

Comparative and ensemble visualization of diffusion tensor fields

Zhang, Changgong

DOI

[10.4233/uuid:5b0da969-6e43-4cfa-bd33-09494ecfe199](https://doi.org/10.4233/uuid:5b0da969-6e43-4cfa-bd33-09494ecfe199)

Publication date

2017

Document Version

Final published version

Citation (APA)

Zhang, C. (2017). *Comparative and ensemble visualization of diffusion tensor fields*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:5b0da969-6e43-4cfa-bd33-09494ecfe199>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

**COMPARATIVE AND ENSEMBLE
VISUALIZATION OF DIFFUSION TENSOR
FIELDS**

COMPARATIVE AND ENSEMBLE VISUALIZATION OF DIFFUSION TENSOR FIELDS

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 12 december 2017 om 12:30 uur

door

Changgong ZHANG

Master of Science, School of Mechanical Engineering and Automation,
Beihang University, Beijing, China,
geboren te Binzhou, Shandong, China.

Dit proefschrift is goedgekeurd door de

promotor: Prof. dr. E. Eisemann

copromotor: Dr. A. Vilanova Bartroli

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. E. Eisemann,	Technische Universiteit Delft
Dr. A. Vilanova Bartroli,	Technische Universiteit Delft

Onafhankelijke leden:

Prof. dr. ir. L.J. van Vliet,	Technische Universiteit Delft
Prof. dr. J.B.T.M. Roerdink,	Rijksuniversiteit Groningen
Prof. dr. I. Hotz,	Linköping University
Prof. dr. T. Schultz,	University of Bonn
Dr. A. Fuster,	Technische Universiteit Eindhoven
Prof. dr. D. Merhof,	RWTH Aachen University
Prof. dr. ir. M.J.T. Reinders,	Technische Universiteit Delft, reservelid



Keywords: Diffusion Tensor Fields, Comparative Visualization, Ensemble Visualization, Glyph Design, Visual Analysis

Printed by: Gildeprint

Front & Back: The adapted © South Park characters are used to illustrate the main research question of this thesis. The *longbow*, which is the literal translation of my given name Changgong, is used to indicate myself. The *tender* glyph, which is placed at the head of the arrow, is one of the main contributions of this thesis.

Copyright © 2017 by C. Zhang

ISBN 978-94-92516-96-1

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>.

*It is not important what people think when you come,
but what they think when you leave.*

Jürgen Klopp

CONTENTS

Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 Motivation	2
1.2 Aim	4
1.3 Outline	6
References	7
2 Magnetic Resonance Imaging and User-guided Compressed Sensing	11
2.1 MRI Principles	12
2.2 Compressed Sensing MRI	14
2.3 User-guided Compressed Sensing MRI	15
2.3.1 Methodology	17
2.3.2 Results and Discussion	19
2.4 Conclusion	23
References	23
3 Background	25
3.1 Brain Anatomy	26
3.2 Diffusion Tensor Imaging	27
3.3 Second-order Symmetric Tensor	30
3.4 State-of-the-Art Visualization	31
3.4.1 Standard Visualization for DTI	32
3.4.2 Comparative/Ensemble Visualization	34
References	37

4	Glyph-based Comparative Visualization for Diffusion Tensor Fields	45
4.1	Introduction	46
4.2	Glyph Design	48
4.2.1	Task Analysis	48
4.2.2	Tensor Dissimilarity Measure	49
4.2.3	Visual Mapping	50
4.3	Glyph Visualization	56
4.3.1	Implementation	58
4.4	Evaluation and Application	59
4.4.1	User Study	59
4.4.2	Case Study	62
4.5	Conclusions	65
	References	66
5	Overview + Detail Visualization for An Ensemble of Diffusion Tensor Fields	71
5.1	Introduction	72
5.2	Statistical Summary	73
5.2.1	Diffusion Tensor Decomposition	73
5.2.2	Diffusion Tensor Ensemble Aggregation	74
5.3	Visual Design Requirement Analysis	79
5.4	Glyph-based Overview Visualization	80
5.4.1	Tensor-based Glyph Design	80
5.4.2	dODF-based Glyph Design	83
5.4.3	Spatial Overview	84
5.5	Detail Visualization	84
5.5.1	Tensor Scale Visualization	85
5.5.2	Tensor Shape Space Visualization	85
5.5.3	Orientation Dissimilarity Matrix Visualization	85
5.5.4	Linked Brushing	86
5.6	Implementation	87
5.7	Evaluation and Application	87
5.7.1	Method Comparison	88
5.7.2	Application to An DTI Group of Healthy Subjects	91

5.8	Conclusions	94
	References	95
6	Comparative Visualization for Ensembles of Diffusion Tensor Fields at Multiple Levels of Detail	99
6.1	Introduction	100
6.2	Comparative Visualization Design	101
6.2.1	Direct Tensor Ensemble Comparison	101
6.2.2	Ensemble Summary Comparison	106
6.2.3	Summary	109
6.3	Application	110
6.3.1	Case Study 1: Fornix	112
6.3.2	Case Study 2: Corpus Callosum	115
6.4	Conclusions and Future Work	116
	References	117
7	Conclusion	121
7.1	Contributions	121
7.2	Future Work	123
	References	124
	Curriculum Vitæ	125
	List of Publications	127

SUMMARY

Scientific visualization of tensor fields is challenging due to the complex and multivariate nature of tensor data. The visualization of multiple tensor fields becomes even more difficult, and still in its infancy. This thesis aims at contributing visual analysis techniques for multiple 3D tensor fields.

We focus specifically on the visual analysis of Diffusion Tensor Imaging (DTI) datasets. DTI is a magnetic resonance imaging (MRI) based modality, which is commonly used in neuroscience to investigate brain white matter in vivo. It requires a long scanning time compared to other imaging modalities. Acceleration of MRI acquisitions has the potential to improve the applicability of DTI. Compressed sensing (CS) is a signal reconstruction technique that is used to accelerate MRI acquisitions. The traditional CS method aims at optimizing the global quality of the reconstructed image. However, in practice, the quality of local structures is often of more interest. Therefore, we investigate CS for this purpose and contribute in this direction by adapting the traditional CS reconstruction method to focus on the quality of local structures.

DTI is able to measure the diffusion profile of water molecules within each voxel, which is influenced by the underlying fibrous structure of white matter, and models it as a second-order symmetric positive-definite tensor (i.e., the so-called diffusion tensor). The diffusion tensor, mathematically expressed as a 3-by-3 symmetric positive-definite matrix, can be decomposed into scale, shape, and orientation, which are its intrinsic properties. It is the intrinsic properties of the diffusion tensor that make DTI a uniquely important imaging modality. Each of them has a biologically meaningful interpretation for the underlying tissue properties. This thesis makes heavy use of tensor intrinsic properties in developing visual analysis techniques for multiple DTI datasets because neuroscientists can put direct interpretations on them.

When comparing two tensor fields in a voxel-wise manner, existing methods are not very effective in visually expressing the differences. This thesis proposes a computationally efficient dissimilarity measure to quantify the pair-wise differences between

diffusion tensors in terms of tensor intrinsic properties. A novel checkerboard-style glyph is designed, which is able to visually convey the local differences between two tensor fields.

When analyzing multiple tensor fields as a whole, i.e., an ensemble of tensor fields, it is impossible to show all the information simultaneously. Therefore, appropriate statistics are needed to summarize essential information. Previous work is mainly based on the Euclidean mean tensor and the complex fourth-order covariance tensor. This type of analysis is performed in a general manner by not considering the intrinsic properties of the diffusion tensor. More importantly, the visual interpretation of the covariance tensor in particular is difficult. This thesis proposes a representative mean tensor and tensor ensemble variations based on tensor intrinsic properties. An overview + detail visual analysis framework is developed to facilitate the visual exploration of ensembles of tensor fields in the 3D physical and feature space.

In cases where two ensembles of tensor fields are compared (e.g., a control group and a patient group), the contradiction between the huge amount of information to be visualized and the limited number of available visual channels becomes much more severe. This thesis resolves this contradiction by carefully combining and extending the checkerboard-style glyph design and the overview + detail framework. A glyph representation is designed to support the comparison of the tensor ensemble summary statistics in the spatial domain. The final major contribution of this thesis is to integrate the level-of-detail concept into the glyph representation, which is able to progressively reveal more information as neuroscientists zoom in.

The validation of the visual analysis techniques proposed in this thesis is carried out in the form of user studies and case studies. The design of a user study in this context is a difficult problem since some domain-specific knowledge is necessary, which restricts the number of available participants. We abstract domain-specific questions into a domain-independent form whenever possible to include as many participants as possible. Case studies are performed by neuroscientists with real-world datasets to confirm already known facts and obtain new findings to generate hypothesis. The overall evaluation results demonstrate the usability of our visual analysis framework.

SAMENVATTING

Wetenschappelijke visualisatie van tensorvelden is uitdagend vanwege de complexe en meerdimensionale eigenschappen van tensor data. De visualisatie wordt zelfs nog moeilijker wanneer meerdere tensorvelden geanalyseerd worden en deze visualisaties staan nog aan het begin van hun ontwikkeling. Deze thesis richt zich op het bijdragen van visuele analysetechnieken voor meerdere 3D tensorvelden.

We focussen vooral op de visuele analyse van Diffusion Tensor Imaging (DTI) data. DTI is een vorm van magnetic resonance imaging (MRI), welke vaak gebruikt wordt in neurowetenschappen om de witte stof hersenen in vivo te onderzoeken. Een langere scantijd is nodig vergeleken met andere scan methodes. Het versnellen van de MRI-scans zou potentieel de toepasbaarheid van DTI kunnen verbeteren. Compressed sensing (CS) is een signaal reconstructie techniek die gebruikt wordt om MRI-scans te versnellen. De traditionele CS-methode probeert de globale kwaliteit van het gereconstrueerde beeld te optimaliseren. Helaas, in de praktijk is vooral de kwaliteit van lokale specifieke structuren belangrijk. Daarom onderzoeken we het gebruik van CS voor dit doelende en dragen we bij in deze onderzoeksrichting door de traditionele CS-reconstruictiemethode te laten focussen op de kwaliteit van lokale structuren.

Met DTI is het mogelijk om het diffusieprofiel van watermoleculen in iedere voxel te meten, deze wordt beïnvloed door de onderliggende vezelbanen van de witte stof en stleert deze als een symmetrische positieve definitie matrix van de tweede orde (de zogeheten diffusie tensor). Die diffusie tensor, wiskundig uitgedrukt als een 3×3 symmetrische positieve definitie matrix, kan worden onderverdeelt in schaal, vorm en oriëntatie, welke de intrinsieke eigenschappen van de matrix zijn. Het zijn deze intrinsieke eigenschappen die DTI een unieke belangrijke beeldvormende methode maken. Ieder van deze intrinsieke eigenschappen heeft een betekenisvolle biologische interpretatie van het onderliggende weefseleigenschappen. Deze thesis maakt veelvuldig gebruik van tensor intrinsieke eigenschappen voor de ontwikkeling van visuele analysetechnieken voor meerdere DTI-datasets zodat neurowetenschappers deze direct

kunnen interpreteren.

Wanneer twee tensorvelden vergeleken worden op voxelniveau zijn bestaande methoden minder effectief in het visueel uitdrukken van de verschillen. Deze thesis stelt een efficiënt berekenbare ongelijkheidsmaat voor om paarsgewijs de verschillen tussen de diffusie tensors te kwantificeren aan de hand van de intrinsieke tensoreigenschappen. Een nieuwe schaakbordstijl-glyph is ontworpen die in staat is om de lokale verschillen tussen twee tensorvelden te visualiseren.

Wanneer meerdere tensorvelden worden geanalyseerd als een geheel, bijvoorbeeld een ensemble van tensorvelden, is het onmogelijk om alle informatie gelijktijdig te laten zien. Daarom zijn geschikte statistieken nodig die de essentiële informatie samenvatten. Eerder werk is voornamelijk gebaseerd op de Euclidische gemiddelde tensor en de complexe covariantie tensor van de vierde orde. Dit soort analyse wordt uitgevoerd in een generieke wijze zonder de intrinsieke eigenschappen van de diffusie tensor mee te nemen. Bovendien is de visuele interpretatie van de covariantie tensor in het bijzonder moeilijk. Deze thesis stelt een representatieve gemiddelde tensor voor en tensor ensemble variaties gebaseerd op de tensor intrinsieke eigenschappen. Een overzicht + detail visuele analyse framework is ontworpen voor de visuele verkenning van ensembles van tensorvelden in de ruimtelijke en data ruimte.

In het geval waar twee ensembles van tensorvelden vergeleken moeten worden, bijvoorbeeld een controle groep en een patiëntengroep, wordt de tegenstelling tussen de enorme hoeveelheid informatie die gevisualiseerd moet worden en het aantal visuele kanalen erger. Deze thesis lost deze tegenstelling op door voorzichtig de schaakbord-glyph en het overzicht + detail framework te combineren en uit te breiden. Een glyphrepresentatie is ontworpen om de vergelijking te ondersteunen van de tensor ensemble samenvattingsstatistieken in het ruimtelijke domein. De laatste bijdrage van deze thesis is de integratie van het level-of-detail concept in de glyphrepresentatie. Deze representatie maakt het mogelijk om progressief meer informatie te laten zien wanneer de neurowetenschapper inzoomt.

De validatie van de visuele analysetechnieken die in deze thesis worden geïntroduceerd is uitgevoerd door middel van gebruikersonderzoeken en gevalsanalyses. Het ontwerp van een gebruikersonderzoek is een lastig probleem omdat bepaalde vakken-nis nodig is, hierdoor is het aantal gebruikers beperkt. We abstraheren vragen waarvoor vakkennis nodig in vragen die geen specifieke vakkennis vereisen zodat we zoveel mo-

gelijk deelnemers konden includeren. Gevalsanalyses zijn uitgevoerd door neurowetenschappers met reële datasets om alreeds bekende feiten te bevestigen en om nieuwe hypothesen te ontwikkelen. De uitkomst van de totale evaluatie demonstreert de bruikbaarheid van ons visuele analyse framework.

1

INTRODUCTION

*The art of war is of vital importance to the State.
It is a matter of life and death, a road either to safety or to ruin.
Hence it is a subject of inquiry which can on no account be neglected.*

Sun Tzu

兵者，国之大事，死生之地，存亡之道，不可不察也。

孙武

1.1. MOTIVATION

The brain is the most important and complex organ in the human body. The cerebrum, as the largest and most developed part of the brain, is roughly made up of two types of tissues, the gray matter and the white matter [1]. A modern imaging technique, called Diffusion Weighted Magnetic Resonance Imaging (DW-MRI) [2, 3], provides a unique way to noninvasively investigate the fibrous structure of white matter in the human brain in vivo by measuring the diffusion profile of water molecules. DW-MRI, as the name suggests, is a specific kind of MRI technique. MRI acquisition is a time-consuming process, and DW-MRI is even more time-consuming since it requires additional scanning sequences. This disadvantage restricts the applicability of DW-MRI. Compressed sensing (CS) [4] is a signal processing technique that enables the reconstruction of MR images with fewer samples in k-space, which is termed as CS MRI. Hence, CS MRI is able to produce high-quality images with relatively short scanning time. However, in practice, researchers are more interested in the quality of relevant local structures, instead of the whole image. In this thesis, we adapt the traditional CS MRI method to take local information of interest into consideration.

In regular MR images the fibrous structure of different white matters is not distinguishable. This is where DW-MRI comes in. DW-MRI enables the generated images to sensitize to different white matter structures by employing a physical phenomenon, i.e., the Brownian motion or diffusion of water molecules. The diffusion behavior of water molecules shows directional dependence (i.e., anisotropy) upon the underlying white matter structure. Specifically, water diffuses more freely (i.e., faster diffusion) along the direction of nerve fiber tracts of white matter than across it. DW-MRI is capable of measuring the extent of water diffusion in many directions for each volume element (voxel) in the brain, which are then used to deduce the underlying fibrous structure of white matter.

It is hard to interpret the generated DW-MRI images without proper data processing. It yields little insight, for example, by directly looking at hundreds of slices of the DW-MRI volumetric images. As a result, one needs to summarize these generated DW-MRI images to make them comprehensible. A widely employed assumption about the water diffusion is that it follows an anisotropic Gaussian distribution. In this scenario the local diffusion behavior can be mathematically modeled with a symmetric second-order tensor \mathbf{D} , which is expressed as a 3×3 symmetric matrix. Consequently, this

specific modeling technique of DW-MRI images is called Diffusion Tensor Imaging (DTI) [5]. Other more advanced modeling techniques such as High Angular Resolution Diffusion Imaging (HARDI) [6] require the acquisition of many more DW-MRI images, which is a time-consuming process. This makes DTI extensively used in clinic and research. This thesis focuses on dealing with DTI datasets, which are 3D tensor fields by nature.

Scientific visualization relies on the high-bandwidth channel of the human visual system to our brains [7]. Its goal is to turn large and complex data into a visual representation so that human experts can *see* the data better. Therefore, scientific visualization is the technique that human experts can resort to in order to obtain a better understanding of what information is expressed in the tensor field. There has been considerable progress in this research direction in last decades (cf. Section 3.4), ranging from the tensor glyphs [8] to fiber tracking [9].

Nowadays, DTI-based analysis is moving beyond a single dataset to two or multiple co-registered datasets. For example, it is necessary to compare two DTI datasets in order to explore the effects of different acquisition parameters [10], or evaluate the registration quality [11]. Figure 1.1a shows that two tensor fields are overlaid with each other, which is a common strategy for voxel-wise comparison. It is also necessary to summarize a group of DTI datasets, normally called cohorts or ensembles, in order to locate variations between subjects that might be caused by natural variations [12] or uncertainty [13]. In neuroscientific studies, it is quite common to compare two groups of DTI datasets in order to investigate the influence caused by pathology [14–16], gender [17], different analysis methods [18], or other factors [19]. Figure 1.1b shows the level-of-detail visualization designed to analyze an ensemble of diffusion tensor fields. The large number of subjects involved in the analysis together with the multivariate nature of diffusion tensors impose challenges on both the analysis and the visualization. Due to the complexity of multivariate statistical analysis, most of the previous approaches are based on derived scalar quantities, such as fractional anisotropy [20] (FA), rather than the multivariate diffusion tensor itself. Well-established visualizations for scalar data (e.g., slice views or line plots) can be used to show the results. The primary reason behind this type of visual analysis approaches is *simplicity*. However, it provides merely one facet of the full tensor information. This means neuroscientists run the risk of missing relevant information. Figure 1.1c shows the most commonly

used visualization, i.e., the slice views, for inter-group comparison. Certain statistical measures, which are color-coded, are overlaid on the FA background images. When analyzing multiple tensor fields, there are few available methods that handle the diffusion tensor in its entirety [19, 21–23]. These previous works assume that diffusion tensors follow a multivariate normal distribution. However, the distribution assumption is not always valid, and more importantly, the summary statistics are difficult to interpret.




Figure 1.1: (a) Overlaying two tensor glyphs is a common strategy to compare two tensor fields in a voxel-wise manner. Image courtesy [24]. (b) shows a relevant work on visualizing a single ensemble of tensor fields at multiple levels of detail. Image courtesy [23]. (c) shows the results of the commonly used t-statistic quantitative analysis and the slice-based visualization for DTI group comparison. Image courtesy [25].

1.2. AIM

In this thesis we investigate the adaption of the traditional CS MRI method to incorporate local information into the reconstruction. However, the main aim of this thesis is to design *understandable* visual analysis techniques for ensembles of diffusion tensor fields while keeping the *full* tensor information. The main aim can be further subdivided into three tasks.

- T1** design a comparative visualization technique so that the pair-wise differences between *two* DTI datasets are presented in a more comprehensive way.
- T2** aggregate *a single ensemble* of DTI datasets so that the summary statistics are easier to understand and design an ensemble visualization technique to facilitate the exploration and analysis.
- T3** develop an effective visualization to convey the huge amount of information with a limited number of visual channels when comparing *two ensembles* of DTI

datasets, based on the solutions to the first two tasks.

We make these three tasks more concrete by using the DTI group study as a concrete example (see Figure 1.2). The first task **T1** corresponds to the scenario where two subjects need to be compared, indicated by the smallest rectangle in Figure 1.2, in order to investigate their *differences*. The key challenges are the quantification of tensor differences that make sense from the perspective of DTI and the visual encoding of such differences. Neuroscientists are not only interested in comparing two specific subjects, but also in analyzing an ensemble of subjects as a whole, indicated by the medium rectangle in Figure 1.2, in order to investigate the *variations*. This scenario corresponds to the second task **T2**. One big challenge is to define suitable statistical measures to quantify ensemble variation. The other is the visual design, which should not only provide understandable visual representations of the summary statistics but also show the original tensors that give rise to the summary statistics. The solutions to **T1** and **T2** serve as inspirations to accomplish the third task **T3**, which corresponds to the inter-group comparison as indicated by the biggest rectangle in Figure 1.2. The major challenge is to resolve the contradiction between the huge amount of information to be visualized and the limited number of available visual channels.




Figure 1.2: The three visualization tasks to accomplish in this thesis are illustrated in the context of DTI group study. Each row indicates an ensemble of subjects. Three gradually expanding rectangles correspond to task **T1**, **T2**, and **T3**, respectively. The size and lightness of the rectangles quantitatively reflect the number of subjects to be analyzed. Image adapted from © South Park.

After the design process, we need to evaluate the developed visualization techniques. In general, the validation of a visual design is a tricky problem [7] because

there are so many variables involved. The validation in this thesis is carried out in the form of user study and case study. The purpose of a user study is to evaluate the effectiveness of the visual encoding. However, the diffusion tensor, the main object to be visualized, is a relatively complex concept that is not familiar to computer scientists or even visualization researchers. Consequently, the basic understanding of the diffusion tensor is compulsory, which severely restricts the number of available participants. The purpose of the case studies is not only to evaluate the usability of our visual analysis framework for the task that we aim at, but also to avoid the creation of any artifacts that do not originate from the data itself. The case studies are performed mainly by the target domain experts – neuroscientists.

1.3. OUTLINE

This dissertation is organized as follows.

Chapter 2 first introduces the basic principles of MRI in Section 2.1, which is the basis of DW-MRI, and compressed sensing Section 2.2, an acceleration technique for MRI acquisition. Our contribution is presented in Section 2.3 which incorporates the local information of interest into MR image reconstruction, and apply to angiography images as a proof-of-concept.

Chapter 3 contains three types of background information. Section 3.1 covers the brain anatomy in a concise way. Section 3.2 presents the fundamental principles of DTI, which is the imaging technique to characterize the macrostructures of brain white matter. Section 3.3 gives the mathematical background of diffusion tensor data. Section 3.4 summarizes state-of-art visualization techniques related to this thesis, including glyph-based visualization, comparative visualization, and ensemble visualization.

Chapter 4 presents a novel glyph design for visually comparing two DTI datasets in a voxel-wise manner. A thorough task analysis is performed to determine what kind of information to compare, and a computationally efficient measure is proposed to quantify the differences between two diffusion tensors in terms of tensor scale, shape, and orientation. Our difference measure can effectively handle the coupling of tensor shape and orientation. The novel glyph design, inspired by the checkerboard-style comparison, is described in detail. The effectiveness of the new glyph design is demonstrated with an informal user study and a case study.

When the total number of DTI datasets to compare goes beyond two to multiple,

it is necessary to resort to statistical tools to summarize them, again in a voxel-wise manner. In Chapter 5, the pair-wise difference measure between two diffusion tensors proposed in Section 4.2.2 is generalized to quantify the variation in an ensemble of diffusion tensors, which are more understandable than the fourth-order covariance tensor used in previous work. More importantly, Chapter 5 describes the overview + detail visual analysis strategy for visualizing a single ensemble of DTI datasets. The glyph-based overview is based on the extension of the superquadric tensor glyph visualization [8] by visually encoding the variation information. The detail views are designed to facilitate the exploration of the original tensors in terms of tensor scale, shape, and orientation. Again, an informal user study and a case study are carried out to justify the effectiveness of the novel two-level visual analysis approach.

Chapter 6 moves forward to the comparison of two ensembles of DTI datasets by extending the overview + detail design in Chapter 5. A new overview glyph, the combination of the checkerboard-style tensor difference glyph proposed in Chapter 4 and the tensor ensemble summary glyph proposed in Chapter 5, aims at facilitating the visual comparison of summary statistics in the spatial domain. More importantly, the concept of levels of detail (LODs) visualization is incorporated into the overview glyph to *progressively* reveal more information as users zoom in. Furthermore, special care is taken into consideration in order to suppress visual clutter when comparing the tensor ensembles in the detail view. This chapter ends with a case study where an HIV positive group and a control group are compared to illustrate the potential of our prototype in helping neuroscientists to obtain new findings.

We conclude this dissertation in Chapter 7 where the main contributions of this dissertation are summarized and possible directions for future work are suggested.

REFERENCES

- [1] S. Standring, *Gray's anatomy: the anatomical basis of clinical practice* (Elsevier, 2015).
- [2] D. L. Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet, *MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders*, *Radiology* **161**, 401 (1986).
- [3] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. D. Chiro, *Diffusion tensor*

- MR imaging of the human brain*, Radiology **201**, 637 (1996).
- [4] M. Lustig, D. Donoho, and J. M. Pauly, *Sparse MRI: The application of compressed sensing for rapid MR imaging*, Magnetic Resonance in Medicine **58**, 1182 (2007).
- [5] P. Basser, J. Mattiello, and D. Lebihan, *Estimation of the effective self-diffusion tensor from the nmr spin echo*, Journal of Magnetic Resonance, Series B **103**, 247 (1994).
- [6] D. S. Tuch, T. G. Reese, M. R. Wiegell, N. Makris, J. W. Belliveau, and V. J. Wedeen, *High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity*, Magnetic Resonance in Medicine **48**, 577 (2002).
- [7] T. Munzner, *Visualization Analysis and Design* (A K Peters/CRC Press, 2014) Chap. 1, pp. 6–7.
- [8] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004) pp. 147–154.
- [9] S. Mori and P. B. Barker, *Diffusion magnetic resonance imaging: Its principle and applications*, The Anatomical Record **257**, 102 (1999).
- [10] K. M. Hasan, D. L. Parker, and A. L. Alexander, *Comparison of gradient encoding schemes for diffusion-tensor MRI*, Journal of Magnetic Resonance Imaging **13**, 769 (2001).
- [11] H. Zhang, P. A. Yushkevich, and J. C. Gee, *Registration of diffusion tensor images*, in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 1 (2004) pp. 842–847.
- [12] S. M. Resnick, D. L. Pham, M. A. Kraut, A. B. Zonderman, and C. Davatzikos, *Longitudinal magnetic resonance imaging studies of older adults: A shrinking brain*, The Journal of Neuroscience **23**, 3295 (2003).
- [13] D. K. Jones, *Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI*, Magnetic Resonance in Medicine **49**, 7 (2003).

- [14] A. Schwartzman, R. F. Dougherty, and J. E. Taylor, *Cross-subject comparison of principal diffusion direction maps*, *Magnetic Resonance in Medicine* **53**, 1423 (2005).
- [15] Y. Wu, P. Storey, B. Cohen, L. Epstein, R. Edelman, and A. Ragin, *Diffusion alterations in corpus callosum of patients with HIV*, *American journal of neuro-radiology* **27**, 656 (2006).
- [16] G. Douaud, S. Jbabdi, T. E. Behrens, R. A. Menke, A. Gass, A. U. Monsch, A. Rao, B. Whitcher, G. Kindlmann, P. M. Matthews, and S. Smith, *DTI measures in crossing-fibre areas: Increased diffusion anisotropy reveals early white matter alteration in MCI and mild Alzheimer's disease*, *NeuroImage* **55**, 880 (2011).
- [17] K. Menzler, M. Belke, E. Wehrmann, K. Krakow, U. Lengler, A. Jansen, H. Hamer, W. Oertel, F. Rosenow, and S. Knake, *Men and women are different: Diffusion tensor imaging reveals sexual dimorphism in the microstructure of the thalamus, corpus callosum and cingulum*, *NeuroImage* **54**, 2557 (2011).
- [18] L. Snook, C. Plewes, and C. Beaulieu, *Voxel based versus region of interest analysis in diffusion tensor imaging of neurodevelopment*, *NeuroImage* **34**, 243 (2007).
- [19] B. Whitcher, J. J. Wisco, N. Hadjikhani, and D. S. Tuch, *Statistical group comparison of diffusion tensors via multivariate hypothesis testing*, *Magnetic Resonance in Medicine* **57**, 1065 (2007).
- [20] P. J. Basser and C. Pierpaoli, *Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI*, *Journal of Magnetic Resonance* **111**, 209 (1996).
- [21] P. J. Basser and S. Pajevic, *A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI*, *IEEE Transactions on Medical Imaging* **22**, 785 (2003).
- [22] P. Basser and S. Pajevic, *Spectral decomposition of a 4th-order covariance tensor: applications to diffusion tensor MRI*, *Signal Processing* **87**, 220 (2007).

- [23] A. Abbasloo, V. Wiens, M. Hermann, and T. Schultz, *Visualizing tensor normal distributions at multiple levels of detail*, IEEE Transactions on Visualization and Computer Graphics **22**, 975 (2016).
- [24] Ø. Bergmann, G. Kindlmann, A. Lundervold, and C.-F. Westin, *Diffusion k-tensor estimation from Q-ball imaging using discretized principal axes*, in *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2006) pp. 268–275.
- [25] T. Su, M. W. Caan, F. W. Wit, J. Schouten, G. J. Geurtsen, J. H. Cole, D. J. Sharp, F. M. Vos, M. Prins, P. Portegies, *et al.*, *White matter structure alterations in hiv-1-infected men with sustained suppression of viraemia on treatment*, AIDS **30**, 311 (2016).

2

MAGNETIC RESONANCE IMAGING AND USER-GUIDED COMPRESSED SENSING

*Let your rapidity be that of the wind,
your compactness that of the forest.*

Sun Tzu

故其疾如风，其徐如林。

孙武

This chapter is partially based on the following publication: “User-guided compressed sensing for magnetic resonance angiography”, **Changong Zhang**, Martijn van de Giessen, Elmar Eisemann, and Anna Vilanova. In *Proceedings of 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2416–2419, 2014 [1].

Magnetic resonance imaging (MRI) is the imaging modality on which DW-MRI is based. Therefore, the presentation of MRI physics in Section 2.1 is needed to understand the principles of DW-MRI in Section 3.2. This chapter is not focused on studying MRI from the perspective of physics so the mechanism of MRI is described in a simple and concise manner. Compared to other imaging modalities like computed tomography (CT), one disadvantage of MRI is that it requires longer scan time. A signal processing technique named compressed sensing (CS) [2, 3], which can efficiently sample and reconstruct the original signal by exploiting its sparsity, is a candidate to solve this problem [4]. This chapter is focused specifically on the application of CS to MRI, rather than the abstract theory of CS. Therefore, a simplified description of CS MRI is presented in Section 2.2 to give a flavor about how CS integrates into the pipeline of MRI acquisition. The traditional CS aims at optimizing the overall quality of the reconstructed images. However, in practice, researchers are normally interested in the quality of specific local areas. Our contribution is presented in Section 2.3, which takes user-specified local information into consideration during image reconstruction.

2.1. MRI PRINCIPLES

Magnetic Resonance Imaging (MRI) is an imaging technique to non-invasively produce tomographic images of the underlying anatomy by exploiting the physical phenomenon called nuclear magnetic resonance (NMR).

The adult human body is around 60% water. The single proton of the hydrogen atom in the water molecule is positively charged, and has an angular momentum (i.e., spin). Each proton then generates a microscopical magnetic field. All of these spins are randomly orientated, and cancel each other out macroscopically. However, in the presence of an external magnetic field \mathbf{B}_0 , these spins align with it, either parallel (i.e., low energy state) or anti-parallel (i.e., high energy state), with a preference for the low energy state. These protons do not simply point with or against the direction of \mathbf{B}_0 , they actually rotate or precess like a spinning top at the Larmor frequency $\omega_L = \gamma B_0$, where γ is the gyromagnetic ratio and B_0 is the strength of the external magnetic field \mathbf{B}_0 . Each proton precesses with a random phase, canceling each other out in the transverse plane, and a net longitudinal magnetization m_z is generated, which is the equilibrium state of the nuclear spin systems.

A radio-frequency (RF) electromagnetic wave at frequency ω_L can flip some pro-

tons from the low energy state to high energy state by energy transmission. Additionally, the RF pulse will synchronize the protons to coherent phases. As a result, the longitudinal magnetization m_z decreases and a rotating transverse magnetization m_{xy} is created in a plane orthogonal to the direction of \mathbf{B}_0 , which in turn triggers a new electromagnetic signal that can be measured by the coils within the MRI scanner. In this case, it is the transverse magnetization m_{xy} that produces the measurable signal. The longitudinal magnetization m_z can be flipped into the transverse plane by a 90° RF pulse, and thus, it is the longitudinal magnetization m_z that produces the measurable signal in this case. m_z recovers to the equilibrium state while m_{xy} decays in an exponential way with respect to time, which can be characterized by the transverse relaxation time T_2 and the longitudinal relaxation time T_1 , respectively. The term relaxation means how signals change with respect to time. Therefore, the contrast between different tissues depends on the proton density P_D , the transverse relaxation time T_2 , and the longitudinal relaxation time T_1 . Interested readers are referred to textbooks such as Prince and Links [5] for more details.

In order to locate the underlying MR quantity $m(x, y)$ for a given voxel in a 2D slice, two additional magnetic gradients G_x and G_y are applied in the x and y direction, respectively, apart from the main magnetic field in the z direction. G_y is used to linearly encode y position to the phase of the magnetization, $\theta(x, y) = G_y y \tau$. G_y is on for a short time τ , and switched off before the application of G_x , which is used to encode the x position into the Larmor frequency. The Larmor frequency is a linear function of x as $\omega(x, y) = \gamma(B_0 + G_x x)$. The received signal is an integral over the 2D slice, given by

$$s(t) = \iint m(x, y) e^{-j(\omega t + \theta)} dx dy = e^{-j\gamma B_0 t} \iint m(x, y) e^{-j\gamma(G_x x t + G_y y \tau)} dx dy,$$

where $m(x, y)$ is proportional to either the transverse magnetization m_{xy} or the longitudinal magnetization m_z . The exponential part $e^{-j(\bullet)}$ is introduced because the used magnetization, expressed as one *complex* quantity, is *rotating* in the xy plane with frequency ω and initial phase θ .

The received signal will be demodulated to yield the baseband signal

$$s_b(t) = e^{+j\gamma B_0 t} s(t) = \iint m(x, y) e^{-j\gamma(G_x x t + G_y y \tau)} dx dy. \quad (2.1)$$

Equation 2.1 can be interpreted as a 2D Fourier transform of $m(x, y)$ with spatial frequencies $u = \gamma G_x t$ and $v = \gamma G_y \tau$ in the x and y direction, respectively. According

to the MRI terminology, this Fourier frequency space is commonly referred to as the *k-space*. In principle, MRI measurements can be viewed as scanning the *k-space*, and the image reconstruction algorithm is the inverse Fourier transform. Again, interested readers are referred to textbooks such as Prince and Links [5] for more details.

2.2. COMPRESSED SENSING MRI

MRI acquisition means a full scanning of the *k-space* by tuning the frequency encoding gradient G_x , phase encoding gradient G_y and its application time τ . The *k-space* can be traversed in several ways, among which the Cartesian trajectory (i.e., row-by-row) is the most commonly used one. Normally the size of the *k-space* is around 512×512 , which indicates that the filling of *k-space* is a time consuming process. On the other hand, we know that MR images have sparse representations in some transform domains. For example, brain MR images have sparse representation in terms of wavelet coefficients, i.e., most of the wavelet coefficients are zero or negligible. Therefore, a natural question to ask is: is it possible to directly measure the *compressed* coefficients, instead of the original complete *k-space*? The theoretical answer to this question is referred to as *compressed sensing* (CS) [2, 3]. Lustig et al. [4] apply the CS technique [2, 3] to the reconstruction of MR image, called CS MRI for brevity, which enables the fast reconstruction of images with fewer measurements in *k-space* than required by the inverse Fourier transform. Here we briefly introduce the basic principles behind CS MRI. Interested readers are referred to [4, 6] for detailed information.

There are three requirements for the application of CS. (a) The image to be reconstructed should have a sparse representation in a known transform domain, and thus be compressible; (b) The aliasing artifacts, caused by *k-space* undersampling, should have noise-like (i.e., incoherent) appearance; (c) The reconstruction algorithm should enforce both the sparse representation and data fidelity with respect to the measurements. Figure 2.1 illustrates the three requirements and the way that CS MRI works. Direct inverse Fourier transform of the partially sampled *k-space* measurements causes artifacts. However, by taking advantage of the transform domain, the reconstruction will achieve similar image quality. The reconstruction process is given as an optimiza-

tion problem:

$$\begin{aligned} \text{minimize} \quad & \|\Psi m\|_1 \\ \text{s.t.} \quad & \|\mathcal{F}_u m - y\|_2 < \epsilon, \end{aligned} \quad (2.2)$$

where m is the image to be reconstructed, Ψ is the operator that transforms m into the sparse domain, \mathcal{F}_u is the undersampled Fourier transform, and y is the actual measurements. Minimizing the ℓ_1 norm can enforce the sparsity while ϵ controls data fidelity.




Figure 2.1: Illustration of the domains and operators of CS MRI as well as the three requirements. The image to be reconstructed should have a sparse representation in a transform domain. The partially sampled k-space measurements (indicated by the red dots) should result in incoherent artifacts. The reconstruction method (indicated by the big curved gray arrow) should enforce the sparsity and be consistent with the partial measurements. \mathcal{F} , \mathcal{F}_u , and Ψ represent the Fourier transform, undersampled Fourier transform, and the sparse transform, respectively. * means the inverse of an operator. Image courtesy [4].

CS MRI actually transfers parts of the cost of MR image generation from hardware (i.e., k-space sampling) to software (i.e., image reconstruction). Therefore, in this way, CS MRI is able to reconstruct MR images at a higher speed without significantly reducing the image quality.

2.3. USER-GUIDED COMPRESSED SENSING MRI

We observe that Equation 2.2 used in CS MRI aims at reconstructing the image as a whole entity. However, in practice, medical researchers are more interested in the quality of specific structures, instead of the whole image. Therefore, we adapt

Equation 2.2 to incorporate the information of user-specified region of interest (ROI), and apply it to angiography as a proof-of-concept.

2

In vessel diagnosis with angiograms, some low-contrast vessels are essential to the reliable assess of pathological abnormalities (e.g., stenosis). Even though good reconstruction results are presented by Lustig et al. [4] for angiograms, which have sparse representations under finite-difference operation, Milles et al. [7] still observe the fading of low-contrast vessels in CS reconstructed angiograms. The fading or breaking of vessels could cause false positives in such diagnostic procedures. In this section, we propose a method aiming at the good reconstruction of specific low-contrast vessels by integrating the local information into the optimization function Equation 2.2. The fading is caused by the fact that ℓ_1 reconstruction employed in CS MRI shrinks the magnitude of the reconstructed sparse coefficients. This results in the reduction of image contrast especially at high undersampling ratios [4]. Then the desired boundaries between the low-contrast vessels and the background are more likely to fade or even vanish. To avoid these artifacts, we propose to add weights to the ℓ_1 reconstruction based on local information. Candes et al. [8] have already demonstrated the performance gain via weighted ℓ_1 minimization in areas of sparse signal recovery and image processing. Chang and Ji [9] extend the weighted ℓ_1 minimization to reconstruct multichannel in-vivo MRI data. Since ℓ_1 reconstruction works well for high-contrast vessels, we weight the ℓ_1 reconstruction for specific vessels to maintain their boundary information.

There are several attempts to improve the quality of certain tissues (i.e., local quality) of the object being imaged, which could be viewed as the ROI, rather than the global reconstruction quality. Sharma and Nayak [10] manage to increase the image contrast within the ROI by imposing the sparsity constraint only outside ROI. However, the ROI cannot be the entire image. Otherwise their method would turn into an ill-posed least-squares problem. Oh and Lee [11] derive the visual weight by incorporating ROI and perceptual characteristics of the human visual system. But their method could reduce to an ill-posed problem if the ROI is really small such as in the case of angiography. In our work, there are no limitations about the size of the ROI which could either be the entire image or empty.

2.3.1. METHODOLOGY

The pipeline of our method is illustrated in Figure 2.2. The input includes the partially measured k-space data and an initial image data. The initial image is used to define the ROI by the users (e.g., radiologists) via semi-automatic or interactive segmentation [12]. The ROI contains low-contrast vessels which cannot be well preserved via traditional CS. The ROI does not need to be precisely defined, and a rough estimation is sufficient. The initial image can either be a zero-filling reverse Fourier reconstructed or a traditional CS reconstructed image. In this chapter, we choose to use the traditional CS reconstructed images but with much fewer iterations. The weights are constructed based on the gradient information within the ROI. The final images are reconstructed via weights-incorporated ℓ_1 reconstruction. In the following we will describe construction of the weights and weighted CS reconstruction in more detail.

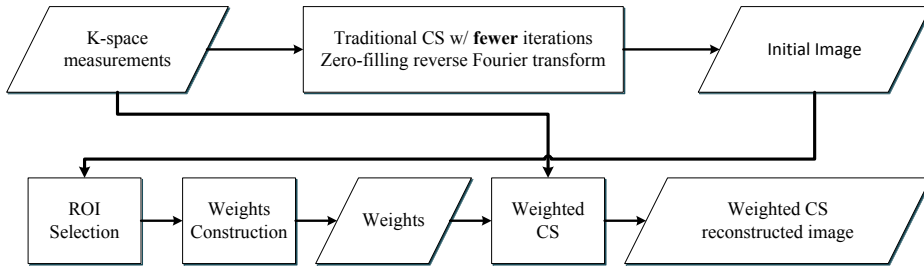


Figure 2.2: The weighted CS reconstruction pipeline. The partially sampled k-space measurements are used to generate an initial image where users can specify ROI. The initial reconstruction can be performed with either the zero-filling reverse Fourier transform or the traditional CS but with much fewer iterations. Gradient-based weights are derived based on the specified ROI, which are then incorporated into CS to reconstruct the final image.

ROI-BASED WEIGHTS CONSTRUCTION

The primary information of angiograms consists of blood vessels from which we derive clinically-related parameters such as the diameters and length of vessels. The boundaries are of special relevance for these parameters. They can be estimated by calculating the gradient in the initial image using central differences. We do not choose a smoothed derivative filter, such as the Gaussian derivative filter, in order to capture low scale vessel features. User-defined ROI is applied to mask the gradients, which are normalized to the range between 0 and 1 by the maximum gradient magnitude. The smoothed weights are constructed by reversing the normalized gradients:

$$M_{i,j} = \|\nabla I_{i,j}\| \cdot ROI_{i,j}$$

$$W_{i,j} = 1 - \frac{M_{i,j}}{\max(M_{i,j}) + \epsilon}, \quad (2.3)$$

Here I is the initial image, ∇ is gradient operator, ROI is the binary mask, M is the gradient magnitude within the ROI, ϵ is a very small numerical value preventing the division by 0, and W is the final weights.

Figure 2.3a shows the fully sampled 3D angiography which will be used as the ground truth. Figure 2.3b is the color-coded initial image reconstructed by traditional CS. Both are under the same rendering settings, and we can clearly observe the breaking and fading of vessels marked within the three white polygons. These regions containing this vessel are selected as ROI. Figure 2.3c shows one slice of the generated weights overlapped with the CS reconstructed data. The weights are color-coded by perceptually linear yellow-to-blue colormap. Yellow represents low weights while blue the high weights. Transparency represents the highest weights (i.e., 1.0) for areas outside of the ROI.




Figure 2.3: ROI-based weights construction. (a) fully sampled 3D angiogram; (b) CS reconstructed angiogram with sampling ratio of 10%; (c) a close-up of the generated weights overlaid with CS reconstructed angiogram.

CS RECONSTRUCTION USING ROI WEIGHT

Lustig et al. [4] present the reconstruction of angiograms using finite-differences as the sparsifying transform, which is referred as total-variation (TV) minimization. Besides, unlike Fourier transform, finite-difference provides spatially local information and has low computational complexity compared to other sparse transforms such as wavelet. We also employ TV minimization to reconstruct the original angiograms. The CS reconstruction are performed by solving the following optimization problem:

$$\begin{aligned} \operatorname{argmin}_m \|F_u m - y\|_2 + \lambda \|m\|_{TV} \\ \|m\|_{TV} = \sum_{i,j} \|(Dm)_{i,j}\|_1, \end{aligned} \quad (2.4)$$

where m is the desired image, F_u is randomly-sampled Fourier operator, y are the k-space measurements, D represents the forward differences operator and λ is the regularization parameter that determines the trade-off between data consistency and TV regularization. The ℓ_2 norm is defined as $\|x\|_2 = (\sum_i |x_i|^2)^{1/2}$ while the ℓ_1 norm as $\|x\|_1 = \sum_i |x_i|$.

The TV term penalizes intensity variations. This penalization, however, shrinks the magnitude of the transform coefficients and causes a reduction of image contrast. The desired boundaries are likely to fade or even vanish. Therefore, we propose to add weights to the TV minimization. Voxels that are inside the ROI with relatively high gradients are candidates for boundaries. The weights should be relatively low for these voxels and vice versa. This means, within the ROI, intensity variations for boundary candidates can be tolerated. The proposed optimization problem is:

$$\begin{aligned} \operatorname{argmin}_m \|F_u m - y\|_2 + \lambda \|m\|_{wTV} \\ \|m\|_{wTV} = \sum_{i,j} W_{i,j} \|(Dm)_{i,j}\|_1, \end{aligned} \quad (2.5)$$

where $W_{i,j}$ is the derived smooth voxel-wise weights in the range of $[0, 1]$.

2.3.2. RESULTS AND DISCUSSION

RESULTS

We simulated k-space data by computing the Fourier transform of a slab (56 slices) of a high resolution 3DFT TOF angiogram of $0.23 \text{ mm} \times 0.23 \text{ mm} \times 0.35 \text{ mm}$. We also used a true fully sampled k-space data of a slab (32 slices) of a low resolution DFT TOF angiogram of $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$. From the full k-space data, five under-sampled data sets with corresponding sampling ratios of 10%, 15%, 20%, 25% and 30% were reconstructed.

We compared our methods qualitatively and quantitatively with the method proposed by Lustig et al. [4] (traditional CS) and the fully sampled data set. The normalized mean squared errors (NMSE) is used to evaluate the performance within the ROI.

Maximum Intensity Projections (MIP) are frequently used by clinical users for the visualization of vascular structures. Often a threshold is used to eliminate noisy background and the visualization of low intensities. For small vessels, if their intensities are low in certain regions, this threshold might generate breaks and discontinuities. We used the masks resulting from several thresholds to evaluate the reconstruction results and compared them with the ground truth. Dice Coefficients (DC) [13] are thus used as an extra metric to compare the reconstruction results.

DISCUSSION

Table 2.1: ROI NMSE of the high resolution TOF angiogram with various sampling ratios

	10%	15%	20%	25%	30%
traditional CS	0.0471	0.0315	0.0241	0.0185	0.0140
ROI-based CS	0.0456	0.0288	0.0208	0.0156	0.0118




Figure 2.4: ROI-based Dice Coefficient comparison. Left: DC variations with different thresholds under the same sampling ratios (20%); Right: DC variations under different sampling ratios at fixed threshold (10%).

As shown in Table 2.1, the ROI-based NMSE of our method are better than that of the traditional CS under each of the sampling ratios for the high resolution angiogram. It means that the vessels within ROI are of better reconstruction quality. Results of the sampling ratios on DC are shown in Figure 2.4 (right) at threshold 10% of the maximum intensity while Figure 2.4 (left) shows DC when varying the threshold with constant sampling ratio of 20%. The DC derived from our method are also better especially at low sampling ratios. With low thresholds, both show high values because the

background voxels are segmented as vessels which are shown in Figure 2.5a. As increasing the thresholds, more background voxels are removed. Both show a decreasing trend when raising the thresholds. But the decline rate is slower for our method. The above evaluation results demonstrate that our method can better maintain the intensity connectivity of vessels as shown in Figure 2.5b.

Table 2.2: ROI NMSE of the low resolution TOF angiogram with various sampling ratios

	10%	15%	20%	25%	30%
traditional CS	0.1628	0.0870	0.0506	0.0366	0.0280
ROI-based CS	0.0999	0.0562	0.0362	0.0256	0.0207


The ROI-based NMSE for the low resolution angiogram are shown in Table 2.2. As shown in Figure 2.6 (left), the vessels contained in the white polygons (i.e., ROI) are marked as the specific reconstruction target. From the close-up views in Figure 2.6 (middle and right), it is clear that our method can maintain the intensity connectivity better in most areas. The reason why we do not include DC in this experiment is that even for the fully sampled data set, the vessels within the ROI cannot be segmented merely based on the thresholds.

Our method does not have limitations on the size of the specified ROI. If no area is assigned as the ROI, the optimization function coincides with traditional CS. If the ROI is the entire image, our reconstruction function just tolerates the intensity variations for voxels with globally large gradient magnitudes rather than reduce to an ill-posed problem.


The difference of overall NMSE between our method and the traditional CS are small. This is due the fact that the vessels within ROI normally occupy a relatively small part of total voxels.

FUTURE WORK

There are still rooms for improvements. We construct the weights from the initial image and keep it constant during the reconstruction. Weights can be adaptively constructed and tuned after each reconstruction iteration. Furthermore, a more extensive evaluation with more datasets and clinically relevant parameters is necessary to show the real applicability of the proposed method in clinic.



(a) MIP renderings with threshold at 6%



(b) MIP renderings with threshold at 10%

Figure 2.5: MIP renderings of the original high-resolution angiogram (left), the traditional CS (middle), our method (right). Sampling ratio is 20%.




Figure 2.6: MIP renderings for the low-resolution angiography. From left to right: the overview of traditional CS, the close-up views of the original image, traditional CS, and our method. Sampling ratio is 20%.

2.4. CONCLUSION

In this chapter, we give an overview of MRI technique, which is the basis of DW-MRI, and one acceleration technique – Compressed Sensing (CS). Furthermore, we propose a weighted CS reconstruction method and apply it to preserve the low-contrast vessels in angiography. Preliminary results show that our method using both simulated and clinically acquired k-space data can maintain the intensity connectivity for ROI-based low-contrast vessels better than the traditional CS.

REFERENCES

- [1] C. Zhang, M. van de Giessen, E. Eisemann, and A. Vilanova, *User-guided compressed sensing for magnetic resonance angiography*, in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2014) pp. 2416–2419.
- [2] E. J. Candes, J. Romberg, and T. Tao, *Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information*, *IEEE Transactions on Information Theory* **52**, 489 (2006).
- [3] D. L. Donoho, *Compressed sensing*, *IEEE Transactions on Information Theory* **52**, 1289 (2006).
- [4] M. Lustig, D. Donoho, and J. M. Pauly, *Sparse MRI: The application of com-*

- pressed sensing for rapid MR imaging*, *Magnetic Resonance in Medicine* **58**, 1182 (2007).
- [5] J. L. Prince and J. M. Links, *Medical imaging signals and systems* (Pearson Prentice Hall, 2006).
- [6] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, *Compressed sensing MRI*, *IEEE Signal Processing Magazine* **25**, 72 (2008).
- [7] J. Milles, M. J. Versluis, A. G. Webb, and J. H. Reiber, *Quantitative evaluation of compressed sensing in MRI: Application to 7T time-of-flight angiography*, in *IEEE International Conference on Information Technology and Applications in Biomedicine (ITAB)* (2010) pp. 1–4.
- [8] E. J. Candes, M. B. Wakin, and S. P. Boyd, *Enhancing sparsity by reweighted ℓ_1 minimization*, *Journal of Fourier Analysis and Applications* **14**, 877 (2008).
- [9] C.-H. Chang and J. Ji, *Improved compressed sensing MRI with multi-channel data using reweighted ℓ_1 minimization*, in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2010) pp. 875–878.
- [10] S. Sharma and K. Nayak, *Region of interest compressed sensing*, in *International Society for Magnetic Resonance in Medicine (ISMRM)* (2009) p. 2816.
- [11] H. Oh and S. Lee, *Visually weighted reconstruction of compressive sensing MRI*, *Magnetic Resonance Imaging* **32**, 270 (2014).
- [12] D. Proksch, J. Dornheim, and B. Preim, *Interaktionstechniken zur korrektur medizinischer 3D-segmentierungen*, *Bildverarbeitung für die Medizin* **574**, 420 (2010).
- [13] K. O. Babalola, B. Patenaude, P. Aljabar, J. Schnabel, D. Kennedy, W. Crum, S. Smith, T. F. Cootes, M. Jenkinson, and D. Rueckert, *Comparison and evaluation of segmentation techniques for subcortical structures in brain MRI*, in *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2008) pp. 409–416.

3

BACKGROUND

*If you know the enemy and know yourself,
you need not fear the result of a hundred battles.*

Sun Tzu

知彼知己，百战不殆。

孙武

This chapter presents the background knowledge which is necessary to understand the context of this thesis and our main contributions. We first introduce the macro/micro structures of the brain white matter in an intuitive way in Section 3.1, and then describe how to measure the fibrous characteristic of white matter in living tissues with MRI scanning in Section 3.2. Section 3.3 is the mathematical background related to tensor calculus while Section 3.4 is the technical background, which is concerned with the overview of the visualization techniques that are most relevant to the contributions of this thesis.

3.1. BRAIN ANATOMY




Figure 3.1: The macroscopic and microscopic view of the brain anatomy. Left: the brain is roughly composed of white matter and gray matter (<https://sampurnabuffalo.wordpress.com/2016/11/01/does-white-matter-matter/>). Right: the structure of the neurons, which are the microscopic unit of the brain (<https://en.wikipedia.org/wiki/Neuron>).

The human brain is the most complex organ of the human body, and can be roughly divided into two parts: gray matter and white matter. Intuitively, if the brain is viewed as a supercomputer in its entirety, the counterparts of gray matter are the central processing units (CPUs) that are responsible for information processing. The white matter is responsible for information transmission. It can be regarded as the cables that connect different parts of the gray matters, The gray matter and white matter together make the brain work like a coordinated complex system. Figure 3.1 left is an example of the macroscopic view of the brain anatomy. Microscopically, the distinction between gray matter and white matter is related to the biological unit of the central nervous system, i.e., the neurons. A neuron is an electrically excitable cell that mainly consists of a cell body and an axon (see Figure 3.1 right). Gray matter contains a large amount of cell bodies and relatively few axons, while white matter contains

relatively very few cell bodies and is mostly composed of long-range axons [1]. The name and color of white matter arises from the *white* substance, i.e., myelin, that surrounds the axon. Axon bundles are organized in a coherent way, and thus at large scale they constitute the so-called *fiber tracts*, which are the pathways of the central nervous system.

3.2. DIFFUSION TENSOR IMAGING

In order to quantify the underlying structure of the white matter in living tissues, we need to resort to a common physical phenomenon, the so-called *water diffusion*. Water diffusion refers to the process of random thermal motion of water molecules (Brownian motion). For example, Figure 3.2 illustrates how the diffusion behaves in the case of kleenex and newspaper, respectively. The diffusion is isotropic (i.e., no directional preference) for the kleenex while anisotropic (i.e., directional preference) for newspaper. Therefore, how the molecules diffuse reflect the fibrous structures of the underlying materials.




Figure 3.2: An illustration of the diffusion patterns when placing a drop of ink on the kleenex and newspaper, respectively (<http://www.cs.utah.edu/~gk/papers/vis99/ppt/slide03.html>).

Einstein in 1905 gave the relationship between the averaged squared displacements of freely diffusing water molecules and the elapsed time t as

$$\langle r^2 \rangle = 2Dt \quad (3.1)$$

, where D is the diffusion constant and t is the diffusion time. Molecules start to spread out from their initial positions as time elapses, and their displacements follow a Gaussian distribution

$$p(r, t) = \frac{1}{\sqrt{4\pi Dt}} e^{-r^2/4Dt} \quad (3.2)$$

, where p gives the probability of water molecules that have moved by r over time interval t . In a word, the motion of freely diffusing water molecules obeys a Gaussian distribution, the variance of which is a function of the diffusion constant D and time t .

One important aspect of water diffusion is that it leads to MR signal loss in the presence of the *diffusion gradient*, as illustrated in Figure 3.3. The application of a horizontal diffusion gradient makes water molecules rotate at different frequencies based on their horizontal positions. The phases for molecules at different horizontal locations are different after this horizontal diffusion gradient is switched off. Yellow and green rectangles indicate the occurrences of horizontal and vertical movements, respectively. Then a reversed horizontal diffusion gradient is applied to correct the introduced phase differences. However, due to the random movements in the horizontal direction, the resultant MR signal is attenuated. Furthermore, the random movements in the vertical direction have no influence on the signal intensity. Therefore, the resulting signal is said to be diffusion-weighted.




Figure 3.3: Schematic of the effects of the diffusion of water molecules on the signal intensity. Colored circles represent water molecules located at different positions. Big black arrows indicate the strengths of magnetic field while small black arrows in circles indicate the phases of each molecular spin. Yellow and green rectangles represent the random movements of water molecules. Image courtesy [2].

By incorporating the probability of diffusion of water molecules presented in Equation 3.2, the attenuated signal intensity due to diffusion can be derived, and expressed as

$$S = S_0 e^{-bD} \quad (3.3)$$

, where S and S_0 are the signal intensities with and without the application of diffusion gradient, respectively. b (the so-called b -value) is the abbreviated aggregate of several parameters (e.g., the magnitude of the applied diffusion gradient) that are controllable via MRI scanners. Detailed mathematic derivation about Equation 3.3 can be found in textbooks such as Mori and Tournier [2].

Equations 3.2 and 3.3 are correct only for 1D diffusion. To model an anisotropic Gaussian distribution in the 3D physical world, Equation 3.2 generalizes to

$$p(\mathbf{r}, t) = \frac{1}{\sqrt{(4\pi t)^3 \det(\mathbf{D})}} e^{-\frac{1}{4t} \mathbf{r}^T \mathbf{D}^{-1} \mathbf{r}} \quad (3.4)$$

, where \mathbf{D} is a second-order symmetric tensor, $\det(\bullet)$ and \mathbf{D}^{-1} are the determinant and the inverse of tensor \mathbf{D} , respectively. Consequently, Equation 3.3 generalizes to

$$S = S_0 e^{-b\mathbf{g}^T \mathbf{D} \mathbf{g}} \quad (3.5)$$

, where \mathbf{g} is the unit vector, indicating the direction of the diffusion gradient mentioned before. The magnitude of the diffusion gradient is encoded into the b -value.

Due to the utilization of tensor \mathbf{D} in the process of diffusion modeling, this imaging and modeling technique is called diffusion tensor imaging (DTI), which was first proposed by Basser et al. [3]. The diffusion tensor \mathbf{D} has six degrees-of-freedom. Therefore, six measurements (i.e., six different \mathbf{g} s) and one baseline measurement S_0 are required to uniquely determine the components of \mathbf{D} . However, in practice, more than six acquisitions in different directions are necessary for the sake of robust estimation of \mathbf{D} . The assumption that the diffusion of water molecules within a voxel obeys a Gaussian distribution is not always valid (e.g., in voxels containing crossing fiber tracts). Therefore, other complex diffusion acquisition techniques have been proposed such as high angular resolution diffusion imaging (HARDI) [4], which normally requires longer scanning time. DTI is still the most commonly used diffusion modeling technique in the clinic and often in medical research. In this thesis, we focus on developing visual analysis approaches specifically for DTI datasets.

3.3. SECOND-ORDER SYMMETRIC TENSOR

In the previous section, we explain the physical origin of diffusion tensors. In this section, we will look at them from the mathematical point of view. Mathematically, diffusion tensors are second-order symmetric positive-definite tensors. They, rather than tensors of general forms, are the focus of this thesis. Furthermore, all the second-order symmetric positive-definite tensors are expressed with respect to an orthonormal basis in \mathbb{R}^3 . Therefore, only Cartesian tensors are considered in the thesis. Interested readers are referred to textbooks [5, 6] or other PhD thesis [7, 8] for further information.

A second-order tensor describes a linear relation between vector spaces. The vector space is \mathbb{R}^3 , representing the physical world. Therefore, a tensor \mathbf{D} takes an input vector \mathbf{v} , and generates an output vector $\mathbf{w} = \mathbf{D}\mathbf{v}$. Both \mathbf{v} and \mathbf{w} are in \mathbb{R}^3 . Linearity is defined by means of the usual vector addition and scalar multiplication. Numerically, a second-order tensor \mathbf{D} can be expressed as a 3-by-3 matrix with respect to an orthonormal coordinate basis $\mathcal{B} = [\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3]$, which is usually determined by the MRI scanner in the context of DTI,

$$[\mathbf{D}]_{\mathcal{B}} = \begin{bmatrix} D_{11} & D_{12} & D_{13} \\ D_{21} & D_{22} & D_{23} \\ D_{31} & D_{32} & D_{33} \end{bmatrix} \quad (3.6)$$

where $D_{ij} = \mathbf{b}_i \cdot \mathbf{D}\mathbf{b}_j$. A tensor \mathbf{D} is symmetric ($\mathbf{D} = \mathbf{D}^T$) if $\mathbf{b}_i \cdot \mathbf{D}\mathbf{b}_j = \mathbf{b}_j \cdot \mathbf{D}\mathbf{b}_i$, i.e., $D_{ij} = D_{ji}$, which is the case for diffusion tensors.

A second-order symmetric tensor \mathbf{D} can be eigen-decomposed to three *real* eigenvalues $\lambda_1, \lambda_2, \lambda_3$, and the corresponding unit-length mutually orthogonal eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$, i.e., $\mathbf{D}\mathbf{e}_i = \lambda_i\mathbf{e}_i$ or $\mathbf{D} = \sum_{i=1}^3 \lambda_i\mathbf{e}_i \otimes \mathbf{e}_i$ where \otimes is the so-called *tensor* or *dyad* product. If all the three eigenvalues are greater than zero $\lambda_i > 0$, a tensor \mathbf{D} is said to be *positive-definite*. Unless otherwise noted, a tensor \mathbf{D} in this thesis will always be second-order symmetric and positive-definite. The eigenvalues are normally ordered as $\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$, which are the *major*, *medium*, and *minor* eigenvalue, respectively. Similarly, $\mathbf{e}_1, \mathbf{e}_2$, and \mathbf{e}_3 are referred to as the *major*, *medium*, and *minor* eigenvector, respectively.

Tensor invariants are scalar-valued properties derived from the original tensor \mathbf{D} , which are intrinsic to \mathbf{D} itself rather than its matrix representation. In other words, tensor invariants are independent of the given orthonormal coordinate system with respect

to which \mathbf{D} is expressed. For example, eigenvalues are tensor invariants while tensor components (e.g., D_{ij}) are not. Apart from the eigenvalues, any scalar-valued functions (e.g., the trace) that take the eigenvalues as input are tensor invariants as well. Tensor trace equals the sum of the eigenvalues $\text{tr}(\mathbf{D}) = \sum_{i=1}^3 \lambda_i$. The Frobenius norm is given by $\|\mathbf{D}\| = \sqrt{\sum_{i=1}^3 \lambda_i^2}$. In the context of DTI, both trace and Frobenius norm represent the overall amount of diffusion. The fractional anisotropy (FA) [9], expressed as $\text{FA} = \sqrt{\frac{1}{2} \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_3 - \lambda_1)^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}}$, reflects the extent of diffusion anisotropy. A value of zero means that diffusion is completely isotropic while a value of one means that diffusion occurs only along one single axis. Westin et al. [10] propose other invariants of anisotropy measures, i.e., c_l (linear anisotropy), c_p (planar anisotropy), and c_s (spherical anisotropy):

$$c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3} \quad c_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3} \quad c_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3} \quad (3.7)$$

These three invariants together describe the tensor shape in the so-called *barycentric* space (see Figure 3.4) since $c_l + c_p + c_s = 1$.




Figure 3.4: The barycentric space of tensor anisotropies. Image adapted from Kindlmann et al. [11].

3.4. STATE-OF-THE-ART VISUALIZATION

This thesis is dealing with the comparative and ensemble visualization of ensembles of 3D tensor fields. In this section, we first present state-of-the-art visualizations

designed for a single tensor field (i.e., second-order symmetric positive-definite tensor fields in the context of DTI). Furthermore, we present the general comparative and ensemble visualization techniques present in literature for scalar, vector, and tensor field.

3.4.1. STANDARD VISUALIZATION FOR DTI

This subsection covers a limited amount of representative visualization techniques for DTI. Interested readers are referred to dedicated book chapters [12, 13] or recent survey paper [14].

The most commonly used method is to reduce the six dimensional tensor into one or more scalar-valued quantities that have biological interpretations, and then traditional scalar visualization such as the slice view can be applied. For example, Fig-

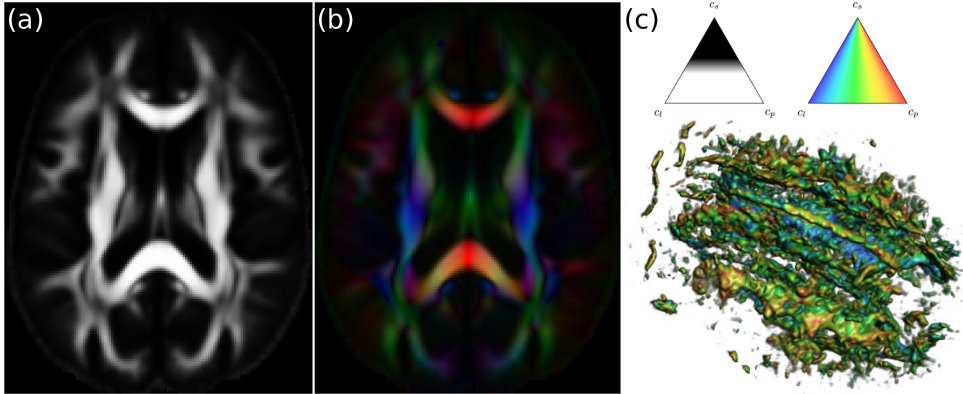


Figure 3.5: Three types of local visualization. (a) slice view of FA scalar map. (b) slice map of the major eigenvector, weighted by c_l . (c) direct volume rendering. Image courtesy [11].

ure 3.5a shows the slice view of FA, which indicates the diffusion anisotropy at different locations. The tensor can also be reduced to the major eigenvector \mathbf{e}_1 , which is empirically assumed to align with the underlying white matter structure. For the visualization, the absolute value of each component of $\mathbf{e}_1 = (x, y, z)$ is mapped to the R, G, and B color components as $(R, G, B) = (|x|, |y|, |z|)$. This color mapping, although ambiguous due to using the absolute value, is able to convey directional information. Additionally, for isotropic and planar-anisotropic tensors, the major eigenvectors are randomly orientated. Some tensor-derived scalar quantities are used to modulate the color saturation (See Figure 3.5b). Kindlmann et al. [11] extend the traditional volume

rendering for scalar fields to visualize tensor fields (see Figure 3.5c). The specification of color and opacity is performed in the barycentric space, which is based on the anisotropy indexes cl , cp , and c_s [10].




Figure 3.6: Superquadric tensor glyph and the placement strategies. (a) glyphs placed on a regular grid. (c) glyph packing. Image courtesy [15].

These kinds of visualizations are based on the reduction of the full tensor information. If to display the full tensor information, glyph-based visualization is a natural choice due to its ability to effectively encode multivariate attributes into various visual channels (e.g., color, size, or shape). Compared to other visualization techniques for multivariate data (e.g., parallel coordinate plots), glyph-based visualization is able to preserve the spatial context. The straightforward way is to employ ellipsoids [16, 17]. Westin et al. [18] design a composite glyph to directly show the linear, planar, and spherical components. Kindlmann [19] propose the superquadric glyph (see Figure 3.6a) to improve the visual perception and reduce image space ambiguity. If tensor glyphs are placed at sampling positions of a regular grid, the structure of the sampling grid, instead of the underlying continuous structure, becomes emphasized. To solve this problem, Kindlmann and Westin [15] propose *glyph packing* to make the continuous structures become more apparent by smartly distributing the glyphs across the field (see Figure 3.6b). Schultz and Kindlmann [20] propose several general principles for glyph-based visualization for symmetric tensor data including the preservation of symmetry, continuity, and disambiguity. High-order tensors can be projected onto a unit sphere via tensor contraction and thus visualized as a deformed sphere, referred to as the radial glyph [21]. Several novel glyph designs are proposed for HARDI data as well. Peeters et al. [22] render HARDI data as a deformed sphere by GPU-based

ray-casting. Schultz and Kindlmann [23] present the higher-order maximum enhancing (HOME) glyph for HARDI data by generalizing the tensor ellipsoid. Interested readers are referred to recent survey papers [24–26] for (tensor) glyph design rules.

Another type of visualization techniques, known as the *fiber tracking* [27], aims at reconstructing and displaying the pathways of neural fibers. Fiber tracking can be divided into two types, deterministic and probabilistic. Deterministic methods intend to extract the most likely fiber trajectories without considering uncertainties. The most commonly used method is *streamline*, the key idea of which is to track along the major eigenvector \mathbf{e}_1 . Streamlines are commonly visualized as plain polylines, illuminated polylines [28], or streamtubes [29] (see Figure 3.7a). Some other researchers employ illustrative rendering [30, 31] to enhance the visual perception. Weinstein et al. [32] present the *tensorline* algorithm in order to stabilize the fiber trajectories through regions with low linear diffusion anisotropies. Hyperstreamlines [33, 34] are proposed to enhance the streamline method by encoding additional information about the tensors into its cross-section. Analogously, hyperstreamline is constructed by sweeping a varying (super)ellipse [33, 35], determined by the second and third eigenvalues at each location, along the streamline.

Probabilistic methods attempt to generate a probability volume (see Figure 3.7b) to indicating the probability that each voxel is connected to a user-specified source voxel. Several most likely fiber tracts can then be derived based on this probability volume. For example, Brecheisen et al. [36] present approaches to visualize noise-induced uncertainty in probabilistic fiber tracking.

This thesis aims at avoiding information reduction and displaying the full information of tensor data, and glyph-based visualization is a promising way to encode multi-variate data in the spatial domain. Therefore, this thesis focus specifically on glyph-based visualization.

3.4.2. COMPARATIVE/ENSEMBLE VISUALIZATION

Comparative Visualization. Comparative visualization refers to the process of understanding the differences or similarities between two or more datasets by making use of an easy-to-interpret visual representation. Gleicher et al. [38] present a general taxonomy that groups visual comparison into three categories: juxtaposition (or side-by-side), superposition (or overlay), and explicit encoding of differences. Juxtaposition

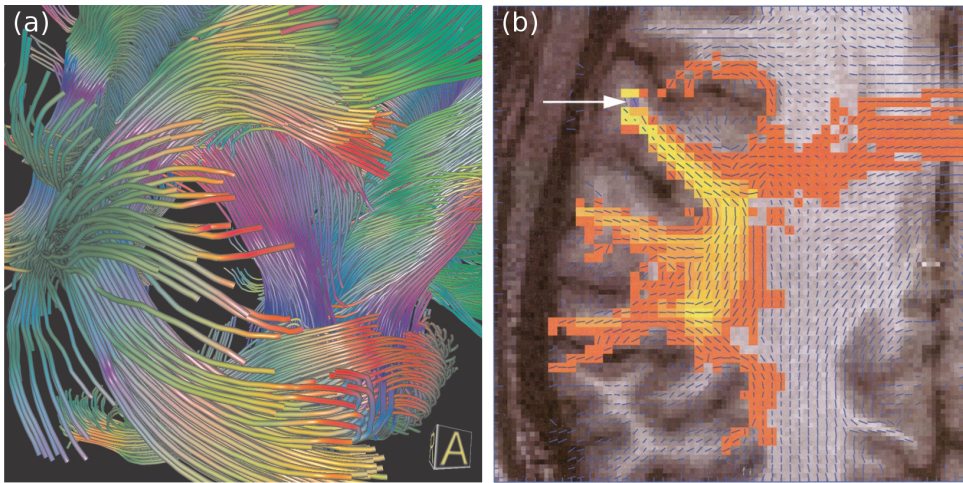


Figure 3.7: Visualization of (a) deterministic and (b) probabilistic fiber tracking results. (a) deterministic fiber tracts are visualized as streamtubes. Image courtesy [29]. (b) probabilistic tractography. Yellow represents high probability and orange lower probability. The source voxel is indicated by the white arrow. Image courtesy [37].

is straightforward to implement but relies on the viewers' memory for comparison. Superposition is effective for comparing objects, due to being in the same frame of reference. However, it is likely to cause occlusions and visual clutter. Explicit encoding is effective for explicitly depicting the differences between objects. However, the information about the original subjects themselves is lost, and it is not trivial to define a metric to quantify the difference.

Comparative visualization has been developed to facilitate data comparison in the spatial domain. For example, Busking et al. [39] propose an image-based implementation to visually compare two intersecting surfaces by local distance cues (e.g., fogging) and contours. Schmidt et al. [40] introduce a comparative visual analysis system for multiple 3D meshes that combines explicit encoding, juxtaposition, and parallel coordinate plots for quantitative measures. van Pelt et al. [41] present an intuitive details-on-demand glyph set for comparatively visualizing wall-shear stresses between different stent configurations.

However, few methods are designed specifically for comparing tensor fields. The most common approach to compare subjects in DTI visualization is to place the resulting images side-by-side. For example, Hotz et al. [42] compare the results of two tensor interpolation approaches by juxtaposition. Overlaying glyphs [43, 44] is also

a commonly used strategy. Da Silva et al. [45] depict the differences between two DTI datasets by overlaying the corresponding extracted streamtubes, rather than the original diffusion tensor fields. In the quantitative DTI image analysis, the comparison is more often performed without considering the spatial information. Line charts or scatter plots of relevant scalar-valued invariants are commonly juxtaposed and/or superposed for comparison in quantitative neuroscientific studies [46]. No glyph-based technique has been presented to facilitate the visual comparison of two DTI datasets, and we handle this case in Chapter 4.

Ensemble Visualization. *Ensemble* data comes from multi-run simulations with different parameter settings [47]. They are commonly used for assessing uncertainty and variability, for instance, in weather forecasting [48]. Ensemble visualization and uncertainty visualization are often discussed together [47, 49]. For scalar fields, uncertainty information is often presented as the variation of isocontours, such as the spaghetti plot [50] or contour boxplot [51], or as isosurfaces as presented by Pöthkow et al. [52, 53] and Pfaffelmoser et al. [54]. Djurcilov et al. [55] directly visualize scalar field uncertainty in direct volume rendering. For vector fields, curve boxplots [56] and streamline variability plots [57] focus on the variation of features extracted from the field, rather than the field itself. Botchen et al. [58] use a 3D noise texture of different frequencies to visualize vector uncertainty in flow fields. Otto et al. [59] analyze the global uncertainty in the 3D vector fields by a topological approach. Brecheisen et al. [36, 60] propose a visualization tool to explore the uncertainty of an ensemble of fiber tracts. These methods can be used to visualize the variation of tensor-derived features. However, we focus on the visualization of the full tensor information without information reduction.

Glyphs are commonly used to visualize multivariate local attributes of ensemble data. Potter et al. [61] present a new hybrid summary plot that incorporates descriptive statistics. Höllt et al. [62] use a glyph, based on the violin plot, to indicate variation in multiple ensemble variables for time series visualization. Pfaffelmoser et al. [63] present uncertainty in gradient orientation within scalar-field ensembles using circular glyphs. Hlawatsch et al. [64] introduce the flow radar glyph to display possible ranges of flow directions. Jarema et al. [65] use glyphs to show modality information of circular distributions in 2D vector fields ensemble. We design a new glyph in Chapter 5

to encode the variation information of ensembles of 3D tensor fields.

There is a limited amount of previous work that focuses on the visual analysis of ensembles of diffusion tensor fields in their entirety. In order to aggregate the ensemble information, multivariate statistics should be used. Basser and Pajevic [21, 66] propose the use of the fourth-order covariance tensor based on the assumption that the set of diffusion tensors follows a multivariate normal distribution. Basser and Pajevic [21] use radial glyphs to visualize the fourth-order tensor and its six orthogonal second-order eigentensors or eigenmodes. In recent work, Abbasloo et al. [43] combine slice views, volume rendering, and superquadric glyphs [19] to visualize the fourth-order covariance tensor at multiple levels of detail. For the visual analysis of an ensemble of orientation diffusion functions, Jiao et al. [67] use direct volume rendering based on the so-called shape inclusion probability function. Our work in Chapter 5 enriches this research direction.

The pair-wise comparison of ensembles of diffusion tensor fields adds more complexity in the visual analysis. The standard approach relies on the univariate statistical analysis of tensor-derived scalar quantities such as fractional anisotropy (FA). Slice views are then used to display the analysis results, which are some specific highly aggregated scalar values (e.g., the t-statistic) [68]. No visualization technique has been developed specifically for inter-ensemble comparison, and we handle this case in Chapter 6.

REFERENCES

- [1] D. Purves, G. J. Augustine, D. Fitzpatrick, W. C. Hall, A.-S. LaMantia, and L. E. White, *Neuroscience*, 5th ed. (Sinauer Associates, 2011).
- [2] S. Mori and J.-D. Tournier, *Introduction to diffusion tensor imaging: and higher order models*, 2nd ed. (Academic Press, 2013).
- [3] P. Basser, J. Mattiello, and D. Lebihan, *Estimation of the effective self-diffusion tensor from the nmr spin echo*, *Journal of Magnetic Resonance, Series B* **103**, 247 (1994).
- [4] D. S. Tuch, *Diffusion MRI of complex tissue structure*, Ph.D. thesis, Harvard University (2002).

- [5] P. C. Kendall and D. Bourne, *Vector analysis and Cartesian tensors* (CRC Press, 1992).
- [6] G. A. Holzapfel, *Nonlinear Solid Mechanics* (John Wiley & Sons, Ltd., 2000) Chap. 1.
- [7] G. Kindlmann, *Visualization and Analysis of Diffusion Tensor Fields*, Ph.D. thesis, The University of Utah (2004).
- [8] T. Schultz, *Feature Extraction for Visual Analysis of DW-MRI Data*, Ph.D. thesis, Saarland University (2009).
- [9] P. J. Basser and C. Pierpaoli, *Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI*, *Journal of Magnetic Resonance* **111**, 209 (1996).
- [10] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz, *Geometrical diffusion measures for MRI from tensor basis analysis*, in *ISMRM '97* (Vancouver Canada, 1997) p. 1742.
- [11] 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 (2000).
- [12] A. Leemans, *Visualization of diffusion MR images*, (Oxford University Press, 2010) Chap. Diffusion MRI: Theory, Methods and Applications, pp. 354–379, first edition ed.
- [13] B. Preim and C. P. Botha, *Visual computing for medicine: theory, algorithms, and applications*, 2nd ed. (Morgan Kaufmann, 2013).
- [14] D. S. Margulies, J. Böttger, A. Watanabe, and K. J. Gorgolewski, *Visualizing the human connectome*, *NeuroImage* **80**, 445 (2013), mapping the Connectome.
- [15] G. Kindlmann and C.-F. Westin, *Diffusion tensor visualization with glyph packing*, *IEEE Transactions on Visualization and Computer Graphics* **12**, 1329 (2006).
- [16] C. Pierpaoli and P. J. Basser, *Toward a quantitative assessment of diffusion anisotropy*, *Magnetic Resonance in Medicine* **36**, 893 (1996).

- [17] D. H. Laidlaw, E. T. Ahrens, D. Kremers, M. J. Avalos, R. E. Jacobs, and C. Readhead, *Visualizing diffusion tensor images of the mouse spinal cord*, in *Proceedings of IEEE Visualization* (1998) pp. 127–134.
- [18] C. F. Westin, S. E. Maier, B. Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis, *Image processing for diffusion tensor magnetic resonance imaging*, in *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MIC-CAI)*, edited by C. Taylor and A. Colchester (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 441–452.
- [19] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004) pp. 147–154.
- [20] T. Schultz and G. L. Kindlmann, *Superquadric glyphs for symmetric second-order tensors*, *IEEE Transactions on Visualization and Computer Graphics* **16**, 1595 (2010).
- [21] P. Basser and S. Pajevic, *Spectral decomposition of a 4th-order covariance tensor: applications to diffusion tensor MRI*, *Signal Processing* **87**, 220 (2007).
- [22] T. H. J. M. Peeters, V. Prckovska, M. van Almsick, A. Vilanova, and B. M. ter Haar Romeny, *Fast and sleek glyph rendering for interactive HARDI data exploration*, in *IEEE Pacific Visualization Symposium* (2009) pp. 153–160.
- [23] T. Schultz and G. Kindlmann, *A maximum enhancing higher-order tensor glyph*, *Computer Graphics Forum* **29**, 1143 (2010).
- [24] A. Kratz, C. Auer, and I. Hotz, *Tensor invariants and glyph design*, in *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*, edited by C.-F. Westin, A. Vilanova, and B. Burgeth (2014) pp. 17–34.
- [25] T. Ropinski, S. Oeltze, and B. Preim, *Survey of glyph-based visualization techniques for spatial multivariate medical data*, *Computers & Graphics* **35**, 392 (2011).
- [26] R. Borgo, J. Kehrer, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen, *Glyph-based visualization: Foundations, design guidelines, techniques and applications*, in *Eurographics State of the Art Reports* (2013) pp. 39–63.

- [27] S. Mori and P. B. Barker, *Diffusion magnetic resonance imaging: Its principle and applications*, *The Anatomical Record* **257**, 102 (1999).
- [28] D. Merhof, F. Enders, F. Vega, P. Hastreiter, C. Nimsy, and M. Stamminger, *Integrated visualization of diffusion tensor fiber tracts and anatomical data*, in *Simulation and visualization* (2005) pp. 153–164.
- [29] D. Merhof, M. Sonntag, F. Enders, C. Nimsy, P. Hastreiter, and G. Greiner, *Hybrid visualization for white matter tracts using triangle strips and point sprites*, *IEEE Transactions on Visualization and Computer Graphics* **12**, 1181 (2006).
- [30] R. Otten, A. Vilanova, and H. Van De Wetering, *Illustrative white matter fiber bundles*, *Computer Graphics Forum* **29**, 1013 (2010).
- [31] M. H. Everts, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg, *Depth-dependent halos: Illustrative rendering of dense line data*, *IEEE Transactions on Visualization and Computer Graphics* **15**, 1299 (2009).
- [32] D. Weinstein, G. Kindlmann, and E. Lundberg, *Tensorlines: advection-diffusion based propagation through diffusion tensor fields*, in *Proceedings of IEEE Conference on Visualization* (1999) pp. 249–530.
- [33] T. Delmarcelle and L. Hesselink, *Visualization of second order tensor fields and matrix data*, in *Proceedings of IEEE Conference on Visualization* (1992) pp. 316–323.
- [34] T. Delmarcelle and L. Hesselink, *Visualizing second-order tensor fields with hyperstreamlines*, *IEEE Computer Graphics and Applications* **13**, 25 (1993).
- [35] V. Wiens, L. Schlaffke, T. Schmidt-Wilcke, and T. Schultz, *Visualizing Uncertainty in HARDI Tractography Using Superquadric Streamtubes*, in *EuroVis Short Papers*, edited by N. Elmqvist, M. Hlawitschka, and J. Kennedy (The Eurographics Association, 2014).
- [36] R. Brecheisen, B. Platel, B. M. ter Haar Romeny, and A. Vilanova, *Illustrative uncertainty visualization of dti fiber pathways*, *The Visual Computer* **29**, 297 (2013).

- [37] M. A. Koch, D. G. Norris, and M. Hund-Georgiadis, *An investigation of functional and anatomical connectivity using magnetic resonance imaging*, *NeuroImage* **16**, 241 (2002).
- [38] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts, *Visual comparison for information visualization*, *Information Visualization* **10**, 289 (2011).
- [39] S. Busking, C. P. Botha, L. Ferrarini, J. Milles, and F. H. Post, *Image-based rendering of intersecting surfaces for dynamic comparative visualization*, *The Visual Computer* **27**, 347 (2011).
- [40] J. Schmidt, R. Preiner, T. Auzinger, M. Wimmer, M. E. Gröller, and S. Bruckner, *Ymca-your mesh comparison application*, in *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014) pp. 153–162.
- [41] R. van Pelt, R. Gasteiger, K. Lawonn, M. Meuschke, and B. Preim, *Comparative blood flow visualization for cerebral aneurysm treatment assessment*, *Computer Graphics Forum* **33**, 131 (2014).
- [42] I. Hotz, J. Sreevalsan-Nair, H. Hagen, and B. Hamann, *Tensor field reconstruction based on eigenvector and eigenvalue interpolation*, in *Scientific Visualization: Advanced Concepts*, Dagstuhl Follow-Ups, Vol. 1 (2010) pp. 110–123.
- [43] A. Abbasloo, V. Wiens, M. Hermann, and T. Schultz, *Visualizing tensor normal distributions at multiple levels of detail*, *IEEE Transactions on Visualization and Computer Graphics* **22**, 975 (2016).
- [44] M. Meuschke, S. Voss, O. Beuing, B. Preim, and K. Lawonn, *Glyph-based comparative stress tensor visualization in cerebral aneurysms*, *Computer Graphics Forum* **36**, 99 (2017).
- [45] M. J. Da Silva, S. Zhang, C. Demiralp, and D. H. Laidlaw, *Visualizing the differences between diffusion tensor volume images*, in *International Society for Magnetic Resonance in Medicine (ISMRM) Workshop on Diffusion MRI* (2000) pp. 237–238.
- [46] M. de Groot, *Cross-Subject Image Analysis in Diffusion Brain MRI*, Ph.D. thesis, Erasmus University Rotterdam (2014).

- [47] G.-P. Bonneau, H.-C. Hege, C. R. Johnson, M. M. Oliveira, K. Potter, P. Rheingans, and T. Schultz, *Overview and state-of-the-art of uncertainty visualization*, in *Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization* (2014) pp. 3–27.
- [48] J. Sanyal, S. Zhang, J. Dyer, A. Mercer, P. Amburn, and R. Moorhead, *Noodles: A tool for visualization of numerical weather model ensemble uncertainty*, *IEEE Transactions on Visualization and Computer Graphics* **16**, 1421 (2010).
- [49] A. L. Love, A. Pang, and D. L. Kao, *Visualizing spatial multivalued data*, *IEEE Computer Graphics and Applications* **25**, 69 (2005).
- [50] K. Potter, A. Wilson, P. T. Bremer, D. Williams, C. Doutriaux, V. Pascucci, and C. R. Johnson, *Ensemble-vis: A framework for the statistical visualization of ensemble data*, in *IEEE International Conference on Data Mining Workshops* (2009) pp. 233–240.
- [51] R. T. Whitaker, M. Mirzargar, and R. M. Kirby, *Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles*, *IEEE Transactions on Visualization and Computer Graphics* **19**, 2713 (2013).
- [52] K. Pöthkow and H.-C. Hege, *Positional uncertainty of isocontours: Condition analysis and probabilistic measures*, *IEEE Transactions on Visualization and Computer Graphics* **17**, 1393 (2011).
- [53] K. Pöthkow, B. Weber, and H.-C. Hege, *Probabilistic marching cubes*, *Computer Graphics Forum* **30**, 931 (2011).
- [54] T. Pfaffelmoser, M. Reitering, and R. Westermann, *Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields*, *Computer Graphics Forum* **30**, 951 (2011).
- [55] S. Djurcilov, K. Kim, P. F. J. Lermusiaux, and A. Pang, *Volume rendering data with uncertainty information*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2001) pp. 243–252.
- [56] M. Mirzargar, R. T. Whitaker, and R. M. Kirby, *Curve boxplot: Generalization of boxplot for ensembles of curves*, *IEEE Transactions on Visualization and Computer Graphics* **20**, 2654 (2014).

- [57] F. Ferstl, K. Bürger, and R. Westermann, *Streamline variability plots for characterizing the uncertainty in vector field ensembles*, IEEE Transactions on Visualization and Computer Graphics **22**, 767 (2016).
- [58] R. P. Botchen, D. Weiskopf, and T. Ertl, *Texture-based visualization of uncertainty in flow fields*, in *IEEE Visualization (2005)* pp. 647–654.
- [59] M. Otto, T. Germer, and H. Theisel, *Uncertain topology of 3D vector fields*, in *IEEE Pacific Visualization Symposium (2011)* pp. 67–74.
- [60] R. Brecheisen, A. Vilanova, B. Platel, and B. ter Haar Romeny, *Parameter sensitivity visualization for dti fiber tracking*, IEEE Transactions on Visualization and Computer Graphics **15**, 1441 (2009).
- [61] K. Potter, J. Kniss, R. Riesenfeld, and C. Johnson, *Visualizing summary statistics and uncertainty*, Computer Graphics Forum **29**, 823 (2010).
- [62] T. Höllt, A. Magdy, P. Zhan, G. Chen, G. Gopalakrishnan, I. Hoteit, C. D. Hansen, and M. Hadwiger, *Ovis: A framework for visual analysis of ocean forecast ensembles*, IEEE Transactions on Visualization and Computer Graphics **20**, 1114 (2014).
- [63] T. Pfaffelmoser, M. Mihai, and R. Westermann, *Visualizing the variability of gradients in uncertain 2D scalar fields*, IEEE Transactions on Visualization and Computer Graphics **19**, 1948 (2013).
- [64] M. Hlawatsch, P. Leube, W. Nowak, and D. Weiskopf, *Flow radar glyphs-static visualization of unsteady flow with uncertainty*, IEEE Transactions on Visualization and Computer Graphics **17**, 1949 (2011).
- [65] M. Jarema, I. Demir, J. Kehrer, and R. Westermann, *Comparative visual analysis of vector field ensembles*, in *IEEE Conference on Visual Analytics Science and Technology (2015)* pp. 81–88.
- [66] P. J. Basser and S. Pajevic, *A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI*, IEEE Transactions on Medical Imaging **22**, 785 (2003).

- [67] F. Jiao, J. M. Phillips, Y. Gur, and C. R. Johnson, *Uncertainty visualization in HARDI based on ensembles of odfs*, in *IEEE Pacific Visualization Symposium* (2012) pp. 193–200.
- [68] T. Su, M. W. Caan, F. W. Wit, J. Schouten, G. J. Geurtsen, J. H. Cole, D. J. Sharp, F. M. Vos, M. Prins, P. Portegies, *et al.*, *White matter structure alterations in hiv-1-infected men with sustained suppression of viraemia on treatment*, *AIDS* **30**, 311 (2016).

4

GLYPH-BASED COMPARATIVE VISUALIZATION FOR DIFFUSION TENSOR FIELDS

*Attack him where he is unprepared,
appear where you are not expected.*

Sun Tzu

攻其无备，出其不意。

孙武

This chapter is based on the following publication: “Glyph-based Comparative Visualization for Diffusion Tensor Fields”, **Changong Zhang**, Thomas Schultz, Kai Lawonn, Elmar Eisemann, and Anna Vilanova. In *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Scientific Visualization)*, 22(1):797–806, 2016 [1].

In the previous chapter, we present the state-of-art comparative visualization strategies for tensor fields. In previous work, tensor-derived scalar-valued quantities are the object of comparison, instead of the tensor data itself. From the perspective of the applied comparison strategy, it is common to use juxtaposition and/or superposition. In this chapter, we present a glyph-based comparative visualization technique at tensor level to show the tensor differences in a more effective manner.

4.1. INTRODUCTION

Diffusion tensors are usually reduced to scalar values that have biological interpretations (e.g., fractional anisotropy, FA). Analysts then carry out comparison on the derived scalar fields, which are compared by juxtaposing slice views [2]. However, scalar level comparison is based on the reduction of the full information of diffusion tensors, and thus gives a limited view on the differences.

Alternatively, the comparison can be performed directly on diffusion tensors. The most commonly employed method for tensor level comparison is to juxtapose the corresponding tensor glyphs [3]. Superposing tensor glyphs [4, 5] is also used but can easily be affected by occlusion. Moreover, Da Silva et al. [6] depict the differences between two DTI datasets by overlaying the corresponding streamtubes, rather than the original diffusion tensor fields. However, the difference between the underlying datasets is not equal to the difference between the extracted streamtubes. Intuitively speaking, two different DTI datasets might produce exactly the same streamtubes. Therefore, we decide to compare two DTI datasets at the tensor level.

In Figure 4.1, we illustrate the results of juxtaposition and superposition using two synthetic DTI datasets. Dataset 1 in Figure 4.1a is defined in the barycentric space of three geometric anisotropy metrics [7]. Dataset 2 in Figure 4.1b is constructed by applying a small amount of random variations to dataset 1. Both are visualized via the superquadric glyphs [8]. With juxtaposition (Figures 4.1a & 4.1b), analysts need to compare them mentally, and thus it is hard to distinguish subtle differences. With superposition (Figure 4.1c), it becomes easier to identify whether two tensors are different or not, if their sizes are comparable. If not, one could be totally enclosed in the other. Occlusion is another problem when directly overlaying two glyphs. An option to address this is transparency. However, subtle differences still remain unclear, as shown in Figure 4.1d. In all the methods mentioned above, the differences are

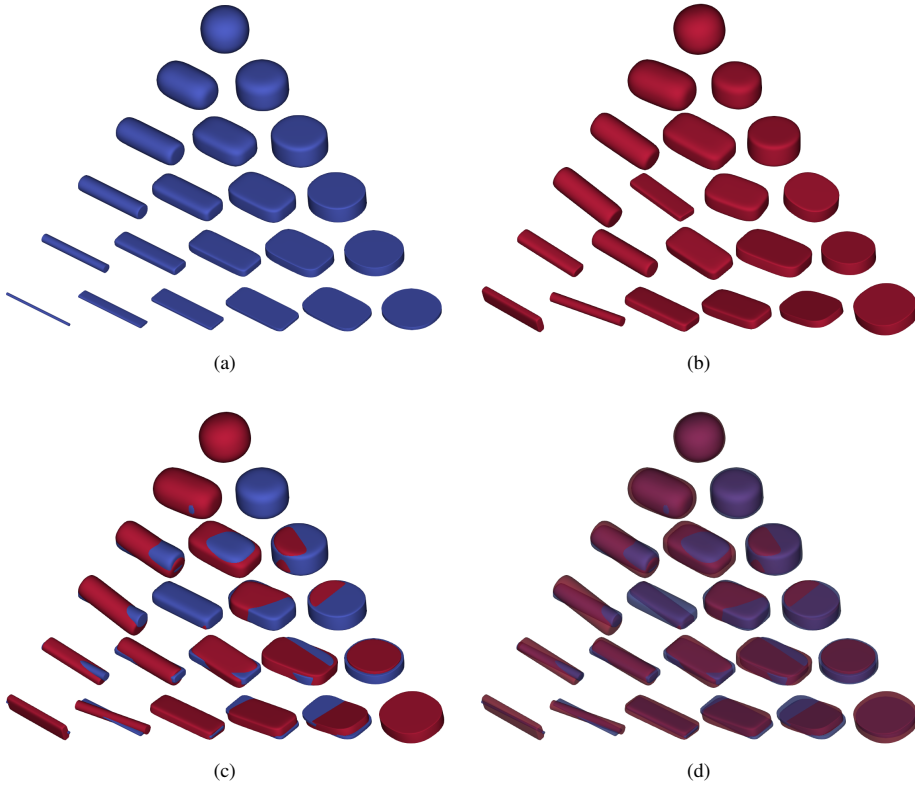


Figure 4.1: Different comparative visualization styles for two synthetic tensor datasets: (a & b) juxtaposition, (c) superposition, and (d) superposition with transparency.

integrally, instead of *separately*, perceived. In other words, even if it is easy to tell whether two tensors are different or not, it is hard to tell, for example, which aspect contributes most to the perceived overall difference.

In this chapter, we focus on the local comparison of diffusion tensors and present the *Tensor difference (Tender)* glyph to efficiently depict the differences between two diffusion tensor fields in all the three tensor intrinsic properties (i.e., scale, shape, and orientation). To this end, we have to design dissimilarity measures to quantify differences in tensor intrinsic properties. The Tender glyph is inspired by the checkerboard style visualization [9] and the superquadric tensor glyph [8]. Checkerboard visualization, a type of side-by-side visual comparison, is frequently used to evaluate registration accuracy [10]. Differences result in obvious discontinuities along the checker-

board edges, which can be readily perceived. The superquadric tensor glyph generates a strong visual cue for tensor shape and orientation via sharp edges, and largely reduces visual ambiguity.

The main contributions of this chapter are:

- A tensor dissimilarity measure to quantify individual differences in tensor scale, shape, and orientation, together with the associated feature space to support linked brushing;
- A novel glyph design, encoding differences between two diffusion tensors in tensor scale, shape, and orientation, to facilitate the visual comparison;

4

4.2. GLYPH DESIGN

The goal of our glyph design is to help analysts gain insight into the differences between two diffusion tensors beyond what a straightforward juxtaposition or superposition can provide. The first step to build the glyph is to analyze what we want to achieve by comparing two DTI datasets. Based on this analysis, we define the dissimilarity measures, and these measures are then encoded into various visual channels of the glyph. In this section, we present the whole design process.

4.2.1. TASK ANALYSIS

Before starting the design process, we must decide what kind of information is useful to compare. The diffusion scale (e.g., mean diffusivity) represents the overall diffusion extent. The diffusion anisotropy type is useful for identifying the underlying fibrous structures. In areas with single fibers such as the corpus callosum, the diffusion profile is of a linear shape, while in the areas with crossing fibers such as the centrum semiovale, the diffusion profile is of a planar shape. Since both diffusion scale and anisotropy do not depend on orientation, they are rotationally invariant. If a tensor is not spherical, its diffusion orientation can be well defined. The orientation is important for deducing the pathway of the underlying fiber tracts. Therefore, we decide to divide the differences into diffusion scale, shape, and orientation. These three components are independent of each other [11], e.g., a change in orientation will neither cause a change in scale nor in shape. However, the major challenge is to deal appropriately with the inherently coupled relation between shape and orientation.

We adopt two requirements for the glyph design. Firstly, the glyph should be able to provide information about the amount of differences of the three components (i.e., scale, shape, and orientation) *separately*. Secondly, the glyph for two distinct tensors should emphasize the visual representation, while that for two nearly equal tensors should not. In this work, we employ the superquadric glyph, which can guarantee unambiguous perception and preserve continuity. Some properties of the superquadric glyph can be found in Kindlmann [8] and Demiralp et al. [12].

4.2.2. TENSOR DISSIMILARITY MEASURE

An important challenge in tensor data comparison approach is to design suitable measures, which is also an important factor in tensor field segmentation [13].

There is no unique way to calculate the distance between tensors. In literature, several measures have been proposed to calculate the similarity or distance between tensors, see Peeters et al. [14] for a detailed discussion. An easy to compute and widely used measure is the Frobenious norm, which can be applied to all types of tensors. Other metrics have been proposed specifically for symmetric and strictly positive definite tensors, such as the Riemannian metric [15] or the log-Euclidean approach [16]. Kindlmann et al. [17] propose a measure built on the geodesic-loxodrome that divides the overall distance into components that have intuitive meanings for DTI (i.e., shape-specific and orientation-specific distances). The fact that these measures produce intuitive distance components makes it an interesting choice for our glyph design.

In order to define the dissimilarity measure, we decompose a diffusion tensor in a similar way as Schultz and Kindlmann [18]. A symmetric second-order positive-definite tensor \mathbf{D} can be decomposed to three ordered real eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$. The trace of a tensor is defined as $|\mathbf{D}| = \sum_{i=1}^3 \lambda_i$. Then we get the normalized tensor $\tilde{\mathbf{D}} = \mathbf{D} / |\mathbf{D}|$ and the normalized eigenvalues $\tilde{\lambda}_i = \lambda_i / |\mathbf{D}|$.

The tensor trace $|\mathbf{D}|$ represents the overall diffusion scale, while the set of three normalized ordered eigenvalues $(\lambda_1, \lambda_2, \lambda_3)$ represents the shape of the diffusion profile. The set of three orthonormal eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ represents the diffusion orientation. We define the tensor dissimilarity measures between two tensors $\mathbf{D}^{(1)}$ and $\mathbf{D}^{(2)}$ as:

$$d_{scl}(\mathbf{D}^{(1)}, \mathbf{D}^{(2)}) = \left| |\mathbf{D}^{(1)}| - |\mathbf{D}^{(2)}| \right| \quad (4.1)$$

$$\tilde{d}(\mathbf{D}^{(1)}, \mathbf{D}^{(2)}) = \|\tilde{\mathbf{D}}^{(1)} - \tilde{\mathbf{D}}^{(2)}\| \quad (4.2)$$

$$d_{shp}(\mathbf{D}^{(1)}, \mathbf{D}^{(2)}) = \sqrt{\sum (\tilde{\lambda}_i^{(1)} - \tilde{\lambda}_i^{(2)})^2} \quad (4.3)$$

$$d_{ori}^2(\mathbf{D}^{(1)}, \mathbf{D}^{(2)}) = \tilde{d}^2 - d_{shp}^2 \quad (4.4)$$

where d_{scl} is the scale difference, \tilde{d} is the normalized tensor difference, d_{shp} is the shape difference, and d_{ori} is the orientation difference.

The main idea is based on the Pythagorean theorem. The difference between two normalized diffusion tensors can only be due to the shape and/or orientation difference. If we subtract the shape difference from the normalized difference, only the orientation difference remains.

The advantage of the difference measures is that they are easy to calculate, and intuitive. More specifically, the orientation difference measure d_{ori} does not measure a spurious non-zero value in cases where one of the tensors is isotropic. As a demonstration, we select $\mathbf{D}^{(1)} = [0.94 \ 0 \ 0; 0 \ 0.24 \ 0; 0 \ 0.24 \ 0]$, $\mathbf{D}^{(2)} = [0.24 \ 0 \ 0; 0 \ 0.94 \ 0; 0 \ 0.24 \ 0]$. We keep $\mathbf{D}^{(1)}$ fixed, and gradually change $\mathbf{D}^{(2)}$ to an isotropic tensor by reducing its FA to 0 while preserving the scale and tensor mode, as shown in Figure 4.2. This experiment confirms the continuity of our orientation difference measure. Furthermore, the orientation difference between a linear tensor and a planar tensor is the largest if the first eigenvector of the linear tensor is perpendicular to the plane spanned by first two eigenvectors of the planar tensor. More importantly, both shape difference and orientation difference are bounded. Using the geodesic-loxodromes [17] as the dissimilarity measure would also provide the desired characteristics to quantify tensor differences in scale, shape, and orientation individually. However, the computational cost of calculating the geodesics is considerably high, so we opt for the computationally simpler alternative presented here.

4.2.3. VISUAL MAPPING

In this subsection, we encode the pair-wise differences in scale, shape, and orientation to various visual channels that constitute the Tender glyph.

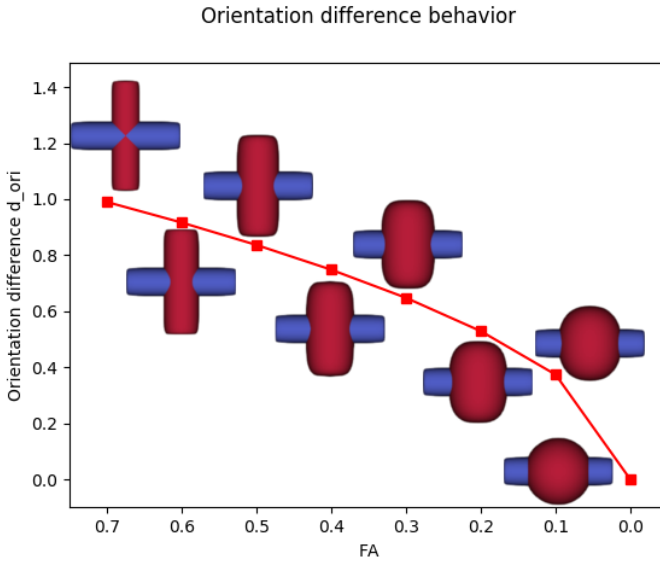


Figure 4.2: The orientation difference d_{ori} gradually decreases to 0 when reducing the FA of one linearly anisotropic tensor (the red one) to 0 while keeping the other (the blue one) fixed. Both tensors are of unit scales.

Shape Difference Encoding. Following the design guideline of intuitive mapping based on semantics [19], the shape difference d_{shp} can be encoded into the shape channel of the glyph. However, summarized as a single scalar value, it is impossible to connect d_{shp} back to the original tensor shapes. For instance, d_{shp} can be the same between an isotropic tensor and a linear tensor as well as a planar tensor and a linear tensor. Therefore, we do not use an explicit encoding of the shape difference. Instead, we preserve the original tensor shapes as much as possible while still facilitating their visual comparison.

Tensor shape is scale-invariant, and hence, we normalize both tensors before shape comparison to remove the influence of tensor scale on shape comparison. We also observe that it is hard to identify the shape difference in terms of three eigenvalues when two tensor glyphs have different orientations. For instance, as shown in Figure 4.3a, two tensors are of unit scale but different orientations and shapes. No matter how we compare them, by juxtaposition (see Figure 4.3a) or by superposition (see Figure 4.3b bottom), it is difficult to identify the shape difference by sight. As soon as we align them to the orientation of one of these two tensor glyphs (see Figure 4.3b top), it becomes easier to tell which tensor has larger or smaller eigenvalues. However,

occlusion still makes it hard to judge the extent of shape differences. Transparency is used to handle occlusion (see Figure 4.3c), but it remains hard to visually quantify the shape differences.

Inspired by the checkerboard visualization, we explain our idea in 2D in Figure 4.4a. One of the superquadric glyphs is first aligned with the other (see Figure 4.4b). Therefore, we can maintain the original orientation of one tensor. Then, we divide the 2D space into four quadrants (i.e., eight parts in 3D), each bounded by two half-axes displayed as dotted black lines in Figure 4.4c, like a checkerboard. We alternate the displaying of the corresponding parts of superquadric glyphs in each quadrant. Figure 4.3d shows the corresponding 3D case. If two tensors have similar shapes, there will be no significant changes along the axes. Otherwise, abrupt changes appear which are easy to recognize. The traditional checkerboard visualization introduces information loss due to its characteristic of alternative displaying. However, there is no shape-specific information loss in our case, since both the tensor and the superquadric glyphs are symmetric and one octant contains all information. Furthermore, we still keep the sharp edges, which serve as strong visual cues for orientation and shape [8]. Note that the use of orthographic projection is necessary since the shape comparison relies on the correct perception of length differences along the axes.

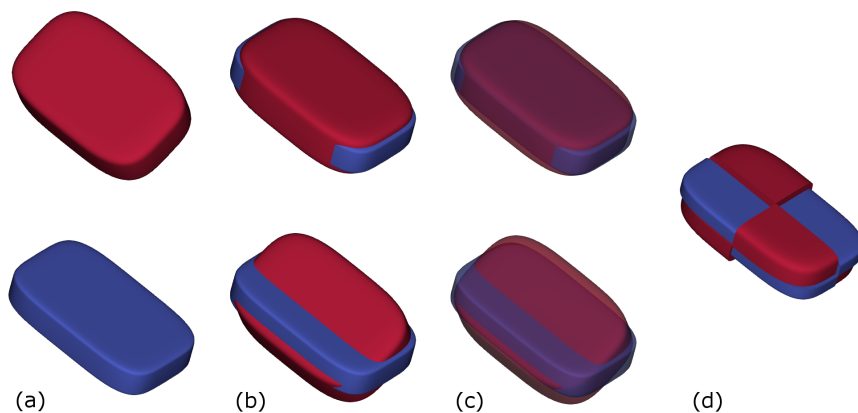


Figure 4.3: The Tender glyph shows the shape difference between two tensors of unit scale but with different orientations. (a) shows two tensor glyphs placed side-by-side. (b) bottom shows them overlaid. (b) top shows them overlaid, but the red glyph is aligned with the blue one. (c) superposition with transparency. (d) the Tender glyph. The differences in terms of three eigenvalues are more obviously manifested by the "staircases" between two parts of the Tender glyph.




Figure 4.4: 2D Illustration of the checkerboard style Tender glyph design. (a) shows two 2D glyphs side-by-side. (b) shows them overlaid but the red glyph is aligned with the blue one. The solid lines represent the visible parts while the dash-dot lines are the invisible parts due to occlusion. (c) shows the result of our design. The two black dot lines divide the 2D space into four parts. Two glyphs are displayed alternatively. There are obvious discontinuities between them.

Scale Difference Encoding. Tensor scale is interpreted as the ‘size’, representing the amount of diffusion. It is intuitive to encode the scale difference d_{scl} to the size channel of the Tender glyph. However, if so, it is impossible to tell which tensor has a larger or smaller scale. An alternative is to directly map the original tensor scale $|\mathbf{D}|$ to the size of its corresponding part of the checkerboard superquadrics. However, these mappings do pose a problem for the perception of tensor scale. The superquadric glyph inherently introduces volume differences for tensors with the same trace, but with different shapes. It is then difficult to isolate the diffusion scale differences. More importantly, if so, it will be hard to compare tensor shapes. Therefore, we decide to encode either the individual tensor scales $|\mathbf{D}^{(i)}|$ or the scale differences d_{scl} into color.

Encoding individual tensor scales enables voxel-wise scale comparison, while encoding scale differences can facilitate the comparison across the field. Additionally, we design two colormaps for scale information encoding. We use the same tensor datasets as shown in Figure 4.1 for illustration. The dual hue colormap is used to show individual tensor scales (see Figure 4.5a). Hue, as an effective categorical cue, is used to distinguish tensor datasets. Scale information is reflected by color luminance. Perceptually, it is easier to compare color luminance than to evaluate size differences when shapes vary. However, it is hard to compare the subtle luminance differences of two colors with different hues for a given voxel (see Figure 4.5a). Thus, we introduce the single hue colormap (see Figure 4.5b). It facilitates detecting subtle luminance differences for both local individual scales and scale differences across the field (see Figure 4.5d). Colors of the same hue can facilitate the comparison of luminance. Glyph

halos are then added for distinguishing datasets. Similarly, Chung et al. [20] use outline color for distinguishing home or opposition teams in the visual analysis of rugby events.

When encoding the scale differences d_{scl} , the Tender glyph has the same hue and luminance for all octants, and it becomes harder to identify the differences in shapes (see Figure 4.5d). Switching to the dual hue colormap improves the combined perception of scale differences and shape differences (see Figure 4.5c). Therefore, the single hue colormap together with the halos is mainly for more accurate scale comparisons, while the dual hue colormap is for the combined perception of scale differences and shape differences simultaneously.

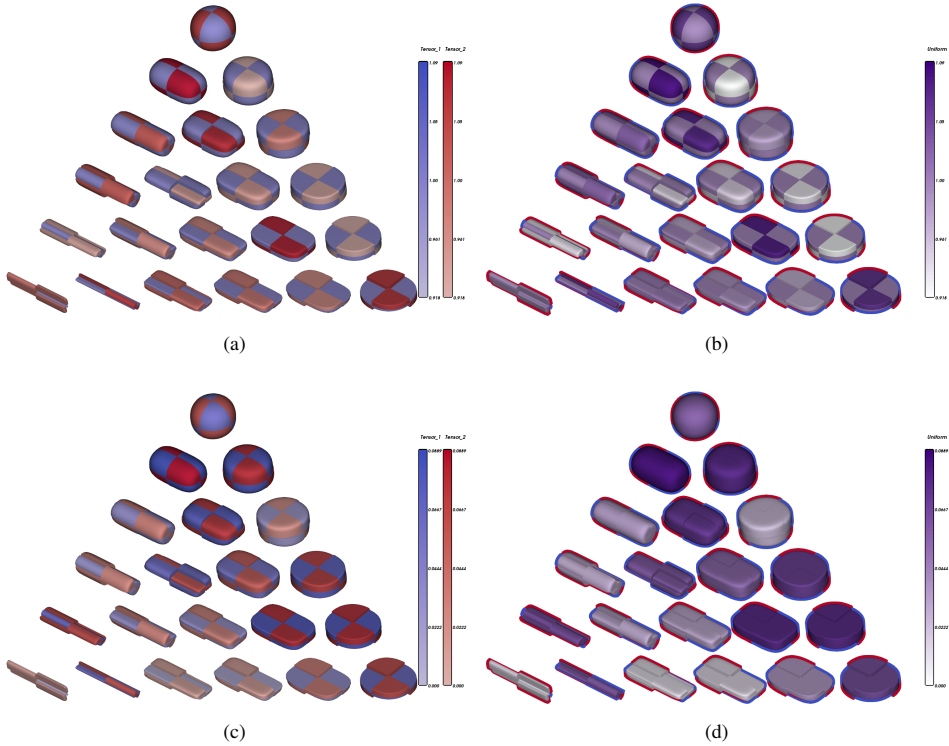


Figure 4.5: Two scale encoding schemes and two colormaps. (a) and (b) show that individual tensor scales $|D^{(i)}|$ are color-encoded with the dual hue and the single hue colormaps, respectively. This encoding scheme is suitable for local voxel-wise scale comparison. (c) and (d) show that tensor scale differences d_{scl} are color-encoded with the dual hue and the single hue colormaps, respectively. This encoding scheme is suitable for scale comparison across the field. The single hue colormap is for more accurate comparison while the dual hue colormap is for combined perception of scale and shape differences simultaneously.

Orientation Difference Encoding The orientation of a diffusion tensor is represented by its three eigenvectors. Straightforward visualization of two sets of eigenvectors as vectors has several disadvantages. The differences between them are not easily perceived due to their 3D nature. More importantly, eigenvectors cannot be uniquely determined for tensors of certain shape types (e.g., isotropic tensors). Therefore, we decide to explicitly encode the scalar-valued orientation difference d_{ori} , which is semantically mapped to the angle of a view-aligned sector.

In order to encode the orientation difference d_{ori} to the angle of sectors, we must define the mapping function. The following example is used to establish a relationship between them. Given two perfectly linear-anisotropic tensors $\mathbf{D}^{(1)} = \mathbf{D}^{(2)} = [1 \ 0 \ 0; 0 \ 0 \ 0; 0 \ 0 \ 0]$ (FA = 1.0, mode = 1.0) aligned with each other, if we gradually rotate tensor one $\mathbf{D}^{(1)}$ around vector (0, 0, 1) from 0° to 180° while keeping tensor two $\mathbf{D}^{(2)}$ fixed, the analytical behavior of the orientation difference is $d_{ori} = \sqrt{2}\sin(\theta)$ with θ the rotation angle, shown as the red curve in Figure 4.6. Therefore, the mapping function from orientation difference to the angle is defined as $\theta = 2\arcsin(d_{ori}/\sqrt{2})$. The scaling factor 2 allows the maximum orientation difference to map to a circle.

If tensors are not perfectly linear, their orientation difference for the same rotation angle will be smaller. As an example, the blue dots in Figure 4.6 indicate the values for $\mathbf{D}^{(1)} = \mathbf{D}^{(2)} = [0.99 \ 0 \ 0; 0 \ 0.1 \ 0; 0 \ 0.1 \ 0]$ (FA = 0.89, mode = 1.0). Figure 4.7a shows the tensors used in this experiment, and the corresponding sectors are shown in Figure 4.7b. Note that the left- and right-most sectors disappear since there are no orientation differences. Figure 4.7c illustrates the sectors for the case of two perfectly linear-anisotropic tensors.

Since the possible maximum of the normalized eigenvalue $\tilde{\lambda}_1$ is one, the radius of the sector is set to be slightly larger than one in order to maintain its visibility. Within certain ranges of viewpoints, those sectors could overlap. Thus, similar to Lie et al. [21], we add an eye-space fixed-width outline to support the depth perception and individual identification. Our final design is achieved by adding the checkerboard style superquadrics, scale-encoded color channels, and the view-aligned sector for orientation differences together, as shown in Figure 4.8a. The Tender glyphs are applied to visualize the differences between the same synthetic datasets used in Figure 4.1. The dual hue colormap is used to encode individual tensor scales.

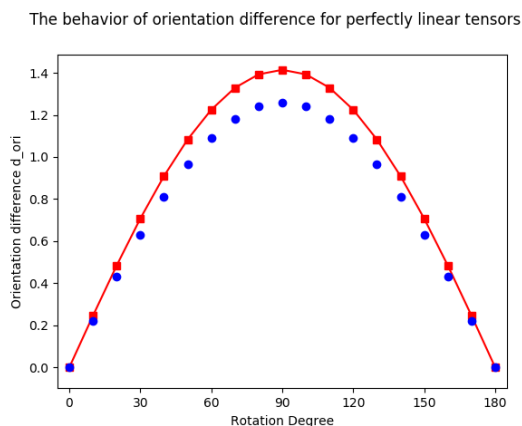


Figure 4.6: The red curve and squares together show the analytic behavior of d_{ori} between two perfectly linear-anisotropic tensors with respect to the rotation angle. Blue dots represent d_{ori} of other two linear-anisotropic tensors with reduced anisotropies.




Figure 4.7: The sectors for encoding orientation differences. (a) shows a group of linear tensors with gradually varying orientations (blue) and a reference tensor (red), and (b) shows the corresponding sectors for the orientation differences between each element tensor with the reference. (c) shows the results for two perfectly linear-anisotropic tensors in the same manner.

4.3. GLYPH VISUALIZATION

If the glyphs are placed throughout the whole volume, they probably occlude each other and cause visual clutter. Thus, glyph-based techniques are more suited for 2D and/or sparse visualization. In this work, we place the glyphs at one or more selected slices. Furthermore, users can restrict the display of the glyphs at anatomically meaningful structures of interest (e.g., corpus callosum). A smart glyph arrangement strategy [22] can remove the distracting effects of the regular sampling grid. However, this approach requires diffusion tensor interpolation. Therefore, we decide to place the glyphs at the original grid points.

Since the Tender glyph does not have an intrinsic orientation, we develop vari-

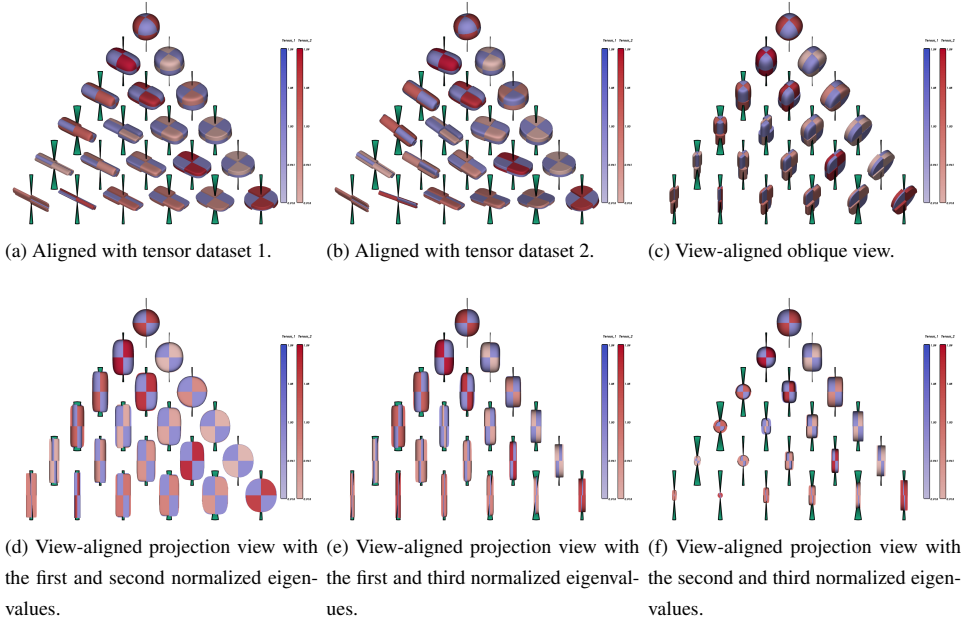


Figure 4.8: Various strategies for orienting the Tender glyphs in 3D space.

ous strategies for its orientation. The options include *a*) aligned with DTI dataset 1; *b*) aligned with DTI dataset 2; *c*) view-aligned oblique view; *d*) view-aligned projection of any pair of two eigenvalues. Aligning the glyphs with respect to one of the datasets is useful when preserving the original orientations is important. An oblique view is good to provide the first impression of the differences of all the three eigenvalues, while the projection of any two eigenvalues can support more accurate comparisons. How the Tender glyphs look like with these orientation strategies can be seen in Figure 4.8. Furthermore, we introduce an interaction for manipulating the glyphs, which is termed as *independent rotation*. The glyphs can be freely rotated around their own centers for a better perception of the shape differences without changing the current camera settings.

To facilitate exploration, we incorporate the visualization of a feature space. It uses as features: scale difference d_{scl} , shape difference d_{shp} , orientation difference d_{ori} , normalized tensor difference \tilde{d} and the Frobenius norm of the difference tensor $d_{Frobenius} = \|\mathbf{D}^{(1)} - \mathbf{D}^{(2)}\|$. The scatter plot matrix (SPLOM) approach is used to arrange the visualizations of these features (see Figure 4.9). Three histograms, placed at the

main diagonal, show the distributions of d_{scl} , d_{shp} , and d_{ori} , while the histograms of \bar{d} and $d_{Frobenius}$ are shown at the upper diagonal. Logarithmic scale is used for all the histograms. Three scatter plots at the lower triangular part show the relationship between any pair of the features d_{scl} , d_{shp} and d_{ori} . For each scatter plot, a density map generated from the whole dataset serves as the context, while the green dots drawn on top represent the distribution of the currently selected tensors. Linked views with filtering are used to support interactive exploration.

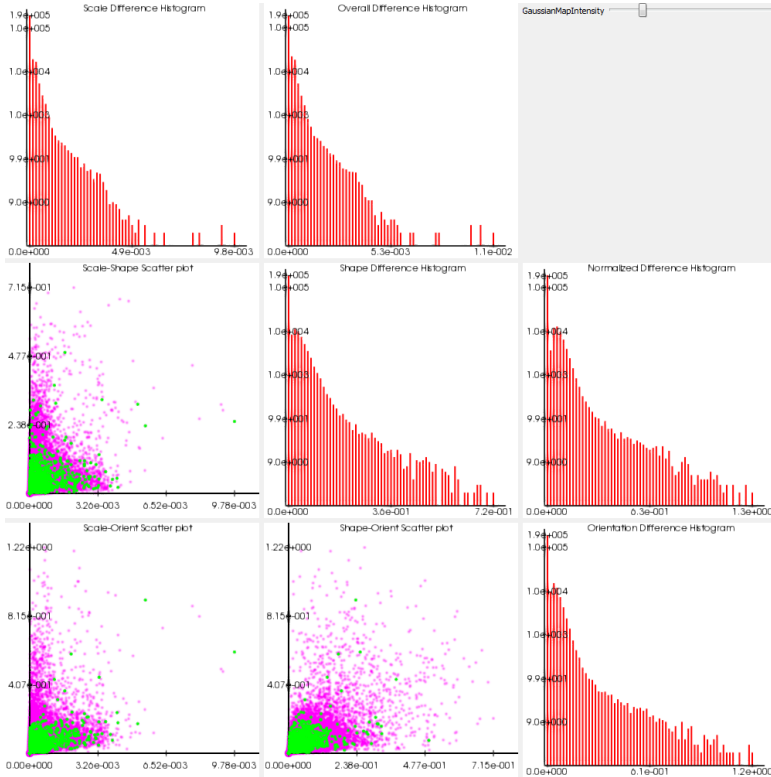


Figure 4.9: The feature space is composed of three scatter plots and five histograms. Three scatter plots are for the relationship of any pair of the features d_{scl} , d_{shp} , and d_{ori} . Five histograms are for d_{scl} , d_{shp} , d_{ori} , \bar{d} , and $d_{Frobenius}$.

4.3.1. IMPLEMENTATION

Our prototype is implemented in C++, combining VTK and OpenGL/GLSL for visualization, and Qt for the GUI design. The geometries of the checkerboard-style superquadrics are constructed via a precomputed palette of base superquadric shapes

[18]. In order to save graphics card memory, we only construct one eighth of each of the possible glyph geometries due to inherent symmetry. The total number of geometries for the precomputed palette is 231, which is enough for our goal. The vertex positions and the normals are stored in card memory. During rendering, the base geometry is used four times to generate the checkerboard-styled superquadrics for each tensor with its corresponding transformation matrix. Phong shading is employed to support the shape perception. The halos are rendered in a similar way to Schultz and Kindlmann [18]. The sector of the Tender glyph is built in the geometry and fragment shaders. A view-aligned quad is generated at each seeding point and texture coordinates are assigned to each corner for subsequent processing to get the expected angle. 4×4 multi-sampling is implemented to relieve the aliasing problem for the outline. Since the sector is analytically generated, the associated attributes for extra samples are computed in the fragment shader. This does not introduce any noticeable performance loss.

4.4. EVALUATION AND APPLICATION

4.4.1. USER STUDY

We conducted an initial user study in order to find out whether our requirements have been achieved by the Tender glyph. Since we intended to evaluate the effectiveness and efficiency of the glyph design on its own, the feature space was not included in the study. To be able to include more users, we abstracted from the diffusion-related concepts. We had 13 participants with various backgrounds: computer science (10; among them, 6 have experiences with computer graphics), DTI image analysis (2), and applied mathematics (1). The evaluation, which took 90 minutes on average per participant, consisted of three phases. In the first phase, we explained the three components on which the Tender glyph is based, i.e., scale, shape, and orientation. We also introduced the juxtaposition/superposition comparison techniques as well as superposition combined with transparency. 6 out of 13 did the test with all the three visualization methods while the rest did without the superposition plus transparency. Finally, we demonstrated our design in a live demo with our prototype. In the second phase, the participants themselves tried both the Tender glyph, the juxtaposed/superposed glyph, and the transparent superposed glyph to get used to the visualization as well as the

available interactions. In the third phase, 22 or 33 controlled tests were performed depending on whether the superposition with transparency test was added or not. During each test, the participants had to perform a comparison task and answer one to three questions as fast and accurately as possible. For all tests except the last three, two datasets each with 8 tensors were present, forming 8 tensor pairs. Participants needed to pick one or more pairs out of these 8 pairs as the answer. For example, they needed to find out which pairs have the smallest shape differences or the largest orientation differences. For the last three tests, they had to answer similar questions but with relatively large fields (9×9). In order to correctly measure their accuracy in answering these questions, we allowed them to give no answer when they were uncertain. One of the three glyph-based methods was randomly assigned to visualize the differences between the test datasets. None of the test datasets was repeated for the same participant. Each test dataset appeared only once during the whole phase. In the end, open questions designed to collect their impression on the utility of our design were asked. An overall ranking on a scale from 0 to 10 about the usefulness of the juxtaposition/superposition, superposition with transparency, and the Tender glyph was asked as the last question.

There were 8, 4 and 8 questions asked for the comparison of orientation, shape, and scale, respectively. The accuracy and efficiency (i.e., time) results of orientation difference comparison tests are shown as two boxplots in Figure 4.10a and 4.10d. In general, they achieved better results with much less time using the Tender glyph. The results confirm that the perception of orientation differences in terms of angles is easy to spot and interpret. The accuracy and efficiency of shape-related questions are shown in Figure 4.10b and 4.10e. In terms of accuracy, the results indicate that the Tender glyph provides more accurate results, while in terms of efficiency there are no significant differences among all the methods. Four participants stated in the open questions that all the methods worked equally well for the shape comparison, so they had no preference. But it is interesting to observe that they improved the accuracy and spent less time with the Tender glyph, despite their perception. There was one participant who correctly answered all the shape-related questions with both transparent superposition and the Tender glyph. But this participant spent less time with the Tender glyph (140s < 177s), and also mentioned in the open questions that the checkerboard design provides stronger visual cues for shape differences, making them easier to recognize.

The results for scale comparison are shown in Figure 4.10c and 4.10f. In terms of accuracy, the results do not clearly indicate whether Tender glyphs or superposition with transparency is superior. The efficiency is clearly better with superposition plus transparency. 5 out of 6 participants mentioned that for small field comparisons they preferred superposition plus transparency since they could roughly estimate the scale differences, as well as decide which tensor has a large or small scale for a specific pair simultaneously. Using the Tender glyphs to obtain the same information requires extra interactions, which causes a more time consuming exploration. For large field comparisons, they all agreed that it is easier with the colormap-based encoding scheme. In practice, many comparisons would be performed on large fields, as shown in the following two case studies. We also realized that the complexity of the interface and the interaction with the prototype creates a steep learning curve for novice users of the Tender glyph. An interesting statement from one participant with DTI image analysis background was that the dual hue colormap is counter-intuitive since blue represents low values in the rainbow colormap that is commonly used in his research.

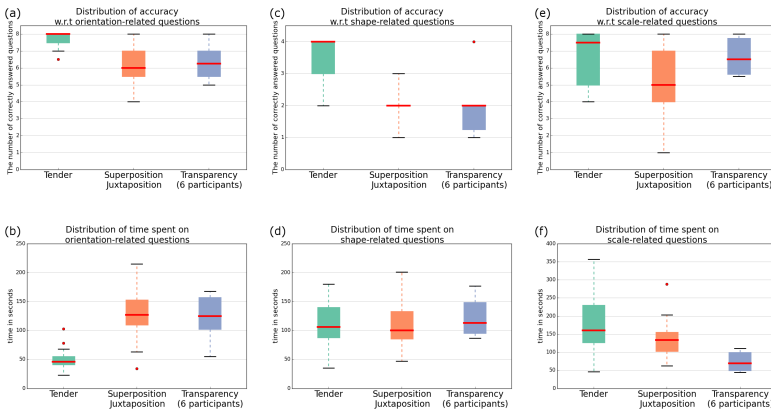


Figure 4.10: The results of the user study in terms of accuracy and efficiency. (a), (c), and (e) represent results in terms of accuracy, while (b), (d), and (f) represent results in terms of efficiency for comparison tasks related to orientation, shape, and scale differences, respectively.

The results of the questions, in relation to the preferences, are as follows. The average ranking for superposition is 4.5. This method could cause occlusions, which makes the comparisons more difficult. They gave 5 on average for juxtaposition. In this way, they could perceive full information which is helpful for rough comparisons

in a small field. For accurate comparisons, especially in a large field, it is difficult to mentally match the corresponding tensor glyphs for comparison. They gave 6.2 on average for superposition with transparency. The advantage of this method is that the occlusion problem and the matching problem are relieved. But there are no strong visual cues for shape comparison. Furthermore, for orientation comparison they had to mentally imagine how to align two glyphs in order to deduce the orientation difference. They gave 7.8 on average for the Tender glyph. With the Tender glyph, the orientation comparison is easy to perform. The Tender glyph also provides stronger visual cues for shape comparison than other glyph-based methods. For scale comparison, several participants suggested that we should simplify the user interface to make it simpler to learn. Furthermore, one participant with DTI image analysis background suggested that the Tender glyph is useful when to avoid presenting and looking at several tensor scalar images next to each other. With the Tender glyph, one image, containing all the information about the tensor differences, would be sufficient.

4.4.2. CASE STUDY

One of the main applications of DTI is the imaging of white matter in the brain. DTI has the potential to provide better understanding of brain connectivity, and improve diagnosis and treatment of diseases such as HIV [23]. Comparison of diffusion tensor fields is part of the analysis to achieve these goals, e.g., to compare the results of different acquisition settings to optimize acquisition sequences, or to compare datasets to find markers that can distinguish pathological tissues.

The main idea of the Tender glyph is to improve the ability to visually reveal local tensor differences as well as large scale patterns in the datasets. Furthermore, by linked selection of the associated feature space, interesting pattern identification is facilitated. In this section, we illustrate the application of Tender glyphs in two scenarios where two DTI datasets are compared.

B-value is an important parameter in DTI acquisition that defines a trade-off between signal-to-noise ratio and contrast. It is often unknown what exact output a b-value change will produce. Two DTI datasets of the same subject, acquired with a b-value of 1000 and 2000 respectively, are compared. The size is $80 \times 80 \times 38$ at $2.875 \times 2.875 \times 3mm$ resolution. A coronal slice is shown in Figure 4.11(a). Blue corresponds to the b-value 1000 DTI dataset, while red represents the b-value 2000

one. It can be immediately observed that only the left bottom and right bottom areas of the slice show large differences in orientation. Figure 4.11(b) shows the superposition of the superquadric glyphs without normalization in the area marked with a rectangle in Figure 4.11(a). Due to large variations in tensor scales, most of the red glyphs are enclosed by the blue ones. Figure 4.11(c) shows the superposition of normalized glyphs. Since they occlude each other, it is hard to distinguish their shape and orientation differences. Moreover, the exact tensor scale differences are impossible to deduce after normalization. Figure 4.11(d) visualizes the tensor differences with the Tender glyphs. The voxels with the largest orientation differences (i.e., the largest angles) can be quickly spotted via the sectors of the Tender glyphs, indicating their pre-attentive capabilities. Most of the angles of the sectors are very small, which reveals that the orientation differences are quite little within this area. The individual tensor scales are encoded via the single hue colormap. By comparing lightness, the extent of scale differences can be estimated. The halos around the glyphs help to identify the corresponding DTI dataset. The single hue colormap together with the halos clearly show that the tensors from the DTI dataset of b-value 1000 have larger scales than that from the b-value 2000 dataset, which was expected. The shape differences are conveyed in terms of the checkerboard style edge differences. With the single hue colormap, it is difficult to perceive shape differences for tensors of similar scales. Figure 4.11(e) shows the Tender glyphs with the dual hue colormap which is more suitable for distinguishing two tensors with similar scales. A closeup of the Tender glyphs within the area marked as the rectangle in Figure 4.11(e) is shown in Figure 4.11(f) in the view-aligned projection style of the first and second normalized eigenvalues. Even the subtle differences, for example for the linear tensors, can be effectively identified.

The feature space selection can be used to find regions with large differences or outliers. The locations with both large shape and orientation differences are selected, as shown in Figure 4.12a, while the corresponding feature space is shown at the top left corner. It can be seen from Figure 4.12b that most of the Tender glyphs are shown at the bottom of the DTI volume, which are probably in uninteresting regions for analysis. However, there are two groups of glyphs located inside the volume, as marked by the white rectangles in Figure 4.12a. Figure 4.12b shows a view of some of these normalized glyphs located on one coronal slice. Figure 4.12c presents a view for this region without encoding the tensor scales. The shape differences become more obvious

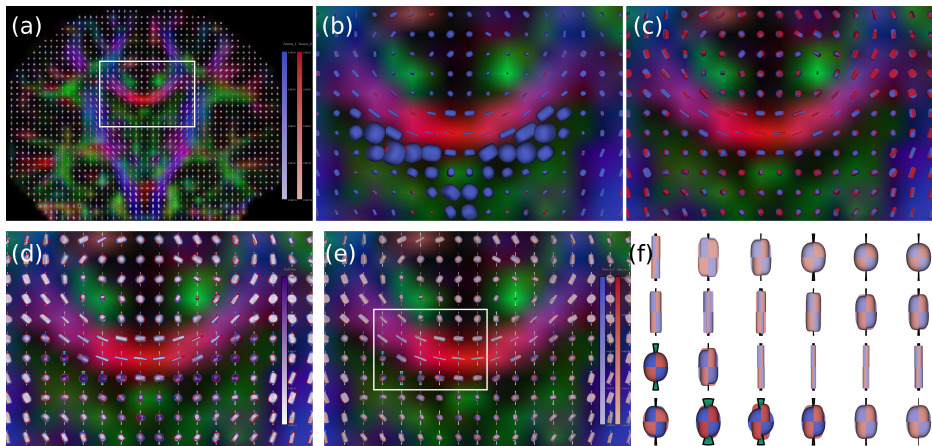




Figure 4.11: The application of the Tender glyphs to compare two DTI datasets acquired with different b-values. (a) Tender glyphs on a coronal slice. (b) Superposition of superquadric tensor glyphs. (c) Superposition of superquadric glyphs after tensor normalization. (d) Tender glyphs with the single hue colormap. (e) Tender glyphs with the dual hue colormap. (f) Tender glyphs displayed in the view-aligned projection style with the first and second normalized eigenvalues.


via the Tender glyphs, and they can be more accurately assessed.



(a) Find interesting areas with both large shape and orientation differences via feature space.



(b) Superposed superquadric tensor glyphs.



(c) The Tender glyphs clearly show the orientation and shape differences in detail.

Figure 4.12: Feather space filtering together with the Tender glyphs permits the detailed comparison at interesting areas.

As another example of DTI dataset comparison the Tender glyph is applied to compare a DTI dataset from a healthy subject and from an HIV-infected subject. The datasets were pre-registered through FSL¹. The size of each dataset is $112 \times 112 \times 55$ with an isotropic resolution of 2mm. Two regions of the corpus callosum (CC) on a sagittal slice are selected as the ROIs shown as white rectangles in Figure 4.13(a).

¹<http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>

Blue represents the DTI dataset of a healthy subject, while red is for the DTI dataset from an HIV-infected subject. Figure 4.13(b) clearly shows the orientation differences, compared to the superposed tensor glyphs shown in Figure 4.13(c). Voxels with consistent orientation differences are more efficiently recognized in the posterior region of the CC. Figures 4.13(d), (e), (f) together show the visualization results for the anterior region of the CC. From Figures 4.13(e) and (f) we can see that the orientations in the anterior region do not change a lot, but their shapes do. The Tender glyphs displayed in the view-aligned style, as shown in Figure 4.13(f), further confirm that tensors of the HIV dataset are less linearly anisotropic in the indicated region. Further studies are necessary to see whether these findings remain consistent between larger populations.

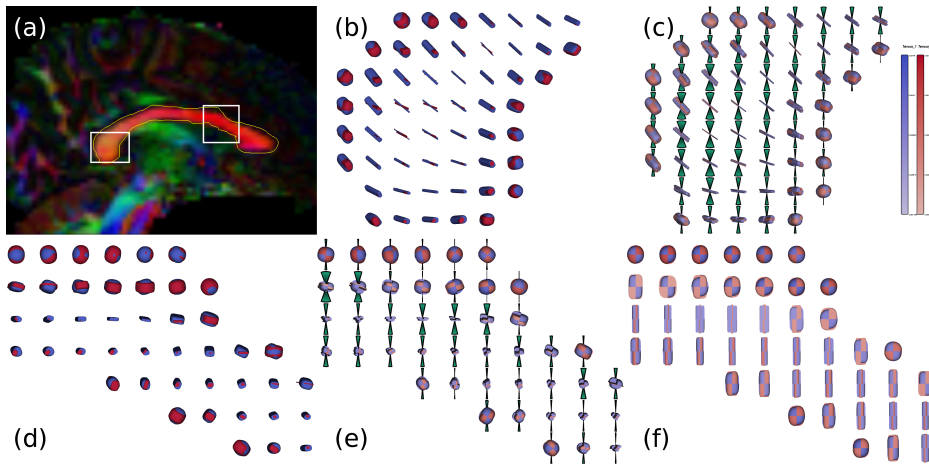


Figure 4.13: The Tender glyphs clearly reveal the change of dominant tensor differences in different regions of CC. (a) Two ROIs selected on the CC. (b) Superposed superquadric tensor glyphs at the posterior area of CC. (c) Tender glyphs reveal consistent orientation differences at the posterior area of CC. (d) Superposed superquadric tensor glyphs at the anterior area of CC. (e) Tender glyphs at the anterior area of CC, aligned with the blue dataset. (f) Tender glyphs at the anterior area of CC in the view-aligned projection style.

4.5. CONCLUSIONS

We present a glyph-based comparative visualization for diffusion tensor datasets. We decompose the overall tensor differences into differences in shape, scale, and orientation, which have biological interpretations for the underlying tissue. We design a novel glyph, i.e., the Tender glyph to encode and visually present them in an easy-to-interpret way. The Tender glyph is composed of two superquadric glyphs arranged in

a checkerboard style that is efficient for the visual comparison of tensor shapes. The individual tensor scales or the scale differences are color-encoded, and we also design two colormaps, each of which has its own advantages for certain purposes. The arcs encode the orientation differences, which are intuitive and pre-attentive for perception. We build a feature space that is helpful for the interactive exploration and selection of relevant features. We present an initial user study that shows that the Tender glyph allows a more accurate and effective analysis for orientation and shape differences, although it does not improve the time performance in all tasks. Two cases based on brain DTI datasets illustrate the potential of the Tender glyph in comparative analysis of real-world datasets.

4

In the future, we would like to carry out a more exhaustive user study in a specific DTI application domain as well as a general evaluation of the Tender glyph in the way presented in Demiralp et al. [12].

Furthermore, the most well-known problem for the single tensor model is that it fails to model complex fiber configurations inside a voxel such as fiber crossings. More sophisticated modeling techniques such as HARDI can be used to describe those complex situations. How to extend the design idea of Tender glyph to HARDI is an open problem. The Tender glyph is currently only for pair-wise comparison. Though in the case study with the HIV dataset we compare two individual subjects from two groups, the Tender glyph can be applied to visualize the differences between the corresponding group means for inter-group comparison. Furthermore, in practice there are situations where multiple datasets need to be compared rather than just two. Therefore, another open problem is how to extend the Tender glyph for comparison beyond two datasets. The underlying uncertainty of DTI [24], which is caused by the noise and the incomplete modeling, is another motivation for multiple DTI datasets comparison.

REFERENCES

- [1] C. Zhang, T. Schultz, K. Lawonn, E. Eisemann, and A. Vilanova, *Glyph-based comparative visualization for diffusion tensor fields*, IEEE Transactions on Visualization and Computer Graphics **22**, 797 (2016).
- [2] D. B. Bashat, L. B. Sira, M. Graif, P. Pianka, T. Hendler, Y. Cohen, and Y. Assaf, *Normal white matter development from infancy to adulthood: Comparing diffu-*

- sion tensor and high b value diffusion weighted MR images*, *Journal of Magnetic Resonance Imaging* **21**, 503 (2005).
- [3] S. Basu, T. Fletcher, and R. Whitaker, *Rician noise removal in diffusion tensor MRI*, in *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2006) pp. 117–125.
- [4] D. K. Jones, L. D. Griffin, D. C. Alexander, M. Catani, M. A. Horsfield, R. Howard, and S. C. Williams, *Spatial normalization and averaging of diffusion tensor MRI data sets*, *NeuroImage* **17**, 592 (2002).
- [5] Ø. Bergmann, G. Kindlmann, A. Lundervold, and C.-F. Westin, *Diffusion k-tensor estimation from Q-ball imaging using discretized principal axes*, in *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2006) pp. 268–275.
- [6] M. J. Da Silva, S. Zhang, C. Demiralp, and D. H. Laidlaw, *Visualizing the differences between diffusion tensor volume images*, in *International Society for Magnetic Resonance in Medicine (ISMRM) Workshop on Diffusion MRI* (2000) pp. 237–238.
- [7] C.-F. Westin, S. E. Maier, B. Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis, *Image processing for diffusion tensor magnetic resonance imaging*, in *MICCAI* (1999) pp. 441–452.
- [8] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVC Symposium on Visualization* (2004) pp. 147–154.
- [9] R. Stokking, I. G. Zubal, and M. A. Viergever, *Display of fused images: methods, interpretation, and diagnostic improvements*. *Seminars in Nuclear Medicine* **33**, 219 (2003).
- [10] P. Kok, M. Baiker, E. Hendriks, F. H. Post, J. Dijkstra, C. W. Löwik, B. P. Lelieveldt, and C. P. Botha, *Articulated planar reformation for change visualization in small animal imaging*, *IEEE Transactions on Visualization and Computer Graphics* **16**, 1396 (2010).

- [11] G. Kindlmann, D. B. Ennis, R. T. Whitaker, and C.-F. Westin, *Diffusion tensor analysis with invariant gradients and rotation tangents*, IEEE Transactions on Medical Imaging **26**, 1483 (2007).
- [12] Ç. Demiralp, C. E. Scheidegger, G. L. Kindlmann, D. H. Laidlaw, and J. Heer, *Visual embedding: A model for visualization*, IEEE Computer Graphics and Applications **34**, 10 (2014).
- [13] A. Kratz, C. Auer, M. Stommel, and I. Hotz, *Visualization and analysis of second-order tensors: Moving beyond the symmetric positive-definite case*, Computer Graphics Forum **32** (2013).
- [14] T. H. J. M. Peeters, P. R. Rodrigues, A. Vilanova, and B. M. ter Haar Romeny, *Analysis of distance/similarity measures for diffusion tensor imaging*, in *Visualization and Processing of Tensor Fields: Advances and Perspectives*, edited by D. Laidlaw and J. Weickert (Springer Berlin Heidelberg, Berlin, Heidelberg, 2009) pp. 113–136.
- [15] P. T. Fletcher and S. Joshi, *Riemannian geometry for the statistical analysis of diffusion tensor data*, Signal Processing **87**, 250 (2007), tensor Signal Processing.
- [16] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, *Log-Euclidean metrics for fast and simple calculus on diffusion tensors*, Magnetic Resonance in Medicine **56**, 411 (2006).
- [17] G. Kindlmann, R. San José Estépar, M. Niethammer, S. Haker, and C.-F. Westin, *Geodesic-loxodromes for diffusion tensor interpolation and difference measurement*, in *MICCAI* (2007) pp. 1–9.
- [18] T. Schultz and G. L. Kindlmann, *Superquadric glyphs for symmetric second-order tensors*, IEEE Transactions on Visualization and Computer Graphics **16**, 1595 (2010).
- [19] R. Borgo, J. Kehrer, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen, *Glyph-based visualization: Foundations, design guidelines, techniques and applications*, in *Eurographics State of the Art Reports* (2013) pp. 39–63.

- [20] D. H. Chung, P. A. Legg, M. L. Parry, R. Bown, I. W. Griffiths, R. S. Laramée, and M. Chen, *Glyph sorting: Interactive visualization for multi-dimensional data*, *Information Visualization* **14**, 76 (2015).
- [21] A. E. Lie, J. Kehrer, and H. Hauser, *Critical design and realization aspects of glyph-based 3D data visualization*, in *Spring Conference on Computer Graphics* (2009) pp. 19–26.
- [22] G. Kindlmann and C.-F. Westin, *Diffusion tensor visualization with glyph packing*, *IEEE Transactions on Visualization and Computer Graphics* **12**, 1329 (2006).
- [23] Y. Wu, P. Storey, B. Cohen, L. Epstein, R. Edelman, and A. Ragin, *Diffusion alterations in corpus callosum of patients with HIV*, *American journal of neuro-radiology* **27**, 656 (2006).
- [24] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith, *Characterization and propagation of uncertainty in diffusion-weighted MR imaging*, *Magnetic Resonance in Medicine* **50**, 1077 (2003).

5

OVERVIEW + DETAIL VISUALIZATION FOR AN ENSEMBLE OF DIFFUSION TENSOR FIELDS

*In all fighting, the direct method may be used for joining battle,
but indirect methods will be needed in order to secure victory.*

Sun Tzu

凡战者，以正合，以奇胜。

孙武

This chapter is based on the following publication: “Overview + Detail Visualization for Ensembles of Diffusion Tensors”, **Changgong Zhang**, Matthan W.A. Caan, Thomas Höllt, Elmar Eisemann, and Anna Vilanova. In *Computer Graphics Forum (Proceedings of EuroVis)*, 36(3):121–132, 2017 [1].

In the previous chapter, we present a glyph-based comparative visualization for two tensor fields. In this chapter, we increase the number of tensor fields to be compared from two to multiple, i.e., an ensemble of tensor fields, and provide an overview + detail framework to facilitate its visual exploration.

5.1. INTRODUCTION

In neuroscientific studies, it is necessary to summarize a group of DTI datasets, normally called cohorts or ensembles, in order to locate variations between subjects that might be caused by natural variations [2] or uncertainty [3].

However, only few available approaches focus on the diffusion tensor in its entirety [4–6]. Basser and Pajevic [4, 5] propose the use of the fourth-order covariance tensor based on the assumption that the set of diffusion tensors follows a multivariate normal distribution. The fourth-order covariance tensor represents deviations from the mean tensor, which is built by component-wise averaging. Basser and Pajevic [5] use radial glyphs to visualize the fourth-order tensor and its six orthogonal second-order eigentensors or eigenmodes. The six mutual orthogonal eigentensors represent all modes of tensor variation. However, this visual encoding is difficult to interpret. In recent work, Abbasloo et al. [6] combine slice views, volume rendering, and superquadric glyphs [7] to visualize the fourth-order covariance tensor at multiple levels of detail. In order to facilitate the interpretation of the covariance tensor, they use the decomposition framework proposed by Kindlmann et al. [8]. The framework is based on the gradients of meaningful tensor invariants. They visualize the six orthogonal eigentensors, which contain a mixture of different tensor intrinsic properties, through animation. Eigentensors allow the identification of correlations but also provide potentially misleading interpretations in some situations (e.g., in the cases shown later in Section 5.7.1).

In this chapter, we integrate concept of the intrinsic properties (i.e., scale, shape, and orientation) of diffusion tensors into the analysis and visualization of DTI ensembles to facilitate their interpretation. A major challenge for this approach is to deal appropriately with the inherently coupled relation between shape and orientation. The main contribution of this work is an integrated overview + detail visual analysis framework for the exploration of DTI ensembles that does not assume any specific distribution of the underlying data.

- We propose a representative tensor that is derived by separately aggregating the intrinsic properties, instead of using the Euclidean mean that treats these properties collectively.
- Variation within the tensor ensemble is defined based on the pair-wise difference measures of the three tensor intrinsic properties (see Section 4.2.2).
- We extend two existing glyph-based visualizations [7, 9] to encode variation uncovered in the ensemble.
- We present a detail visualization that enables intuitive exploration of the tensor ensemble, separately for the tensor intrinsic properties, without assuming a specific distribution.

5.2. STATISTICAL SUMMARY

In order to aggregate or summarize an ensemble of diffusion tensors, we resort to statistical tools. In this section, we introduce our statistical analysis approaches.

We define the space of symmetric tensors in \mathbb{R}^3 as Sym_3 and the corresponding tensor field as $\mathbf{F}: \mathbb{R}^3 \rightarrow \text{Sym}_3$. An ensemble of n diffusion tensor fields is denoted by $\{\mathbf{F}_i\}_{i=1}^n$. Since our focus is voxel-wise analysis, the main object of interest is a set of tensors at a specific location \mathbf{x} within the field, denoted by $\{\mathbf{F}_i(\mathbf{x})\}_{i=1}^n$. For simplicity, we denote the set as $\{\mathbf{D}_i\}_{i=1}^n$, and refer to it as a *tensor ensemble*. Notice that most of the solutions presented in this work are also valid for ensembles of any second-order positive-definite tensors independent of their origins.

5.2.1. DIFFUSION TENSOR DECOMPOSITION

The exact way of decomposing a tensor is dependent on the application. For diffusion tensors, an informative and intuitive decomposition consists of three intrinsic properties (i.e., scale, shape, and orientation). Each of these properties carries a specific interpretation about the underlying fiber structures. Decomposition into these three intrinsic properties is well studied in the previous chapter and Ennis and Kindlmann [10]. Hence, we briefly introduce the decomposition as proposed in Section 4.2.

A symmetric and positive-definite diffusion tensor \mathbf{D} can be eigen-decomposed into three sorted real eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ where λ_1 is often referred to as

the *main* or *major* eigenvalue. The orthonormal eigenvectors are ordered accordingly $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$. The tensor trace, $\text{tr}(\mathbf{D})$, is defined as the sum of the eigenvalues $\text{tr}(\mathbf{D}) = \sum_{j=1}^3 \lambda_j$, and is proportional to a commonly used tensor invariant called mean diffusivity [11]. Finally, the sorted normalized eigenvalues $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$ are given by $(\lambda_1, \lambda_2, \lambda_3) / \text{tr}(\mathbf{D})$.

A diffusion tensor has six degrees-of-freedom (DOFs) in total. Tensor scale, one DOF, can be represented by the trace $\text{tr}(\mathbf{D})$ and has a range of $[0, \infty)$. It indicates the overall amount of diffusion. There is more diffusion in regions without fibers (e.g., ventricle) than with fibers (e.g., corpus callosum). Tensor shape, two DOFs, can be represented by the sorted normalized eigenvalues $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$ with range $[0, 1]$, reflecting the underlying fiber configurations (e.g., a single fiber population or two crossing fiber populations). Tensor orientation is given by the set of eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$. The major eigenvector \mathbf{e}_1 is empirically assumed to be aligned with the underlying fiber pathways. Shape and orientation are inherently coupled. For example, if the tensor shape is isotropic (i.e., three equal eigenvalues), the orientation is completely undefined. This fact complicates the decomposition and the subsequent statistical analysis.

5.2.2. DIFFUSION TENSOR ENSEMBLE AGGREGATION

We characterize a tensor ensemble by *separately* deriving the mean and standard deviation for scale, shape, and orientation. Given its special characteristics, we also propose an alternative that does not rely on the mean and standard deviation of the orientation. Instead, we estimate the diffusion orientation distribution function (dODF) according to the normal distributions implied by each tensor of the ensemble. This approach enables the depiction of multiple orientations within the ensemble.

TENSOR ENSEMBLE MEAN

We first consider the special case of calculating the mean on an ensemble consisting of just two tensors. Here, the mean tensor can simply be viewed as an interpolation between two tensors. Kindlmann et al. [12] illustrate that Euclidean as well as Log-Euclidean interpolation cannot guarantee a monotonic interpolation of the tensor shape invariants, and as such can introduce spurious tensor shapes. Extending the Euclidean mean to more than two ensemble members corresponds to computing a component-wise average. Although it has been used in previous work [4, 6], the Euclidean mean

has undesired properties [12]. Let us consider an ensemble of linear anisotropic tensors with exactly the same eigenvalues but different eigenvectors, as shown in Figure 5.1a. It can be observed that the Euclidean mean tensor (Figure 5.1b) has a spurious shape that is not present in the original ensemble. Such a shape can mislead the analyst to make false assumptions about the underlying ensemble. Therefore, we propose to derive the mean tensor by computing the average of the three intrinsic properties *separately* and then combine them to form the mean tensor. Therefore, the characteristics are preserved and can be analyzed in the mean tensor. We denote this new mean tensor as $\bar{\mathbf{D}}$.

The trace of $\bar{\mathbf{D}}$ is defined as the average of individual tensor traces,

$$\text{tr}(\bar{\mathbf{D}}) = \frac{1}{n} \sum_{i=1}^n \text{tr}(\mathbf{D}_i). \quad (5.1)$$

Note that $\text{tr}(\bar{\mathbf{D}})$ is identical to the trace of the Euclidean mean.

Because of normalization, sorting, and positive definiteness, all possible eigenvalue tuples $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$, representing the shape, form a triangle in \mathbb{R}^3 (see Figure 5.2). Therefore, the shape of $\bar{\mathbf{D}}$ is defined as

$$(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)^{\bar{\mathbf{D}}} = \frac{1}{n} \sum_{i=1}^n (\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)^{\mathbf{D}_i}. \quad (5.2)$$

The superscript indicates the corresponding tensor from which the normalized eigenvalues are derived. Our idea is similar to Gahm et al. [13], but we aim at summarizing the tensor ensemble members while they focus on pair-wise tensor interpolation. Furthermore, we work directly with the eigenvalues instead of specific tensor invariants.

So far we have defined scale and shape of the mean tensor, which represent the mean diffusion amount and mean diffusion shape, respectively. However, the definition of the mean orientation is not as straightforward as the mean scale and shape because




Figure 5.1: Two types of mean tensors for (a) 10 linear anisotropic tensors with the same eigenvalues (0.7, 0.15, 0.15) but different eigenvectors. (b) is the Euclidean mean and (c) is the ensemble mean calculated according to Equation (5.3).




Figure 5.2: Tensor shape space is defined by the sorted and normalized eigenvalues $\bar{\lambda}_j$. Three glyphs are placed at the corners as reference, colored according to their $\bar{\lambda}_j$.

5

the eigenvectors are not always well-defined (i.e., in the case of two or three equal eigenvalues), and do not have a unique sign. In order to construct the complete mean tensor, a set of eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)^{\bar{\mathbf{D}}}$ is required to orient the mean tensor in the 3D domain. Here, we use the eigenvectors of the Euclidean mean $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)^{\text{Eu}}$ as the eigenvectors of our mean tensor, as suggested by Gahm et al. [13]. The orientation of the mean tensor can provide anatomical reference that is relevant for the understanding of the context. However, strongly diverging orientations in the ensemble can get lost. Therefore, the orientation has to be analyzed with care.

Finally, the separate components are assembled into the mean tensor as,

$$\bar{\mathbf{D}} := \text{tr}(\bar{\mathbf{D}}) \sum_{j=1}^3 \bar{\lambda}_j^{\bar{\mathbf{D}}} \mathbf{e}_j^{\bar{\mathbf{D}}} \otimes \mathbf{e}_j^{\bar{\mathbf{D}}}. \quad (5.3)$$

By separately handling scale, shape, and orientation, we effectively avoid the introduction of tensor intrinsic properties that are not present in the tensor ensemble into the mean tensor. Figure 5.1c shows the mean tensor of the ensemble introduced in Figure 5.1a, calculated according to Equation (5.3). This new mean tensor preserves the shape of the tensors in the ensemble.

However, $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)^{\text{Eu}}$ can be ill-defined and do not give a good representation of the underlying ensemble orientation. For example, consider a second ensemble of linear tensors, consisting of orthogonal pairs in Figure 5.3a. The Euclidean mean in Figure 5.3b of this ensemble has a planar shape. Thus its eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)^{\text{Eu}}$ cannot be uniquely defined. The mean tensor, according to Equation (5.3), will have a clear linear shape, but an *arbitrary* and *meaningless* orientation.

Alternatively, we propose to summarize the orientation based on the diffusion orientation distribution function (dODF) that can be analytically expressed [14] as

$$p(\mathbf{r}) = \frac{1}{\sqrt{(4\pi\tau)^3 \det(\mathbf{D})}} e^{-\frac{1}{4\tau} \mathbf{r}^T \mathbf{D}^{-1} \mathbf{r}},$$

$$\text{dODF}(\mathbf{u}) = \int_0^\infty P(r\mathbf{u}) r^2 dr = \frac{1}{4\pi \sqrt{\det(\mathbf{D})} (\mathbf{u}^T \mathbf{D}^{-1} \mathbf{u})^{\frac{3}{2}}}. \quad (5.4)$$

$p(\mathbf{r})$ is the probability density function of a 3D normal displacement distribution, which is the fundamental assumption underlying DTI. It represents the probability of a water molecule ending up with displacement vector \mathbf{r} after a certain amount of time τ . dODF is the radial integral of the normal distribution along a given direction \mathbf{u} (i.e., the unit vector). Furthermore, $\det(\bullet)$ is the tensor determinant and $(\bullet)^{-1}$ is the tensor inverse. The dODF integrates to one over the unit sphere \mathbb{S}^2 , i.e., $\int_{\mathbb{S}^2} \text{dODF}(\mathbf{u}) d\mathbf{u} = 1$. It is dimensionless, and invariant with respect to $\det(\mathbf{D})$ [14]. Note that dODF responds to the tensor orientation as well as the tensor shape. The mean dODF, denoted $\overline{\text{dODF}}$, is defined as the average of individual dODFs,

$$\overline{\text{dODF}}(\mathbf{u}) = \frac{1}{n} \sum_{i=1}^n \text{dODF}_i(\mathbf{u}). \quad (5.5)$$

Similar to HARDI (e.g., Q-Ball imaging [15]), which models multiple diffusion orientations for the case that a single voxel contains multiple fiber orientations (e.g., fiber crossing [16]), the $\overline{\text{dODF}}$ models multiple orientations within the ensemble, and as such provides a faithful representation of the actual diffusion directions. Figure 5.4a shows the $\overline{\text{dODF}}$ glyph for the tensor ensemble presented in Figure 5.3a, preserving the two major orientations.




Figure 5.3: (a) A group of linear tensors with two distinct orientations. (b) The Euclidean mean tensor. A mean tensor generated according to Equation (5.3) would comprise the same scale and shape as the tensors in a) but with a random orientation (thus not shown).

TENSOR ENSEMBLE VARIATIONS

To quantify the variation in the ensemble separately for scale, shape, and orientation, we generalize the standard deviation σ to the pair-wise tensor difference measures (see Section 4.2.2), denoted as d_{scale} , d_{shape} , and $d_{\text{orientation}}$, respectively. The pair-wise difference measures have several desirable properties, especially in relation to the link between shape and orientation. The orientation difference measure $d_{\text{orientation}}$ is designed to smoothly converge to zero as one of the compared tensors becomes isotropic, ensuring continuity. Furthermore, the difference measures are considerably faster to compute, compared to the *geodesic-loxodromes* approach [12]. Using the pair-wise differences, the standard deviation for each of the intrinsic tensor properties is defined as

$$\sigma_p = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (d_p(\overline{\mathbf{D}}, \mathbf{D}_i))^2}, \tag{5.6}$$

where p is one of the three intrinsic properties {scale, shape, orientation}. The statistical summary contains a mean tensor and three scalar-valued quantities that represent the variation within each individual property. Although conceptually similar to the *aggregate covariance measures* proposed by Kindlmann et al. [8], the {scale, shape, orientation} variations are defined independently of the covariance tensor. Furthermore, since both d_{shape} and $d_{\text{orientation}}$ are bounded, the corresponding variations σ_{shape} and $\sigma_{\text{orientation}}$ are also bounded.

The standard deviation for dODFs, describing the variation in the diffusion prob-




Figure 5.4: (a) The corresponding $\overline{\text{dODF}}$ glyph for the tensor ensemble in Figure 5.3a. A conventional direction-encoded colormap is used. (b) The glyph with the threshold set to 60% of the maximum variation. (c) The glyph of (b) min-max normalized to enhance the direction with a large diffusion probability density.

ability density in a given direction \mathbf{u} , is then given by

$$\sigma_{\text{dODF}}(\mathbf{u}) = \sqrt{\frac{1}{n-1} \sum_{i=1}^n \left(\overline{\text{dODF}}(\mathbf{u}) - \text{dODF}_i(\mathbf{u}) \right)^2}. \quad (5.7)$$

The σ_{dODF} is conceptually similar to the *variance in Apparent Diffusion Coefficient* (Equation (13) in Abbasloo et al. [6]), and provides a direct interpretation compared to the fourth-order covariance tensor.

Both $\overline{\text{dODF}}$ and σ_{dODF} are functions defined on the unit sphere \mathbb{S}^2 . Hence, $\overline{\text{dODF}}$ and σ_{dODF} can be approximated with spherical harmonics of up to fourth-order [17], resulting in two sets of spherical harmonics coefficients per voxel, each containing 15 coefficients. Higher-order approximation would provide more accurate results but at increased computational complexity. For the purpose of numerical computation, an icosahedron-based fourth-order tessellation is employed to discretize $\overline{\text{dODF}}$ and σ_{dODF} .

5.3. VISUAL DESIGN REQUIREMENT ANALYSIS

Following Shneidermans mantra “Overview first, zoom and filter, then details on demand” [18], we propose a two-level overview + detail visualization. The overview is used to present the aggregate information on the ensemble, presented in Section 5.2.2. The detail views allow the exploration of the original tensor data.

The goal of the overview visualization is to enable the analyst to inspect the aggregate ensemble information (i.e., mean and variation) at discrete locations in the domain, as well as to reveal potential large-scale coherent patterns. Therefore, we define three requirements for the overview visualization:

- O1** It shall show the mean tensor and tensor variations simultaneously in the 3D domain through a single view, to avoid attention shifts and reduce reliance on memory during visual exploration.
- O2** The mean and variations in each of the three intrinsic properties should be separable and, therefore, individually identifiable.
- O3** When no variation is present in the ensemble, the visual representation shall converge to a well-known base representation, such as a glyph-based visualization for a single DTI dataset, to decrease the entry burden for analysts.

To identify how variations in the ensemble arise and to inspect tensors for different ensemble members, we define the following requirements for the detail visualization:

- D1** It shall show the complete tensor ensemble for voxels of interest without assumptions on the ensemble distribution.
- D2** It shall allow direct access to the three intrinsic properties separately and facilitate interpretations.
- D3** It shall allow to identify correlations between the different intrinsic properties.

5.4. GLYPH-BASED OVERVIEW VISUALIZATION

Glyphs principally allow the simultaneous visualization of multiple attributes. Therefore, it is a natural choice to employ glyphs to encode the multivariate summary information (**O1**). Well studied glyph representations (e.g., [7] for DTI, [9] for HARDI) form the basis to which our new design should gracefully reduce in case no variation is present in the ensemble (**O3**).

5.4.1. TENSOR-BASED GLYPH DESIGN

Superquadric tensor glyphs are commonly employed as base geometry for their ability to reduce ambiguity and preserve continuity [7, 19, 20]. A glyph G is constructed following

$$G(\mathbf{D}) = s(\text{tr}(\mathbf{D})) \mathbf{R} \tilde{\Lambda} B(\tilde{\lambda}_j), \quad (5.8)$$

where the superquadric base geometry $B(\tilde{\lambda}_j)$ is determined according to the sorted normalized eigenvalues. $\tilde{\Lambda}$ is a diagonal matrix composed of the normalized eigenvalues for scaling the base geometry. \mathbf{R} is a rotation matrix, defined by the column matrix of the eigenvectors $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$. s is an overall scaling function, which takes the tensor trace $\text{tr}(\mathbf{D})$ as input. Unlike the previous work [19, 20], we can use $\text{tr}(\mathbf{D})$ instead of the Frobenius norm $\|\mathbf{D}\|$ because we normalize eigenvalues with respect to tensor trace as presented in Section 5.2.1.

We indicate the major diffusion direction by the conventional spherical colormap [21], determined by the absolute values of the components of the major eigenvector \mathbf{e}_1 . The color is desaturated according to the tensor invariant c_l [22]. Unless noted otherwise, we use this colormap throughout this chapter. In the following, we extend

the superquadric glyph to incorporate the tensor variation information based on the fundamental construction rule Equation (5.8).




Figure 5.5: Visualization of the same mean tensor with different types/extents of variations (a): scale, (b&c): shape with Perline noise and halftone pattern, (d): orientation, (e): all combined. The variation gradually increases from left to right.

Scale Variation. The mapping of data attributes into visual channels of glyphs should be semantic [23], thus easing the learning process. The influence of the tensor scale on the glyph appearance is reflected in the scaling part $s(\text{tr}(\mathbf{D}))$ of Equation (5.8). Therefore, we decide to encode the tensor scale variation into glyph halos, similar to [19]. The larger the scale variation, the thicker is the halo (Figure 5.5a). Extra benefits brought by halos are the enhancement of depth perception, which is especially important in the context of orthographic projections, and identification of individual glyphs [24]. The halos are generated by rendering the same glyph with a larger scaling

factor, by adding the standard deviation to the trace, $s(\text{tr}(\mathbf{D}) + \sigma_{scale})$.

Shape Variation. Shape perception is roughly based on two factors, the underlying geometry, reflected in $\tilde{\Lambda}B(\tilde{\lambda}_j)$ of Equation (5.8), and the shading information [25]. A way to show tensor shape variation would be distorting the geometry and/or shading. However, introducing additional, irregular geometry would lead to a deviation from the commonly used glyphs. Thus it would potentially increase the learning burden. Instead we encode σ_{shape} by introducing a texture on top of the geometry. We propose two different approaches with distinct advantages and disadvantages, and allow the user to select one, depending on the use case.

The first approach adds Perlin noise of increasing frequency to indicate increasing σ_{shape} in object space, similar to previous work [26, 27], as shown in Figure 5.5b. Glyphs encoding large variations will show strong luminance contrast, to which the human visual system is sensitive [25]. The resulting texture will be invisible for $\sigma_{shape} = 0$ and, therefore, resembles the appearance of the glyph for a single dataset. The second approach employs a screen-space halftone pattern [28], as shown in Figure 5.5c. Here, σ_{shape} is encoded into the distances between neighboring dots on the glyph surface. The distance is inversely proportional to σ_{shape} . Smaller variations result in much larger distances while higher variations result in more densely packed dots.

The halftone pattern provides a more precise estimation of tensor shape variations, however, at the expense of artifacts during interaction, due to its screen space nature. The Perlin noise does not suffer from these artifacts, but is harder to interpret and quantify due to its irregularity. A downside for both techniques is their interference with color, which means that properties encoded by the color (e.g., the major eigenvector or FA) will become less legible.

Orientation Variation. The orientation information is reflected in the rotation matrix \mathbf{R} of Equation (5.8). One option to visualize $\sigma_{orientation}$ is to superimpose several glyphs with perturbed rotations [29]. However, this is likely to cause strong occlusions. Instead, we explicitly show $\sigma_{orientation}$ with a sector as shown in Figure 5.5d. Since there is no inherent orientation for the 2D sectors in the 3D domain, we align them with the view direction [24] to maintain optimal visibility. The orientation variation $\sigma_{orientation}$ is then encoded into the angle of the sector by the same mapping function as introduced in Chapter 4, with larger sectors indicating larger $\sigma_{orientation}$.

As shown in Figure 5.5, the halo, the surface texture, and the sector, visually encoding tensor variations in scale, shape, and orientation, respectively, are individually recognizable (**O2**). Furthermore, our design converges to the original superquadric glyph representation when no variation is present (**O3**, left column of Figure 5.5).

5.4.2. dODF-BASED GLYPH DESIGN




Figure 5.6: Three dODF-based glyphs for (a) an ensemble with gradually varying shape and orientation and two ensembles of linear tensors with crossing angle of (b) 60° and (c) 45° , respectively. The variation threshold is set to 60% of the maximum variation. The ensembles are illustrated by the small black icons on top.

We use the conventional spherical plot, which is commonly used to visualize HARDI data [9], to visualize $\overline{\text{dODF}}$. The glyph geometry represents the average diffusion probability density for each direction. Additionally, to emphasize the directional information, the *de facto* direction-coded colormap is used. This colormap can influence the shape perception [30], however, it is familiar to the target users. In order to incorporate the variation information, we employ a threshold-based method. Directional variation is encoded into neutral gray if it is greater than the user-specified threshold. We choose a neutral gray, with lightness value 0.5 in the CIE Lab color space, which has similar lightness as, but is not part of the direction-coded colormap to minimize the interference with the perception of the directional and shape information. Figure 5.4b shows the dODF glyph with the variations incorporated, derived from the ensemble of linear tensors in Figure 5.3a. This representation preserves the two major diffusion orientations that are present in the ensemble and the large variation is visible at the four lobes. Moreover, min-max normalization [9] is used to enhance the visual perception of directional maxima (e.g., see Figure 5.4c). The extended dODF glyph

meets the requirements **O2** and **O3**. Figure 5.6 further illustrates the dODF glyph representation for three different ensembles.

5.4.3. SPATIAL OVERVIEW

To fulfill requirement **O1**, glyphs are laid out according to their location in the data in a spatial view. We construct the tensor glyph (Section 5.4.1) and the dODF glyph (Section 5.4.2) for each voxel. Displaying both simultaneously would result in occlusion and visual clutter. Therefore, we show only a single, user selected glyph type at a time. The distinct properties of the two glyphs make them suitable for different tasks. If the analyst is interested in variations in individual tensor properties, the tensor glyph is the best choice; for detailed information on the orientation variation the dODF glyph is more suitable.

5

5.5. DETAIL VISUALIZATION

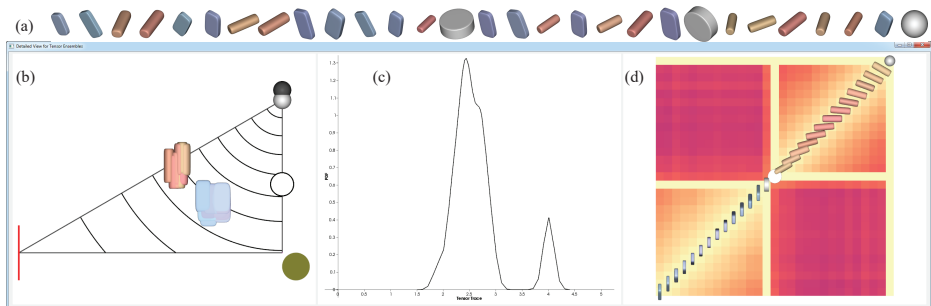


Figure 5.7: Our framework for the linked detail views. (a) shows a direct visualization of a synthetic set of tensors containing one isotropic tensor and two planar-anisotropic tensors, similar to small multiples. (b) shows the shape space view, (c) the scale space view, and (d) the orientation dissimilarity matrix view with a perceptually linear magma colormap for all the tensors in (a).

Based on the overview glyphs, the analyst can identify regions of interest and select few individual voxels for detail inspection of the ensemble. For further inspection of selected voxels, we provide linked detail views. A straightforward way to display all tensors (**D1**) within the ensemble are small multiples. However, simply placing the glyphs next to each, e.g. ordered by their id in the ensemble, hardly brings usable insights, as shown in Figure 5.7a. For effective comparison, data visualized as small multiples need to be ordered. Due to their multivariate nature, however, the tensor

glyphs presented above have no intrinsic order, and imposing an arbitrary order violates the *representation invariance principle* [31]. Following the idea of separating the intrinsic properties, we propose to show scale, shape, and orientation in three linked views, as shown in Figure 5.7. The separation fulfills **D2**. Furthermore, linking these views fulfills **D3**. To facilitate quantitative analysis, we also display all the statistical information (e.g., σ_{shape}) via graphical user interfaces.

5.5.1. TENSOR SCALE VISUALIZATION

Visualizing the scalar-valued tensor scale is straightforward. We estimate the probability density function of the tensor scales within the ensemble using a 1D kernel density estimate. The function is visualized directly in a 1D line plot. The exemplary plot in Figure 5.7c clearly shows the distribution of scale within the ensemble in Figure 5.7a.

5.5.2. TENSOR SHAPE SPACE VISUALIZATION

Figure 5.7b shows the proposed shape space visualization using the ensemble presented in Figure 5.7a. The tensor shape space is two dimensional and maps to a right triangle in \mathbb{R}^3 in Figure 5.2. We show the triangle in the background and draw the three interpretable extreme glyphs on the corners of the triangle. Additionally we show the iso-contours of other tensor invariants (e.g., FA) for contextual information, as the curved lines shown in Figure 5.7b. In principal the shape space visualization is a scatterplot. However, a mental reconstruction of the actual shape, just from the position and the extreme glyphs, is challenging. Therefore, we show the glyphs with unit scale and identical orientation to represent the tensor shape information in the plot. To reduce clutter, the size of the glyphs is adjustable. In the smallest setting, the shape space visualization becomes a standard scatterplot. Finally, the color channel is left for encoding additional information such as the age of the individuals.

5.5.3. ORIENTATION DISSIMILARITY MATRIX VISUALIZATION

The tensor orientation, given by the rotation matrix $[\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3]$, is an element of the manifold of the group $\text{SO}(3)$, and as such, has three dimensions. However, if a tensor has equal eigenvalues its eigenvectors are ill-defined. For example, the orientation of an isotropic tensor is undefined. Due to this coupling of shape and orientation, it is

not straightforward to define a space to show the distribution of tensor orientations. To sidestep this problem, we propose to compute pair-wise differences of the orientation using $d_{\text{orientation}}$ [32]. We visualize the pair-wise differences as a dissimilarity matrix, as shown in Figure 5.7d, which would allow us to identify patterns in the orientation behavior. A good ordering is essential to identify clusters of tensors with similar orientations. Thus we arrange the rows and columns of the matrix view using hierarchical clustering [33].

The orientation difference measure $d_{\text{orientation}}$ is not a *metric*. For example, a planar tensor has no orientation difference with any tensor that has the same minor eigenvector \mathbf{e}_3 . We refer to such a tensor \mathbf{D}^{oc} , where $d_{\text{orientation}}(\mathbf{D}^{oc}, \mathbf{D}_i) = 0$ for $i = 1 \dots n$, as an *oricentric* tensor. Considering the orientation differences $d_{\text{orientation}}$ only, oricentric tensors have no inherent order and, hence, can be placed at random positions during ordered hierarchical clustering [34]. Therefore, we propose a custom, *top-down* hierarchical clustering approach. We first divide the ensemble into oricentric and non-oricentric tensors. The oricentric tensors are clustered according to d_{shape} , which is a metric, using hierarchical clustering with complete-linkage. The group of non-oricentric tensors is then bisected, using the k-medoids algorithm [35] with $k = 2$, based on $d_{\text{orientation}}$. The bisection can cause tensors that were non-oricentric in relation to the complete ensemble to be oricentric in relation to the new subsets. Therefore, we recursively apply the separation into non-oricentric and oricentric subsets, followed by the bisection of the non-oricentric subset, until non-oricentric subset contains at most one or two tensors.

As suggested by Gruvae and Wainer [34], we concatenate two subsets such that elements on the edge of different subsets that are most similar are placed next to each other to show the relation between subsets. To concatenate two non-oricentric subsets, we choose the pair of edge elements with the smallest orientation difference while for a non-oricentric subset and an oricentric subset, we choose the pair with the smallest shape difference.

5.5.4. LINKED BRUSHING

Each of the three views presents only one aspect of the tensor properties in the ensemble. We use linked brushing to connect all the three views and to reveal correlations between different tensor properties ($\mathbf{D3}$). Therefore, we propagate selections defined

in one view to the other linked views. Thus, all the selected tensors in all the three views will be red highlighted. We also provide the option to render only the silhouettes of glyphs that are not selected in order to reduce occlusions.

5.6. IMPLEMENTATION

Our prototype is implemented in C++, combining VTK and OpenGL/GLSL for visualization. We use Qt for the GUI and Teem¹ for tensor data processing.

The glyph rendering is implemented by computing base geometries and storing them in GPU memory, as proposed by Schultz and Kindlmann [19]. We use three renderpasses to render the glyphs and corresponding variation information. In the first pass we draw the mean tensor glyph, including the texture indicating the shape variation. We linearly scale $[0, \sigma_{shape}]$ to $[0, 1]$. Since the theoretical maximum of tensor shape difference d_{shape} is $\sqrt{2/3}$, which corresponds to the length of the hypotenuse of the right triangle in Figure 5.2, typical σ_{shape} values are significantly smaller than 1 and therefore require the scaling. We render the Perlin noise texture by first assigning a random value *rand* between -1.0 and 1.0 to every fragment, with the fragments' object space coordinates and σ_{shape} as the input for the Perlin noise generator. To guarantee that the color-encoded information is legible to some extent, we only apply the texture to fragments with $|rand| < (1 - \sigma_{shape})$. We darken the fragments with negative sign while we brighten those with positive sign. We then render the halos by redrawing the geometry with increased size in black in the second pass, and finally the sectors in the third pass.

The dODF-based glyphs are rendered via ray casting [9]. After determining the intersection point \mathbf{p} of the ray with the glyph, we use the unit directional vector, $\mathbf{u} = \mathbf{v} / \|\mathbf{v}\|$ with $\mathbf{v} = \mathbf{p} - \mathbf{c}_{glyph}$ and the glyph center \mathbf{c}_{glyph} , to compute the diffusion probability density variation according to Equation 5.7.

5.7. EVALUATION AND APPLICATION

In this section, we compare our design to the work of Abbasloo et. al. [6], which is the closest to our work. We consider that domain knowledge is essential to understand the presented visual analysis. Unfortunately, this limits the number of possible

¹teem.sourceforge.net

participants. Here, we present a qualitative user study with two neuroscientists. We also present a case study illustrating the potential of our prototype with a real-world ensemble of 46 DTI datasets.

5.7.1. METHOD COMPARISON

Abbasloo et al. [6] use small multiples for the six eigenmodes, \mathbf{E}^k with $k = 1, \dots, 6$, of the fourth-order covariance tensor. To illustrate the effects of each eigenmode, they build a pair of tensors $\mathbf{D}(t; k) = \mathbf{D}^{\text{Eu}} \pm t\mu_k \mathbf{E}^k$, where \mathbf{D}^{Eu} is the Euclidean mean tensor and μ_k is the square root of the corresponding eigenvalues of the covariance tensor. The parameter t can be freely adjusted by users. The pair of tensors is then rendered

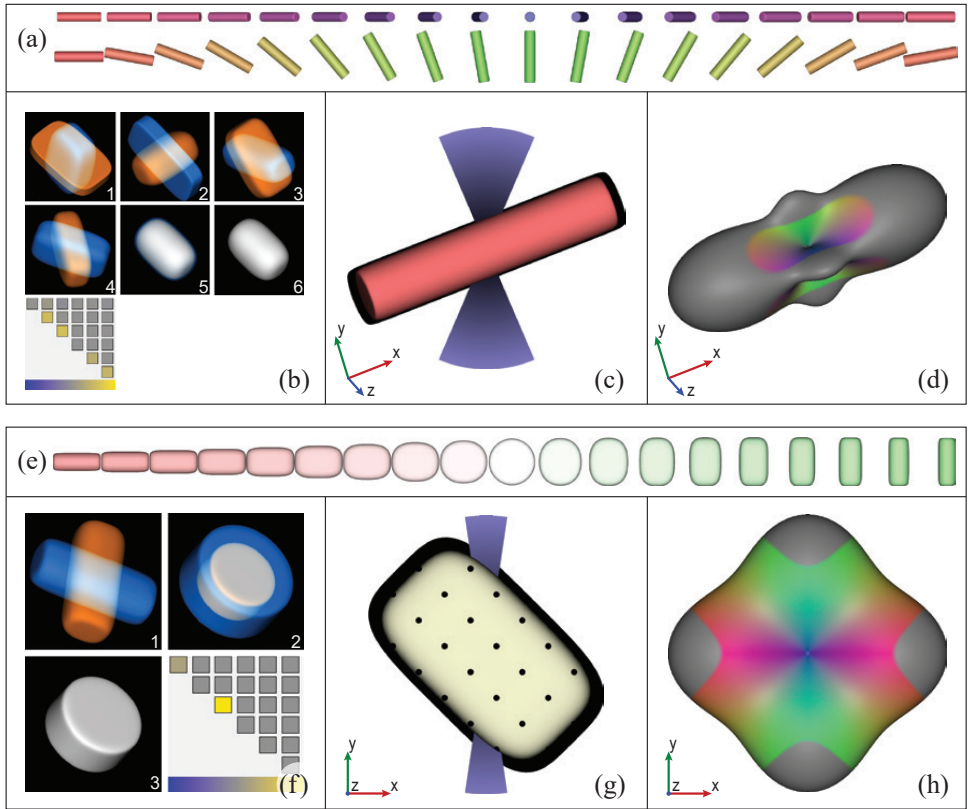


Figure 5.8: Comparison of two different aggregate visualizations for two different synthetic ensembles ((a) and (e)). (b) and (f) show the glyph-based eigenmodes visualization and IGRT matrix view. (c) and (g) show our tensor glyphs, depicting the variations in intrinsic properties. (d) and (h) show the dODF-based glyphs. The directions with large variations are colored in gray. For the purpose of clear illustration, images are generated from different viewpoints.

with two complementary colors, blue and orange. The normalized covariance tensor is decomposed with respect to the spherical invariants set (i.e., Frobenius norm, FA, and tensor mode) and rotation tangents [8], and displayed via a blue-gray-yellow colormap. Figures 5.8b and 5.8f show their visualizations. Figures 5.8c and 5.8g show our tensor glyphs while Figures 5.8d and 5.8h show our dODF-based glyphs. The variation threshold of $\overline{\text{dODF}}$ can be freely adjusted by users as well.

During the experiment, we first gave an introduction with a live demo of the two visualizations to be compared and the corresponding interaction methods. At this time, the participants interacted with the prototype to get themselves familiar with the visualizations. The objects of comparison are the aggregated glyphs in Section 5.4 and the visualizations proposed by Abbasloo et al. [6] of three synthetic datasets. We presented them to the participants, together with interaction capabilities, in *random* order. The participants freely analyzed the visualizations, and then we asked them to describe the original ensemble based on the aggregated visualizations they were observing in the form of a questionnaire. The participants had *no* information on the configuration of the ensembles apart from the provided aggregated visualizations, beforehand.

The first ensemble in Figure 5.8a comprises most variation in the orientation, some variation in scale, and no variation in shape. The scale variation is captured by all visual representations as suggested by the fourth eigenmode in Figure 5.8b and the visible halo in Figure 5.8c, respectively. The IGRT (i.e., invariant gradients and rotation tangents [8]) matrix of Abbasloo et al. shows high values in the 2nd and 3rd diagonal entries, corresponding to FA and mode respectively. The first and second eigenmodes, shown in the small multiples of Figure 5.8b, however, falsely indicates variation in shape. Our tensor glyph in Figure 5.8c shows no variation in shape. Both participants could not describe the variation in tensor shapes using the method of Abbasloo et al., while with our visualization they readily confirmed that only linear tensors are included in the ensemble. In Abbasloo et al., the third and fourth eigenmodes, together with the rotation tangents part of the IGRT matrix, i.e., the bottom three diagonal entries, show orientation variation in two orthogonal directions. The first and second eigenmodes show clear simultaneous variation in orientation and shape. One participant commented that it is really hard for him to combine all the eigenmodes to deduce information about the ensemble. Both participants deliberately ignored the rotations manifested in the first and second eigenmodes in Figure 5.8b when deducing

the orientation, and came to the correct conclusion that there are two main directions of rotation in the ensemble. Our aggregated tensor glyph in Figure 5.8c makes the orientation variation directly visible as a large sector. Furthermore, the dODFglyph in Figure 5.8d shows the principle diffusion direction as three orthogonal directions with high diffusion variation. Combining both glyphs, they came to the conclusion that a group of linear tensors spreads around the major eigenvector \mathbf{e}_1 of the mean tensor, which, while not describing the orientation variation exactly, is very close to the correct description.

The second ensemble in Figure 5.8e shows some variation in scale, shape, and orientation. This simulates the interface between two orthogonally aligned fiber tracts [8]. In Abbasloo et al., the last four eigenmodes correspond to zero eigenvalues. Therefore, we only keep the first three views in Figure 5.8f. The orientation and scale variation are manifested in the first and second eigenmode, respectively. The IGRT matrix in Figure 5.8f, however, indicates no variation in orientation (see the bottom three rows) but large variation in tensor mode (see the 3rd diagonal entry), which is part of the tensor shape. This contradiction between the eigenmode views and the IGRT matrix view causes considerable confusion for both participants. This is caused by the fact the variation in mode and minor eigenvector, shown as the 3rd and 6th diagonal entries respectively, should be interpreted integrally [8] near a planar mean tensor. Our tensor glyph in Figure 5.8g shows variation for all the three properties. The dODF glyph in Figure 5.8h provides additional information, indicating two main orthogonal directions with high diffusion variation. Both participants speculated that there may be fiber crossing effects in the underlying ensemble.

After the participants finished the overview analysis we showed them the underlying ensemble with our detail views and informally discussed the application. With regard to the tensor scale, the visualization of Abbasloo et al. and ours work well. However, our technique is found to be more accurately describing the shapes in the ensemble. In terms of tensor orientation, both participants remarked that the non-linearity of tensor orientations makes the covariance-based visualization harder to understand. With our method they could readily identify orientation variation but both thought it requires a considerable learning curve to interpret the dODF and to effectively combine the information of the two aggregated glyphs. Another insight we learn from the user study is that the visual representation of the shape variation is not optimal. While par-

ticipants found the noise texture less readable than the half tone texture, they preferred the stability of the object-space Perlin noise. We intend to explore other visual design idioms to optimize this in the future.

We do not show correlations between different tensor properties explicitly in our glyph designs. Correlations can be shown via eigenmodes, and the off-diagonal entries of the IGRT matrix. However, it is unclear whether it can be reliably interpreted. For instance, the first and second eigenmodes in Figure 5.8b indicate some relation in shape and orientation but there is no shape variation in the ensemble. Generally, both participants deem the proposed mean tensor, based on the separation of tensor intrinsic properties, as well as our visual design more intuitive.

5.7.2. APPLICATION TO AN DTI GROUP OF HEALTHY SUBJECTS

In this section, we present an exemplary analysis of a cohort DTI study, conducted with a neuroscientist, who is specialized in DTI-based group analysis and a co-author of this work. The dataset used for this case study is an ensemble, consisting of 46 DTI scans of healthy volunteers, age 47 to 78, serving as the control group in clinical research. Each scan has a resolution of $224 \times 224 \times 144$ with isotropic voxel size of 2mm. The diffusion weighting is $b = 1000 \text{ s/mm}^2$ along 64 uniformly distributed directions with four non-diffusion averages. The numerical range of the tensor values is scaled by 1000. The datasets are registered with a non-rigid registration approach by DTI-TK, which takes the whole tensor information into consideration [36].

During the visual exploration with the glyph-based overview in Figure 5.9b, we find a region of interest (ROI) with both high scale and shape variations. This region (see green highlight in Figure 5.9a) corresponds to the border between the ventricle, corpus callosum (CC), and inferior fronto-occipital fasciculus (IFO). The glyph-based overview in Figure 5.9b shows small variation in orientation (i.e., small sector angle). The glyphs in the bottom-right have thicker halos, indicating a larger scale variation, compared to the glyphs corresponding to the fibrous regions, i.e., bright areas in Figure 5.9a. In contrast, shape variation shown as half tone patterns is stronger in the top region. We now make a selection and use the detail views for further exploration. We select one of the voxels with larger shape and scale variation (see orange highlight in Figure 5.9b). In the detail views, the glyph color encodes the subject age. In the shape space view in Figure 5.9c, we see a transition from anisotropic tensors to

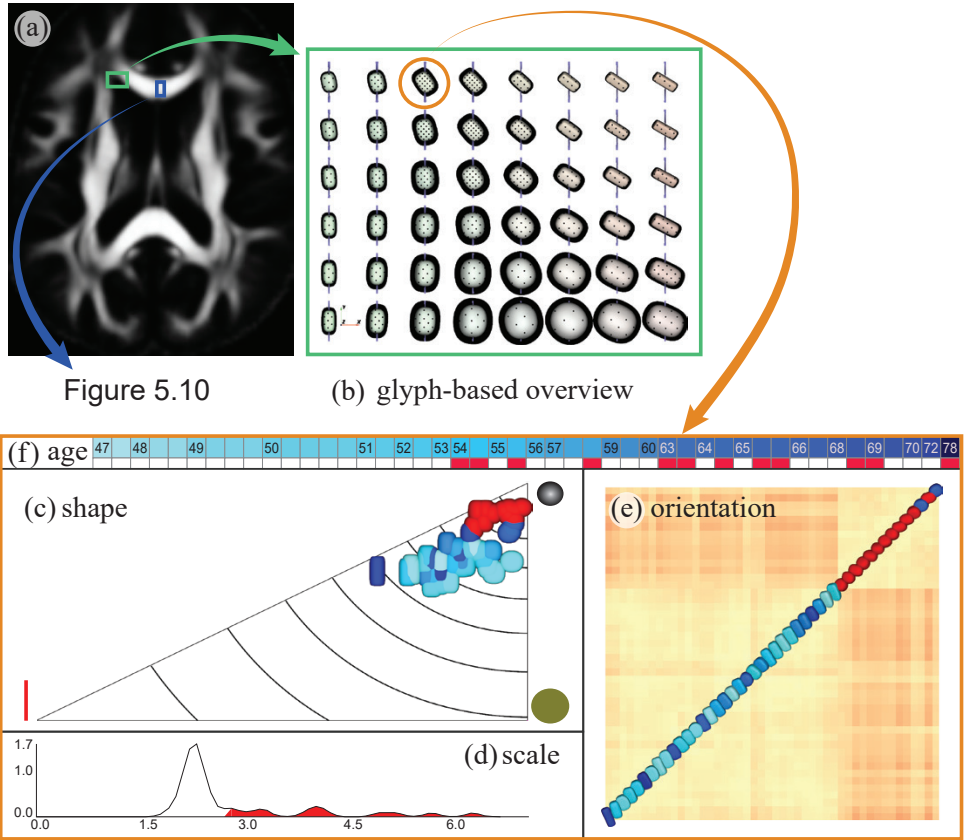


Figure 5.9: Visual analysis results for a border region of ventricle, CC, and IFO (a). The glyph-based overview (b) guides the selection of a voxel of interest for further exploration via shape space (c), scale line plot (d), orientation difference matrix (e), and age information (f).

nearly isotropic tensors. The scale view in Figure 5.9d shows a right-skewed distribution with a strong peak at lower scale values, as well as several large values. Selecting the right tail of the distribution in the scale view reveals the nearly isotropic tensors in the shape-space view via linked brushing, i.e., red highlighted in the different detail views. We also highlight the selection in a simple table view in Figure 5.9f, clearly identifying the older subjects. This confirms the expectation that older people have larger ventricles [2, 37], resulting in tensors with lower diffusion anisotropies but high diffusion scales. Somewhat unexpected is a clear separation in the orientation view in Figure 5.9e, even though differences in orientation are generally small, as already indicated by the glyph-based overview. Interestingly, there seems to be further correlation

with age, placing the previously selected tensors all in one group on the top right of the diagonal.

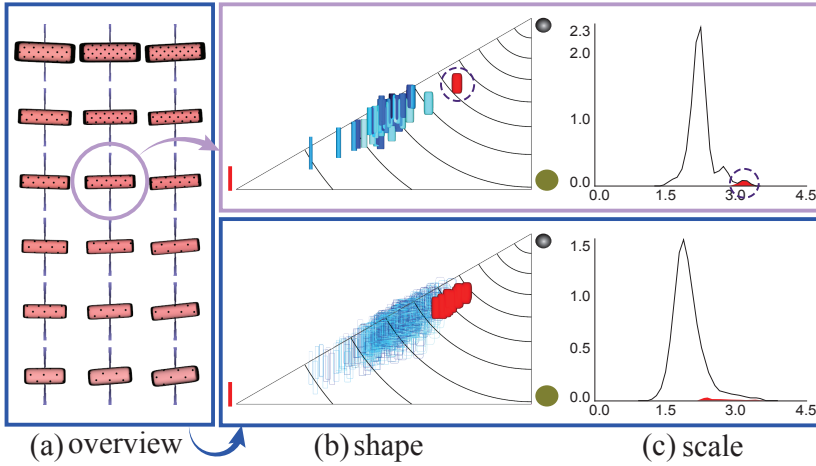


Figure 5.10: Identification of potential outliers in a region of the corpus callosum, highlighted in blue in Figure 5.9a. (a) shows the overview glyph visualization. (b) and (c) show the shape space view and scale line plot for (top) the selected voxel in (a) and (bottom) the complete ROI, respectively.

For the second analysis, we choose a different ROI (see blue highlight in Figure 5.9a) in the CC to illustrate that our detail views can be used to identify potential outliers. The glyph-based overview in Figure 5.10a shows that the variations in the ensemble are similar throughout the selected region. Slight variation in orientation and some variation in scale and shape can be identified. We select a voxel in the center, highlighted in Figure 5.10a, for further inspection in the detail views. An outlier tensor with lower anisotropy, encircled in Figure 5.10b top, becomes apparent. Through the linked selection we also find this member exhibits a larger scale value (see the encirclement in Figure 5.10c top) compared to the rest of the ensemble members. To verify that this behavior is consistent in the complete ROI, we select the 18 voxels in Figure 5.10a for visualization in the shape space view (Figure 5.10b bottom). We now select the subject corresponding to the outlier, identified in Figure 5.10b top, through the table view, highlighting only the corresponding tensors in the shape space view in Figure 5.10b bottom. The selection highlight clearly shows that all the corresponding diffusion tensors have coherently lower anisotropies (see Figure 5.10b bottom) and generally larger scale values (see Figure 5.10c bottom). This indicates that the selected ensemble member as a whole could be regarded as an outlier for this area.

5.8. CONCLUSIONS

We present an interactive approach for the analysis, visualization, and exploration of diffusion tensor ensembles. At the core of our visual analysis approach is a novel way to compute a representative tensor and quantify tensor variation through tensor intrinsic properties (i.e., scale, shape, and orientation), each of which has a biologically meaningful interpretation. By adapting the dODF, we provide an alternative for aggregating orientation variation, for cases where a single representative tensor cannot properly summarize different orientations in the ensemble. Another core contribution is the overview + detail visual design to facilitate per-voxel inspection of variation of tensor intrinsic properties free of assumptions about the underlying distribution. Finally, we demonstrate the effectiveness of our approach by comparing it to previous work by means of a qualitative user study with two domain experts and show the potential of our design by a case study.

The conducted user study requires knowledge about diffusion tensor analysis as well as second-order tensor normal distribution and as such limits the pool of possible participants. We consider a larger user study that would allow for strong conclusions on the effectiveness of the presented techniques future work.

We test our detail view system for ensembles of less than hundred members, a reasonable number for current DTI group studies. However, scaling to much larger numbers of ensemble members would require other solutions, specifically with respect to the visualizations of tensor shape space and orientation difference matrix.

Apart from data aggregation, another main task when conducting DTI group studies is the detection of potential outliers. While we demonstrate outlier detection using our detail views, our glyph representation currently does not provide any information on potential outliers. To ease the process of outlier detection a careful adaption of the glyph design that does not increase clutter is needed.

Finally, our design is targeted at second-order positive-definite tensors. However, similar rules apply to the glyph construction of second-order symmetric tensors [19]. Investigating the applicability of our design to those or even general second-order tensors would be an interesting future direction.

REFERENCES

- [1] C. Zhang, M. Caan, T. Höllt, E. Eisemann, and A. Vilanova, *Overview + detail visualization for ensembles of diffusion tensors*, Computer Graphics Forum **36**, 121 (2017).
- [2] S. M. Resnick, D. L. Pham, M. A. Kraut, A. B. Zonderman, and C. Davatzikos, *Longitudinal magnetic resonance imaging studies of older adults: A shrinking brain*, The Journal of Neuroscience **23**, 3295 (2003).
- [3] D. K. Jones, *Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI*, Magnetic Resonance in Medicine **49**, 7 (2003).
- [4] P. J. Basser and S. Pajevic, *A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI*, IEEE Transactions on Medical Imaging **22**, 785 (2003).
- [5] P. Basser and S. Pajevic, *Spectral decomposition of a 4th-order covariance tensor: applications to diffusion tensor MRI*, Signal Processing **87**, 220 (2007).
- [6] A. Abbasloo, V. Wiens, M. Hermann, and T. Schultz, *Visualizing tensor normal distributions at multiple levels of detail*, IEEE Transactions on Visualization and Computer Graphics **22**, 975 (2016).
- [7] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004) pp. 147–154.
- [8] G. Kindlmann, D. B. Ennis, R. T. Whitaker, and C.-F. Westin, *Diffusion tensor analysis with invariant gradients and rotation tangents*, IEEE Transactions on Medical Imaging **26**, 1483 (2007).
- [9] T. H. J. M. Peeters, V. Prckovska, M. van Almsick, A. Vilanova, and B. M. ter Haar Romeny, *Fast and sleek glyph rendering for interactive HARDI data exploration*, in *IEEE Pacific Visualization Symposium* (2009) pp. 153–160.
- [10] D. B. Ennis and G. Kindlmann, *Orthogonal tensor invariants and the analysis of diffusion tensor magnetic resonance images*, Magnetic Resonance in Medicine **55**, 136 (2006).

- [11] P. J. Basser and C. Pierpaoli, *Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI*, *Journal of Magnetic Resonance* **111**, 209 (1996).
- [12] G. Kindlmann, R. San José Estépar, M. Niethammer, S. Haker, and C.-F. Westin, *Geodesic-loxodromes for diffusion tensor interpolation and difference measurement*, in *MICCAI (2007)* pp. 1–9.
- [13] J. K. Gahm, N. Wisniewski, G. Kindlmann, G. L. Kung, W. S. Klug, A. Garfinkel, and D. B. Ennis, *Linear invariant tensor interpolation applied to cardiac diffusion tensor MRI*, in *MICCAI (2012)* pp. 494–501.
- [14] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel, *Reconstruction of the orientation distribution function in single- and multiple-shell q -ball imaging within constant solid angle*, *Magnetic Resonance in Medicine* **64**, 554 (2010).
- [15] D. S. Tuch, *Q -ball imaging*, *Magnetic Resonance in Medicine* **52**, 1358 (2004).
- [16] T. Schultz and H.-P. Seidel, *Estimating crossing fibers: A tensor decomposition approach*, *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE Visualization)* **14**, 1635 (2008).
- [17] C. P. Hess, P. Mukherjee, E. T. Han, D. Xu, and D. B. Vigneron, *Q -ball reconstruction of multimodal fiber orientations using the spherical harmonic basis*, *Magnetic Resonance in Medicine* **56**, 104 (2006).
- [18] B. Shneiderman, *The eyes have it: A task by data type taxonomy for information visualizations*, in *IEEE Symposium on Visual Languages* (1996) pp. 336–343.
- [19] T. Schultz and G. L. Kindlmann, *Superquadric glyphs for symmetric second-order tensors*, *IEEE Transactions on Visualization and Computer Graphics* **16**, 1595 (2010).
- [20] N. Seltzer and G. Kindlmann, *Glyphs for asymmetric second-order 2D tensors*, *Computer Graphics Forum* **35**, 141 (2016).
- [21] S. Pajevic and C. Pierpaoli, *Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber*

- tract mapping in the human brain*, *Magnetic Resonance in Medicine* **42**, 526 (1999).
- [22] C.-F. Westin, S. E. Maier, B. Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis, *Image processing for diffusion tensor magnetic resonance imaging*, in *MICCAI* (1999) pp. 441–452.
- [23] R. Borgo, J. Kehrer, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen, *Glyph-based visualization: Foundations, design guidelines, techniques and applications*, in *Eurographics State of the Art Reports* (2013) pp. 39–63.
- [24] A. E. Lie, J. Kehrer, and H. Hauser, *Critical design and realization aspects of glyph-based 3D data visualization*, in *Spring Conference on Computer Graphics* (2009) pp. 19–26.
- [25] C. Ware, *Information Visualization: Perception for Design*, 3rd ed. (Elsevier, 2013).
- [26] R. P. Botchen, D. Weiskopf, and T. Ertl, *Texture-based visualization of uncertainty in flow fields*, in *IEEE Visualization* (2005) pp. 647–654.
- [27] T. Höllt, F. de Matos Ravanelli, M. Hadwiger, and I. Hoteit, *Visual analysis of reservoir simulation ensembles*, in *Workshop on Visualization in Environmental Sciences* (2016) pp. 1–4.
- [28] N. Smit, K. Lawonn, A. Kraima, M. DeRuiter, H. Sokooti, S. Bruckner, E. Eise-mann, and A. Vilanova, *Pelvis: Atlas-based surgical planning for oncological pelvic surgery*, *IEEE Transactions on Visualization and Computer Graphics* **23**, 741 (2017).
- [29] D. K. Jones, L. D. Griffin, D. C. Alexander, M. Catani, M. A. Horsfield, R. Howard, and S. C. Williams, *Spatial normalization and averaging of diffusion tensor MRI data sets*, *NeuroImage* **17**, 592 (2002).
- [30] D. Borland and R. M. Taylor II, *Rainbow color map (still) considered harmful*, *IEEE Computer Graphics and Applications* **27**, 14 (2007).

- [31] G. Kindlmann and C. Scheidegger, *An algebraic process for visualization design*, IEEE Transactions on Visualization and Computer Graphics **20**, 2181 (2014).
- [32] C. Zhang, T. Schultz, K. Lawonn, E. Eisemann, and A. Vilanova, *Glyph-based comparative visualization for diffusion tensor fields*, IEEE Transactions on Visualization and Computer Graphics **22**, 797 (2016).
- [33] M. Behrisch, B. Bach, N. Henry Riche, T. Schreck, and J.-D. Fekete, *Matrix re-ordering methods for table and network visualization*, Computer Graphics Forum **35**, 693 (2016).
- [34] G. Gruvaeus and H. Wainer, *Two additions to hierarchical cluster analysis*, British Journal of Mathematical and Statistical Psychology **25**, 200 (1972).
- [35] L. Kaufmann and P. Rousseeuw, *Clustering by means of medoids*, Statistical Data Analysis Based on the L1-Norm and Related Methods , 405 (1987).
- [36] H. Zhang, P. A. Yushkevich, D. C. Alexander, and J. C. Gee, *Deformable registration of diffusion tensor MR images with explicit orientation optimization*, Medical Image Analysis **10**, 764 (2006).
- [37] C. Lebel, M. Gee, R. Camicioli, M. Wieler, W. Martin, and C. Beaulieu, *Diffusion tensor imaging of white matter tract evolution over the lifespan*, NeuroImage **60**, 340 (2012).

6

COMPARATIVE VIS. FOR ENSEMBLES OF DIFFUSION TENSOR FIELDS AT MULTIPLE LEVELS OF DETAIL

If our troops are no more in number than the enemy, that is amply sufficient; it only means that no direct attack can be made. What we can do is simply to concentrate all our available strength, keep a close watch on the enemy, and obtain reinforcements.

Sun Tzu

兵非贵益多，惟无武进，足以并力料敌取人而已。

孙武

This chapter is based on the following publication: “Comparative Visualization for Diffusion Tensor Imaging Group Study at Multiple Levels of Detail”, **Changgong Zhang**, Thomas Höllt, Matthan W.A. Caan, Elmar Eisemann, and Anna Vilanova. In *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, 53–62, 2017 [1].

In the previous two chapters, we compare two diffusion tensor fields, and summarize a single ensemble of diffusion tensor fields. In this chapter, the size of the problem to solve keeps increasing. Specifically, we will visually compare two ensembles of tensor fields in the context of DTI group studies. The comparative visualization technique developed in this chapter is inspired by those proposed in Chapters 4 and 5.

6.1. INTRODUCTION

DTI group studies aim at spatially locating diffusion-related differences between two groups in order to discover potential markers for white matter structures, and are normally carried out in a voxel-wise manner. For example, several studies [2–5] compare a control group and a patient group to search for differences caused by certain pathologies. Compared to the problem we solve in Chapter 5, the amount of information to be visualized is hugely increased. Therefore, inter-group comparison adds much more complexity for the design of comprehensive visual analysis tools.

The standard approach in literature relies on the univariate statistical analysis of tensor-derived scalar quantities such as fractional anisotropy (FA), rather than the multivariate diffusion tensor itself. Slice views are commonly used to display the analysis results, which are some specific highly aggregated scalar values (e.g., the t-statistic). This type of scalar-based comparative analysis provides merely one facet of the full tensor information, and leaves a large amount of information unused. If to keep the full tensor information, one has to resort to the multivariate analysis tools such as multivariate normal distribution [6, 7]. However, the intrinsic structure of diffusion tensors is not taken into consideration, and it is also unclear how to effectively compare the generated fourth-order covariance tensors. More importantly, to the best of our knowledge, there is no comparative visualization technique developed specifically for inter-group comparison.

Our work extends Chapter 5 to facilitate the visual comparison of two DTI groups at multiple levels of detail. Our comparative visualization framework is built based on the Shneiderman mantra “Overview first, zoom and filter, then details-on-demand” [8]. The spatial overview is for comparing the multivariate statistical summary information of tensor ensembles, i.e., mean tensor & tensor variations (see Section 5.2). We combine the checkerboard style pair-wise tensor difference (Tender) glyph (see Section 4.2) and the tensor ensemble summary glyph (see Section 5.4) to support the

visual comparison in the spatial domain. More importantly, we integrate the details-on-demand scheme into the glyph visualization such that, as analysts zoom in, more statistical information will be progressively revealed. We refer to the visual comparison of statistical information in the spatial space as *summary comparison*. The purpose of the spatial overview is to guide analysts to locate voxels of interest for the further detailed comparison of the original tensors without aggregation. We extend the detail visualization of Chapter 5 by utilizing superposition and explicit difference encoding to support the comparison in terms of tensor intrinsic properties as the finest level, and refer to it as *feature comparison*. The main contributions of this chapter can be summarized as:

- We present a level-of-detail glyph representation to support the visual comparison of tensor ensemble summary statistics between two ensembles of DTI datasets at different zoom levels in the spatial domain;
- We facilitate the visual comparison of two groups of the diffusion tensors in terms of tensor intrinsic properties by superposition and explicit encoding.

6.2. COMPARATIVE VISUALIZATION DESIGN

Neuroscientists normally apply voxel-wise analysis to locate diffusion related differences between two groups of DTI datasets. Therefore, the main objects to be compared in this work are two groups of diffusion tensors, the so-called *tensor ensembles*, at the same spatial location. Each tensor within the ensemble corresponds to one individual subject. We denote them as $\{\mathbf{D}_i\}_A$ and $\{\mathbf{D}_j\}_B$ with $i \in [0, m)$ and $j \in [0, n)$, where m and n are the total number of subjects in the corresponding groups A and B , respectively. We use two complementary and colorblind-safe colors, i.e., orange and blue [9, 10], to identify different groups.

In the following sections, we first develop the visualization for *directly* comparing $\{\mathbf{D}_i\}_A$ and $\{\mathbf{D}_j\}_B$ but not considering the field information, and then design the visualization for comparing the ensemble summary information in the 3D spatial domain.

6.2.1. DIRECT TENSOR ENSEMBLE COMPARISON

In Chapter 5, we visualize a single tensor ensemble in terms of three tensor intrinsic properties (scale, shape, and orientation), each of which has a biologically meaning-

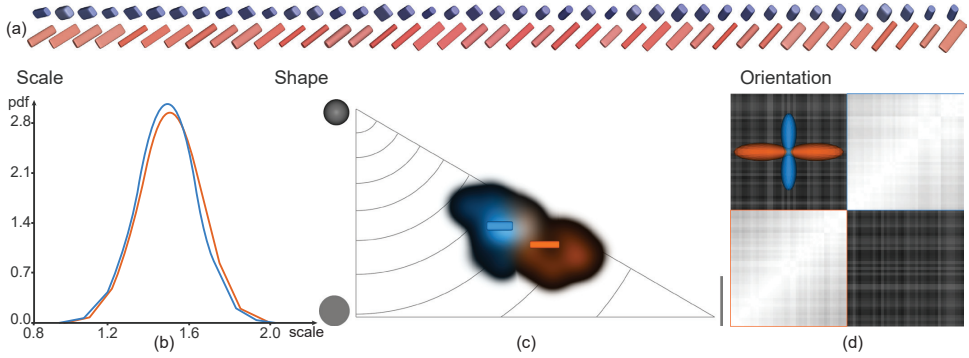


Figure 6.1: **Direct visual comparison of two tensor ensembles.** (a) each row represents a synthetic tensor ensemble, which is the subject of comparison. All the tensor glyphs are colored according to the major eigenvectors with the conventional spherical colormap [11]. The color is desaturated based on the tensor invariant c_1 [12]. The major eigenvectors of these two ensembles are perpendicular to each other. (b) shows the line plots of the probability density function of the one dimensional tensor scale for the two ensembles. (c) shows the blended 2D density plot in the triangular tensor shape space, overlaid with the mean tensor glyphs for each ensemble. (d) shows the pair-wise orientation dissimilarity matrix with the two dODF [13] glyphs. (b–d) blue and orange colors indicate the tensor ensemble at the upper and lower row in (a), respectively. Note that the glyphs in (a, c, d) are visualized with different viewing settings.

ful interpretation. We extend the *detail views* (see Section 5.5) to facilitate the visual comparison of two tensor ensembles, also from the perspective of the tensor intrinsic properties. Two synthetic tensor ensembles (see Figure 6.1a) are used for illustration in the following sections.

Tensor Intrinsic Properties. A symmetric positive-definite diffusion tensor \mathbf{D} can be eigen-decomposed into three ordered real eigenvalues ($\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$), and the corresponding orthonormal eigenvectors $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$. λ_i and \mathbf{e}_i ($i = 1, 2, 3$) are referred to as the major, medium, and minor eigenvalue and eigenvector, respectively. The tensor trace $|\mathbf{D}|$ is defined as the sum of the eigenvalues $|\mathbf{D}| = \sum_{i=1}^3 \lambda_i$. Finally, the ordered normalized eigenvalues $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$ are given by normalizing the eigenvalues with respect to the trace $(\lambda_1, \lambda_2, \lambda_3) / |\mathbf{D}|$.

Tensor intrinsic properties include scale, shape, and orientation. Tensor scale is mathematically expressed as the trace $|\mathbf{D}|$, which represents the amount of diffusion. Tensor shape is two dimensional, and mathematically expressed as the set of normalized eigenvalues $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$. Tensor shape reflects the underlying fiber configurations. For example, for regions with coherently aligned fiber tracts, the diffusion tensors are of linear shape. Tensor orientation is given by the set of eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$, which

can be used to deduce the pathway of the fiber tracts. Interested readers are referred to Chapter 4 for more details.

Scale Comparison. In Section 5.5, a 1D probability density function (PDF) is employed to depict the distribution of tensor scale, which is visualized as a line plot. It is straightforward to compare two distributions of tensor scale. We use the *superposition* strategy to overlay two line plots within the same frame of reference, each of which is colored according to their group identifier color. For example, as shown in Figure 6.1b, two line plots largely overlap, indicating that two corresponding distributions are quite similar to each other.

Shape Comparison. Tensors with all possible shapes are located on a right triangle (see Figure 5.2), and hence we compare the shapes of two tensor ensembles in this triangular shape space. However, displaying all the tensor glyphs at once or drawing them as a scatter plot generates clutter or does not allow an easy understanding of the actual shapes, respectively. We tackle these problems by combining 2D probability density estimation [14] with glyph-based visualization. First, we draw the 2D density plot for each tensor ensemble separately. We then blend the two plots in image space as $c_{i,j} = p_{i,j}^A * c_A + p_{i,j}^B * c_B$, where (i, j) is the pixel coordinate, p is the density value, c_A and c_B are the corresponding group identifier colors (see Figure 6.1c). Abbasloo et al. [10] employ the same strategy to avoid occlusions for glyph-based visualization. Blending the colors as described will create an impression of neutral gray, where the density values are similar. If one of the values is larger, the corresponding group color will become obvious at this position, and thus allow the identification of differences in the distribution of the tensor shape between ensembles. However, the mental reconstruction of the actual tensor shapes is not easy. Therefore, we draw several representative glyphs to assist the perception of the blended density plot without introducing much occlusion. We provide two options. First, the user can select to simply show the mean tensor for each group, resulting in just two glyphs in total rendered on top of the density plot (see Figure 6.1c). However, the shape distribution can be multi-modal. Therefore, the second option is based on automatic peak finding in the density and subsequent clustering. To this end, we apply GPU-based mean shift clustering based on the generated density plot, similar to previous work by Höllt et al. [15]. The mean shift algorithm [16] has several useful properties. The density clusters

will closely match the visual clusters that analysts would identify. Furthermore, since we compute the density plot for visualization, we can use the same density estimate as input for the mean shift algorithm. Finally the mean shift algorithm can also function as a peak finding algorithm on the density plot and the identified peaks can be used to place one representative glyph for each cluster. To obtain the clusters, based on the density estimate, we first compute the gradient of the density estimate by calculating the central differences for each pixel in image space. Then, starting with the initial position for each glyph, we ascend along the gradient direction until we reach a peak. All glyphs that ascend to the same peak are considered as belonging to the same cluster and the peak can be used for placing the representative glyph for that cluster. Finally, we reconstruct the shapes for the representative glyphs, based on the position of the peak in image space, and render the glyphs on top of the density plots. An example with tensor glyphs at cluster centers can be seen later in Figure 6.4d top. By combining the 2D density plot and the tensor glyphs, we effectively reduce the visual clutter and provide a straightforward comparison in the tensor shape distribution.

Orientation Comparison. It is challenging to define a space to explicitly convey the orientation distributions due to the inherent coupling of tensor shape and orientation. As shown in Section 5.5, the pair-wise dissimilarity matrix of tensor orientations proves to be useful in identifying similar orientation behaviors. Therefore, we order the elements for each tensor ensemble, in the same way as in Section 5.5, and then concatenate these two ensembles together to build the final dissimilarity matrix. The orientation matrix view is therefore composed of three parts (see Figure 6.1d). A white-to-black perceptually linear colormap is employed to color the matrix. The lower diagonal part depicts the orientation variation of one ensemble, while the upper diagonal part shows that of the other. The remaining two parts, which are symmetric with respect to the diagonal, depict the inter-group orientation differences. The diagonal parts allow exploring the orientation variation within the ensemble, and the off diagonal parts allow inspecting the inter-ensemble orientation differences, which belongs to the comparison category of *explicit encoding*.

In order to support the interpretation of tensor orientation variation of each ensemble, we convert each tensor into the so-called diffusion orientation distribution function (dODF) [13], and obtain an average dODF glyph for each ensemble. In principle the




Figure 6.2: **Level-of-detail glyph representations** for tensor ensemble summary comparison. (a) displays the corresponding tensor summary glyphs for the synthetic ensembles shown in Figure 6.1a. (b) shows the new LOD glyph at zoom level 3 (i.e., the fine level), designed to compare the tensor ensemble summary information. (c) shows the new LOD glyph at zoom level 2 (i.e., the medium level), used to support the visual comparison of the corresponding mean tensors alone. (d) shows our LOD glyph at zoom level 1 (i.e., the coarse level) with plain color coding to indicate the tensor shape difference. (b,c) orange and blue color are for the identification of different tensor ensembles.

dODF glyph is a directional probability density function defined on a unit-sphere. We apply the *superposition* strategy to compare two dODF glyphs, and provide the option to add transparency. We place the dODF glyphs on top of one of the two symmetric off diagonal parts, which will not cause information loss. The dissimilarity matrix is two dimensional while the dODF glyphs are three dimensional objects. Therefore, we offer the option to freely rotate the dODF glyphs to have a comprehensive perception without affecting the viewing settings. Hereby, we not only show the orientation variations for a single ensemble, but also facilitate the visual comparison between two ensembles.

As shown in Figure 6.1d, both ensembles exhibit a small intra-ensemble orientation variation (i.e., brighter colors), while the inter-ensemble orientation differences are rather large (i.e., darker colors). Additionally, the superposed dODF glyphs clearly shows the differences in dODF for every directions.

Linked Brushing. Each view is extended to support the inter-group comparison in one tensor intrinsic property. Therefore, linked brushing is used to connect all of the three views. To minimize visual clutter, we only highlight the selected tensor glyphs with red halos (e.g., see Figure 6.4d bottom). A filled line plot (e.g., see Figure 6.4e) is drawn to reflect the contributions of selected tensors to the overall scale distributions, and the corresponding diagonal entries of the matrix view (e.g., see Figure 6.4f) are also colored red. We use a table view to display the metadata (e.g., age) of each individual subject, which can also be used for linked selection.

6.2.2. ENSEMBLE SUMMARY COMPARISON

The extended comparative visualization, described in the previous section is useful for directly comparing the tensor ensembles for voxel(s) of interest. However, this visualization cannot provide spatial information. Therefore, we use a separate spatial visualization, described in this section, to show the ensemble differences in the 3D spatial domain. Before performing voxel-wise ensemble comparison information needs to be aggregated. Chapter 5 summarize the tensor ensemble $\{\mathbf{D}_i\}$ into a mean tensor $\overline{\mathbf{D}}$ and the variations in scale σ_{scale} , shape σ_{shape} , and orientation $\sigma_{\text{orientation}}$. Furthermore, they extend the superquadric tensor glyph [17] to visualize the summary information.

We employ the same aggregation method, which handles the tensor intrinsic properties separately, and design a new glyph to visually compare the summary information in 3D. Our new design is a combination of the checkerboard style Tender glyph (see Section 4.2), designed for comparing two diffusion tensors, and the tensor ensemble summary glyph (see Section 5.4), designed for visualizing the summaries of a single tensor ensemble. Both types of glyphs are suitable for detailed inspection at close-up views. In this work, we incorporate the level-of-detail concept by designing different visual representations at different zoom levels. For a small field of view, it is more suitable to encode the full information for a limited number of voxels, while for a large field of view it is more suitable to encode a limited amount of information for a large amount of voxels to provide an overview. The zoom level is reversely proportional to the extent of information aggregation; the higher the zoom level, the less aggregation is needed. Therefore, in the following, we describe the levels from the most detailed to the coarsest by gradually removing information.

Aggregation Level 1 (Zoom Level 3) : Mean tensors and tensor variations. At the finest zoom level, the aggregate information to be compared is the mean tensor and three scalar-valued tensor variations. We apply the checkerboard style Tender glyph to compare the mean tensors. The original design uses the color channel to compare tensor scales. However, we decide to reserve the color channel for later usage at the coarsest zoom level. Therefore, we have to resort to other visual channels to encode the tensor scale information. We find that the original sector representation, which is shown both above and below the glyphs, for the orientation difference is redundant. Hence, we divide the sector into two parts, i.e., the upper and lower parts. Then, the

angles of the upper and lower part are used to encode orientation-related and scale-related differences, respectively. Figure 6.3 illustrates the idea of sector representation by analogy with the 1D normal distribution. At zoom level 3, two bands are drawn with the identification color. The center of the band is aligned with the mean value (i.e., the vertical dashed line) while the length of the band is determined by the standard deviation. As shown in Figure 6.3a, if these two bands overlap, the overlapping region will be colored white (i.e., the addition of orange and blue). Otherwise, if there is a gap between the two bands, the in-between region is colored to with a dark gray (Figure 6.3b). We create the sectors for the glyph visualization by replacing the length with the angle of the arc. Figures 6.3a and 6.3b at zoom level 3 are analog to the lower and upper sectors in Figure 6.2b, respectively, with the standard deviation corresponding to scale and orientation variations (see Section 5.2). The major benefit with this visual design is that it facilitates the direct evaluation of the overlap between the two distributions in scale and orientation.

For the comparison of mean tensor shapes, we place them in a checkerboard style as presented for the Tender glyph. For the visual comparison of tensor shape variations σ_{shape} , we employ the halftone textures due to their ability to provide a more precise

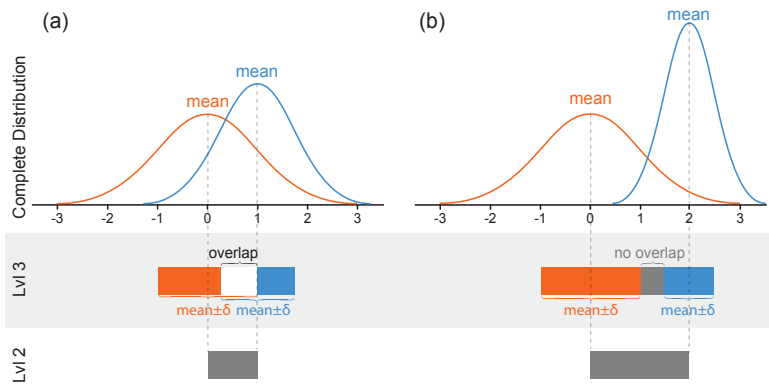


Figure 6.3: **Illustration of the sector construction** by analogy with one dimensional normal distributions for (a) a large and (b) small distribution overlap. Blue and orange colors are used to identify the corresponding normal distributions. At level 3, two colored bands are drawn, the length of which are proportional to the standard deviation of the distribution. The center of the bands are based on the mean value. (a) if these two bands are overlaid, white color is generated. (b) Otherwise, the color in-between is kept to gray. At level 2, a single gray band is drawn, the length of which indicates the difference in the mean values of the distributions. (a) and (b) are analog to the lower and upper sectors, respectively, in Figures 6.2b and 6.2c.

estimation (see Section 5.4). In this way, we convert the comparison of tensor shape variations into that of the density of the colored dots.

Figure 6.2b shows the final glyph representation at this zoom level, which combines the two ensemble summary glyphs in Figure 6.2a. Please note that Figures 6.1a and 6.2a are generated with different viewing settings. The shape differences between the two mean tensors are manifested as the discontinuities between the glyphs. The orange ensemble exhibits a slightly larger shape variation (i.e., the dots are packed slightly more dense). The upper gray sector shown in Figure 6.2b indicates a large orientation difference between the mean tensors while the two small colored sections on both sides indicate little orientation variation within both ensembles. From the lower sector shown in Figure 6.2b, we can deduce that two ensembles have similar scale variations (i.e., two small orange and blue ends on both sides) and largely overlap with each other (i.e., the white region).

Aggregation Level 2 (Zoom Level 2) : Mean Tensors. As users zoom out, more voxels are shown, and hence fewer pixels are assigned per voxel. Accordingly, we reduce the amount of information to be displayed. Prior to comparing the associated tensor variation information, neuroscientists are normally interested in comparing the mean tensors. If the mean tensors are different, neuroscientists prefer to dive further to assess the overlap between the distributions of the tensor ensembles. Therefore, at this zoom level, we decide to compare the mean tensors alone. More specifically, we leave out the visual channels that are used to encode tensor variation information, and show the glyphs to only support the comparison of the mean tensors. Figure 6.2c shows that the glyph at this level is a simplified representation of the glyph at zoom level 3. Specifically, the textures are removed. The sectors are simplified to encode only the pair-wise scale/orientation differences to the angles, which are analog to the visual representation at level 2 in Figure 6.3. Furthermore, the sectors are colored using a neutral gray. At zoom levels 2 and 3 the glyph color is employed to explicitly encode the difference of the ensembles, as described in the next paragraph.

Aggregation Level 3 (Zoom Level 1): Difference Measure between Mean Tensors. At the coarsest zoom level, each voxel occupies a very limited number of pixels, and therefore requires a higher level of aggregation. The goal of this level is to help analysts to gain a first overview and identify areas of interest, into which they can dig deeper.

Consequently, we decide to reduce these two mean tensors to single scalar values that can then be encoded into the color channel, which is still visible even if only few pixels are available per voxel. At this level, it is most important to choose an appropriate difference measure between the mean tensors. Such measures are strongly application-dependent. For example, Abbasloo et al. [10] use slice views to show the overall variance and generalized FA of the fourth-order covariance tensor to provide a first impression of the DTI ensemble. We allow analysts to choose among a set of pre-defined difference measures between the mean tensors $d(\overline{\mathbf{D}}_A, \overline{\mathbf{D}}_B)$, depending on the task at hand. For example, the analyst can choose the difference in a specific tensor intrinsic property such as the tensor shape difference d_{shape} . Furthermore, statistical measures like the t-statistic, with which neuroscientists are familiar in their routine research, can be used. We decide to use the shape of the glyphs instead of simple view-aligned quads to show the color at this zoom level. This guarantees a smooth transition between different glyph representations as much as possible. Since usually only a few pixels are used for the glyph at this zoom level, the shape of individual glyphs is barely noticeable, making it unnecessary to apply an illumination model. Figure 6.2d shows the glyph representation at the coarse level, which is colored according to the shape differences between the mean tensors.

6.2.3. SUMMARY

We have introduced the three zoom levels in reverse order for the sake of the easy explanation of the aggregation of tensor ensembles. The actual visual exploration workflow starts from zoom level 1. The coarse zoom level utilizes the *explicit encoding* strategy to depict application-dependent difference measures between two mean tensors. This level provides a first impression for a large field of view, and hence can guide analysts to voxels of interest (e.g., Figure 6.4a). Analysts can then zoom into level 2, which is based on *juxtaposition* and *explicit encoding*. At this level, analysts can visually compare the mean tensors in all the three aspects of tensor intrinsic properties. If analysts would like to further compare the tensor ensemble variations, they can continue to zoom level 3, which is also based on *juxtaposition* and *explicit encoding*. At this level, analysts can scrutinize a small amount of voxels, harnessing the full tensor ensemble summary information. Finally, this zoom level serves as the bridge between the comparison of the tensor ensemble summaries in spatial domain and the

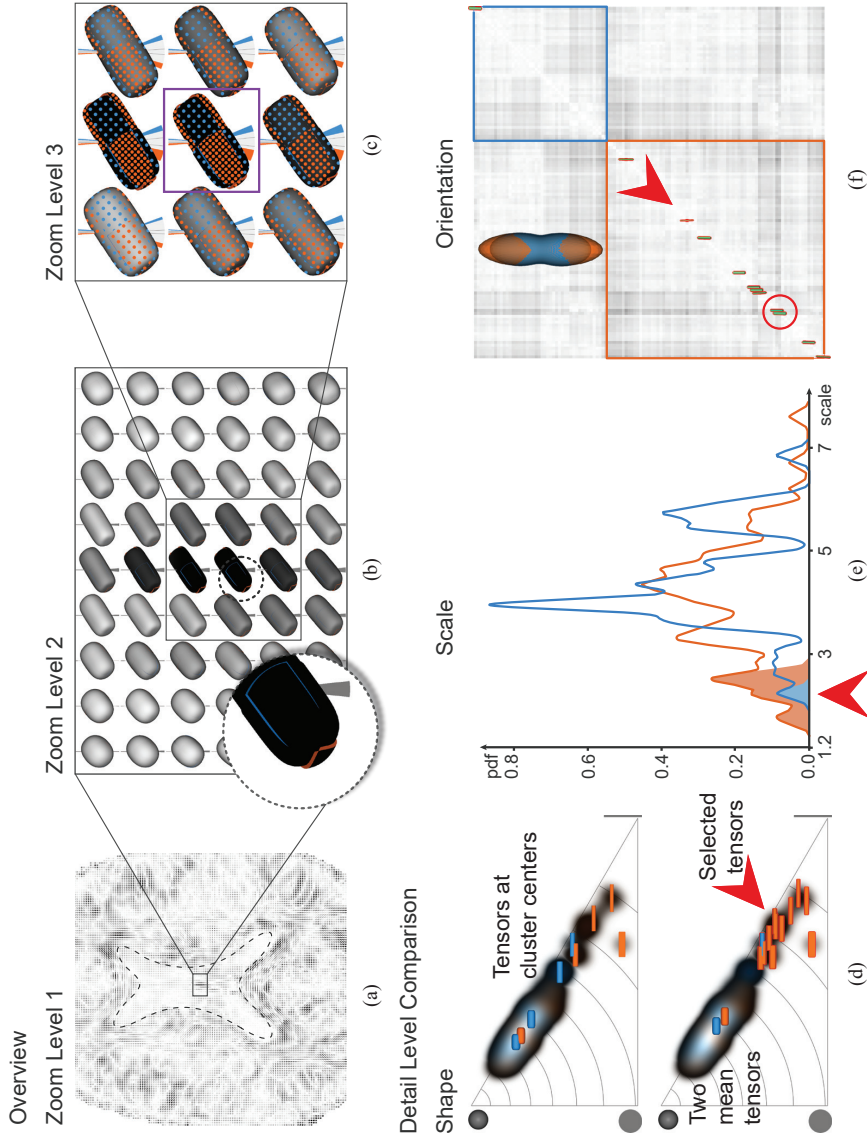
comparison in tensor intrinsic properties for voxels of interest. This level-of-detail glyph visualization enables the gradual exposure of the differences between two tensor ensembles in a voxel-wise manner.

6.3. APPLICATION

We demonstrate the usefulness of our visual analysis prototype by an exemplary comparison of two groups of diffusion tensor fields. The two groups consist of a group of 75 HIV positive (HIV+) patients and a control group with 46 subjects. Each DTI dataset has an original resolution of $224 \times 224 \times 144$ with an isotropic voxel size of 2mm. We uniformly crop each dataset to $121 \times 161 \times 111$ voxels in order to remove those outside of the skull. The diffusion weighting is $b = 1000 \text{ s/mm}^2$ along 64 uniformly distributed directions with four non-diffusion averages. The numerical range of the tensor values is scaled by 1000. The case studies were conducted and interpreted together, by the first and the third author, who is a neuroscientist specialized in DTI group analysis. Throughout the examples, we indicate the HIV+ group by orange color and the control group by blue color.

All datasets are registered with a non-rigid approach provided by DTI-TK [18], which takes the whole tensor information into consideration. The registration quality is important for voxel-wise analysis, however, here, we focus on the comparative visualization of DTI group studies only. Therefore, the influence of the registration quality on the analysis results is beyond the scope of this work. Although we do not explicitly show the detail level comparison for several selected voxels, it is possible to do this with our prototype. Therefore, this type of local averaging can to some extent reduce the influence of potential mis-registration. Please note that for our detail level comparison it makes no difference whether a single voxel or multiple voxels are selected.

Figure 6.4: **Exemplative exploration of the fornix** (roughly indicated by the small rectangle in a). Overview visualization in the spatial domain at different zoom levels, from coarse to fine are shown in (a–c). At zoom level 2 (b), the dark glyphs indicate that the average tensors of the patient group have larger normalized first eigenvalues (e.g., see the inset). At zoom level 3 (c), it is observed that the shape variation for the HIV+ group than that of the control group. The voxel enclosed in the purple rectangle (c) is chosen to be further analyzed in the feature space (d–f). Here a selected group of subjects is indicated by red arrows and shown as additional glyphs with red halos (d, bottom), area under the pdf of tensor scale (e), and glyphs on the diagonal of the orientation dissimilarity matrix (f). The dODF glyphs are used to assist the comparison in tensor orientations.



6.3.1. CASE STUDY 1: FORNIX

The neuroscientist starts the analysis with the visualization of the coarsest zoom level (Figure 6.4a) to identify voxels of interest. This zoom level is constructed based on the tensor shape difference d_{shape} , which represents the difference in the mean diffusion anisotropy profiles between two groups. We can see that the highlighted region is largely white, indicating that the voxels exhibit only small differences in shape. This region roughly matches the ventricle of the brain (see the dashed line in Figure 6.4a). Therefore, these small differences in shape are not surprising, since the diffusion tensors in the ventricles are nearly isotropic. We can see a small region in the center, highlighted by the rectangle in Figure 6.4a, which is notably darker, indicating a larger shape difference. The neuroscientist selects this region, which turns out to be the body of the fornix, for further analysis in zoom level 2 (Figure 6.4b). The fornix is a C-shaped bundle of fibers that roughly travel anteriorly. Note that in zoom level 2 the glyphs are rotated without changing the viewing settings for the sake of a clearer illustration. Here, we can see that the more elongated glyphs exhibit larger pair-wise shape differences (i.e., darker colors). For example, in the central voxels, the discontinuities between the checkerboard sides of the glyphs can be readily perceived (due to size limitations of the figure, see the inset in Figure 6.4b for a clearer view). Glyphs with orange sidefaces have larger normalized major eigenvalues $\tilde{\lambda}_1$ but smaller normalized medium $\tilde{\lambda}_2$ and minor eigenvalues $\tilde{\lambda}_3$. This implies that, on average, the HIV+ group has larger linear anisotropies in the fornix. This finding surprises the neuroscientist, as it conflicts with the common hypothesis that diseases of the nervous system reduce the extent of diffusion anisotropy by damaging the underlying fiber tracts. Furthermore, we can observe that the orientation differences are consistently small for this region, while scale differences (i.e., the angles of the lower sectors) are proportional to the shape differences (i.e., the darkness of the color). The neuroscientist further zooms into the finest level to inspect variations in a three by three voxel area in the center, as marked in Figure 6.4b. In Figure 6.4c we can see that variation in orientation is small and similar for both groups, as indicated by the angles of the upper sectors. The glyphs corresponding to the HIV+ group exhibit larger scale variations (i.e., larger angles in the orange part of the lower sectors) and larger shape variations (i.e., more densely packed orange dots). The fornix on the axial slice consists of only a few voxels and, therefore, partial voluming effects and/or mis-registration can lead to a contamination

by voxels from the surrounding ventricles which normally have larger diffusion scales. The larger variation in the diffusion scales could be an indicator that the HIV+ group is more vulnerable to such contamination. More surprising is the large difference in shape variation between two groups (see the purple rectangle in Figure 6.4c). Therefore, the neuroscientist selects this voxel to compare the original tensor ensembles at the detail level.

First, the neuroscientist inspects the shape distribution. Figure 6.4d top shows the density plot of the shape distribution, overlaid by the glyphs corresponding to the automatically generated density cluster centers. We can see that the tensor shapes are mainly linear for both groups. The overall span is quite large, with FA values ranging from 0.2 to 0.9, as indicated by the glyph locations relative to the background FA contours. Therefore, we select all the tensor glyphs corresponding to large anisotropies (i.e. the glyphs marked by the red arrow in Figure 6.4d, bottom). It can be seen that most of the selected glyphs originate from the HIV+ group (orange color) while only one is from the control group (blue). The neuroscientist speculates that this could be attributed to the scanner shift, as some subjects were scanned with two different scanners. However, we could exclude this hypothesis by inspecting the origin of each scan, which showed that 1) both groups contain subjects scanned with either scanner and 2) the selection includes tensors originating from either scanner (see the supplementary video). Inspecting the detail visualization for the scale (Figure 6.4e), we can see that all of the selected tensors also exhibit small values for tensor scale (see the colored region in Figure 6.4e). The orientation matrix view (Figure 6.4f) does not exhibit such a strong clustering. Overall, we see that, based on the average dODF glyphs with added transparency, the HIV+ group has larger diffusion probability in the vertical direction. We also find two tensors in the HIV+ group (red circle in Figure 6.4f) exhibiting larger orientation difference compared to the rest of the group and some subjects from the control group.

This example demonstrates that our prototype is able to help the neuroscientist to obtain new findings, i.e, the HIV+ group presents both larger linear anisotropies and tensor shape variations for the fornix structure. While the interpretation of the described findings is out of the scope of this article, they inspired the neuroscientist to further verify these findings with a quantitative analysis.

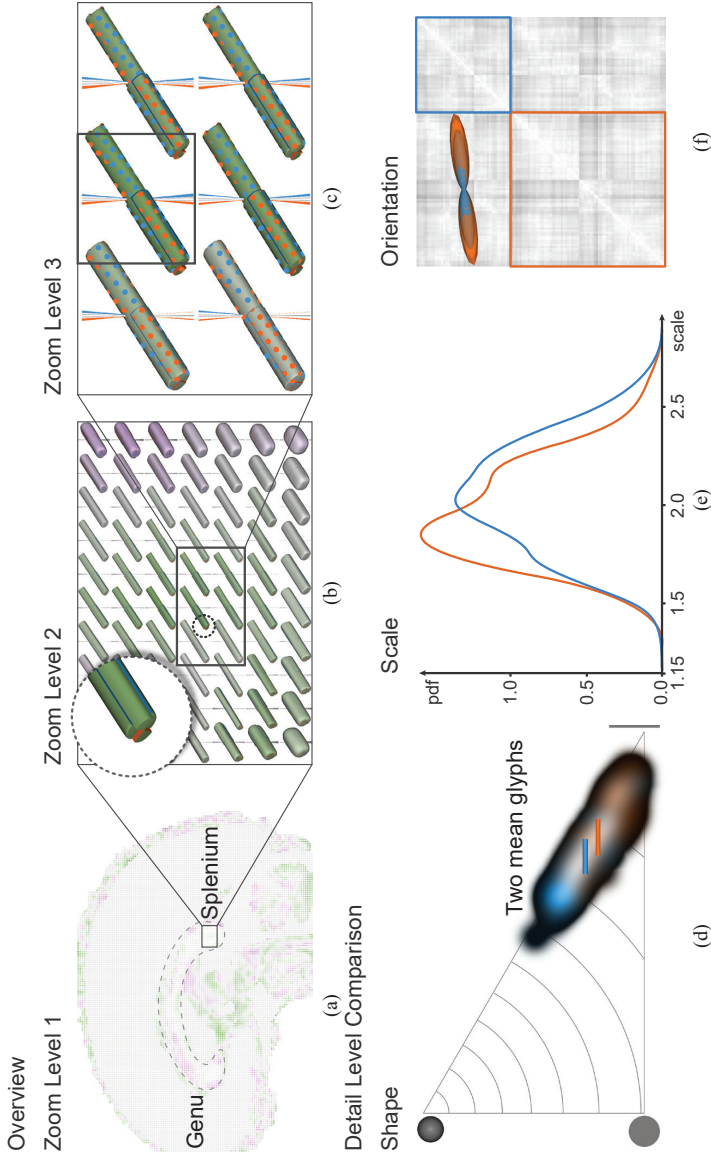


Figure 6.5: **Comparative visualization for the corpus callosum (CC)**, colored according to the t-statistic values with a green-white-violet colormap. (a) shows the coarsest zoom level with the CC marked with the dashed line. (b) shows the mean tensor glyphs at zoom level 2, corresponding to the splenium that is marked by the rectangle in (a). The green color means negative t-statistic values for the splenium. The inset shows the HIV+ group has larger normalized major eigenvalues, as shown by the discontinuity, exposing the orange sideface. (c) shows the visual comparison of 6 selected voxels at zoom level 3 indicating rather small difference between the two groups. (d, e, f) show the detail level comparisons for the shape, scale, and orientation, respectively.

6.3.2. CASE STUDY 2: CORPUS CALLOSUM

In the second case study we inspect the corpus callosum (CC), which is a flat bundle of fibers connecting the left and right cerebral hemispheres. It is indicated by the dashed line in a sagittal slice shown in Figure 6.5a. The neuroscientist starts with the overview visualization on zoom level 1, employing the commonly used scalar-valued statistical measure, i.e., t-statistic. The neuroscientist hypothesizes that the FA values are higher in the control group than in the HIV+ group. A positive/negative t-statistic provides evidence for/against the hypothesis. The absolute value of the t-statistic indicates the extent to which the FA values are significantly different between two groups. *Significant difference* intuitively means large difference in the FA values of the mean tensors and small distribution overlaps. The voxel-wise t-statistic is constructed following the standard routine in FSL based on the FA values. We use a diverging green-white-violet colormap to encode t-statistic results (green=negative, violet=positive), since the values for the t-statistic are centered around zero and specifically this green to violet one, as it does not interfere with the blue and orange colors used for the different groups. In zoom level 1 (Figure 6.5a), it can be seen that the anterior part of CC (i.e., the genu) appears to be more violet while the posterior part (i.e., the splenium) exhibits some green. The neuroscientist is specifically interested in these voxels with negative t-statistic values and zooms into level 2 for this region (Figure 6.5b). We observe that the HIV+ group, indicated by orange sidefaces (see the inset in Figure 6.5b), has larger normalized major eigenvalues $\tilde{\lambda}_1$ and smaller normalized medium and minor eigenvalues, similar to what we have found in the fornix. Furthermore, the scale and orientation differences are quite small, as indicated by the small lower and upper sectors, respectively. For further comparison of the group variations, we then zoom into level 3 (see Figure 6.5c). It can be seen that the variations in scale, shape, and orientation are quite similar between the two groups. We now inspect the complete distributions for a single voxel, indicated by the square in Figure 6.5c. In the shape space visualization (Figure 6.5d), we can immediately see that the two mean tensor glyphs are placed at some distance. In combination with the emphasized blue and orange tones on the ends of the density plot this confirms that the FA is indeed larger for the HIV+ group, as indicated by the green color in Figures 6.5a–c. As discussed before, diseases like HIV damage the fiber structures, which then makes the water diffuse more freely. Consequently, we expect the HIV+ patients to exhibit a larger amount of

diffusion [19]. However, inspecting Figure 6.5e, we find that the distribution of the tensor scale for the HIV+ group (orange line) is slightly more skewed towards a lower scale value compared to the control group (blue). This means that a subject from the HIV+ group has a higher probability to present a small amount of diffusion, compared to a subject from the control group. This finding contradicts the expectation that HIV+ patients would exhibit a larger amount of diffusion. While we cannot present a solid interpretation for this finding in the scope of this chapter, it shows that our prototype is able to reveal more information compared to the standard FA-based statistical analysis of DTI inter-group comparison.

6.4. CONCLUSIONS AND FUTURE WORK

In this work, we present a level-of-detail comparative visualization technique to support the comparison between two groups of diffusion tensor fields. The comparison is divided into two parts, i.e., the comparison of tensor ensemble summaries in the spatial domain and that of the original tensor ensembles in terms of tensor intrinsic properties. For the comparison of tensor ensemble summaries, our glyph representation is based on the combination of two existing glyph designs in Chapters 4 and 5, respectively. More importantly, we integrate the details-on-demand concept to present more information as users zoom in. The coarse zoom level is designed to show large scale patterns of user specified scalar-valued difference measures, mimicking slice views. The medium zoom level facilitates the comparison of the mean tensors of the two groups in terms of their tensor intrinsic properties. The glyph representation at this level is a variant of the Tender glyph. The finest zoom level additionally supports comparison of variation information of each group. The glyph representation at this level is an extension of the tensor ensemble summary glyph. Here, the glyph-based comparative visualization helps analysts to locate voxels of interest for further comparison in tensor intrinsic properties at the detail level. Therefore, we extend the detail visualization in Chapter 5 to support the visual comparison of two original tensor ensembles by exploiting the superposition and explicit encoding strategy. We performed two case studies together with a neuroscientist to illustrate the usefulness of our prototype by comparing an HIV+ patient group and a control group. Our prototype successfully reveals new findings that are otherwise unknown with standard scalar-based comparison. The case studies triggered the neuroscientist's interest to perform a formal study to further

inspect these new findings.

There is still room for improvement as future work. Regarding the glyph design, the sector representation is not intuitive since it relies on the users' memory of the correspondence to tensor scale or orientation. Furthermore, although the image-space textures, used to indicate shape variation, provide a good estimation of the variation, their image-space nature can cause annoying effects during interaction. Research regarding these issues could further improve the glyph representation. In the spatial visualizations, the glyphs are placed on a regular grid, which is likely to hide the underlying continuous structures of the field. The combination of a smart sampling scheme [20] with the level of detail glyph representations could further improve the overview visualizations. Finally, the next logical step would be the comparison of more than two groups.

REFERENCES

- [1] C. Zhang, T. Höllt, M. W. Caan, E. Eisemann, and A. Vilanova, *Comparative visualization for diffusion tensor imaging group study at multiple levels of detail*, in *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)* (2017) pp. 53–62.
- [2] A. Schwartzman, R. F. Dougherty, and J. E. Taylor, *Cross-subject comparison of principal diffusion direction maps*, *Magnetic Resonance in Medicine* **53**, 1423 (2005).
- [3] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens, *Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data*, *NeuroImage* **31**, 1487 (2006).
- [4] G. Douaud, C. Mackay, J. Andersson, S. James, D. Quested, M. K. Ray, J. Connell, N. Roberts, T. J. Crow, P. M. Matthews, S. Smith, and A. James, *Schizophrenia delays and alters maturation of the brain in adolescence*, *Brain* **132**, 2437 (2009).
- [5] W. Van Hecke, G. Nagels, G. Emonds, A. Leemans, J. Sijbers, J. Van Goethem, and P. M. Parizel, *A diffusion tensor imaging group study of the spinal cord in*

- multiple sclerosis patients with and without T2 spinal cord lesions*, *Journal of Magnetic Resonance Imaging* **30**, 25 (2009).
- [6] P. J. Basser and S. Pajevic, *A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI*, *IEEE Transactions on Medical Imaging* **22**, 785 (2003).
- [7] B. Whitcher, J. J. Wisco, N. Hadjikhani, and D. S. Tuch, *Statistical group comparison of diffusion tensors via multivariate hypothesis testing*, *Magnetic Resonance in Medicine* **57**, 1065 (2007).
- [8] B. Shneiderman, *The eyes have it: A task by data type taxonomy for information visualizations*, in *IEEE Symposium on Visual Languages* (1996) pp. 336–343.
- [9] O. Dzyubachyk, J. Blaas, C. P. Botha, M. Staring, M. Reijnders, J. L. Bloem, R. J. van der Geest, and B. P. F. Lelieveldt, *Comparative exploration of whole-body mr through locally rigid transforms*, *International Journal of Computer Assisted Radiology and Surgery* **8**, 635 (2013).
- [10] A. Abbasloo, V. Wiens, M. Hermann, and T. Schultz, *Visualizing tensor normal distributions at multiple levels of detail*, *IEEE Transactions on Visualization and Computer Graphics* **22**, 975 (2016).
- [11] S. Pajevic and C. Pierpaoli, *Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain*, *Magnetic Resonance in Medicine* **42**, 526 (1999).
- [12] C.-F. Westin, S. E. Maier, B. Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis, *Image processing for diffusion tensor magnetic resonance imaging*, in *MICCAI* (1999) pp. 441–452.
- [13] I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel, *Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle*, *Magnetic Resonance in Medicine* **64**, 554 (2010).

- [14] O. D. Lampe and H. Hauser, *Interactive visualization of streaming data with kernel density estimation*, in *2011 IEEE Pacific Visualization Symposium* (2011) pp. 171–178.
- [15] T. Höllt, N. Pezzotti, V. van Unen, F. Koning, E. Eisemann, B. P. Lelieveldt, and A. Vilanova, *Cytosplore: Interactive immune cell phenotyping for large single-cell datasets*, *Computer Graphics Forum* **35**, 171 (2016).
- [16] K. Fukunaga and L. Hostetler, *The estimation of the gradient of a density function, with applications in pattern recognition*, *IEEE Transactions on Information Theory* **21**, 32 (1975).
- [17] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004) pp. 147–154.
- [18] H. Zhang, P. A. Yushkevich, D. C. Alexander, and J. C. Gee, *Deformable registration of diffusion tensor MR images with explicit orientation optimization*, *Medical Image Analysis* **10**, 764 (2006).
- [19] Y. Wu, P. Storey, B. Cohen, L. Epstein, R. Edelman, and A. Ragin, *Diffusion alterations in corpus callosum of patients with hiv*, *American Journal of Neuro-radiology* **27**, 656 (2006).
- [20] G. Kindlmann and C.-F. Westin, *Diffusion tensor visualization with glyph packing*, *IEEE Transactions on Visualization and Computer Graphics* **12**, 1329 (2006).

7

CONCLUSION

We have to change from doubters to believers. Now!

Jürgen Klopp

7.1. CONTRIBUTIONS

The analysis of multiple 3D tensor fields is an important but challenging task in the context of DTI-based neuroscientific studies. However, not much work has yet been done due to the complex and multivariate nature of tensor data. In this thesis, we have presented several novel visual analysis methods to assist the exploration and analysis of multiple 3D tensor fields.

The acquisition of DTI datasets relies on MRI technique. One major disadvantage of MRI is that it is a time-consuming process, which restricts the applicability of DTI. Compressed sensing (CS) has the potential to accelerate MRI acquisition. Therefore, a brief introduction of MRI and CS is given in Chapter 2. In Section 2.3, this thesis contributes by adapting the traditional CS MRI reconstruction method to incorporate

local information, and applies it to preserve low-contrast vessels in MR angiography as a proof-of-concept.

The main aim of this thesis is to design understandable visual analysis techniques for multiple 3D tensor fields while keeping the full tensor information. The main contributions are summarized as follows.

- The dissimilarity measures proposed in Chapter 4 enable the quantification of pair-wise differences of diffusion tensors in tensor scale, shape, and orientation. Compared to previous method [1], the new dissimilarity measures are computationally more efficient. More importantly, these dissimilarity measures can be readily interpreted by neuroscientists. To visually encode such differences in scale, shape, and orientation, a checkerboard-style glyph is designed by exploiting the fact that human eyes are more sensitive to discontinuity. Thus, this new glyph is able to facilitate the visual comparison between two DTI datasets in a voxel-wise manner compared to straightforward juxtaposition and/or superposition.
- We extend the pair-wise differences measures in Chapter 5 to quantify the variation in tensor scale, shape, and orientation of an ensemble of diffusion tensors. Compared to the fourth-order covariance tensor [2] that is commonly used to characterize tensor variations, our measures provide a more direct interpretation. An overview + detail framework is proposed in Chapter 5 to visualize an ensemble of diffusion tensor fields. We extend the superquadric glyph [3] to additionally encode the variation information in order to provide an overview of the voxel-wise tensor ensemble in the spatial domain. The glyph-based overview comprehensively manifests the multi-variate ensemble summary statistics (i.e., mean tensor and tensor variations). The detail visualization is composed of three linked sub-views, each of which depicts the original tensor ensembles for user specified locations from exactly one aspect of the tensor intrinsic properties (scale, shape, or orientation). By doing so, we manage to support the exploration and understanding of the tensor ensemble without assuming a specific distribution.
- The checkerboard-style glyph and the overview + detail framework are combined and extended in Chapter 6 to support the visual comparison of two ensembles of

DTI datasets. In particular, the level-of-detail concept is integrated into the glyph design to progressively reveal information as users zoom in so that the glyph-based visual comparison is improved. In this way, tensor ensemble summary statistics can be effectively compared at multiple levels.

- We carefully design the user studies and case studies to evaluate the visualization techniques presented in Chapters 4, 5, and 6 to validate the effectiveness of the visual encodings and show the potential usability of the whole visual analysis framework. Since not much work has yet been done concerning the design of a user study for glyph-based tensor field visualization, the user studies in this thesis can serve as a reference for other tensor field visualization work.

7.2. FUTURE WORK

Possible directions of future work related to specific visualization techniques are listed at the end of each individual chapter, and all the possible directions are summarized here to provide a compact overview.

In this thesis, the main object of research is the diffusion tensor, which mathematically is a second-order symmetric positive-definite tensor. Therefore, one natural question to ask is whether the analysis and visualization techniques proposed in the thesis can be extended to second-order symmetric positive-definite tensors but not in the context of DTI, second-order symmetric tensors (not necessarily positive-definite), general second-order tensors (not necessarily symmetric), or even higher-order tensors. Since second-order symmetric tensors have six degrees-of-freedom, the pairwise dissimilarity measure can be directly extended. The major issue is that whether the quantification of tensor dissimilarity in the three aspects (i.e., scale, shape, and orientation) is useful or not is application-dependent. The tensor ensemble summary glyph design in Chapter 5 can be directly extended to second-order symmetric tensors. However, the extension of the checkerboard-style comparison strategy to second-order symmetric tensors requires additional efforts. For general second-order tensors and higher-order tensors, the visualization of a single field itself is already quite challenging [4], not to mention that of multiple fields. One well-known disadvantage of the diffusion tensor model is that it fails to model complex diffusion pattern for regions with crossing fibers while HARDI data is able to do. However, HARDI data generally

has more degrees-of-freedom than second-order symmetric tensors. This makes the extension of comparative and ensemble visualization techniques much more difficult.

From the perspective of neuroscientific studies, multiple group comparisons are commonly conducted. The next logical step would be the comparison of multiple ensembles of diffusion tensor fields. From the perspective of validation, the user/case studies conducted in this thesis normally require knowledge about diffusion tensor analysis, which limits the pool of available participants. Therefore, we consider an extensive evaluation that would allow for strong conclusions on the effectiveness of the presented visual analysis techniques as future work.

We believe the visual analysis techniques presented in this thesis will promote research in the direction of visual analysis of multiple tensor fields.

REFERENCES

- [1] G. Kindlmann, R. San José Estépar, M. Niethammer, S. Haker, and C.-F. Westin, *Geodesic-loxodromes for diffusion tensor interpolation and difference measurement*, in *MICCAI* (2007) pp. 1–9.
- [2] P. Basser and S. Pajevic, *Spectral decomposition of a 4th-order covariance tensor: applications to diffusion tensor MRI*, *Signal Processing* **87**, 220 (2007).
- [3] G. Kindlmann, *Superquadric tensor glyphs*, in *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004) pp. 147–154.
- [4] T. Gerrits, C. Rössl, and H. Theisel, *Glyphs for general second-order 2D and 3D tensors*, *IEEE Transactions on Visualization and Computer Graphics* **23**, 980 (2017).

CURRICULUM VITÆ

Changgong ZHANG

24-01-1988 Born in Binzhou, China.

EDUCATION

2002–2005	Beizhen Senior High School Binzhou, Shandong, China
2005–2009	Undergraduate in Mechanical Engineering and Automation Beihang University, Beijing, China
2009–2012	Graduate in Aerospace Manufacturing Engineering Beihang University, Beijing, China
2012-2017	PhD. Scientific Visualization Delft University of Technology, Delft, The Netherlands <i>Thesis:</i> Comparative and Ensemble Visualization of Diffusion Tensor Fields <i>Promotor:</i> Prof. dr. E. Eisemann <i>Copromotor:</i> Dr. A. Vilanova

LIST OF PUBLICATIONS

4. **C. Zhang**, T. Höllt, M.W.A. Caan, E. Eisemann, and A. Vilanova, *Comparative Visualization for Diffusion Tensor Imaging Group Study at Multiple Levels of Detail*, [Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine \(VCBM\)](#), 53 (2017).
3. **C. Zhang**, M.W.A. Caan, T. Höllt, E. Eisemann, and A. Vilanova, *Overview + Detail Visualization for Ensembles of Diffusion Tensors*, [Computer Graphics Forum](#) **36**, 121 (2017).
2. **C. Zhang**, T. Schultz, K. Lawonn, E. Eisemann, and A. Vilanova, *Glyph-based Comparative Visualization for Diffusion Tensor Fields*, [IEEE Transactions on Visualization and Computer Graphics](#) **22**, 797 (2015).
1. **C. Zhang**, M. van de Giessen, E. Eisemann, and A. Vilanova, *User-guided compressed sensing for magnetic resonance angiography*, [Proceedings of Annual International Conference of the IEEE Engineering in Medicine and Biology Society \(EMBC\)](#), 2416 (2014).