part lies in investigating how these functional regions interconnect to form complex brain networks. These are commonly considered to be the basis of the brain's phenomenal processing capabilities. A large part of these networks, especially the long range connections, are formed by brain white matter which consists of myelinated axons extending from brain cells or neurons in the brain gray matter. One of the most promising medical imaging technologies currently available for the analysis of brain white matter connections is Diffusion Tensor Imaging, commonly abbreviated as DTI. It is an imaging technique based on magnetic resonance imaging (MRI) which allows quantification of local fiber orientation in living tissues through measurements of water diffusion. The orientational information captured by DTI in each point of the tissue can be used to reconstruct fiber pathways in three dimensions through a procedure commonly called *fiber tracking*. Visualization of these pathways offers unique insight into the 3D layout of white matter fiber bundles and provides great potential for both neuro-scientific research and clinical applications such as neurology, neuroradiology and neurosurgery. For example, fiber tracking allows neurosurgeons to assess the spatial relation between a brain tumor and surrounding critical fiber bundles. This provides useful information for planning the angle of approach toward the tumor during resection.

1.2 Problem background

Despite the many advantages of complex imaging technologies such as DTI, there is also a downside. The amount and complexity of the generated image data poses considerable challenges for analysis and interpretation. For example, MRI is one of the most flexible and versatile imaging technologies available today. As already explained above, it allows the characterization of tissue structure through water diffusion but can also measure brain function, cell metabolism, blood flow, blood perfusion and many other properties. However, this wealth of information is not immediately available to the end user. It needs to be extracted from the

potentially thousands of raw images through a complex pipeline of mathematical modeling, analysis and visualization. As a consequence, a direct link between what is being measured and what is being visualized is no longer present. The processing pipeline behaves like a black box whose inner workings remain hidden from the user and which makes reliable interpretation of the output difficult or even impossible. To make matters worse, this black box can be a major source of errors and uncertainties itself.

DTI and fiber tracking are a prime example of this problem. The 3-dimensional insight they provide has great advantage over 2-dimensional representations such as provided by conventional MR imaging. However, the DTI imaging and processing pipeline, of which fiber tracking is one of the final stages, can be subject to considerable error and uncertainty. This pipeline is necessary in order to extract useful information from the high-dimensional image data. However, the errors and uncertainties arising from the different pipeline stages can lead to unpredictable variations in the final output. Since the main purpose of fiber tracking is to increase our knowledge about brain function or to use it as a clinical tool, it is essential that we try to characterize and quantify the errors and uncertainties as completely as possible and visually communicate their effect to the user.

This is not a trivial matter though. Uncertainty can arise for an almost infinite number of reasons. Measuring devices may suffer from noise and bias, have limited resolution or may be incorrectly configured. If the measured data is too complex to handle, it may have to be simplified using mathematical models. However, selecting the best model may be difficult and many models over-simplify the data thereby ignoring important characteristics or even producing incorrect results. Further processing may be required to extract additional features of interest from the data or the model. However, the selected features may give an incomplete or even incorrect representation of the phenomenon of interest. They may also be highly sensitive to parameter settings and errors occurring earlier in the pipeline. In medical applications additional uncertainties arise. For example, tissue pathology may disturb the measured signal in unexpected ways. Furthermore, there may be lack of knowledge about the true relation between tissue properties and the measured signal. As we will see in Chapter 4, when we discuss errors and uncertainties in more detail, this lack of knowledge is also present in DTI and fiber tracking.

Basically, the list of potential sources of uncertainty is endless. The more you zoom in on details of the pipeline, the more elements you will find that can potentially give rise to errors and uncertainties. This observation also illustrates several limitations and constraints for analysis and visualization. A large part of the uncertainties arising from the DTI processing pipeline cannot be quantified or represented in a way that allows their analysis and visualization. For example, it is very difficult to determine the exact error introduced by measuring a certain

property at the millimeter scale while the phenomenon of interest takes place at the micrometer scale. In DTI there is also no established ground truth against which to validate measurement and post-processing results. Fiber tracking can only be validated through indirect methods that have considerable limitations of their own. For example, one can create phantoms that mimick brain white matter using artificial materials and place these phantoms inside a DTI scanner. However, such phantoms are always limited in their ability to capture the complex physiology of real, living brain tissue. Other methods, such as chemical tracers, can highlight selected fibers very accurately but are invasive, highly toxic and cannot reveal more than a few, microscopically thin, fibers at a time. Validation can also be done on the basis of anatomical knowledge, for example, in neurosurgery. This can be problematic, however, in the presence of tissue pathology, such as brain tumors, which can distort or otherwise affect the anatomy in unexpected ways.

Without good validation it is difficult to make definitive statements about the *accuracy* of fiber tracking. This should be kept clearly in mind in the remainder of this thesis as well. Accuracy is defined as the extent to which a measured value matches the true value. In the context of DTI, it refers to the correspondence between the pathways reconstructed with fiber tracking and the underlying white matter anatomy. As explained above, it is very difficult to establish accuracy without validation against a gold standard. For this reason, most methods currently available for the analysis of DTI uncertainty focus on precision or *reproducibility*. This refers to the extent in which repeated measurements or calculations produce similar output results. For example, a wide range of probabilistic fiber tracking algorithms has been proposed that attempt to describe the reproducibility of fiber tracking as affected by image noise.

Despite the fact that existing methods for uncertainty analysis leave many sources of uncertainty unexplored, they can still provide very useful information, especially for users who are less familiar with the underlying technology and its limitations. For example, neurosurgeons primarily consider fiber tracking to be a visualization tool which can potentially assist clinical decision-making. They cannot be expected to have detailed knowledge about DTI-specific image artifacts, model differences and limitations, tracking errors, parameter settings or rendering techniques. Uncertainty visualization can help such users to become aware of these issues. Unfortunately, most tracking algorithms commonly used in clinical settings do not calculate nor visualize uncertainty in any way. Fiber pathways are all visualized as equally reliable, even though some may be false positives that have no correspondence to the underlying anatomy. Similarly, fiber pathways may be missing in the visualization while in reality they do exist. Such false negatives can be caused by imaging or processing errors and pose a considerable problem for neurosurgical planning because they can be easily mistaken for tissue pathology effects.

1.3 Challenges

In the previous paragraphs we have described the central problem of uncertainty in fiber tracking. We have also explained that, if fiber tracking is to be used for scientific research or clinical decision-making, it is important to visually communicate the uncertainty to the user. Finally, we pointed out that there are fundamental limitations to what we can achieve in terms of quantifying and representing uncertainty. Since this thesis deals primarily with visualization, we rely on existing methods for uncertainty analysis. Here, we will describe the main challenges related to uncertainty visualization and how these apply to DTI and fiber tracking.

The first challenge is related to *data complexity*. DTI data is high-dimensional, potentially consisting of dozens of measured values for each voxel. By means of mathematical modeling data dimensionality can be significantly reduced but the result is still difficult to visualize and has kept many researchers busy over the last few years. Adding uncertainty information, even as a simple scalar value, further complicates this problem and can easily lead to visual clutter and information overload. The second challenge of uncertainty visualization is finding strategies for evaluating its *effectiveness*. There are different goals involved when considering this. For example, one could evaluate whether *awareness* of the presence of uncertainty is useful to the user. In some cases, users may suspect reliability issues but under- or overestimate the magnitude of the problem. In other cases, users may have been completely unaware of the presence of uncertainty. Another goal would be to show that visualizing uncertainty actually affects and even improves *decision-making*, that is, leads to better outcomes as compared to decisions made without taking uncertainty into account. As we will see in later chapters, this is a very difficult thing to do, especially since surgical decision-making is affected by many factors. It is far from trivial to prove that showing uncertainty in fiber tracking is the main factor responsible for a change in decision. In this thesis we will primarily focus on eliciting an awareness of uncertainty in the user and evaluating whether this provides useful information. To accomplish this we develop several visualization tools for this purpose and discuss the output of these tools extensively with clinical experts.

1.4 Thesis outline

In this section we provide a short overview of our contributions. We start the main body of this thesis with two background chapters. Chapter 2 presents a medical overview of brain anatomy and the main applications for fiber tracking discussed in this thesis. Chapter 3 focuses on the general principles behind diffusion tensor imaging and fiber tracking which can help readers less familiar with these topics to follow the remaining chapters. In Chapter 4 we provide an overview of research performed in uncertainty analysis for DTI and fiber tracking. The

first part focuses on the different sources of uncertainty in the DTI processing pipeline. The second part discusses probabilistic fiber tracking methods which serve as a basis for many of the uncertainty visualization techniques presented in later chapters. In Chapter 5 we present an overview of related work in the field of uncertainty visualization. Our first contribution for visualizing parameter sensitivity in fiber tracking algorithms is presented in Chapter 6. Next, we proceed with an approach to visualize pathway confidence intervals on the basis of probabilistic fiber tracking in Chapter 7. Chapter 8 presents a general framework for the interactive exploration of distances between objects with uncertain spatial extent. We apply this framework in the context of neurosurgical planning where distances are assessed between a brain tumor and surrounding white matter tracts calculated from probabilistic fiber tracking. In Chapter 9 we present a number of clinical cases involving brain tumor patients and discuss both the benefits and limitations of uncertainty visualization in neurosurgical settings. We conclude our thesis with Chapter 10 where we discuss our contributions in a wider context and present suggestions for future work.

Change is good. You go first.

- Dilbert

2 Medical background

2.1 Introduction

Visualization research typically operates on the boundary of multiple application domains and fields of expertise. This is primarily what makes visualization interesting, but it can also make it difficult because we have to know something, at least at a basic level, about all the application domains involved. In this thesis we present visualization techniques for fiber tracking based on diffusion tensor imaging and we apply these techniques in the context of brain connectivity research and neurosurgical planning. The primary interest of both these application domains is the living human brain. For this reason, we present a comprehensive overview of brain anatomy in this chapter. We also describe shortly what brain connectivity research and neurosurgical planning involve and how DTI fiber tracking is used in these domains.

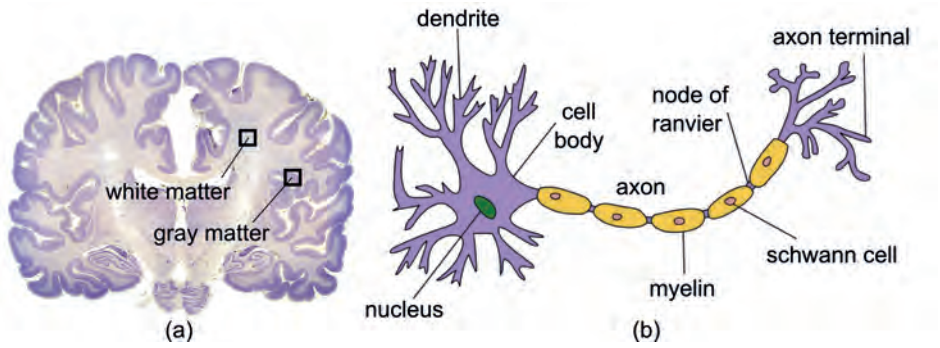

Figure 2.1: (a) Gray and white matter in cell-stained slice (http://www.brainmuseum.org). (b) Architecture of a neuron (Quasar Jarosz, http://en.wikipedia.org).

2.2 Brain anatomy

The brain is the center of the nervous system and the most complex organ in the human body. It consists of billions of *neurons*, which are electrically excitable cells that interconnect in complex networks through electrical and chemical signaling. A typical neuron possesses a cell body and dendrites (Figure 2.1(b)) which are thin branches extending from the cell body and connecting to hundreds of other neurons. The collection of neuron cell bodies and interconnecting dendrites forms a dense network and is generally called *gray matter* (Figure 2.1(a)). Besides dendrites, each neuron also has a single *axon* which is a special extension of the cell body transporting electrical signals across large distances to remote parts of the body. Axons are encapsulated by a *myelin sheath* whose primary function is to speed up signal transmission. The collection of myelinated axons is generally

called *white matter* and is responsible for the transfer of information between different gray matter regions (Figure 2.1(a)).

2.2.1 Gray matter

Gray matter is primarily located on the outer surface of the brain, called the *cerebral cortex*, where it forms a layer of neuronal tissue. This layer is intricately folded in order to obtain a maximum surface area inside the confined space of the skull (Figure 2.2(a)). The folds, or valleys, themselves are called *sulci* (singular: sulcus) while the smooth areas between folds are called *gyri* (singular: gyrus). The pattern of gyri and sulci shows striking similarities between different subjects which allows them to be classified and named (Figure 2.2(b)). The cerebral cortex is commonly subdivided into four *lobes*: frontal, occipital, parietal and temporal (Figure 2.2(a)). Gray matter is commonly associated with the 'processing' functions of the brain, such as memory, language, motor and vision.

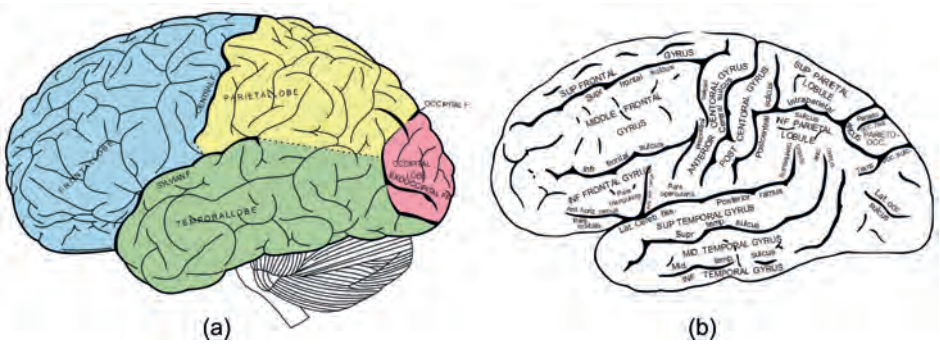

Figure 2.2: (a) Four major cortical lobes. (b) Major gyri and sulci.

2.2.2 White matter

White matter bundles interconnect different gray matter regions and are also associated with specific brain functions. For the purpose of this thesis we limit ourselves to a description of only the major white matter bundles. Because these bundles are also discussed in the following chapters we consider them to be most relevant. Three different *types* of white matter bundles can be found in the brain:

- **Projection fibers** - These mostly run vertically and interconnect functional regions in the cerebral cortex with the rest of the body (Figurer 2.3(a)).
- **Association fibers** - These interconnect cortical regions within the same hemisphere (Figure 2.3(b)).
- **Commissural fibers** - These interconnect the two brain hemispheres (Figure 2.3(c)).

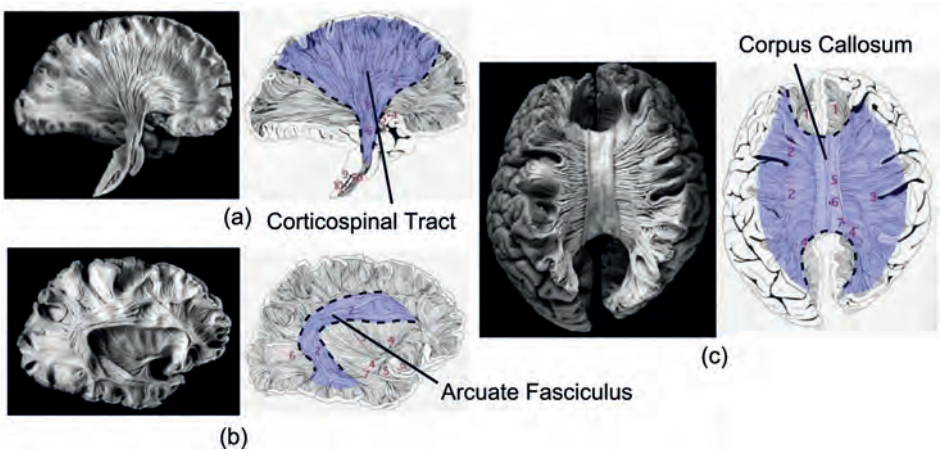

Figure 2.3: (a) Projection fibers. (b) Association fibers. (c) Commissural fibers. Images adapted from Williams et al. [167].

The three bundle types can be further subdivided into specific bundles that are anatomically distinguishable in shape. A single, identifiable fiber bundle, however, can serve multiple brain functions. A clear example is the cortico-spinal tract (see details below) which transports information from a wide area in the cortex down the spinal cord to the rest of the body. Only the fibers originating from the primary motor cortex are involved in specific motor functions, e.g. hand movement. Below, we will describe three major fiber bundles to which we will refer in the following chapters. These are illustrated and highlighted in blue in Figure 2.3. Many more exist and for a complete overview we refer the reader to the white matter atlas published by Oishi et al. [115].

- **Corticospinal tract** - The cortico-spinal tract (CST) is the collection of myelinated axons that run between cerebral cortex and the spinal cord. It consists mostly of axons involved with motor function, but especially the fibers originating from the primary motor cortex are of neurosurgical interest. Damage to these fibers can result in permanent functional deficits (Figure 2.3(a)). The primary motor cortex is located in the precentral gyrus, just anterior to the central sulcus in the frontal lobe (Figure 2.2(a)).
- **Arcuate fasciculus** - The arcuate fasciculus is the fiber bundle connecting the posterior part of the temporo-parietal lobe with the frontal cortex. Recently, it is considered to be part of the superior longitudinal fasciculus, a larger complex of fiber bundles involved with a wide range of brain functions (Figure 2.3(b)). For neurosurgical purposes, however, the arcuate fasciculus has specific interest because of its involvement with language function.
- **Corpus callosum** - The corpus callosum is the largest white matter struc-

ture in the brain and belongs to the category of commissural fibers. It primarily facilitates communication between the two cerebral hemispheres (Figure 2.3(c)). Its posterior part is called the *splenium*, its anterior part the *genus*.

2.3 Applications for fiber tracking

The ability to visualize the 3D layout of fibrous tissues offers great potential for various application domains, especially those involving neuronal tissue such as encountered in brain white matter. In this section, we will provide a comprehensive overview of two such applications: brain connectivity research and neurosurgery. We will discuss (1) their main goals, (2) how fiber tracking can be used to reach those goals and (3) how uncertainties in the fiber tracking output affect this process.

2.3.1 Brain connectivity research

The primary goal of brain connectivity research is to map out structural networks of brain connections in as much detail as possible. This is considered to be essential for understanding brain function. Diffusion tensor imaging and fiber tracking can play an important role in this research, along with other methods for studying brain connections such as chemical tract tracing, optical imaging and basic histology. An example of fiber tracking used for brain connectivity research is presented by Hagmann et al. [57] who subdivided the white matter/gray matter interface into 500-4000 equally sized regions of interest (ROI's). Based on whole-brain fiber tracking they constructed weighted networks of interconnected cortical ROI's (Figure 2.4(b)). Another example is Track Density Imaging, originally proposed by Calamante et al. [28], a DTI-based imaging technique which uses the global information obtained from millions of fiber tracts reconstructed with whole-brain fiber tracking to create images of white matter structures at a resolution far beyond that of the acquired images (Figure 2.4(a)).

As we will see in Chapter 4 on DTI uncertainty, fiber tracking cannot provide the same confidence in the existence of a brain connection as, for example, chemical tract tracing. However, using probabilistic fiber tracking (Chapter 4) it is possible to make statements about the reproducibility of the tracking results. Also, unlike chemical tracing, fiber tracking can be done in-vivo and for large sets of voxels simultaneously. For this reason, a large part of brain connectivity research consists in combining different sources of connectivity information and cross-validating the results of fiber tracking with alternative tract reconstruction methods [12].

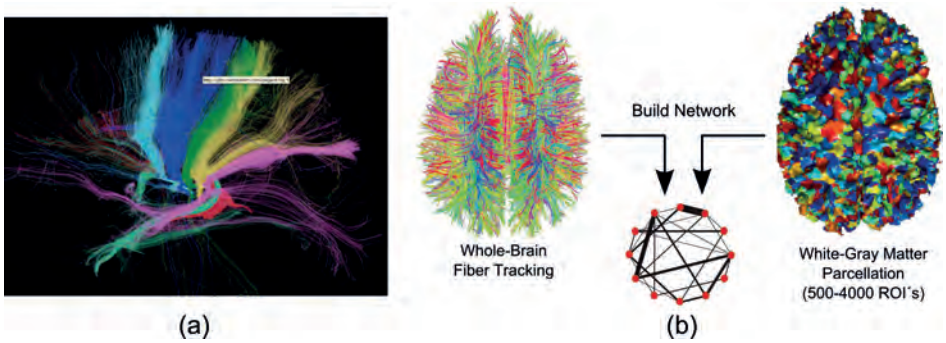

Figure 2.4: (a) Fiber tracts constructed from Track Density Imaging [28]. (b) Construction of brain network on the basis of whole-brain fiber tracking and a 500-4000 ROI white-gray matter parcellation [57].

2.3.2 Neurosurgical planning

The ability of DTI and fiber tracking to non-invasively reveal anatomical connections in brain white matter also offers great potential for neurosurgical planning. By highlighting functionally relevant fiber pathways in 3-dimensions neurosurgeons are able to improve preoperative patient evaluation, surgical planning, intraoperative navigation and patient outcome [6]. Such information cannot be easily obtained from conventional MR imaging because it represents brain white matter as a homogeneous mass without details about individual fiber bundles.

Despite these advantages, a reliable interpretation of fiber tracking results in neurosurgical settings can be difficult. The main reason is the presence of pathological tissues, such as a brain tumor. These may not only affect the behavior of water diffusion, which is the basic property being measured by DTI, but also affect the underlying fiber structure itself. For example, a brain tumor may exert varying effects on fiber bundles, such as displacement, infiltration or disruption. Furthermore, certain tumor types are surrounded by a layer of fluid, called edema, which can cause fiber tracking algorithms to fail, even though the underlying fiber structure is intact. In Chapter 4 we will discuss probabilistic fiber tracking algorithms which attempt to deal with such problems.

Brain tumors

An important part of neurosurgery is dedicated to the resection of brain tumors. In the following chapters we will regularly refer to the planning of such procedures. Here, we give a short overview of the different tumor grades as defined by the World Health Organization (WHO). Tumor grades are based on the microscopic appearance of tumor cells. The *less* these cells look like normal, differentiated

cells, the more aggressive the tumor is likely to be. The grades are defined as follows:

- **Grade I** - Well-differentiated (low grade)
- **Grade II** - Moderately differentiated (medium grade)
- **Grade III** - Poorly differentiated (high grade)
- **Grade IV** - Undifferentiated (high grade)

In Chapter 9 we will describe five clinical cases of brain tumor patients. Most of these have a low-to-medium-grade glioma (WHO-I/II). These are slow-growing tumors with a relatively good prognosis (10+ years) *if* the tumor can be resected as completely as possible. For this, tumor tissue may need to be resected *beyond* the boundaries visible on the medical images. The benefits of advanced techniques such as DTI and functional MRI are particularly appreciated here because they assist the neurosurgeon in finding the exact boundaries of functionally critical brain regions and white matter bundles. This is especially important for low-grade glioma patients because their functional anatomy may be altered. Due to the slow growth of these tumors, the brain has time to reorganize and move certain brain functions elsewhere.

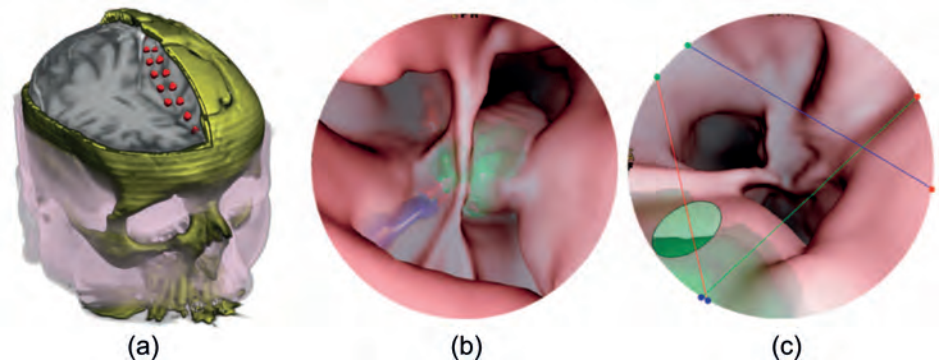

Figure 2.5: (a) 3D multimodal visualization using GPU-based volume rendering [19]. (b,c) Two endoscopic visualizations of the nasal cavity as accessed in pituitary tumor resection [112].

Visualization strategies for neurosurgical planning

Neurosurgical procedures are often complex and associated with considerable risks. The margin for error is small and obtaining a good view of the target

lesion and surrounding brain structures is difficult on the basis of 2D imaging alone. For this reason, neurosurgical planning has received wide interest in past visualization research. As the following chapters will show, we apply and evaluate our visualization techniques also in neurosurgical settings. We will give a short overview of related work in this area here.

Neurosurgery typically involves multiple imaging modalities such as structural MRI, CT, functional MRI and DTI. Combining datasets from these modalities in a single 3D visualization is challenging and several solutions have been proposed for it, some using CPU-based methods [141] and others GPU-accelerated techniques [19]. Different systems often focus on different aspects. Some focus on specific neurosurgical procedures such as endoscopic pituitary tumor resections [112] (see Figure 2.5). Others propose more general visualization systems for neurosurgical purposes, combining anatomical MRI, functional MRI and DTI in a single view [19, 21, 51, 86]. Still other systems address the general problem of rendering multiple volumetric datasets without explicitly focusing on a neurosurgical context [137, 136, 25]. Most of these neurosurgically-oriented systems rely on deterministic fiber tracking. An explicit indication of errors and uncertainty in these systems is not common yet.

2.4 Conclusions

In the previous sections we have provided a short overview of basic medical background information relevant to the remainder of this thesis. We will now proceed with a more technical background chapter explaining the basic principles of diffusion tensor imaging and fiber tracking.

Let's form proactive synergy restructuring teams.

- Dilbert

3 Technical background

3.1 Introduction

In the previous chapter we introduced the reader to basic medical background related to brain anatomy and two applications for DTI and fiber tracking that are of particular interest to this thesis: brain connectivity research and the neurosurgical planning of brain tumor resections. In this chapter we provide more technical background as we discuss the general principles behind diffusion tensor imaging, fiber tracking and 'traditional' visualization strategies that are currently used in this area. We will limit ourselves primarily to deterministic fiber tracking algorithms which, given the same input, always produce the same output. In chapter 4 we discuss another class of tracking algorithms, called *probabilistic* fiber tracking, which form the basis of many visualizations presented in this thesis.

3.2 Characteristics of diffusion

Diffusion is an essential transport mechanism in living organisms. For example, transport of metabolites into cells is mainly facilitated by means of diffusion. Diffusion is the process of thermally-driven, random movement of molecules over time, also called Brownian motion. The diffusion process is commonly described by Fick's Law of Diffusion:

$$\frac{\partial}{\partial t}P(\mathbf{r},t) = D\nabla^2 P(\mathbf{r},t) \tag{3.1}$$

In this equation, $P(\mathbf{r},t)$ is a probability density function (PDF) describing the probability that a molecule is displaced by a vector $\mathbf{r}$ in time t. The coefficient D is the diffusion coefficient (m^2/s), and depends on molecule type and properties of the medium. ∇^2 is the Laplacian. If diffusion is unrestricted in 3D space, D is a scalar constant and Eq. (3.1) is solved by a 3D Gaussian distribution:

$$P(\mathbf{r},t) = \frac{1}{\sqrt{(4\pi Dt)^3}}\exp\left(\frac{-\mathbf{r}^2}{4Dt}\right) \tag{3.2}$$

In the next paragraph we will describe how measurements of water diffusion can be used to extract information about tissue physiology and structure. We will pay special attention to the possibilities of deriving orientational information in fibrous tissues such as the myelinated axons encountered in the brain white matter.

Unrestricted, apparent and anisotropic diffusion

In free, unrestricted fluids the diffusion coefficient D is a scalar constant. In pure water at body temperature (37 degrees Celcius) the diffusion coefficient is 3×10^{-3} $\mathrm{mm}^2\mathrm{s}^{-1}$ (Figure 3.1(a)). In tissues, however, water diffusion is restricted by physical barriers such as cell membranes, cytoskeletal structures and macromolecules. In this case, the diffusion coefficient becomes time-dependent and can be described as $D(t)$. If the observation time t is small, the diffusing molecules

will not have encountered any barriers yet and $D(t) = D$. If the observation time is large, molecules will have encountered barriers and their movement is restricted (Figure 3.1(b)). In this case, $D(t) < D$. The reduced diffusion coefficient is commonly called the *apparent diffusion coefficient* or ADC [90, 144] and can be used to detect the presence of physical barriers in tissue.

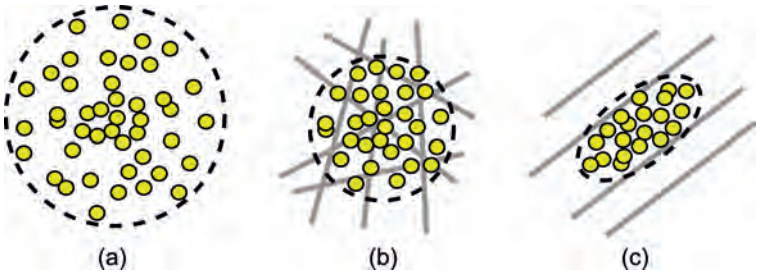

Figure 3.1: (a) Isotropic diffusion without physical barriers. (b) Restricted isotropic diffusion without preferred orientation. (c) Restricted anisotropic diffusion with preferred orientation.

In fibrous tissues, water diffusion shows orientational behavior. For example, in brain white matter, diffusion of water molecules *along* the fibers will be less restricted than *perpendicular* to the fibers. In the next paragraphs we will explain how diffusion tensor MRI can make use of this orientational behavior to reconstruct local fiber orientations in tissues.

3.3 Diffusion tensor MRI

Diffusion tensor MRI, or diffusion tensor imaging (DTI), is an imaging technique based on magnetic resonance (MR) that allows measurement of diffusion in living tissues. One of the first MR pulse sequences that allowed to do this was proposed by Stejskal and Tanner [144] in 1965. They proposed the pulse gradient spin echo sequence that, apart from the standard magnetic gradients used for spatial encoding of the image, applies additional magnetic gradients to measure diffusion along multiple gradient directions $\mathbf{g}_i$, $(i \in 1, ..., N)$. Stejskal and Tanner derived the following relation between diffusion and the measured MR signal $S(\mathbf{g}_i)$, based on exponential signal decay and a Gaussian diffusion distribution:

$$S(\mathbf{g}_i) = S_0 \exp(-b\ D_i) \tag{3.3}$$

where S_0 is the T2-weighted MR signal without diffusion-weighting, b is an MR acquisition parameter, called b-value, which depends on gradient strength and diffusion time t. D_i is the apparent diffusion coefficient (ADC) along gradient direction $\mathbf{g}_i$. It should be noted that this formulation assumes exponential signal

decay and Gaussian-distributed diffusion. In practice, this requires a b-value < 2000 s/mm^2.

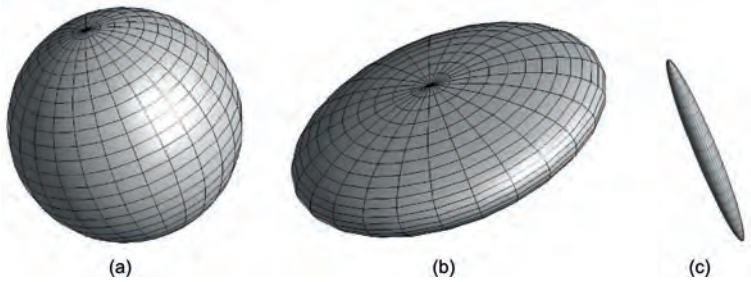

Figure 3.2: Shapes of the 2nd-order tensor: (a) spherical or isotropic, (b) planar and (c) ellipsoidal or linear.

3.3.1 Mathematical description of diffusion

As explained earlier, apparent diffusion is measured along a given gradient direction $\mathbf{g}_i$. The 3D shape of water diffusion can be probed by measuring diffusion in many different gradient directions. In this thesis we will commonly refer to the 3D shape of diffusion as the *diffusion profile*.

In DTI, it is assumed that water diffusion in each voxel can be described by a 3D Gaussian distribution. A second-order tensor can be used to model such a distribution. This model can describe spherical (isotropic), planar and ellipsoid shaped diffusion profiles, as illustrated in Figure 3.2. The diffusion tensor $\mathbf{D}$ is a symmetric, positive definite (all eigenvalues ≥ 0) tensor represented by a 3×3 matrix whose scalar elements are denoted by D_{ij} as follows:

$$\mathbf{D} = \begin{pmatrix} D_{11} & D_{12} & D_{13} \\ D_{12} & D_{22} & D_{23} \\ D_{13} & D_{23} & D_{33} \end{pmatrix} \tag{3.4}$$

3.3.2 Fitting the diffusion tensor

Based on the tensor representation of diffusion in each voxel, we can now rewrite Equation (3.3) to incorporate the diffusion tensor as follows:

$$S(\mathbf{g}_i) = S_0 \exp\left(-b\ \mathbf{g}_i^T\ \mathbf{D}\ \mathbf{g}_i\right) \tag{3.5}$$

where $D_i = (\mathbf{g}_i^T\ \mathbf{D}\ \mathbf{g}_i)$ is the ADC value in direction $\mathbf{g}_i$ and $\mathbf{D}$ is the diffusion tensor. Since the diffusion tensor is symmetric we need at least 6 diffusion measurements $S(\mathbf{g}_i)$ in each voxel to reconstruct it. To estimate the diffusion tensor

$\mathbf{D}$, we can take N instances of Eq. (3.5), one for each gradient direction $\mathbf{g}_i$, and rewrite them in matrix form as follows:

$$\mathbf{s} = \mathbf{B}\ \mathbf{d} \tag{3.6}$$

where $\mathbf{s}$ is a vector containing the signal values for each gradient direction, $\mathbf{B}$ is the B-matrix containing direction-dependent b-value information, and $\mathbf{d}$ is a vector containing the 6 unique values of the diffusion tensor. If $N = 6$, Equation (3.6) has the exact solution $\mathbf{d} = \mathbf{s}\ \mathbf{B}^{-1}$. To reduce noise effects, however, DTI scans commonly apply gradients in more than 6 directions. In this case, solving the matrix equation becomes a least-squares problem:

$$\mathbf{d}' = (\mathbf{B}^T\mathbf{B})^{-1}\mathbf{B}^T\ \mathbf{s} \tag{3.7}$$

3.3.3 Tensor eigenanalysis

As described in earlier sections, measurement of the 3D shape of water diffusion in living tissues can provide information about the underlying tissue structure. In pure water, or when tissue structure is randomly oriented in a voxel, the diffusion tensor will have spherical shape, what we call isotropic diffusion. In voxels where tissue structure has a single preferred orientation, e.g., in muscle tissue or certain fiber bundles of the brain white matter, the diffusion tensor will have an elongated, ellipsoid shape. In the living brain, more complex fiber configurations arise as well, such as crossing, kissing, bending and diverging fibers. Unfortunately, the 2nd-order diffusion tensor cannot distinguish between such complex configurations and will tend to a disk-like or planar shape for all such configurations. The different shapes are illustrated in Figure 3.2. Another thing to bear in mind is that the diffusion tensor only specifies an orientation, not a direction.

From the above, it is clear that tensor shape is a meaningful property. For the purpose of analysis of diffusion tensor data, several geometric measures exist to quantify tensor shape [164] and many of them rely on tensor eigenanalysis. This involves finding eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ and eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ such that the following equation is satisfied:

$$\mathbf{D}\ \mathbf{e}_i = \lambda_i\ \mathbf{e}_i \quad , i \in \{1, 2, 3\} \text{ and } \mathbf{e}_i \neq \mathbf{0},\ \lambda_i \geq 0 \tag{3.8}$$

where $\mathbf{D}$ is the diffusion tensor. Since the null vector is not allowed, Equation (3.8) can be rewritten as $(\mathbf{D} - \lambda_i\ \mathbf{I})\mathbf{e}_i = 0$, where $\mathbf{I}$ is the identity tensor. This implies that the matrix $\mathbf{D} - \lambda_i\ \mathbf{I}$ is singular and its determinant is zero. Each eigenvalue λ_i can then be calculated from:

$$\det(\mathbf{D} - \lambda_i\ \mathbf{I}) = 0 \tag{3.9}$$

Given the eigenvalues, each eigenvector $\mathbf{e}_i$ can be calculated by solving:

$$(\mathbf{D} - \lambda_i \mathbf{I})\mathbf{e}_i = 0 \tag{3.10}$$

When diffusion anisotropy is linear, we have $\lambda_1 >> \lambda_2 \approx \lambda_3$, and the main diffusion direction is clearly defined by the largest eigenvector λ_1, which is commonly called the *main eigenvector* or MEV. In case of planar anisotropy, we have $\lambda_1 \approx \lambda_2 >> \lambda_3$ and there is no main diffusion direction.

3.3.4 Tensor shape and anisotropy measures

As explained in the previous paragraphs, the 3D tensor shape can provide useful information about the underlying tissue structure. However, the second-order tensor is a 6-dimensional object which is difficult to present in a quantitative manner to a user. For this reason, several shape metrics have been proposed that simplify the tensor to a single scalar value. One of the most popular metrics is *fractional anisotropy* or FA. This metric describes the normalized variance of tensor eigenvalues [7] and is defined as follows:

$$\begin{aligned} FA &= \frac{1}{\sqrt{2}} \frac{\sqrt{(\lambda_1-\lambda_2)^2+(\lambda_2-\lambda_3)^2+(\lambda_1-\lambda_3)^2}}{\sqrt{\lambda_1+\lambda_2+\lambda_3}} \\ &= \frac{\sqrt{3}}{\sqrt{2}} \frac{\left|\mathbf{D}-\frac{1}{3}\mathrm{tr}(\mathbf{D})\ \mathbf{I}\right|}{|\mathbf{D}|} \end{aligned} \tag{3.11}$$

where $\mathbf{I}$ is the identity tensor and $\mathrm{tr}(\mathbf{D})$ is the tensor trace. An additional advantage of this measure is that it can be calculated using only the tensor trace and determinant (as indicated in the 2nd equation), without explicitly calculating the eigenvalues first. Westin et al. [164] proposed additional anisotropy measures to distinguish between linear, planar and spherical diffusion:

$$\begin{aligned} C_L &= \frac{\lambda_1-\lambda_2}{\lambda_1+\lambda_2+\lambda_3} \\ C_P &= \frac{2(\lambda_2-\lambda_3)}{\lambda_1+\lambda_2+\lambda_3} \\ C_S &= \frac{3\lambda_3}{\lambda_1+\lambda_2+\lambda_3} \end{aligned} \tag{3.12}$$

Each of these measures has a range between $[0, 1]$ and $C_L + C_P + C_S = 1$. Other measures, such as mean diffusivity ($MD = \mathrm{tr}(\mathbf{D})/3 = (\lambda_1 + \lambda_2 + \lambda_3)/3$) do not describe the shape of diffusion but only its magnitude. For a more complete overview of anisotropy measures we refer the reader to Vilanova et al. [156]. It should be noted, that with a planar or disk-like diffusion profile, fractional anisotropy will remain high. This can have unexpected effects when performing deterministic fiber tracking on the basis of FA thresholds, as will be explained in Section 3.4.

3.4 Fiber tracking

One of the most exciting benefits DTI has to offer for the analysis and understanding of brain white matter, is the reconstruction of 3D pathways by means

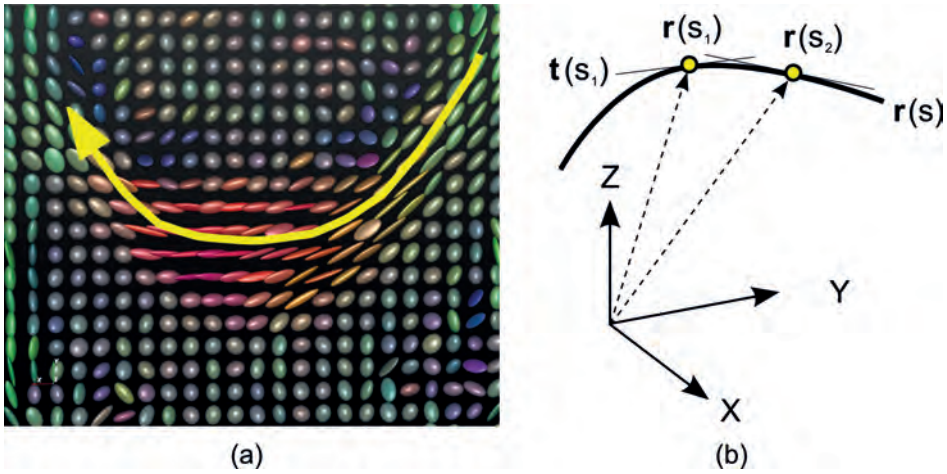

Figure 3.3: (a) Conceptual overview of fiber tracking as tracing pathways through a diffusion tensor field. (b) Mathematical representation of fiber tracking where pathway location is parameterized by a vector $\mathbf{r}$ which is a function of length s along the pathway. The tangent $\mathbf{t}(s_1)$ represents the local fiber orientation estimated from the local diffusion tensor.

of fiber tracking. It involves tracing 3D paths through the diffusion tensor field. Fiber tracking algorithms can be subdivided into two general classes: deterministic and probabilistic ones. Deterministic algorithms, given the same input, always produce the same output. Probabilistic algorithms, on the other hand, introduce an element of randomness into the tracking process in order to simulate the effects of data noise on the resulting pathways. Because they play a special role in this thesis, we defer a discussion of probabilistic algorithms until Chapter 4. In this section, we focus primarily on *streamline tracing*, one of the most common and widely used deterministic tracking algorithms available. Other deterministic algorithms exist, such as the Tensor Deflection (TEND) algorithm [89, 162], but a thorough discussion of these is outside the scope of this thesis.

Streamline tracing was originally designed to trace pathways through vector fields. In case of 2nd-order diffusion tensor fields, it is used to trace paths following local tensor main eigenvectors. Mathematically, a streamline can be represented by a 3D space-curve. The equation describing this curve is based on the fact that the tangent to the streamline must be parallel with the underlying fiber orientation, as estimated by the tensor MEV. Location on the streamline $\mathbf{r}$ is a function of the arc length s (the distance along the streamline measured from the starting point). The evolution of the streamline is described by the following differential equation:

$$\frac{\mathrm{d}\mathbf{r}(s)}{\mathrm{d}s} = \mathbf{e}_1(\mathbf{r}(s)) \tag{3.13}$$

where $\mathbf{e}_1(\mathbf{r}(s))$ is the tensor MEV at position $\mathbf{r}(s)$ [12]. This equation is solved by the following integral:

$$\mathbf{r}(s) = \int_0^s \mathbf{e}_1(\mathbf{x})\ \mathrm{d}\mathbf{x} \tag{3.14}$$

It should be noted that $\mathbf{e}_1$ only specifies an orientation and not a (signed) direction. Implementation of streamline techniques requires 3 basic elements: (1) definition of seed points, (2) path integration and (3) definition of stopping criteria.

1. **Seed points** - These define the initial conditions for solving the streamline integration and represent the 3D points where streamline tracing should start. They are commonly defined by the user but they can also be generated automatically, e.g., by placing a seed point in each non-empty voxel (whole-brain seeding). User-defined seeds are specified by 2D or 3D regions of interest (ROI) whose interior is sampled to generate a set of seed points. These can be evenly spaced [155] or randomly positioned.
2. **Numerical integration** - Several numerical schemes exist to implement Equation (3.14), using Euler forward, 2nd- or 4th-order Runge-Kutta interpolation.
3. **Stopping criteria** - These prevent the algorithm from tracing into regions where the vector field is not reliably defined. In DTI, this can occur in areas of planar or isotropic diffusion, such as brain gray matter. In this case, calculations of the main eigenvector $\mathbf{e}_1$ are not well-defined and depend largely on noise and round-off errors. Stopping criteria are usually based on some anisotropy measure such as FA or C_L. In many cases, thresholds on maximum pathway curvature are also used.

3.4.1 Resolving multiple fiber orientations

In the previous paragraphs we have focused primarily on the diffusion tensor model. It is the most widely used model in clinical settings. A disadvantage of this model is its inability to describe voxels with multiple fiber populations (Figure 3.4). To accurately capture the diffusion profile arising from such complex fiber configurations, many more gradient directions are required (> 100). Such imaging techniques are collectively called High Angular Resolution Diffusion Imaging, or HARDI, and were first introduced by Tuch et al. [151].

Whereas DTI reconstructs a Gaussian distribution function for each voxel, HARDI reconstructs more complex distribution functions, such as illustrated in Figure

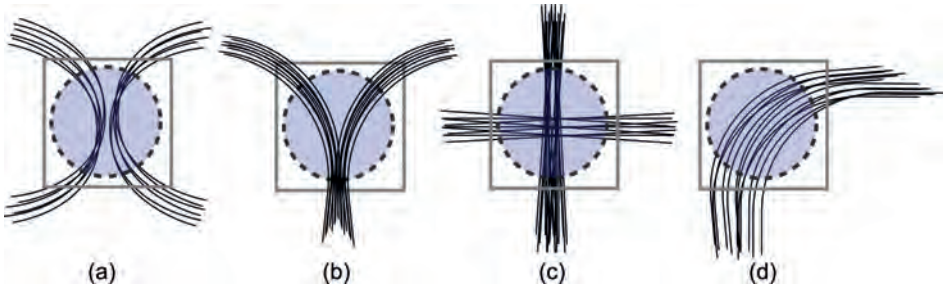

Figure 3.4: Examples of voxels containing complex fiber structures and which are represented by a planar diffusion tensor (blue circle): (a) kissing, (b) splitting, (c) crossing and (d) bending.

3.8(a) and (b). The angular portions of these functions are conveniently described by spherical harmonics (SH). For this reason, they have been widely used in the literature to model the apparent diffusion coefficient [3, 40, 119, 41, 68]. Despite the advantages of HARDI, it requires hundreds of diffusion measurements per voxel as well as complex post-processing steps. Until now, this makes it less practical for clinical settings. In this thesis we will primarily focus on DTI.

3.5 Visualization strategies

In the previous sections we have described the basic concepts behind diffusion tensor analysis and fiber tracking. This discussion is mainly concerned with *data*, such as tensors or pathways consisting of connected points. In this section we present a short overview of techniques for *visualizing* this data on graphical display devices. We will start with visualization techniques for tensor data, after which we proceeed with tensor-derived scalar measures. We conclude the discussion with strategies for visualizing the fiber pathways generated by fiber tracking.

3.5.1 Tensor data

Second-order diffusion tensors are complex, 6-dimensional objects. One of the few methods to display them without loss of information, is *glyph visualization*. Glyphs are compact graphical objects located at specific points in space which use a combination of color, size, shape and texture to encode properties in the underlying data. Pierpaoli and Basser [125] first introduced ellipsoid-shaped glyphs to visualize 2nd-order tensors obtained from DTI. They used glyph size to encode mean diffusivity, while the long axis of the ellipsoid corresponds to the direction of the tensor MEV. Several extensions of this idea have since been proposed [87, 79]. An example are superquadrics which are a family of shapes defined by similar formulas as ellipsoids except for using arbitary powers instead of squares. Box-shaped glyphs, ellipsoids and superquadrics are illustrated in Figure 3.5.

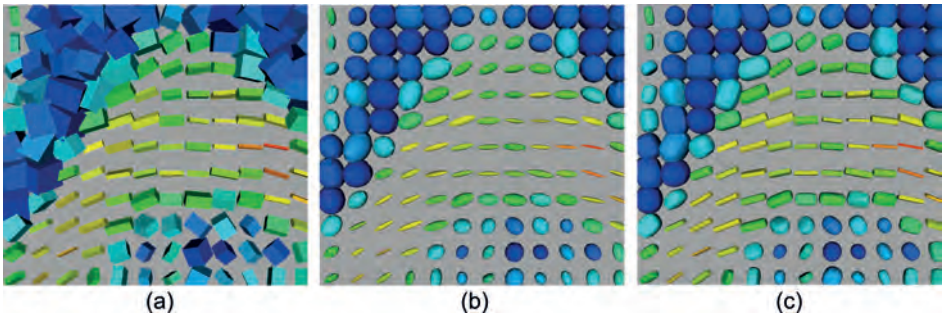

Figure 3.5: Examples of tensor glyph shapes: (a) boxes, (b) ellipsoids and (c) superquadrics.

Tensor glyphs have several disadvantages. Usually, a glyph is placed in each data point. If many data points are shown the density of glyphs will lead to cluttered visualizations and make it difficult to distinguish between the different visual properties of each glyph. Also, glyphs can only represent local information at discrete data points which makes it difficult to infer global information from them.

3.5.2 Tensor eigenvectors

As explained earlier, the tensor MEV is an estimate of the main diffusion direction and the underlying fiber orientation. A popular approach to visualize fiber orientation in 2D maps is to convert the tensor MEV $\mathbf{e}_1$ to an RGB color. This can be done as follows:

$$\begin{aligned} R &= |\mathbf{e}_1 \cdot \mathbf{x}| \\ G &= |\mathbf{e}_1 \cdot \mathbf{y}| \\ B &= |\mathbf{e}_1 \cdot \mathbf{z}| \end{aligned} \tag{3.15}$$

where $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ defines the coordinate frame of the acquired dataset. To reduce the visual prominence of isotropic regions, color saturation or opacity can be modulated on the basis of a tensor scalar measure, such as FA. It should be noted that manipulation of color saturation is best performed in HSV space (hue, saturation, value) instead of RGB space.

3.5.3 Tensor-based scalar measures

As explained earlier, tensor-derived scalar measures often allow for easier interpretation of the data. Since these measures are scalar-valued, standard techniques for 2D and 3D scalar field visualization can be used. Figure 3.6(a) and (b) show examples of such scalar maps for mean diffusivity (MD) and fractional anisotropy

(FA). Figure 3.6(c) illustrates how multiple scalar measures can be incorporated into a single 2D image by assigning different colors to different measures.

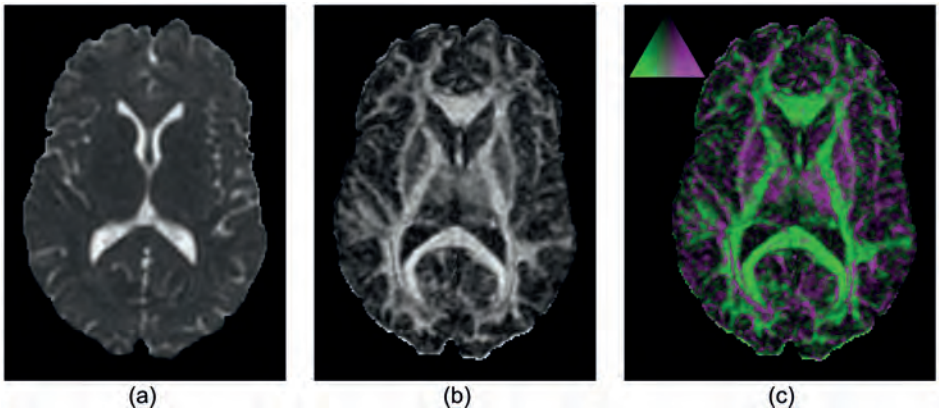

Figure 3.6: Scalar maps of (a) mean diffusivity, (b) fractional anisotropy and (c) combination of C_L (green) and C_P (magenta) (see also Equation (3.12).

Scalar fields can also be visualized in 3D using standard volume rendering techniques. Using iso-surface raycasting it is possible to visualize 3D surfaces of equal anisotropy, possibly color-coded on the basis of the underlying main eigenvector orientation [80]. Another approach was suggested by Wenger et al. [163] who use volumetric line integral convolution (LIC) to create a 3D texture of thread-like structures on the basis of tensor MEV's. This 3D texture is then visualized using direct volume rendering. An additional 3D texture with the same, but thicker, thread structures is used to create halos around the threads in the first volume. This provides an impression of shadowing and allows the threads to be better distinguishable.

3.5.4 Deterministic fiber pathways

The output of fiber tracking is generally represented as 3D pathways consisting of connected points. Many different visualization strategies exist for such data. First, we can distinguish between different shapes used to represent each path, such as streamlines or streamtubes. Streamlines are the simplest approach but, since they have undefined surface normals, it is not possible to perform standard illumination calculations on them. As a result their 3D shape and orientation can be difficult to perceive. Streamtubes do allow straightforward illumination techniques because they have well-defined surface normals. An example of a streamtube visualization for standard fiber tracking is given in Figure 3.7(a). Figure 3.7(b) shows streamtubes where the radius depends on some underlying data property, such as fractional anisotropy. Hyperstreamlines are an extension

of streamtubes which are specifically tailored for 2nd-order tensor fields [39, 172]. Hyperstreamlines have varying cross-sections depending on the local tensor eigenvectors and -values.

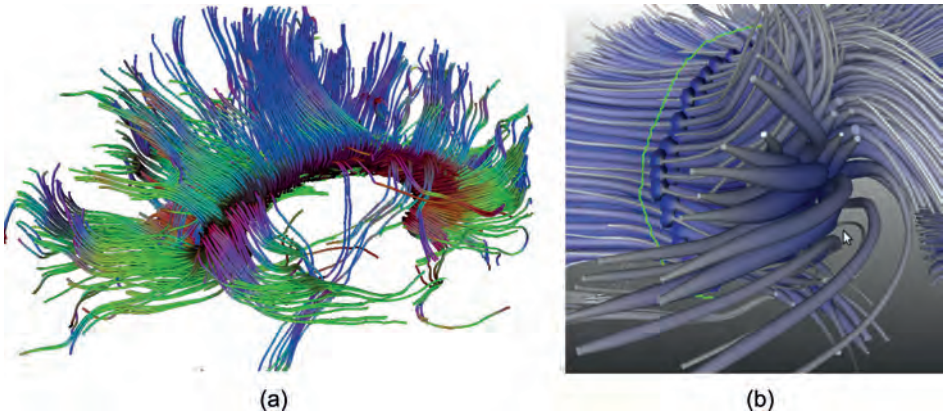

Figure 3.7: (a) Corpus callosum represented by streamtubes color-coded with main eigenvector orientation. (b) Streamtubes with fractional anisotropy encoded in streamtube radius.

Apart from a distinction between different *shapes* used to represent fiber pathways, there also exist different approaches for *rendering* these shapes. A common technique used to improve the orientational nature of pathways is to encode local (x, y, z) orientation with RGB colors. Figure 3.7(a) provides an example of this. Such RGB coloring can be applied to streamlines, streamtubes and hyperstreamlines. Because of undefined surface normals, streamlines cannot be easily illuminated to improve their depth perception. Even so, they are light-weight graphics primitives which can be rendered in great numbers without serious performance issues. This is more problematic with streamtubes or hyperstreamlines whose geometry is much more complex. To overcome the illumination problems of streamlines, Peeters et al. [124] propose an alternative method to calculate normals perpendicular to the local streamline direction. Using GPU-acceleration techniques this allows them to interactively apply hair-like rendering effects which significantly improve spatial depth perception.

3.5.5 HARDI visualization

HARDI data can be visualized using glyphs. However, the complex shape of the corresponding diffusion profile makes standard geometry-driven rendering unsuitable in terms of performance. Peeters et al. [123] proposed a GPU-accelerated ray casting method which allows rendering of glyphs defined by spherical harmonics at interactive framerates. Figure 3.8(a) and (b) illustate examples of this

approach. This method can also be used to render DTI tensor glyphs such as described in Section 3.5.1.

3.6 Conclusions

In the previous sections we have discussed the basic principles of diffusion tensor imaging and fiber tracking. We also gave an overview of the 'traditional' methods of visualizing DTI data and fiber tracking results. These techniques do not take errors and uncertainty into account however. In the next chapter we will address these issues in more detail and discuss existing methods for the analysis of uncertainty in DTI and fiber tracking.

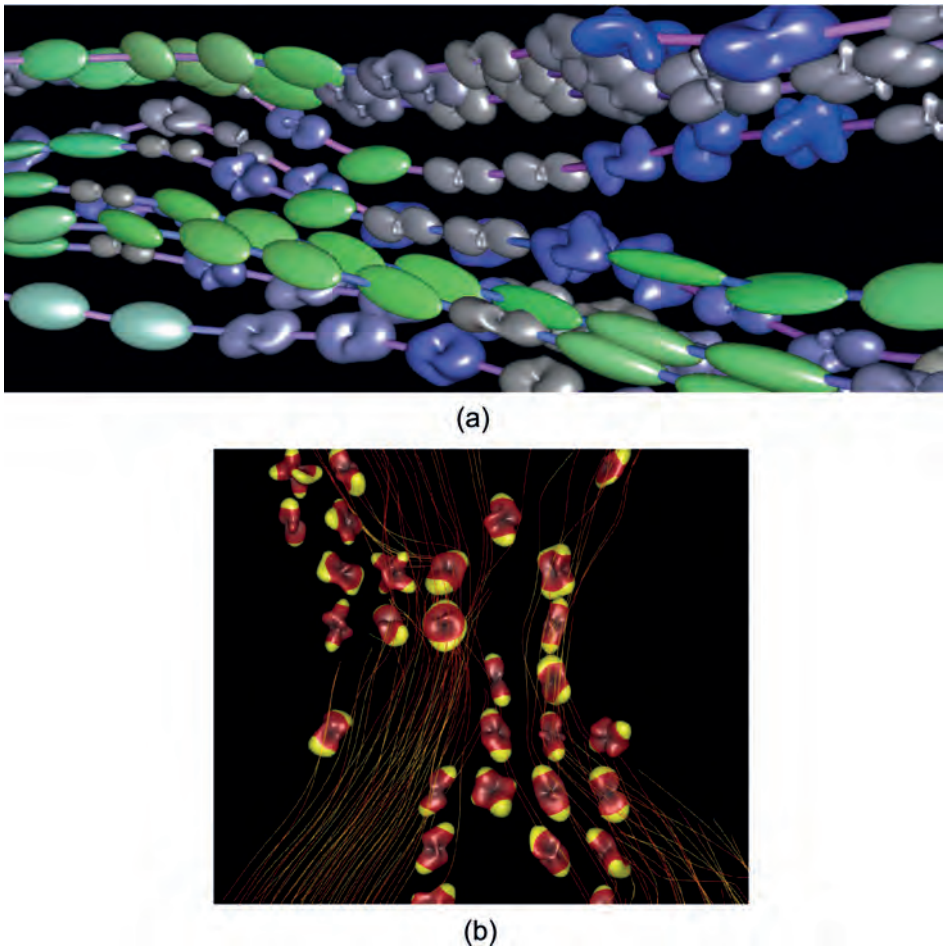

Figure 3.8: (a) Fused visualization of streamlines and glyphs where 2nd-order tensor glyphs are rendered in regions of linear anisotropy and HARDI glyphs are rendered in regions of complex fiber architecture. (b) Another example of combining streamlines with HARDI glyphs [129].

The creator of the universe works in mysterious ways. But he uses a base ten counting system and likes round numbers.

- Dilbert

4 Uncertainty analysis in DTI and fiber tracking

4.1 Introduction

In this thesis we deal primarily with the *visualization* of uncertainty. This means that we have to rely on existing methods for the *analysis* of uncertainty in DTI and fiber tracking. This has been an active area of research ever since DTI and fiber tracking became popular. The purpose of this chapter is to provide an overview of such analysis methods. The first part of this chapter describes the different sources of uncertainty present in the DTI imaging and processing pipeline. This overview will serve as a basis for further discussion of uncertainty in this thesis. In the second part of this chapter we discuss probabilistic fiber tracking methods. These methods attempt to simulate the effects of image noise on the tracking output by incorporating an element of randomness in the tracking process. Many such algorithms have been proposed and we will provide an overview of the main types. As will become clear from this chapter, not all sources of uncertainty can be quantified in a manner that allows their analysis and visualization. As a result, many sources of error and uncertainty remain unexplored.

4.2 Definition of uncertainty

Before we continue, we should more precisely explain what we mean with 'uncertainty'. The term is rather vague and has different meanings in different application domains. It can apply to predictions of future events, to physical measurements already made or simply to the unknown. In this thesis we do not wish to add to the confusion and introduce completely new definitions. We simply consider uncertainty as a state of ignorance about the true outcome of a measurement or computational process. For example, DTI and fiber tracking are used to make inferences about white matter pathways on the basis of physical measurements of water diffusion. Uncertainty arises from our ignorance about the true outcome of fiber tracking, that is, the difference between the reconstructed fiber pathways and the underlying white matter anatomy. In this definition, uncertainty arises as a consequence or result of something else, which we will call *sources* of uncertainty. Examples of sources of uncertainty are image noise, magnetic distortions, partial volume effect and model errors. These and more will be discussed in detail in the next sections.

4.3 The DTI uncertainty pipeline

In order to discuss the different sources of uncertainty affecting the output of fiber tracking we propose a 4-stage pipeline, which we call the *DTI uncertainty pipeline*. This pipeline consists of 4 stages, each responsible for a specific function. The four stages are (1) image acquisition, (2) diffusion modeling, (3) fiber tracking and (4) visualization. A diagram of the uncertainty pipeline is illustrated in Figure 4.1. Besides performing its intended function, each stage is also a source

of potential errors and uncertainties.

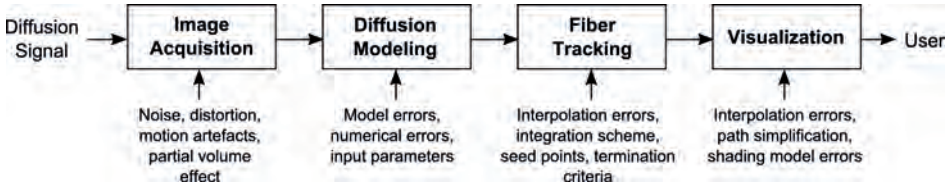

Figure 4.1: The DTI uncertainty pipeline consisting of four stages: (1) image acquisition, (2) diffusion modeling, (3) fiber tracking and (4) visualization.

Before we discuss the individual stages and the associated uncertainties, we need to say something about the *input* of the DTI uncertainty pipeline: the diffusion signal. Even though this does not constitute a 'processing' stage in any sense, there is a fundamental type of uncertainty associated with the interpretation of this signal. As explained earlier, the primary goal of fiber tracking is to gain insight into the 3D layout of brain white matter. What is being measured with DTI, however, is diffusion of water. The presence and orientation of white matter fibers is merely *inferred* on the basis of anisotropic diffusion of water molecules. Even though the correspondence between diffusion and fiber orientation has been validated to a certain extent in muscle tissue [154, 153], spinal cord tissue [29, 103] and phantom studies [98, 150], the exact physiological basis of diffusion anisotropy and its relation to fiber orientation remains a topic of debate. Different components of the white matter axons contribute to diffusion anisotropy, such as the cytoskeletal lattice of neurofilaments and microtubules, axonal membranes and myelin. However, there is still much debate about the individual contributions of these components to the overall diffusion signal measured in each voxel [9].

We will now proceed with a discussion of the different pipeline stages and the uncertainties affecting their outputs.

4.3.1 Image acquisition

DTI is based on magnetic resonance and, for this reason, suffers from many of the same image artifacts. However, typical DTI pulse sequences take more time, are much more complex and involve strong magnetic gradients that are constantly being switched on and off. This enhances typical MRI artifacts and gives rise to additional, DTI-specific artifacts. We will not provide an exhaustive list of these artifacts but instead focus on those that have been investigated in the context of fiber tracking and visualization of its output.

Noise - The relatively long duration of DTI pulse sequences results in increased signal loss and decreased signal-to-noise ratio (SNR). The effect of noise on fiber

tracking output has been widely investigated in the neuro-imaging literature [88, 4, 65, 150] and has given rise to a wide range of probabilistic tracking algorithms that attempt to take image noise into account [11, 50, 122]. We will defer an extensive discussion of probabilistic tracking methods to Section 4.4.

Partial volume effect - Partial volume effect, or PVE, results in loss of information and detail due to averaging of tissue properties inside a single image voxel. It leads to contamination of the diffusion measurement, especially in voxels located at the boundary of two or more tissue types, e.g., white matter and cerebral spinal fluid. Due to the low image resolution of DTI, PVE is generally more pronounced than in conventional MRI. The effect of PVE on deterministic fiber tracking has been investigated by a number of authors [59, 72, 158].

Motion artifacts - With the growing trend to increase the number of gradient directions in DTI, acquisition time increases as well. This makes it more likely that subjects move during the scan which results in mis-alignment of the images. Fortunately, mis-alignments involve only rigid-body transformations and can be corrected easily using image registration. However, care should be taken to also re-align the gradient encoding vectors (stored in the **B**-matrix) before calculating the tensor and performing fiber tracking [92].

Magnetic distortions - The quality of spatial encoding of MR images relies on the homogeneity of the static magnetic field. However, differences in magnetic susceptibility between adjacent tissues types can cause disturbances in the magnetic field and thereby result in image distortions. These can be quite severe, for example wherever brain tissue is close to air-filled cavities in the head. Susceptibility artifacts can be corrected to some extent but this is not common practice and the correction algorithms can introduce biases of their own [72]. The effect of susceptibility distortion on fiber tracking, in some cases, may result in fibers ending up in the wrong sulci [5, 91]. Eddy currents are another source of image distortion but these can be easily corrected on a slice-by-slice basis.

Scanner settings - DTI pulse sequences are configured on the basis of many parameters, such as number of gradient directions, b-value, number of non-diffusion-weighted images, etcetera. In principle, these settings are fixed between different scan sessions. However, the effect of changes in these parameters (e.g., when comparing data between different institutions) can be difficult to predict. Also, some drift in these parameters can be expected due to warming up of the scanner. Especially the choice of the b-value can have a large impact on the resulting data. It is defined as follows:

$$b = \gamma^2 \, G^2 \, \delta^2 \, \left(\Delta - \frac{\delta}{3} \right) \tag{4.1}$$

where γ is the gyro-magnetic ratio of water, G is the gradient amplitude, δ is the

gradient pulse duration and Δ is the time between consecutive gradient pulses. Pulse duration also determines the diffusion time. For short durations, molecules will not have had time to run into physical barriers. The measured diffusion profile will therefore be isotropic. Conversely, long durations lead to increasing numbers of molecules hitting physical barriers. This results in more pronounced diffusion anisotropy in fibrous tissues.

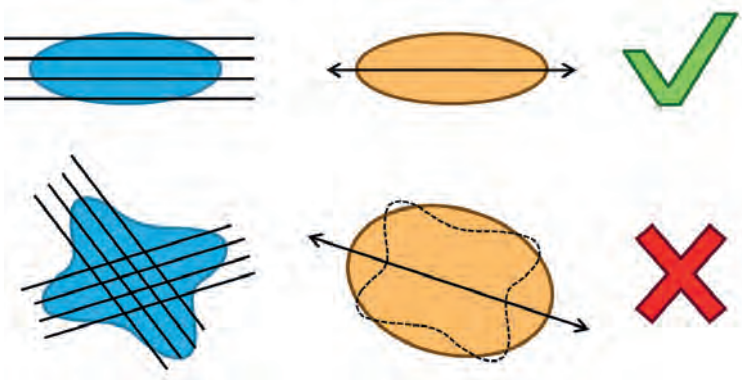

Figure 4.2: Diffusion profiles for single-orientation fibers (top row) and crossing fibers (bottom row). Crossing fiber populations are incorrectly described by the 2nd-order tensor model. As a result, the direction of main diffusion as predicted by this model does not correspond with either fiber population.

4.3.2 Diffusion modeling

DTI can only measure diffusion in a single direction at a time. To determine the direction of *greatest* diffusion, many measurements have to be performed for each voxel along many different gradient directions. To deal with such high-dimensional data mathematical models are fitted to the data in each voxel. This reduces the set of data values to its essential characteristics. In DTI a 2nd-order tensor model is used for this purpose. However, this model can describe only a *single* average fiber orientation per voxel. If fibers are crossing or diverging inside a given voxel, the 2nd-order tensor model will be incorrect and result in unreliable results. This is illustrated in Figure 4.2. Even if the tensor model is correct, variations can arise from the fitting procedure. Ordinary least squares fitting is most commonly applied but other regression procedures exist that give slightly different results [83, 84].

4.3.3 Fiber tracking

As explained in Chapter 3, fiber tracking involves the integration of voxel-wise fiber orientations into a pathway that connects different brain regions [12]. Several methods to implement this exist but the most popular and widely used method remains streamline tracing. Each streamline is described by a differential equation that allows calculation of a position on the streamline by integrating all previous position changes. However, the iterative nature of the numerical solvers, makes streamline tracing highly sensitive to local errors in the underlying vector field, which in case of DTI is the tensor MEV. As tracking proceeds away from the starting point, directional errors accumulate. As explained in Chapter 3, each fiber tracking algorithm consists of 4 basic elements: (1) numerical integration, (2) interpolation, (3) stopping criteria and (4) seeding and filtering. The purpose of these elements have already been described earlier. Here, we will focus on the uncertainties arising from them.

Numerical integration - Numerical integration in fiber tracking can be implemented in different ways, such as Euler forward integration [110] and 2nd- or 4th-order Runge-Kutta integration [8]. The quality and smoothness of pathways also depends on the integration step size [150]. TEND, or tensor deflection, uses the full diffusion tensor to guide the direction of step vectors. If the tensor is planar or isotropic, tracking proceeds in a straight line [89]. This method may be less sensitive to noise and will therefore result in slightly different pathways.

Interpolation - Also, when calculating the next integration step vector, some form of interpolation is needed when the current position lies in between data points. Different options will lead to different results, e.g., nearest-neighbor, linear [110] or higher-order interpolation [8, 2] of (1) the main eigenvector [110] or (2) the tensor itself [8, 2]. The different choices that can be made for any given implementation of fiber tracking can result in considerable overall differences between the output of different algorithms, even if the starting conditions (data, thresholds, seed points) are kept constant. The variability may have a negative effect on the confidence users have in the reliability of fiber tracking, especially in risk-related situations such as neurosurgery [27]

Stopping criteria - Another major source of variation, between different users or different sessions, are parameter settings such as stopping criteria. Many fiber tracking algorithms are highly sensitive to small changes in the associated thresholds, meaning that a small change in threshold value can lead to a large change in pathway length [150]. We discuss this issue in more detail in Chapter 6 where we present techniques for visualizing the effect of parameter changes on the fiber tracking result.

Seeding and filtering - Seed points define the initial conditions for numeri-

cal integration. In most cases, seeding has to be done manually by the user, thereby introducing user bias. Bias can be reduced by detailed written instructions [33, 159] or combining seed regions with filter regions in different Boolean combinations (AND, OR, NOT) [65, 60]. The purpose of seeding and filtering is to extract specific fiber bundles, e.g., the motor part of the corticospinal tract (CST) or the optic radiation. An alternative approach could be to generate pathways on the basis of whole-brain seeding and apply a pathway clustering algorithm to selectively subdivide the total set of pathways into subsets that, hopefully, represent anatomically recognizable fiber bundles [157, 26]. A similar result might be achieved by using a pre-classified white matter atlas and warping the current dataset to it in order to find out which pathways belong to which bundle [114].

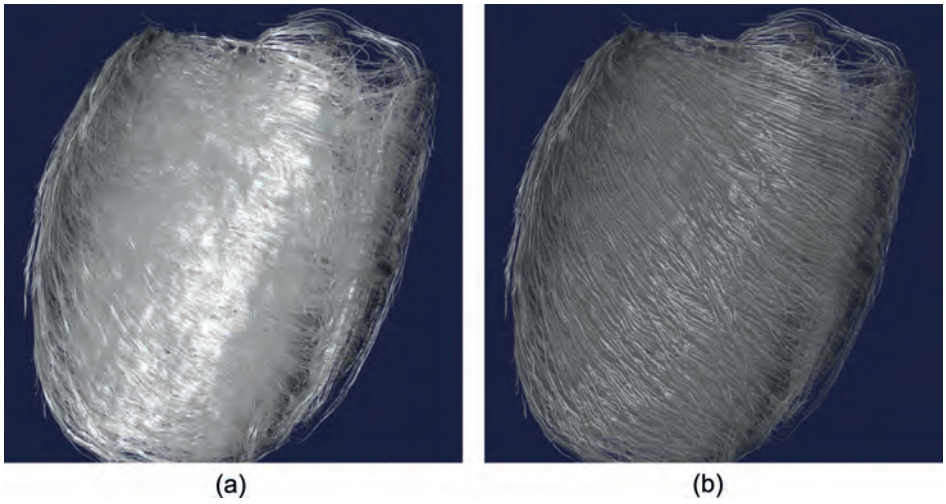

Figure 4.3: GPU-based streamline illumination and shadowing [123]. (a) Hair-like illumination of muscle fibers in an ex-vivo heart. (b) Similar illumination, now with shadowing applied to improve depth perception of individual fibers.

4.3.4 Visualization

Even though we use visualization techniques to show uncertainty, it can be a source of errors and uncertainty itself. For example, the pathways calculated from fiber tracking may be simplified (by removing vertices) for the purpose of rendering performance. Also, different lighting models can be chosen to render fiber pathways. Lighting (and shadowing) plays an important role in depth perception. Many existing fiber tracking tools do not apply lighting to the fiber pathways or rely on very simple lighting models. Photo-realistic rendering attempts to simulate real-world lighting as closely as possible but the algorithms involved are computationally expensive and often not suitable for rendering at

interactive framerates. One counter-example is the work proposed by Peeters et al. [123] who present a GPU-based rendering algorithm for applying lighting and shadowing to streamlines. An example rendering is illustrated in Figure 4.3.

4.4 Probabilistic fiber tracking

As explained in the beginning of this chapter, we rely on existing methods for the quantification and characterization of selected sources of uncertainty. Probabilistic fiber tracking algorithms address a widely investigated source of uncertainty, namely *image noise*. In this section we give a brief overview of such algorithms. Whereas deterministic algorithms typically produce a single estimate of orientation for each fiber population in a voxel, probabilistic algorithms derive a *distribution* of fiber orientations in each voxel based on an estimation of noise. For multiple fiber populations, each with a different average orientation, such a distribution can be multimodal. Probabilistic tracking algorithms can roughly be subdivided into the following categories:

- Tensor-derived uncertainty
- Bootstrapping
- Bayesian modeling

4.4.1 Tensor-derived uncertainty

Tensor-based approaches rely on a single DTI scan and calculate a 2nd-order diffusion tensor for each voxel. Anisotropy and tensor shape properties are then used to derive probability distribution functions (PDF's). One of the earlier methods was proposed by Koch et al. [85] and relied on Monte Carlo random walks, where fiber tracking proceeds from voxel to (neighboring) voxel without interpolation. The probability of choosing a particular neighboring voxel depends on the angle between the tensor main eigenvector in the initial voxel and the vector connecting the initial voxel with the neighboring voxel. Another approach for determining step vectors along a fiber pathway is to assume the main eigenvector in each voxel to be a correct average but perturb its direction based on the linear anisotropy of the underlying tensor [20, 122]. Monte Carlo random walks have also been used to simulate random particle movement [58]. To determine each step along the particle's path, a random unit vector is selected from a uniformly distributed orientation distribution function (ODF) and multiplied by the underlying diffusion tensor. This approach is similar to the tensor deflection method proposed by Lazar et al. [89].

4.4.2 Bootstrapping

A disadvantage of the tensor-shaped methods is that they are biased toward the tensor originally fitted to the data. The tensor MEV is taken to be a correct average and uncertainty arises only from variations in orientation about the MEV. A

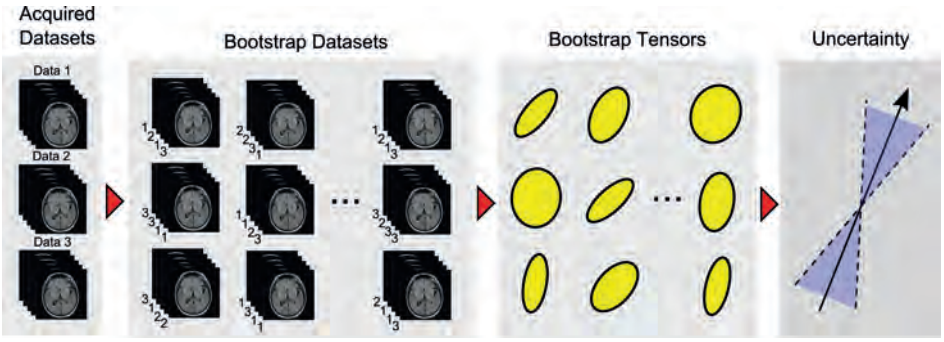

Figure 4.4: Example of bootstrap method: three datasets are acquired and used to generate a large number of bootstrap datasets which consist of random slice selections with replacement from the original datasets. For each bootstrap dataset a tensor dataset is created. In each voxel a cone of uncertainty can be derived on the basis of main eigenvector variations across all bootstrap tensors for that voxel [12].

better, and less biased, approach to characterize uncertainty due to noise would be to repeat the measurement many times. However, due to time constraints, diffusion MRI does not allow acquisition of the hundreds of datasets that would be required for a robust estimation of uncertainty. Bootstrapping is a well-known technique in statistics to estimate the distribution of a given statistic on the basis of only a small set of samples. In DTI, bootstrapping has been used to characterize uncertainty in various tensor scalar measures [120] as well as fiber orientation [70].

Regular bootstrap - The general principle of DTI bootstrapping is illustrated in Figure 4.4. Based on a small set of image acquisitions, a large set of bootstrap data sets is created by randomly selecting slices, with replacement, from the original datasets. For this, consider a single DTI dataset consisting of N diffusion measurements per voxel V, one measurement per gradient direction. Bootstrapping involves repeating the diffusion-weighted image (DWI) scan M times thereby resulting in M realizations V_i^* $(0 < i \leq M)$ for each original voxel V. We can now construct a new voxel W by, for each gradient direction $\mathbf{g}_j$ $(1 < j \leq N)$, randomly selecting a voxel V_i^*, with replacement, and taking the value corresponding to $\mathbf{g}_j$. The new voxel can be used to re-fit the diffusion model (whether 2nd-order tensor or high-order, multi-fiber models). Assuming a 2nd-order tensor model, repeating this procedure many times allows us to approximate the distribution of tensors at each voxel position.

Wild bootstrap - Regular bootstrapping provides an intuitive way to estimate uncertainty in diffusion tensors. However, it requires at least 4 to 5 repeats of the DWI scan. In clinical settings, this can take too much time, especially with current trends for using medium to high angular resolution imaging (e.g. > 32

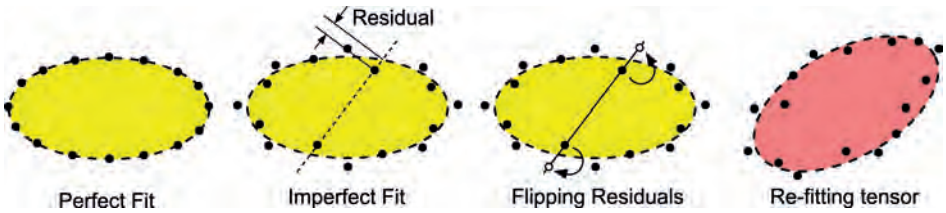

Figure 4.5: Wild Bootstrap: In the ideal case we have a perfect model fit. In practice, the fit is imperfect leading to model residuals (differences between measured and model-predicted signals). Residuals are randomly flipped, after which the tensor model is fitted again. This process is repeated many times to obtain a large collection of tensor datasets [71].

directions). Wild bootstrapping requires only a single DTI scan at the expense of certain simplifying assumptions [71, 165]. The general principle behind Wild Bootstrapping is illustrated in Figure 4.5. It works as follows: a given voxel V consists of N diffusion measurements. As explained in Chapter 3, the vector $\mathbf{d}'$ containing the 6 unique elements of the diffusion tensor $\mathbf{D}$ can be estimated (using ordinary least squares optimization) as follows:

$$\mathbf{d}' = (\mathbf{B}^T\mathbf{B})^{-1}\mathbf{B}^T\ \mathbf{s} \tag{4.2}$$

where $\mathbf{B}$ is the B-matrix. Once we have $\mathbf{d}'$ we can compute the model-predicted signal values $\mathbf{s}' = \mathbf{B}\ \mathbf{d}'$ and subtract these from the original signal values $\mathbf{s}$ to obtain a vector $\mathbf{r} = (\mathbf{s}' - \mathbf{s})$ of residuals. The Wild bootstrap procedure involves repeating the estimation of $\mathbf{d}'$ in each voxel multiple times as follows:

$$\mathbf{d}' = (\mathbf{B}^T\mathbf{B})^{-1}\mathbf{B}^T\ (\mathbf{s} + \mathbf{q}(\mathbf{r})) \tag{4.3}$$

where the function $\mathbf{q}$ randomly multiplies each separate residual of $\mathbf{r}$ by 1 or -1. By perturbing residuals randomly each tensor fit will be slightly different from the previous one. The end result of applying the Wild Bootstrap procedure K times for all voxels in the original dataset, is a set of K tensor volumes. Conceptually, this amounts to having a dataset where each voxel contains a distribution of tensors. The uncertainty in anisotropy and main eigenvector direction predicted by Wild bootstrap data has been found to be very similar to the original, regular bootstrap method [71]. It should be noted that residuals are flipped for each gradient direction independently. They are *not* exchanged between directions because tensor fitting is done, not directly on the diffusion signal, but on the *log-transformed* diffusion signal. This causes variations in the diffusion signal due to noise to become dependent on direction.

4.4.3 Bayesian modeling

Bootstrapping methods do not require an explicit mathematical formulation of uncertainty. They also require minimal assumptions regarding the estimation of uncertainty. Bayesian methods, on the other hand, provide a formal framework for modeling uncertainty, defining prior assumptions and predicting error propagation. Bayes's Theorem states that the posterior probability distribution of the model parameters Θ, given the data D and the model M, is defined as follows:

$$P(\Theta|D,M) = \frac{P(D|\Theta,M)\ P(\Theta|M)}{P(D|M)} \tag{4.4}$$

A possible disadvantage of Bayesian methods is that they rely on assumptions about the prior and likelihood distributions of model parameters and noise in the data. Bootstrapping methods are less dependent on such assumptions. Even so, the explicit mathematical formulation of assumptions as separate PDF's allows them to be changed and updated as knowledge about the underlying processes increases. Together with the wide range of techniques for solving Bayesian problems, this has made Bayesian inference a popular approach for developing probabilistic fiber tracking algorithms. Different diffusion models have been proposed in this context such as the 'ball-and-stick' model [11], tensor [11, 50, 74], and models supporting multiple fiber orientations per voxel [13, 66, 64].

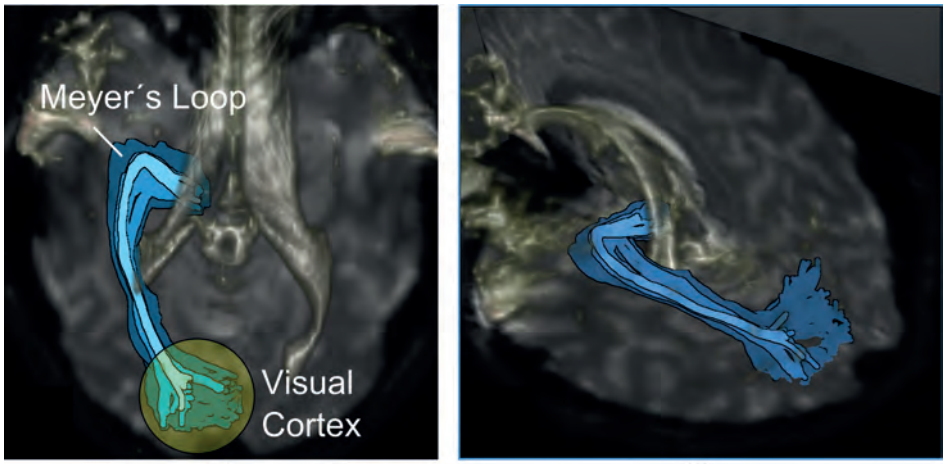

Figure 4.6: Optic radiation: (a) Cranial-to-caudal view indicating Meyer's Loop and visual cortex. (b) Side view of optic radiation.

4.4.4 Hybrid methods - ConTrack

Not all probabilistic fiber tracking methods can be easily classified according to the categorization used in the previous examples. Some algorithms combine elements of Bayesian inference and bootstrapping, such as the ConTrack algorithm proposed by Sherbondy et al. [142]. This algorithm is specifically designed to reconstruct pathways that are known to exist and therefore relies on at least two regions of interest. The main difference between ConTrack and other probabilistic fiber tracking methods is its separation of (1) pathway sampling and (2) pathway scoring. Most algorithms combine these two steps. They tend to ignore pathways evolving in unlikely directions, even though such pathways may be perfectly valid anatomically. This can occur, for example, in the optic radiation. This bundle runs between the geniculate nucleus in the mid-brain towards the visual cortex at the back of the head. However, it does this by first running in lateral-anterior direction and then sharply bending back to continue its path posteriorly. The sharply bending part of the optic radiation is commonly called Meyer's Loop. Many tracking algorithms, even probabilistic ones, have problems reconstructing Meyer's Loop [13, 64, 66]. The optic radiation, including Meyer's Loop, is illustrated in Figure 4.6. The main problem ConTrack tries to address is the failure of many probabilistic algorithms to reconstruct certain fiber pathways that are known to exist.

Pathway sampling - ConTrack performs pathway sampling by generating a large set of potential pathways that have end-points in two ROIs. The directions in which these pathways are allowed to evolve is only loosely constrained. This allows pathways to be sampled which would otherwise have been ignored.

Pathway scoring - Pathway scoring involves computing a confidence value for each pathway sample based on its anatomical validity. A pathway score $Q(s)$ is defined as:

$$Q(s) = q(D|s)\ q(s) \tag{4.5}$$

where s is the pathway and D the data values encountered along the pathway. The data-dependent part of the score determines how well the pathway is supported by the diffusion tensors encountered along its length. For a pathway consisting of N nodes, it can be computed as follows:

$$q(D|s) = \prod_{i=1}^{N} q(\mathbf{D}_i|\mathbf{t}_i) \tag{4.6}$$

where $\mathbf{D}_i$ and $\mathbf{t}_i$ are the diffusion tensor and tangent vector at node i. The term $q(\mathbf{D}_i|\mathbf{t}_i)$ is computed using a Bingham distribution [35]. For details we refer the reader to Sherbondy et al. [142]. The data-independent part of the

score represents prior knowledge about the shape of neuronal connections and is defined as follows:"

$$q(s) = q_{end}(s_1)\ q_{end}(s_N)\ \prod_{i=2}^{N-1} q_{curve}(\Theta_i)\ q_{length}(s_i) \tag{4.7}$$

where $q_{end}(s_i)$ returns 1 if s_i is inside a ROI, and zero otherwise. The function $q_{length}(s_i)$ returns a non-zero value if s_i is in white matter, and zero otherwise. The function $q_{curve}(\Theta_i)$ returns a score depending on the angle of the pathway as it enters and exits a given node. Again, for details we refer the reader to Sherbondy et al. [142].

4.5 Visualization strategies

In the previous chapter we presented a number of 'traditional' approaches for the visualization of DTI tensor data and deterministic fiber tracking. Visualization of the output of probabilistic fiber tracking has received less attention. It is also more difficult because of the large quantity of fiber pathways potentially generated by such methods. Direct visualization of these pathways using, for example, streamlines is mostly ineffective because of the resulting visual clutter. Currently, the most popular method to visualize probabilistic tracking results is to create 3D scalar maps, similar to those described in Section 3.5.3 of the previous chapter. Each voxel stores a fiber density depending on how many fiber pathways intersect it. This provides an indication of the likelihood or probability of each voxel being part of the fiber bundle of interest. These 3D probability maps can then be visualized as 2D color-coded maps or using volume rendering such as illustrated in Figure 4.7.

Even though 2D and 3D fiber density maps are the most popular way to present probabilistic tracking results, they do not provide insight into the 3D layout of fiber pathways and variations in their shape due to uncertainty. The visualization techniques presented in Chapters 6, 7 and 8 are specifically designed to deal with this issue. However, before we proceed with these chapters, we will first present an overview of existing research in uncertainty visualization in the next chapter.

4.6 Conclusions

In the previous sections we have described the different sources of uncertainty arising from the DTI acquisition and processing pipeline. We also outlined existing methods for characterizing uncertainty, especially probabilistic fiber tracking methods which attempt to describe the effects of image noise on the fiber tracking output. In the next chapter we will focus on visualization of these uncertainties.

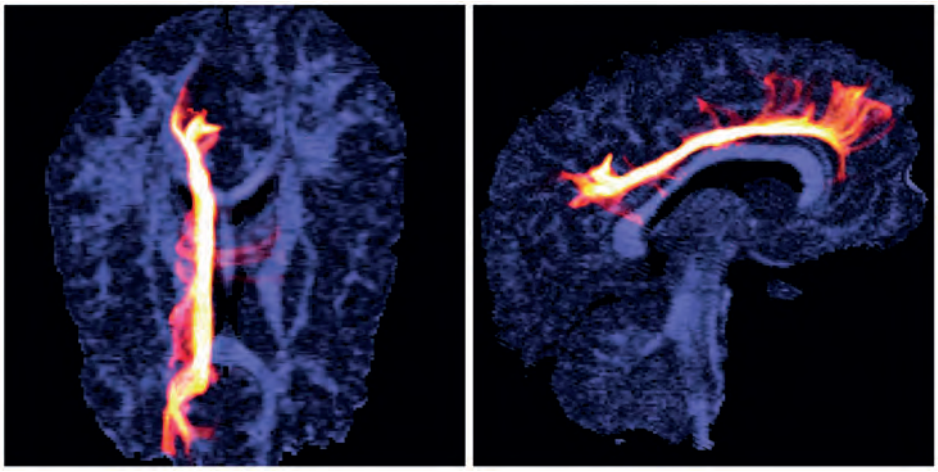

Figure 4.7: Color-coded scalar probability maps based on probabilistic fiber tracking [105].

63% of all statistics are made up... including this one.

- Dilbert

5 Uncertainty visualization

5.1 Introduction

In the previous chapter we have given an overview of uncertainty analysis in DTI and fiber tracking. In this chapter we focus on the visualization of uncertainty and previous work done in this area.

Different ways exist to classify uncertainty visualization techniques. As explained in Chapter 1, uncertainty can arise from a wide variety of sources, only some of which can be characterized in a quantifiable manner. Such quantifiable measures of uncertainty are standard deviations, percentiles, min/max ranges, errors, confidence intervals or complete distributions [121]. Other descriptions of uncertainty are not easily quantified nor measured such as the quality of the data source, the degree of subjective influence in the data recording process or the logical consistency between relations present in the data [54, 102]. As explained in Chapter 1, the lack of knowledge about the true relation between the measured signal (water diffusion) and the physical phenomenon of interest (fiber orientation) is also a source of uncertainty that is difficult to characterize.

Data visualization allows us to analyze large quantities of data by presenting in a way that makes optimal use of the pattern recognition capabilities of the human visual system. This means the data, or measures of its uncertainty, has to be recorded or transformed to a format that allows mapping to visual properties, e.g. color, shape, size or texture. Without such a visual mapping, visualization itself is not possible. For this reason, in our discussion of related work in uncertainty visualization we always consider three aspects which each technique needs to take into account:

- **Characterization** - Uncertainty has to be described in some way, by measurements or models which reflect prior knowledge about the phenomenon of interest. For example, uncertainty about temperature in a specific location could be described by a probability distribution derived from repeated measurements. Or, temperature variations are assumed to follow a Gaussian-distributed pattern and measurements are performed to estimate its mean and standard deviation.
- **Representation** - Once uncertainty has been characterized, we have to decide which features of this description we wish to visualize. For example, given a Gaussian model for temperature variations we may select the standard deviation as our representation of uncertainty. Or, we determine a 95% confidence interval and use it as a quantitative measure.
- **Visualization** - After choosing how we quantify uncertainty, we can decide how to visualize it. For example, information about the range of temperatures at different locations in a certain land area can be visualized using color mapping or iso-contours.

In the following sections we will present an overview of uncertainty visualization

techniques. Choosing a good classification for this is not straightforward because, as explained earlier, uncertainty arises in many different forms. Several authors have proposed classification schemes for uncertainty visualization. For example, Pang et al. [121] classify techniques on the basis of what they *change* in the display, such as adding glyphs, adding continuous geometry (lines and surfaces), modifying geometry, modifying visual attributes, animation and sound. Griethe and Schumann [54] propose a similar classification but add interactive manipulation (e.g., of parameter settings) as a means of visualizing/exploring uncertainty. In the following, we propose a classification on the basis of the following categories:

1. **Value** - This represents uncertainty about measured or calculated data values. This category can be further subdivided according to data type, e.g., scalars, vectors, tensors and multivariate data.
2. **Location** - This represents uncertainty about the location where data values or objects exist. This can refer to the position of individual objects such as particles or it can refer to surfaces, such as reconstructed from height measurements.
3. **Parameter** - This represents uncertainty about the effect of variations in parameter values of models or computational algorithms.

5.2 Value uncertainty

Value uncertainty has received the most attention in the visualization research community. In many cases, data values are obtained from physical measurements. Modern data acquisition techniques allow the collection of enormous quantities of data and hardware performance often allows many repeated measurements in order to characterize uncertainty. We can distinguish between different data types, such as scalars, vectors, tensors and multivariate data. We will shortly describe different approaches for visualizing uncertainty for these different types.

5.2.1 Scalars

Scalar data is the most common type of data available from physical measurements. For example, Djurcilov et al. [42] use volume rendering to visualize 3D measurements of water temperature and salinity in ocean data. Uncertainty is characterized using Monte Carlo simulations based on ocean models and represented as a scalar-valued variance. Visualization was done using 1D and 2D transfer functions mapping combinations of data values and uncertainties to color and opacity. Cedilnik et al. [31] propose a general method for visualizing normalized, scalar representations of uncertainty in 2D scalar maps. To visualize uncertainty they overlay the 2D map with grid lines and apply varying levels of distortion to the grid lines based on the uncertainty of the underlying data values (Figure 5.1(b) and (c)). Uncertainty in scalar data can also be represented more fully, for example by means of probability density functions (PDFs). Visualization

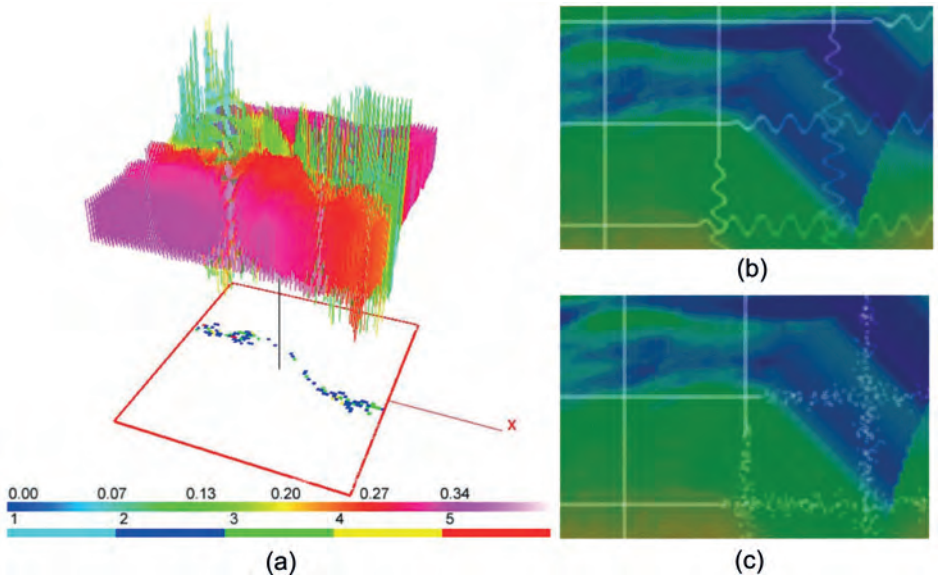

Figure 5.1: (a) Scalar uncertainty described by PDF features [77]. (b,c) Scalar-valued uncertainty used to introduce disturbances in grid lines such as (b) distortion and (c) noise [31].

of such PDF's can be accomplished with different types of graphs and plots, such as encountered in traditional exploratory data analyis [152] or extensions thereof [128, 16, 47, 166, 62]. A disadvantage of these techniques is the amount of screen space they require. Unless the data is very sparse, they do not easily translate to 2D or 3D maps without visual clutter and occlusion. A solution to this problem can be to reduce the PDF to a limited set of features, such as the mean or standard deviation. Kao et al. [77] take this approach to visualize such PDF features in 2D geographic maps, such as mean, standard deviation, interquartile distance and mean-median differences (Figure 5.1(a)). In some cases, the distribution of a scalar value may be unknown but a min/max range may be available. Olston et al. [116] propose specific visualization techniques for such *bounded* uncertainty.

5.2.2 Vectors

Vector fields are common in many application domains such as physics, weather prediction and oceanography. Even though vectors can represent any N-tuple of data values, they mostly represent directional information. Chapter 3 already mentioned that vector fields play an important role in DTI and fiber tracking. DTI data is initially represented as a 2nd-order tensor field but fiber tracking is mostly performed on the field of tensor main eigenvectors. We can distinguish different graphics primitives used for the visualization of vector field uncertainty:

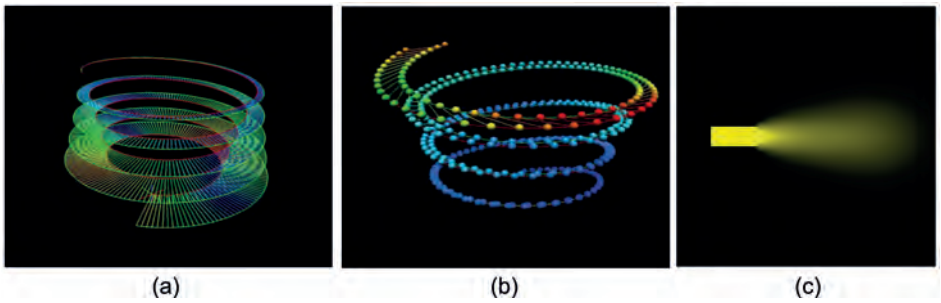

Figure 5.2: Uncertainty arising from different streamline tracing algorithms. (a) Line segments. (b) Barbell glyphs [95]. (c) GPU-based texture cross-advection [24].

Streamlines - The shape of the streamlines derived from streamline tracing depends on several factors. In Chapter 4 we already indicated that the numerical integration scheme and integration step size can be a source of streamline variation. Lodha et al. [95] investigated precisely this issue and propose various visualization strategies, such as barbell glyphs, ribbons and animations, to highlight streamline variations as a result of different tracing algorithms (Figure 5.2(a, b)). In Chapter 6 we will discuss how user-defined parameters can have similar effects. Another approach for streamline-based vector visualization is proposed by Boller et al. [23] who calculate trajectories of air flow by interpolating points inside so-called hypercubes (3D + time). Uncertainty is represented by the distance between each trajectory and an estimated worst-case trajectory in each hypercube. Visualization of uncertainty is done by mapping distance to streamline thickness.

Glyphs - Streamlines provide global information about vector fields at the expense of computational biases and errors. Alternatively, glyphs represent data values directly but only locally. This approach was taken by Wittenbrink et al. [168] who use glyphs to visually encode uncertainty in the direction and magnitude of wind velocities captured by repeated radar measurements. Uncertainty is represented as standard deviations or min/max ranges. The authors propose various glyph designs encoding different data values and uncertainties (Figure 5.3(a)). Although not explicitly designed for vector field uncertainty, Hlawatsch et al. [63] propose Flow Radar Glyphs that show directional changes in a local vector value over time (Figure 5.3(b)).

Texture-based approaches - Another common approach to visualize vector fields is to use texture advection such as proposed by Botchen et al. [24]. They investigate particle movement captured by particle image velocimetry. Directional uncertainty of the particles is captured by multiple measurements and represented

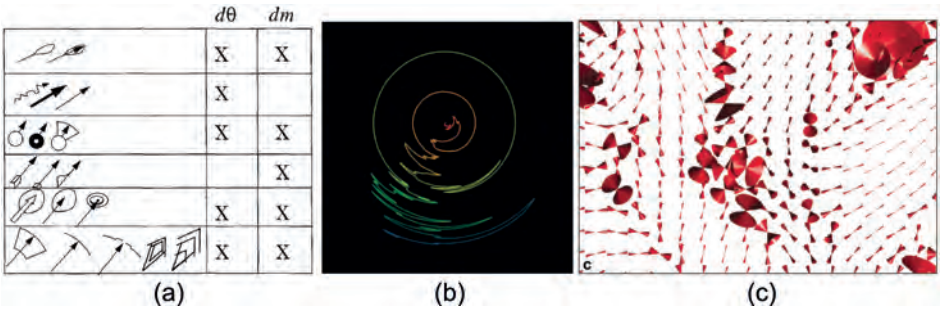

Figure 5.3: (a) Uncertainty glyphs [168]. (b) Flow Radar Glyph for representing time-varying changes of vector directions [63]. (c) Uncertainty cones [70].

as the root mean square of these measurements. Short streaklines are shown that are aligned with the average direction of particle movement. The root mean square is used to apply Gaussian cross advection to each streakline. Osorio et al. [117] propose a variation of line integral convolution (LIC) describing uncertainty for flow fields in both direction and magnitude. The LIC itself is derived from the mean vector in each location while uncertainty is represented by the standard deviation and visually encoded by color, blur and fog-like effects.

5.2.3 Tensors

Tensor data has received very little attention in uncertainty visualization research. In Chapter 3 we saw that 2nd-order diffusion tensors can be visualized using locally positioned glyphs. However, these do not provide global information about the tensor field. As explained in Chapter 4, different methods exist to capture uncertainty in the tensor field, such as bootstrapping [165, 71]. These methods result in a type of distribution of tensors in each data point. Visualization of this information has been limited to variations in each tensor's main eigenvector, for example, by means of uncertainty cones (Figure 5.3(c)) [70].

5.2.4 Multivariate data

Multivariate data are becoming increasingly common. It is defined as data where each data point contains not a single variable (scalar, vector or tensor) but multiple variables, possibly of different type. For example, each data point in a 3D dataset could store both temperature and pressure (both scalars), or temperature and ocean current (a scalar and a vector). Each variable may be associated with uncertainty represented by a single scalar, a full PDF or selected PDF features. For example, to visualize bivariate data distributions, bivariate box plots have been proposed by various authors as an extension to traditional techniques in exploratory data analysis [10, 140, 149, 53]. If more than two variables need to

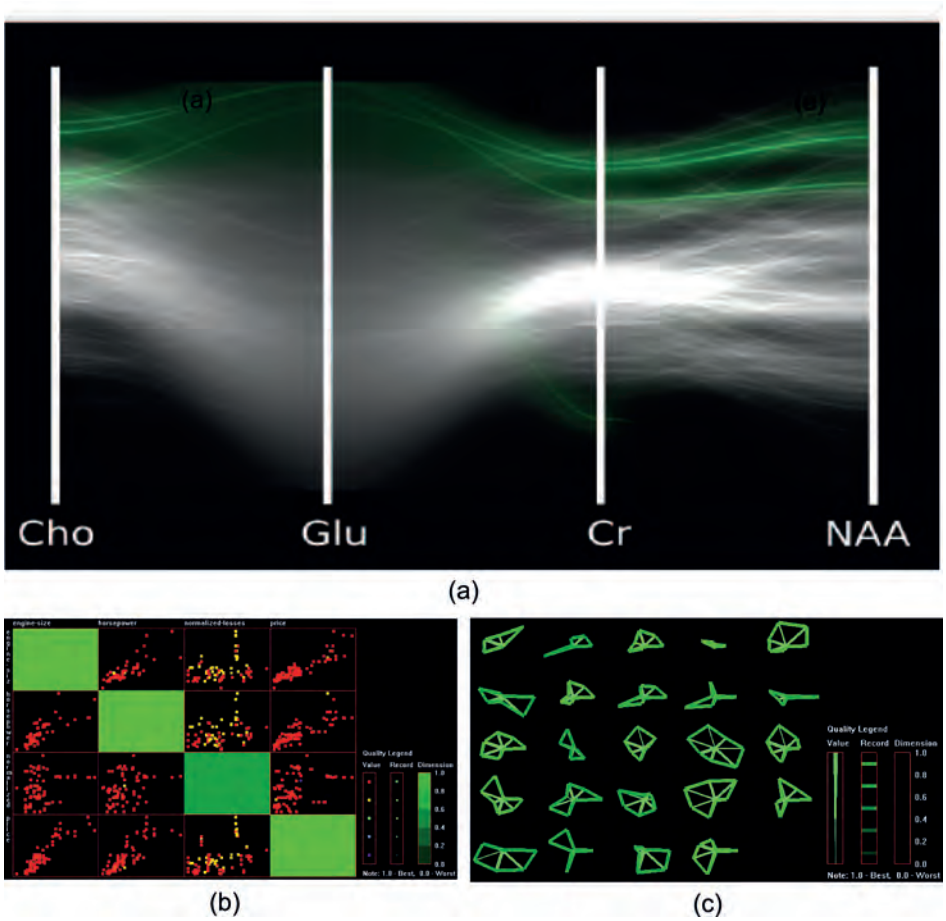

Figure 5.4: (a) Parallel coordinate plot [49]. (b) Scatterplot matrix and (c) star glyphs with uncertainty encoding [170].

be analyzed, other strategies are needed. Feng et al. [49] use a combination of parallel coordinate (PC) plots and scatterplots to visualize correlations between multiple variables. For scatterplots, each variable's uncertainty is determined by calculating a 2D PDF based on kernel density estimates. The scatterplot is rendered by linking probability to the level of blurring. A similar approach is taken with PC plots. The authors apply their visualization method to the analysis of tumor tissue on the basis of MR spectroscopy. Love et al. [100] do not explicitly mention uncertainty but suggest a number of approaches to deal with spatially varying PDF data in a general way (Figure 5.4(a)). They propose (1) parametric methods if the PDF can be appropriately parameterized, (2) PDF shape descriptors or (3) PDF operators that allow traditional algorithms to be re-defined in

terms of PDF data, such as color mapping, contouring and streamline and pathline tracing. Xie et al. [170] propose two methods for visualizing multivariate datasets together with scalar-valued uncertainty measures for each variable. First, they extend the dataset with additional variables representing the uncertainty and use PC plots, scatter plot matrices and star glyphs to visualize and explore relations between the dataset variables (Figure 5.4(c)). Second, based on the same visualization techniques, they encode uncertainty information as visual properties, such as color and opacity.

5.3 Location uncertainty

In the previous section we have discussed uncertainty in data value. However, measurements of spatial properties such as size, shape, and position are also subject to uncertainty. We summarize these as data *location* uncertainties. There are many application domains where uncertainties of location play a role, such as particle movement, star positions in astrophysics and surface reconstructions. The two most common primitives for which visualizations of position uncertainty have been investigated are (1) points and (2) surfaces.

Points - Lodha et al. [94] investigated particles moving at a constant speed. They describe uncertainty in terms of PDFs on position, direction and speed. To encode uncertainty in the rendering of each particle three methods are proposed. A particle is rendered as (1) a spherical glyph where the radius scales with probability or (2) as a point cloud where the cloud density and opacity depends on probability. The last method, uses a combination of color and transparency to encode uncertainty. Li et al. [93] focus on position uncertainty of stars in large-scale astrophysics data. The 3D position of stars is determined by standard polar coordinates, which can be measured very accurately, and distance, for which the parallax method can be used. For both an error can be estimated which provides a measure of uncertainty. Position uncertainty is visualized using either line segments along the line of sight (which ignore polar coordinate errors) or ellipsoid shapes which incorporate uncertainty both in distance as well as polar coordinates (Figure 5.5(b)).

Surfaces - Besides particle position uncertainty, Lodha et al. [96] also investigated variations in the output of surface reconstruction methods caused by different interpolation schemes. They present a wide range of techniques to highlight differences between two surfaces, such as texture mapping, bump mapping, pseudo-coloring, glyphs, transparency, animation, or plain side-by-side comparison [96]. Ehlschlaeger et al. [45] investigate landscape elevation data and propose the use of animation to highlight uncertainty in the corresponding height measurements. Multiple realizations of the elevation data are obtained by means of stochastic modeling.

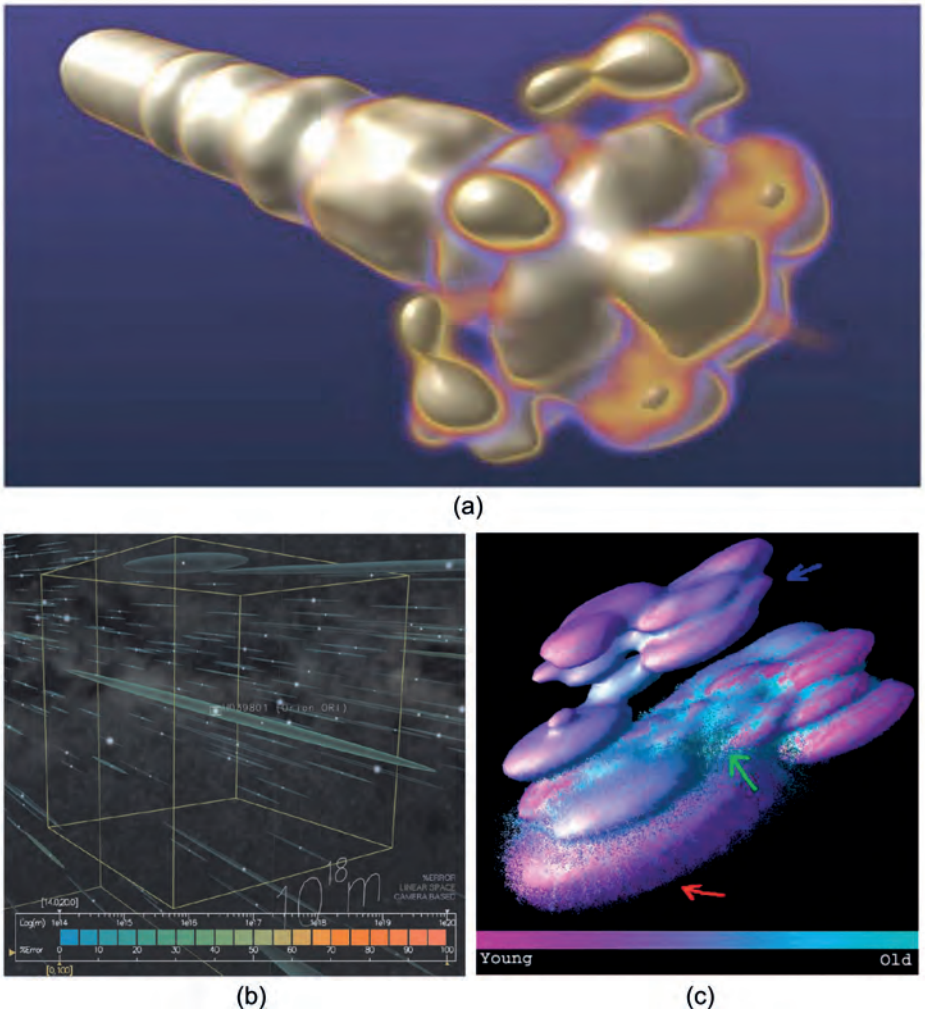

Figure 5.5: (a) 3D uncertainty isocontours of fuel density as it gets injected into the combustion chamber of an engine [126]. (b) Uncertainty in star positions [93]. (c) Point-based rendering of uncertainty in surface reconstruction of tumor tissue where color indicates the tumor age [55].

Once a probabilistic model of a landscape has been calculated, several scenarios can be investigated such as which land parts run the risk of flooding when sea levels rise. Grigoryan et al. [55] present a method for rendering surfaces using point primitives instead of polygons. Surface uncertainty is encoded by displacing point primitives along the surface normal. For highly uncertain surfaces this results in fuzzy point clouds that intuitively convey the underlying uncertainty while leaving

the use of color open for other data properties (Figure 5.5(c)). Joseph et al. [73] present a system, called UISURF, for comparative visualization of isosurfaces calculated using different algorithms and parameters. Uncertainty, in this case, is defined by the level of difference between two algorithms and visualized using different techniques, such as wireframe overlays, pseudo-coloring, transparency and box glyphs. Pothkow et al. [126] propose mathematical formulations of uncertain equivalents of 2D and 3D isocontours based on traditional probability theory and statistics. They apply these, among others, to fuel density fields during injection into an engine combustion chamber (Figure 5.5(a)). In a follow-up study [127], they apply their probabilistic formulation of isocontours to the Marching Cubes algorithm [97]. Rhodes et al. [134] address the problem of visualizing isosurfaces in very large datasets. A common method to increase interactivity is to create lower resolution versions of the dataset. However, this can result in unacceptable errors and information loss. The authors propose a method to calculate differences between high resolution and low resolution isosurfaces and visualize these differences together with the low resolution rendering. This allows assessment of which areas in the dataset are particularly sensitive to resolution reduction.

5.4 Parameter uncertainty

In order to deal with the complexity of many real-world phenomena, scientists and engineers often rely on models and transformations that simplify a given phenomenon to its essential characteristics. These models typically depend on input parameters that modify the model's behavior, sometimes in unexpected ways. Parameter uncertainty means that the reliability or stability of the model with respect to variations in parameter settings is unknown. If the model exhibits non-linear behavior, small perturbations in the input values or the parameter settings can give rise to large changes in the model's output values. In Chapter 6 we present work of our own in this area where we look at the sensitivity of anisotropy and curvature thresholds on the output of streamline tracing in DTI. Parameter sensitivity can also be investigated in other contexts. For example, Hadwiger et al. [56] propose a visualization framework for exploring different parameter settings in image segmentation for the purpose of feature detection in industrial CT. Berger et al. [17] proposed a system for the visual analysis and exploration of the parameter and output space in the context of car engine design. Users can select regions of parameter space and interactively see to which region in the space of output values this selection maps. Viceversa, the user can select desired output values and find out which parameter combinations are responsible for these output values (Figure 5.6(b)). Correa et al. [37] present a general framework for exploring the propagation of data uncertainty, modeled by a given distribution model, through the different stages of the visual analytics pipeline (Figure 5.6(a)). Also in medical volume rendering parameters play an important

role.

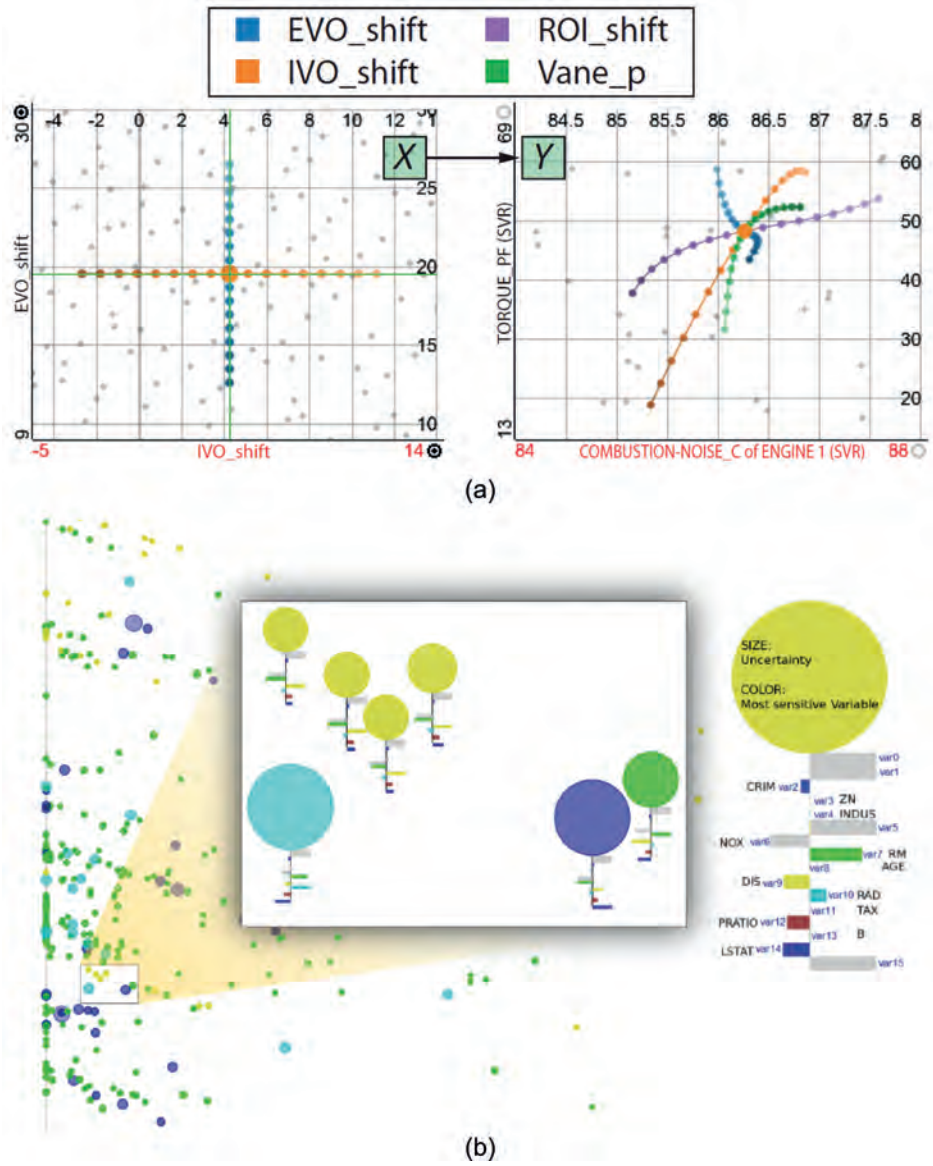

Figure 5.6: (a) Scatterplot indicating sensitivities for multiple parameters in each data point [37]. (b) Variations in input parameters X related to variations in output parameters Y [17].

Transfer functions (TF) are commonly used to highlight different tissue types.

However, finding good TF settings is very time consuming and unintuitive, sometimes leading to rendering artifacts which may be confused with real pathology. Lundstroem et al. [101] propose a method to find and explore such artifactual TF settings using animation. Malik et al. [104] present a comparative visualization system that allows the comparison of industrial 3D CT series acquired with varying acquisition parameters. They specifically focus on dimensional measurements for the purpose of detecting manufacturing artifacts in of the object of interest.

5.5 Conclusions

In the previous sections we have provided a comprehensive overview of previous work in uncertainty visualization research. We have attempted to structure the discussion by categorizing work on the basis of value uncertainty, location uncertainty and parameter uncertainty. As will be clear from this chapter, the visualization of uncertainty in the context of neuro-imaging has not received much attention so far. This is, of course, the main topic of this thesis and in the next chapter we will start with our first major contribution, the visualization of parameter sensitivity in deterministic fiber tracking algorithms.

I get mail; therefore I am.

- Dilbert

6 Parameter sensitivity

This chapter is based on:

"Parameter Sensitivity Visualization in DTI Fiber Tracking." Ralph Brecheisen, Anna Vilanova, Bram Platel, and Bart ter Haar Romeny. In *IEEE Transactions on Visualization and Computer Graphics* (2009), 15(6):1441-1448.

6.1 Introduction

Chapter 4 has provided a detailed overview of the different sources of uncertainty occurring in the DTI processing pipeline. In this chapter we take a look at the effect of user-defined parameters on the visual output of deterministic fiber tracking. Specifically, we look at the sensitivity of the fiber tracking algorithm with respect to stopping criteria, that is, the threshold values that determine when fiber tracking should terminate. The two most commonly used stopping criteria are the *anisotropy threshold* and the *curvature threshold*. As already explained in Chapter 3, different types of anisotropy threshold exist such as fractional (FA) and linear anisotropy (CL).

In many cases, a fixed set of thresholds is selected and re-used for different patient datasets. The threshold values are usually based on experience or empirical measurements of known anatomy. However, such measurements are not patient-specific. Furthermore, a small change in the threshold values can lead to very different output results. The user remains completely unaware of such effects and is therefore likely to either underestimate or overestimate the tracking results. Little or no effort is spent on evaluating the stability of the threshold values or what the effect of threshold variations is on tract-specific features, such as average fiber length or tract volume. Such features play an important role in group studies where they provide quantitative measures for comparison [38]. If one is planning to draw conclusions based on such quantitative features, it is important to know whether the feature is sensitive to threshold variations or not. This has also been recently confirmed in a clinical trial by Taoka et al. [145].

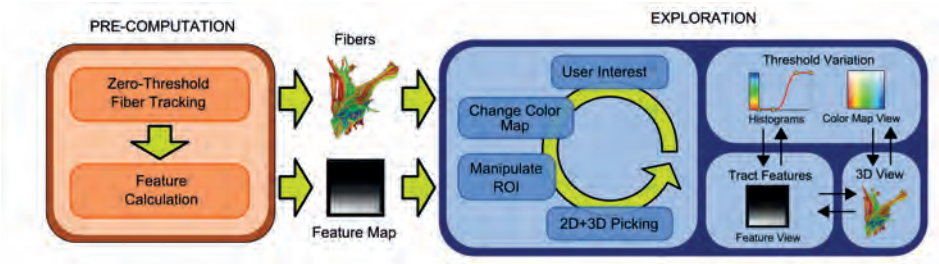

Figure 6.1: Pipeline overview: (1) The pre-computation stage consists of two steps: Zero-threshold fiber tracking computes stopping criteria along streamlines. Tract feature calculation involves computing a quantitative metric for a discrete grid of sample points on the continuous parameter space. The feature values are stored in a feature map. (2) The exploration stage consists of tools and views to (i) investigate variations in the stopping criteria along the fibers and (ii) how these variations affect tract features such as the average fiber length. Selected threshold combinations can be picked in both 2D maps and a 3D fiber view. Picking operations and manipulation of regions-of-interest are linked between the different views.

In this chapter, we have selected standard streamline tracing [110, 8] based on the FACT algorithm with 2nd-order Runge-Kutta integration. It is fast, simple to implement and the most widely used method for diffusion tensor imaging. Our visualization method is

not limited to this particular fiber tracking algorithm. The requirements we place on the algorithm are (1) that it makes use of stopping thresholds and (2) that variations in threshold values affect only streamline *length*. It should not affect streamline shape. With existing fiber tracking tools, investigating how different threshold values in the stopping criteria affect the output results would require manually trying out many different threshold combinations. This is a very time-consuming process that may have to be repeated for every new dataset. Even if parameter values can be modified interactively, it would still leave the user wandering around blindly in an unknown parameter space. Furthermore, just seeing how a single parameter combination affects the output, is not sufficient. To give true insight into the algorithm's behavior across a *range of threshold combinations*, requires specific visualization techniques. To deal with these issues, we propose a visual exploration tool that allows users to investigate the behavior and sensitivity of DTI fiber tracking for stopping criteria. This tool relies on a combination of 3D and 2D views linked together by means of point selections and brushing techniques, similar to Doleisch et al. [43]. We will show that our tool can assist users to gain more insight into the sensitivity of the algorithm for these criteria. This should result in more reliable output results and an improved ability to compare groups of patients based on quantitative tract features. We also present a small user evaluation based on the experiences of three domain experts who use DTI fiber tracking on a regular basis in their own research.

6.2 Processing pipeline overview

An overview of our processing and exploration pipeline is illustrated in Figure 6.1. The pipeline consists of two main stages: (1) an unattended pre-computation stage, and (2) an exploration stage where the user visually explores the results of the pre-computation stage. The pre-computation stage (Section 6.3) consists of generating streamlines for the full space of threshold values. We explain how these results can be computed in a fast and simple manner. We also compute a number of quantitative *tract features* on a grid of sample points in the threshold space. This results in a feature map that allows the user to see the effect of threshold variations on such features. Finally, the exploration stage (Section 6.4) provides a number of views and interaction tools to visually explore threshold variations and their effect on tract features.

6.3 Pre-computation

The pre-computation stage generates the information to interactively explore the threshold variations in fiber tracking. The only user-defined settings required at startup are the geometry of the seed region and which anisotropy measure to use (fractional, linear, etc.). However, our approach is not restricted to the use of seed regions. It is possible to perform whole-brain fiber tracking and then select a subset of fibers to be investigated [1, 22].

6.3.1 Zero-threshold fiber tracking

The basis for our exploration pipeline is the theoretical set of fiber bundles or *fiber tracts* resulting from streamline tracing with all possible combinations of threshold val-

ues. Once we have this set of tracts, we can visualize them interactively, compute tract features and show how the feature values change as we modify threshold values.

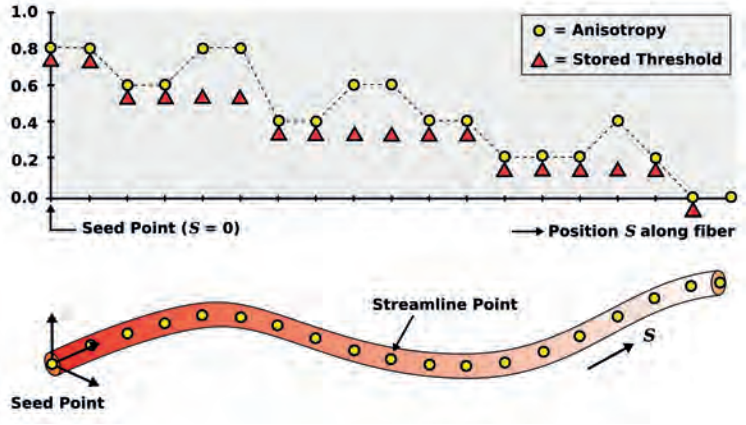

Figure 6.2: Anisotropy threshold profile along single example fiber starting from the fiber seed point on the far left. The yellow circles indicate a variation of tensor anisotropy along the fiber. The red triangles indicate the anisotropy threshold that is stored for each point. As can be seen, the anisotropy itself does not need to be monotonically decreasing. The threshold value is always monotonically decreasing. The same holds for streamline curvature and curvature threshold.

A fast and simple way to obtain the required tracking results for the whole space of threshold combinations is to perform so-called *zero-threshold* streamline tracing, meaning that we allow tracing to proceed in an uninhibited manner without thresholds. We do this only once for a given seed region and this results in a collection of streamlines with maximum length and maximum curvature, the so-called *zero-threshold tract*. While performing zero-threshold tracing, we inspect the local tensor anisotropy and curvature in each streamline point and store, as additional point attributes, the combination of threshold values that would have terminated tracing at this point. This approach works because the thresholds we are considering only affect streamline *length*, not shape. Figure 6.2 illustrates the concept of zero-threshold tracking for the anisotropy threshold. As can be seen in the graph the anisotropy along the fiber is not always monotonically decreasing. The anisotropy *threshold* on the other hand, by definition, is always monotonically decreasing. After all, once the anisotropy drops below the threshold, streamline tracing terminates regardless of whether there are higher anisotropy values further along the fiber. Since we define curvature as the dot product between two subsequent step vectors, the same reasoning holds for the curvature threshold. To visualize specific output results for a given threshold combination we apply a simple filtering procedure on the streamline points of the zero-threshold tract. Streamline points with threshold values *above* the selected threshold combination are simply discarded.

6.3.2 Calculation of quantitative tract features

DTI fiber tracking is most widely used in population studies that require statistics to show abnormalities in either brain or muscle tissue (e.g., in the heart). To be able to quantitatively compare fiber tracking results between patients and healthy controls, users often define so-called tract features (also called tract metrics) that compute some property of a given fiber tract or muscle. Examples are total and average fiber length [38]. A general problem of such quantitative features is their dependence on fiber tracking parameters such as the stopping criteria. It is important to check whether a particular tract feature remains stable for any threshold variations, and if so, for which regions of the threshold space this holds. The anisotropy and curvature threshold combinations we are considering in this chapter span a continuous 2D space. Obviously we cannot compute a tract feature for each point in this space since there are infinitely many of them. We can however define a regular grid of sample points and compute tract features for those. We call such a grid of computed feature values a *feature map* and it allows us to visualize how a given tract feature varies as a function of threshold combination. In the following paragraphs, we will discuss a number of tract features that we implemented in our tool. This list is by no means complete since many different, application-specific tract features can be defined. In principle, our tool can be extended with any scalar-valued tract feature that can subsequently be analyzed for stability.

Total, Average and Standard Deviation Fiber Length These tract features are the most widely used. Correia et al. [38] define as many as nine different tract features based on fiber length, each tuned for specific white matter conditions. For example, old age is generally accompanied by a significant reduction in both total and average fiber length. In conditions where fiber anisotropy is interrupted by multiple lesions, the total fiber length may not differ from a healthy person. The average fiber length however may reveal this condition quite clearly.

Average Fiber Density per Voxel We compute this feature by counting, for each non-empty voxel, the number of streamlines intersecting that voxel. The total count is then divided by the number of non-empty voxels to obtain the average fiber density. Of course, this feature depends on both dataset resolution and seed point density. However, this is not problematic, because for analyzing feature stability we are only looking for changes (or lack thereof) in a single dataset. In certain compact and elongated fiber tracts (such as the cingula or optic radiations) the streamlines remain closely packed together. If at some threshold values the streamlines start running off in random directions we expect the average fiber density per voxel to drop.

Tract Volume This is another popular tract feature whose reduction can point to neurodegenerative diseases such as Alzheimer's. In studies of muscle tissue, tract volume can be used, together with average fiber length, to create models of specific muscles for biomechanical simulations. Tract volume can be computed by taking the volume of a single voxel and multiplying this by the total number of non-empty voxels. It is also possible to take the fiber density of each voxel into account as a weighting factor.

Mean of End-Point Distances We derived this feature from the area of fiber clustering. In fiber clustering one starts with a single fiber and tries to find neighboring fibers

that look similar according to some similarity measure. In this case, we already have a cluster (a single fiber tracking result) and we wish to compute the average similarity between its streamlines. The average of end-point distances is a similarity measure adapted from Moberts et al. [109]. Fibers are considered to be similar if they have endpoints that lie close together. This reflects the fact that fibers from the same anatomical structure connect the same areas of the brain. The end point distance d_E between fibers F_i and F_j is defined as,

$$d_E = \min(d_1, d_2) \tag{6.1}$$

where,

$$d_1 = ||F_{i,1} - F_{j,1}|| + ||F_{i,end} - F_{j,end}|| \tag{6.2}$$

$$d_2 = ||F_{i,1} - F_{j,end}|| + ||F_{i,end} - F_{j,1}|| \tag{6.3}$$

Here, $F_{i,1}$ and $F_{i,end}$ refer to the first and last points on fiber F_i. The end-point distances are computed for each fiber pair and averaged for the given tract. This is then repeated for all threshold combinations.

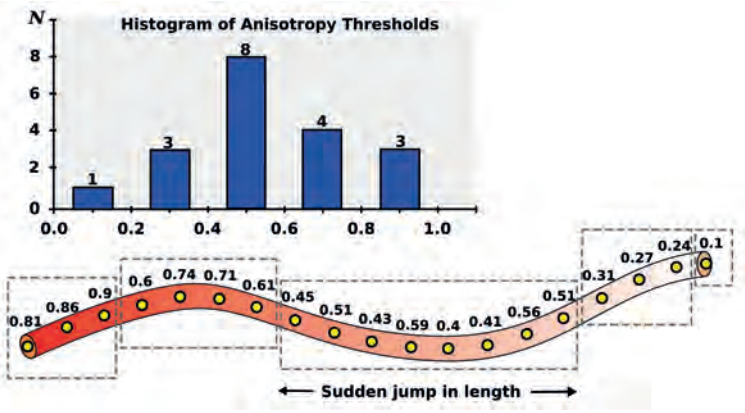

Figure 6.3: Histogram of continuously varying anisotropy thresholds along single fiber. The peak between $[0.4 - 0.6]$ corresponds to a sudden jump in fiber length indicating a possibly critical threshold value.

6.3.3 Cumulative Histograms

In addition to the feature map we also provide a number of histogram views as an alternative way to show the variation of threshold values along streamlines and the stability of tract features within the threshold space. To make it easier to detect stable regions, the histograms can be viewed as cumulative distributions.

Threshold histograms If we consider the zero-threshold tract that was computed in the pre-computation step (Section 6.3), we can create a two-dimensional threshold histogram by subdividing the threshold space into bins. We can then assign each streamline

point in the zero-threshold tract to its corresponding bin according to the threshold values stored at this point. Technically, the 2D histogram is not the same as a feature map but we can use the feature map view to display this 2D histogram. We can also compute a histogream for each threshold separately. This is illustrated in Figure 6.3 for the anisotropy threshold. A peak in the histogram corresponds to a sudden jump in overall fiber length, possibly pointing to a critical combination of threshold values. An alternative approach to obtain a 2D histogram is to normalize the individual 1D histograms combining them to form a 2D joint probablility distribution. This would assume however that anisotropy and curvature are independent variables, which need not be the case. As Figure 6.4 illustrates it is quite reasonable to believe that the average anisotropy in a voxel depends on the average curvature of fibers.

Feature histogram The feature histogram provides an alternative view on the distribution of feature values in the threshold space. The user can choose whether to show a normal histogram or the cumulative distribution. As explained in Section 6.3.2 this can help detect stable regions where threshold changes have little effect.

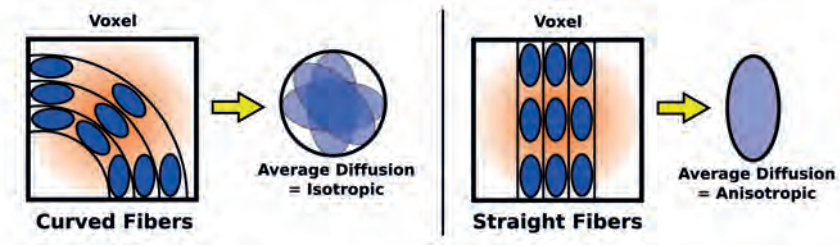

Figure 6.4: Example of possible correlation between curvature and anisotropy. High curvature inside a voxel averages anisotropies with different orientations, which may result in a measurement with (near) isotropic diffusion. Zero curvature in a voxel with the identical anisotropy values results in greater average anisotropy.

6.4 Exploration of thresholds and tract features

After the pre-computation stage has finished the user can start to interactively explore threshold variations and their effect on different tract features. For this purpose, we provide a 3D view for displaying streamline output, 2D plots capturing the threshold space and allowing the user to select and manipulate threshold ranges interactively, and quantitative histograms for detecting stable regions in the feature map. The different elements of our tool's user interface (shown in Figure 6.5) are described in the following paragraphs.

3D viewing and interaction: In the upper-left of Figure 6.5 the 3D view is shown where the different tracking results can be visualized. The fibers can be represented either by streamlines or streamtubes. Each vertex on the streamline or streamtube is associated with a set of threshold values calculated in that vertex and which are stored as 2D texture coordinates. This allows us to use 2D texture mapping to apply different

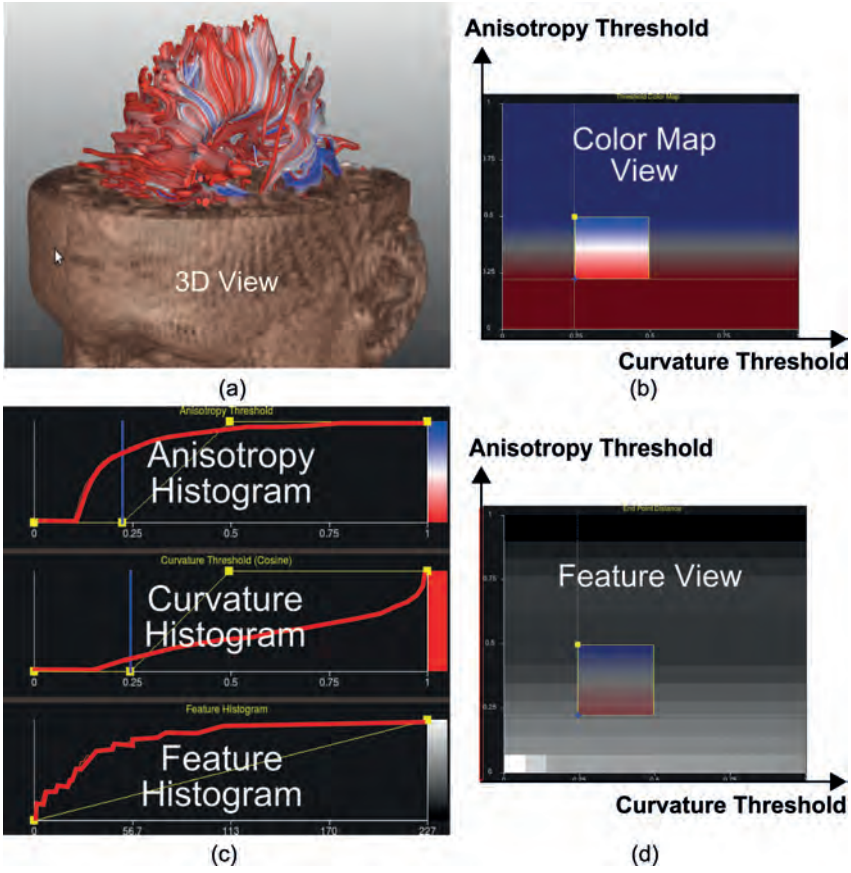

Figure 6.5: (a) Main viewports of our exploration tool. 3D visualization of fiber tract together with anatomical context and axial fractional anisotropy slice. (b) Color map view used for selecting individual threshold combinations and definition of color detail regions. (c) Cumulative histograms of both threshold and feature values. (d) Feature map view showing changes in quantitative tract features as a function of threshold combination at discrete sample points of the parameter space.

color maps to the fibers. We also provide the option to map one threshold parameter (e.g., curvature threshold) to the radius of a streamtube and the other to color. If the user finds color transitions on the fibers that may point to interesting threshold changes, he or she can pick specific points on the fiber and see in the color map view (described next) what the exact threshold value is in that location.

Color map view: Although this view is also used to display the color map that is currently active, it is primarily meant as an interaction tool. It allows the user to pick points in the 2D threshold space and immediately see what the fiber tracking result would be in the 3D view. In this case, the color map is applied to the full range of

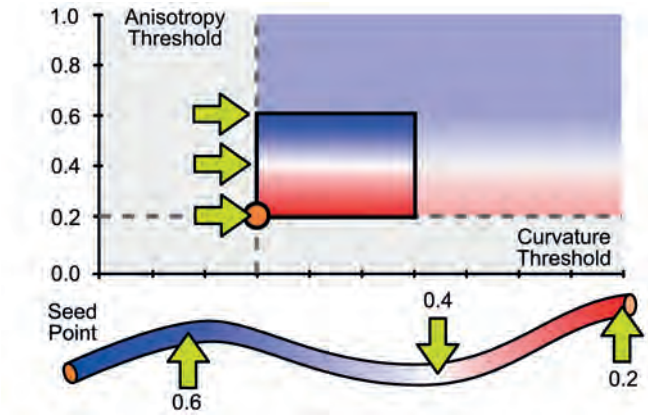

Figure 6.6: Relation between region-of-interest (ROI) selection and application of color map to subrange of threshold values. The streamlines displayed in the 3D view correspond to the lower-left corner point of the ROI. The color mapping along the streamlines would show how they shorten due to increasing threshold values.

threshold values ($[0, 1] \times [0, 1]$). For example, an anisotropy threshold of zero is mapped to a red color, while an anisotropy threshold of one is mapped to a blue color with intermediate values going through white. However, if the user wishes to study variations in the local Region-Of-Interest (ROI) surrounding a given threshold combination, the color map should be applied to this region only. Figure 6.6 shows the relation between ROI selection and fiber coloring/geometry. The color map view allows the user to specify such a ROI, and move/resize it interactively with immediate updates in the 3D view. Figure 6.7 illustrates how ROI coloring improves the distinction between threshold variations along the fibers.

1D and 2D color maps: We support a number of uni- and bivariate color maps to highlight threshold variations along the fibers, either to show the anisotropy threshold, curvature threshold or both. There has been a substantial amount of research in the area of uni- and bivariate color mapping [143, 133]. Figure 6.9 shows the different color maps we support. Gray scales are well suited for detecting subtle changes in value and preserves the order of data values. For this reason we use it primarily to highlight tract feature values. Double-ended color scales specifically highlight low, medium and high values which make them effective for visualizing threshold variations. This is quite effective in combination with defining a ROI in the color map view because the ROI focusses the center and extrema of the color scale to a smaller threshold range. The 2D complementary color scale is defined by colors that lie on opposite sides of the hue circle. For example, one variable X of range $[X_{max} - X_{min}]$ is mapped to brightness of green. Another variable Y of range $[Y_{max} - Y_{min}]$ is mapped to brightness of the complementary color, which is purple [48]. The contribution of both variables is obtained by adding the colors in RGB space. This color scale highlights the absence or presence of correlations between two variables [143, 133]. As explained in Section 6.3.2 there may

exist a correlation between curvature and anisotropy. However, when we applied this color scale to the fibers this correlation was not visually detectable.

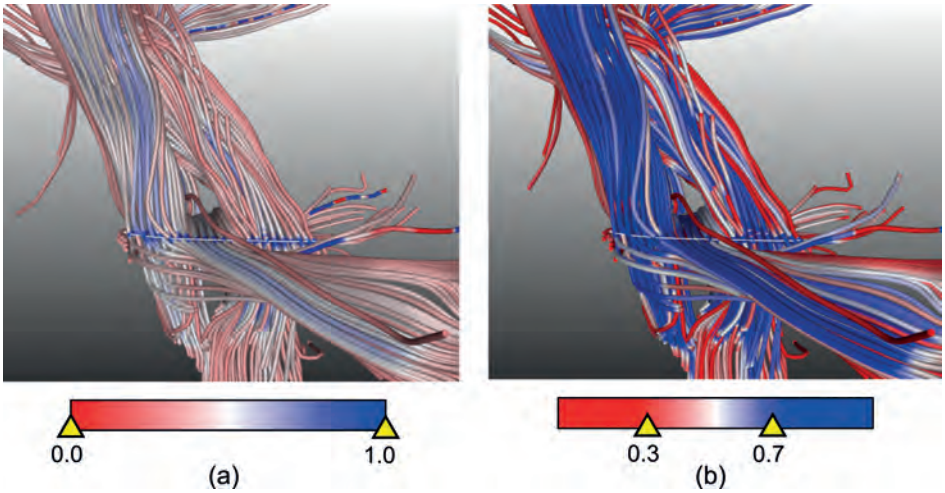

Figure 6.7: (a) Color scale mapped to full threshold domain $[0, 1]$. (b) Color scale mapped to threshold region-of-interest $[0.3, 0.7]$. This highlights threshold variations along the fibers with full color detail.

Two color scales remain, namely the 1D hue color scale and the 2D hue-saturation color scale. Despite its limitations, hue remains one of the most popular and widely used color scales. It is difficult to interpret without the use of a color legend and the scale introduces discontinuities that may not be present in the data itself. Nevertheless, we include it for completeness. The 2D hue-saturation scale suffers from similar drawbacks, even though the problem is somewhat alleviated if we link hue to curvature threshold (which varies much less than the anisotropy threshold). The saturation scale has similar properties as the gray scale but when combined with hue it loses contrast and is less effective for detecting changes.

Feature view: When DTI fiber tracking is applied in comparative studies of white matter between patients, quantitative tract features are often used. It is important to verify whether a given tract feature is sensitive to threshold variations or not. The feature map view can be used specifically for this purpose. As mentioned previously each point in the threshold space corresponds to a single fiber tracking result. Section 6.3 discussed the different features that can be computed for such a tract. Obviously, we cannot compute a tract feature for each point in the threshold space because there are infinitely many of them. However, we can sample the space on a rectilinear grid and compute a tract feature for each sample point. We can then display the feature values as a 2D image where each sample point is represented by a pixel. We call this image the *feature map* and by using gray scale coloring we can highlight subtle changes in the tract feature values as a function of threshold combination. The user can select different

features from a pull-down menu. We decided not to merge the feature map view with the color map view because the feature map itself would lose too much gray value contrast when combined with the RGB/HSV color maps. Instead, we provide the feature map as a separate view that can be overlaid with a user-defined opacity on top of the color map ROI.

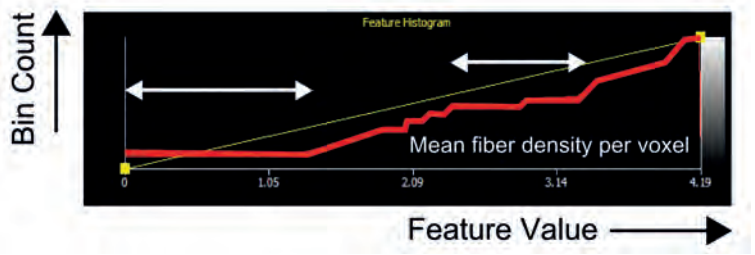

Figure 6.8: Cumulative feature histogram (indicated by red line) for mean fiber density per voxel. The white arrows indicate areas where the tract feature is relatively stable for threshold variations.

Histogram views: We provide several histograms as an alternative way of displaying threshold sensitivity. The threshold histograms were already discussed in Section 6.3, except here they are visualized as separate 1D histograms. They can be converted to cumulative histograms, which makes it easier to detect flat (and therefore stable) regions in the threshold space. We also show a feature histogram which is simply a histogram constructed from the list of feature values computed on the grid sample points described in the previous paragraph. The feature histogram shows the currently selected feature.

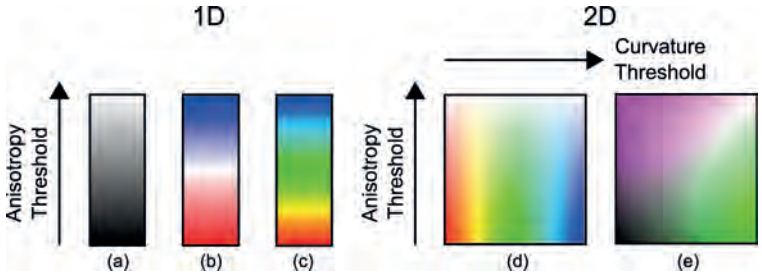

Figure 6.9: 1D color maps: (a) gray scale, (b) red-white-blue, (c) hue. 2D color maps: (d) Hue-saturation, (e) green-purple complementary colors.

6.5 Results and discussion

In this section we present a number of visualizations that were created for three clinical DTI datasets. The first dataset represents a healthy human brain (Siemens 3 Tesla, voxel

dimensions $1 \times 1 \times 1$ mm^3, resolution $231 \times 172 \times 131$, 72 directions). The second dataset also represents a healthy human brain (Philips 3 Tesla, voxel dimensions $2 \times 2 \times 2$ mm^3, resolution $128 \times 128 \times 66$, 32 directions). The third dataset represents muscle tissue of a right human forearm (Philips 3 Tesla, voxel dimensions $1.79 \times 1.79 \times 6$ mm^3, resolution $112 \times 112 \times 54$, 15 directions). Figure 6.10 illustrates how our visualization method can help find critical threshold variations along the streamlines. One can see a clear white band running across a group of neighboring streamlines (indicated by the white arrow). The colors show a sharp transition from blue through white to red corresponding to a sudden drop in the fractional anisotropy threshold. If this drop was purely caused by noise it would be unlikely to form such a regular, localized band pattern across multiple streamlines. Another phenomenon which our approach can help identify is crossing fiber structures. Figure 6.11 illustrates how a sudden drop in curvature threshold can result in the tracking of erroneous fiber tracts such as the cingulum, which runs along the 'gutter' of the corpus callosum complex. Such a region may contain more disk-shaped tensors where the main diffusion direction is not well-defined ($\lambda_1 \approx \lambda_2$). In this case, fiber tracking becomes very sensitive to noise and can start tracing the wrong tract. Using our tool, one can see exactly the point at which this wrong turn occurs. If this location is picked with the mouse, the exact threshold values will be indicated in the color map.

To further evaluate the added value of our visualization approach, we have asked three expert users to give feedback on the practical utility of our tool using these three datasets. All of them use DTI fiber tracking on a regular basis albeit for quite different research purposes. We performed the evaluation by first demonstrating the tool to the user. We then let the user work with the tool him or herself. After this, we presented the user with a list of questions designed to obtain structured feedback about the practical utility of our tool and parameter sensitivity visualization in general.

6.5.1 Brain development in premature neonates

Our first user is a clinical physician who uses DTI fiber tracking to study the effect of hypoxic ischemia on white matter development in premature neonates. To show the difference between neonates with white matter deficiencies and healthy controls, this user employs quantitative tract-based features. However, because such features are highly dependent on stopping criteria, this has been problematic. In the past, the user investigated the sensitivity of stopping criteria on certain tract features by repeatedly performing fiber tracking with different combinations of thresholds. With existing DTI tools this is extremely time-consuming, so the possibility to show this information in a single feature map of our tool was very much appreciated. Especially regions of the feature map that are relatively constant are of particular interest because they point to subranges of the threshold space where the feature is relatively stable and can be safely used for inter-subject comparison.

Typical tract features the user is interested in are tract volume and average fiber length. These are available in our tool. We briefly discussed the average fiber density per voxel feature. The user is not employing this feature at the moment but thought that it could be combined with tract volume by weighting each voxel's contribution with its fiber density. Some features that the user suggested to be included are the average

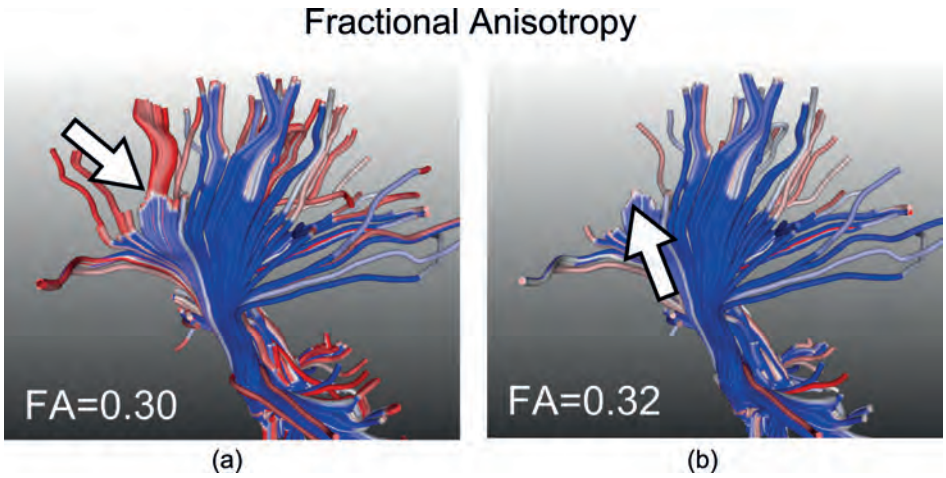

Figure 6.10: Effect of small variations in anisotropy threshold on fiber tracts of the corona radiata (running from the brainstem towards the brain cortex). If fractional anisotropy is increased from (a) 0.30 to (b) 0.32 the red tract parts (indicated by the white arrow) completely disappear. The cut-off point is already visible in (a) in the white regions of the fibers (indicated by the white arrow).

anisotropy (either fractional or linear) along the fiber tract and the average apparent diffusion coefficient (ADC). We do not support these features at the moment but they can be easily included. Because so many different, application-specific tract features are available [38], we did not provide a complete list.

6.5.2 Muscle tissue

The second user is a researcher who uses DTI and fiber tracking methods to investigate injury and repair of muscle tissue. With DTI it is possible to detect subtle tissue differences that are not visible on normal MR. After demonstrating our visualization tool, we let the user interact with one of his own datasets. Showing DTI-based parameters along the streamlines is a feature that is completely missing in the tools the user is currently working with. He encouraged us to extend the number of parameters that can be shown along the streamlines (now we show only anisotropy and curvature thresholds) but even with just the anisotropy color coding he saw, for the first time in 3D, transitions from muscle tissue to connective tissue (tendons). Figure 6.12 illustrates these transitions (red = connective tissue, white = muscle, blue = bone attachment) for a selected set of muscles in the forearm.

We also discussed the various tract features with this user. He is particularly interested in tract volume and average fiber length because these, together with the attachment angle between muscle and bone, are sufficient to model the muscle for biomechanical simulations. However, since these features are highly dependent on user-defined stopping criteria, it is difficult to find values that are reproducible. Our visualization method was

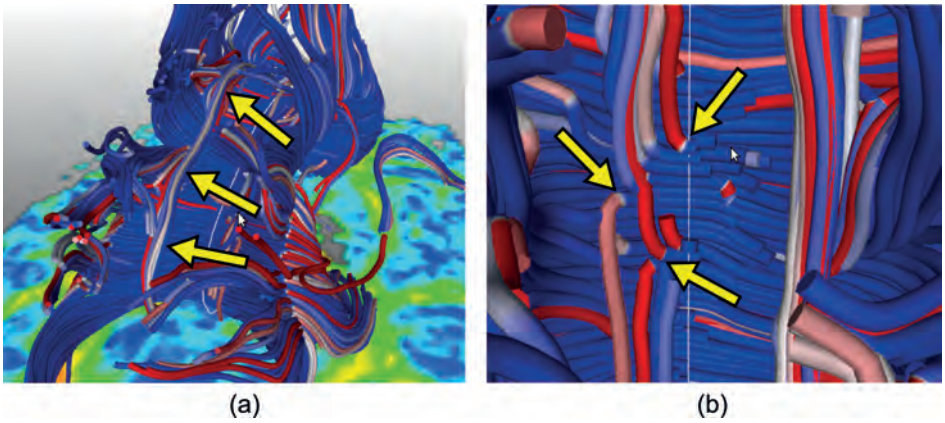

Figure 6.11: (a) Reconstruction of the corpus callosum with unintended tracing of the cingulum (indicated by yellow arrows). (b) Close-up of area where unintended tracing originates. Color variations (blue to red with narrow band of white) show the sudden drop in curvature threshold which seems to be responsible for the deviation of the pathway into the cingulum.

considered to be very useful for visualizing the variation of feature values as a function of stopping thresholds. Another feature which the user found interesting is the mean of end-point distances. Muscle fibers have a strongly parallel orientation. As long as the fiber tracking output contains only muscle fibers, this feature is expected to remain relatively constant. However, when muscle fibers converge to form connective tissue the streamline end-points are expected to come closer together, resulting in a *decrease* of the mean of end-point distances. If the goal is to segment only muscle fibers, this can help to find the exact thresholds where muscle fibers move into connective tissue. The threshold histogram provides similar information because the anisotropy threshold of connective tissue is generally lower than muscle tissue. As soon as many connective fibers are being traced by the algorithm, a sudden jump in fiber length is expected which results in a peak in the threshold histogram. The average fiber density per voxel is expected to *increase* as muscle fibers converge to form more dense tendon fibers.

The overall feedback from this user was very positive because our tool directly showed tissue transitions in 3D that had not been seen previously. The tract features we provided are by no means complete but the general approach of showing the behavior of any tract-based feature as a function of threshold was considered very useful, especially for group studies. Tract features are highly dependent on user settings however, which is why some people are reluctant to use them. Our user also strongly encouraged us to extend the parameter-space concept to other tensor-based properties such as the apparent diffusion coefficient and the tensor's main eigenvalue.

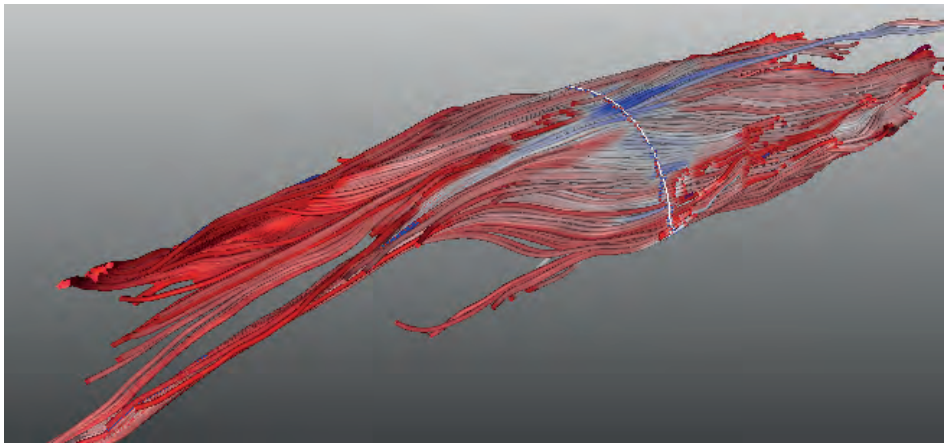

Figure 6.12: Fiber tracking of muscle in a human forearm. The elbow is located in the top-right corner, the hand is located in the bottom-left corner, palm facing upward. Red colored tissue corresponds to tendons running towards the hand. White colored tissue is muscle. The blueish colorations indicate attachment of muscle to bone.

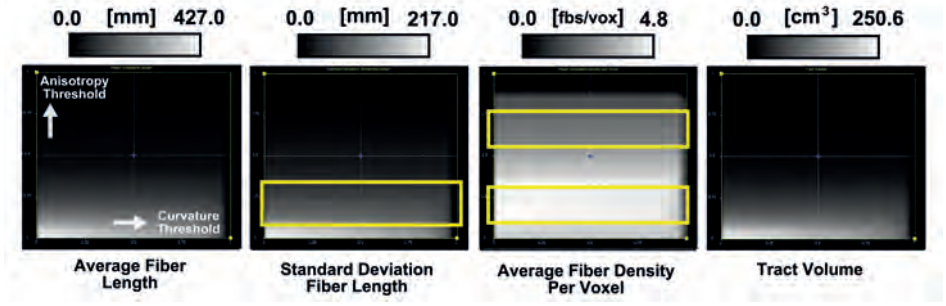

Figure 6.13: Four example feature maps (anisotropy on vertical axis, curvature on horizontal axis). Average fiber length shows gradual increase as both the anisotropy and curvature thresholds become lower, indicating a strong dependency on these thresholds. Standard deviation of fiber length has a region of constant value (yellow box) where at least the anisotropy threshold seems to have little effect. Average fiber density per voxel has two regions that are relatively stable (yellow boxes). The tract volume feature also shows a significant dependence on anisotropy threshold variation. The color legends indicate the min/max range and unit of measurement for each feature.

6.5.3 Brain connectivity

The third and last user is a researcher involved in the study of brain connectivity who uses DTI fiber tracking together with functional MR to investigate a patient with Landau-Kleffner syndrome. This patient acquired aphasia (loss of receptive and expressive language skills) due to epileptic seizures in the temporal lobe during early childhood. After

more than 15 years of intensive training in sign language she is able to communicate quite well, although she has trouble understanding people without actually seeing them. The goal is to investigate whether she has formed brain connections that are not present in healthy controls. Especially connections between the remaining language areas and the motor area dealing with hand coordination are of interest. The main problem when using DTI fiber tracking for this type of research is finding threshold values that are sufficiently low to confirm a suspected connection between two regions of interest, but are also sufficiently high to be plausible. Our visualization was considered very useful in helping to make a more reliable choice of threshold values. As with the second user, this user suggested to extend the parameter-space concept to include parameters that describe tensor shape. For example, it would have been interesting to show the *linear* anisotropy threshold on one axis and the *fractional* anisotropy threshold on the other axis. Using 2D color coding along the streamlines, one could then see where the tensor shape suddenly changes from linear to planar or vice versa. If this occurs at regular positions across multiple streamlines, this might point to true tissue changes (e.g., fiber crossing) instead of just noise. This user was rather sceptical about using tract-based features because of their sensitivity to user-defined parameters. We showed how our visualization method can be used to investigate how stable a tract feature really is within the range of threshold values relevant for the user's application. He agreed that our approach would encourage the use of quantitative tract features more often. They can provide additional information for showing that certain brain connections only exist in patients and not in healthy controls.

6.6 Conclusions

In the previous sections we have presented a tool for the visual exploration of parameter sensitivity in fiber tracking. We also allow users to investigate the behavior of certain tract-based features. In the next chapter we proceed with our second major contribution, the visualization of fiber confidence intervals on the basis of probabilistic fiber tracking.

I respectfully decline the invitation to join your hallucination.

- Dilbert

7 Illustrative confidence intervals

This chapter is based on:

"Illustrative Uncertainty Visualization of DTI Fiber Pathways." Ralph Brecheisen, Bram Platel, Bart ter Haar Romeny, and Anna Vilanova. In *The Visual Computer* (2012), conditional accept.

7.1 Introduction

In the previous chapter we looked at techniques for visualizing and exploring parameter sensitivity in deterministic fiber tracking. In this chapter we focus on two, slightly related, sources of uncertainty which are potentially responsible for many false positives and random variations in pathways reconstructed with fiber tracking. These are *image noise* and *model errors*. These sources of uncertainty have been discussed in detail in Chapter 4. The two are related because noise and round-off errors can affect tensor eigenanalysis such that the two largest eigenvectors are arbitrarily ranked. This can result in random fiber orientation estimates which manifest themselves as a variation in pathway trajectories during fiber tracking. In regions where the diffusion model is incorrect, noise effects will be particularly noticeable and fiber pathways will tend to spread out in random directions.

Existing visualizations of DTI fiber tracking largely ignore the effects of noise and model errors. This gives an impression of certainty that can be misleading. This is unacceptable in applications such as neurosurgery, where fiber tracking may be used for surgical risk assessment and decision-making. If false positives or false negatives in the fiber tracking output are not correctly identified this can result in sub-optimal tumor resection or damage to healthy brain tissue. To deal with risks of healthy tissue damage neurosurgeons may take into account safety margins around critical brain structures [161]. However, such safety margins are mostly fixed and have no relation to the uncertainty in the underlying data. Furthermore, safety margins enclose only one out of many possible fiber configurations. Figure 7.1 illustrates that, due to noise and modelling errors, there are many such configurations possible, each one slightly different. It is clear from this example that a fixed safety margin may not adequately cover the possible variations in fiber shape.

With the visualization framework described in this chapter we attempt to visually communicate to the neurosurgeons that the fiber tracking algorithm they use may produce a sub-optimal reconstruction of the tracts of interest. We are specifically focusing on variations such as captured by probabilistic tracking algorithms. However, intuitively showing such information is not a trivial task. Even without considering uncertainty, diffusion tensor data presents considerable visualization challenges. For this reason, each tensor is often reduced to a single vector describing the principal direction of diffusion. Streamline visualization can then be used to show pathways through the tensor field. This gives a good impression of the global structure of fiber tracts. However, if there are too many streamlines this approach can lead to highly cluttered visualizations, such as illustrated in Figure 7.1(b). Probabilistic fiber tracking algorithms generally sample thousands of potential pathways in order to compute connection probabilities between different regions of the brain. In that case, it is no longer feasible, nor informative, to render each individual streamline.

To deal with this issue, we propose *Illustrative Confidence Intervals*, a rendering technique based on illustrative silhouettes and outlines that provides information about the variation in fiber pathways while at the same time reducing the visual clutter associated with standard streamline visualizations. We define a silhouette to be the interior area of an object's shape when projected to the viewing plane. An outline is defined as the

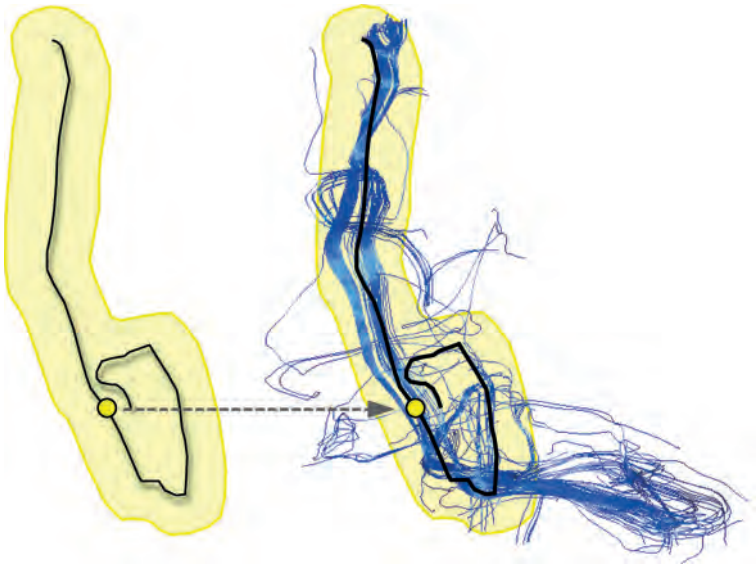

Figure 7.1: (a) Single fiber obtained from the original tensor volume. The yellow silhouette represents a fixed safety margin. (b) 100 variations of the same fiber based on the Wild Bootstrap method [71] with the same safety margin. It is clear that the safety margin does not cover the possible variations in the tensor data

border of a silhouette. Our method is based on the work by Otten et al. [118] but, as will be clear from Section 7.3, we extend this work in several ways. We use silhouettes and outlines to visualize variations in the output of probabilistic fiber tracking algorithms. Furthermore, we provide several interaction widgets that allow the specification of different fiber confidence intervals and visual styles for each interval. Employing a focus+context uncertainty lens the user can highlight uncertainty only in a specific region of interest, while the remaining view is rendered in a traditional manner.

Our framework does not depend on a specific algorithm to compute the variation in fiber pathways. It only requires a set of fibers, each associated with a so-called *confidence value*. A confidence value is a scalar-valued measure which represents the reproducibility of a given fiber pathway as calculated on the basis of probabilistic fiber tracking. It can be computed in different ways depending on the application. Our framework does not depend on a specific choice of confidence value and to demonstrate the flexibility of our approach we have selected two such measures, (1) based on fiber distance and (2) based on the connection probability between different brain regions.

We conclude this chapter with an informal user evaluation based on both live and video demonstrations, questionnaires and interviews with three neurosurgeons. For a short overview of existing visualization systems focused on neurosurgical applications we refer the reader to Chapter 2. Many of these systems provide extensive support to the neu-

rosurgical workflow but mostly do not take uncertainty into account. The techniques presented in this chapter can serve as a useful extension to fill this gap.

7.2 Computing fiber confidence

The visualization framework we propose requires a set of fibers (represented by streamlines) where each fiber is associated with a confidence value. Figure 7.2 gives an overview of the processing pipeline to compute fiber confidence from a set of diffusion-weighted images. Confidence values are stored in a table, normalized and sorted from high (1.0) to low (0.0) confidence. How we use the table to visualize fiber confidence intervals is described in the next section. As explained in Chapter 4, different algorithms exist to do probabilistic fiber tracking. All of these algorithms generate a large set of fiber pathways but, except for the ConTrack algorithm [142], none of them explicitly compute a confidence value for each fiber. For this reason, we compute such confidence values as a post-processing step after probabilistic fiber tracking. In this chapter, we use both the ConTrack algorithm as well as the Wild Bootstrap method [71] combined with a fiber distance measure. Details about the Wild Bootstrap and ConTrack algorithms can be found in Chapter 4.

7.2.1 Computing fiber distances

In this section we describe how to calculate fiber confidence values on the basis of Wild Bootstrap fiber tracking. For a given set of seed points, the end result of N repetitions of the Wild Bootstrap method is a set of N fiber pathways for each seed point. Each fiber pathway describes a single possible outcome of fiber tracking and the total set represents pathway variations originating from that seed point. We can consider each seed point independently because we do not change the seed region. To be able to assign a confidence value to each fiber we assume that the distribution of fibers in each seed point is uni-modal. In that case, we can choose a distance measure between pairs of fibers and use it to define a 'mean' fiber which is most central, given the distance measure, in the set of fibers at each seed point. We subsequently use each fiber's distance to the mean fiber as a representation of confidence. A confidence interval can then be expressed as: (1) a percentage (e.g., 50%) of the most confident fibers or (2) all fibers with a confidence higher than a given threshold confidence (e.g., 0.95). Which option is most appropriate depends on the user task and the distribution of confidence values in the total set of fibers. Our framework supports both options. Fiber distance measures have been extensively described in earlier work on fiber clustering [109] and comparison of fiber tracking algorithms [69]. The choice of distance measure is entirely dependent on the application within which it is used. Each measure has its advantages and disadvantages. For this reason, our framework is not dependent on the particular choice of distance measure in any way. To demonstrate its flexibility in this respect, we implemented two widely used measures: (1) the *Minimum End-Point Distance* and (2) the *Mean of Closest-Point Distances* [109, 36]. The minimum end-point distance d_E between two fibers F_i and F_j is defined as follows,

$$d_E(F_i, F_j) = \min(d_1, d_2) \tag{7.1}$$

where,

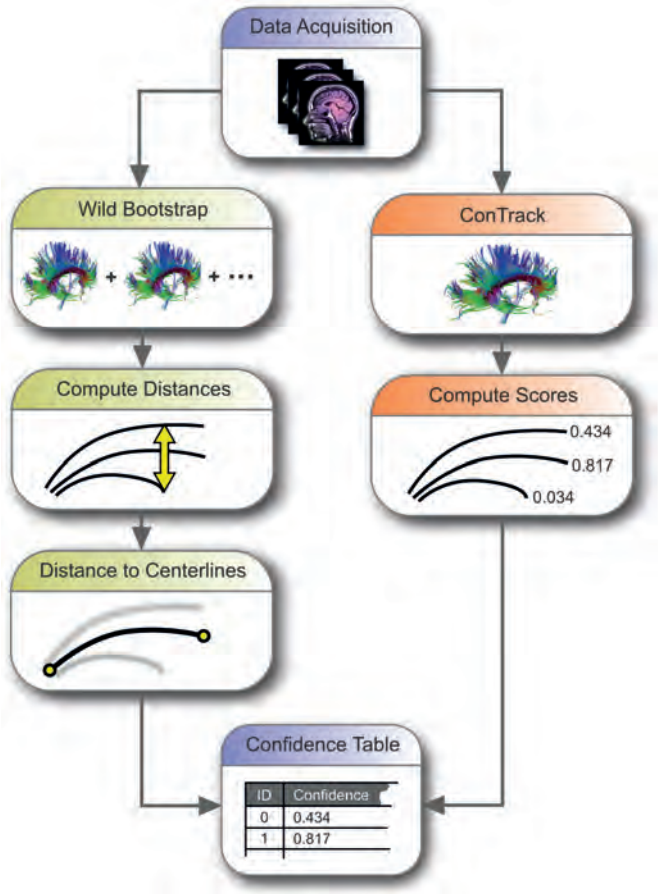

Figure 7.2: Pipeline for computing fiber confidence using (1) Wild Bootstrap and (2) ConTrack. The Wild Bootstrap method generates a large collection of unlabeled fibers which are assigned a confidence value by means of a fiber distance metric. The fibers produced by ConTrack already have a confidence value assigned to them.

$$d_1 = ||F_{i,1} - F_{j,1}|| + ||F_{i,end} - F_{j,end}|| \tag{7.2}$$

$$d_2 = ||F_{i,1} - F_{j,end}|| + ||F_{i,end} - F_{j,1}|| \tag{7.3}$$

where $F_{i,1}$ and $F_{i,end}$ refer to the first and last points on fiber F_i. The mean of closest-point distances d_M is defined as follows,

$$d_M(F_i, F_j) = \text{mean}(d_m(F_i, F_j), d_m(F_j, F_i)) \tag{7.4}$$

where, given that $\mathbf{p}_r$ and $\mathbf{p}_s$ are the points on each fiber,

$$d_m(F_i, F_j) = \underset{\mathbf{p}_r \in F_i, \mathbf{p}_s \in F_j}{\text{mean min}} ||\mathbf{p}_r - \mathbf{p}_s|| \tag{7.5}$$

Given the assumptions described previously, we define, for each seed point, the mean fiber to be the fiber with the smallest sum of distances to all other fibers at that seed point. This definition will result in an existing fiber and, given the chosen distance measure, this will be the most central fiber of the set. Our approach is similar to O'Donnell et al. [113] in that they also select an existing fiber instead of trying to compute an exact center line. The latter approach has been attempted in previous work [71, 46, 32], however, this can lead to ill-defined results when streamlines strongly diverge or even turn back on themselves. Merhof et al. [107] circumvent this problem by reverting to a smoothed, voxel-based representation of the fibers which allows them to generate a tight fitting surface mesh around the tracts. However, their approach does not take uncertainty into account.

Our definition of fiber distance assumes a uni-modal distribution of fibers originating from a seed point. Although this may be a valid assumption for some fiber tracts, it may be incorrect for others. In that case a different approach to compute fiber confidence must be chosen. Our framework allows easy integration of alternative methods. One such method will be described in the next section.

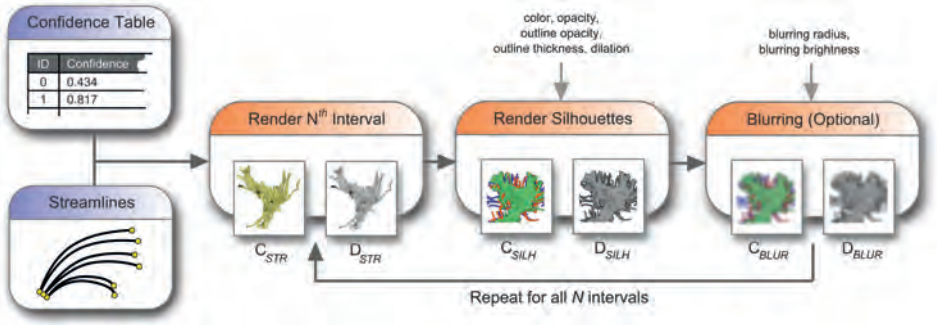

Figure 7.3: Pipeline for visualizing our illustrative confidence intervals

7.3 Visualization of fiber confidence

In the current work we have chosen to describe the uncertainty in fiber pathways in terms of discrete confidence intervals and represent the fibers in each interval by means of silhouettes and outlines. Confidence intervals provide a high-level overview of the distribution of fibers in the set, similar to confidence intervals used in descriptive statistics. Detail is lost by aggregating multiple fibers in a single interval but if multiple confidence intervals are shown, which our framework allows, the user can obtain a good impression of which parts of the fiber tract are most reliable or reproducible and which parts are highly variable. Showing the fibers by means of streamlines, even if each confidence interval would be colored differently, will result in a large amount of visual clutter. Also,

users are often not interested in individual fibers but rather in fiber *bundles*. A silhouette and outline representation fulfils these requirements by reducing visual clutter while preserving the overall shape of the fiber bundle. As explained before, we define a silhouette as the interior area of an object's shape when projected to the viewing plane. An outline is defined as the border of a silhouette. Otten et al. [118] presented an approach to render silhouettes and outlines for sets of clustered fibers using a GPU-accelerated algorithm. We extend this approach in the following ways:

- Extending the set of visual styles (e.g., transparency, color schemes, blurring, etc.)
- Making each visual style adaptive to interval confidence
- Combining an arbitrary number of confidence intervals in multiple render passes
- Adding anatomical context by means of 3D orthogonal slices, volume rendering and surface models (e.g., tumor, brain ventricles and cortical surface) and other types of streamline visualizations

Figure 7.3 gives an overview of the different stages in our visualization pipeline.

7.3.1 Generating silhouettes and outlines

To create a silhouette and outline representation we render the streamlines directly into an off-screen buffer image. Next, we apply a dilation operator to each non-empty pixel of the streamline image. The dilation operator is based on a circular structuring element with a user-defined silhouette radius R_s. This widens the area covered by the streamlines and closes holes between neighbouring streamlines. A second, larger structuring element with a user-defined outline radius R_o is used to add an outline of a given thickness to the silhouette. The general mechanism is illustrated in Figure 7.4.

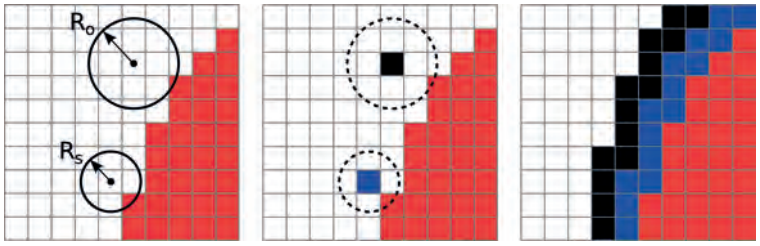

Figure 7.4: Creating a silhouette and outline. Blue pixels represent the added silhouette. Black pixels represent the added outline.

7.3.2 Rendering confidence intervals

To visualize the illustrative confidence intervals we transfer the streamlines and associated confidence table to our rendering pipeline (see Figure 7.3). We then specify a confidence interval $\Delta = [a, b]$, where $[a, b] \subseteq [0, 1]$, to control which streamlines are rendered to the screen as silhouettes and outlines. If the confidence interval is $[0, 1]$ we render *all* streamlines at once and the resulting silhouette represents all possible variations of streamlines. If we have N intervals, then each interval Δ_i where $0 \leq i < N$, can

be rendered with a different visual style. Figure 7.3 illustrates the different steps in our visualization pipeline. For each interval Δ_i these steps are executed as follows:

Step 1: We start with the interval of *highest* confidence, that is Δ_{N-1} and render the corresponding streamlines to a color buffer C_{STR} and depth buffer D_{STR} using OpenGL framebuffer objects. Depth testing is set to *GL_LESS*.

Step 2: In this step, stencil testing is enabled (function: *(GL_EQUAL,0,1)*, operation: *GL_INCR*) to ensure that subsequent intervals (with lower confidence) do not overwrite the current one (with higher confidence). We could have solved this by setting the depth test to *GL_ALWAYS* and rendering the low confidence intervals first. However, reducing the opacity of a given silhouette would then result in unwanted color mixing of the different intervals. If the opacity is reduced, we want to see the *anatomy* lying behind it and not other confidence intervals. The buffers C_{STR} and D_{STR} from the previous step are passed to a GPU shader program that creates the silhouette and outline representation described previously. The output result of the shader program is written to the color buffer C_{SILH} and depth buffer D_{SILH}. In this step the visual style parameters are applied (see also Section 7.3.3).

Step 3 (Optional): This step is only performed if silhouette blurring is enabled. This allows the user to show intervals of lower confidence with increased blurring. The stencil test is now performed at this step instead of Step 2. The buffers C_{SILH} and D_{SILH} from Step 2 are passed to a second GPU shader program that applies Gaussian blurring to the color buffer. The shader program takes a blurring radius and brightness offset as parameters (see also Section 7.3.3). The brightness offset allows the user to adjust brightness in case it is too low after the blurring operation.

7.3.3 Confidence histogram widget

Depending on the distribution of confidence values within the fiber set, the user may wish to choose the confidence intervals differently. To make this possible we introduce the *confidence histogram widget* that shows a histogram of confidence values and allows the user to define and manipulate confidence intervals graphically. Additionally, it is possible to identify deviations from a uni-modal distribution in the fibers. In Section 7.2 we explained how we use distance to compute fiber confidence. This approach assumes that the distribution of fibers originating from a seed point is uni-modal. If this is not the case and the fibers split up into two, equally-sized bundles, the mean fiber may be assigned randomly to one or the other bundle. All fibers *not* in the same bundle as the mean fiber will have a larger distance and are therefore assigned a lower confidence. In that case, the histogram is likely to show two peaks. Figure 7.5(a) illustrates an example histogram computed from output scores of the ConTrack algorithm. The confidence intervals are displayed as a set of semi-transparent rectangles overlaid on top of the histogram. Each rectangle represents an interval. Rectangle colors map to silhouette or outline colors, depending on the current selection mode. Rectangle width maps to interval width. Rectangle height maps to a selected visual property value, such as opacity or blurring radius. The full list of visual properties available in our framework is given in Table 7.1.

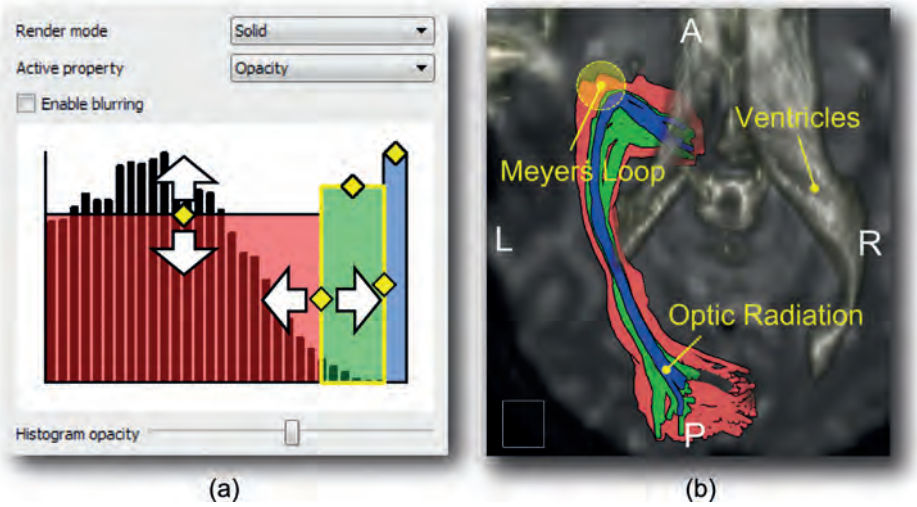

Figure 7.5: (a) Confidence Histogram Widget with confidence values increasing from left to right. White arrows indicate possible manipulation directions. Selected active property is 'Opacity'. (b) Corresponding confidence intervals of optic radiation fibers looping around ventricles.

Visual properties can be set for each confidence interval separately or they can be assigned automatically through a number of pre-sets. For example, the user can assign a number of standard color scales to the different intervals, such as warm to cool, light to dark or decreasing saturation. Scalar properties such as opacity can be assigned in a staircase or inverted staircase pattern. Intervals can be automatically subdivided in equal widths, e.g. $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$. Alternatively, a subdivision can be selected where each interval contains the same number of fibers. These two options reflect the two different perspectives we discussed in Section 7.2.1 where a confidence interval can be expressed as either (1) a percentage of the most confident fibers or (2) all fibers with a confidence between a certain range. As explained previously, which option is most appropriate depends on the user task.

Finally, the confidence histogram widget allows saving and loading of visual styles. From a practical point of view we do not expect the neurosurgeon to set these visual properties each time he or she uses the tool. Automatic pre-sets, saving and loading of visual properties are specifically designed to make this process easier and less repetitive. However to be able to properly evaluate the benefits and drawbacks of many different visual styles (as described in Section 7.4.2) we also provide low-level control of these properties.

7.3.4 Uncertainty lens

Uncertainty in DTI fiber tracking algorithms is an important issue in neurosurgical applications. However, some fiber tracts can be more reliably reconstructed than others, even with deterministic approaches. Also, not all regions of the brain are affected by

Table 7.1: User-adjustable properties

Property	Description
Color	Color of inner silhouette area
Outline color	Color of silhouette outline
Opacity	Opacity of inner silhoutte area
Outline opacity	Opacity of silhouette outline
Outline thickness	Thickness in pixels of silhouette outline
Dilation	Half-width in pixels of silhouette
Blurring	Blurring enabled/disabled
Blurring radius	Gaussian blurring kernel (2σ)
Blurring brightness	Scaling factor to adjust brightness of blurred regions

the surgical approach. In such regions, showing uncertainty may be unnecessary or even confusing. For this reason, we provide the *Uncertainty Lens* as a Focus+Context approach to show uncertainty only within a user-defined region-of-interest (ROI). Outside the lens standard streamline visualization can be used. For example, the surgeon can place the uncertainty lens over the tumor and its immediate surroundings and see the fiber variations due to the uncertainty within a relevant context. If necessary, the user can interactively move and resize the uncertainty lens. Figure 7.6(c) gives an example visualization.

7.4 Results and discussion

Figure 7.6 illustrates uncertainty visualizations created with our framework. Figure 7.6(a) depicts a transversal view of the optic radiation (running from the thalamus to the visual cortex) with brain ventricles rendered using DVR. Figure 7.6(b) provides a sagittal view of the corticospinal tract surrounding a tumor. Figure 7.6(c) illustrates our *Uncertainty Lens* together with a standard streamtube visualization, tumor and semi-transparent cortical surface. Figures 7.6(d-i) illustrate a close-up of the optic radiation with most of the visual styles we support (warm-to-cool outlines, warm-to-cool surfaces, decreasing opacity, light-to-dark, increasing dilation, increasing blur).

7.4.1 Datasets and performance

These visualizations were created with datasets from three suppliers. **Kempenhaeghe Epilepsy Center, Heeze, NL:** DTI ($112 \times 112 \times 60$, $2 \times 2 \times 2\text{mm}^3$, b-value 1000, 30 grad. dirs.). **St. Elisabeth hospital, Tilburg, NL:** DTI ($128 \times 128 \times 60$, $1.75 \times 1.75 \times 2$ mm^3, b-value 800, 30 grad. dirs.), T1 MRI ($288 \times 288 \times 175$, $1 \times 1 \times 1\text{mm}^3$). **VisContest 2010:** DTI ($128 \times 128 \times 72$, $1.8 \times 1.8 \times 1.8\text{mm}^3$, b-value 1000, 30 grad. dirs.), T1 MRI ($512 \times 512 \times 176$, $0.488 \times 0.488 \times 1\text{mm}^3$). We measured the framerate performance of our rendering pipeline using a GeForce 8800 GTX graphics card from NVIDIA and a screensize of 720×760 pixels: Figure 7.6(a) = 6.6 fps, Figure 7.6(b) = 10.2 fps and Figure 7.6(c) = 17.5 fps.

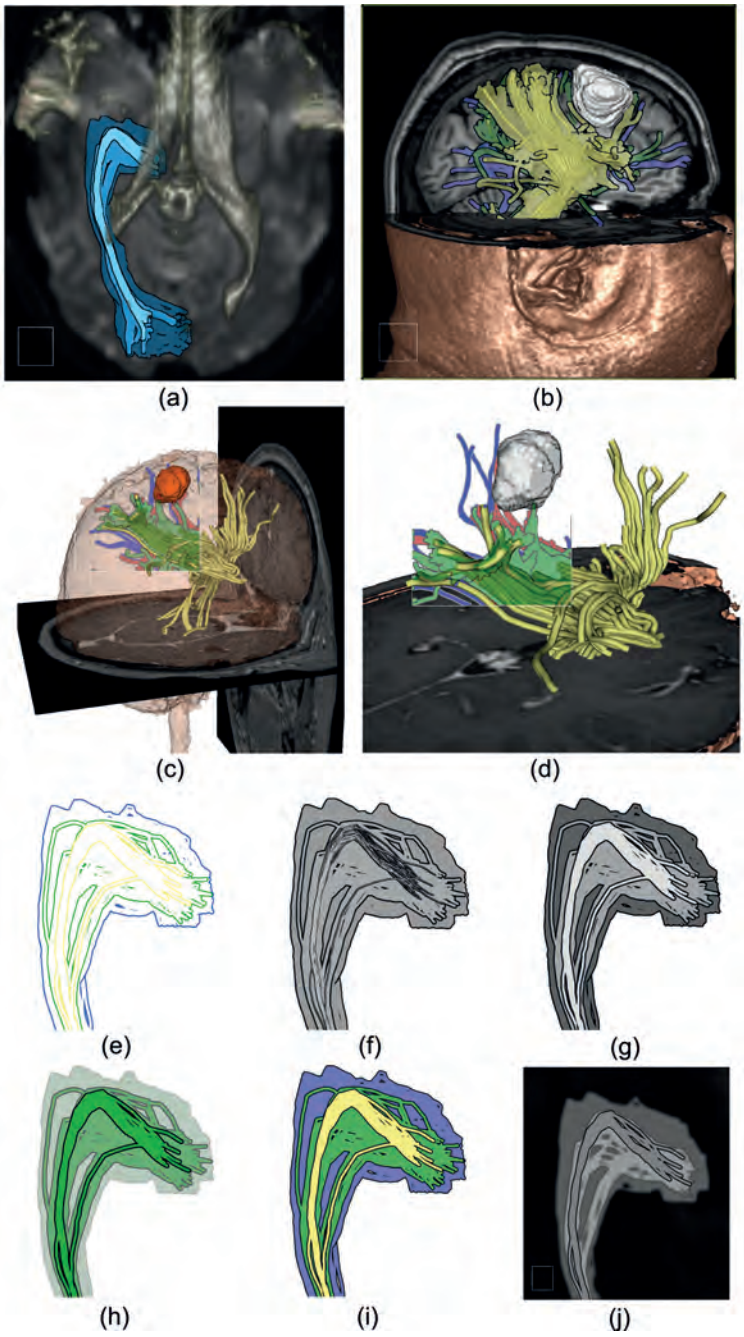

Figure 7.6: (a) Optic radiation, (b) pyramidal tract with tumor, (c) uncertainty lens, (d-i) different visual styles applied to optic radiation (warm-to-cool outlines, warm-to-cool surfaces, decreasing opacity, light-to-dark, increasing dilation, increasing blur).

7.4.2 User evaluation

To evaluate our visualization approach we conducted an informal user study with three neurosurgeons. First, we wanted to know their general opinion about the need and benefits of uncertainty visualization for their work. Second, we wanted to evaluate the specific rendering options that our visualization framework provides. We presented the neurosurgeons with a set of screen-shots, videos and live demonstrations. Furthermore, we asked them to fill out a questionnaire with respect to the visual material. Three clinical scenarios were presented: (1) anterior temporal lobe resection (ATLR) for the treatment of focal epilepsy, (2) resection of a low-grade glioma close to motor cortex and (3) resection of an intra-cerebral metastasis with surrounding edema (fluid), also close to motor cortex. For each scenario we showed a standard fiber visualization (stream tubes) and our illustrative confidence intervals. We used three intervals each containing 10%, 25% and 100% of the most confident fibers. Table 7.2 shows the general opinion of the neurosurgeons on the questions listed below. Table 7.3 shows the specific ratings for each question. Ratings are given on a $[1-5]$ Likert scale where 1 means either 'not useful at all' or 'low' and and 5 means 'highly useful' or 'high', depending on the context of the question.

Q1.1 What is the risk of a visual field deficit in anterior temporal lobe resection?
Q1.2 How important is uncertainty visualization for anterior temporal lobe resections?
Q1.3 What is the risk of resecting this particular low-grade glioma?
Q1.4 How important is uncertainty visualization for low-grade glioma resection?
Q1.5 What is the risk of resecting this particular intracerebral metastasis?
Q1.6 To what extent would you discuss the uncertainty with the patient?

Table 7.2: General opinion about DTI uncertainty visualization. Ranking of visual styles rating is done on a $[1-5]$ Likert scale.

	User A	**User B**	**User C**
Q1.1	2	2	2
Q1.2	4	4	3
Q1.3	4	1	1
Q1.4	4	4	5
Q1.5	2	2	2
Q1.6	5	4	4

Q2.1 To what extent does the standard fiber visualization give you more confidence in assessing risk of anterior temporal lobe resections?
Q2.2 To what extent does the standard visualization give more confidence in assessing risk of resecting the given low-grade glioma?
Q2.3 To what extent does the standard visualization give more confidence in assessing risk of resecting the given intracerebral metastasis?
Q2.4 Which two visual styles do you prefer?
Q2.5 Does our representation of uncertainty give you more confidence?
Q2.6 How useful is it to show uncertainty only in a selected ROI?

Q2.7 What is the overall rating of the potential use of our visualization?

The standard fiber visualization was considered somewhat helpful for scenario 1 and 2. It gives at least a rough indication of the location of the fiber tract. For scenario 3 all users gave a low rating because fibers seemed to be missing inside the edema that, in their opinion, should have been present. This is caused by a failure of the fiber tracking algorithm to trace into low anisotropy regions such as edematous fluid. We presented the users with different visual styles for the illustrative silhouettes and asked them which style or combination of styles most clearly communicated the different levels of confidence. Initially, two out of three users rated decreasing opacity (with decreasing confidence) to be most intuitive. Light-to-dark, non-gray coloring came second. After they finished the questionnaire we discussed this first choice with the users and explained that reducing opacity actually removes information instead of showing it. They agreed but indicated that they selected decreasing opacity mainly because it prevents occlusion of the underlying anatomical slices (which is important for context). After some discussion, they proposed that a combination of light-to-dark coloring and a reduced, but fixed, opacity would be a good alternative for representing fiber confidence. The third user, interestingly, preferred equal colors for all intervals but with decreasing outline thickness (going from fat to thin to zero). All other visual styles (silhouette dilation, blurring and warm-to-cool colors) were less appealing to the users. Our Focus+Context uncertainty lens was considered to be useful. The initial uncertainty visualization (without uncertainty lens) looked rather intimidating to them. The lens reduces visual clutter and allows attention to be focused on the tumor and its immediate surroundings. Also, the fully opaque visual styles were considered to be much more acceptable this way.

Table 7.3: Opinions about the proposed uncertainty visualization. Except for the ranking of visual styles rating is done on a $[1-5]$ Likert scale. Scores are averages over all the three scenarios.

	User A	**User B**	**User C**
Q2.1	4	2	3
Q2.2	3	2	2
Q2.3	1	1	1
Q2.4	Opacity Lightness	Outline -	Opacity Lightness
Q2.5	4	4	4
Q2.6	4	4	3
Q2.7	4	4	4

7.4.3 General discussion

As noted in the user evaluation, the neurosurgeons were skeptical about the fiber tracking results for the intra-cerebral metastasis. The edematous fluid surrounding the tumor causes the tracking algorithm to prematurely terminate thereby resulting in false negatives. Our particular implementation of the Wild Bootstrap method will also suffer from such false negatives (failing to show something that is actually there) because it uses

the same fiber tracking algorithm with the same termination criteria. If diffusion in a voxel, because of edematous fluid, is almost isotropic, the tensor model will fit the data well and have very small residuals. Random perturbation of the residuals will not result in large shape variations, so on average the tensors will remain isotropic. Because the Wild Bootstrap method uses the same thresholds as standard streamline tracing, none of the generated streamlines will get past these isotropic tensors. With crossing fibers however, the tensor can become disk-like while the underlying diffusion profile might actually be cross-shaped. In this case, the fractional anisotropy of the tensor remains high which may cause standard algorithms to continue tracking even though the model fit is bad (large residuals). The Wild Bootstrap will expose this problem by showing a large amount of variation in fiber pathways originating from the crossing fiber region. This indicates that the standard tracking algorithm may be generating random results. In order to deal with false negatives, alternative methods for probabilistic fiber tracking can be used. These are often based on Monte Carlo sampling of pathways and less dependent on stopping criteria [50, 11]. Even in regions of low anisotropy these algorithms can detect the presence of fiber pathways.

In this chapter we have applied the Wild Bootstrap approach using the 2nd-order tensor model and deterministic streamline tracing. However, these choices are not specific to the Wild Bootstrap. It is a general technique that can be applied to any model that tries to fit data and results in non-zero residuals. Even with higher-order models that are able to detect crossing fibers, you can still use the Wild Bootstrap to investigate the sensitivity of these models to noise in the data. Our choice to use streamline tracing is also not the only choice. Other fiber tracking algorithms, such as tensor deflection (TEND) [89], can be used in combination with the Wild Bootstrap as well.

With respect to the computation of confidence based on fiber distance we wish to point out that the choice of distance measure may have a large effect on the confidence intervals. As Jiao et al. [69] point out distance measures that are based on averages, such as the *mean of closest-point distances*, may over- or under-estimate the true distance due to streamline discretization problems or complex streamline configurations. The *closest end-points distance* may surely over-simplify the situation. However, for certain cases it may provide relevant information. For example, if apriori knowledge is available that fibers start and end in the same anatomical regions, they could still be considered 'close' even though they may take wildly different routes to get to their destination. In general we regard distance measures to be application-specific without any particular measure being the best one for all situations.

The user evaluation, although not conclusive, has provided us with useful information. Whether DTI fiber tracking, and therefore DTI uncertainty, is relevant for neurosurgical planning depends on many factors, such as tumor type. High-grade gliomas have a bad prognosis and are commonly associated with functional deficits in the patient. If there is any risk of additional damage due to resection, tumor tissue is simply left in place. Low-grade gliomas, on the other hand, have a relatively good prognosis. The tumor grows slowly allowing brain functions to reorganize if they get invaded (brain plasticity). In these cases neurosurgeons are willing to use all the information available in order to maximize resection while at the same time minimizing damage to healthy tissue. DTI

fiber tracking and uncertainty visualization are especially relevant here. For anterior temporal lobe resections DTI uncertainty information is important because part of the optic radiation, Meyer's Loop, is often damaged during the procedure, resulting in a partial loss of vision. The extent of Meyer's Loop varies significantly between patients (between 3 and 6 cm's from the temporal pole). This makes it difficult to predict in advance whether visual defects will result or not. Meyer's Loop however is a difficult fiber tract to reconstruct due to strong curvature and proximity to other fiber tracts (resulting in partial volume effects at the tract boundaries). Uncertainty visualization can provide helpful information in this case.

7.5 Conclusions

In the previous sections we have presented a novel method for visualizing the output of probabilistic fiber tracking in the form of fiber confidence intervals. Our method is based on an illustrative rendering algorithm using silhouettes and outlines to reduce the visual clutter commonly associated with the large numbers of fibers produced by probabilistic tracking algorithms. In the next chapter we will present our third major contribution, the exploration of distances between objects of varying shape such as white matter fibers reconstructed from probabilistic tracking.

If you have trouble sounding condescending, find a Unix user to show you how it's done.

- Dilbert

8 Distance measurements

This chapter is based on:

"Visualization of the Effects of DTI Uncertainty on Distance Measurements in Neurosurgical Planning." Ralph Brecheisen, Bram Platel, Bart ter Haar Romeny, and Anna Vilanova. Submitted for publication.

8.1 Introduction

The previous chapter presented techniques for the visualization of the effects of noise and model errors calculated from probabilistic fiber tracking algorithms. Such algorithms effectively reconstruct white matter fiber bundles whose spatial extent is defined in a probabilistic manner. Even though most existing visualizations of probabilistic fiber tracking are based on 3D scalar-valued fiber density maps, we have shown that it is feasible to work with the fiber pathways directly and render them as illustrative confidence intervals. However, since these confidence intervals are based on image-space rendering they are less suitable for calculating 3D distances. Scalar density fields have the advantage that they allow the reconstruction of 3D isosurfaces. Using GPU-based raycasting techniques these isosurfaces can be interactively changed. In case of scalar fields resulting from probabilistic fiber tracking, the isosurfaces represent different confidence levels in the spatial extent of the fibers. This representation allows an interesting feature which has received very little attention so far, namely distance measurements between volumetric objects with uncertain, or spatially varying, extents. For example, when volumetric objects are considered at different confidence levels, associated with different isosurfaces, not only the magnitude of the minimal distance between the objects may change but also the location where a minimal distance occurs. This is illustrated in Figure 8.1.

In neurosurgical planning it is very common to perform measurements or visual assessments of distance, for example between a tumor and surrounding brain structures. Minimal distances between such objects are an important factor in the assessment of surgical risk. Normally, distance assessments are performed on the basis of 2D anatomical images and knowledge of classical functional anatomy. If the target is judged to be close to critical areas the neurosurgeon can refine judgement by means of additional functional neuroimaging techniques such as functional MRI (f-MRI) and DTI. Although these non-invasive techniques give promising results, the gold standard remains electrocortical stimulation mapping (ESM). This is an invasive surgical technique where small electrical currents are applied at different locations in the brain (e.g. on the cortical surface or inside the resection cavity) while the patient is monitored. For instance, stimulating the primary motor cortex can induce muscular contractions. ESM also has a number of drawbacks. It requires a surgical procedure and may cause epileptic seizures [15]. Furthermore, it only allows probing the white matter at a limited number of locations. This makes it difficult to mentally visualize the shape of fiber tracts. Also, the elicited response of the patient is dependent on the individual current applied [81] and allows only a rough estimate of distance between electrode tip and the involved fiber tracts [76]. Incorporating quantitative distance measurements in the preoperative planning can complement the ESM procedure and make it more efficient because locations of minimal distance are known beforehand [14].

In this chapter we present a general visualization framework that allows the exploration of distances between volumetric objects while taking the uncertainty in their spatial extent into account. Existing approaches for measuring and visualizing 3D distances do not take uncertainty into account. Neither do they allow the user to interactively explore the effects of uncertainty on such measurements. In this chapter, we propose a distance analysis framework that offers a number of visualization and interaction features that

deal with these issues. Our main contribution is the incorporation of uncertainty in the exploration of distances between 3D objects. To allow interactive exploration we implemented a fast, GPU-based algorithm for computing distance transforms on large volumes. We provide an integrated user interface that combines 2D and 3D views by means of brushing and linking.

We will start this chapter with an analysis of specific clinical questions and requirements with respect to distance measurements in neurosurgical practice. We then proceed with an explanation of the computational features of our framework, followed by visualization features such as minimal distance graph, tumor distance map and assisted camera positioning.

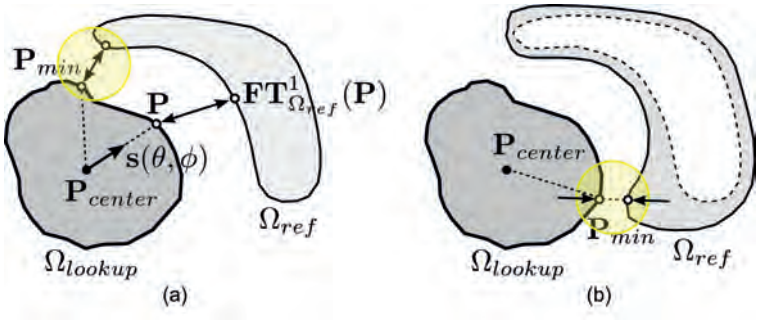

Figure 8.1: (a) Distance relation between lookup object Ω_{lookup} and reference object Ω_{ref}. The point of minimal distance $\mathbf{P}_{min}$ is highlighted in yellow. (b) If the reference object Ω_{ref} changes shape as indicated, the point of minimal distance $\mathbf{P}_{min}$ moves to a different location.

8.2 Related work

For related work on fiber tracking, the characterization of fiber uncertainty and visualization systems for neurosurgical planning, we refer the reader to Chapters 2 and 4. Here, we provide a short overview of work done in the area of quantitative measurements. Quantitative analysis of 3D spatial relations is an important task in both diagnostic and surgical practice. Preim et al. [130] propose a number of 3D measurement widgets and discuss strategies for interacting with these widgets using a mouse and keyboard. Reitinger et al. [131] propose an immersive virtual reality environment where 3D measurements can be performed using hand and arm movements. Both approaches rely on manual positioning of the measurement widgets in 3D space which can introduce a considerable amount of error and variation between users. Roessling et al. [138] address this issue by proposing an automatic algorithm for computing distances between presegmented objects. The framework is based on the construction of a spatial search tree which hierarchically decomposes the mesh surfaces into subsets of geometric primitives. It is highly flexible and can be used for a variety of multi-object configurations. It does not, however, support real-time changes in the spatial extent of objects.

8.3 Clinical questions and requirements

To understand how distance measurements impact surgical risk assessment we presented the neurosurgeons with a set of problem scenarios [139] describing potential ways of working with our software tool. Based on their feedback regarding feasibility and clinical relevance we defined a set of questions that the tool should help to answer.

Q1 Which tumor parts are close to surrounding critical structures?
Q2 What is the magnitude and location of minimal distance to these structures?
Q3 What is the effect of fiber uncertainty on the distance measurement?
Q4 What is the effect of variations in tumor segmentation?

Our framework provides the following set of features for answering these questions:

- **Real-time feature transform** (Q3, Q4) Allows interactive exploration of the effect of object uncertainty on distance measurements (Section 8.5).
- **Minimal distance graph** (Q3, Q4) Shows minimal distance as a function of object uncertainty (Section 8.6.1).
- **3D distance arrow** (Q2) Shows magnitude and location of distances in the 3D scene (Section 8.6.1).
- **3D tumor distance map** (Q1, Q2) Shows distances on the 3D tumor surface using different color maps (Section 8.6.2).
- **2D tumor distance map** (Q1, Q2, Q3) Shows distances on the tumor surface unfolded to a 2D map (Section 8.6.2).
- **Automatic camera positioning** (Q1, Q2) Selecting a point in the 2D tumor distance map automatically re-orients the camera in 3D such that the corresponding point on the 3D tumor surface is brought to the front (Section 8.6.3).

8.4 Pre-processing pipeline

Our computation pipeline starts with a pre-processing stage that is specific for the application domain. It involves tumor segmentation, fiber tractography and, if necessary, conversion of fiber pathways to a scalar *reproducibility map*. We define this term as follows. Fiber tract uncertainty can be expressed in different ways depending on the computational method used. Even though the term probability is often used, we prefer the term reproducibility because it better describes what is actually being computed. We distinguish two types of tractography methods: voxel-centric and fiber-centric. Voxel-centric methods are the most common approach and assign a reproducibility value to each voxel describing how often this voxel is found to be connected to the seed region. They do not explain how or by which pathway a voxel is connected. Fiber-centric methods assign a reproducibility value to complete pathways thereby describing how often a pathway interconnects the seed and target regions. An example is the ConTrack algorithm explained in Chapter 4 [142]. Our choice of volumetric input format allows us to support data obtained from both voxel- and fiber-centric methods. It also allows us to explore fiber uncertainty by visualizing different iso-surfaces in the scalar map. The tumor is represented by a single binary map or multiple binary maps (e.g. segmentations obtained from multiple users or image modalities). The remainder of the pipeline is independent of the particular application and interactively computes distance information between a given pair of volumetric objects, provided they are represented

as reproducibility maps.

To characterize fiber uncertainty we tested three well-known probabilistic tractography algorithms. The PICo [122] and Wild Bootstrap [71] algorithms are voxel-centric approaches. Their output is a scalar map with values that can be directly interpreted as reproducibilities and loaded into our framework without conversion. As mentioned earlier, the ConTrack algorithm [142] is a fiber-centric method. It generates fiber pathway geometry and requires conversion to a reproducibility map. This is done as follows:

Algorithm 1 Generate fiber-centric reproducibility map

```
initVoxelValuesToZero()
for all fibers F with reproducibility R_F do
  for all voxels V along fiber F and with reproducibility R_V do
    if R_F > R_V then
      R_V = R_F
    end if
  end for
end for
```

Once datasets are loaded as reproducibility maps, we can explore variations in the objects' spatial extent by visualizing iso-surfaces for different reproducibility thresholds.

8.5 Real-time feature transforms

Let $U : \mathbb{R}^3 \to \mathbb{R}$ be a function representing an uncertainty volume. For a given threshold T, let $\Omega(T) = \{\mathbf{P} \in \mathbb{R}^3 | U(\mathbf{P}) > T\}$ be an object of interest. A well-established approach to obtain distance information surrounding an object $\Omega(T)$ is the Euclidian distance transform $DT_{\Omega(T)} : \mathbb{R}^3 \to \mathbb{R}^+$ defined as:

$$DT_{\Omega(T)}(\mathbf{P}) = \min_{\mathbf{Q} \in \Omega(T)} \|\mathbf{P} - \mathbf{Q}\| \tag{8.1}$$

An extension of the distance transform is the so-called *feature transform* $\mathbf{FT}_{\Omega(T)} : \mathbb{R}^3 \to \mathcal{P}(\partial\Omega(T))$, where $\mathcal{P}(\partial\Omega(T))$ represents the (power) set of points on the boundary of $\Omega(T)$ and is defined as:

$$\mathbf{FT}_{\Omega(T)}(\mathbf{P}) = \{\mathbf{Q} \in \partial\Omega(T) | \|\mathbf{P} - \mathbf{Q}\| = DT_{\Omega(T)}(\mathbf{P})\} = \arg\min_{\mathbf{Q} \in \partial\Omega(T)} \|\mathbf{P} - \mathbf{Q}\| \tag{8.2}$$

On a voxel grid, $DT_{\Omega(T)}$ gives the distance from each voxel $\mathbf{P}$ in a dataset to the nearest voxel $\mathbf{Q}$ of the surface $\partial\Omega(T)$ of some object $\Omega(T)$. $\mathbf{FT}_{\Omega(T)}$ gives the set of surface voxels $\mathbf{Q}$ closest to each voxel $\mathbf{P}$, also called the *features* of $\mathbf{P}$. The *one-point* feature transform $\mathbf{FT}^1_{\Omega(T)} : \mathbb{R}^3 \to \partial\Omega(T)$ is a simplification of the feature transform in Eqn. 8.2 where a single feature is recorded per point. In the following, we shall use $\mathbf{FT}^1_{\Omega(T)}$ since it gives us *both* the identity of feature points, which we need, and also the distance to these feature points. Furthermore, it is faster to compute and requires less memory than the full feature transform $\mathbf{FT}_{\Omega(T)}$.

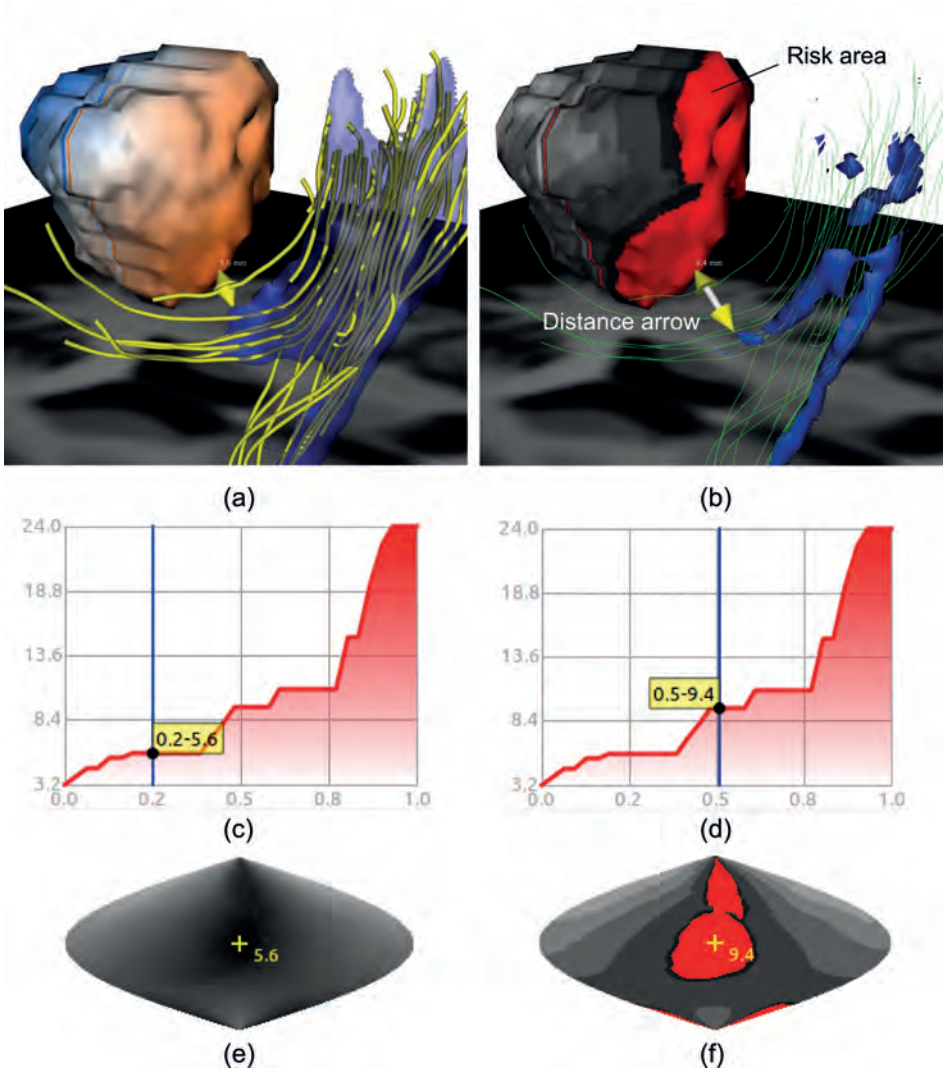

Figure 8.2: (a) Tumor with warm-to-cool distance coloring and fibers visualized as streamtubes (yellow) and a 20% reproducibility surface. (b) Tumor with grayscale distance color and highlighted risk area (red). Fibers are visualized as streamlines (yellow) and a 50% reproducibility surface. (c,e) Minimal distance graph and tumor distance map for tumor displayed in (a). (d,f) Minimal distance graph and tumor distance map for tumor displayed in (b).

After pre-processing our framework takes two inputs: (1) a *reference* object $\Omega_{ref}(T_{ref})$, abbreviated Ω_{ref}, for which the feature transform is computed, and (2) a *lookup* object $\Omega_{lookup}(T_{lookup})$, abbreviated Ω_{lookup}, for which we look up feature points in Ω_{ref}.

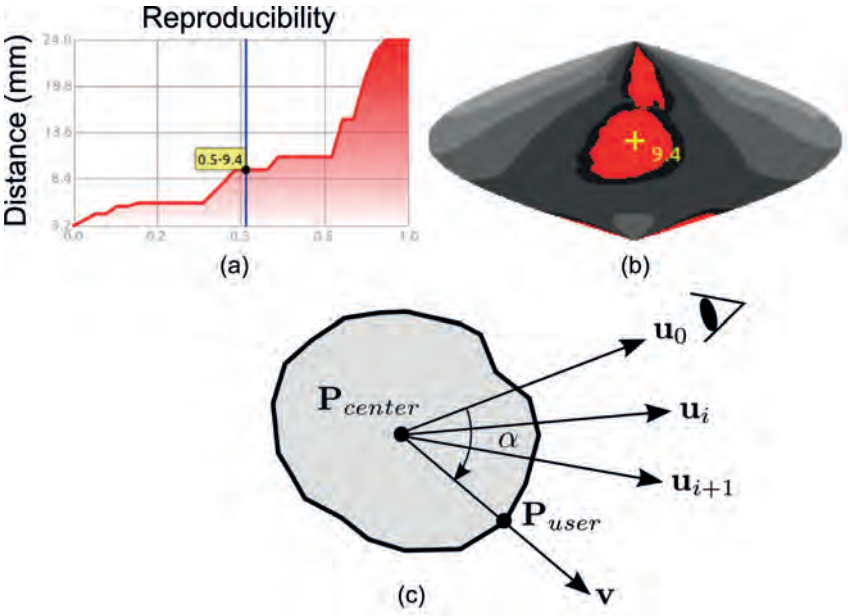

Figure 8.3: (a) Minimal distance graph showing minimal distance as a function of reproducibility. (b) Sinusoidally projected tumor distance map. (c) Repositioning camera view from $\mathbf{u}_0$ to $\mathbf{v}$ in incremental steps.

Distance and (one-point) feature transforms can be computed using several algorithms. On the CPU, we refer to the Augmented Fast Marching Method (AFMM) [147], the scanline-based algorithm of Meijster et al. [106], the raster scanning method of Mullikin et al. [111], the Euclidean tolerance feature transform [132], the graph-searching method of Lotufo et al. [99] and the recent Euclidian scanning algorithm of Wang and Tan [160]. Although these methods have low computational complexity (typically $O(N \log N)$ and at best $O(N)$ [106] for N lookup voxels), they are still too slow for our purposes. We need interactive, on-the-fly recomputation of these transforms in order to allow the user to explore distances to objects with changing spatial extent. For this reason we use a GPU-based algorithm proposed by Telea et al. [146], which extends the exact Euclidean distance transform method of Cao et al. [30] to compute one-point feature transforms of voxel objects. To speed up computation, this implementation uses only the voxels representing the boundary $\partial\Omega(T)$ rather than all voxels in $\Omega(T)$ itself. Boundary voxels are those voxels having at least one background neighbor voxel, using 6-connectivity. This results in closed voxel boundaries, i.e., it guarantees that the centers of lookup voxels can never be closer to interior (non-boundary) voxels of $\Omega(T)$ than to boundary voxels. In that case, the computed distance and feature transforms are identical to the more expensive approach where we would consider all voxels in Ω_T. Note that this assumes non-intersecting objects. The boundary voxels in both Ω_{lookup} and Ω_{ref} are selected based on reproducibility iso-surfaces defined by thresholds T_{lookup} and T_{ref}. In the neurosurgical setting we are considering in this chapter, we use the tumor as lookup object Ω_{lookup} and the fibers as reference object Ω_{ref}. This decision is based on question

Q1 described in Section 8.3, which requires that we show *parts of the tumor* that are close to the fibers. This makes sense because it is the tumor that is being resected. We therefore show distances on the tumor surface which implies that we compute the feature transform on the fibers.

8.6 Visualization tools

This section introduces the main visualization tools included in the framework. These are (1) the minimal distance, (2) the tumor distance map, (3) assisted camera positioning and (4) additional interaction features.

8.6.1 Minimal distance graph

The ability to compute feature transforms in real-time allows the user to interactively explore the effects of fiber (or tumor) uncertainty on minimal distance (question Q3 and Q4). For this purpose we propose the *Minimal Distance Graph* which shows minimal distance as a function of reproducibility. An example is illustrated in Figure 8.3A. The graph is constructed by sampling the range of reproducibility values at regular intervals. For each sample, we compute:

$$\mathbf{P}_{min}(\Omega_{ref}) = \underset{\mathbf{P}\in\partial\Omega_{lookup}}{\arg\min} \|\mathbf{FT}^1_{\Omega_{ref}}(\mathbf{P}) - \mathbf{P}\| \tag{8.3}$$

where $\mathbf{P}_{min}$ is the point on $\partial\Omega_{lookup}$ where a minimal distance to $\partial\Omega_{ref}$ occurs. The minimal distance itself is computed from:

$$D_{min}(\mathbf{P}_{min}) = \|\mathbf{FT}^1_{\Omega_{ref}}(\mathbf{P}_{min}) - \mathbf{P}_{min}\| \tag{8.4}$$

The graph is constructed from a fixed number of sampled reproducibility thresholds T_{ref}. To query thresholds that lie between samples, the user can select a point anywhere on the horizontal reproducibility axis. This will trigger re-execution of the computation pipeline. The recomputed minimal distance is displayed as a label on the graph curve. A distance arrow in shown in 3D with end-points corresponding to $\mathbf{P}_{min}$ and its corresponding feature point. Point selection also triggers recomputation of the *tumor distance map* described in the next section.

8.6.2 Tumor distance map

One of the aims of our distance analysis framework is to show which parts of the tumor, represented by Ω_{lookup}, are closest to the surrounding fibers, represented by Ω_{ref}. To do this, we compute distance information across the object surface $\partial\Omega_{lookup}$. We have chosen to represent Ω_{lookup} as a 3D mesh and apply the distance information as a color-coded 2D texture. For this, the texture has to be parametrized with texture coordinates (u, v). We assume that $\partial\Omega_{lookup}$ has a star convex shape, meaning each point on $\partial\Omega_{lookup}$ has a one-to-one mapping to a unit sphere positioned at its center of gravity $\mathbf{P}_{center}$ (see Figure 8.1).

Based on the 2D texture we can easily highlight areas where distance values are below a pre-defined threshold. The resulting areas are called *risk areas* because they show, at a single glance, which parts of the tumor surface are closest to the fibers. The pre-defined

threshold, called the *risk distance*, is a user-defined parameter. For example, the neurosurgeon may wish to see which areas of the tumor are at a distance of 10.0 mm or less from any surrounding fibers.

To create the 2D texture, we compute a distance map which stores closest-point distances from $\partial\Omega_{lookup}$ (the tumor) to $\partial\Omega_{ref}$ (the fibers). This requires a function $Q : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ which transforms map coordinates (u, v) to coordinates (θ, ϕ) on the unit sphere. Applying the 2D texture to the 3D mesh requires only a trivial mapping $\theta = u$ and $\phi = v$. For each coordinate (θ, ϕ), we start at $\mathbf{P}_{center}$ and perform ray-casting along a ray $\mathbf{R}(\lambda) = \mathbf{P}_{center} + \lambda\ \mathbf{s}(\theta, \phi)$ where

$$\mathbf{s}(\theta, \phi) = \begin{bmatrix} \cos\theta \sin\phi \\ \sin\theta \sin\phi \\ \cos\phi \end{bmatrix} \tag{8.5}$$

and λ is the ray-casting step size. We follow the ray in direction $\mathbf{s}$ until we find a point $\mathbf{P} \in \partial\Omega_{lookup}$. This is illustrated in Figure 8.1. We then calculate its distance D to $\partial\Omega_{ref}$ based on the feature transform $\mathbf{FT}^1_{\Omega_{ref}}(\mathbf{P})$ and store the distance in the distance map at position (u, v). This procedure is repeated for all points in the distance map and is sufficiently fast to allow interactive recalculation of the map for different thresholds T_{ref} selected by the user.

The computed 2D texture can be color-coded in various ways (see Section 8.6.4) and, when applied to the 3D mesh of Ω_{lookup}, can provide distance information in 3D between different areas on the tumor and the fibers. However, manipulating the camera view to search for specific points of interest on the tumor surface (e.g., point of minimal distance) may become tedious. To allow inspection of the complete tumor surface at a single glance, without the need for view-point manipulation, we also provide an unfolded, 2D display of the tumor distance map. This idea is similar to the tumor map proposed by Rieder et al. [135] who use it to show tumor coagulation zones after radio-frequency ablation. However, thanks to the GPU-based feature transform, our distance map can be interactively updated as the spatial extent of the object changes. When projecting the 3D tumor surface to a 2D plane care should be taken to preserve area as much as possible. The initial mapping $(\theta = u, \phi = v)$ will lead to severe area distortions when displayed as a 2D image. We have chosen a *sinusoidal* mapping for its computational efficiency but other mappings, such as the Mollweide projection [82] are also possible. The sinusoidal mapping is defined as follows:

$$\phi = v + \Delta\phi \tag{8.6}$$

$$\theta = u\ \cos^{-1}(v) + \Delta\theta \tag{8.7}$$

where $\Delta\phi$ and $\Delta\theta$ are angular offsets used to rotate the distance map such that the point of minimal distance (and its surrounding risk area) is placed in the center of the map where distortions are minimal. Even with a sinusoidal mapping distortions near the top and bottom of the map are inevitable. This issue is not specifically addressed in Rieder et al. [135]. The angular offsets are obtained by first computing:

$$\mathbf{w} = \mathbf{P}_{min} - \mathbf{P}_{center} \tag{8.8}$$

$$\mathbf{w}_N = \mathbf{w}/||\mathbf{w}|| \tag{8.9}$$

where $\mathbf{w}_N = [x_N\ y_N\ z_N]^T$ is the normalized vector originating from $\mathbf{P}_{center}$ and pointing towards $\mathbf{P}_{min}$ (already computed for the 3D map). $\Delta\theta$ and $\Delta\phi$ define angles between $\mathbf{w}_N$ and the vector $(0, 0, 1)$ corresponding to the center of the distance map at $(\theta = 0, \phi = 0)$. The offsets are computed by applying a standard Cartesian-to-spherical coordinate transformation on vector $\mathbf{w}_N$:

$$\Delta\theta = \begin{cases} \tan^{-1}\frac{y_N}{x_N} & \text{if } x_N \neq 0 \\ 0 & \text{if } x_N = 0 \end{cases} \tag{8.10}$$

$$\Delta\phi = \cos^{-1} z_N \tag{8.11}$$

The star convexity assumption of the tumor allows a further speed-up in ray casting. We can pre-compute its minimum radius R_{min} and start ray casting at a distance $R_{min} - \epsilon$ from $\mathbf{P}_{center}$, where ϵ is a small offset to handle round-off errors in case the iso-surface lies exactly at distance R_{min}. We use a relatively large step size and iterative step-size refinement [56] to obtain an accurate estimation of the iso-surface.

8.6.3 Assisted camera positioning

The tumor distance map provides an overview of the full tumor surface at a single glance. This feature comes at the cost of reduced spatial awareness in 3D. For this reason the 2D distance map and 3D view are linked by means of point selections. Clicking a point in the tumor distance map can, if enabled, automatically update the 3D camera position such that the corresponding point $\mathbf{P}_{user}$ on the 3D tumor surface is oriented towards the user. Due to possible irregularities in the tumor surface reconstruction we have chosen not to align the camera view vector with the tumor surface normal but, instead, align it with the vector $\mathbf{P}_{user} - \mathbf{P}_{center}$. To allow the user to follow the change in view-point, we apply a speed-adjustable animation that smoothly interpolates the camera view vector using quaternion rotation. Figure 8.3(c) illustrates the setup. Initially, the camera is looking down a vector $\mathbf{u}_0$. It should be rotated such that it looks down a vector $\mathbf{v}$. In the following, $\mathbf{u}_0$ and $\mathbf{v}$ are assumed to be quaternion vectors. Rotation takes place in small angular increments α about the axis represented by $\mathbf{w} = \mathbf{u}_0 \times \mathbf{v}$. A rotation of angle α about $\mathbf{w}$ is represented by the quaternion function:

$$\mathbf{q}(\alpha) = \cos\frac{\alpha}{2} + \mathbf{w}\ \sin\frac{\alpha}{2} \tag{8.12}$$

We define a function $\mathbf{F}$ that uses $\mathbf{q}(\alpha)$ to rotate quaternion vector $\mathbf{u}_i$ to $\mathbf{u}_{i+1}$ over an angle α as follows:

$$\mathbf{u}_{i+1} = \mathbf{F}(\mathbf{u}_i) = \mathbf{q}(\alpha)\ \mathbf{u}_i\ \mathbf{q}^{-1}(\alpha) \tag{8.13}$$

where $0 \leq i < N$, $\mathbf{u}_{N-1} = \mathbf{v}$ and $\mathbf{q}^{-1}$ is the complex conjugate of $\mathbf{q}$. The speed of the animation can be controlled by adjusting the angular increment α.

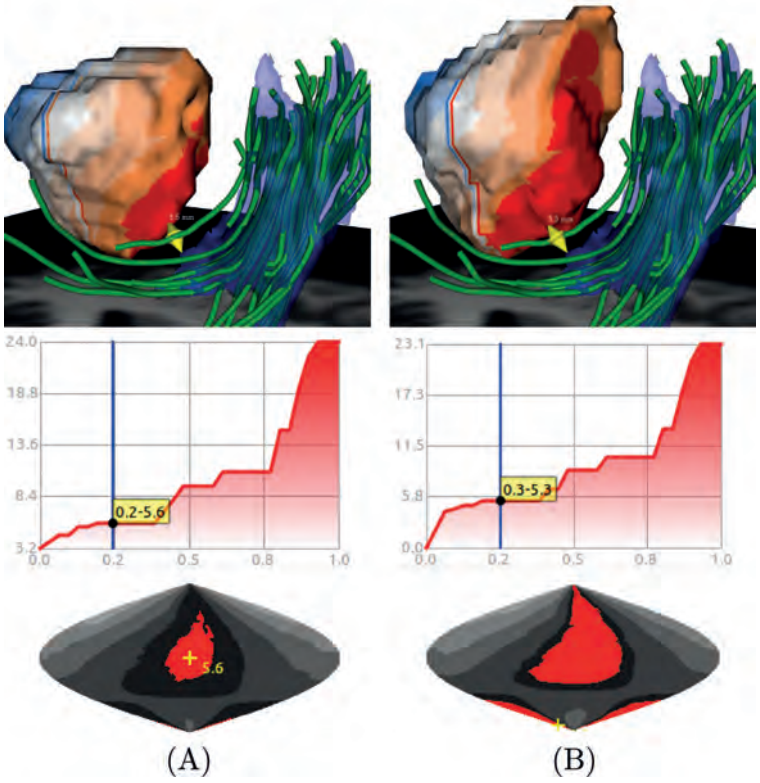

Figure 8.4: Distance analyses performed on different segmentations of the same tumor. Tumor segmented from MRI (A). Tumor segmented from DTI (B).

8.6.4 Additional exploration features

Besides the visualization of distances across the tumor surface, our tumor distance map provides a number of other features discussed below.

Point selection - The tumor distance map can be used directly for exploration of distances. The minimal distance is highlighted by a cross-hair and corresponding label indicating its magnitude (see tumor distance map Figure 8.2(c,d)). The user can click on the map to query distance values at different positions on the tumor surface. The selected point is highlighted on the map as well as in 3D using a spherical glyph. The principle of linking selections in a 2D map to a 3D view is similar to work done by Termeer et al. [148].

User-adjustable risk areas - The user can interactively change risk distance and see how it affects the size and shape of risk areas in both the 2D and 3D tumor distance maps. Because distance maps can be recomputed in real-time, the user can also keep the

risk distance fixed and see how the risk area changes size and shape due to changes in fiber probability iso-value. Figure 8.2(b,f) shows an example of a risk area highlighted on both the tumor distance map and 3D tumor mesh.

Distance contours - The tumor distance map, both in 2D and 3D, can be displayed as a contour map where distances inside a given interval are mapped to a single color. This makes it easier to quantify different surface areas with small, medium and large distances. This is also illustrated in Figure 8.2(b,f).

Color mapping - Using gray values to encode distances may interfere with the CT or MR images that are used for anatomical context (e.g. in Figure 8.2(b)). In that case, the user may prefer to encode distances using a different color scale. We implemented a warm-to-cool color map for this purpose where warm colors are associated with short distances and cool colors with long distances (see Figure 8.2(a)).

8.7 Results and discussion

Tumor segmentation was performed using ITK-SNAP [171] on the 1 mm^3 T1-weighted MRI dataset, resulting in a binary mask. Using FSL [67] we calculated the affine transformation from the MRI coordinate frame to the DTI coordinate frame. The tumor binary mask was transformed and resampled to the DTI coordinate frame. The MRI dataset was loaded as a 3D texture into our visualization tool without resampling. The 3D texture bounding box was scaled, rotated and translated based on the calculated transformation.

DTI fiber tractography was performed on the motor part of the corticospinal tract (m-CST) and the arcuate fasciculus (AF) which is associated with language. These tracts are of main interest to the neurosurgeons, especially the AF when the tumor is located in the language-dominant hemisphere. For the voxel-centric tractography algorithms we used ITK-SNAP to create seed regions. For the fiber-centric ConTrack algorithm we used the VistaSoft package [142]. We followed the steps described by Wakana et al. [159] to define the position and shape of seed regions. MRI and DTI datasets were acquired using a Philips Integra 3T scanner. DTI data dimensions are $128 \times 128 \times 60$, voxel dimensions are $1.75 \times 1.75 \times 2$ mm.

Performance

The CUDA-based distance/feature transform computation required to find closest tumor points to given fibers requires 8 bytes GPU RAM per voxel, i.e. 128 MB for 256^3 volumes and 1 GB for 512^3 volumes, which fits current high-end graphics cards. The speed of the algorithm is roughly linear with the voxel volume size (number of lookup points), since the number of reference voxels is kept small by considering only surface voxels of the reference shapes (see Sec. 8.5). Performance figures are given in Table 8.1, for several volume sizes, on three different NVIDIA GPU configurations, under Linux, and also Mac OS X for the GT 330M.

The implementation of the distance/feature transform is optimized for power-of-two volumes. Performance is almost constant regardless of the number of fiber voxels, and is

Volume (voxels)	GRAM (MB)	DT/FT (GT 330M) (sec.)	DT/FT (GTX 280) (sec.)	DT/FT (C1060) (sec.)
128x128x128	16	0.07	0.04	0.05
256x256x256	128	0.40	0.28	0.30
512x512x512	1024	n/a	2.10	2.20

Table 8.1: DT/FT performance for various volumes and GPU configurations.

roughly linear in the number of voxels of the embedding space, i.e., independent of the actual number of reference voxels where we are interested to evaluate the distance to voxels. For the GT 330M GPU, volumes of 512^3 have not been treated, as our configuration for this card had only 512 MB graphics RAM. CPU-based optimized distance/feature transforms, such as the AFMM [147] or the Euclidean TFT [132], are tens of times slower due to the lack of parallelization *and* the fact that the number of lookup points is large, i.e., the entire volume minus the relatively thin fiber structures. For completeness, note that in applications involving considerably fewer lookup points, e.g., one is interested in feature-and-distance values only within a thin narrow band around a reference shape, fast-marching methods such the AFMM are faster than GPU-based methods, as the GPU-to-CPU transfer overheads dominate computational costs.

The preliminary results of our work suggest that interactive visualization of distances between objects with uncertain spatial extent can help to answer relevant clinical questions. This information would be difficult, or even impossible, to obtain solely by means of a visual inspection of 2D images. Neither can existing approaches for 3D distance analysis provide information about the reliability of the distance measurements due to uncertainty.

The uncertainty in a given object's spatial extent can be characterized in many different ways. Our main focus has been on white matter pathways and the reproducibility with which they can be reconstructed by various probabilistic tractography algorithms. This reproducibility largely depends on sensitivity to data noise and is modelled in different ways by different algorithms. The PICo method models noise explicitly while the Wild Bootstrap method does not. ConTrack internally also uses a Wild Bootstrap approach to estimate certain parameters. We have shown that our framework is not dependent on any particular tractography method. It only requires that the output data is presented in a simple volumetric format which we have labeled the *reproducibility* map. Each voxel is assigned a reproducibility value depending on how often it is classified as being part of the object of interest (i.e., the fiber tract) after a certain number of iterations of the algorithm.

The distinction between voxel-centric and fiber-centric methods may be helpful when considering the application domain. We have illustrated the use of our distance analysis framework in a neurosurgical context. Discussions with the neurosurgeons led to the conclusion that voxel-centric approaches may be less suitable for neurosurgical purposes. First, voxel-centric algorithms can tell you how reproducibly a voxel is connected to the seed region but not *how* or through which pathways. Second, the shape of the iso-surface

is likely to be different from what a neurosurgeon expects based on knowledge of white matter anatomy (see Figure 8.2(b) and (c)). Third, due to the sampling strategy of voxel-centric algorithms, voxel reproducibilities always decrease with distance from

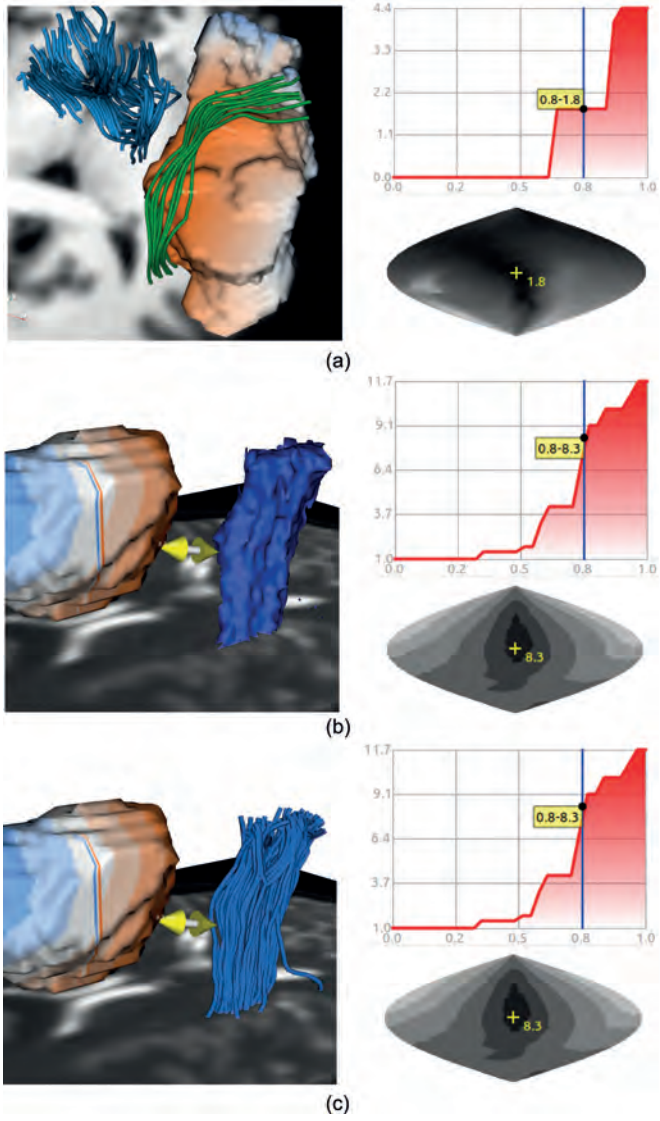

Figure 8.5: (a) Grade II glioma in left temporal lobe touching language tracts (green). Motor tracts are indicated in blue. (b) Fiber-centric reproducibility surface. (c) Streamlines corresponding to fiber-centric reproducibility surface.

the seed region. This may not be a realistic representation of fiber uncertainty, especially if it is known from physical examinations that fiber tracts are present and intact. For all these reasons, a voxel-centric representation of uncertainty may not be the best choice in a neurosurgical setting. A fiber-centric algorithm, on the other hand, assigns reproducibility values to complete pathways instead of single voxels, thereby preserving pathway information. Iso-surfaces generated from its output represent all voxels intersected by pathways with the same reproducibility. For this reason, the iso-surface shape will more closely resemble the shape of the underlying pathways (see Figure 8.5(b) and (c)).

8.8 Conclusions

In the previous sections we have presented a framework for the interactive exploration of distances between volumetric objects of uncertain spatial extent. This framework is particularly useful in a neurosurgical context where distances are an important concept for surgical planning and decision-making. In the next chapter we do not introduce novel visualization techniques but, instead, focus even more on the neurosurgical context by applying the techniques presented in this and previous chapters to 5 real-world, clinical cases of patients with brain tumors.

No one believes forecasts, but we all want to hear them.

- Dilbert

9 Applications in neurosurgical planning

9.1 Introduction

In Chapter 2 we mentioned that DTI and fiber tracking are of great interest for neurosurgical applications, such as planning and intra-operative navigation. The ability to non-invasively reconstruct the 3D layout of brain white matter is especially useful for the planning of tumor resections involving critical brain areas [169, 34, 61]. Such information is not adequately provided by conventional imaging such as contrast-enhanced T1, T2 or FLAIR imaging [34]. 3D visualization of fiber tracts allows neurosurgeons to preoperatively evaluate optimal access trajectories toward the lesion. It can also assist in decision-making. For example, if the tumor is suspected to be close to critical fiber tracts the neurosurgeon can use fiber tracking results to decide whether or not to use intra-operative stimulation mapping [78, 44]. If a tumor shows infiltration of important fiber tracts that are still functionally intact, a complete resection may not be possible which is important information to convey during patient counseling. Even if intra-operative stimulation is decided upon, fiber tracking can still point out the 3D locations where stimulation is applied most effectively and thereby help to make the overall procedure more efficient [52].

As can be assessed from the multi-stage DTI processing pipeline described in Chapter 4, the output of fiber tracking can be subject to a wide range of sources of uncertainty. There is also substantial *clinical* evidence for the reliability issues of standard fiber tracking methods. These issues have been reported by several authors [34, 81, 27]. Also, there have been numerous attempts to clinically validate fiber tracking results on the basis of intra-operative navigation and electrical stimulation mapping [75, 18, 15, 108]. The sensitivity of fiber tracking to user-defined parameters, such as stopping criteria and seed point placement, has also been investigated [159, 33, 60].

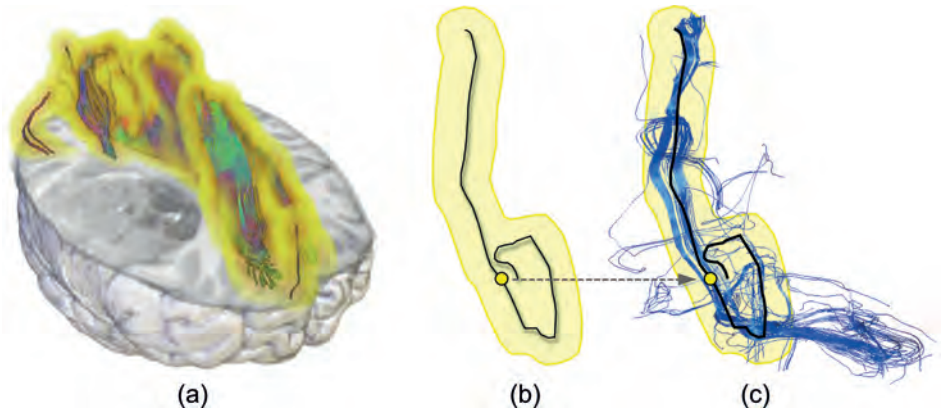

Figure 9.1: Fixed-distance, visual safety margin (image courtesy of Florian Weiler, Fraunhofer MeVis, Bremen) (a). Single seed point and fiber calculated using deterministic fiber tracking (b). Fiber pathway variations for a single seed point calculated using Wild Bootstrap fiber tracking [71] (c).

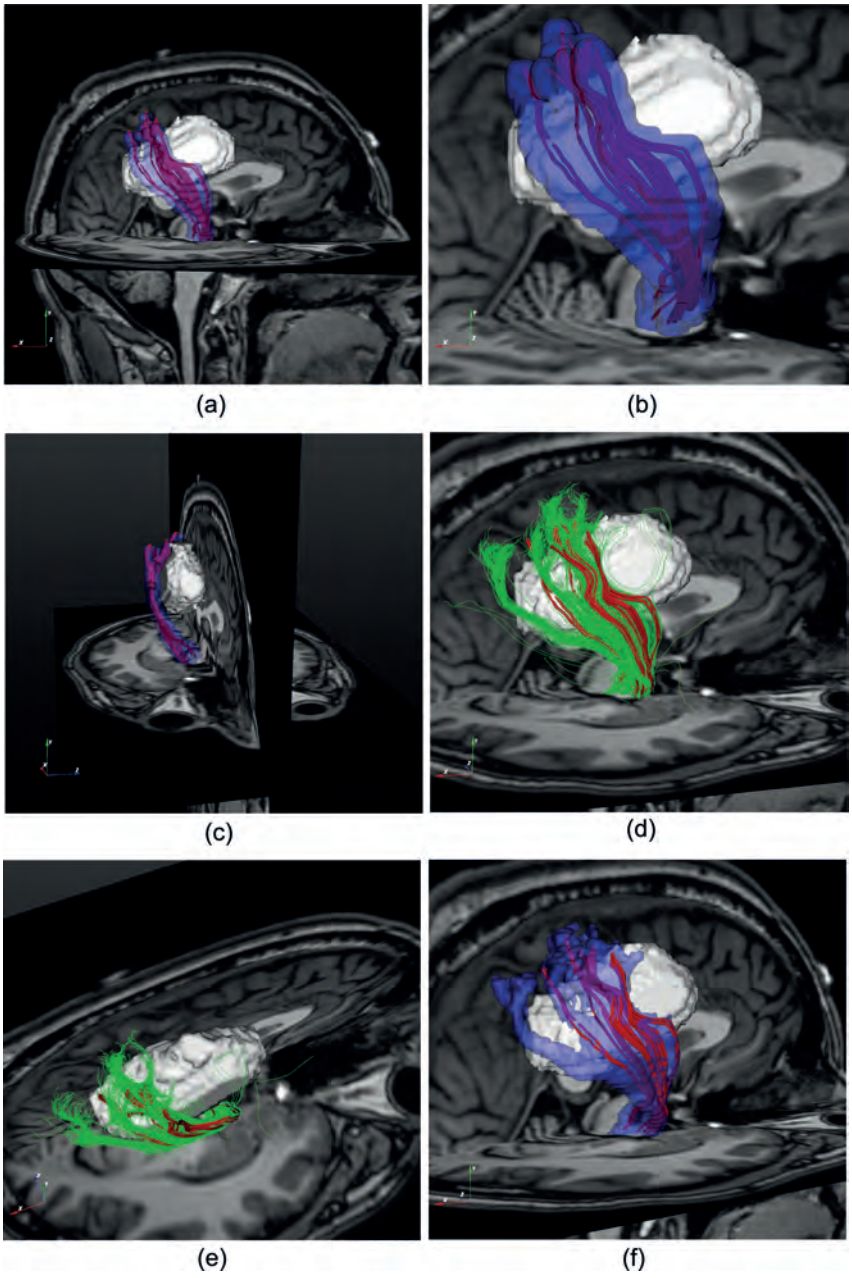

Figure 9.2: Patient 1. Deterministic fibers with fixed safety margin (a-c). Deterministic fibers (red) and probabilistic fibers (green) (d-e). Deterministic fibers with uncertainty-driven safety margin (f).

In the previous chapters, we have shown that the visualization techniques described have potential applications in neurosurgical settings. In Chapter 7 we performed an informal user study to evaluate the usefulness of our visualizations for neurosurgeons. Chapter 8 focused on distance measurements, which form an important part of the neurosurgical planning routine. The visualization techniques presented in these chapters rely on the output of probabilistic fiber tracking algorithms in order to capture fiber pathway variations due to noise and model errors. However, this seems to suggest that such tracking algorithms are commonly available to neurosurgeons, which they are not. Neurosurgeons typically work with commercially available viewing stations. The fiber tracking algorithms implemented in these systems are mostly of a deterministic nature and, therefore, provide support for neither the characterization and representation of uncertainty, nor its visualization. It is left entirely up to the user to estimate the reliability of what is being presented. In some cases, the user has the option to illustrate this estimation by defining a visual safety margin (Figure 9.1(a)) at a fixed distance to the tracked fiber bundle [161]. However, this does not hide the fact that the system itself provides no information whatsoever about reliability issues. The safety margins only visualize the opinion of the user and have no relation to the underlying data or the tracking algorithm. Figures 9.1(b) and (c) clearly demonstrate that a fixed, user-defined safety margin may not adequately capture the possible variations in fiber pathways predicted by probabilistic tracking algorithms [71, 50, 122]. Similarly, the differences between deterministic tracking algorithms, as compared by Buergel et al. [27], cannot be captured by simply taking one of the tracking results and expanding it with a fixed safety margin (unless the safety margin is too large to be of practical use).

In this chapter we apply the visualization techniques presented in Chapter 6 and 8 to five clinical cases. To provide the neurosurgeon with a clear comparison of the differences between deterministic and probabilistic algorithms we render deterministic fiber tracts in the traditional manner, e.g. using streamtubes such as described in Chapter 3, while the probabilistic tracking information is represented as illustrative confidence intervals (Chapter 7) or reproducibility iso-surfaces (Chapter 8). We call these probabilistic boundaries *uncertainty-driven safety margins* which are calculated from probabilistic fiber tracking and represent uncertainty in the spatial extent of fiber pathways due to image noise and DTI model errors. We combine the uncertainty-driven safety margins with real-time calculation of distance information between tumor and surrounding fiber tracts.

9.2 Pre-processing steps

In this section we will shortly describe the image acquisition and post-processing steps taken to create the visualizations described in the next section.

- **Image acquisition** - All images were acquired using a Philips Achieva 3T scanner (Philips Healthcare, The Netherlands) with a 8-channel head coil. T1-weighted images were acquired at 288 x 233 pixels, 176 slices and 1 mm^3 voxels. Diffusion-weighted images were acquired using single-shot EPI at 128 x 128 pixels, 60 slices, 1.75 x 1.75 x 2 mm voxels, 30 gradient directions and a b-value of 1000 s/mm^2.
- **Image registration** - The T1-weighted images were aligned with the DTI images using a rigid-body transformation calculated with the FLIRT tool of the FSL

package [67].

- **Tumor segmentation** - Tumors were segmented from the T1-weighted images using ITK-SNAP, a general purpose, open-source image segmentation tool [171]. The segmented tumor data was subsequently aligned with the DTI data using the previously calculated rigid-body transformation.
- **Fiber tracking** - Deterministic fiber tracking (FACT algorithm, 2nd-order Runge-Kutta integration) was performed on diffusion tensor data calculated from the original diffusion-weighted images. Probabilistic fiber tracking was performed using the Wild Bootstrap method (for details on both tracking algorithms see Chapter 4). We reconstructed two major fiber bundles for each clinical case: (1) the corticospinal tract associated with motor function and (2) the arcuate fascicle associated with language function (see also Chapter 2). Seed and target regions were defined according to the recipes recommended by Wakana et al. [159].

9.3 Clinical cases

In the following paragraphs we will shortly provide a description of each patient's medical background and present visualization results emphasizing the tumor, fiber tracts, safety margins and distance information. For an overview of tumor grades we refer the reader to Chapter 2.

Patient 1 - Patient 1 is a 48-year old male who presented with focal sensomotoric seizures of short duration, multiple times per week and starting a year earlier. The patient had no neurological deficits. The tumor was diagnosed to be a high-grade glioma (WHO-III), located in the right gyrus cingulus and subcortically underneath the sensomotor region. Figure 9.2 shows images from the investigation of patient 1. The tumor is located in the right cingulate gyrus underneath the sensomotor cortex. Given the patient's history of focal sensomotor seizures, special attention was given to fibers of the corticospinal tract (CST). We loaded pathways reconstructed using deterministic FT and visualized them using stream tubes. The fibers are surrounded by a fixed, user-defined safety margin (Figure 9.2(a), (b) and (c)) to illustrate what the visualization would look like when using existing, commercial viewing stations. Figure 9.2(d) and (e) show the deterministic pathways (red) together with pathways calculated using the Wild Bootstrap method, a well-known probabilistic tracking algorithm proposed by Jones et al [71]. As can be seen the fixed safety margin does not correctly capture the pathway uncertainty predicted by the probabilistic algorithm. Figure 9.2(f) illustrates the probabilistic pathways again, now visualized using our uncertainty-driven safety margins with a semi-transparent surface style.

Patient 2 - Patient 2 is a 41-year-old male who presented with a single generalized seizure without neurological deficits. The tumor was diagnosed to be a low-grade glioma (WHO-II), located in the left premotor area (F1) of the gyrus frontalis superior. Figure 9.3 shows images from the investigation of patient 2. The tumor is located left premotor area of the gyrus frontalis superior. Given the patient's history of a single seizure and the location of the tumor as visible on the anatomical MR images, primary attention

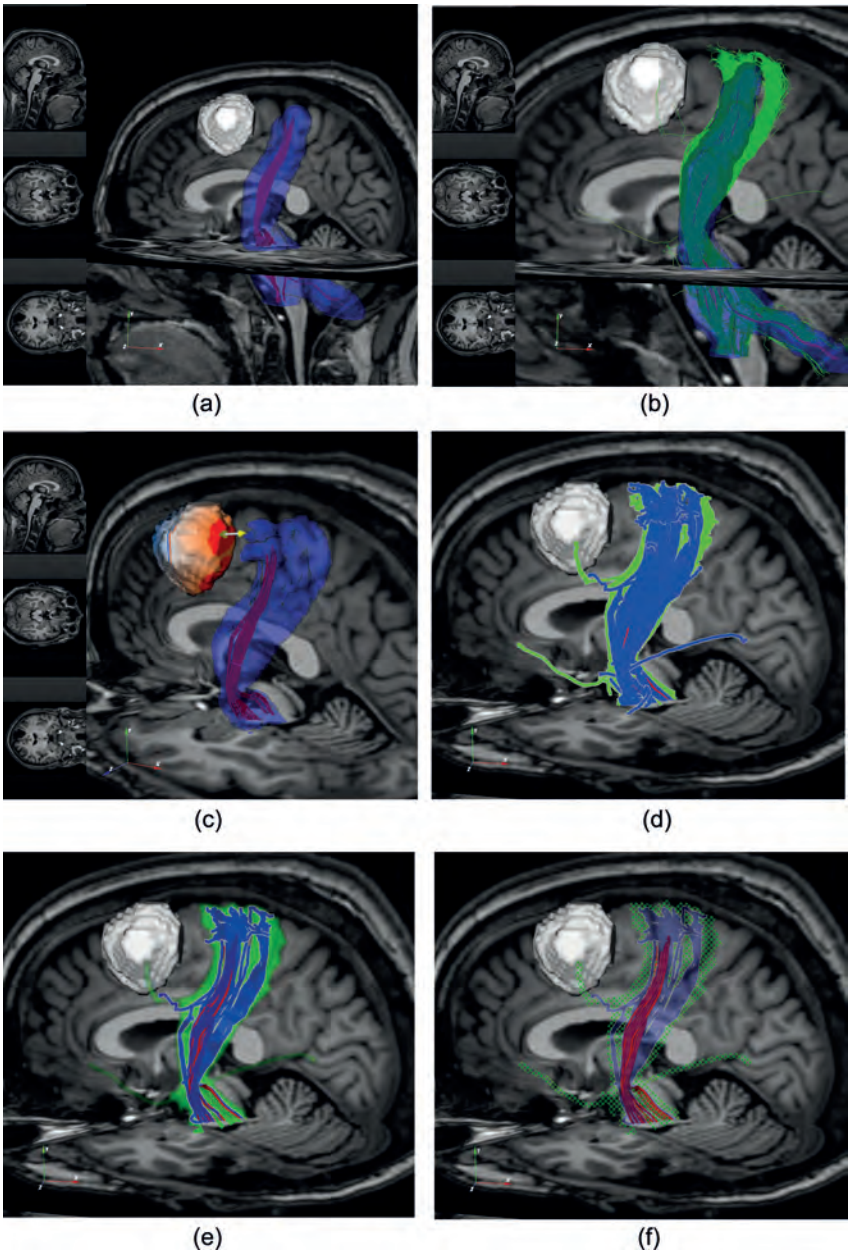

Figure 9.3: Patient 2. Deterministic fibers with fixed safety margin (a). Comparison fixed safety margin and probabilistic fiber extent (b). Deterministic fibers and uncertainty-driven safety margin (c). Different visual styles for uncertainty-driven safety margin (d-f).

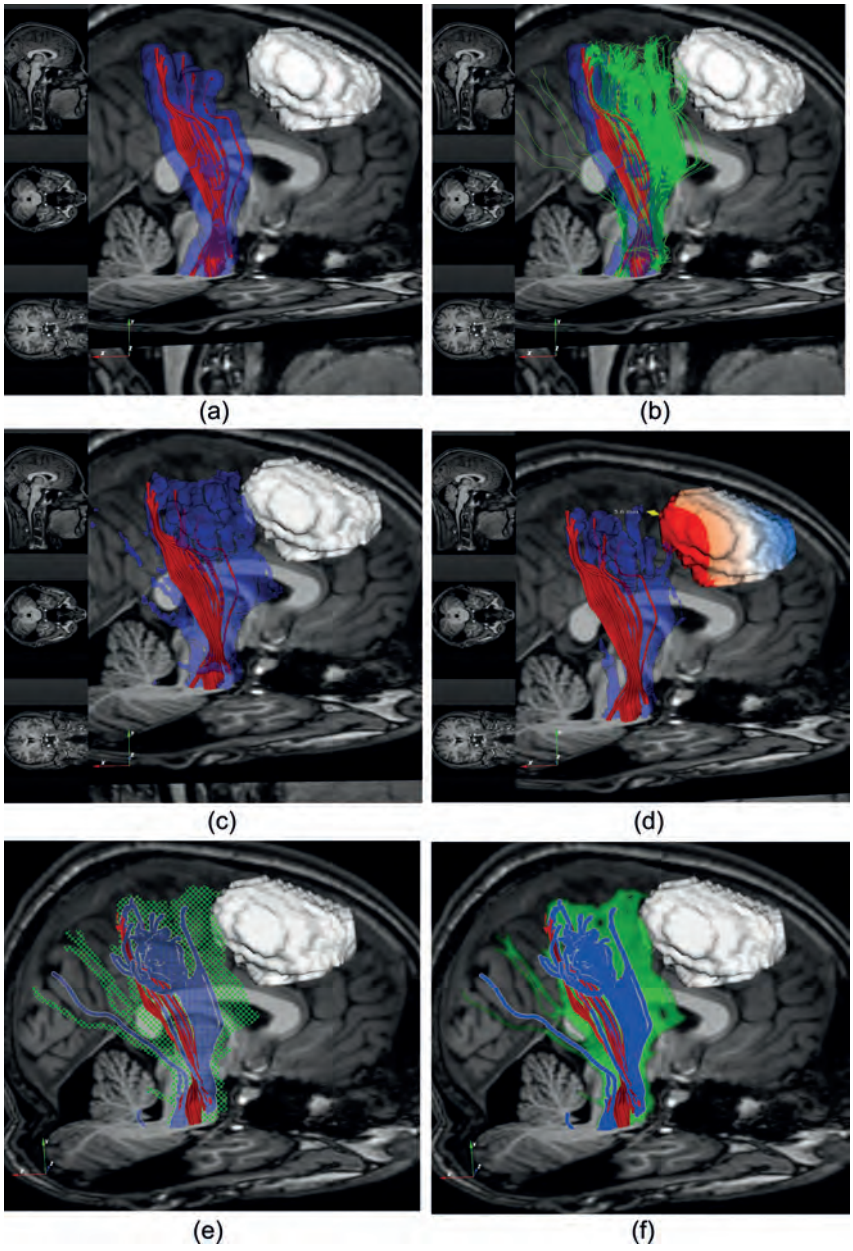

Figure 9.4: Patient 3. Deterministic fibers and fixed safety margin (a). Comparison fixed safety margin and probabilistic fibers (b). Deterministic fibers and uncertainty-driven safety margin (c). Idem with tumor distance map (d). Different visual styles for uncertainty-driven safety margin (e-f).

was given to the CST. Figure 9.3(a) shows the output of deterministic FT together with a fixed safety margin. Figure 9.3(b) shows the fixed safety margin in combination with a streamline representation of the probabilistic pathways in order to highlight the difference in shape. Figure 9.3(c) provides an example of automatic calculation of 3D distances between tumor and our uncertainty-driven safety margin calculated from the probabilistic pathways. The tumor surface is color-coded based on distance to the fiber pathways. Warm orange colors represent small distances, cold blue colors represent large distances. The bright red areas on the tumor indicate those locations where the tumor lies within 15 mm's of the fiber pathways. The point of minimal distance is indicated by the double-headed arrow and is 12.4 mm. Figure 9.3(d), (e) and (f) show examples of uncertainty-driven safety margins with a silhouette and outline style. This particular visual style allows visualization of multiple boundaries, each representing pathways of a given confidence interval. Figure 9.3(d) shows the spatial extent of the 50% most reproducible pathways (blue) and the 50% least reproducible pathways (green). Figure 9.3(e) and 9.3(f) provide examples of alternative ways to render each confidence interval. For example, in Figure 9.3(e) silhouette boundaries are blurred with decreasing confidence while in Figure 9.3(f) this is done by increasing the coarseness of the checkerboard pattern inside the silhouette. The blue checkered region represents 25% of the most confident pathways.

Patient 3 - Patient 3 is a 31-year-old male who presented with multiple seizures without neurological deficits. The tumor was diagnosed to be a low-grade glioma (WHO-II), located in the right premotor area (F1) of the gyrus frontalis superior. Figure 9.4 shows images from the investigation of patient 3. The tumor is located in the right premotor area of the gyrus frontalis superior. Given the patient's history of multiple seizures, the primary focus was on the CST. Figure 9.4(a) shows these tracts as reconstructed using deterministic FT. Again, a fixed safety margin is shown as would be possible in existing, commercial DTI analysis systems. Figure 9.4(b) and (c) show pathway information calculated from the probabilistic Wild Bootstrap algorithm, both as streamlines (Figure 9.4(b)) and as an uncertainty-driven safety margin with semi-transparent surface style (Figure 9.4(c)). The difference with the fixed safety margin is again noticeable. Figure 9.4(d) illustrates distance information projected on the tumor surface. The bright red areas, again, highlight points on the tumor within 15 mm of the uncertainty-driven safety margin. Figure 9.4(e) and (f) show alternative examples of the silhouette representation of our safety margins.

Patient 4 - Patient 4 is a 39-year-old male who presented with multiple complex partial seizures. The patient suffered from memory loss and mild speech problems. The tumor was diagnosed to be a low-grade glioma (WHO-II), located in the right temporal lobe. Figure 9.5 shows images from the investigation of patient 4. The tumor is located in the right temporal lobe. Figure 9.5(e) clearly shows it pushing against the CST. Given the combination of complex partial seizures and language problems, both the CST and arcuate fasciculus (AF) were considered. Figure 9.5(a) shows a reconstruction of the language pathways based on deterministic FT together with a fixed safety margin. Figure 9.5(b) and (c) show the result of probabilistic fiber tracking. In Figure 9.5(d) distance information was calculated between the tumor surface and surrounding AF. The bright red areas indicate locations on the tumor where the distance < 15 mm. The minimal

distance was calculated to be 2.7mm. Figure 9.5(e) shows the interaction of the tumor with CST. The distance color-coding on the tumor surface was calculated with respect to the CST. The point of minimal distance (3.5 mm) is visible in the bottom left part of the tumor. Figure 9.5(f) and (g) show uncertainty-driven safety margins surrounding the AF. Figure 9.5(h) and (i) show similar visualizations for the CST.

Patient 5 - Patient 5 is a 37-year-old male who presented with a single seizure and mild speech problems. The tumor was diagnosed to be a low-grade glioma (WHO-II), located in the right temporal lobe.Figure 9.6 shows images from the investigation of patient 5. The tumor is located in the right temporal lobe. Given the patient's history of a single seizure and speech problems, both the CST and AF were considered. Deterministic FT results are illustrated in Figure 9.6(a). Figure 9.6(b) shows the CST with a fixed safety margin. Compared to patient 4, there is little or no distortion of fiber geometry. Distance calculations performed between tumor and CST result in a minimal distance of 8.2 mm based on the uncertainty-driven safety margin calculated from probabilistic fiber tracking. The results are shown in Figure 9.6(c). Figure 9.6(a), (d), (e) and (f) indicate that the tumor is close to the AF. Distance calculations between tumor and the uncertainty-driven safety margin result in a minimal distance of 0 mm. Figure 9.6(e) and (f) show examples of our safety margins using silhouette-based confidence intervals.

9.4 Discussion

In the previous section we have shown that the visualization techniques proposed in earlier chapters can be usefully applied to real clinical cases. We have presented five such cases involving patients with brain tumors. Techniques such as illustrative confidence intervals (Chapter 7) and reproducibility isosurfaces (Chapter 8) can serve as a type of uncertainty-driven safety margin surrounding fiber pathways created using deterministic streamline tracing. The advantage of combining probabilistic and deterministic tracking in this way, is that it allows the user to compare the two methods and see in which locations their outputs coincide and in which locations their outputs deviate. Probabilistic tracking methods are able to detect reproducible pathways in low-anisotropy regions where deterministic algorithms would just produce random results. To prevent such false positives they rely on strict anisotropy and curvature thresholds. However, this can potentially ignore true positives and considerably reduce the spatial extent of the visualized fiber bundle because only a limited number of pathways actually reach their destination. This is one of the reasons behind the limitations reported by Kinoshita et al. [81]. This phenomenon can also be seen when comparing Figure 9.2(b) with 9.2(d) of patient 1, Figure 9.3(a) with 9.3(b) of patient 2 and Figure 9.4(a) with 9.4(b) of patient 3.

Differences between deterministic and probabilistic output results are not always large. In patient 4 the tumor is located in the lower, temporal part of the brain. It is very large and pushes the corticospinal tracts to the contra-lateral side. We experimented with placing seed points both in the cerebral peduncle and the cortex but found no significant difference in the resulting set of pathways. As can be seen in Figure 9.5(e), the deterministic pathways (red) and the probabilistic surface hull (blue) are quite similar, especially in the area closest to the tumor.

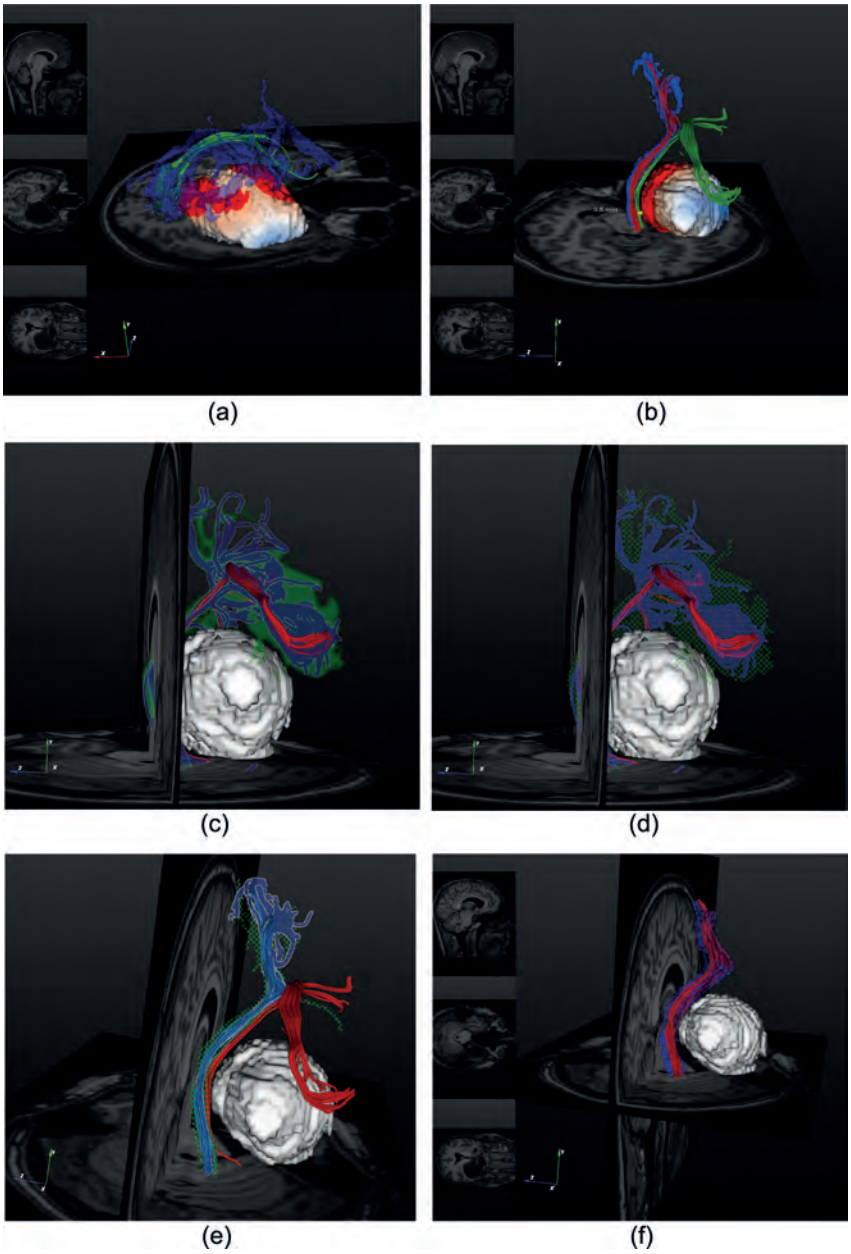

Figure 9.5: Patient 4. Deterministic fibers, uncertainty-driven safety margin and tumor distance map (a). Tumor distance map with deterministic motor fibers (red), uncertainty-driven safety margin (blue) and deterministic language fibers (green) (e). Idem with different visual styles for uncertainty-driven safety margins (c-f).

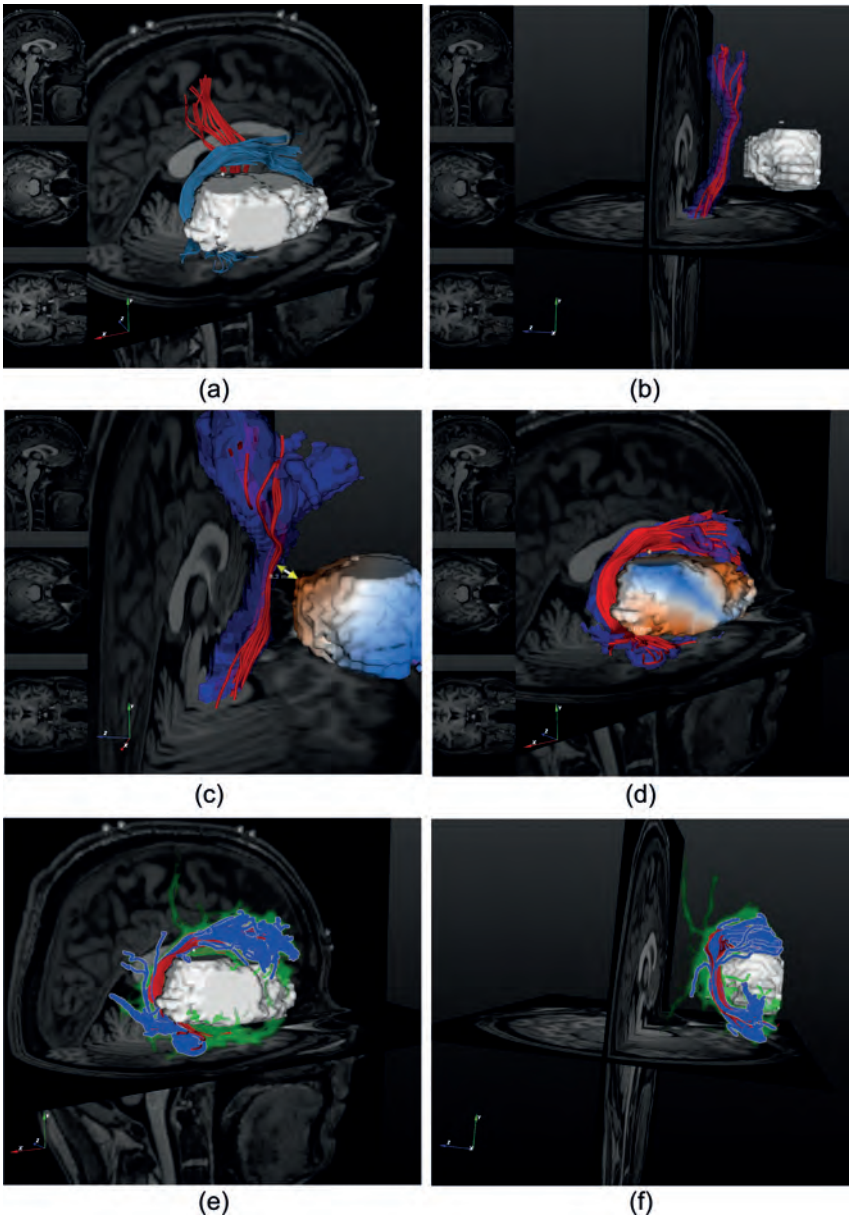

Figure 9.6: Patient 5. Deterministic language (blue) and motor fibers (red) (a). Deterministic motor fibers with fixed safety margin (b). Deterministic fibers with uncertainty-driven safety margin and tumor distance map (c). Idem for language fibers (d). Different visual styles and viewpoints (e-f).

The informative power of the visualizations presented in this chapter is primarily achieved by combining deterministic tracking results with the uncertainty-driven safety margins. In this respect, these visualizations are complementary. The combination of deterministic and probabilistic methods also has some educational value, especially for users who have less experience with fiber tracking and its interpretation. Our visualization method allows easy comparison between deterministic and probabilistic results. The safety margins are not fixed but adapt themselves to changes in confidence along a given fiber bundle as well as between different fiber bundles. This allows the clinical user to discover, in an intuitive manner, how the performance of a fiber tracking algorithm depends on the region of the brain where it is applied.

Despite the advantages of probabilistic fiber tracking, there are also limits to the interpretation of its output and its subsequent visualization. In this thesis we have primarily discussed uncertainty in terms of *reproducibility*, meaning how similar fiber pathways are when tracking is repeated many times. This does not mean, however, that the pathways always represent the underlying anatomy accurately. As pointed out before, a gold standard for validation of the tracking results is lacking. Furthermore, these algorithms can also suffer from image artifacts, model errors, incorrect seed and target regions or tissue pathology effects. Still, when deterministic tracking fails for some reason, probabilistic methods may be the only way to detect reproducible fiber structures.

9.5 Conclusions

In this chapter we have presented 5 clinical cases of patients with brain tumors and have shown that the visualization techniques presented in this thesis can provide interesting and useful information for neurosurgical planning. This chapter represents the final main contribution of this thesis. In the next chapter we will summarize our contributions and provide suggestions for future work.

Large corporations welcome innovation and individualism in the same way the dinosaurs welcomed large meteors.

- Dilbert

10 General discussion

10.1 Introduction

In the previous chapters we have presented several techniques for the visualization and exploration of uncertainty in fiber tracking on the basis of diffusion tensor imaging. In the following paragraphs we will briefly recapitulate our main contributions and highlight the potential added value for neuro-scientific and neurosurgical purposes.

Our first contribution was presented in Chapter 6 and attempts to address the problem of parameter sensitivity in deterministic fiber tracking. We have focused on visualization techniques that highlight stability problems in the stopping criteria. An important aspect of finding stable threshold values for these stopping criteria is the ability to interactively change them and see how the fiber tracking output reacts. With GPU-acceleration interactively changing tracking parameters becomes easier, but this still allows only visualization of individual thresholds at a time. The zero-threshold tracking and color-coding we propose in Chapter 6 allows to view pathways across a *range* of such threshold values. One of the advantages of our technique is that it directly shows which parts of a traced fiber bundle are likely to be included or excluded for a range of threshold values. This could serve as a first visual check in neurosurgical settings where false negatives may be the result of either tissue pathology (e.g., perifocal edema) or an overly strict anisotropy threshold.

In Chapter 7 we presented a visualization technique for showing the effects of image noise on fiber pathways. Uncertainty in this context is commonly characterized by means of probabilistic tracking methods. However, visualization of their output is mostly limited to 2D fiber density maps which do not show the 3-dimensional layout of the fibers and their associated uncertainty. We have proposed illustrative rendering techniques, based on silhouettes and outlines, to render the large numbers of fiber pathways resulting from probabilistic tracking. These techniques simplify the dense collection of streamlines to their essential shape characteristics. This reduces visual clutter and maintains a sense of 3D shape and extent.

Chapter 8 presents a first attempt to integrate uncertainty visualization of fiber pathways in a practical task common in neurosurgical settings: the measurement of distances. Distance is a very important concept in surgical planning since it directly impacts decision-making. A tumor that is close to critical brain regions or fibers may require resection under the guidance of electrical stimulation. A tumor that is far from such structures, may not require such measures. However, the assessment whether distances are large or small, depends on the spatial extent of the objects involved, in this case the tumor and surrounding fiber tracts. This is difficult to do on the basis of 2D images, which is current practice, but becomes virtually impossible if uncertainty in the spatial extent of the object has to be taken into account. The exploration framework which we propose in Chapter 7 allows (1) visualization of distance in 3D and (2) interactive exploration of distances with changing object shape. Furthermore, we provide several tools, such as the tumor distance map, that present the distance information in a clinically relevant way.

Chapter 9 attempts to apply the visualization and distance analysis techniques presented in Chapters 7 and 8 to a set of real-world clinical datasets. The discussion presented

in this chapter focuses more on the clinical benefits of probabilistic fiber tracking and its visualization. We propose to use our visualizations as an alternative for the fixed-distance safety margin commonly featured in commercial fiber tracking systems. We show that fixed margins do not adequately capture the potential reliability issues in fiber tracking output. The visualizations shown in this chapter can, for this reason, be considered to be a type of uncertainty-driven safety margins where the spatial extent of the margin actually reflects uncertainty derived from the underlying data.

10.2 Remaining challenges and future outlook

DTI and fiber tracking are promising tools that can bring a wealth of information about brain connectivity and function. However, in order to correctly interpret the output they also require detailed knowledge of (1) the technical details and (2) the underlying white matter anatomy. Depending on the type of user a different balance between these two factors exist. For example, neuroscientists may be looking for previously unknown brain connections and, therefore, cannot always rely on anatomical knowledge. However, they mostly have detailed knowledge of image acquisition and post-processing methods, including their limitations, which can help them to assess the reliability of the output. Given the wide range of tracking algorithms to choose from, uncertainty visualization can play a useful role in the comparison of such algorithms. For example, our illustrative confidence intervals could be used to visualize the outputs of multiple tracking algorithms where overlap between the resulting fiber sets (i.e., where the algorithms agree) represents a higher confidence.

In neurosurgery, the added value of uncertainty visualization is more difficult to assess. Fiber tracking forms only a single part in the overall investigative procedure. The primary goal is to improve patient outcome but this is dependent on a wide range of factors, many of them not related to imaging. Patient history, physical fitness, type of lesion, medication, intra-operative navigation and surgical expertise may have a larger impact on clinical decision-making than fiber tracking. Even if fiber tracking would be used routinely, it would be difficult to establish its exact contribution to patient outcome as compared to the contribution of other factors. Given the wide range of issues neurosurgeons have to deal with when planning and executing surgical procedures, it is understandable that they have less opportunity to familiarize themselves with the technical details of DTI and fiber tracking. On the other hand, they can rely on extensive knowledge of structural and functional anatomy to assess the accuracy of the fiber tracking results. In many cases, deterministic fiber tracking suffices to provide the neurosurgeon with confirmation of the spatial layout of fiber bundles and their distance to the target lesion. A problem, and opportunity, arises when standard fiber tracking methods fail, for example, when edematous fluid surrounding a tumor lowers diffusion anisotropy to such an extent that deterministic tracking fails to reconstruct fibers between brain regions that are known to be connected. This may have been determined on the basis of a physical examination. If, as a consequence, it cannot be assessed on which side fibers pass the tumor, this can be problematic for planning the angle of approach toward the tumor. In such cases, probabilistic tracking methods may have a better chance of detecting reproducible fiber structures in such areas. A similar scenario was described in a book chapter by Bartsch et al. [6] and provides an excellent opportunity

for uncertainty visualization techniques, such as the fiber confidence intervals and distance exploration methods we presented in this thesis.

This brings us to a future outlook for uncertainty visualization on the basis of DTI and fiber tracking. As will be clear, much remains to be done in this area, especially for neurosurgical applications where the benefits of uncertainty visualization require careful selection of relevant clinical cases. A very interesting line of research which could be pursued in the near future is to find previous clinical cases with complex pathology, where deterministic fiber tracking methods either failed, or would be expected to fail if they had been applied. Probabilistic tracking methods may still result in useful information and it would be interesting to evaluate the effectiveness of our uncertainty visualization techniques as compared to standard methods for probabilistic fiber tracking, such as 2D fiber density maps. This is more difficult to do with cases that are relatively straightforward.

Another interesting challenge in neurosurgical applications would be to combine our fiber tracking visualization with uncertainty in the spatial extent of the tumor. The distance exploration framework which we presented in Chapter 8 allows both objects of interest to have varying shapes. However, it is far from trivial to create a useful probabilistic segmentation of tumor tissue. It might be interesting to combine the 3D uncertainty visualization of the white matter fibers with a 2D visualization of the anatomical images. The latter gives the best, unprocessed representation of the tumor boundaries, e.g., as captured by T2-weighted, FLAIR or contrast-enhanced T1-weigted MRI. Also, a combination of 2D and 3D visualizations may be easier to adopt since neurosurgeons are already familiar with 2D imaging. Another requirement in this area would be to develop a user interface for the software that is easy-to-use for the neurosurgeons. This allows them to work with the visualization tools themselves and is likely to improve the quality of their feedback.

Many interesting challenges and opportunities also remain in the area of neuroscientific research, especially brain connectivity studies. Our visualization techniques are quite suitable for visual comparison of information obtained from large groups of individuals. For example, fiber tracking results from such a large group could be collated into a probabilistic fiber atlas which could subsequently be visualized using different confidence intervals as described in Chapter 7. Another area where uncertainty visualization can play a role is in the development of new fiber tracking algorithms. Our techniques provide an intuitive way to visually assess the pathway output of such algorithms and compare it to established approaches.

Final thoughts

To conclude this thesis, we point out that uncertainty visualization is a highly interesting but also a very tough subject. Intuitively, it makes sense to show uncertainty in the data. Data is always measured, processed and visualized for a specific purpose, be it exploration, problem solving, decision-making or communication. In all cases, the essence is to get at the truth about the phenomenon being measured. This will allow you to increase your knowledge about it and use this knowledge to make better decisions. However, all measurements and processing methods are subject to errors, assumptions

and biases which can have unpredictable effects on the data being investigated. If such issues can be characterized in some way, it makes sense to visually communicate them to the user. However, as we already pointed out in the introduction of this thesis, many sources of uncertainty remain hidden and cannot be easily characterized or quantified. This raises a number of critical and understandable questions, such as 'How valuable is a visualization that shows only part of the uncertainty and not all of it?', or 'How valuable is a visualization that shows uncertainty of a kind I do not understand?'. Again, we come back to the experience and background knowledge of the user. Some users realize that they are seeing only part of the uncertainty, but still appreciate and understand the visualization of those uncertainties that can be quantified. Other users may not fully understand the source of uncertainty and find its visualization confusing. It is likely that the latter scenario applies to a much larger group of users, especially in non-scientific settings. In the opinion of the author, this is where the main challenge of uncertainty visualization lies: in finding approaches to communicate uncertainty in an accessible and meaningful way to 'lay' people who use visualization only as a tool for problem solving and decision-making.

Bibliography

[1] D. Akers, A. Sherbondy, R. Mackenzie, R. Dougherty, and B. Wandell. Exploration of the Brain's White Matter Pathways with Dynamic Queries. In *Proceedings of IEEE Visualization '04*, pages 377–384, 2004.

[2] A. Aldroubi and P.J. Basser. Reconstruction of Vector and Tensor Fields from Sampled Discrete Data. In *Contemporary Mathematics*, pages 1–15. American Mathematical Society, 1999.

[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*, 48:331–340, 2002.

[4] A.W. Anderson. Theoretical Analysis of the Effects of Noise on Diffusion Tensor Imaging. *Magnetic Resonance in Medicine*, 46:1174–1188, 2001.

[5] J.L. Andersson, M. Richter, W. Richter, S. Skare, R.G. Nunes, M.D. Robson, and T.E. Behrens. Effects of Susceptibility Distortions on Tractography. In *Proceedings of the 12th Annual Meeting of the Internal Society for Magnetic Resonance in Medicine (ISMRM '04)*, page 87, 2004.

[6] A.J. Bartsch, A. Biller, and G.A. Homola. *Tractography for Surgical Targeting*, pages 415–444. Academic Press, 2009.

[7] P. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. Microstructural and Physiological Features of Tissues Elucidated by Quantitative Diffusion Tensor MRI. *Journal of Magnetic Resonance*, 111:209–219, 1996.

[8] P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In Vivo Fiber Tractography using DT-MRI Data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.

[9] C. Beaulieu. The Basis of Anisotropic Water Diffusion in the Nervous System. *NMR in Biomedicine*, 15:435–455, 2002.

[10] S. Beckette and W. Gould. Rangefinder Box Plots. *The American Statistician*, 41:149, 1987.

[11] T.E. Behrens, M.W. Woolrich, M. Jenkinson, H. Johansen-Berg, R.G. Nunes, S. Clare, P.M. Matthews, J.M. Brady, and S.M. Smith. Characterization and Propagation of Uncertainty in Diffusion-Weighted MR Imaging. *Magnetic Resonance Medicine*, 50:1077–1088, 2003.

[12] T.E.J. Behrens and S. Jbabdi. *MR Diffusion Tractography*, pages 333–351. Academic Press, 2009.

[13] T.E.J. Behrens, H. Johansen-Berg, S. Jbabdi, M.F.S. Rushworth, and M.W. Woolrich. Probabilistic Diffusion Tractography with Multiple Fibre Orientations: What Can We Gain? *NeuroImage*, 34:144–155, 2007.

[14] L. Bello, A. Castellano, E. Fava, G. Casaceli, M. Riva, G. Scotti, S.M. Gaini, and A. Falini. Intraoperative Use of Diffusion Tensor Imaging Fiber Tractography and Subcortical Mapping for Resection of Gliomas: Technical Considerations. *Neurosurgical Focus*, 2:E6, 2010.

[15] L. Bello, A. Gambini, A. Castellano, G. Carrabba, F. Acerbi, E. Fava, C. Giussani, M. Cadioli, V. Blasi, A. Casarotti, C. Papagno, A.K. Gupta, S. Gaini, G. Scotti, and A. Falini. Motor and Language DTI Fiber Tracking Combined With Intraoperative Subcortical Mapping for Surgical Removal of Gliomas. *NeuroImage*, 1:369–382, 2008.

[16] Y. Benjamini. Opening the Box of a Box Plot. *The American Statistician*, 42:257–262, 1988.

[17] W. Berger, H. Piringer, P. Filzmoser, and E. Groeller. Uncertainty-Aware Exploration of Continuous Parameter Spaces Using Multivariate Prediction. *Computer Graphics Forum*, 30:911–920, 2011.

[18] J.I. Berman, M.S. Berger, S.W. Chung, S.S. Nagarajan, and R.G. Henry. Accuracy of Diffusion Tensor Magnetic Resonance Imaging Tractography Assessed Using Intraoperative Subcortical Stimulation Mapping and Magnetic Source Imaging. *Journal of Neurosurgery*, 107:488–494, 2007.

[19] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Buehler. High-Quality Multimodal Volume Rendering for Preoperative Planning of Neurosurgical Interventions. *IEEE Transactions on Computer Graphics and Visualization*, 13:1696–1703, 2007.

[20] M. Bjornemo, A. Brun, R. Kikinis, and C.F. Westin. Regularized Stochastic White Matter Tractography Using Diffusion Tensor MRI. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI '02)*, pages 435–442, 2002.

[21] J. Blaas, C.P. Botha, C. Majoie, A. Nederveen, F.M. Vos, and F.H. Post. Interactive Visualization of Fused fMRI and DTI for Planning Brain Tumor Resections. In *SPIE Medical Imaging 2007*, 2007.

[22] J. Blaas, C.P. Botha, B. Peters, F.M. Vos, and F.H. Post. Fast and Reproducible Fiber Bundle Selection in DTI Visualization. In *Proceedings of IEEE Visualization '05*, pages 59–64, 2005.

[23] R.A. Boller, S.A. Braun, J. Miles, and D.H. Laidlaw. Application of Uncertainty Visualization Methods to Meteorological Trajectories. *Earth Science Informatics*, 3:119–126, 2010.

[24] R.P. Botchen, D. Weiskopf, and T. Ertl. Texture-Based Visualization of Uncertainty in Flow Fields. In *Proceedings of IEEE Visualization '05*, pages 647–654, 2005.

[25] R. Brecheisen, A. Vilanova, B. Platel, and B.M. ter Haar Romenij. Flexible GPU-Based Multi-Volume Ray-Casting. In *Vision, Modelling and Visualization 2008 (VMV 2008) : Proceedings of the 13th International Fall Workshop*, pages 1–6, 2008.

[26] A. Brun, H. Knutsson, H-J. Park, M.E. Shenton, and C-F. Westin. Clustering Fiber Traces Using Normalized Cuts. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention, (MICCAI '04)*, pages 368–375, 2004.

[27] U. Buergel, B. Maedler, C.R. Honey, A. Thron, J. Gilsbach, and V.A. Coenen. Fiber Tracking with Distinct Software Tools Results in a Clear Diversity in Anatomical Fiber Tract Portrayal. *Zentralblatt fuer Neurochirurgie*, 70:27–35, 2009.

[28] F. Calamante, J.D. Tournier, G.D. Jackson, and A. Connelly. Track-Density Imaging (TDI): Super-Resolution White Matter Imaging Using Whole-Brain Track-Density Mapping. *NeuroImage*, 53:1233–1243, 2010.

[29] J.S.W. Campbell, K. Siddiqi, V.V. Rymar, A.F. Sadikot, and G.B. Pike. Flow-Based Fiber Tracking with Diffusion Tensor and Q-Ball Data: Validation and Comparison to Principal Diffusion Direction Techniques. *NeuroImage*, 27:725–736, 2005.

[30] T-T. Cao, K. Tang, A. Mohamed, and T-S. Tan. Parallel Banding Algorithm to Compute Exact Distance Transform with the GPU. In *Proceedings of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games '10*, pages 83–90, 2010.

[31] A. Cedilnik and P. Rheingans. Procedural Annotation of Uncertain Information. In *Proceedings of IEEE Visualization '00*, pages 77–84, 2000.

[32] W. Chen, S. Zhang, S. Correia, and D.S. Ebert. Abstractive Representation and Exploration of Hierarchically Clustered Diffusion Tensor Fiber Tracts. *Computer Graphics Forum*, 27:1071–1078, 2008.

[33] O. Ciccarelli, G.J. Parker, A.T. Toosy, C.A. Wheeler-Kingshott, G.J. Barker, P.A. Boulby, D.H. Miller, and A.J. Thompson. From Diffusion Tractography to Quantitative White Matter Tract Measures - A Reproducibility Study. *NeuroImage*, 18:348–359, 2003.

[34] C.A. Clark, T.R. Barrick, M.M. Murphy, and B. Anthony Bell. White Matter Fiber Tracking in Patients with Space-Occupying Lesions of the Brain - A New Technique for Neurosurgical Planning? *NeuroImage*, 20:1601–1608, 2003.

[35] P.A. Cook, D.C. Alexander, and G.J.M. Parker. Modelling Noise-Induced Fibre-Orientation Error in Diffusio-Tensor MRI. In *Proceedings of International Symposium on Biomedical Engineering '04*, pages 332–335, 2004.

[36] I. Corouge, S. Gouttard, and G. Gerig. Towards a Shape Model of White Matter Fiber Bundles using Diffusion Tensor MRI. In *International Symposium on Biomedical Imaging '04*, pages 344–347, 2007.

[37] C.D. Correa, Y.H. Chan, and K.L. Ma. A Framework for Uncertainty-Aware Visual Analytics. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)*, pages 51–58, 2009.

[38] S. Correia, S.Y. Lee, T. Voorn, D.F. Tate, R.H. Paul, S. Zhang, S.P. Salloway, P.F. Malloy, and D.H. Laidlaw. Quantitative Tractography Metrics of White Matter Integrity in Diffusion-Tensor MRI. *NeuroImage*, 42:568–581, 2008.

[39] T. Delmarcelle and L. Hesselink. Visualizing Second-Order Tensor Fields with Hyperstreamlines. *IEEE Computer Graphics and Applications*, 13:25–33, 1993.

[40] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, Fast and Robust Analytical Q-Ball Imaging. *Magnetic Resonance in Medicine*, 56:497–510, 2007.

[41] M. Descoteaux, R. Deriche, T.R. Knoesche, and A. Anwander. Deterministic and Probabilistic Tractography Based on Complex Fibre Orientation Distributions. *IEEE Transactions on Medical Imaging*, 28:269–286, 2009.

[42] S. Djurcilova, K. Kima, P. Lermusiaux, and A. Pang. Visualizing Scalar Volumetric Data with Uncertainty. *Computers and Graphics*, 26:239–248, 2002.

[43] H. Doleisch, M. Gasser, and H. Hauser. Interactive Feature Specification for Focus+Context Visualization of Complex Simulation Data. In *Proceedings of Symposium on Data Visualization '03*, pages 239–248, 2003.

[44] H. Duffau, L. Capelle, N. Sichez, D. Denvil, M. Lopes, and J.P. Sichez. Intraoperative Mapping of Subcortical Language Pathways Using Direct Stimulations: An Anatomo-Functional Study. *Brain*, 125:199–214, 2002.

[45] C.R. Ehlschlaeger, A.M. Shortridge, and M.F. Goodchild. Visualizing Spatial Data Uncertainty Using Animation. *Computers in Geoscience*, 23:387–395, 1997.

[46] F. Enders, N. Saubers, D. Merhof, P. Hastreiter, C. Nimsky, and M. Stamminger. Visualization of White Matter Tracts with Wrapped Streamlines. In *Proceedings of IEEE Visualization '05*, pages 51–58, 2005.

[47] W.W. Esty and J.D. Banfield. The Box-Percentile Plot. *Journal of Statistical Software*, 8, 2003.

[48] J.R. Eyton. Complementary-Color Two-Variable Maps. *Annals of the Association of American Geographers*, 74:477–490, 1984.

[49] D. Feng, L. Kwock, Y. Lee, and R.M. Taylor. Matching Visual Saliency to Confidence in Plots of Uncertain Data. *IEEE Transactions on Visualization and Computer Graphics*, 16:980–989, 2010.

[50] O. Friman, G. Farneback, and C-F. Westin. A Bayesian Approach for Stochastic White Matter Tractography. *IEEE Transactions on Medical Imaging*, 25:965–978, 2006.

[51] G.T. Gering, A. Nabavi, R. Kikinis, N. Hata, L.J. O'Donnell, E.L. Grimson, F.A. Jolesz, P.M. Black, and W.M. Wells. An Integrated Visualization System for Surgical Planning and Guidance Using Image Fusion and Open MR. *Journal of Magnetic Resonance Imaging*, 13:967–975, 2001.

[52] A.J. Golby, G. Kindlmann, I. Norton, A. Yarmarkovich, S. Pieper, and R. Kikinis. Interactive Diffusion Tensor Tractography Visualization for Neurosurgical Planning. *Neurosurgery*, 86:496–505, 2011.

[53] K.M. Goldberg and B. Iglewicz. Bivariate Extensions of the Boxplot. *Technometrics*, 34:307–320, 1992.

[54] H. Griethe and H. Schumann. The Visualization of Uncertain Data: Methods and Problems. In *Proceedings of SimVis '06*, pages 143–156, 2006.

[55] G. Grigoryan and P. Rheingans. Point-Based Probabilistic Surfaces to Show Surface Uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 10:564–573, 2004.

[56] M. Hadwiger, C. Sigg, H. Scharsach, K. Buehler, and M. Gross. Real-Time Ray-Casting and Advanced Shading of Discrete Isosurfaces. *Computer Graphics Forum*, 24:303–312, 2005.

[57] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, V.J. Wedeen, R. Meuli, and J-P. Thiran. Mapping Human Whole-Brain Structural Networks with Diffusion MRI. *PLoS ONE*, 2:1–9, 2007.

[58] P. Hagmann, J-P. Thiran, L. Jonasson, P. Vandergheynst, S. Clarke, P. Maeder, and R. Meuli. DTI Mapping of Human Brain Connectivity - Statistical Fibre Tracking and Virtual Dissection. *NeuroImage*, 19:545–554, 2003.

[59] H.K. Hahn, J. Klein, C. Nimsky, J. Rexilius, and H-O. Peitgen. Uncertainty in Diffusion Tensor Based Fibre Tracking. *Acta Neurochirugica Supplement*, 98:33–41, 2006.

[60] E. Heiervang, T.E.J. Behrens, C.E. Mackay, M.D. Robson, and H. Johansen-Berg. Between Session Reproducibility and Between Subject Variability of Diffusion MRI and Tractography Measures. *NeuroImage*, 33:867–877, 2006.

[61] T. Hendler, P. Pianka, M. Sigal, M. Kafri, D. Ben-Bashat, and S. Constantini. Delineating Gray and White Matter Involvement in Brain Lesions: Three-Dimensional Alignment of Functional Magnetic Resonance and Diffusion-Tensor Imaging. *Journal of Neurosurgery*, 99:1018–1027, 2007.

[62] J.L. Hintze and R.D. Nelson. Violin Plots: A Box Plot-Density Trace Synergism. *The American Statistician*, 52:181–184, 1998.

[63] M. Hlawatsch, W. Leube, P. Nowak, and D. Weiskopf. Flow Radar Glyphs - Static Visualization of Unsteady Flow with Uncertainty. *IEEE Transactions on Visualization and Computer Graphics*, 17:1949–1958, 2011.

[64] T. Hosey, G. Williams, and R. Ansorge. Inference of Multiple Fiber Orientations in High Angular Resolution Diffusion Imaging. *Magnetic Resonance in Medicine*, 54:1480–1489, 2005.

[65] H. Huang, J. Zhang, P.C.M. Van Zijl, and S. Mori. Analysis of Noise Effects on DTI-Based Tractography Using the Brute-Force and Multi-ROI Approach. *Magnetic Resonance in Medicine*, 52:559–565, 2004.

[66] S. Jbabdi, M.W. Woolrich, J.L.R. Andersson, and T.E.J. Behrens. A Bayesian Framework for Global Tractography. *NeuroImage*, 37:116–129, 2007.

[67] M. Jenkinson and S.M. Smith. A Global Optimisation Method for Robust Affine Registration of Brain Images. *Medical Image Analysis*, 5:143–156, 2001.

[68] B. Jeurissen, A. Leemans, D.K. Jones, J-D. Tournier, and J. Sijbers. Probabilistic Fiber Tracking Using the Residual Bootstrap with Constrained Spherical Deconvolution. *Human Brain Mapping*, 32:461–479, 2011.

[69] F. Jiao, J.M. Philips, J. Stinstra, J. Krüger, R. Varma, E. Hsu, J. Korenberg, and C.R. Johnson. Metrics for Uncertainty Analysis and Visualization of Diffusion Tensor Images. In *Proceedings of Medical Imaging and Augmented Reality*, pages 179–190, 2010.

[70] D.K. Jones. Determining and Visualizing Uncertainty in Estimates of Fiber Orientation from Diffusion Tensor MRI. In *Proceedings of International Society for Magnetic Resonance Medicine (ISMRM '02)*, volume 49, pages 7–12, 2003.

[71] D.K. Jones. Tractography Gone Wild - Probabilistic Fibre Tracking Using the Wild Bootstrap with Diffusion Tensor MRI. *IEEE Transactions on Medical Imaging*, 27:1268–1274, 2008.

[72] D.K. Jones and M. Cercignani. Twenty-Five Pitfalls in the Analysis of Diffusion MRI Data. *NMR in Biomedicine*, 23:803–820, 2010.

[73] A.J. Joseph, S.K. Lodha, J.C. Renteria, and A. Pang. UISURF: Visualizing Uncertainty in Isosurfaces. In *Proceedings of Computer Graphics and Imaging*, pages 184–191, 1999.

[74] E. Kaden, T.R. Knoesche, and A. Anwander. Parametric Spherical Deconvolution: Inferring Anatomical Connectivity Using Diffusion MR Imaging. *NeuroImage*, 37:474–488, 2007.

[75] K. Kamada, T. Todo, Y. Masutani, S. Aoki, K. Ino, T. Takano, T. Kirino, N. Kawahara, and A. Morita. Combined Use of Tractography-Integrated Functional Navigation and Direct Fiber Stimulation. *Journal of Neurosurgery*, 102:664–672, 2005.

[76] K. Kamada, T. Todo, T. Ota, Y. Masutani, F. Aoki, S. Takeuchi, K. Kawai, and N. Saito. The Motor-Evoked Potential Threshold Evaluated by Tractography and Electrical Stimulation. *Journal of Neurosurgery*, 4:785–795, 2009.

[77] D. Kao, A. Love, J.L. Dungan, and A. Pang. Visualizing Spatially Varying Distribution Data. In *Proceedings of the 6th International Conference on Information Visualization*, pages 219–225, 2002.

[78] G.E. Keles, D.A. Lundin, K.R. Lamborn, E.F. Chang, G. Ojemann, and M.S. Berger. Intraoperative Subcortical Stimulation Mapping in Hemispherical Perirolandic Gliomas Located Within or Adjacent to the Descending Motor Pathways: Evaluation of Morbidity and Assessment of Functional Outcome in 294 Patients. *Journal of Neurosurgery*, 100:369–375, 2004.

[79] G. Kindlmann. Superquadric Tensor Glyphs. In *Proceedings of IEEE Visualization '04*, pages 147–154, 2004.

[80] G. Kindlmann, D. Weinstein, and D. Hart. Strategies for Direct Volume Rendering of Diffusion Tensor Fields. *IEEE Transactions on Visualization and Computer Graphics*, 6:124–138, 2000.

[81] M. Kinoshita, K. Yamada, N. Hashimoto, A. Kato, S. Izumoto, T. Baba, M. Maruno, T. Nishimura, and T. Yoshimine. Fiber-Tracking Does Not Accurately Estimate Size of Fiber Bundle in Pathological Condition - Initial Neurosurgical Experience Using Neuronavigation and Subcortical White Matter Stimulation. *NeuroImage*, 25:425–429, 2005.

[82] S. Klein, U.A.V.D. Heide, I.M. Lips, M.V. Vulpen, M. Staring, and J.P.W. Pluim. Automatic Segmentation of the Prostate in 3D MR Images by Atlas Matching Using Localized Mutual Information. *Medical Physics*, 35:1407, 2008.

[83] C.G. Koay, J.D. Carew, A.L. Alexander, P.J. Basser, and M.E. Meyerand. Investigation of Anomalous Estimates of Tensor-Derived Quantities in Diffusion Tensor Imaging. *Magnetic Resonance in Medicine*, 55:930–936, 2006.

[84] C.G. Koay, L.C. Chang, J.D. Carew, C. Pierpaoli, and P.J. Basser. A Unifying Theoretical and Algorithmic Framework for Least Squares Methods of Estimation in Diffusion Tensor Imaging. *Journal of Magnetic Resonance*, 182:115–125, 2006.

[85] M.A. Koch, D.G. Norris, and M. Hund-Georgiadis. An Investigation of Functional and Anatomical Connectivity Using Magnetic Resonance Imaging. *NeuroImage*, 16:241–250, 2002.

[86] A. Koehn, F. Weiler, J. Klein, O. Konrad, H.K. Hahn, and H-O. Peitgen. State-of-the-Art Computer Graphics in Neurosurgical Planning and Risk Assessment. In *Proceedings of Eurographics 2007*, pages 117–120, 2007.

[87] D.H. Laidlaw, E.T. Ahrens, D. Kremers, M.J. Avalos, C. Readhead, and R.E. Jacobs. Visualizing Diffusion Tensor Images of the Mouse Spinal Cord. In *Proceedings of IEEE Visualization '98*, pages 127–134, 1998.

[88] M. Lazar and A.L. Alexander. An Error Analysis of White Matter Tractography Methods - Synthetic Diffusion Tensor Field Simulations. *NeuroImage*, 20:1140–1153, 2003.

[89] M. Lazar, D.M. Weinstein, J.S. Tsuruda, K.M. Hasan, K. Arfanakis, M.E. Meyerand, B. Badie, H.A. Rowley, V. Haughton, A. Field, and A.L. Alexander. White Matter Tractography Using Diffusion Tensor Deflection. *Human Brain Mapping*, 18:306–321, 2003.

[90] D. Le Bihan. MR Imaging of Intravoxel Incoherent Motions: Application to Diffusion and Perfusion in Neurologic Disorders. *Radiology*, 116:401–407, 1986.

[91] J. Lee, M. Lazar, J. Lee, J. Holden, E. Terasawa-Grilley, and A.L. Alexander. Correction of B0 EPI Distortions in Diffusion Tensor Imaging and White Matter Tractography. In *Proceedings of the 12th Annual Meeting of the Internal Society for Magnetic Resonance in Medicine (ISMRM '04)*, page 2172, 2004.

[92] A. Leemans and D.K. Jones. The B-Matrix Must Be Rotated When Motion Correcting Diffusion Tensor Imaging Data. *Magnetic Resonance in Medicine*, 61:1336–1349, 2009.

[93] H. Li, C.W. Fu, Y. Li, and A.J. Hanson. Visualizing Large-Scale Uncertainty in Astrophysical Data. *IEEE Transactions on Visualization and Computer Graphics*, 13:1640–1647, 2007.

[94] S.K. Lodha, N.M. Faaland, and A.P. Charaniya. Visualization of Uncertain Particle Movement. In *Proceedings of the Computer Graphics and Imaging Conference*, pages 226–232, 2002.

[95] S.K. Lodha, A.T. Pang, R.E. Sheehan, and C.M Wittenbrink. UFLOW: Visualizing Uncertainty in Fluid Flow. In *Proceedings of IEEE Visualization '96*, pages 249–254, 1996.

[96] S.K. Lodha, B. Sheehan, A. Pang, and C. Wittenbrink. Visualizing Geometric Uncertainty of Surface Interpolants. In *Proceedings of the Conference on Graphics Interface '96*, pages 238–245, 1996.

[97] W.E. Lorensen and H.E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *Computer Graphics*, 21:163–169, 1987.

[98] N.F. Lori, E. Akbudak, J.S. Shimony, T.S. Cull, A.Z. Snyder, R.K. Guillory, and T.E. Conturo. Diffusion Tensor Fiber Tracking of Human Brain Connectivity - Acquisition Methods, Reliability Analysis and Biological Results. *NMR in Biomedicine*, 15:494–515, 2002.

[99] T. Lotufo, A. Falcao, and F. Zampirolli. Fast Euclidian Distance Transform Using a Graph Search Algorithm. In *Proceedings of 13th Brazilian Symposium on Computer Graphics and Image Processing*, pages 269–275, 2000.

[100] A.L. Love, A. Pang, and D.L. Kao. Visualizing Spatial Multivalue Data. *IEEE Computer Graphics and Applications*, 25:69–79, 2005.

[101] C. Lundstroem, P. Ljung, A. Persson, and A. Ynnerman. Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation. *IEEE Transactions on Visualization and Computer Graphics*, 13:1648–1655, 2007.

[102] A.M. MacEachren, A. Robinson, S. Hopper, S. Gardner, R. Murray, M. Gahegan, and E. Hetzler. Visualizing Geospatial Information Uncertainty: What We Know and What We Need to Know. *Cartography and Geographic Information Science*, 32:139–160, 2005.

[103] S. Madi, K.M. Hasan, and P.A. Narayana. Diffusion Tensor Imaging of In Vivo and Excised Rat Spinal Cord at 7T with an Icosahedral Encoding Scheme. *Magnetic Resonance in Medicine*, 53:118–125, 2005.

[104] M.M. Malik, C. Heinzl, and M.E. Groeller. Comparative Visualization for Parameter Studies of Dataset Series. *IEEE Transactions on Visualization and Computer Graphics*, 16:829–840, 2010.

[105] T. McGraw and M. Nadar. Stochastic DT-MRI Connectivity Mapping on the GPU. *IEEE Transactions on Visualization and Computer Graphics*, 13:1504–1511, 2007.

[106] A. Meijster, J. Roerdink, and W. Hesselink. A General Algorithm for Computing Distance Transforms in Linear Time. In *Mathematical Morphology and Its Applications to Image and Signal Processing*, pages 331–340. Springer, 2000.

[107] D. Merhof, M. Meister, E. Bingöl, C. Nimksy, and G. Greiner. Isosurface-Based Generation of Hulls Encompassing Neuronal Pathways. *Stereotactic and Functional Neurosurgery*, 87:50–60, 2009.

[108] N. Mikuni, T. Okada, R. Enatsu, Y. Miki, J.A. Urayama, S. Takahashi, K. Nozaki, H. Fukuyama, and N. Hashimoto. Clinical Significance of Preoperative Fibre-Tracking to Preserve the Affected Pyramidal Tracts During Resection of Brain Tumours in Patients with Preoperative Motor Weakness. *Journal of Neurology, Neurosurgery, and Psychiatry*, 78:716–721, 2007.

[109] B. Moberts, A. Vilanova, and J.A. van Wijk. Evaluation of Fiber Clustering Methods for Diffusion Tensor Imaging. In *Proceedings of IEEE Visualization '05*, pages 65–72, 2005.

[110] S. Mori and P.C.M. van Zijl. Fiber Tracking - Principles and Strategies. *NMR in Biomedicine*, 15:468–480, 2002.

[111] J. Mullikin. The Vector Distance Transform in Two and Three Dimensions. *Graphical Models and Image Processing*, 54:526–535, 1992.

[112] A. Neubauer, S. Wolfsberger, M.T. Forster, L. Mroz, R. Wegenkittl, and K. Buehler. Advanced Virtual Endoscopic Pituitary Surgery. *IEEE Transactions on Visualization and Computer Graphics*, 11:497–507, 2005.

[113] L.J. O'Donnell, C-F. Westin, and A.J. Golby. Tract-Based Morphometry for White Matter Group Analysis. *NeuroImage*, 45:832–844, 2009.

[114] L.J. O'Donnell and C.F. Westin. Automatic Tractography Segmentation Using a High-Dimensional White Matter Atlas. *IEEE Transactions on Medical Imaging*, 26:1562–1575, 2007.

[115] K. Oishi, A.V. Faria, P.C.M. Van Zijl, and S. Mori. *MRI Atlas of Human White Matter*. Academic Press, 2011.

[116] C. Olston and J.D. Mackinlay. Visualizing Data with Bounded Uncertainty. In *Proceedings of IEEE Symposium on Information Visualization '02*, pages 37–40, 2002.

[117] R.S. Osorio and K.W. Brodlie. Contouring With Uncertainty. In *Proceedings of the 6th Theory and Practice of Computer Graphics Conference*, pages 59–66, 2008.

[118] R. Otten, A. Vilanova, and H. van de Wetering. Illustrative White Matter Fiber Bundles. *Computer Graphics Forum*, 29:1013–1022, 2010.

[119] E. Ozarslan, 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*, 36:1086–1103, 2006.

[120] S. Pajevic and P.J. Basser. Parametric and Non-Parametric Statistical Analysis of DT-MRI Data. *Journal of Magnetic Resonance*, 161:1–14, 2003.

[121] A.T. Pang, C.M. Wittenbrink, and S.K. Lodha. Approaches to Uncertainty Visualization. *The Visual Computer*, 13:370–390, 1997.

[122] G.J.M. Parker, H.A. Haroon, and C.A.M. Wheeler-Kingshott. A Framework for a Streamline-Based Probabilistic Index of Connectivity (PICo) using a Structural Interpretation of MRI Diffusion Measurements. *Journal of Magnetic Resonance Imaging*, 18:242–254, 2003.

[123] T.H.J.M. Peeters, V. Prckovska, M.A. Van Almsick, A. Vilanova, and B.M. Ter Haar Romeny. Fast and Sleek Glyph Rendering for Interactive HARDI Data Exploration. In *Proceedings of IEEE Pacific Visualization Symposium*, pages 153–160, 2009.

[124] T.H.J.M. Peeters, A. Vilanova, G.J. Strijkers, and B.M. Ter Haar Romeny. Visualization of the Fibrous Structure of the Heart. In *Proceedings of Vision, Modeling and Visualization (VMV '06)*, pages 309–316, 2006.

[125] C. Pierpaoli and P.J. Basser. Toward a Quantitative Assessment of Diffusion Anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, 1996.

[126] K. Pothkow and H.C. Hege. Positional Uncertainty of Isocontours: Condition Analysis and Probabilistic Measures. *IEEE Transactions on Visualization and Computer Graphics*, 17:1393–1406, 2010.

[127] K. Pothkow, B. Weber, and H.C. Hege. Probabilistic Marching Cubes. *Computer Graphics Forum*, 30:931–940, 2011.

[128] K. Potter, J. Kniss, R. Riesenfeld, and C.R. Johnson. Visualizing Summary Statistics and Uncertainty. *Computer Graphics Forum*, 29:823–831, 2010.

[129] V. Prckovska, T.H.J.M. Peeters, M. van Almsick, B.M. ter Haar Romeny, and A. Vilanova. Fused DTI/HARDI Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17:1407–1419, 2011.

[130] B. Preim, C. Tietjen, W. Spindler, and H-O. Peitgen. Integration of Measurement Tools in 3D Medical Visualization. In *Proceedings of IEEE Visualization '02*, pages 21–28, 2002.

[131] B. Reitinger, D. Schmalstieg, A. Bornik, and R. Beichel. Spatial Analysis Tools for Virtual Reality-Based Surgical Planning. In *Proceedings of IEEE Symposium on 3D User Interfaces '06*, pages 27–44, 2006.

[132] D. Reniers and A.C. Telea. Tolerance-Based Feature Transforms. In *Proceedings of Advances in Computer Graphics and Computer Vision (VISAPP '06/GRAPP '06)*, pages 187–200, 2008.

[133] P. Rheingans. Task-Based Color Scale Design. In *Proceedings of Applied Image and Pattern Recognition (SPIE '99)*, pages 35–43, 1999.

[134] P.J. Rhodes, R.S. Laramee, R.D. Bergeron, and T.M. Sparr. Uncertainty Visualization Methods in Isosurface Rendering. In *Proceedings of Eurographics*, pages 83–88, 2003.

[135] C. Rieder, A. Weihusen, C. Schumann, S. Zidowitz, and H-O. Peitgen. Visual Support for Interactive Post-Interventional Assessment of Radiofrequency Ablation Therapy. *Computer Graphics Forum*, 29:1093–1102, 2010.

[136] F. Roessler, R.P. Botchen, and T. Ertl. Dynamic Shader Generation for GPU-Based Multi-Volume Ray Casting. *IEEE Transactions on Computer Graphics and Visualization*, 28:66–77, 2008.

[137] F. Roessler, E. Tejada, T. Fangmeier, and T. Ertl. GPU-Based Multi-Volume Rendering for the Visualization of Functional Brain Images. In *Proceedings of SimVis 2006*, pages 3015–318, 2006.

[138] I. Roessling, C. Cyrus, L. Dornheim, A. Boehm, and B. Preim. Fast and Flexible Distance Measures for Treatment Planning. *International Journal for Computer Assisted Radiology and Surgery*, 5:633–646, 2010.

[139] M.B. Rosson and J.M. Carroll. *Usability Engineering: Scenario-Based Development of Human-Computer Interaction*. Morgan Kaufmann, 2002.

[140] P.J. Rousseeuw, I. Ruts, and J.W. Tukey. The Bagplot: A Bivariate Boxplot. *The American Statistician*, 53:382–287, 1999.

[141] L. Serra, R.A. Kockro, C.G. Guan, N. Hern, E.C.K. Lee, Y.H. Lee, C. Chan, and W.L. Nowinski. Multimodal Volume-Based Tumor Neurosurgery Planning in the Virtual Workbench. In *Proceedings of MICCAI '98*, pages 1007–1015, 1998.

[142] A.J. Sherbondy, R.F. Dougherty, M. Ben-Shachar, S. Napel, and B. Wandell. ConTrack: Finding the Most Likely Pathways Between Brain Regions Using Diffusion Tractography. *Journal of Vision*, 8:1–16, 2008.

[143] S. Silva, J. Madeira, and B.S. Santos. There is More to Color Scales Than Meets the Eye - A Review on the Use of Color in Visualization. In *Proceedings of Information Visualization '07*, pages 943–950, 2007.

[144] E.O. Stejskal and J.E. Tanner. Spin Diffusion Measurements - Spin Echoes in the Presence of a Time-Dependent Field Gradient. *Journal of Chemical Physics*, 42:288–292, 1965.

[145] T. Taoka, M. Morikawa, T. Akashi, T. Miyasaka, H. Nakagawa, K. Kiuchi, T. Kishimoto, and K. Kichikawa. Fractional Anisotropy - Threshold Dependence in Tract-Based Diffusion Tensor Analysis: Evaluation of the Uncinate Fasciculus in Alzheimer Disease. *American Journal of Neuroradiology*, 30:1700–1703, 2009.

[146] A.C. Telea and A.C. Jalba. Voxel-Based Assessment of Printability of 3D Shapes. In *Proceedings of Mathematical Morphology and Its Applications to Image and Signal Processing (ISMM '10)*, pages 393–404, 2011.

[147] A.C. Telea and J.J. van Wijk. An Augmented Fast Marching Method for Computing Skeletons and Centerlines. In *Proceedings of Symposium on Data Visualization '02*, pages 64–72, 2002.

[148] M. Termeer, J. Olivan Bescos, A. Vilanova, F. Gerritsen, and M.E. Groeller. CoViCAD: Comprehensive Visualization of Coronary Artery Disease. *IEEE Transactions on Visualization and Computer Graphics*, 13:1632–1639, 2007.

[149] J. Tongkumchum, B. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. Two-Dimensional Box Plot. *Songklanakarin Journal of Science and Technology*, 27:859–866, 2005.

[150] J.D. Tournier, F. Calamante, M.D. King, D.G. Gadian, and A. Connelly. Limitations and Requirements of Diffusion Tensor Fiber Tracking: An Assessment Using Simulations. *Magnetic Resonance in Medicine*, 47:701–708, 2002.

[151] D.S. Tuch, T.G. Reese, and M.R. Wiegell. High Angular Resolution Diffusion Imaging Reveals Intravoxel White Matter Fiber Heterogeneity. *Magnetic Resonance in Medicine*, 48:577–582, 2002.

[152] J.W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[153] C.C. Van Donkelaar, L.J.G. Kretzers, P.H.M. Bovendeerd, L.M.A. Lataster, K. Nicolay, J.D. Janssen, and M.R. Drost. Diffusion Tensor Imaging in Biomechanical Studies of Skeletal Muscle Function. *Journal of Anatomy*, 194:79–88, 1999.

[154] A. Van Doorn, P.H.M. Bovendeerd, K. Nicolay, M.R. Drost, and J.D. Janssen. Determination of Muscle Fibre Orientation Using Diffusion-Weighted MRI. *European Journal of Morphology*, 34:5–10, 1996.

[155] A. Vilanova, G. Berenschot, and C. van Pul. DTI Visualization with Streamsurfaces and Evenly-Spaced Volume Seeding. In *Proceedings of IEEE TCVG Symposium on Visualization '04*, pages 173–182, 2004.

[156] A. Vilanova, S. Zhang, G. Kindlmann, and D.H. Laidlaw. *Visualization and Image Processing of Tensor Fields*, chapter An Introduction to Visualization of Diffusion Tensor Imaging and Its Applications, pages 121–153. Mathematics and Visualization. Springer Verlag, 2004.

[157] A.N. Voineskos, L.J. O'Donnel, N.J. Lobaugh, D. Markant, S.H. Ameis, M. Niethammer, B.H. Mulsant, B.G. Pollock, J.L. Kennedy, C.F. Westin, and M.E. Shenton. Quantitative Examination of a Novel Clustering Method Using Magnetic Resonance Diffusion Tensor Tractography. *NeuroImage*, 45:370–376, 2009.

[158] S.B. Vos, D.K. Jones, M.A. Viergever, and A. Leemans. Partial Volume Effect as a Hidden Covariate in DTI Analyses. *NeuroImage*, 55:1566–1576, 2011.

[159] S. Wakana, A. Caprihan, M.M. Panzenboeck, J.H. Fallon, M. Perry, R.L. Gollub, K. Hua, J. Zhang, H. Jiang, and P. Dubey. Reproducibility of Quantitative Tractography Methods Applied to Cerebral White Matter. *NeuroImage*, 36:630–644, 2007.

[160] J. Wang and Y. Tan. Efficient Euclidian Distance Transform Using Perpendicular Bisector Segmentation. In *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR '11)*, pages 1625–1632, 2011.

[161] F. Weiler, H.K. Hahn, A. Koehn, O. Friman, J. Klein, and H-O. Peitgen. Dealing with Inaccuracies in Multimodal Neurosurgical Planning - A Preliminary Concept. In *Proceedings of the 22nd Internal Congress and Exhibition of Computer Assisted Radiology and Surgery (CARS)*, pages 77–78, 2008.

[162] D. Weinstein, G. Kindlmann, and E. Lundberg. Tensorlines - Advection-Diffusion Based Propagation Through Diffusion Tensor Fields. In *Proceedings of IEEE Visualization '99*, pages 249–254, 1999.

[163] A. Wenger, D.F. Keefe, S. Zhang, and D.H. Laidlaw. Interactive Volume Rendering of Thin Thread Structures within Multivalued Scientific Data Sets. *IEEE Transactions on Visualization and Computer Graphics*, 10:664–672, 2004.

[164] C-F. Westin, S.E. Maier, H. Mamata, A. Nabavi, F.A. Jolesz, and R. Kikinis. Processing and Visualization of Diffusion Tensor MRI. *Medical Image Analysis*, 6:93–108, 2002.

[165] B. Whitcher, D.S. Tuch, J.J. Wisco, A.G. Sorensen, and L. Wang. Using the Wild Bootstrap to Quantify Uncertainty in Diffusion Tensor Imaging. *Human Brain Mapping*, 29:346–362, 2008.

[166] L. Wilkinson. Dot Plots. *The American Statistician*, 53:276–281, 1999.

[167] T.H. Williams, N. Gluhbegovic, and J.Y. Jew. *The Human Brain: Dissections of the Real Brain*. Virtual Hospital, The University of Iowa, College of Medicine, 1997.

[168] C.M. Wittenbrink, A.T. Pang, and S.K. Lodha. Glyphs for Visualizing Uncertainty in Vector Fields. *IEEE Transactions on Visualization and Computer Graphics*, 2:266–279, 1996.

[169] J.S. Wu, Y. Mao, L.F. Zhou, W.J. Tang, and J.S. Hu. Clinical Evaluation and Follow-Up Outcome of Diffusion Tensor Imaging-Based Functional Neuronavigation. *Neurosurgery*, 61:935–949, 2007.

[170] Z. Xie, S. Huang, M. Ward, and E. Rundensteiner. Exploratory Visualization of Multivariate Data With Variable Quality. In *Proceedings of IEEE Symposium on Visual Analytics and Technology*, pages 183–190, 2006.

[171] P.A. Yushkevich, J. Piven, C. Hazlett, H. Smith, G. Smith, R. Ho, S. Ho, J.C. Gee, and G. Gerig. User-Guided 3D Active Contour Segmentation of Anatomical Structures. *NeuroImage*, 31:1116–1128, 2006.

[172] S. Zhang, C. Demiralp, and D.H. Laidlaw. Visualizing Diffusion Tensor MR Images Using Streamtubes and Streamsurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 9:454–462, 2003.

Summary

Diffusion tensor imaging (DTI) is an imaging technique based on magnetic resonance that describes, in each point of the tissue, the distribution of diffusing water molecules. The distribution is mathematically modelled using a second-order tensor. In fibrous tissues the diffusion tensor will have an elongated, ellipsoid shape whose main axis is assumed to be aligned with the underlying fiber structure. Fiber tractography traces paths through the tensor field by following each tensor's main direction thereby resulting in a three-dimensional reconstruction of the fibers. This is particularly interesting for the exploration and visualization of neuronal connections in brain white matter and has great potential for applications in neuroscience and neurosurgery.

DTI and fiber tractography are unique in that they provide insight into white matter structures in vivo and non-invasively. However, despite these capabilities the application of DTI and fiber tractography in clinical practice remains limited. The image acquisition and post-processing pipeline is complex and consists of many stages. At each stage errors and uncertainties are introduced due to image noise, magnetic distortions, partial volume effects, scanner settings, diffusion model assumptions and user parameters. These uncertainties are propagated through the pipeline and possibly enhanced in subsequent stages thereby leading to potentially unreliable results in the final tractography output. To the user the processing pipeline behaves like a black box whose internal details remain hidden and whose quality of output cannot be reliably assessed. Contrary to standard CT and MR images it is not possible to look at the "raw" diffusion-weighted images. Without further processing the images are practically meaningless. This means the user either has to accept (and trust) the processing output or refrain from using fiber tracking all together.

In this thesis we assume that the user has certain reservations about the quality of the tractography output. Unfortunately, there is no gold standard against which the output of tractography can be validated. Consequently, we cannot make definitive statements about the "true" certainty or uncertainty of fiber reconstructions. We can, however, discuss tractography output in terms of stability and reproducibility. The output of tractography algorithms can be subject to large variations. In this thesis we present a number of visualization strategies that make these variations visible to the user and allow a better assessment of the reliability of fiber reconstructions obtained from any given tractography algorithm.

Samenvatting (summary in Dutch)

Diffusie tensor imaging (DTI) is een beeldvormingstechniek gebaseerd op magnetische resonantie (MR) welke, in ieder punt van het weefsel, de diffusie van water moleculen beschrijft. De distributie van diffunderende water moleculen wordt wiskundig gemodelleerd met behulp van een tweede-orde tensor. In vezelachtige weefsels zal de diffusie tensor een uitgerekte, ellipsoide vorm hebben waarvan aangenomen wordt dat de lange as parallel loopt aan de richting van de onderliggende vezelstructuur. Vezeltraktografie, of kortweg traktografie, reconstrueert de 3-dimensionale geometrie van vezelbanen door in ieder punt de lokale richting van het tensorveld te volgen. Dit is vooral interessant voor het exploreren en visualiseren van neuronale verbindingen in de witte stof van het brein en daarmee van groot potentieel nut voor toepassingen binnen de neurowetenschap en de neurochirugie.

DTI en traktografie zijn uniek in zoverre dat ze op niet-invasieve wijze, en in-vivo, inzicht verschaffen in de 3-dimensionale geometrie van witte stof structuren. Echter, ondanks deze mogelijkheden blijft de toepassing van DTI en traktografie in de klinische praktijk beperkt. De beeldacquisitie en -verwerkingspijplijn van DTI is complex en bestaat uit vele stappen. Bij iedere stap worden fouten en onzekerheden geintroduceerd, bijvoorbeeld door beeldruis, magnetische vervorming, partiele volume-effecten, foutieve modelaannames en variabele scanner- en gebruikersinstellingen. Deze onzekerheden propageren door de beeldacquisitie- en verwerkingspijplijn en worden mogelijk versterkt na iedere bewerkingsstap. Dit leidt potentieel tot een aanzienlijke onzekerheid in het uiteindelijke resultaat van de traktografie. Voor de gebruiker fungeert de pijplijn als een zwarte doos waarvan de interne werking verborgen blijft en de kwaliteit van het uitgangsresultaat niet betrouwbaar ingeschat kan worden. In tegenstelling tot CT en standaard MRI beelden, is het niet mogelijk om de ruwe diffusie-gewogen beelden te inspecteren. Het aantal beelden is te groot en zonder verdere beeldverwerking zijn de beelden praktisch onbegrijpelijk. Dit betekent dat de gebruiker, hetzij het bewerkte eindresultaat moet vertrouwen of geheel moet afzien van het gebruik van traktografie.

In dit proefschrift nemen we aan dat de gebruiker enige twijfels heeft over de kwaliteit van het traktografieresultaat. Helaas bestaat er geen alom geaccepteerde gouden standaard waartegen het resultaat gevalideerd kan worden. Om die reden kunnen we geen absolute uitspraken doen over de accuraatheid van de vezelreconstructies. Echter, we kunnen het traktografieresultaat wel beschouwen in termen van stabiliteit en reproduceerbaarheid. Er zijn vele bronnen van variatie aanwezig in de DTI beeldacquisitie- en verwerkingspijplijn. In dit proefschrift presenteren we visualizatietechnieken die deze variaties en onzekerheden zichtbaar maken voor de gebruiker. Hiermee kan een betere inschatting gemaakt worden van de betrouwbaarheid van vezelreconstructies zoals deze door een verscheidenheid aan traktografiealgoritmes geproduceerd worden.

List of publications

Refereed journals

- **R. Brecheisen**, B. Platel, B.M. ter Haar Romeny, A. Vilanova. Illustrative Uncertainty Visualization of DTI Fiber Pathways. In *The Visual Computer*, to appear, 2012.
- **R. Brecheisen**, B. Platel, A. Vilanova, B.M. ter Haar Romeny. Parameter Sensitivity Visualization in DTI Fiber Tracking. In *IEEE Transactions on Visualization and Computer Graphics*, pages 1441–48, 2009.

Refereed conference proceedings

- **R. Brecheisen**, A. Vilanova, B. Platel, B.M. ter Haar Romeny. Flexible GPU-Based Multi-Volume Ray-Casting. In *Proceedings of Vision, Modeling, and Visualization (VMV)*, pages 1–6, Konstanz, Germany, 2008.
- **R. Brecheisen**, A. Vilanova, B. Platel, B.M. ter Haar Romeny. Hardware-Accelerated 3D Multimodal Visualization Techniques for Tumor Resection in Neurosurgery. In *Proceedings of Computer Assisted Radiology and Surgery (CARS)*, Barcelona, Spain, 2008.

Refereed abstracts

- **R. Brecheisen**, A. Vilanova, B. Platel, B.M. ter Haar Romeny. Visualizing Parameter Variations and Their Effect on Quantitative Tract Metrics in Diffusion Tensor Tractography. In *Proceedings of the 2nd Annual Meeting of the International Society of Magnetic Resonance in Medicine (ISMRM)*, Utrecht, Netherlands, 2010.

Conference posters

- **R. Brecheisen**, B. Platel, B.M. ter Haar Romeny, A. Vilanova. Illustrative Uncertainty Visualization in DTI Fiber Tracking. *IEEE Symposium on Visualization (EuroVis)*, Bergen, Norway, 2011.
- **R. Brecheisen**, B. Platel, B.M. ter Haar Romeny, A. Vilanova. Illustrative Confidence Intervals for Visualizing DTI Uncertainty. *Biomedical Engineering (BME) Research Day*, Eindhoven, Netherlands, 2011 (2nd poster prize).
- **R. Brecheisen**, Parameter Sensitivity Visualization for DTI. *2nd Annual Meeting of the International Society of Magnetic Resonance in Medicine (ISMRM)*, Utrecht, Netherlands, 2010.
- **R. Brecheisen**, Visualizing Parameter Sensitivity in DTI Fiber Tracking. *Biomedical Engineering (BME) Research Day*, Eindhoven, Netherlands, 2009.

Submitted / under revision

- **R. Brecheisen**, B. Platel, A.C. Telea, G.J. Rutten, P.L. Kubben, B.M. ter Haar Romeny, A. Vilanova. Visualization of the Effects of DTI Uncertainty on Distance

Measurements in Neurosurgical Planning. Submitted to *IEEE Transactions on Visualization and Computer Graphics*, 2012

Acknowledgements

Many people have contributed to the realization of this thesis, either in a direct and undeniable way or indirectly through interesting and fruitful discussions. Special thanks go to my promotor Bart ter Haar Romeny and co-promotors Anna Vilanova and Bram Platel. But I also thank, in random order, Roy van Pelt, Luc Florack, Geert-Jan Rutten, Pieter Kubben, Remco Duits, Erik Franken, Bart Janssen, Vesna Prckovska, Paulo Rodrigues, Tim Peeters, Neda Sepasian, Alessandro Becciu, Pauly Ossenblok, Alexandru Telea, Meister Edi Groeller, Bernhard Preim, Veerle Visser-Vandewalle, Jack van Wijk, Ellen Brunenberg, Marieke Heisen, Hans van Assen, Frans Kanters, Evgeniya Balmashnova, Andrea Fuster, Rocco Gasteiger, Alexandra Baer, Mathias Neugebauer, Steffen Oeltze, Martin Haidacher, Markus Hadwiger, Johanna Beyer and Maurice Termeer. Less for research-related topics, but all the more so for fun and pleasure, I thank Bart Rienties, Joost Mulders and my good friend and business partner (hopefully for a long time to come) Roland Beuker. Thanks also to Jos, Lucie and Laurens Coolsen for many happy moments during the last years.

For helping me conclude the final phase of my PhD and standing by my side at the public defense, I thank my two 'paranimfen', Leo van Lippen and Wilco Kroon.

If you think to yourself "who designed this nice thesis cover?". Well, the credits for that all go to my sister, Kim. Thanks very much for taking the trouble to help me out with this. I'd also like to thank her husband Arjen Rauwerdink and, of course, the children Jon and Fen. I really love being their uncle :)

I'd also like to thank my girlfriend Marielle Coolsen for her support during my study years and my PhD. Marielle, I'm really happy to be with you. As they say, opposites attract and this seems to work incredibly well for us.

And finally, last but not least, the two people who are absolutely at the top of my list to thank from the bottom of my heart: my parents, Frans and Ivonne. Without you I would never have taken the step to go back to the university and start a PhD. Having a 30-year old son moving back into your house must have taken some getting used to. Still, you always thought that I should do what I enjoy most and I have done exactly that. It has made me a very happy man also because my choice to go back to school resulted directly in meeting the love of my life, Marielle Coolsen. I sincerely hope that, when the time comes, I can be as open-minded and supportive to my children as you have been with me.

Biography

Ralph Brecheisen was born on the 16th of July 1973 in Eersel, the Netherlands. He completed VWO (preparatory scientific education) in 1992 at the Rythovius College in Eersel, the Netherlands. After receiving his propedeuse (first-year) diploma in 1993 at the Hogere Zeevaartschool (Nautical Officer School) in Vlissingen, he switched to Mechanical Engineering and received his bachelor degree at Fontys Hogescholen in Eindhoven. He continued with a post-bachelor study in Computer Science, also at Fontys Eindhoven, after which he worked as a software engineer for six years at various companies such as Thales B.V. in Hengelo, UPAY Systems in Zwolle, ICT Embedded Software in Eindhoven and Philips Healthcare in Best.

In 2003, Ralph decided to go back to school and started a study in Biomedical Engineering at Eindhoven University of Technology. He graduated cum laude in 2007 and continued with a PhD project in medical visualization from 2008 until 2012. The current thesis is the end result of this fruitful period. Ralph is currently working as a part-time software developer in the Mathematics department of Eindhoven University of Technology. In his spare time, he runs a software company with friend and business partner Roland Beuker, developing cloud applications for small-to-medium businesses.