

VISUAL ANALYTICS METHODS FOR RETINAL LAYERS IN
OPTICAL COHERENCE TOMOGRAPHY DATA



Dissertation
submitted for the academic degree of

Doktoringenieur (Dr.-Ing.)

Faculty of Computer Science and Electrical Engineering,
University of Rostock, Germany

submitted by
Martin Röhlig
born on 21.01.1984 in Lichtenstein (Saxony)
resident in Rostock, Germany

Rostock, September 30, 2020

REVIEWERS:

Prof. Dr.-Ing. habil. Heidrun Schumann, University of Rostock
Prof. Dr.-Ing. habil. Bernhard Preim, University of Magdeburg

DATE OF SUBMISSION:

September 30, 2020

DATE OF DEFENSE:

January 07, 2021

LOCATION:

Rostock

Martin Röhlig: *Visual Analytics Methods for Retinal Layers in Optical Coherence Tomography Data*, 2020

This work is licensed under a [Creative Commons "Attribution 4.0 International"](https://creativecommons.org/licenses/by/4.0/) license.



ABSTRACT

Optical Coherence Tomography (OCT) is an important non-invasive imaging technology for the early detection of various eye diseases. Despite the diagnostic value of OCT, the identification of substructural defects in the high-resolution 3D images of the retina is challenging. In this thesis, we therefore introduce novel Visual Analytics (VA) methods for the localization and analysis of early retinal alterations. Specifically, we present solutions for two fundamental problems in the VA of complex 3D spatial data: (1) ensuring the visibility of relevant information in the 3D spatial context and (2) reducing the data complexity to meaningful characteristics. These solutions not only help to cope with the problems of complex 3D spatial data in general, but especially to gain diagnostically valuable insights from the retinal OCT data. Together with ophthalmologists and industry experts, we evaluated our VA methods in two cross-sectional studies with patients suffering from diabetes mellitus. The results show that they improve the detection of small and localized retinal defects and thus contribute to a deeper understanding of the condition of the retina at an early stage of disease.

ZUSAMMENFASSUNG

Die optische Kohärenztomographie (OCT) ist eine wichtige nichtinvasive Bildgebungstechnologie für die Früherkennung verschiedener Augenkrankheiten. Die Identifizierung von substrukturellen Defekten in den hochauflösenden 3D-Bildern der Netzhaut ist jedoch eine Herausforderung. In dieser Arbeit stellen wir daher neuartige Visual-Analytics-Methoden (VA-Methoden) zur Lokalisierung und Analyse früher Netzhautveränderungen vor. Konkret präsentieren wir Lösungen für zwei grundlegende Probleme bei der VA von komplexen 3D-Raumdaten: (1) die Sicherstellung der Sichtbarkeit relevanter Informationen im 3D räumlichen Kontext und (2) die Reduzierung der Datenkomplexität auf aussagekräftige Merkmale. Diese Lösungen helfen nicht nur generell bei der Auswertung von komplexen 3D-Raumdaten, sondern auch besonders bei der Gewinnung diagnostisch wertvoller Erkenntnisse aus den OCT-Daten der Netzhaut. Unsere VA-Methoden haben wir zusammen mit Augenärzten und Industrieexperten in zwei Querschnittsstudien mit Patienten, die an Diabetes mellitus leiden, evaluiert. Die Ergebnisse zeigen, dass wir die Erkennung von kleinen und lokal begrenzten Netzhautdefekten verbessern und damit zu einem tieferen Verständnis des Zustandes der Netzhaut in frühen Krankheitsstadien beitragen konnten.

To my family.

SELECTED PUBLICATIONS

This thesis is based on the following journal publications (in chronological order). A full list of the author's published work is available starting with page [229](#).

1. M. Luboschik, **M. Röhlig**, A. T. Bittig, N. Andrienko, H. Schumann, and C. Tominski. "Feature-Driven Visual Analytics of Chaotic Parameter-Dependent Movement." In: *Computer Graphics Forum* 34.3 (2015). Best paper award, pp. 421–430. DOI: [10.1111/cgf.12654](https://doi.org/10.1111/cgf.12654).
2. **M. Röhlig**, M. Luboschik, and H. Schumann. "Visibility Widgets for Unveiling Occluded Data in 3D Terrain Visualization." In: *Journal of Visual Languages & Computing* 42.Supplement C (2017), pp. 86–98. DOI: [10.1016/j.jvlc.2017.08.008](https://doi.org/10.1016/j.jvlc.2017.08.008).
3. **M. Röhlig**, C. Schmidt, R. K. Prakasam, P. Rosenthal, H. Schumann, and O. Stachs. "Visual Analysis of Retinal Changes with Optical Coherence Tomography." In: *The Visual Computer* 34.9 (2018), pp. 1209–1224. DOI: [10.1007/s00371-018-1486-x](https://doi.org/10.1007/s00371-018-1486-x).
4. R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Götze, A. Jünemann, H. Schumann, and O. Stachs. "Deviation maps for understanding thickness changes of inner retinal layers in children with type 1 diabetes mellitus." In: *Current Eye Research* 44.7 (2019). DOI: [10.1080/02713683.2019.1591463](https://doi.org/10.1080/02713683.2019.1591463).
5. **M. Röhlig**, R. K. Prakasam, J. Stüwe, C. Schmidt, O. Stachs, and H. Schumann. "Enhanced Grid-Based Visual Analysis of Retinal Layer Thickness with Optical Coherence Tomography." In: *Information* 10.9 (2019), p. 266. DOI: [10.3390/info10090266](https://doi.org/10.3390/info10090266).
6. R. K. Prakasam, A. Matuszewska-Iwanicka, D.-C. Fischer, H. Schumann, D. Tschöpe, B. Stratmann, H.-J. Hettlich, R. F. Guthoff, O. Stachs, and **M. Röhlig**. "Thickness of intraretinal layers in patients with type 2 diabetes mellitus depending on a concomitant diabetic neuropathy: results of a cross-sectional study using deviation maps for OCT data analysis." In: *Biomedicines* 8.7 (2020). DOI: [10.3390/biomedicines8070190](https://doi.org/10.3390/biomedicines8070190).

ACKNOWLEDGMENTS

Although this thesis was written by one, the work presented here would not have been possible without the support and participation of others. I am deeply grateful for the opportunity to have met, learned from, and worked with so many great people.

First and foremost, I thank Heidrun Schumann for being a truly exceptional mentor on a personal and professional level. Really special to me are your family-like supervision and your unparalleled working environment, which gave me the necessary security and confidence in my abilities. Your invaluable advice and everlasting patience guided me through this thesis, while keeping motivation high and pressure low in times of struggle. As someone who tends to have the work in his mind 90 % ready at any time, I thank you especially for your support in completing this thesis 100 %.

I further thank Bernhard Preim in particular for agreeing to evaluate this thesis. It is a great honor to have such an outstanding scientist as a reviewer and to have received your valuable feedback.

My special thanks also go to our collaboration partners at the Rostock University Medical Center, in particular Dr. Ruby Kala Prakasam, Prof. Dagmar-Christiane Fischer, and Prof. Oliver Stachs, who not only believed in my work, but were willing to break out of established work routines and actually try something new with their data. Without your profound knowledge and close participation, it would have been impossible to achieve the included results.

I am very grateful to my colleagues in and around the visual computing and computer graphics group at the University of Rostock. A special thanks to Martin Luboschik, Christian Tominski, Steve Dübel, and Christoph Schmidt for the many open-minded and in-depth discussions far beyond the topics covered here. You are my dear friends, who have shaped my way of thinking and ultimately me as a person. It has been great working with you, and with regard to this thesis, I thank you especially for helping me to “focus” and “finish it”.

Finally, I thank those who are most important in my life – my family. You have helped me realign the thoughts in my head and focus on what really matters, at times when it was necessary, and even on some occasions when I thought it was not, at least initially. Mizuho, Karl, and Hanna, you are my source of joy, happiness, love, and life.

CONTENTS

I THE INTRO

1	INTRODUCTION	3
1.1	Motivation	3
1.2	Medical background	4
1.3	Challenges	6
1.4	Contribution	7
1.5	Outline	9
2	FUNDAMENTAL PROBLEMS IN VISUAL ANALYTICS OF COMPLEX 3D SPATIAL DATA	11
2.1	Design options for the visualization of 3D spatial data	11
2.2	Visibility widgets for unveiling occluded data in 3D visualization	15
2.3	Feature visualization for reducing data complexity	23
2.4	Summary	30
3	NOVEL APPROACHES FOR VISUAL ANALYTICS OF RETINAL OCT DATA	31
3.1	Improvements in the visibility of retinal layers	31
3.2	Reductions in the complexity of retinal layers	37
3.3	Summary	47
4	APPLICATION OF VISUAL ANALYTICS METHODS IN OPHTHALMOLOGY	49
4.1	Evaluation of the analysis procedure performed with visual analytics methods	49
4.2	Evaluation of the analysis results obtained with visual analytics methods	52
4.3	Summary	58
5	SUMMARY	59

II THE CORE

6	VISIBILITY WIDGETS FOR UNVEILING OCCLUDED DATA IN 3D TERRAIN VISUALIZATION	63
6.1	Introduction	64
6.2	Background and problem statement	66
6.3	Related work	69
6.4	Visibility widgets	72
6.5	Implementation	84
6.6	User feedback	86
6.7	Generalization	88
6.8	Discussion and conclusion	89
7	FEATURE-DRIVEN VISUAL ANALYTICS OF CHAOTIC PARAMETER-DEPENDENT MOVEMENT	93

7.1	Introduction	94
7.2	Motivation and related work	95
7.3	Feature extraction	98
7.4	Feature visualization and interaction	101
7.5	Application to systems biology	106
7.6	Discussion and conclusion	110
8	VISUAL ANALYSIS OF RETINAL CHANGES WITH OPTICAL COHERENCE TOMOGRAPHY	113
8.1	Introduction	114
8.2	Background	115
8.3	Related work	118
8.4	Requirements	120
8.5	Unify, structure & enhance data	122
8.6	Visualize, emphasize, select & compare data	124
8.7	Application	132
8.8	User feedback	134
8.9	Discussion and conclusion	136
9	ENHANCED GRID-BASED VISUAL ANALYSIS OF RETINAL LAYER THICKNESS WITH OPTICAL COHERENCE TOMOGRAPHY	139
9.1	Introduction	140
9.2	Background	142
9.3	Grid design	145
9.4	Grid exploration	149
9.5	Grid-based visual analytics of retinal layer thickness	155
9.6	Application	161
9.7	Discussion and conclusions	166
10	DEVIATION MAPS FOR UNDERSTANDING THICKNESS CHANGES OF INNER RETINAL LAYERS IN CHILDREN WITH TYPE 1 DIABETES MELLITUS	169
10.1	Introduction	171
10.2	Materials and methods	172
10.3	Results	176
10.4	Discussion	179
11	THICKNESS OF INTRARETINAL LAYERS IN PATIENTS WITH TYPE 2 DIABETES MELLITUS DEPENDING ON A CONCOMITANT DIABETIC NEUROPATHY: RESULTS OF A CROSS-SECTIONAL STUDY USING DEVIATION MAPS FOR OCT DATA ANALYSIS	183
11.1	Introduction	184
11.2	Materials and methods	186
11.3	Results	189
11.4	Discussion	193
11.5	Conclusion	196
III THE OUTRO		
12	CONCLUSION AND FUTURE WORK	199
12.1	Concluding remarks	199

12.2 Topics for future work 200

BIBLIOGRAPHY 203

LIST OF SCIENTIFIC CONTRIBUTIONS 229

ACRONYMS

VA	Visual Analytics
OCT	Optical Coherence Tomography
B-scan	Cross-sectional Tomographic Depth Scan
ETDRS	Early Treatment Diabetic Retinopathy Study
AMD	Age-related Macular Degeneration
T ₁ DM	Type 1 Diabetes Mellitus
T ₂ DM	Type 2 Diabetes Mellitus
TR	Total Thickness of the Retina
RNFL	Retinal Nerve Fiber Layer
GCL	Ganglion Cell Layer
IPL	Inner Plexiform Layer
INL	Inner Nuclear Layer
RPE	Retinal Pigment Epithelium
SOM	Self-Organizing Map
MA	Map-based Analysis
TM	Thickness Map
RM	Reference Map
DevM	Deviation Map
CA	Current Analysis
DM	Diabetes Mellitus
DR	Diabetic Retinopathy
DN	Diabetic Peripheral Neuropathy
DN _{neg}	Diabetic Peripheral Neuropathy Negative
DN _{pos}	Diabetic Peripheral Neuropathy Positive
FC	Foveal Center
IMR	Inner Macular Ring
OMR	Outer Macular Ring
MD ± SE	Mean Deviation and Standard Error
MD ± SD	Mean Deviation and Standard Deviation
NSS	Neuropathy Symptom Score
NDS	Neuropathy Deficient Score
cSLO	Confocal Scanning Laser Ophthalmoscope

HSD	Honestly Significant Difference
UOCTE	Unified OCT Explorer
UOCTML	Unified OCT Markup Language
XML	Extensible Markup Language

Part I

THE INTRO

This first part summarizes the key ideas and results of this work. A common understanding is established by developing a structured view on relevant open research questions. On this basis, novel visual analytics solutions for complex 3D spatial data in general and optical coherence tomography data of the retina in particular are presented.

INTRODUCTION

This work deals with the development of novel Visual Analytics (VA) methods for retinal data acquired via Optical Coherence Tomography (OCT) [1]. Starting from basic investigations regarding the complexity and presentation of the data, a VA tool is proposed that helps to analyze exactly the information that is relevant for a medical diagnosis. The implementation is evaluated in two cross-sectional studies on patients with diabetes mellitus and compared with current analysis methods.

1.1 MOTIVATION

Advances in medical imaging technology enable the acquisition of data about the human body with ever higher resolution, quality, and accuracy. In ophthalmology, one of the rapidly advancing technologies is OCT. OCT is a noninvasive technique to capture 3D images of the human retina [2]. Ophthalmologists interpret the OCT to understand a variety of retinal and systemic disorders [3]. This includes common eye diseases in the focus of public health interest, such as Age-related Macular Degeneration (AMD) [4], diabetic retinopathy [5], or glaucoma [6], and other pathologies with ocular signs, such as multiple sclerosis [7]. Over the years, the number of ophthalmic OCT examinations has steadily increased, rising globally to about 30 million per year by 2011 [8] or to about 6.7 million in the US Medicare population alone by 2017 [9]. Likewise, the funding and publishing of OCT-based research in ophthalmology have grown continuously [8], e.g., with over 2800 new bibliographic references in the MEDLINE database and related biomedical literature in 2019 [10]. This progress in OCT image acquisition and availability is accompanied by increasing efforts to effectively process and analyze the data.

One OCT dataset can consist of hundreds of cross-sections of retinal tissue with micrometer resolution. Looking at all OCT examinations of patients in clinics worldwide, the available datasets today accumulate to an enormous amount of data. Ophthalmologists are using the data for their analytical tasks. However, due to their size and complexity, it is difficult to translate the growing number of volumetric images into insights. In particular, it is a challenge to identify and consider exactly the information within the overall data that is relevant for the assessment of a specific patient or the evaluation of a particular research study. Current analysis methods are based on the inspection of individual cross-sections to examine details, or on the application of automated data simplifications to obtain a quick overview of the data.

This is accompanied by a high manual analysis effort and the risk of losing valuable information.

In this context, VA [11] has become a key technology that helps experts understand complex data, make decisions, and generate knowledge. VA is one of the top research areas when it comes to the analysis of medical data and beyond. Generally, it is considered a combination of:

...automated analysis techniques with interactive visualizations for an effective understanding, reasoning and decision making on the basis of very large and complex datasets [12]

This combination of automated, visual, and interactive means provides a unique scientific foundation for the development of novel techniques that meet the needs of ophthalmologists. So far, however, visualization research on the peculiarities of retinal OCT data and the associated requirements in ophthalmology has been limited. The actual design of suitable VA solutions is difficult and the requirements cannot be met with a single technique. In order to contribute to the understanding of today's volumes and diversity of retinal OCT data, further research on the data complexity and the visibility of diagnosis-relevant information is mandatory. In this work, our goal is to develop new VA methods to address these problems and to provide them in a comprehensive tool for application in ophthalmology.

1.2 MEDICAL BACKGROUND

The diagnostic value of OCT arises out of the high axial and lateral resolution of modern scanning devices. Nowadays, OCT scanners provide the unique possibility of differentiating and presenting fine-grained substructures of the retina – *the retinal layers* – in only few seconds [13, 14], close to an in-vivo 'optical biopsy' of the retina [15]. Figure 1.1 gives a first impression of the retinal layers in the OCT image data. Using the images, ophthalmologists can assess photoreceptor defects, abnormal changes in axonal thickness, or macular degeneration and edema [16]. Consequently, OCT examinations have become an essential means for assessing and monitoring retinal conditions in daily clinical practice and in ophthalmic research [17].

Across ophthalmic clinics, OCT examinations are widely adapted for triage of individual patients requiring rapid non-elective assessment of acute and chronic sight loss [18, 19]. Expert clinicians utilize OCT findings and clinical history to diagnose patients with a pathology affecting the macula, the central part of the retina that is required for high-resolution color vision. Furthermore, OCT has become a standard tool for comprehensive baseline assessment of the retina prior to initiation of therapy and monitoring therapeutic effect [20]. On top of that, in ophthalmic research, OCT examinations play a fundamental role in

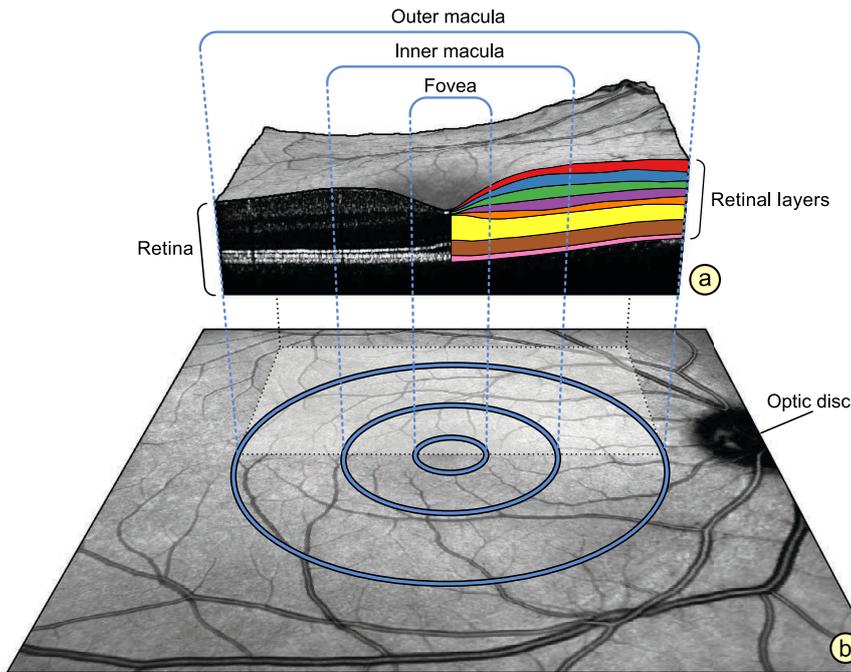


Figure 1.1: Illustration of the multi-layered retina in the human eye. A section of a 3D OCT image of the macula with highlighted retinal layers (a) is depicted above an image of the inner surface of the eye (b).

studies on the basic principles of different ocular and neurodegenerative diseases. Ophthalmic researchers evaluate OCT data to identify disease-induced alterations of the retina and to derive respective ocular biomarkers or surrogate measures. The high sensitivity of OCT enables them to reveal subtle retinal defects that occur at a very early stage of diseases even before any signs of complications in results of other examination methods [21, 22]. The gathered research insights are eventually fed back into clinical practice to improve the diagnosis and staging of retinal diseases and to implement appropriate interventions.

Both in clinical practice and in ophthalmic research, ophthalmologists need to analyze retinal OCT data for their diagnoses. The data analysis requirements vary greatly depending on the type and stage of retinal disease. At an advanced stage of a disease, pathological findings can be noticeably different from a healthy retina. Determining where exactly and to what degree alterations occur in the complex 3D image data is complicated, however. This makes the early detection of retinal diseases all the more important. Yet, in the early stages of a disease it is even harder to gain insights, as pathological alterations of the retina tend to appear more subtle, localized, and scattered throughout the data. Correct identification of these characteristics is necessary for precise early management in retinal diseases and guiding decisions of whether a certain therapy should be initiated or not. That is why subsidiary data measurements are carried out to extract further details on the retinal layers, which would be difficult to detect from the raw OCT images

alone. Deriving additional measures from the OCT images also adds to the complexity of the data analysis because the different aspects of the extended data must be linked.

All in all, the continuing use of OCT technology is accompanied by a growing complexity of analytical tasks and the associated expectations in medicine. Ophthalmologists need to gain insight into the very structures and properties of the data that are relevant for their diagnosis. Data exploration plays a central role, with the emphasis on the visual analysis of disease-related morphological patterns. The ophthalmologists have to see alterations in their spatial context to be able to understand the retinal condition, draw conclusions, and relate their insights to other parameters. In this respect, new techniques for evaluating retinal OCT data are needed that support interactive data exploration and the analysis of patient-specific data characteristics.

1.3 CHALLENGES

Although VA of medical data has proven to be effective in many applications [23], there are several challenges with regard to retinal OCT data. Many of them are concerned with the complexity of the data, particularly in terms of data size and structural composition. OCT data have an inherent spatial embedding. This means that the data consist of two types of components: characteristic components and referential components [24]. The characteristic components reflect specific measurements, e.g., axonal thickness, and the referential components reflect the spatial context, e.g., the specific location where the measured values were obtained. When handling such data, effective visual communication of these spatial properties must be ensured. Furthermore, usually only a part of the total data is relevant for assessing the retina. In order to support individualized diagnoses, the data must be reduced to the information specific to a patient or patient group. This already complicates the analysis of a single OCT dataset. Attempting to include several datasets simultaneously in the analysis further exacerbates the problem. Current analysis methods based on single image viewing or generalized data simplification offer only limited support in this respect. To come up with a new VA approach, we consider two major challenges:

VISUALIZATION OF 3D SPATIAL DATA: From the perspective of VA, retinal OCT data are volumetric data with a spatial embedding. This entails that the two types of spatial data components – the characteristic components and the referential components – have to be shown together. It must be made clear which characteristics belong to a given spatial reference and vice versa. This complicates the visualization design. The development of suitable visualizations therefore requires an understanding and careful consideration of the mutual effects in 2D and 3D encoding of each type of data component. Since the OCT data

are also volumetric, at least part of the data must usually be displayed in 3D. This ensures a faithful representation of the data. One of the biggest problems of 3D presentations is the handling of occlusion during data display. Care must be taken to avoid leaving any important information unseen. Thus, new visualization strategies are required that adequately present the spatial properties of retinal OCT data in 2D and 3D and ensure the visibility of all available information.

HANDLING OF DATA COMPLEXITY: OCT examinations produce complex retinal images with high spatial resolution. When several examinations are involved, a considerable amount of data accumulates. Attempting to simultaneously treat all data in all details can render the analysis ineffective. In addition, not only the image data but also derived measurements of the retinal layers have to be considered. To scale down the amount of data, strong data aggregations are currently being made prior to further processing and analysis. This generalized data reduction, however, can lead to a loss of patient-specific information. It must therefore be decided which parts of the data are to be included in which fidelity. To this end, the main characteristics of the data that are meaningful for a diagnosis must be identified. Once they are determined, the selection of spatial subsets needs to be facilitated in order to concentrate on the details relevant to the patients. Thus, new concepts are necessary that allow for a targeted reduction of retinal OCT data based on key data characteristics and their analysis at an appropriate level of abstraction.

Given these challenges, our goal is to develop a VA approach that on the one hand communicates the spatial properties of retinal OCT data and on the other hand allows a reduction of the data to relevant information. In doing so, we have to consider the requirements of ophthalmologists and the evaluation of our goals with applications in the ophthalmic field.

1.4 CONTRIBUTION

We address the challenges associated with the analysis of retinal OCT data using a systematic VA approach. Our basic idea is to tackle the challenges in three parts: (1) fundamental problems in VA of complex 3D spatial data, (2) new VA concepts for retinal OCT data, and (3) application of VA methods in ophthalmology. In the first part, we investigate the fundamental problems of ensuring the visibility and reducing the complexity of 3D spatial data. In the second part, we adapt the basic concepts to the specific problems associated with the analysis of retinal OCT data. In the third part, we transfer our solutions to the ophthalmic field to reflect on their practical use in cross-sectional studies. Given this overall structure, our contributions within the three parts are:

FUNDAMENTAL PROBLEMS IN VA OF COMPLEX 3D SPATIAL DATA: To create a conceptual basis, we investigate two fundamental problems related to VA of complex 3D spatial data: ensuring the visibility of all relevant information and reducing the data complexity to a level that is manageable for visual analysis. With respect to the first problem, we investigate the challenges of 3D visualizations in more detail. In particular, we deal with the implications of occlusion and present solutions for the computation of hidden information. We show that multiple coordinated views are effective for revealing different parts of 3D spatial data in overviews and details. Our new visualization design was published under the title “*Visibility Widgets for Unveiling Occluded Data in 3D Terrain Visualization*” in the *Journal of Visual Languages & Computing* [25]¹.

Having the concepts for the visualization of 3D spatial data available, we next deal with the second fundamental problem – the reduction of data complexity to a suitable level. For this purpose, we focus on the visualization of features that capture the most important information in the data. New strategies are developed by investigating the computation of features for complex spatial data with a use case in systems biology. The main characteristics of multiple moving entities are extracted and linked to the raw data to unveil spatio-temporal patterns. Our novel feature-based VA approach was published under the title “*Feature-Driven Visual Analytics of Chaotic Parameter-Dependent Movement*” in the journal *Computer Graphics Forum* [26].

NEW VA CONCEPTS FOR RETINAL OCT DATA: Based on the investigations of complex 3D spatial data, we develop new concepts for the analysis of retinal OCT data. To improve the visibility of diagnostically relevant information, we propose a novel visual design to present, highlight, and compare different aspects of the volumetric data. We present multiple interlinked 2D and 3D views that ensure effective communication of retinal layer measurements together with raw OCT images. Coordinated interaction facilitates exploration, selection of spatial subsets, and inspection of details on demand. The visual design is complemented by a unified data access to simplify the concurrent visual analysis of multiple datasets.

In order to cope with the data complexity, we introduce various types of features that provide patient- and group-specific information. A new visualization design allows to explore the features at different levels of abstraction, from precise feature maps to adaptive feature grids. By computing deviations of the maps or grids, our design enables a number of interactive comparisons between patients or between patients and normative data. To enhance the assessment of the comparisons, we combine the interactive visualization of features with statistical tests. We show that results of this combination allow a targeted reduction of

¹ References to self-authored publications are printed in bold.

the complex data and help to direct the analysis to parts that are likely to be relevant for a medical diagnosis.

All developed concepts are integrated into a comprehensive VA tool with a user interface that assists ophthalmologists in their analytical tasks. The tool design including the analysis of feature maps was published under the title “*Visual Analysis of Retinal Changes with Optical Coherence Tomography*” in the journal *The Visual Computer* [27]. The extended design and analysis of adaptive feature grids was published under the title “*Enhanced Grid-Based Visual Analysis of Retinal Layer Thickness with Optical Coherence Tomography*” in the journal *Information* [28].

APPLICATION OF VA METHODS IN OPHTHALMOLOGY: With the developed VA solutions at hand, we investigate their practical application in ophthalmology. In particular, we demonstrate how the developed concepts support concrete analysis tasks in ophthalmic research. Based on a participatory design with ophthalmologists, we propose a coordinated VA procedure for the evaluation of single-patient and multi-patient data. Our tool and procedure are applied in two cross-sectional studies focusing on early retinal changes in patients with diabetes mellitus type 1 and type 2 relative to healthy controls. The results obtained with our approach are compared with the results of the methods currently used in ophthalmology.

Beyond that, we further assess the results using pair analytics [29] and questionnaires together with medical and industry experts. With regard to the feedback collected, we argue that our problem-driven design has enabled more medical insights, reduced analysis effort, and expanded the software portfolio for analyzing retinal OCT data. To reach a wider audience in ophthalmology, the application of our VA solutions in the studies was published in two articles under the title “*Deviation maps for understanding thickness changes of inner retinal layers in children with type 1 diabetes mellitus*” in the biomedical journal *Current Eye Research* [30] and under the title “*Thickness of intraretinal layers in patients with type 2 diabetes mellitus depending on a concomitant diabetic neuropathy: results of a cross-sectional study using deviation maps for OCT data analysis*” in the journal of *Biomedicines* [31].

1.5 OUTLINE

This thesis consists of three parts. **Part i** contains this introduction (**Chapter 1**) and continues in **Chapter 2** with the investigation of fundamental problems related to VA of complex 3D spatial data. In **Chapter 3**, new VA concepts for the analysis of retinal OCT data are introduced. In **Chapter 4**, the corresponding analysis procedure is presented and results of two cross-sectional studies are reported. In **Chapter 5**, **Part i** is summarized and concluded.

Chapters 6 to 11 of Part ii collect the original publications on which this work is based. The chapters contain the details for the approaches, which are outlined in Chapters 2 to 5. Apart from minor corrections of typographical errors, the text in Part ii is identical to the original publications. The layout of the figures and tables, the typesetting, and the citations are harmonized to conform to the style of this thesis.

Part iii contains an overall summary and conclusion. Key insights about the greater picture behind this work are derived and discussed. The thesis is completed with research topics for future work on the VA of retinal OCT data in ophthalmology.

FUNDAMENTAL PROBLEMS IN VISUAL ANALYTICS OF COMPLEX 3D SPATIAL DATA

The VA of complex 3D spatial data, including retinal OCT data, requires dealing with two fundamental problems: (1) the visibility of relevant information and (2) the reduction of the data complexity. In this chapter, we present two new solutions to these problems. First, in Section 2.1, we investigate principal design options for the visualization of 3D spatial data [32]. Building on this, in Section 2.2, we develop a new approach for ensuring the visibility of occluded data in 3D visualizations. Finally, in Section 2.3, we address the data complexity and introduce a new approach in which we reduce the 3D spatial data to high-level features. Our original publications on the problems of data visibility and data complexity are included in Chapters 6 and 7 of Part ii, respectively.

2.1 DESIGN OPTIONS FOR THE VISUALIZATION OF 3D SPATIAL DATA

Designing visual representations that make 3D spatial data easy to understand is difficult. Visualization researchers are confronted with a number of principal options for encoding data. A key question is whether to use 2D or 3D presentations. For retinal OCT data, it is essential to address this question at the very beginning, as it will influence any further design considerations. To this end, we contribute a new systematization for 2D and 3D visualizations and discuss the advantages and disadvantages of the different design options [32].

2.1.1 Problem analysis

Before discussing the design options, we must first identify the basic properties of 3D spatial data. From a visualization perspective, the data consist of two different parts: the referential components and the characteristic components [24]. The referential components define the spatial context and the characteristic components (or attributes) represent results of measurements. The spatial context connects the attributes to different locations. Hence, we differentiate between the *reference space* \mathcal{R} (spatial coordinates) and the *attribute space* \mathcal{A} (measured values) [33]. For retinal OCT data, \mathcal{R} is three-dimensional and corresponds to the retina and its layers. \mathcal{A} can be multi-dimensional with one (univariate) or more (multivariate) attributes [34], e.g., reflectivity in cross-sectional images, computed axonal thickness, and so on. A major challenge in the visualization of such 3D spatial data is to convey the relationship between \mathcal{R} and \mathcal{A} .

Based on the data properties, we can now look at the fundamental differences in the visual communication of the relationship between \mathcal{R} and \mathcal{A} . In general, we consider presentations in which \mathcal{A} is directly embedded in the visual representation of \mathcal{R} . An important consideration here is whether to draw \mathcal{R} and \mathcal{A} in 2D or 3D. Both 2D and 3D presentations have different advantages and disadvantages. While 2D presentations are arguably more abstract and easier to read, 3D presentations allow to encode additional information. The third dimension also brings difficulties, such as occlusions and distortions. Hence, there is no universal answer as to whether 2D or 3D should be used. To make this decision for retinal OCT data, we characterize the design options in order to better understand their effects and implications [32].

2.1.2 *Related work*

The suitability of a presentation influences the speed at which information is absorbed, the number of errors made, the comprehension, and the visual working memory [35]. To understand the available design options, several categorizations of visualizations, such as photorealistic and non-photorealistic rendering [36] or static and animated images [37], were suggested. Related research also indicates that the choice of using 2D or 3D has a particularly great influence on the quality of a visualization. Previous studies showed that the decision depends on various factors, including data complexity and display technology [38–40] or the task and application context [41]. Interestingly, the combination of 2D and 3D often achieved good to superior performance and increased confidence in solving problems.

So far, however, only a general distinction was made between 2D and 3D presentations. That is, 2D presentations were assumed to consist only of two-dimensional graphical elements (e.g., points or lines) and 3D presentations only of three-dimensional graphical elements (e.g., solids or surfaces). Such a distinction limits the characterization of design options, especially considering the properties of 3D spatial data. Thus, we propose to differentiate between the dimensionality of the representation of \mathcal{A} and the dimensionality of the representation of \mathcal{R} .

2.1.3 *Characterization of 2D and 3D visualizations*

In order to investigate the design options for 3D spatial data, we distinguish between the presentation of \mathcal{R} and the presentation of \mathcal{A} . To this end, we developed a new systematization [32]. Figure 2.1 shows a matrix of the resulting categories. The horizontal axis compares 2D and 3D presentations of the spatial context and the vertical axis illustrates differences between 2D and 3D presentations of the embedded attributes. Characterizing these four combinations helps us to better understand the different options.

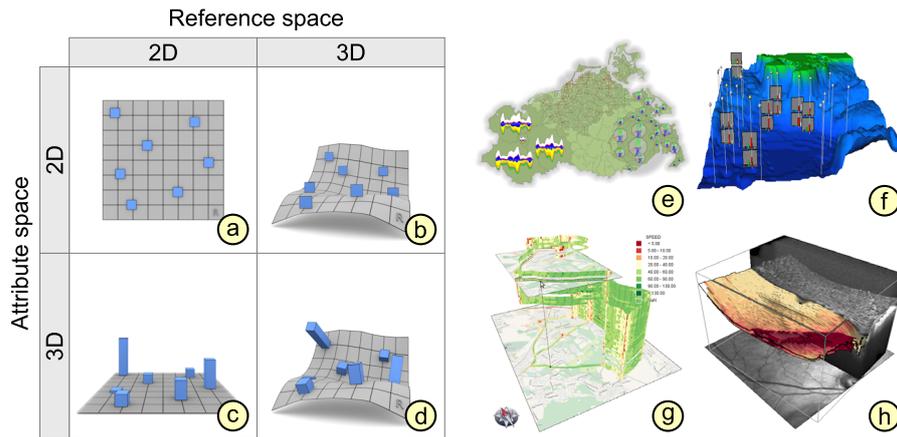


Figure 2.1: Systematic categorization of 2D and 3D presentations of the spatial context and embedded attribute (a to d). The examples are 2D diagrams on a 2D map [42] (e), 2D diagrams on billboards and 3D ocean floor [43] (f), 3D stacked trajectories over a 2D map [44] (g), and 3D retinal layers in a 3D OCT image [27] (h).

2D ATTRIBUTES IN 2D SPATIAL CONTEXT: Showing both \mathcal{A} and \mathcal{R} with 2D graphical elements is a well-established design option for 3D spatial data (Fig. 2.1 (a) and (e)). The attributes are displayed directly in the spatial context without the need for projections or visibility calculations [45]. This makes it straightforward to assign the attribute values to their respective locations. Several attributes are often combined into more complex diagrams [46, 47]. However, care must be taken when placing the graphical elements of attributes on a map to avoid visual clutter and overdrawing [48, 49]. Otherwise, sophisticated countermeasures, such as spatial generalization [50], need to be applied. Overall, the 2D visual abstraction provides an overview of the 3D spatial data and usually makes the resulting images easy to interpret.

2D ATTRIBUTES IN 3D SPATIAL CONTEXT: This category combines 2D and 3D graphical elements to represent \mathcal{A} and \mathcal{R} (Fig. 2.1 (b) and (f)). While the spatial context is depicted in its full extent, the attributes are usually integrated via diagrams on 2D billboards. The billboards are mainly used to display abstract attributes, e.g., measurements at selected points of interest. The use of 3D makes projections and visibility calculations necessary. Perspective foreshortening can cause billboards that are far away from the viewer to be difficult to perceive. Occlusion caused by the 3D representation of \mathcal{R} is another problem, which may require alternative views or interaction techniques. Altogether, this category helps to study attribute values in relation to certain spatial structures, such as mountains or valleys in 3D terrain.

3D ATTRIBUTES IN 2D SPATIAL CONTEXT: For this category, \mathcal{A} is encoded by 3D graphical elements and \mathcal{R} is abstracted as a map using

2D graphical elements (Fig. 2.1 (c) and (g)). The 3D elements are usually used to display more attribute values or to encode the altitude of the values to convey their 3D distribution. Another possibility is the depiction of time to support the analysis of spatio-temporal correlations [51, 52]. However, occlusion and perspective distortion can make it difficult to decipher such 3D and 2D combinations. In addition, the higher the attribute values are displayed above the map, the more difficult is their spatial assignment. Hence, advanced projection methods [53] or additional cues may need to be included. In summary, the third dimension in this category offers greater flexibility in the presentation of attributes.

3D ATTRIBUTES IN 3D SPATIAL CONTEXT: This category shows both \mathcal{A} and \mathcal{R} with 3D graphical elements (Fig. 2.1 (d) and (h)). The distribution of attribute values and the structure of the spatial context are faithfully communicated. As a general rule, the number of perceptible elements is higher than in 2D presentations [38]. Yet, the 3D graphical elements of \mathcal{A} and \mathcal{R} also lead to rather large occlusions, which can obstruct the communication of data parts [54]. This main problem must be taken into account. Moreover, adequate data resolution and appropriate lighting and shading are required to avoid misleading visual representations of attribute values [55, 56]. Overall, this category support realistic representations of 3D spatial data and emphasizes three-dimensional spatial relationships.

2.1.4 Discussion

In general, all combinations of 2D and 3D are suited for presenting 3D spatial data. Which combination of graphical elements to choose for \mathcal{A} and \mathcal{R} depends on the specific data and application scenario. Our characterization is a first step to support this decision by highlighting the effects and implications of the four categories. This promotes problem awareness and helps to develop an appropriate visualization. On the basis of our characterization, we were able to design prioritized visualizations of 3D terrain, geo-spatial attributes, and uncertainties [57].

As for retinal OCT data, we conclude from our discussion that a combination of 2D and 3D presentations is advantageous. Since we are dealing with 3D spatial data, 3D presentations (3D attributes and 3D spatial context) enable us to faithfully convey the spatial structure of the OCT volume and the distribution of disease-related defects therein. By additionally abstracting the 3D spatial data in 2D presentations (2D attributes and 2D spatial context), we can also provide an overview of retinal layers and support the assignment of measured attribute values to their exact spatial location. However, apart from the combined advantages, the use of 3D and 2D also means that we have to deal with the associated problems. Our first fundamental challenge is to ensure the visibility of all relevant information. In the case of 3D graphical

elements involved in the visualization, this means that it is definitely necessary to deal with the problem of occlusion and prevent any part of the 3D spatial data from going unseen. The complexity of the data further complicates the occlusion problem, as not all details within the 3D spatial context can be displayed in a single image. The problem of data complexity must also be taken into account in 2D presentations. With 2D graphical elements involved in the visualization, it is essential to focus on the main characteristics of the 3D spatial data to avoid visual clutter and allow accurate decoding of attribute values. In line with our second fundamental challenge, this means that we have to reduce the data complexity to a level that is manageable for visual analysis.

The presented characterization of visualization options for 3D spatial data helped us to consider the identified problems from the beginning. Building on this, in this thesis, we contribute two novel approaches to solve these problems for an effective VA of retinal OCT data.

2.2 VISIBILITY WIDGETS FOR UNVEILING OCCLUDED DATA IN 3D VISUALIZATION

Occlusion is a fundamental problem in the 3D visualization of retinal OCT data. The reason for this is that retinal microstructures or measurements of local defects are easily obscured within the volumetric images. This can lead to considerable loss of relevant information. To ensure the visibility of this information, we contribute a new VA approach that supports the detection, presentation, and interaction with occluded data. The original publication on our approach received a best paper award [58]. Our subsequently invited and extended journal article [25] is included in Chapter 6 of Part ii.

2.2.1 Problem analysis

Whether parts of retinal OCT data become occluded in a visualization depends on various factors. We distinguish three principal types of occlusion: (1) occlusion by viewpoint, (2) occlusion by 3D spatial context, and (3) occlusion by attributes. Figure 2.2 illustrates this differentiation.

OCCLUSION BY VIEWPOINT: Independent of the other two types of occlusion, it is often not possible to present the volumetric OCT data in their entirety from a selected viewpoint. Hence, in 2D and 3D visualizations, parts of the retinal details in the 3D spatial context end up outside the field of view (Fig. 2.2 (a)). We consider these data parts to be occluded because they are not visible with common rendering methods.

OCCLUSION BY 3D SPATIAL CONTEXT: In 3D visualizations of retinal OCT data, the graphical elements of the attributes can be hidden by the 3D graphical elements of the spatial context (Fig. 2.2 (d)).

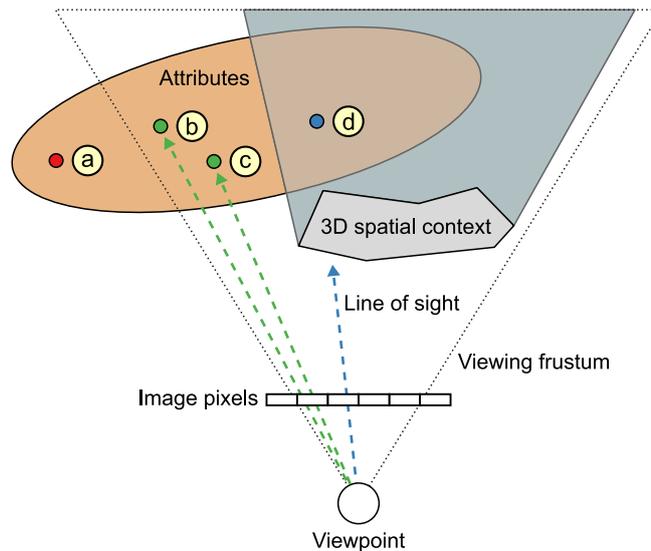


Figure 2.2: Three types of occlusion of retinal OCT data (adapted from [54]). Data point (a) is outside and data points (b), (c), and (d) are located inside the viewing frustum. Points (b) and (c) are obscured because they contribute to the same image pixel, while point (d) is hidden by the 3D spatial context.

If the attributes are mapped directly to the spatial context, the values may additionally end up on parts of the 3D scene that are facing away from the viewer. This is the case, for example, when measurement results are textured on the uneven surface of the retina but are not visible in depressions or behind elevations.

OCCUSION BY ATTRIBUTES: Occlusion also occurs when multiple retinal attribute values compete for encoding the same image pixel (Fig. 2.2 (b) and (c)). This results from values of different attributes that have exactly the same spatial location. For instance, measurements of both volume and thickness in the center of the retina. In 3D visualizations, these values can even have slightly different positions, but still be mapped to the same pixel due to the perspective projection.

In conjunction with our characterization of design options (Sect. 2.1), the three types of occlusion can be further summarized with respect to the dimensionality of the presentation of the 3D spatial data. Specifically, occlusion by viewpoint occurs regardless of whether the attributes and the spatial context are displayed in 2D or 3D. Occlusion by 3D spatial context, on the other hand, only occurs if the spatial context is actually presented in 3D. Occlusion by attributes, in turn, can occur in both 2D and 3D presentations of the attributes and the spatial context. In 3D presentations of attributes, occlusion problems can be reduced compared to 2D presentations by distributing the values more freely in the 3D scene. If, however, the spatial context is also presented in 3D, the

problems increase because with greater distance from the viewpoint, more spatial locations of attribute values are mapped to the same pixel due to perspective distortion.

2.2.2 *Related work*

Occlusion is one of the biggest problems in 3D visualization of 3D spatial data. In a comprehensive survey and taxonomy [54], five design patterns for dealing with occlusion were identified: tour planners, volumetric probes, multiple viewports, virtual X-ray tools, and projection distorters. The first two patterns are hardly applicable for our purpose, as they require intensive preprocessing or strongly affect invariances of the 3D spatial data. The last three patterns are more suitable for our work. For each of them different occlusion-aware visualization techniques exist (e.g., [59–64]). However, with regard to the visualization of 3D spatial data and especially retinal OCT data, we found that the existing approaches use the patterns mostly in isolation and address only one of the three types of occlusion at a time (cf. Sect. 6.3). A comprehensive solution for our application context has not yet been introduced. Our contribution therefore is an integrated approach that combines the strengths of the three design patterns and thus leverages their full potential to deal with all three types of occlusion.

2.2.3 *Detecting, presenting, and interacting with occluded data*

We aim to ensure that relevant information is actually visible during visual analysis. To this end, our novel VA approach [25] helps to deal with occlusion in the 3D visualization of 3D spatial data, including retinal OCT data. Our comprehensive solution not only presents hidden parts of data, but also supports their automated detection and interactive access. In this respect, we consider three principal steps. In the first step, the detection automatically determines data parts that are hidden and therefore could have gone unnoticed. In the second step, the presentation draws the user’s attention to the fact that these data parts are currently not visible and shows where they are located. In the third step, the interaction enables the user to access the hidden parts of the data and integrate them into the analysis together with the already visible parts. By supporting all three steps, we create an awareness that information can be obscured in an image, while also enabling the user to take action and make this information visible in a targeted manner.

DETECTION OF OCCLUDED DATA: The first step of our approach is to determine the part of the 3D spatial data that is hidden in a visualized image from a given viewpoint. In an interactive 3D visualization of retinal OCT data, this means that in each rendered image, all three types of occlusion must be accurately detected. According to our character-

ization of 3D visualizations (Sect. 2.1), this further entails analyzing the proximity between 3D graphical elements for various representations of attributes and spatial context. For this purpose, we developed a computational occlusion detection based on the general model of 3D occlusion [54]. Our pipeline-based computation is performed during rendering, without any preprocessing and with most calculations done in image space. This has three main advantages. First, the detection during rendering offers a high degree of flexibility and is independent of the design patterns that are subsequently used to unveil the detected occlusion. Second, the calculations in image space are largely independent of the actual visual representation of the 3D spatial data, i.e., the graphical elements used to display the attributes and the spatial context. Third, the pipeline-based procedure is well-suited for implementation on graphics hardware to ensure interactive computation rates. This enables us to capture all three types of occlusion. The procedure can be summarized in four steps.

1. *Data culling*: To enable an interactive occlusion calculation, it is necessary to reduce the 3D spatial data to the actually necessary amount. Hence, a reduced set of data points is generated by selecting only points of attributes that are of current interest. All these data points are initially marked as visible.
2. *Detection of occlusion by viewpoint*: Next, the part of the reduced data outside the 3D viewing frustum must be computed. To simplify this calculation, the selected data points are projected into the image space. Those that are outside the image boundaries are not visible and are marked as occluded.
3. *Detection of occlusion by 3D spatial context*: For data points within the image boundaries, those partially or completely obscured by graphical element of the spatial context must be determined. To ease this step, the graphical elements are abstracted to pixels. The data points are then marked as occluded if they are behind pixels of the spatial context according to the depth buffer [65, 66].
4. *Detection of occlusion by attributes*: For the remaining data points, all those covered by graphical elements of attributes need to be identified. This is again done by abstracting the graphical elements to pixels. In case a pixel in the image gets filled more than once, the associated data points are marked as occluded.

PRESENTATION OF OCCLUDED DATA: After the occlusion detection, the next step is to make the user aware of the fact that the visualization is affected by occlusion. On the one hand, this means that retinal OCT data must be presented in a suitable way and on the other hand, the visibility of the data must be communicated simultaneously.

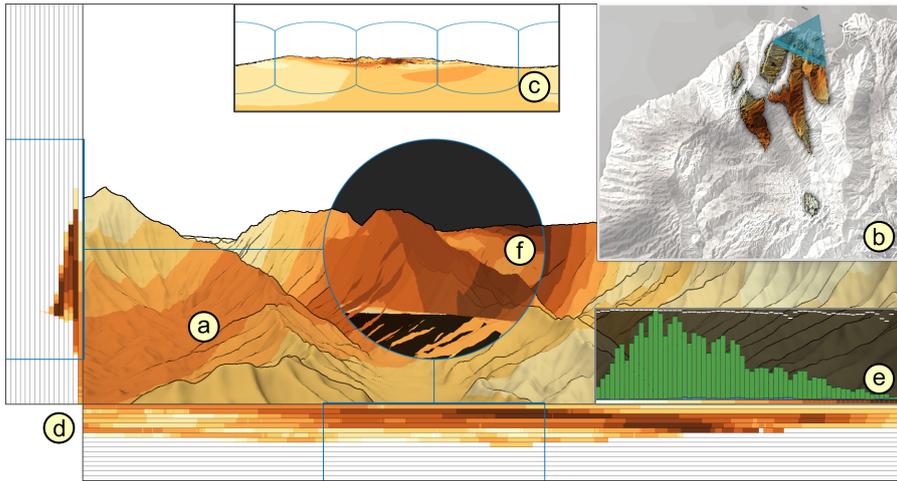


Figure 2.3: Visibility widgets help users identify occluded information in 3D visualizations [25]. In this example, the main 3D view (a) and a secondary 2D map (b) show color-coded weather attributes in 3D terrain. Three widgets highlight hidden information: panorama (c), bands (d), and histogram (e). An interactive lens (f) visualizes data hidden on the backside of the spatial context.

Presenting both visible and occluded parts of the data together is essential to ensure that no relevant information goes unseen. Thus, we introduced a new visualization design based on multiple linked 2D and 3D views. To explain the ideas behind our design, we would first like to illustrate it with a concrete application example, the visualization of weather data in 3D terrain. In this example, the distinction between attributes (weather data) and spatial 3D context (terrain) is easy to understand, which facilitates the discussion of our design decisions regarding the three types of occlusion. Figure 2.3 shows an overview of our design with the application example. We distinguish between two basic views for representing the 3D spatial data and three occlusion-specific views to unveil hidden information. All views combine different design patterns to treat all three types of occlusion.

The basic views are a main 3D view and a secondary 2D map view. In line with our characterization of the design options (Sect. 2.1), this combination enables both a faithful 3D representation of spatial details and a 2D overview of the entire spatial context. To also handle occlusion, the design pattern of multiple viewports is used to simultaneously provide the two alternative perspectives (3D detail and 2D overview) on the 3D spatial data.

3D DETAIL: For a faithful representation of 3D spatial data, it is necessary to show them in detail. The main 3D view delivers the details about the spatial distribution of the attributes and the structure of the 3D spatial context (Fig. 2.3 (a)). The general spatial structure is illustrated using shading. Depressions and elevations are further emphasized using enhanced edges and spatial perception is

reinforced via stylization, e.g., aerial perspective. Using color, the attributes are displayed directly in the grayscale representation of the context. This provides visual contrast and conveys the 3D spatial relationships.

2D OVERVIEW: In addition to 3D details, it is also important to grasp all the 3D spatial data at a glance. On demand, the 2D map provides an overview of the attributes and the spatial context from a vantage point (Fig. 2.3 (b)). Similar to the main view, the attributes are color-coded and the spatial structure is shaded gray. Hidden parts of the main view are further highlighted in the map together with the location of the main viewpoint. In this way, the visual abstraction in the 2D map serves spatial orientation and already helps to a certain extent to treat occlusion by viewpoint and occlusion by 3D spatial context.

Nevertheless, the occlusion in the main 3D view must be treated further. When displaying occluded data, it is hardly possible to encode visible and hidden parts together directly in the main view due to overlaps. For this reason, the occlusion-specific views are designed as compact visibility widgets that are activated as needed. The widgets keep the main view largely unchanged and the 3D representation of the 3D spatial data retained. They consist of a panorama, two bands, and a histogram, which help uncover and locate partially or fully hidden data parts both inside and outside the main view. To address all three types of occlusion, the widgets show different aspects of the occluded data with variations of the projection distorter pattern.

PRESENTATION OF OCCLUSION BY VIEWPOINT: To make data outside the main view visible, more of the 3D scene must be integrated into an image. The panorama widget provides a 360° overview, allowing to look at the main 3D scene with an extremely wide angle (Fig. 2.3 (c)). This way, parts of data outside the main view can be easily identified. It can also be set to show only visible or only occluded parts to ensure the visibility of all information in the the immediate vicinity of the viewpoint.

PRESENTATION OF OCCLUSION BY 3D SPATIAL CONTEXT: To create an awareness of data behind the spatial context, visual cues are needed about their position in the 3D scene. The bands widget shows where the majority of obscured attribute values is found along the horizontal and vertical axes of the main view (Fig. 2.3 (d)). To indicate the distance to these values, the bands are divided into layers representing different depths from the viewpoint. This allows to approximate the 3D position of the hidden information.

PRESENTATION OF OCCLUSION BY ATTRIBUTES: Maintaining the spatial position of data covered in pixels with multiple assigned

values is hardly possible. The histogram widget instead abstracts from space and summarizes which attribute values in the main view are more or less affected by occlusion (Fig. 2.3 (e)). To convey the relationship between the hidden and visible values, its bins are divided by the percentage of occlusion in each value range. That way, it can be determined whether all values of interest are actually visible in the image.

INTERACTION WITH OCCLUDED DATA: After raising awareness through the presentation of occluded data, the final step is to enable the user to become active and interact with the hidden information. The aim of the interaction here is to make the occluded parts selectively visible in the main view so that they can be visually analyzed in the full 3D representation. However, using the main view directly for interaction is not an option, as it is barely possible to interact with non-visible parts of the 3D presentation. For this reason, we take the other views in which the parts are visible and link the interaction performed with them to the main view. By adding interaction to the presentation of occluded data, we can also integrate the third design pattern of virtual X-ray tools into our approach. This pattern helps to locally relate the occluded parts to the visible parts of the 3D spatial data. With the advantages of all three patterns combined, we are able to support access to information hidden due to all three types of occlusion.

INTERACTION WITH OCCLUSION BY VIEWPOINT: We need to allow access to information outside the view boundaries. The most straightforward way to do this is to change the viewpoint. Yet, it is difficult to change it manually in such a way that information is made visible in a targeted manner while maintaining orientation in the 3D scene. On this account, we provide automated navigation support. To auto-rotate the main view to information outside the field of view, points of interest are selected in the panorama. To auto-move the main view to places with hidden information, corresponding destinations are selected in the 2D map overview. A visibility history eases the search for destinations in the map by highlighting previously visited and unexplored places with hidden information.

INTERACTION WITH OCCLUSION BY 3D SPATIAL CONTEXT: We have to provide access to information that is within the view boundaries but behind the parts of the spatial context. For this purpose, we utilize the X-ray design pattern in the form of an interactive lens in the main view (Fig. 2.3 (f)). To reveal hidden information in place, the lens allows to see locally through the geometry of the 3D spatial context. To improve the readability of the revealed information, the lens additionally abstracts from spatial details and reduces the influence of shading on colored attribute values.

Using the lens together with the bands widget further helps to steer the interactive search towards the majority of information hidden in the 3D scene.

INTERACTION WITH OCCLUSION BY ATTRIBUTES: We must also allow access to information covered by the representation of other attributes. This requires, however, a way to specify either where or which hidden attribute values are of interest. Therefore, we support spatial and value-based selection methods and automatically reveal hidden values accordingly in the main 3D view. To inspect occluded values at specific points or within regions, respective parts of the spatial context are selected in the 2D map. To examine occluded value ranges, the value distributions visualized in the bands and histogram widgets are brushed. In this way, the abstracted information of the widgets is put in relation to the space again in the main view.

2.2.4 Discussion

Our novel VA approach not only integrates three different design patterns for the presentation of occluded 3D spatial data, but also supports automated detection and interaction with all three types of occlusion. In [Section 6.5](#), we show that this is possible with interactive refresh rates by implementing our design concepts on consumer graphics hardware using common rendering methods. With the resulting flexibility, we are able to control the visibility of 3D spatial data and thereby support users in exploring 3D visualizations. We conclude that our solution can be a useful extension of existing approaches to occlusion management.

We successfully tested our approach in a demonstrating example in the field of avionics (cf. [Fig. 2.3](#)). Here, ensuring the visibility of weather data in 3D terrain was important. Especially, as the textured weather attributes on the terrain surface meant that not all relevant information was always visible. Positive feedback from domain experts indicates that our design was indeed helpful in mitigating the adverse effects of occlusion when visualizing such 3D spatial data ([Sect. 6.6](#)). Beyond that, we applied the solutions to volumetric OCT data of the human retina ([Sect. 6.7](#)). With the help of the widgets, hidden attribute values within retinal layers could be uncovered.

Overall, the core ideas of our approach are an important step towards meeting our first fundamental challenge: ensuring the visibility of relevant information in the 3D visualization of retinal OCT data. The biggest advantage is that several aspects of the 3D spatial data are displayed simultaneously. With the help of our characterization ([Sect. 2.1](#)), we devised a combination of 3D and 2D views that supports both a faithful representation in the necessary detail (3D attributes in 3D spatial context) and an abstract overview of the distribution of the 3D spa-

tial data (2D attributes in 2D spatial context). The additional visibility widgets made it possible to detect, present, and interact with hidden information in these linked 3D and 2D views.

Having this solution to the occlusion problem available, we now consider our second fundamental challenge. In the next section, we show how we solve the complexity problem by focusing on the main characteristics of 3D spatial data.

2.3 FEATURE VISUALIZATION FOR REDUCING DATA COMPLEXITY

The complexity of retinal OCT data complicates the visual data analysis. The reason for this is that relevant information on small defects of retinal layers is mixed up within many irrelevant parts of the data. Not only must these spatial details be displayed in their 3D spatial context, but they also need to be visually related across datasets of numerous patients in an understandable way. Without handling the data complexity, excessive clutter and overdrawing make it difficult to show even basic properties, not to mention insight into higher-level patterns. A common approach to address this problem is the extraction and visualization of features. We adapt this idea and contribute novel feature-based VA concepts that help to concentrate on relevant information within complex 3D spatial data. The original article on our work received a best paper award [26] and is included in Chapter 7 of Part ii.

2.3.1 Problem analysis

The feature-based visual analysis of retinal OCT data is itself a challenge. First and foremost, there is no universal definition of a feature that allows the complex 3D spatial data to be reduced to relevant information. Rather, the criteria for specifying a feature depend heavily on the data to be analyzed. Thus, different data require different features. For certain types of data, tailored feature definitions exist. For example, in the visualization of flow data, various features are used to describe critical points, vortices, or shock waves [67]. Another example is the visualization of movement data, in which different features are applied to characterize individual or multiple moving entities [68, 69]. These examples show that a single feature alone is usually not enough for a thorough visual analysis. Instead, a whole range of feature definitions is required to cover the richness of complex 3D spatial data. Moreover, once suitable features are defined, they must also be extracted and then presented to the user interactively. With respect to retinal OCT data, for all these steps – specification, extraction, presentation, and interaction with matching features – no coherent solution is available.

Therefore, in order to come up with ideas to solve these steps, we first developed a new feature-based VA approach in the context of another difficult but known problem: the analysis of parameter-dependent

movements. In this example from systems biology, we were faced with thousands of simulations of chaotic movements. To make sense of the complex 3D spatial data, the spatio-temporal characteristics of individual entities and dynamically forming groups had to be extracted from the highly entangled movement data. On top of that, the dependence of the movement characteristics on parameters that control the simulation needed to be understood. Based on the solutions to this problem presented below, we adapt the key ideas in [Section 3.2](#) to arrive at a feature-based VA approach dedicated to retinal OCT data.

2.3.2 *Related work*

The research on feature-based visual analysis has its roots in flow visualization [67] and has since expanded to various application contexts [70]. The general goal is the automated extraction of features that capture meaningful data characteristics [71]. For retinal OCT data, different measurements, e.g., the thickness or volume [72, 73], have been used to characterize the condition of the retina and its layers. Yet, an actual feature-based solution that reduces the complexity of retinal OCT data based on such characteristics does not exist.

Switching to our demonstrating example from systems biology, an overview of feature-based techniques applied to movement data can be found in [68] and a conceptual framework for visual parameter space analysis in [74]. So far, however, we have only seen several individual solutions, but none that directly meet the needs of analyzing parameter-controlled simulations of chaotic movements with dynamic groups. Although feature-based approaches were previously used to analyze group movements [69, 75], most of them exhibited fixed membership and non-chaotic behavior. Features of groups with dynamic membership have been rarely addressed [76] and then again only for rather few simple movements. Approaches are also missing to analyze spatio-temporal features of all movements belonging to a certain parameter configuration and to explore them with regard to all alternative configurations. Our contribution therefore is a new feature-based VA approach, which covers not only chaotic movements and dynamic groups but also the analysis of parameter dependencies.

2.3.3 *Feature-based visual analysis*

We aim to reduce the complexity of 3D spatial data. To this end, our novel feature-based VA approach [26] helps to focus on the main data characteristics rather than on the individual data elements. Our design includes both the systematic definition and automated extraction of matching features on the one hand and their presentation and interactive investigation on the other hand. Following the basic procedure of feature-based visual analysis [71], we divided our approach into

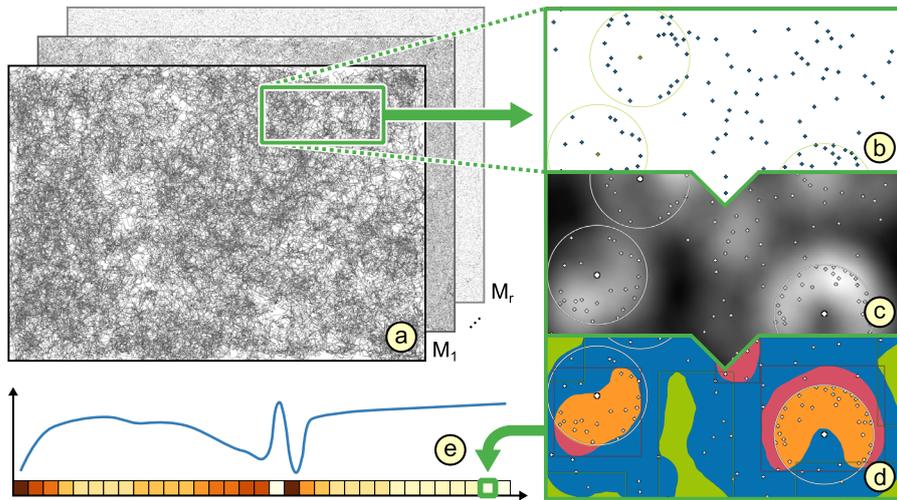


Figure 2.4: Drawing thousands of entangled trajectories of simulated movements leads to excessive clutter, making relevant information indecipherable (a). Feature extraction reduces the chaotic 2D movement to meaningful characteristics by converting the entity positions (b) per time step to a density map (c) and then regions (d), which are finally assembled into 1D time series of feature values (e).

four principal steps. In the first step, we look at how to formally specify which characteristics are relevant for the interpretation of the data. In the second step, we deal with how these characteristics can be automatically extracted from the raw data into features. In the third step, we consider how only the extracted features can be visualized to convey the main characteristics of the data. In the fourth step, we investigate how to adjust interactively, which features are to be visualized, and how these can be related to the raw data again. By covering all four steps, we achieve a higher level of abstraction of the complex 3D spatial data and make the visual representation much clearer and more concise.

COMPOSITION OF THE DATA: Before we describe the key ideas behind each step of our approach, we must first examine the basic composition of the given data. This is necessary because all steps of a feature-based visual analysis must correspond to this composition in order to actually reduce the data complexity to relevant information.

In our demonstrating example, the data consist of r simulation runs. Each run $R_i = (P_i, M_i) : 1 \leq i \leq r$ includes a parameter configuration P_i (simulation input) and a corresponding movement M_i (simulation output). Each movement M_i consists of the trajectories T_1, \dots, T_m of m entities moving in a 2D spatial context. The trajectories are sampled uniformly at time steps t_1, \dots, t_n . Each trajectory point stores various attributes, e.g., position or speed. Points in close proximity to each other define temporary groups of entities.

While the movement of a single simulation run is 2D, the parameters that control the simulation add another dimension to the data.

Thus, all simulated movements are interpreted together as 3D spatial data. The complexity of the data results from thousands of parameter configurations, each leading to chaotic movements with thousands of trajectories. As shown in [Figure 2.4](#) (a), the direct drawing of all data leads to images with extreme clutter, making it impossible to study individual movements or the influence of parameters.

FEATURE SPECIFICATION AND EXTRACTION: After knowing the composition of the complex 3D spatial data, the first step of our approach is to specify what characteristics are relevant. In our context, this means that we need to work out feature definitions that describe meaningful aspects of the chaotic movement behavior of individual entities and dynamic groups. Here it is not enough to look only at the attributes of the trajectories themselves. Rather, we have to also take into account where (space) and when (time) the movements take place. A single feature is obviously not sufficient to gain a deeper insight into all these aspects of the data. Hence, we approached the definition of matching features by categorizing existing definitions in literature (e.g., [69, 76]) and systematically extending these definitions with custom features for our demonstrating example (see supplementary material of [26]). In total, we distinguished four categories of features:

- *Basic features* describe the general behavior of moving entities via aggregated attribute values for all trajectories in a movement. Examples are the average speed or average acceleration of entities.
- *Group features* characterize the behavior of dynamic groups by summed attribute values of the trajectories of their current members. Examples include the number of groups per time step or the ratio of group members and entities not belonging to a group.
- *Region features* consider the spatial distribution of the moving entities beyond their individual trajectories. Interesting examples include the number and average size of regions with low or high density of entities of certain types.
- *Advanced features* capture the common development or specific variation in previously extracted feature values. Such features of features can, for example, detect points in time with significant changes in movement behavior.

Based on the specification, the next step is to automatically extract the features from the complex 3D spatial data. For a comprehensive visual analysis, several features from all four categories are required. In order to match the different categories, we have to provide different extraction mechanisms. Basic and group features are extracted by aggregating individual trajectories or sets of them. Regional features are obtained by computing and quantizing 2D density maps for each time step of

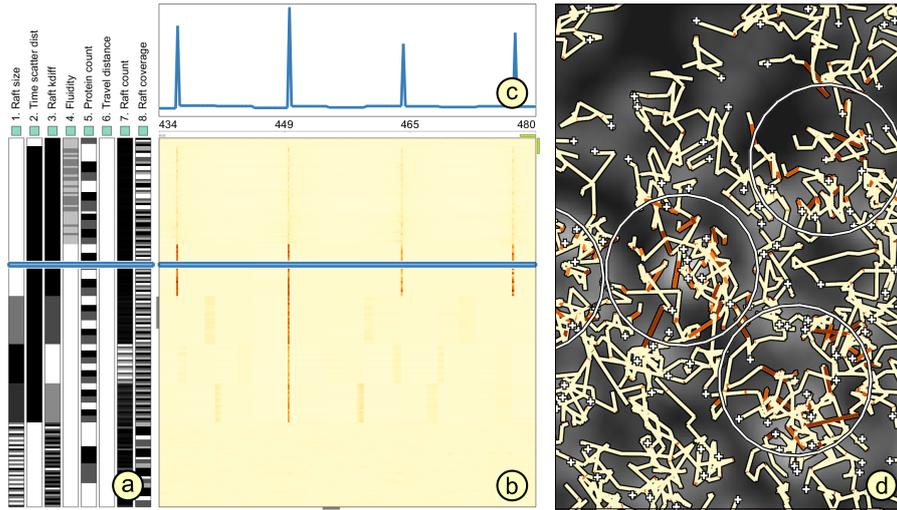


Figure 2.5: Visualization of features and detail information for selected parts of the data. Shown are color-coded parameter configurations (a) and feature values over time (b), a chart with selected time series (c), and a trajectory view with selected trajectory segments (d).

a movement (Fig. 2.4 (b) to (d)). Advanced features are derived by temporal clustering of previously extracted features.

Given the variety of the features and the complexity of the 3D spatial data, it is hardly possible to perform the different extraction mechanisms on the fly. We therefore integrated an extensive pre-processing that computes as many features as possible in advance. This considerably reduces the complex 3D spatial data. In our example, each chaotic 2D movement with thousands of entities is abstracted to a 1D time series of feature values (Fig. 2.4 (e)). The main advantage is that once pre-computed, the features are readily available for visual analysis. This allows users to include or exclude certain features as needed and thus determine which of them are most helpful in characterizing the data. Moreover, since the number of extracted feature values is much smaller than the number of raw data elements, the time series for all simulation runs are now easier to visualize than the respective movements.

FEATURE VISUALIZATION AND INTERACTION: Having the extracted features available, the last two steps of our approach are to visualize the features and make them interactively accessible. The visualization must present not only the temporal development of the features, but also the spatial context of the movements and the dependencies on the parameter configurations. According to our characterization of design options (Sect. 2.1), 2D visualizations can make it easier to interpret all these complex aspects of the data. However, not all the information can be shown in a single 2D image. We therefore introduced a new visualization design with linked overview and detail views. Figure 2.5 illustrates the view arrangement in the user interface. The multiple

views ensure that all relevant aspects of the features, parameters, and spatial context are actually visible (cf. Sect. 2.2). Moreover, since we reduced the complexity of the 3D spatial data, the 2D visualization of the extracted features does not suffer from the visual clutter seen in the raw movement images (cf. Fig. 2.4).

OVERVIEW: To get an overall idea of the simulated movement behavior, the extracted features and simulation parameters must be displayed together. Therefore, the overview shows them as stacked rows for all simulation runs in a pixel-oriented visualization. For each simulation run, the parameter configuration is depicted on the left (Fig. 2.5 (a)) and the associated time series of a selected feature is displayed in the center (Fig. 2.5 (b)). The rows are color-coded, with lighter colors indicating lower values and darker colors higher values. Horizontally, the temporal feature development is reflected in changing colors. Vertically, the dependencies between parameters and features are mirrored in patterns across the colored rows. While this design only shows one feature at a time, an interactive switch between different features is easily possible thanks to the extraction pre-process.

DETAIL: Temporal and spatial dependencies must also be examined in detail. For this purpose, two detail views present additional information for selected parts of the data. The temporal aspects are communicated in a chart view that shows time series of one or more selected simulation runs (Fig. 2.5 (c)). This enables precise comparisons of movement characteristics of different parameter configurations. The spatial aspects are detailed in a trajectory view that shows the movement trajectories of a selected simulation run (Fig. 2.5 (d)). The trajectory segments are colored according to the time series of feature values to illustrate where the movement has certain characteristics. Optionally, the 2D density map of a selected time step is overlaid to further highlight the connection between the trajectories and specific regions in the spatial context.

To complement the feature visualization, the interaction must help the user to investigate interesting patterns in both the parameter dependencies and spatio-temporal characteristics. Given the size of the 3D spatial data, such patterns can occur globally over all simulation runs and time steps, but also locally only for subsets of simulation runs and limited time intervals. Hence, we integrated coordinated interaction techniques to search for and investigate global and local patterns in the linked views. The interaction is also essential to deal with the data complexity, which, although reduced, can make it difficult to concentrate on the relationships between parameters, features, time, and space.

INVESTIGATION OF GLOBAL PATTERNS: To make patterns visible across all simulation runs, the respective parameter configurations and

extracted characteristics must be related to each other. The ability to visually identify these relationships depends heavily on how the stacked rows are arranged in the overview visualization. The simulation runs themselves have no fixed order. Hence, we support their automated arrangement by sorting them on the basis of either features or parameters. The first option allows an undirected search for similarities or differences in the time-series of feature values and then relates them to noticeable patterns in the visualized parameter configurations. Conversely, the second option facilitates a directed search by specifying a certain sequence of parameter configurations and then checking the extracted features for corresponding patterns. To explore different arrangements using both sorting options, various sorting metrics and algorithms can be applied interactively.

INVESTIGATION OF LOCAL PATTERNS: Certain spatio-temporal characteristics are limited to a few simulation runs and specific time intervals. In order to focus on such local patterns, the overview allows to zoom and pan along the axis of simulation runs and along the axis of time points. During this interaction the visibility of the data has to be ensured (cf. [Sect. 2.2](#)). Hence, two indicators show whether one of the visualization axes is affected by overdrawing of the attributes. If this is the case, the zoom level can be automatically adjusted so that the patterns are fully visible. In addition, miniature scroll bars help with orientation by showing where the current view is in relation to the entire data. The zooming and panning along the time axis is also linked to the detail views. By only highlighting those trajectory segments that are within the displayed time interval in the overview, the visual clutter in the trajectory view is greatly reduced. Only in this way is it possible to locally relate the high-level temporal features down to individual steps of the low-level chaotic movements.

2.3.4 *Discussion*

Our novel feature-based [VA](#) approach for complex 3D spatial data supports not only the systematic specification of relevant data aspects, but also the automated extraction of matching features. This allows us to concentrate on the main data characteristics, while leaving irrelevant data parts out of consideration. We demonstrated that this is even possible for parameter-dependent chaotic movements with dynamic groups, for which no coherent solution existed before.

Researchers from systems biology successfully applied our solution to study stochastic simulations of chaotic movements on cell surfaces [77]. They performed about 2000 simulation runs, controlled by 8 parameters with an average of 5000 moving entities (max. 25000) over 4000

time steps. Using our approach, they were able to reduce the 180 GB of chaotic movement data to about 60 features (each ≈ 80 MB). While exploring the features, the experts made several new observations regarding parameter dependencies and the movement behavior of temporary groups (cf. Sect. 7.5). Their positive feedback suggests that our feature-based approach is indeed helpful in analyzing movement data that are not only large but also complex. Besides movement data, we adapted the interactive visualization part of our solution for the analysis of parameter dependencies of segmented time-series data [78].

In summary, the concepts developed as part of our approach bring us closer to solving our second fundamental challenge: reducing the complexity of retinal OCT data. Our discussion of design options (Sect. 2.1) again helped to create 2D visualizations that convey meaningful data characteristics in their spatial context. Although we demonstrated our approach with chaotic movements from systems biology, the design concepts are generic enough to be applicable to retinal OCT data (cf. Sect. 3.2).

2.4 SUMMARY

In this chapter, we presented new solutions for two fundamental challenges in the VA of complex 3D spatial data: (1) the visibility of relevant information and (2) the reduction of the data complexity. We started by characterizing principal design options for the visualization of 3D spatial data. Building on this, we introduced visibility widgets that allow attributes to be effectively communicated directly in 2D and 3D presentations of the spatial context, even when parts are hidden due to different types of occlusion. We then presented a feature-based visual analysis approach that enables the main characteristics of complex 3D spatial data to be explored without visual clutter affecting the visualization. The reduction to features, i.e., to 1D time series, even made it possible to display the complex attributes exclusively in 2D. For both approaches, we established a close integration of computation, visualization, and interaction. In this way, our comprehensive designs extend the range of solution for the occlusion problem and the complexity problem in the VA of complex 3D spatial data.

So far, we illustrated our approaches to both problems with exemplary data from different application areas. For the occlusion problem, we chose weather data in 3D terrain and for the complexity problem, simulations of chaotic parameter-dependent movements. Our next step is to transfer the key ideas developed here to the VA of retinal OCT data. In this connection, we will combine feature-based visual analysis with multiple linked 2D and 3D views to ensure that relevant information on local defects of retinal layers is not obscured by irrelevant parts of the data. To meet the requirements of ophthalmologists, we will further test how well the adapted concepts work with actual patient data in different analysis scenarios.

NOVEL APPROACHES FOR VISUAL ANALYTICS OF RETINAL OCT DATA

Building on our fundamental investigations of complex 3D spatial data, we now focus specifically on data acquired by OCT examinations of the human eye. Here, we adapt the basic design concepts developed in Chapter 2 to arrive at new VA solutions dedicated to retinal OCT data. First, in Section 3.1, we take up the idea of extracting main data characteristics to reduce the complex 3D retinal layers to 2D attribute maps. Based on this, we improve their visibility in the 3D spatial context of one retinal OCT dataset by displaying them in linked 2D and 3D views. Then, in Section 3.2, we go one step beyond and consider also the comparison of retinal layers. This requires an even stronger data abstraction. We therefore transfer further ideas of our feature-based VA approach to capture and simultaneously analyze diagnostically relevant information from multiple OCT datasets. Our original publications on the novel VA approaches for retinal OCT data are included in Chapters 8 and 9 of Part ii.

3.1 IMPROVEMENTS IN THE VISIBILITY OF RETINAL LAYERS

The visual analysis of the retinal images acquired with OCT is essential for the detection of early signs of eye diseases or the investigation of disease progression. Yet, the volumetric data cannot be displayed in a single picture without obscuring the retinal layers. This hinders the visibility of small and local defects of the retina. As more and more image data are acquired with modern OCT devices, such details become even harder to see. Therefore, we developed a novel VA approach that improves the visibility of diagnostically relevant information about the layers in retinal OCT data. Our initial publication on this work received a best paper award [79]. The subsequently invited and extended journal article [27] is included in Chapter 8 of Part ii.

3.1.1 *Problem analysis*

To explain the visibility problem of retinal layers, we start with an overview of how the OCT data are currently usually analyzed. As illustrated in Figure 3.1, the eye is first scanned with an OCT device. The result is a highly detailed 3D tomogram of the multi-layered retina in the posterior eye segment (Fig. 3.1 (a) and (b)). The 3D tomogram consists of multiple 2D Cross-sectional Tomographic Depth Scans (*B-scans*). In the *B-scans*, the retina is represented by gray bands, with each band

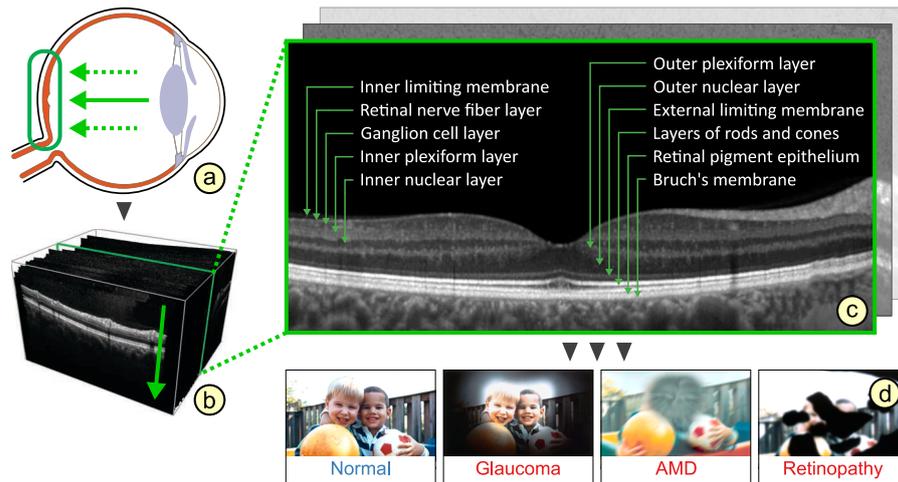


Figure 3.1: OCT examinations of the eye. The posterior segment of the eye is scanned (a). The result is a 3D tomogram (b) consisting of multiple 2D B-scans (c). Analyzing the retinal layers (gray bands in B-scans) allows to detect various eye diseases (d). Images of simulated vision loss © by National Eye Institute, National Institutes of Health.

corresponding to a retinal layer (Fig. 3.1 (c)). To detect the various retinal diseases (Fig. 3.1 (d)), the layers must be checked for defects, e.g., increased or decreased thickness. The most basic and widely used method is to visually analyze each B-scans by looking for variations in the width of the gray bands along the horizontal image axes.

In practice, however, the identification of such diagnostically relevant information within the volumetric OCT image data can be quite complicated. Particularly in the early stage of a disease, the defects tend to be locally confined and only affect certain regions of one or a few layers. Often, they also appear initially as small morphological alterations scattered throughout the 3D spatial data. The corresponding gray value changes of the bands are subtle and thus difficult to recognize and keep track of when viewing the image data. Moreover, since not just one but dozens of B-scans in a single OCT dataset have to be checked one after the other, the relevant information gets easily obscured and can go unnoticed. Doing the same for hundreds or thousands B-scans across several OCT datasets, further reduces the visibility of the information and renders the data analysis time-consuming and error-prone.

3.1.2 Related work

Current visualization approaches for retinal OCT data are based on commercial or non-commercial software. Commercial OCT software is mainly distributed by OCT device manufactures, e.g., Zeiss, Topcon, Heidelberg Engineering, and others. As for non-commercial OCT software, only few visualization approaches exist (e.g., [80–83]). To support the examination and measurement of retinal layers, most commercial and

non-commercial approaches integrate an automated boundary segmentation of the layers [84, 85]. Yet, the available visualization techniques for OCT images and segmented layers often differ. We identified three basic view types across existing OCT software tools:

- *2D cross-sectional views* show the B-scans individually, optionally with overlaid layer boundaries. This allows to view details, but as noted above, information is obscured across multiple images.
- *3D views* show the OCT image data as a volume rendering. This provides an overview but combined visualizations of the 3D tomogram and the layers are usually not available.
- *2D top-down views* show the segmented layers over a fundus image of the interior surface of the eye. This helps to relate the layers to the fundus, but only one layer can be examined at a time.

In summary, the available view designs address the visibility of retinal layers only to a limited extent. What sets our novel VA approach apart is that we combine and enhance all three types of views to convey the very information that is relevant for a diagnosis.

3.1.3 Attribute-driven visual analysis of retinal layers

We aim to make visible where in the 3D retinal OCT data which layers are affected by a disease and to what extent. Showing the raw OCT image data alone, however, makes it difficult to detect small and localized defects of the retinal layers. With our novel VA approach [27], we therefore focus on the visualization of meaningful attributes derived from segmented layer boundaries. This attribute-driven visualization helps us to determine which layers in the 3D spatial data are of interest and then improve their visibility accordingly. Another advantage of the layer attributes compared to the raw OCT images is that we can abstract from irrelevant dataset properties, e.g. acquisition-related artifacts in the orientation of the 3D tomograms. That way, we not only enable an accurate visualization of the layers within one dataset, but also a consistent presentation when switching to layers from other datasets.

COMPUTATION OF ATTRIBUTES OF RETINAL LAYERS: In order to capture the different defects of the retinal layers caused by various eye diseases, we have to consider a number of layer attributes. The most common attribute is the layer thickness (e.g. [72]). In addition, we also include other known attributes, e.g., the layer volume (e.g., [73]), and several new attributes, e.g., the curvature of the layer boundaries, to characterize the condition of the retina. We compute all these attributes for all data points and compile them into attribute maps for each layer. During visualization, we then do not show the 3D layers directly, but

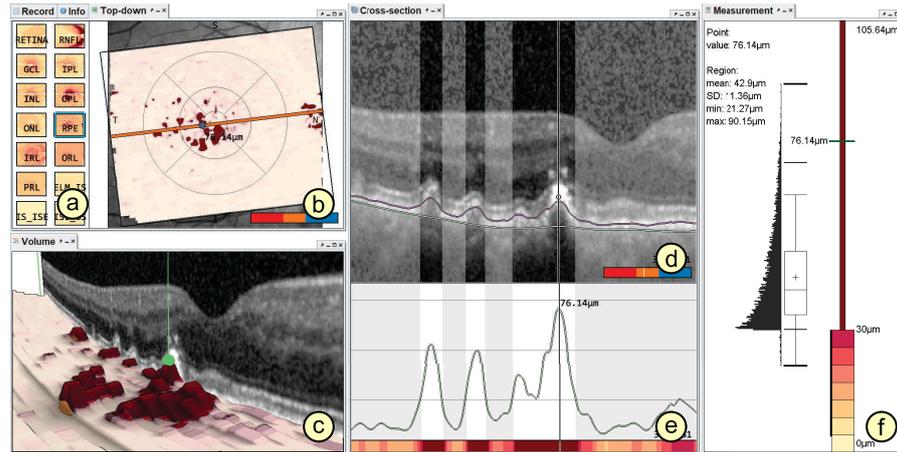


Figure 3.2: Visualization, emphasis, and selection of retinal layers. Shown are the 2D top-down view ((a) and (b)), the 3D view (c), the 2D cross-sectional view ((d) and (e)), and the measurement view (f). Parts of a layer are selected and highlighted in the linked views.

represent them using the various 2D attribute maps. This reduces the complexity of the layers and clarifies their structure and properties in 2D and 3D presentations of the OCT data. In addition, we offer an interactive access to the attributes to examine them in detail. This is in contrast to other approaches, which mainly considered the analysis of thickness maps in mostly static variants of the 2D top-down view.

VISUALIZATION OF RETINAL LAYERS IN 3D SPATIAL CONTEXT:

In our visualization design we must take care that no relevant information about the retinal layers gets overlooked in the 3D spatial context. To this end, we build on the key ideas from our characterization of 2D and 3D visualizations (Sect. 2.1) and our concepts for handling occlusion of 3D spatial data (Sect. 2.2). Specifically, we adapt and combine the established three basic 2D and 3D views for retinal OCT data (Sect. 3.1.2). We enhance each of these views by showing the derived layer attributes together with the segmented layer boundaries in the context of the OCT image data. Furthermore, we incorporate a new measurement view, which allows to inspect the layer attributes in detail. Our view design permits to start from an overview and then selectively narrow down the data to those layers that are of interest. This considerably improves the visibility of relevant information compared to the basic method, where only individual sections of the layers are viewed in one B-scans after the other (Sect. 3.1.1). As illustrated in Figure 3.2, each view conveys different aspects of the layers.

2D TOP-DOWN VIEW: To get a first idea of which retinal layers are of interest, it is necessary to start with an overview. Hence, the top-down view depicts all layers in a dataset, represented by their color-coded attribute maps (Fig. 3.2 (a)). For spatial reference, the

attribute map of a selected layer is enlarged over a fundus image of the inner eye surface (Fig. 3.2 (b)). This design extends existing displays by allowing all layers to be viewed at once without having to manually switch through them. Thus, the identification of layers with relevant characteristics is simplified.

3D VIEW: After a layer of interest has been determined, its spatial structure must be examined. Thus, the 3D view shows the boundaries of the enlarged layer in the top-down view as 3D surfaces colored by the attribute map (Fig. 3.2 (c)). The spatial context is further represented by a blended volume visualization of the OCT image data. Showing all data aspects in 3D allows for a faithful layer representation. Now, variations in the structure or attribute values of the layers can be viewed in 3D, which was not possible before.

2D CROSS-SECTIONAL VIEW: After the 3D overview of the layer boundaries, their exact shape can be inspected at close range. The cross-sectional view shows them section by section (Fig. 3.2 (d)). The boundaries of the selected layer are superimposed on the B-scans. At the bottom, the attribute values are displayed in an aligned line chart (Fig. 3.2 (e)). Using the linked overviews, we can jump directly to the interesting layer sections and analyze their profile alongside the various attributes. This would hardly be possible if only the individual B-scans of the OCT image data were shown.

MEASUREMENT VIEW: Finally, for a thorough examination, the visualized layer attributes must be quantified. To this end, the measurement view abstracts from space and summarizes the attribute values (Fig. 3.2 (f)). Measurements of individual points are displayed as lines and numerical text labels above a color legend. Next to the legend, measurements of all attribute values within a layer or selected part are shown as box-and-whisker plots. These measurements help to assess the layers in detail.

EMPHASIS AND SELECTION OF RETINAL LAYERS: The visualization of layer attributes, layer boundaries, and OCT images improves the overall visibility of the retinal layers in the linked views. However, when switching between datasets, layers, or attributes, it is not always easy to immediately decipher which parts of the layers are of interest. To make them more accessible, we therefore integrate methods for visual emphasis and interactive selection of layer parts.

The visual emphasis brings spatial relationships and distinctive attribute values to the attention of the user. In terms of space, it allows the structure of the layer parts to be recognized from the spatial context, e.g., by illumination and enhanced edges of the layer surfaces in the 3D view (Fig. 3.2 (c)). With regard to the attributes, it makes it easier to distinguish between layer parts with relevant or irrelevant values, e.g., by color accentuation of the highest or lowest layer values (Fig. 3.2 (b))

to (e)). In this way, interesting parts of the layers can be easily identified in the context of the overall data.

The interactive selection provides access to the identified layer parts for further investigations. Yet, since the linked views visualize the layers in different ways, coordinated selection methods are required. We consider two principal options: spatial selections and attribute-oriented selections. Both options help to specify layer parts that would be difficult to describe with the respective other option. They allow either to start from a certain part of the space to get to corresponding attribute values or vice versa.

SPATIAL SELECTION: The first option is needed to select certain spatial areas, e.g., anatomically defined parts of the retina. Hence, we offer various geometric selection methods in the spatial visualizations that are independent of the attribute values. To create a rough sketch of the layer parts, the top-down view offers a selection polygon. The sketch can be adjusted by fitting the polygon corners to the layer boundaries in the cross-sectional view and in the 3D view. This allows to refine their exact spatial location. The attribute values associated with the selected layer parts are then displayed in the measurement view.

ATTRIBUTE-ORIENTED SELECTION: The second option is needed to select value ranges of interest independent of the visualized spatial structures. We support this, for example, by specifying a single starting point with a specific value in the attribute maps. All neighboring points within the same value range are then automatically added to the selection. Alternatively, value ranges in the line chart (Fig. 3.2 (e)) or color legend (Fig. 3.2 (f)) can be brushed to automatically select all layer parts they are enclosed by. This helps to relate the value summary displayed in the measurement view to the possibly very complex shapes in the spatial context.

3.1.4 Discussion

Our novel attribute-driven VA approach effectively improves the visibility of diagnostically relevant information about the retinal layers. Based on our fundamental investigations (Sects. 2.1 and 2.2), we show meaningful layer attributes in 2D and 3D representations of the spatial context. With our combination of 2D and 3D views, the attributes can be examined together with the segmented layer boundaries and the OCT image data. We further put emphasis on interesting layer parts and offer interactive direct access to them. This ensures that all relevant parts of a layer are actually visible. Based on the reduction of the 3D layers to 2D attribute maps, we can even display all layers of a dataset side by side and allow the user to browse through them. In this respect,

our approach goes beyond what was previously possible in the visual analysis of retinal OCT data.

During tests of our solutions together with ophthalmologists, the experts pointed out the high quality and potential of the developed methods. The discussions with them also sparked the development of several extensions. For instance, to handle occlusions between layers, we introduced a new compound map that displays attributes for all layers in a dataset simultaneously [86]. In addition, we proposed new ideas to reveal the influence of acquisition parameters on the quality of retinal OCT data [87].

So far, all of our designs focused on the interactive visualization of one retinal OCT dataset. In the next section, we consider how to automatically extract relevant information about retinal layers in multiple datasets. To this end, we complement our attribute-driven VA approach presented here with the ideas we have developed in our feature-based visualization of 3D spatial data.

3.2 REDUCTIONS IN THE COMPLEXITY OF RETINAL LAYERS

The complexity of retinal OCT data makes it difficult to concentrate on the main characteristics during data visualization. Especially, when not just one but several datasets have to be considered, the individual layer attributes, layer boundaries, and OCT images become easily overwhelming. With such an amount of 3D spatial detail, it is not practical to display all the information of each dataset in a dedicated view. The comparison of several OCT datasets is nevertheless absolutely necessary to assess the condition of the retinal layers of individual patients and patient groups in relation to the normative data of healthy controls. This is why the layer attributes are currently strongly reduced using standardized spatial grids. Although the resulting grid data are easier to compare, they generally show only about 0.0034 % of the values compared to an attribute map for a typical dataset in our investigations. This can lead to a loss of diagnostically important information. Therefore, we developed a novel feature-based VA approach that achieves both: the presentation of local layer regions with relevant information at an appropriate level of detail and the abstraction of irrelevant information in the complex retinal OCT data. The original publication on our approach received a best paper award [88]. Our subsequently invited and extended journal article [28] is included in Chapter 9 of Part ii.

3.2.1 Problem analysis

As discussed in our fundamental investigations (Sect. 2.3.1), no feature-based visual analysis approach is available that allows to reduce the complex retinal OCT data to relevant information. To adapt our previous

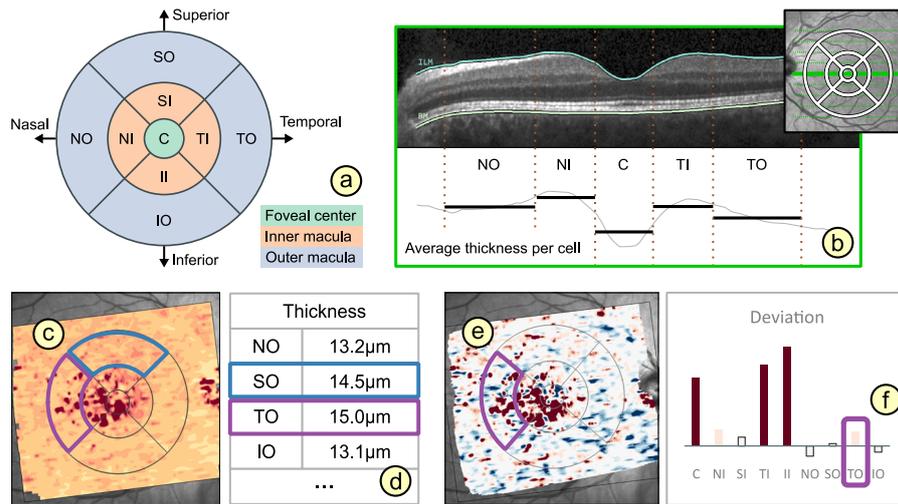


Figure 3.3: The ETDRS grid and related problems. The ETDRS grid cells divide the retina into nine regions (a), for each of which one aggregated attribute value is computed per layer across all B-scans (b). Cells with and without local variations in attribute values (red spots) are almost indistinguishable (c) based on the aggregated values (d). Local positive (red spots) and negative (blue spots) differences in attribute values within a cell (e) nearly cancel each other out (f).

ideas (Sect. 2.3.3), we therefore first consider how multiple retinal OCT datasets are currently analyzed and what problems are involved.

In general, the visual inspection of the raw images of more than one or two OCT datasets is not practical. This is why at present statistical methods are mainly used for the analysis of numerous datasets. Typically, selective measurements of retinal layers, e.g., thickness or volume [72, 73], are summarized across all datasets. To retrieve the measurements from each dataset in a standardized way, predominantly grids based on the Early Treatment Diabetic Retinopathy Study (ETDRS) [89] are applied. As shown in Figure 3.3 (a), the cells of the ETDRS grid divide the retina into nine large regions defined by three rings, i.e., a central foveal ring with 1 mm diameter, an inner macular ring (pericentral) with 3 mm diameter, and an outer macular ring (peripheral) with 6 mm diameter. The inner and outer rings are further divided into four quadrants (nasal, temporal, superior, and inferior). For each grid cell and retinal layer, one aggregated value, usually the arithmetic mean of the measured attribute, is computed across all B-scans (Fig. 3.3 (b)).

This current ETDRS grid-based approach has the advantage that it allows a compact representation of retinal layers from several complex OCT datasets. The aggregated grid values provide a quick overview of measurements in anatomically predefined retinal areas. However, there are also several disadvantages. On the one hand, the strong data aggregation does not accurately capture local variations of attribute values within the coarse grid cells (Fig. 3.3 (c) and (d)). On the other hand, the fixed grid layout does not always match irregularly shaped regions of

interest, so that relevant attribute values between cell boundaries may be obscured by other irrelevant values within the predefined cells. Both problems can lead to loss of information. Furthermore, aggregation artifacts may lead to false negative results when computing differences between datasets, since local positive and negative differences in attribute values within a grid cell can be nullified (Fig. 3.3 (e) and (f)). Our feature-based approach must therefore aim for a solution that not only shows attributes over an entire layer, but features over selected local regions where meaningful layer characteristics are present.

3.2.2 *Related work*

The **ETDRS** grid-based analysis is the most common way to deal with the complexity of multiple retinal **OCT** datasets. It has been widely used in ophthalmic research to study retinal layers from patients with a certain eye disease in relation to data from healthy controls (e.g., [90, 91]). Here, the standardized grid layout helps to compare results from different studies. Although other grid types are available, they were mostly designed for special applications, e.g., rectangular grids for asymmetry analysis of retinal layers in glaucoma diagnosis [92].

Since most existing **OCT** software does not support the analysis of multiple datasets by default, the grid data must typically first be exported and then analyzed with general-purpose statistics software. The results are usually reported in the form of tables or statistical plots (Fig. 3.3 (d) and (f)). This, however, provides little information about what the results refer to in the 3D spatial context. The **ETDRS** grids themselves are typically shown in separate static images to demonstrate the layout and labeling of the cells. These presentations are related to the layer-centric top-down view described in the last section (Sect. 3.1.2).

In addition to grid-based data analysis, machine learning approaches have recently been proposed for the diagnosis and referral of retinal diseases (e.g., [93]). They support the automated extraction of diagnostically relevant information from complex **OCT** data. So far, however, the available methods mainly target single **OCT** datasets of individual patients with intermediate or advanced defects. Furthermore, the interactive exploration of the extracted information is limited beyond the static presentation of the computed results.

In summary, all related approaches either focus on single **OCT** datasets or rely on the generalized reduction of retinal layers by **ETDRS** grids. What is missing is a selective reduction that also supports multiple datasets and the interactive presentation of the results in their 3D spatial context. The contribution of our feature-based **VA** approach therefore consists in an analytical abstraction of the retinal layers that allows to explore relevant information at an appropriate level of detail.

3.2.3 *Feature-based visual analysis of local regions of retinal layers*

Our aim is to reduce the complex retinal OCT data to the main characteristics of the retinal layers. This means on the one hand, we have to find a representation of the layers that is compact in terms of the amount of data to be analyzed. On the other hand, this also means that the layer representation must be expressive with regard to relevant details. With our novel VA approach [28], we meet both of these requirements by bringing together all our previous concepts. Our solution allows to abstract from irrelevant parts of the layers, while preserving all relevant parts without loss of information. We achieve this by applying our ideas of a feature-based visual analysis (Sect. 2.3.3) and focusing on local regions with meaningful characteristics in both the spatial context and the attribute space. First, the regions in the spatial context are determined by computing appropriate grid representations for each layer. This dataset-specific abstraction already reduces the overall data complexity to the necessary amount. We then formally describe and extract the regions in the attribute space of the layers. This provides us with information that is most likely relevant for a diagnosis. Finally, we adapt our interactive visualization design with 2D and 3D views (Sect. 3.1.3) to ensure that the computed regions are visible and explorable in the spatial 3D context of one or more complex retinal OCT datasets.

LOCAL LAYER REGIONS IN THE SPATIAL CONTEXT: We want to balance how much data need to be analyzed and how well relevant spatial details are presented. This balance may be different for the layers in each individual retinal OCT dataset or collection of datasets. The reason for this is that advanced layer defects often affect large connected areas of the layers, which permits a representation with a high spatial abstraction. Conversely, early layer defects may only involve small and locally confined areas scattered over the layers, which requires a representation with low spatial abstraction. In this connection, we have so far discussed layer representations with two different abstraction levels: the ETDRS grids (Sect. 3.2.1) and the attribute maps (Sect. 3.1.3). The ETDRS grids have a high spatial abstraction, which provides an overview, but with the potential loss of relevant details. The attribute maps have a low spatial abstraction, which offers details, but also for irrelevant parts of the data. In a sense, the ETDRS grids and the maps are at opposite ends of an abstraction spectrum.

We extend these two established representations to enable a spatial abstraction that corresponds to the given data. We start from the layout of the well-established ETDRS grid (highest abstraction with largest regions) and subdivide it incrementally to derive alternative grid layouts that capture more information (intermediate abstractions with smaller regions). In case none of the grids provides an adequate representation, we resort to the attribute maps (lowest abstraction with individual

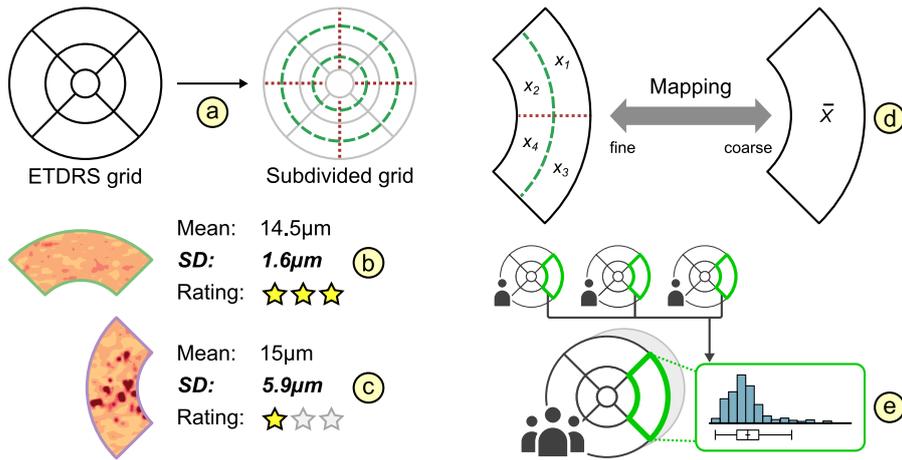


Figure 3.4: Computation of dataset-specific grids and their comparison. Alternative grids are derived from the **ETDRS** grid using radial (green) and sector-wise (red) subdivisions (a). The grids are rated based on the variation of attribute values within their cells (b and c). Grids with different layouts are compared by mapping fine cells to coarse cells and vice versa (d). The data of patient or control groups are compiled into an aggregated grid per retinal layer (e).

points). On the one hand, this allows us to match the local regions with relevant characteristics in the spatial context. On the other hand, the coherent abstractions ensure the comparability of the resulting grid representations. [Figure 3.4](#) shows the two steps of our approach:

COMPUTATION OF DATASET-SPECIFIC GRIDS: To capture local regions in the spatial context, we need to provide grids that represent the layer attributes on different spatial abstraction levels. Here, it is important to take the **ETDRS** grid as a basis, so that all derived grid layouts also refer to anatomically distinct areas of the retina. In principal, we obtain these layouts by radial and sector-wise subdivision of the **ETDRS** grid cells ([Fig. 3.4](#) (a)). The radial subdivision enables a more detailed analysis with regard to the macular rings and the sector-wise subdivision with regard to the macular quadrants. By incrementally doubling the radii or sectors, the size of the regions per cell decreases and the amount of stored grid values simultaneously increases in constant steps.

With the various possible grid layouts available, we then have to consider how the subdivision should actually be performed for a given dataset. In other words, we must find a way to determine which layout best matches the layer attributes, considering the cell sizes and the amount of information stored. To do so, we evaluate the correspondence between the layouts and attribute values based on the variation within each grid cell. Specifically, we assume that a cell with low variation is consistent with the underlying attribute values, while high variation indicates a po-

tential loss of information (Fig. 3.4 (b) and (c)). On this basis, we rate the grids and automatically suggest the best fit with the highest acceptable variation and lowest cell count.

SUPPORT FOR THE COMPARISON OF GRIDS: The computed grid layouts allow an appropriate representation of the layer attributes specific to one dataset. Now, we have to consider the comparability of different layouts from several datasets. For this purpose, we support mapping a fine grid to a coarse grid and vice versa (Fig. 3.4 (d)). For the mapping from fine to coarse, the finer cells are merged and the values aggregated. For the mapping from coarse to fine, the coarser cells are divided and the values recomputed. After mapping the grids, the corresponding cell values can be compared directly. Both ways are needed, e.g., to relate different subdivided grids of follow-up examinations to each other or to the **ETDRS** grid of the baseline examination when monitoring the progression of a patient's disease. Mapping the grids of multiple datasets is also required to determine a common layout for the layer attributes of a patient group or the normative data of a control group.

LOCAL LAYER REGIONS IN THE ATTRIBUTE SPACE: Our computation of grids and maps allows to capture local regions in the spatial context of the given data. Based on this, we now concentrate on features over local regions in the attribute space of the layers. These features must not only further reduce the data complexity, but also provide relevant information for the diagnosis of eye diseases. Furthermore, they also have to facilitate the comparison of this information in several datasets. In this respect, our approach covers the specification and extraction of the features as well as their comparison.

For the feature specification, we collaborated with ophthalmologists. Using their domain knowledge, we found that the features must describe regions with abnormal attribute values. To classify the values as normal or abnormal, two types of features with different thresholds were considered: (1) clinically defined attribute ranges and (2) attribute intervals specific to control data. For the first feature type, a fixed minimum and maximum per layer are predefined for all given data. This allows to generally detect abnormal regions of all individual datasets. For the second feature type, the limits are specified per grid cell or map point of a layer. These features are more spatially specific and help to identify abnormal regions relative to the control data used.

For the feature extraction, we require different methods. The first feature type involves a simple check of the attribute values against the predefined limits of each layer. This can be done on the fly during visualization. The computation of the second feature type is more elaborate, as the control data are not necessarily fixed. The control intervals per grid cell or map point must rather be derived from numerous control datasets in advance. Hence, we first determine a common alignment

and layer representation. Next, we precompile the layer attributes of all control datasets into one aggregated grid or map per layer (Fig. 3.4 (e)). Finally, the control intervals and a summary of all source values are computed for each aggregated cell or map point. With this preprocessing, the attribute values can be classified during visualization also for the second feature type.

For the feature comparison, we can then directly relate the local regions in the attribute space of the layers. This has two advantages. First of all, we are able not only to compute differences between the attribute values of the individual datasets, but also with regard to a common reference. This helps to assess how far each dataset is from the normative ranges defined by the feature thresholds. The second advantage is that the local attribute regions are extracted on top of the computed regions in the spatial context. This further narrows down which subsets of grid cells and map points actually contain abnormal values and should therefore be considered for the comparison between datasets. Depending on the type of feature used, comparisons of layers of individual patients and individual patient to control data are thus easier. By applying the described preprocessing to patient data, we also enable the comparison of average attribute values between the aggregated layers of patient groups and patient and control groups. That way, we can focus the visual analysis on the diagnostically relevant information in one or even multiple retinal OCT datasets.

VISUALIZATION OF LOCAL LAYER REGIONS IN THE SPATIAL CONTEXT: Our interactive visualization design must support the exploration of the local regions in the attribute space and spatial context of the retinal layers. Here, we first explain how the spatial regions defined by the various grids are presented for the layers of individual patients and aggregated layers of patient groups. Since these regions considerably reduce the complexity of the overall data, we must take care that the abstracted information is conveyed in an understandable way. On the one hand, this means that we have to ensure the visibility of the regions when they are presented in their 3D spatial context. For this purpose, we adapt the four linked views described in the previous section (Sect. 3.1.3). On the other hand, as the regions may vary from dataset to dataset, we need to make the grids adjustable so that more or less details can be retrieved on request. This way, the local regions in the spatial context can be investigated interactively.

VISUALIZATION OF GRIDS IN 2D AND 3D VIEWS: The local regions defined by the grids have to be made visible together with the other aspects of the retinal OCT data. To this end, the biggest change has been made to the 2D top-down view. Besides the attribute maps, it now also shows an overview of the corresponding grids for each layer of a patient or aggregated layer of a patient group. The grid cells are colored based on the stored attribute values.

The measurement view has been adjusted to provide additional information about these cells. For grid cells selected in the top-down view, the average cell value and a statistical summary of all contained attribute values are displayed. The 3D view and the 2D cross-sectional views were also slightly modified. For selected grid cells of individual patients, they highlight the respective spatial regions on the layer surfaces and the individual *B-scans*. This helps to relate the cell values and their locations to the layer boundaries and *OCT* images visible in these views.

INTERACTIVE ADAPTATION OF GRIDS: To start the exploration of local regions in the spatial context, suitable grids are automatically suggested based on the computed grid ratings for each retinal layer. These initial grids provide an overall spatial abstraction of the layers. To investigate certain cell values or spatial regions in detail, however, it must be possible to adjust this abstraction locally if necessary. Thus, we allow an interactive browsing through alternative grids with more or less cells (lower or higher spatial abstraction). In case cells of a selected grid do not match the underlying attribute values (high variation within the cells), we mark their border to indicate a possible loss of information. These or other cells of interest can then be examined in more detail by subdividing them or replacing them with corresponding sections of the attribute map. The other cells of less relevant spatial areas can be merged to obtain context information. This not only reduces the overall data complexity, but also allows to selectively focus on the spatial regions that are relevant for a layer.

COMPARISON OF LOCAL LAYER REGIONS WITH ABNORMAL VALUES: So far, we explained the interactive visualization of grids in the spatial context of the layers of individual patients and aggregated layers of patient groups. However, the condition of these layers can only be accurately assessed if the local regions are considered in both the spatial context and the attribute space and two or more datasets are related to each other. This typically includes comparisons between several individual patients, between several patient groups, and between individual patients and patient groups. While the first attribute feature type allows the patient data to be classified in a common way, the difficulty here is that the grid cells or map regions with abnormal values can occur at different locations in the retinal layers. Thus, we have to find ways to visually bring these locations close together using the four linked views. On top of that, considering the second attribute feature type, all these comparisons of patients must also be made in relation to data from healthy controls. In this respect, care must be taken to ensure that grid cells and map regions with abnormal deviations from the control intervals are clearly visible in the spatial context. With our comparative visualization design, we enable all of these comparison of local regions

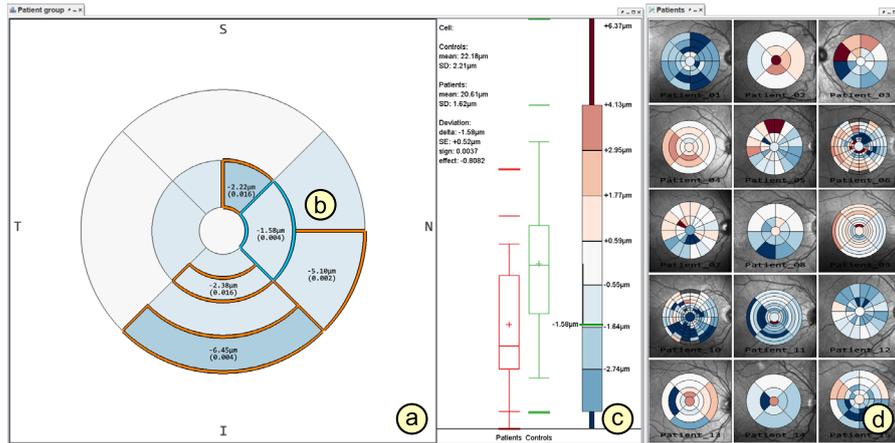


Figure 3.5: Comparison of multiple patients to control data. The top-down view (a) shows the deviation grid of an aggregated layer of a patient group. Individual grid cells were subdivided to investigate regions with significant differences (orange cell border) in detail (b). The measurement view (c) shows the patient values of a selected cell (blue cell border) in relation to the control interval. The deviation grids of the layers of individual group members are displayed next to each other in separate top-down views (d).

in the attribute space and spatial context of retinal layers. [Figure 3.5](#) illustrates the comparison of multiple patients to control data.

COMPARISON OF ABNORMAL REGIONS IN MULTIPLE PATIENTS: To compare abnormal regions in the layers of multiple patients, they must be shown in their own spatial context, but brought visually close together. For this purpose, we use the first attribute feature type to highlight the grid cells and map regions with abnormal attribute values using distinct colors. We then display the data of the patients in separate views. For example, as shown in [Figure 3.5](#), several 2D top-down views are added to the user interface to compare the grid of an aggregated layer of a patient group and the grids of the layers of individual group members. To ensure that matching parts of these layers are visible, we link the views and synchronize the interactive navigation. This helps to compare the location of the abnormal regions, especially when multiple 3D views of individual patients are added. To also compare the values of the abnormal regions, we synchronize the selection of grid cells or map regions in the different views. The measurement view can then be used to display the respective values for each compared layer side by side. This functionality offers a distinct advantage over existing [OCT](#) software, which usually only permits the visualization of a single patients at a time and without the computed abnormal regions.

COMPARISON OF ABNORMAL REGIONS TO CONTROL DATA: To relate abnormal regions in the layers of individual patients or patient groups to control data, a spatially accurate comparison of the attribute values stored in the grids and maps is required. Therefore, we use the second attribute feature type to compute the patient deviation per grid cell or map point from the control intervals. The resulting deviation grids and maps are then displayed in the linked views using a diverging color palette (Fig. 3.5 (a)). Grid cells and map regions with values outside the control intervals are highlighted. This allows to identify the exact location of the abnormal regions. To further quantify the relevance of the observed differences between patient and control groups, we apply statistical tests to every grid cell and map point. The resulting areas with significant deviations are marked by emphasized borders in the grids and maps (Fig. 3.5 (b)). The measurement view can in turn be used to display additional statistical information, e.g., p -values and effect sizes, for selected cells or map regions (Fig. 3.5 (c)). In this way, abnormal layer regions of aggregated patient groups (Fig. 3.5 (a)) and individual group members (Fig. 3.5 (d)) can be assessed with high spatial specificity.

3.2.4 Discussion

Our novel feature-based VA approach with different levels of spatial abstraction achieves a balance between the compactness and the expressiveness in the representation of retinal OCT data. Based on the reduced data complexity, we can focus the visual analysis on local layer regions with diagnostically relevant information that are specific not only to individual patients but also to whole patient groups. The combination of our interactive visualization of with statistical quantification makes it possible to compare patients to control data with high spatial specificity. This is a significant improvement over previous analysis approaches, which rely solely on the statistical evaluation of multiple OCT datasets, using only the ETDRS grid and no reference to the spatial context.

Together with experts from ophthalmology, we successfully tested our approach with data from two cross-sectional studies. All experts noted the effectiveness of our solution for the early detection of abnormal regions with small and localized defects of the retinal layers. Overall, more medical findings could be obtained with less effort than with the current analytical approach. More results can be found in the next chapter (Ch. 4).

Based on the expert feedback, we also introduced a new measurement method that supports the comparison of grid cells and irregular map regions with significant differences [94]. In addition, we extended our solutions to provide unified access to retinal OCT data from different sources [27]. This is a first step towards a common representation and

feature-based comparison of several datasets that were acquired with different OCT devices and have different data formats.

3.3 SUMMARY

In this chapter, we described new solutions for the VA of retinal OCT data. In Section 3.1, we introduced improvements to the visibility of retinal layers in their 3D spatial context. In Section 3.2, we then focused on reductions in the complexity of retinal OCT data and presented a novel feature-based solution with local regions in both the attribute space and the spatial context of the layers.

On the basis of the design concepts described here, we developed a prototypical VA tool for retinal OCT data. At the back end, its modular software architecture combines all the described methods in three functional components for computation, visualization, and interaction. At the front end, a flexible user interface permits uniform access and interactive visual analysis not only for one, but for numerous complex retinal OCT datasets. The functionality offered in this way is unique, as no equally coherent solution existed before. More technical details on the tool design can be found in Chapters 8 and 9 of Part ii of this thesis.

The tool and its functionality complete the conceptual and technical part of our VA approach to retinal OCT data. In the next chapter, we describe how the design concepts can be applied in practice and outline what results can be achieved with our tool compared to the data analysis approach currently used in ophthalmology.

APPLICATION OF VISUAL ANALYTICS METHODS IN OPHTHALMOLOGY

Having developed our new VA concepts for retinal OCT data, it is necessary to test how well they work in practice with actual patient data. In this chapter, we therefore apply and evaluate our approach in the context of cross-sectional studies. First, in Section 4.1, we introduce a new VA procedure for the analysis of single-patient and multi-patient data and compare its individual analysis steps with the current analysis practice. Then, in Section 4.2, we apply our and the current analysis approach to the data of patients with diabetes mellitus. We take a closer look at the results obtained and reflect on the advantages and disadvantages of both approaches. The full description of the VA procedure is contained in Chapter 9, and our original publications on both studies are included in Chapters 10 and 11 of Part ii.

4.1 EVALUATION OF THE ANALYSIS PROCEDURE PERFORMED WITH VISUAL ANALYTICS METHODS

Our new VA approach supports the analysis of retinal OCT data from both individual patients and patient groups. So far, however, we described the computation and interactive visualization methods mostly individually. This does not immediately make it clear how they are to be applied in practice. Especially, if we take as a concrete analysis example the evaluation of cross-sectional studies, which requires several methods to handle multiple analysis tasks with the data of different patients. To get a general idea not only of the individual methods, but also of the effort involved with our overall VA approach, we have to consider how the methods work together when applied in such an application example. Therefore, we present a novel VA procedure for the evaluation of single-patient and multi-patient data and, based on this, discuss differences in the way the analysis is performed with our approach and current analysis practice. Our original work on the VA procedure was published in [28] and is included in Section 9.5 of Part ii.

4.1.1 *An analysis procedure for single-patient and multi-patient data*

We aim to achieve both: (1) to demonstrate the interplay of our VA methods and (2) to assess them in terms of the effort required to analyze single-patient and multi-patient data. For the first part, a complementary analysis procedure is required that shows step by step when and how which VA method is to be applied. For the second part, the VA

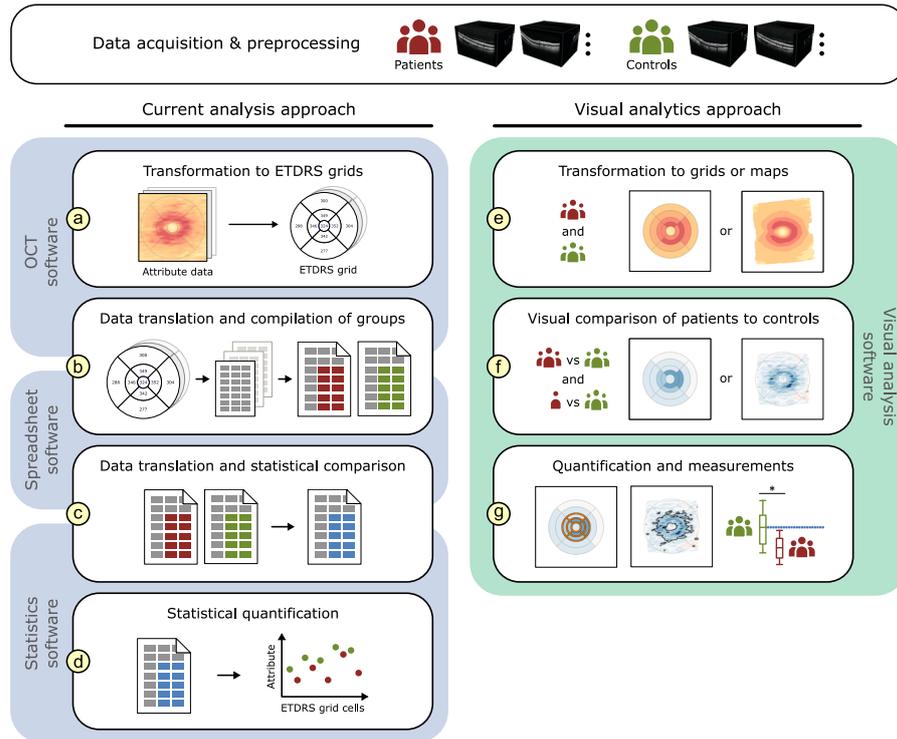


Figure 4.1: Comparison of the current **ETDRS** grid-based analysis (**CA**) procedure and our **VA** procedure for the evaluation of cross-sectional study data. After a common data preparation stage, **CA** comprises four analysis steps (a to d) with three software tools and **VA** involves three analysis steps (e to g) with our visual analysis tool.

methods must be compared for each step of the procedure with the methods currently used in ophthalmology. In this regard, the evaluation of cross-sectional study data is a suitable application example. That is because it covers the entire process, starting with the preparation of retinal **OCT** datasets, through the visualization of individual patients and patient groups in comparison to healthy controls, to the measurement of findings. In order to assess our approach in this context, we collaborated with ophthalmologists. First, we conducted interviews and observations on the job to better understand their current analysis practice based on the **ETDRS** grid. By participating in the evaluation of two studies [90, 95], we gained further insights into the course of the statistical analysis of the grid data with a mixture of commercial and non-commercial **OCT** software. Together, we then compiled the steps connected with the new **VA** approach. To do so, we demonstrated our methods, jointly re-evaluated the data of the two studies, and gathered feedback. This way, we identified four stages common to the Current Analysis (**CA**) approach with **ETDRS** grids and our new **VA** approach. Figure 4.1 shows the stages of both approaches side by side.

1. *Data preparation*: The first stage involves **OCT** image acquisition and layer segmentation. Afterwards the data quality is checked.

This is necessary to prevent artifacts from interfering with the data analysis. In *CA*, often each dataset of each subject must be validated individually. In *VA*, we reduce this effort by displaying several datasets simultaneously and point out missing values or insufficient signal strength. We also allow to correct segmentation errors, as they are directly visible in layers represented by subdivided grids or maps. This is not always the case in *CA* due to the higher spatial abstraction in larger cells of the *ETDRS* grid.

2. *Data transformation*: In the second stage, all datasets are registered and the the layer attributes are computed. In *CA*, position and laterality between subjects must often be matched by hand. The layer attributes are then reduced via the *ETDRS* grid (Fig. 4.1 (a)). The limited number of cell values is easier to handle in the following data translation and statistical analysis. In *VA*, we offer semi-automated registration of datasets and the layer attributes are aggregated to appropriate grids or maps (Fig. 4.1 (g)). This promotes a spatially more precise analysis of individual patients, patient and control groups, and optional subgroups.
3. *Comparison of patients and controls*: In the third stage, the patient and control data are assembled and compared. In *CA*, the *ETDRS* grid values are first copied, often manually, into spreadsheet software to compile patient and control groups (Fig. 4.1 (b)). Statistics software is then used to compute deviations and present them as tables or diagrams (Fig. 4.1 (c)). In *VA*, patients are visually compared to controls via deviation grids and maps (Fig. 4.1 (f)). This shows where exactly and to what extent abnormal regions exist in the spatial context. Our interactive exploration allows the analysis to start with the entire patient group to get an overview, then move on to patient subgroups to identify influencing factors, and finally to individual patients to detect special cases.
4. *Quantification of results*: In the fourth stage, the findings are quantified and further investigated. In *CA*, statistical tests are computed for each standardized *ETDRS* grid cell (Fig. 4.1 (d)). In case of significant differences, the patient values are plotted in relation to the controls. While this helps to test predefined hypotheses, localized attribute variations remain unmeasured. In *VA*, all statistical significant cells or map regions are highlighted in the deviation grids and maps (Fig. 4.1 (g)). This way, the spatial distribution of abnormal regions can be detected precisely. Interactive measurements allow to further investigate local regions of patient groups and to display corresponding layer parts for individual patients together with layer boundaries and *OCT* images.

4.1.2 Discussion

The comparison of our new VA procedure with the CA procedure helped us to identify advantages and disadvantages of the different methods in each analysis stage. Most importantly, our coherent solution simplifies the overall analysis process. This is because we provide our VA methods in one holistic tool. In CA, on the other hand, it is necessary to switch between the methods of various special and general-purpose software to evaluate the study data. In their feedback, the ophthalmologists therefore emphasized the reduced analysis effort as a major advantage of our approach. They further liked that the interplay of the methods in our tool made it possible to go back and forth between comparisons, measurements, and adaptations of grids and maps for individual patients and patient groups. They felt that this flexible analysis process would help them to reveal more subtle deviations of the retinal layers. Such study-specific results are essential for the early detection of eye diseases. This flexibility is in contrast to CA, where the separate software packages mean that if, after some analysis steps, it turns out that a previous step needs to be adjusted, the entire procedure must be reset to make the desired change and then repeated from that point on.

4.2 EVALUATION OF THE ANALYSIS RESULTS OBTAINED WITH VISUAL ANALYTICS METHODS

The discussion of our new VA procedure in relation to the current ETDRS grid-based procedure allowed us to identify general differences in computation, presentation, and interaction with retinal OCT data. To further evaluate our solution, it is necessary to also take a look at the analysis results. For this purpose, the complex evaluation of data from cross-sectional studies is again a suitable application example. This time, especially qualitative and quantitative differences in the presentation and measurement of findings with our and the current analysis approach have to be considered. Therefore, we apply both approaches to the data from two cross-sectional studies and reflect on their advantages and disadvantages based on the results obtained. The original publications on the evaluation of both studies are included in Chapters 10 and 11 of Part ii.

4.2.1 Study scenarios with early and advanced defects of retinal layers

Our goal is to evaluate the effectiveness of our new VA approach in obtaining diagnosis-relevant analysis results. We are particularly interested in how well disease-related variations in the attributes of the retinal layers can be detected. To evaluate our approach in this respect, we need to compare the results of our methods with the results of the current practice of using ETDRS grids. Especially, qualitative and quanti-

tative differences in the presentation and measurement of the respective analysis results must be taken into account.

Given this context, we again collaborated with ophthalmologists. Together, we first determined two different scenarios to compare both analysis approaches in a realistic setting. On the one hand, we considered patients in an early stage of disease, where the ophthalmologists expected mostly small and local layer defects. On the other hand, we included patients in a more advanced stage of disease, where usually larger layer defects are present. We then matched these scenarios with data from two cross-sectional studies in pediatric and adult patients with diabetes mellitus and healthy controls of comparable age. Based on the analysis procedures described in the last section (Sect. 4.1), we prepared the study data and then defined individual objectives for each study before performing the analysis steps for the comparison between patient and controls with both methods.

4.2.2 Objectives and design of two cross-sectional studies

Previous studies have shown alterations in the thickness of retinal layers in pediatric and adult patients with diabetes mellitus type 1 or 2 compared to healthy controls (e.g., [21, 90]). However, given the current analysis practice, we hypothesized that these alterations are not always adequately captured by the CA approach based on ETDRS grids. Thus, the general goal of both of our studies was to test whether the detection of local regions with variations in layer thickness can be improved by applying our VA approach. In addition, we defined one secondary objective for each study. For the first study, we wanted to assess the visual identification of local regions with abnormal thickness variations, as this is particularly important in an early stage of the disease. For the second study, we planned to evaluate the measurements of abnormal regions, since we assumed that quantitative differences in the results of CA and VA are more pronounced at a later stage of the disease.

The study data were obtained by OCT examinations of the two patient groups and corresponding control groups. In both studies, the inner retinal layers were segmented and the Total Thickness of the Retina (TR) and the thickness of the Retinal Nerve Fiber Layer (RNFL), the Ganglion Cell Layer (GCL), the Inner Plexiform Layer (IPL), and the Inner Nuclear Layer (INL) were examined. To compare the patients to the controls, we visualized the differences in layer thickness using ETDRS-based deviation grids in CA and deviation maps in VA (Sect. 3.2.3). Statistical tests were performed on the grid cells and map points to mark local regions with statistically significant deviations ($p < 0.05$; orange cell borders in grids and black outlines in maps). For the study evaluation, a pair analytics approach [29] was used to steer the analysis and reflect on the results obtained. One visualization expert and at least one or two ophthalmologists participated.

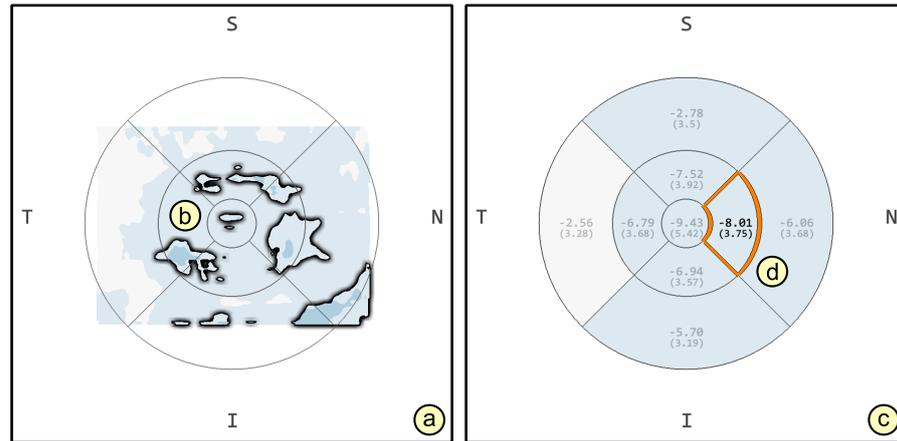


Figure 4.2: Patients with type 1 diabetes mellitus compared to controls. The map (a) depicts several irregularly shaped regions with significant thinning (black outlines) in the inner macular ring of TR (b), while the grid (c) shows only one significant cell (d; orange border).

4.2.3 Identification of early layer defects with ETDRS grids and maps

In the first study, the ophthalmologists wanted to investigate how well defects of retinal layers can be identified in the visualized grids of CA and the maps of VA. In particular, they were interested in whether the methods are capable of capturing small variations in layer thickness at an early stage of diabetes mellitus, which are assumed to be limited to local regions scattered over the layers. For this purpose, we analyzed the retinal layer thickness of 26 pediatric patients with type 1 diabetes mellitus (mean age of 14.5 years) in relation to the data of 29 age-matched controls (mean age of 14.5 years). The patients had a mean diabetes duration of 5.6 years without clinical signs of diabetic retinopathy in result of other examination methods. Here, we summarize the qualitative differences between the two analysis approaches. Our complete statistical analysis of CA can be found in [90] and the VA of the study data in Chapter 10.

Regarding the comparison between patients and controls, the ophthalmologists generally noted considerable thinning in several retinal layers in the results of both approaches. The highlighting of statistically significant areas in grids and maps helped them to identify relevant parts of the layers. Significant grid cells and map regions were found primarily in the inner macular ring of TR and RNFL, and selectively in nasal direction of the outer macular ring. The least amount of significant deviation was detected for INL. These results are in accordance with results of other studies that reported thickness reduction of inner retinal layers in patients with type 1 diabetes mellitus as an early event of neurodegenerative change (e.g., [22, 96]).

Regarding the comparison of both analysis methods, the ophthalmologists pointed out a key difference in the results. The maps revealed

additional details of the spatial distribution of local regions with abnormal deviations, while the ETDRS grids could not detect significant thinning in certain areas. This was particularly noticeable in the maps of the entire patient group, on which more small and isolated regions with significant thinning were visible. The ability to see these regions in the spatial context was new to the ophthalmologists. They also observed a tendency of CA to produce false negative results in the study. For example, several irregular-shaped map regions with significant thinning that cross all quadrants of the inner macular ring were detected by VA (Fig. 4.2 (a) and (b)), but only in one significant ETDRS grid cell of CA (Fig. 4.2 (c) and (d)). If the relevance of the observed deviation is judged by statistical significance alone, the local thinning within the other ETDRS grid cells could go unnoticed in the current analysis practice. Using the flexibility of the VA approach, the ophthalmologists further explored patient subgroups and individual patients. By relating the spatial patterns in the maps to clinical parameters, they were able to identify a patient subgroup of severe cases (high deviation, long duration of disease, and high blood glucose) and a subgroup of mild cases (low deviation, short duration of disease, and low blood glucose). This helped them to consider the influencing factors and data consistency in their conclusions from the first study.

4.2.4 *Measurement of advanced layer defects with ETDRS grids and maps*

From the results of the first study, the ophthalmologists concluded that abnormal variations in retinal layer thickness do actually occur in the early stages of diabetes mellitus. They also verified that small deviations can be made visible with our methods. In addition to their identification, however, it is necessary to measure them accurately. Such measurements are especially important for monitoring the degree of retinal alteration observed during the progression of a disease. This is because at a more advanced stage of diabetes mellitus, the defects are expected to lead to larger local regions with higher variations in layer thickness than the early defects visible in the first study. Therefore, with the second study, the ophthalmologists now wanted to investigate how well our interactive VA methods can measure such thickness variations compared to the CA method. To this end, we used both methods to analyze the quantitative differences in the layer thickness of 33 adult patients with type 2 diabetes mellitus (mean age of 60.4 years) and 40 controls of comparable age (mean age of 56.0 years) [31]. The patient group had a mean diabetes duration of 13.3 years and included several patients with mild to moderate diabetic peripheral neuropathy. Here, we summarize the main differences between the two analysis approaches. The full evaluation of the study data is included in Chapter 11.

The ophthalmologists began the study evaluation by reviewing the grids and maps for statistically significant deviation of patients from

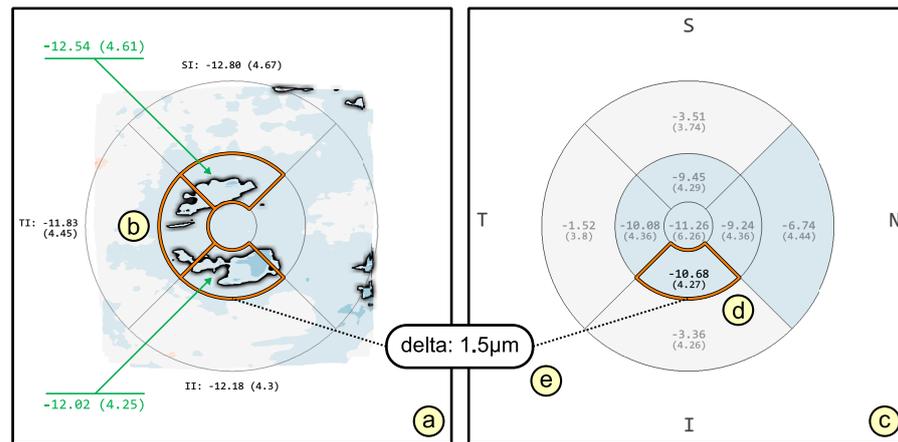


Figure 4.3: A patient subgroup with type 2 diabetes mellitus and mild to moderate diabetic peripheral neuropathy compared to controls. The map (a) depicts several regions with significant thinning in TR (b; black outlines). The grid (c) shows only one significant cell (d; orange border). The measured delta (e) of the region part in the map (a) and the corresponding grid cell (c) shows an underestimation of the deviation in the grid. Two regions of interest are selected in the map and measured as single defects (green arrows).

controls. Similar to the first study, they identified considerable thinning in the inner macular ring and selective quadrants of the outer macular ring for TR and RNFL. Different in this study, however, was that the thinning extended to larger connected regions, which were also found in GCL and IPL. This is in accordance with results of other studies in patients with type 2 diabetes mellitus (e.g., [21, 97]). Using the VA approach, they could again reveal additional local regions with significant thinning and their exact spatial extension in the maps. For example, Figure 4.3 shows the comparison between a patient subgroup with mild to moderate diabetic peripheral neuropathy and controls. The upper irregular map region with significant thinning crossing the superior-temporal quadrants of the inner macular ring was detected only by VA (Fig. 4.3 (a) and (b)), but not as significant ETDRS cells by CA (Fig. 4.3 (c) and (d)).

After checking the deviation grids and maps, the ophthalmologists continued with measurements of the identified local regions with statistically significant deviation. For CA, they followed the common practice of using the mean deviation and standard error values of the corresponding ETDRS grid cells. To match these values in VA, we introduced a special measuring method for the irregular regions marked on the maps [31, 94]. Specifically, we divided the map regions based on the areas defined by the ETDRS grid cells and automatically computed the mean deviation and standard error values of only the significant parts within each cell area. This enabled a direct comparison of the measurements of the grids and the maps. The divided map regions further

helped the ophthalmologists to evaluate the measurements in relation to the anatomically different areas of the retina, i.e., the foveal center and the inner and outer macular rings, and to interpret them with respect to the nasal, temporal, inferior, and superior directions.

When comparing the measurements, the most important difference between *CA* and *VA* noted by the ophthalmologists was a tendency of the *ETDRS* grid values to underestimate the patients' deviation (Fig. 4.3 (e)). The reason for this is that the maps allow to measure only the statistically significant parts of the cell areas without also averaging the non-significant parts, as is done in *CA*. In addition, the *VA* method allowed the ophthalmologists to interactively measure also regions that cross several quadrants in a macular ring as a single defect. Examples are the irregular-shaped regions highlighted on the deviation map in Figure 4.3 (a; green arrows). They argued that such regions need to be closely monitored to assess the degree and spatial extent of thickness variations in relation to clinical findings and disease progression. In *CA*, this was not possible with the predefined *ETDRS* grid.

4.2.5 Discussion

The evaluation of the two cross-sectional studies has shown that our *VA* methods are not only able to improve the identification of retinal layer defects in their spatial context, but also to measure them more accurately. With our comprehensive *VA* tool this is now possible for whole patient groups, patient subgroups, and individual patients in early or advanced stages of a disease. Consequently, new insights could be gained from the study results, as additional local regions with irregular layer alterations could be identified and measured. Especially in patients with early diabetes mellitus, these regions were previously not recognized by the *CA* and remained unnoticed. The ophthalmologists therefore concluded that our solutions have expanded their capabilities in the early detection of eye diseases and deepened their understanding of early retinal defects induced by diabetes mellitus. The opportunity to obtain more qualitative and quantitative findings was seen as the main advantage of our approach. For this reason, they emphasized in their feedback the importance of choosing an appropriate data analysis method for the evaluation of cross-sectional studies. They strongly recommended further studies together with the *VA* approach and a safe transition of the *VA* methods into the clinical settings for efficient investigation of retinal and other systemic pathologies.

The two studies and feedback from collaborating ophthalmologists suggest that *VA* is a useful extension of *CA*. Here, we compared both approaches using *ETDRS* grids (highest spatial abstraction) and maps (lowest spatial abstraction). This helped us to demonstrate respective advantages and disadvantages. Since we actually gained new medical insights in the studies, we made them available to the scientific commu-

nity in ophthalmology together with an overview of our approach. To do so, we first introduced the general ideas on using VA for the analysis of retinal OCT data in two articles in ophthalmic journals [98, 99] and presented them at major national and international ophthalmic conferences [100–103]. We then jointly published the detailed study results in two biomedical journals [30, 31]. We also tested the subdivided grids (intermediate levels of spatial abstraction; Sect. 3.2) for the evaluation of cross-sectional study data. The initial results showed that they are equally helpful in detecting abnormal layer regions, but with the benefit of an adjusted spatial resolution compared to the maps (cf. Ch. 9).

4.3 SUMMARY

In this chapter, we compared our new VA approach to the current analysis practice for retinal OCT data. On a procedural level, we showed in Section 4.1 that our methods reduce the effort for the analysis of single-subject and multiple-subject data from cross-sectional studies. On the level of the results, we then demonstrated in Section 4.2 that they also help to gain more insights from the evaluated study data, especially when patients in an early stage of the disease are involved. This indicates that our design concepts described in the previous chapters actually achieve the desired results when applied in practice.

To arrive at these solutions, we closely worked together with ophthalmologists from the Rostock University Medical Center, Rostock, Germany. In addition, we collaborated with experts from our industrial project partner Heidelberg Engineering GmbH, Heidelberg, Germany, which is one of the leading manufacturers of modern OCT devices and associated analysis software. Together, we continuously evaluated our designs throughout the development. At the beginning and during the development, we followed a participatory approach to develop suitable algorithms, visualizations, and interactions. Between the development phases, we jointly tested the developed methods using pair analytics to gather details of the practical application. After major development milestones (in the middle and at the end of the project), we conducted demonstration and feedback sessions and collected structured feedback with questionnaires. The aim of these evaluations was to assess the applicability, suitability, sensitivity, and specificity of our developed VA methods. All experts judged these aspects positively, emphasized the novelty of our solutions, and expressed their interest in using our tool in their own field of work. They further highlighted the variety and scope of the functionality provided, the faithful 2D and 3D presentations, and the ease with which regions of interest can be isolated. While the ophthalmologists saw their analysis requirements met, the industry experts concluded that our approach is a useful extension of the existing OCT software portfolio. Hence, they deemed a future integration into a commercial front end highly desirable.

SUMMARY

In this first part of the thesis, we presented our novel **VA** approach for retinal **OCT** data. In total, we summarized and discussed 6 original research works. Even though some of them initially had a different application context, we were able to successfully transfer the key ideas from one work to the next. Specifically, we have demonstrated how the individual works build on each other and how they have contributed to the development of the new **VA** methods for retinal **OCT** data.

On our way to the presented solutions, we discussed three consecutive parts in the chapters of **Part i**. They are our initial investigations of fundamental problems in the **VA** of complex 3D spatial data (**Ch. 2**), based on them, the development of the novel concepts for the **VA** of retinal **OCT** data (**Ch. 3**), and the application of our **VA** concepts in ophthalmology (**Ch. 4**). In each of these parts, we made several contributions. **Figure 5.1** shows them in an overview.

BASIC INVESTIGATIONS: As part of our basic investigations of complex 3D spatial data, we have contributed new **VA** solutions to two fundamental problems: (1) the visibility of relevant information and (2) the reduction of data complexity. We started by characterizing principal design options for the visualization of the attributes and the spatial context in 2D and 3D. On this basis, we proposed visibility widgets that combine different design patterns to deal with three different types of occlusion in the visualization 3D spatial data. Then, we introduced a feature-based **VA** approach, where we reduced several complex 3D spatial datasets to their main characteristics, thus enabling their comparison. Both of these

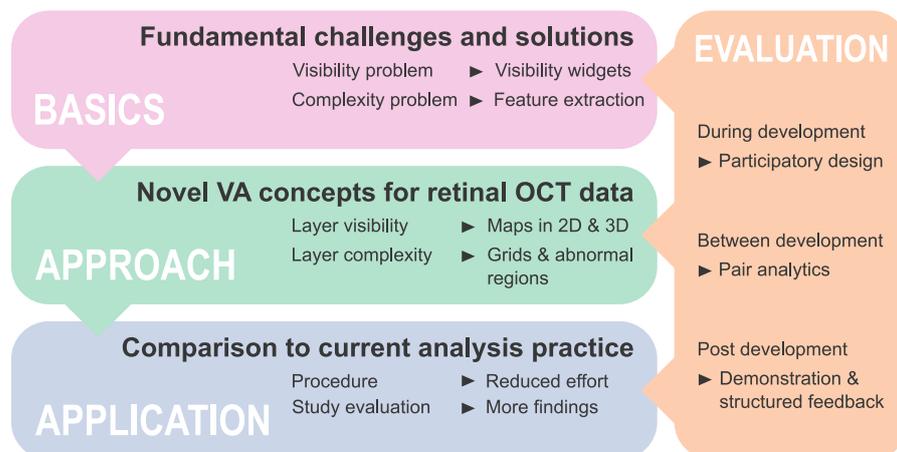


Figure 5.1: Overview of the three thesis parts and the contributions made.

solutions extended the existing approaches to the problems of visibility and complexity.

NOVEL APPROACHES: For the development of our novel **VA** approaches for retinal **OCT** data, we built on the ideas from our basic investigations. In particular, we adapted the concept of extracting high-level features to reduce the complex 3D retinal layers to 2D attribute maps. Besides reducing the data complexity, these maps helped to improve the visibility of the layers in 2D and 3D visualizations of an **OCT** dataset. We then looked at the comparison of retinal layers. For this, we computed features only over local regions in their attribute space and spatial context, which contained relevant information for the comparison of individual patients and patient groups in relation to control data. Achieving both improved visibility and reduced complexity of the retinal layers was previously not possible with existing methods for retinal **OCT** data.

APPLICATION IN OPHTHALMOLOGY: With our developed **VA** methods available, we tested how well they work in practice with actual patient data. To this end, we first introduced a complementary **VA** procedure for single-subject and multi-subject data. The comparison with the current analysis practice demonstrated the lower analysis effort of our methods. We also showed that our solutions support the entire complex process of analyzing data from multiple patients and controls. We then applied our and the current procedure to the data of two cross-sectional studies. The comparison of the results revealed that our methods not only improve the identification of abnormal layer regions, but also enable more accurate measurements of these regions. Overall, more diagnostically relevant insights could be gained with our approach.

Throughout the entire development of all our solutions, we worked together with different domain experts. In order to integrate their expertise into our **VA** designs, we followed a participatory design approach, applied pair analytics, and collected informal and structured feedback. By doing so, we were able to both meet their requirements and achieve meaningful results. This enabled us to publish our approaches not only in the visualization community, but also in the field of ophthalmology.

All the individual approaches of the last three chapters were described in a compact manner. We focused especially on their key aspects and on establishing connections between them with regard to the **VA** retinal **OCT** data. The following **Part ii** of this thesis contains the original publications that provide the details behind the approaches outlined so far. Afterwards, an overall conclusion of the topics covered and an outlook on future work is given in **Part iii**, which starts on page 199.

Part II

THE CORE

This second part contains the original publications on which this work is based. The following chapters supplement the compact descriptions from the first part with additional details. Each chapter begins with a short summary of the scientific and author's contribution, the original abstract, and references to the original and related publications.

VISIBILITY WIDGETS FOR UNVEILING OCCLUDED DATA IN 3D TERRAIN VISUALIZATION

SCIENTIFIC CONTRIBUTION: In this chapter, we contribute novel VA solutions to the occlusion problem in 3D visualizations of spatial data. The introduced concepts extend existing approaches by not only integrating three different design patterns for 2D and 3D presentation of hidden information in a 3D spatial context, but also supporting automated detection and interaction with three different types of occlusion. These concepts formed the basis for improving the visibility of retinal layers in OCT data. The journal article [25] included here is an extended version of [58], which received a best paper award.

AUTHOR CONTRIBUTION: The author of this thesis developed major parts of the presented VA approach, including conceptualization, implementation, performance testing, and evaluation with domain experts, and wrote the main body of both publications [25, 58].

ABSTRACT: In 3D visualization of terrain, occlusion has detrimental impact on discovering, accessing, and spatially relating embedded data. This can lead to significant loss of information. To address this issue, we introduce visibility widgets: compact auxiliary views that visualize occluded data, their locations, and related properties, and thus support users in revealing hidden information. The widgets comprise different occlusion management strategies for detecting and analyzing invisible data inside as well as outside of the field of view. In addition, each widget is designed complementary and transient to facilitate interactive exploration and navigation while not monopolizing the user's attention. We demonstrate the application of visibility widgets focusing the visual exploration of weather data in 3D terrain.

ORIGINAL PUBLICATION: [25] — M. Röhlig, M. Luboschik, and H. Schumann. "Visibility Widgets for Unveiling Occluded Data in 3D Terrain Visualization." In: *Journal of Visual Languages & Computing* 42. Supplement C (2017), pp. 86–98. DOI: [10.1016/j.jvlc.2017.08.008](https://doi.org/10.1016/j.jvlc.2017.08.008)

FURTHER READING: The VA approach is based on our characterization of design options for the visualization of 3D spatial data:

- ▷ S. Dübel, M. Röhlig, H. Schumann, and M. Trapp. "2D and 3D Presentation of Spatial Data: A Systematic Review." In: *Proceedings of the IEEE VIS International Workshop on 3DVis: Does 3D really*

make sense for Data Visualization? 2014. DOI: [10.1109/3DVis.2014.7160094](https://doi.org/10.1109/3DVis.2014.7160094)

The introduced concepts further led to the development of a new prioritized visualization design that ensures the visibility of geo-spatial attributes and uncertainties in 3D terrain:

- ▷ S. Dübel, M. Röhlig, C. Tominski, and H. Schumann. “Visualizing 3D Terrain, Geo-spatial Data, and Uncertainty.” In: *Informatics 4.1* (2017). DOI: [10.3390/informatics4010006](https://doi.org/10.3390/informatics4010006)

6.1 INTRODUCTION

Visually analyzing geo-spatial data in virtual environments is a necessary task in many domains. Prominent examples include wayfinding in digital terrain models [104], assessing photovoltaic potential and residential quality in virtual city models [105], or investigating oceanographic data [43, 60]. Another application domain is avionics where terrain-related weather data are important, e.g., in flight planning. Here, analysts typically have to explore the virtual environments to *discover* and to *access* relevant information, and to *spatially relate* it to the geometry of the frame of reference. In this context, 3D visualization allows for faithful representations of the geo-spatial data and the reference geometry, facilitates orientation via landmarks, and enables analysts to recognize the virtual surroundings in the real world [106]. At the same time, analysis tasks are often severely hampered by occlusion effects inherent to 3D visualizations [54].

In case of digital terrain models, geo-spatial data are often visualized by mapping data values onto the terrain’s surface. This is a common visualization approach, especially for quantitative data, such as meteorological scalar fields for precipitation or wind speed, but also for qualitative data, such as different types of weather zones. It allows users to relate data values directly to topological features. Yet, it also implies that relevant information gets partly or fully occluded by the geometry of the terrain. These problems are most prevalent in near-surface views of the environment. As a result, analysts may have difficulties detecting, identifying, and distinguishing partly occluded data. If analysts are even unaware of the existence of hidden information, e.g., because it is fully occluded by a mountain range, they might lose important information. Particularly in security-critical scenarios, such as analyzing hazardous weather phenomena, this can have major implications.

Several approaches have been developed for managing occlusion. For instance, multi-perspective panoramas have been successfully applied for navigating, wayfinding, and investigating objects in virtual 3D city and landscape models [107]. However, open issues are how to manage occlusion of embedded data, to efficiently compute and visualize the visibility of such information, and in this way, to support exploration

of information-rich geo-virtual environments. To close this gap, we introduce visibility widgets: compact auxiliary views that facilitate revealing and analyzing partly or fully occluded information, inside as well as outside of the field of view. Our contributions are:

DETECTING OCCLUDED DATA: We identify different types of occlusion for embedded data in 3D terrain. We discuss their computation in general and present efficient implementations based on common rendering techniques.

PRESENTING OCCLUDED DATA: We visualize hidden information via novel visual designs instantiated by four complementary visibility widgets. Each widget comprises different occlusion management strategies to (i) overview visible and occluded data, and to (ii) locate and analyze that data in a main 3D visualization.

INTERACTING WITH OCCLUDED DATA: We provide coordinated interaction techniques for (i) selecting and highlighting information in interlinked views, (ii) investigating details on-demand, and (iii) facilitating discovery, access, and spatial relation of occluded data.

In this paper, we present an extended version of our work on visibility widgets [58]. We add details and consider new aspects. Our extensions include:

- *New data:* In addition to quantitative data, we incorporate qualitative data that are mapped onto the terrain surface (Sect. 6.4).
- *New functionality:* We aid navigation by indicating previously unexplored regions in a map widget and enhance the readability of encoded data by showing alternative visual representations in an interactive lens (Sect. 6.4.3). We enhance our approach with regard to multiple occluders (Sect. 6.5).
- *New details:* We provide an explanation of the application background (Sect. 6.2) and a detailed description of our tool's implementation together with performance measurements for different devices (Sect. 6.5).
- *New feedback:* We report on user feedback gathered in a post-development demonstration session together with domain experts (Sect. 6.6).
- *New applications:* We present a solution for visualizing occluded data in the context of isosurfaces (Sect. 6.7).

The remainder of the paper is organized as follows. We start with introducing the application domain and discuss different occlusion-induced problems (Sect. 6.2). With regard to these problems, the related

work is reviewed (Sect. 6.3). Due to open issues, we introduce a novel visual-interactive design (Sect. 6.4) and give details on its prototypical implementation (Sect. 6.5). We report on the gathered feedback from aviation experts (Sect. 6.6) and demonstrate the generalizability of our approach (Sect. 6.7). Finally, we summarize our solution and give an outlook on future work (Sect. 6.8).

6.2 BACKGROUND AND PROBLEM STATEMENT

In this section, we briefly describe the background of our work and identify problems caused by different aspects of occlusion.

6.2.1 Background

In our work, we focus on usage scenarios from the field of avionics. This field strongly depends on three-dimensional topographic data and various related geo-spatial information. For instance, during start and landing of aircrafts, the 3D geometry of terrain and information about weather conditions, e.g., precipitation, wind speeds, or visibility ranges, are of particular interest. In 3D visualizations, geo-spatial information is traditionally encoded by color-mapping it on the 3D surface of the terrain (e.g., [108] or [32]). Yet, such encodings result in information loss, since relevant information may be located on back-facing surface parts, or on front-facing surface parts that are occluded by other parts of the terrain. This leads to the fundamental question: how to communicate data that are *occluded by the terrain*?

Considering the actual flight, crucial weather phenomena, e.g., hazardous weather zones, must be observed, as their avoidance is vital for security. However, such information may be located outside the field of view and consequently, is not visible. This scenario raises a second question: how to communicate data that are *occluded because of a given viewpoint*?

Finally, if certain weather phenomena conjoined with the 3D terrain geometry require re-planning of flight routes, several data must be checked simultaneously. But, data points from different data sources might be related to the same spatial locations, e.g., continuous wind speeds in categorical save zones. This poses a third question: how to communicate data that are *occluded by other data* without losing relevant information?

To address these problems, we introduce visibility widgets that help to communicate relevant information. Our investigations focus on terrains of up to 300 km². The terrain geometry originates from digital elevation models (DEM), which are widely-used in the aviation domain. In addition, we consider data points from gridded quantitative meteorological data, e.g., precipitation, temperatures, wind speeds, or visibility ranges, and region-based qualitative data, e.g., no-fly zones,

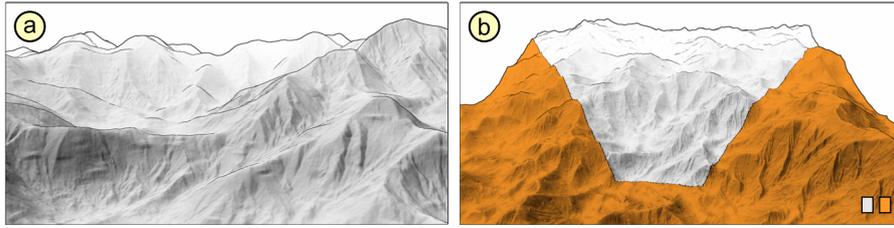


Figure 6.1: Occlusion by viewpoint. The images show terrain from a viewpoint near the terrain’s surface (a) and from a top-down perspective for reference (b). In the close-up view (a), about 60% of the terrain is outside the field of view, as illustrated by the orange colored parts in (b).

hazardous zones, or safe zones. The visibility widgets aim at revealing occluded information for all three problems described above: (i) occlusion by viewpoint, (ii) occlusion by terrain, and (iii) occlusion by data. This differentiation is driven by the problems themselves. For instance, data outside the field of view cannot be displayed *in place* and thus, requires *additional* views. Another reason is the technical implementation of occlusion detection, which differs for each type of occlusion.

6.2.2 Problem: occlusion by viewpoint

Digital terrain models are typically large and cannot be displayed in their entirety due to limitations of available display space. Choosing different viewpoints, i.e., viewing positions and orientations, is therefore necessary to either review the encoded data in overviews or to inspect details in close-up views. Yet, no matter what viewpoint is chosen, none will show all data values at once and thus, occlusion prevails. In this regard, we distinguish between data located (i) outside and (ii) inside the field of view. Data outside the field of view are always invisible, as they are not covered by common rendering methods. Hence, we consider such data as occluded. Data inside the field of view may be visible or invisible in the rendered image, depending on further aspects. [Figure 6.1](#) illustrates occlusion caused by different viewpoints.

6.2.3 Problem: occlusion by terrain

Terrain-related data may be located in valleys behind mountains. Therefore, data often end up occluded by the surrounding parts of the terrain. Dealing with the resulting implications is the primary focus of our work. Interestingly, from a spatial interaction point of view, we are dealing with a singleton environment containing solely the 3D terrain. Thus, only self-occlusion can occur depending on individual geometric features, e.g., mountains and valleys. Other types of spatial interactions, such as intersection, enclosement, and containment (cf. [54]), play a

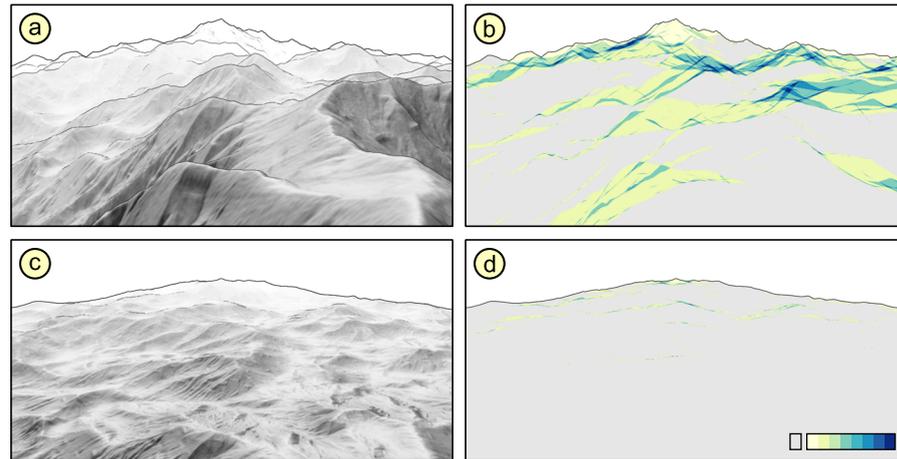


Figure 6.2: Occlusion by terrain. The images show terrains with high mountains (a and b) and a rather flat ground with few small hills (c and d). Occluded parts of the terrain are color-coded with yellow (one layer of occlusion) to blue (more than eight layers of occlusion), and gray for no occlusion (b and d). The shape of the first terrain results in 72% occlusion with a maximum of 19 occlusion layers (a and b), while the second terrain only yields 15% occlusion with a maximum of 16 occlusion layers.

secondary role in our setup. The main influencing factor is the terrain's geometric configuration. With that, we refer to general surface characteristics, including the existence and amount of high mountains or deep valleys and their respective arrangement. Generally, the amount of occluded data directly depends on the geometric structure. For example, detailed terrains with numerous high mountain ranges and steep valleys result in many occlusion-affected viewpoints. In contrast, occlusion of data embedded in rather flat terrain mostly occurs due to back-facing surface parts that can be easily visualized from different viewpoints. [Figure 6.2](#) illustrates different amounts of occlusion for such geometric configurations.

6.2.4 Problem: occlusion by data

Occlusion may also occur due to multiple data values competing for the encoding of a single pixel on screen. Such situations arise because of: (i) data values of different data points that share the same spatial location or (ii) data values of data points with different spatial locations that are mapped onto the same pixel due to perspective projection. In the first case, data points from different data tables cover the same spatial region, e.g., a quantitative data point such as wind speed that is given within a hazardous weather zone. This results in a compelled exclusive data selection, as generally only one data point can be encoded in each pixel. Special encodings, e.g., based on bivariate and trivariate color palettes, may postpone this limitation towards two or three data points.

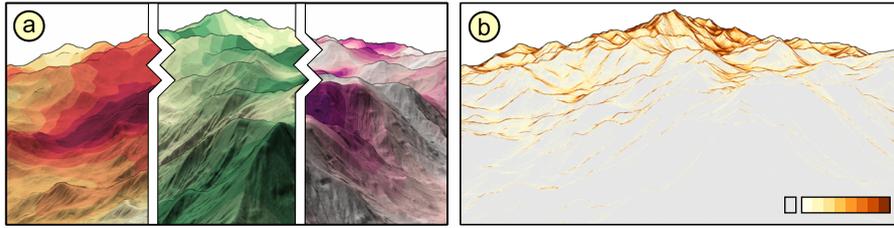


Figure 6.3: Occlusion by data. The images show terrain with three data tables mapped onto the surface (a) and over-plotted parts of the data for a single data table (b). An exclusively shown data table hides all others due to the commonly used color-coding (a). Even with only one data table visible (b), 44% of the data values are affected by over-plotting with a maximum of 114 values per pixel, color-coded by yellow (two values per pixel) to brown (more than eight values per pixel).

In the second case, occlusion by over-plotting particularly impacts data points that are distant to the viewer, and thus are presented in a small screen area. A greater distance of these data points and a denser spatial distribution of them increases the over-plotting.

To avoid adverse occlusion-related effects in both situations, special attention and care is required. Otherwise, common rendering pipelines may produce unexpected results or follow unwanted default behavior, such as displaying nearest samples of the data. As a result, relevant information might get occluded by less important data values. [Figure 6.3](#) shows occlusion by data due to multiple data points at the same location and due to over-plotting in distant regions.

In sum, we are dealing with proximity-related occlusion inside and outside of the field of view, resulting from the geometrical configuration of the terrain, and the density of the embedded data.

6.3 RELATED WORK

Occlusion is one of the main challenges in information-rich virtual environments, such as geo-spatial data embedded in 3D terrain [109]. Various approaches have been proposed to overcome its implications in related 3D visualizations. Elmqvist and Tsigas [54] present a comprehensive survey and taxonomy for occlusion management. They identify five distinct design patterns: tour planners, volumetric probes, multiple viewports, virtual X-ray tools, and projection distorters. Each design pattern comprises a long history of research with different occlusion-aware visualization techniques. Here, we focus on their general properties and outline few relevant examples with regard to our usage scenario and the identified types of occlusion. For a full discussion and profound overview, please refer to [54].

TOUR PLANNERS AND VOLUMETRIC PROBES: The first two patterns are hardly applicable for our purpose, as they either require knowledge and preprocessing of all locations of interest in advance, or strongly affect the spatial reference and other invariances of the embedded data. The last three patterns, on the other hand, are particularly related to our work and will be discussed in greater detail.

MULTIPLE VIEWPORTS: The widely-used multiple viewports pattern provides several alternative views of a 3D environment simultaneously. Typically, one main viewport is used for first-person or detail views, whereas one or more secondary viewports show overviews of the environment. In virtual reality, common examples are worldlets [110] and worlds-in-miniature [59]. These techniques embed miniaturized 3D representations of specific objects or regions of interest in the main 3D visualization. Interacting with these representations allows users to inspect the environment from a vantage point and facilitates later recognition from an arbitrary viewpoint. The GeoZui3D system assists navigation and exploration of ocean data via frame-of-reference interaction [60]. Multiple windows can be created and linked to seamlessly change focus between different locations or to inspect hidden parts of an underwater environment.

Navigating in a large virtual city model is addressed by Fukatsu et al. [111]. An overhead map helps users in maintaining orientation while walking or driving through streets under limited visibility conditions. Such techniques are also powerful means for discovering and accessing occluded information in terrain visualization, as terrain is essentially 2D in nature and lends itself to overhead maps. Multiple viewports allow to address both, occlusion by viewpoint and by terrain. Yet, none of the existing terrain-specific approaches allows for identifying occlusion by data. Such issues have been addressed in the information visualization domain (e.g., [49]) with some approaches in principle being applicable to terrain-related data (e.g., small multiples [112]). Also, supporting 3D navigation through secondary views has only been sparsely addressed so far. The general concept of providing different views on the terrain and its related data is the basic idea of our widget solution presented in [Section 6.4.2](#).

VIRTUAL X-RAY TOOLS: The virtual X-ray pattern helps to discover and access hidden information by reducing the opacity of occluding objects. Generally, two types of approaches exist: *active* tools function as a user-controlled searchlight [62], and *passive* techniques automatically uncover invisible objects based on semantic information [61]. Especially active interaction allows for selectively removing occluding objects and thus, facilitates exploration. Viola et al. [113] combine virtual X-ray with stylization and smooth animation for emphasizing selected features in medical volume data. Managing occlusion in massive 3D point clouds

is addressed by Discher et al. [114]. An interactive X-ray lens allows the user to inspect buildings or streets that would otherwise be covered by vegetation. Albeit virtual X-ray tools are appropriate for handling occlusion by terrain, applying them can have a negative impact on depth cues, which might impair depth perception and can hinder spatial relation. Moreover, using virtual X-ray techniques for occlusion of quantitative data and qualitative data has not been thoroughly addressed so far. In this case, special care is required if multiple semitransparent objects are to be displayed. For instance, in terrain visualizations with high depth complexity, e.g., due to consecutive mountains, unintended mixing of colors might impair the interpretation of color-coded data. Again, techniques from other visualization domains (e.g., magic lenses [115]) might be adapted to X-ray tools to handle terrain-specific data occlusion. In Section 6.4.2.2, we introduce an interactive lens that is inspired by X-ray tools.

PROJECTION DISTORTERS: This pattern integrates two or more different views of an environment into a single view, typically via nonlinear projection. The basic idea originated from multi-perspective panoramas [63] and artistic multi-projection rendering [116], and has been subsequently applied for static and animated geographical visualizations [104, 117]. The individual base views can often be actively selected by the user, which is particularly relevant in scenarios that focus on exploration. For example, multi-perspective 3D panoramas allow for switching between degressive, progressive, and hybrid perspectives [64, 107, 118], or applying advanced camera models to circumvent complex occluders [119]. Each choice supports a different navigation task, reduces occlusion, and improves utilization of screen space. Automatically determining the visibility of objects for avoiding occlusion has been studied for different application scenarios. Examples are urban context-aware visualizations [106, 120–122] and disocclusion of driving routes in rural environments [104, 123, 124]. Based on the computed results, local projection techniques are applied to generate views that ensure visibility and that introduce minimal distortion in the spatial surroundings. Pasewaldt et al. [125] utilize a multi-perspective overview+detail visualization for exploring 3D buildings. A secondary panoramic overview shows all facades of a building and facilitates navigation within a main 3D perspective view. In general, projection distorters are mostly used for discovery, rarely for access, and almost never for spatial relation [54]. Depending on the applied distortion, not only occlusion by terrain but also occlusion by viewpoint can be managed. Occlusion by data is not addressed, as distortion alone generally does not solve the overlap of multiple data points at the same spatial location. A distortion-inspired panoramic widget is part of our solution (Sect. 6.4.2.2).

	Occlusion by		
	Viewpoint	Terrain	Data
Tour planners		●	
Volumetric probes		●	
Multiple viewports	●	●	○
Virtual X-ray		●	○
Projection distorters	●	●	

Table 6.1: Summary of related work. The table depicts the five design patterns for occlusion management [54] with regard to the occlusion problems. A cell is marked based on the support of the design pattern (full: ●, partial / potential: ○) for a certain problem.

OTHER APPROACHES: Several other approaches for supporting exploration and navigation in 3D visualizations have been proposed. For instance, additional visual cues can be added for hinting at points of interest outside of the field of view [126]. Other examples include highlighting specific landmarks or modifying their appearance to aid orientation, e.g., by applying geometric scaling [127] or semantic levels-of-abstraction [128]. On top of that, 3D visibility analysis has been applied to assess the visual impact of occluders with respect to a large number of different viewpoints in the context of urban planning [129].

In summary, the five design patterns [54] offer different approaches for managing occlusion. Still, each technique covers only a certain aspect of occluded data in 3D terrain and does not match all problems in combination (Tab. 6.1). Although multiple views represent a potential solution, a comprehensive design for terrain-related data has not been introduced so far. Our goal is to develop an integrated approach that combines the strengths of the individual solutions and thus, brings them to their full potential. To achieve this goal, we (i) introduce tailored concepts for capturing different types of occlusion, (ii) incorporate multiple occlusion management strategies, (iii) visualize occluded data via coordinated widgets, and (iv) augment each widget with appropriate interaction techniques to ease navigation and to allow users to inspect details. With the resulting flexibility, we are able to tackle the peculiarities of embedded geo-spatial data and to assist users in exploring 3D terrain visualizations.

6.4 VISIBILITY WIDGETS

We aim at supporting users in three main tasks for visually exploring virtual environments: *discovery*, *access*, and *spatial relation* of hidden information. To this end, we address three fundamental questions related

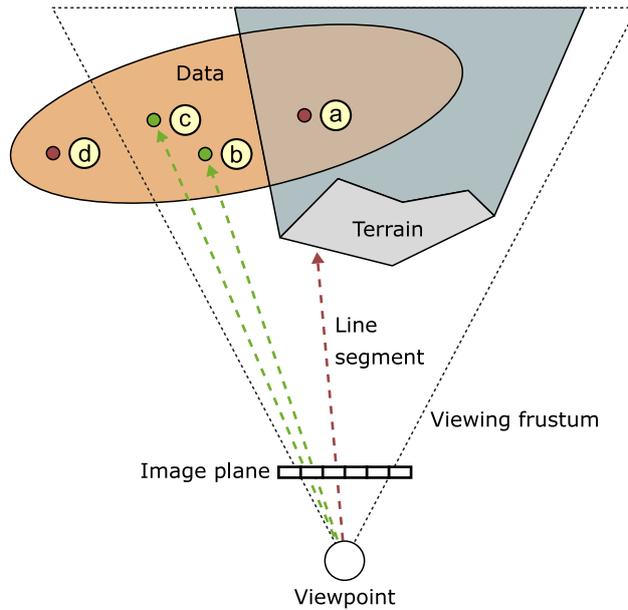


Figure 6.4: Basic model for occlusion in 3D terrain (adapted from [54]). From the viewpoint the data points (a), (b), and (c) are located inside the viewing frustum, while data point (d) is located outside the viewing frustum. Data point (a) is occluded by terrain, whereas data points (b) and (c) are conditionally visible since they contribute to the same image pixel. Whether or not data points (b) and (c) are actually visible depends on how occlusion by data is resolved.

to occlusion in 3D terrain: (i) how to detect occluded data, (ii) how to present occluded data, and (iii) how to interact with occluded data.

6.4.1 Detecting occluded data

Before we introduce our presentation and interaction with occluded data, we describe a basic model for occlusion in 3D terrain and explain our conceptual approach for occlusion detection on modern graphics hardware.

OCCUSION MODEL: We apply a general pipeline-based computation for detecting the different types of occlusion. Following the basic model for 3D occlusion in [54], we consider data of the following form (Fig. 6.4). A terrain T is represented by its surface in Cartesian space $(x, y, z) \in \mathbb{R}^3$. A data table $D = \{d_1, \dots, d_n\}$ consists of n data points. Each data point d_k with $1 \leq k \leq n$ stores one data value and a location on the surface of T . In our case, multiple data tables are available describing different types of data, e.g., wind speeds or weather zones. A viewpoint $v = (M, P)$ is given by a view matrix M and a projection matrix P . Both matrices describe the position and orientation of the viewing frustum v_f of a virtual camera. A data point d_k is considered

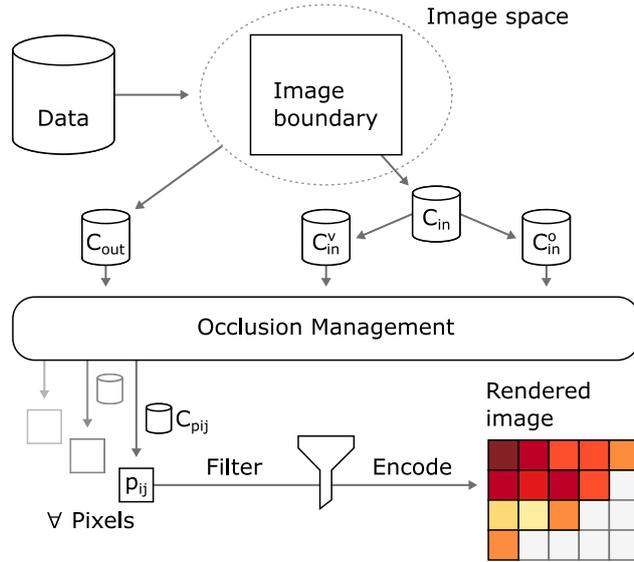


Figure 6.5: Overview of our computational occlusion detection. First all data values are projected into image space. The values are then subdivided into *inside* (C_{in}) and *outside* (C_{out}) with respect to the image boundaries. Data values inside the image are further classified into *conditionally visible* (C_{in}^v) or *occluded* (C_{in}^o). After applying occlusion management, the remaining values ($C_{p_{ij}}$) are assigned to pixels (p_{ij}), filtered in case of occlusion by data, and finally encoded in the respective views.

inside the field of view if its location is contained in v_f , and *outside the field of view* otherwise. In addition, d_k is treated as *conditionally visible* from a viewpoint v if there exists a line segment between v and d_k that does not intersect any part of T , and *occluded* if no such line segment exists. Whether d_k is finally classified as *visible* or as *occluded* depends on whether or not other data points contribute to the same pixel in the rasterization stage of the rendering process.

COMPUTATIONAL OCCLUSION DETECTION: Figure 6.5 provides a conceptual overview of our computational occlusion detection. The procedure consists of four primary steps. Implementation details on how we compute each step using a combination of customized rendering techniques are outlined in Section 6.5.

- (S_1) *Data reduction*: We generate a reduced data set D_r by selecting only data points that are of particular interest. All other data points that are not currently considered by the user are filtered out accordingly.
- (S_2) *Detection of occlusion by viewpoint*: We project D_r into image space according to M and P of the viewpoint v . As a result, we get a set C of potentially contributing data points for a target image I .

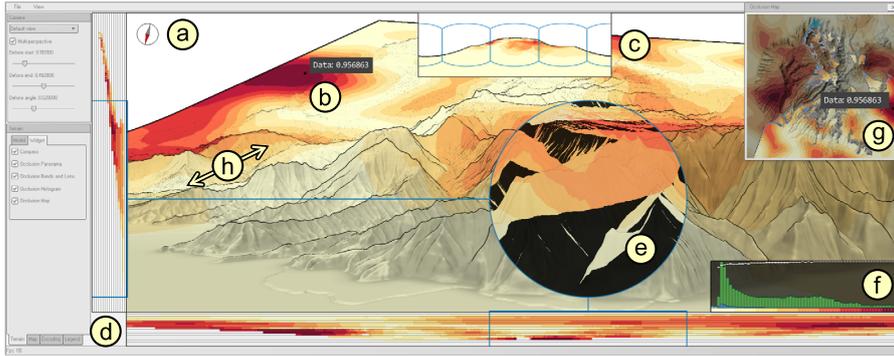


Figure 6.6: Overview of our visualization tool. The main 3D view shows color-coded data embedded in terrain (middle). The user interface supports loading different terrain models, adjusting the visual encoding, and switching between multiple data tables (left). An interface skin with light colors helps to focus on the main visualization. Our tool integrates four widgets and a rich set of coordinated interactions for showing hidden information and details on demand: (a) compass, (b) tooltips, (c) panorama view, (d) bands view, (e) interactive lens, (f) histogram view, (g) 2D overview map, and (h) multi-perspective 3D view.

According to the boundaries of I , we subdivide C into C_{in} and C_{out} , containing data points inside or outside of the image, respectively.

- (S₃) *Detection of occlusion by terrain:* We sort C_{in} by distance to v and classify the data points into conditionally visible C_{in}^v or occluded C_{in}^o with regard to the surface of T . Altogether, we obtain four subsets: for inside C_{in} and outside C_{out} in general, and for conditionally visible C_{in}^v or occluded C_{in}^o within the image.
- (S₄) *Detection of occlusion by data:* To determine which data points are to be considered in each view, we forward the four subsets to the different occlusion management strategies instantiated by our widgets (Sect. 6.4.2). In return, we get a set of data points $C_{p_{ij}}$ for every pixel p_{ij} in $I = \{p_{11}, \dots, p_{wh}\}$ with $1 \leq i \leq w$ and $1 \leq j \leq h$. If $C_{p_{ij}}$ contains only a single data point, it is directly encoded in the respective view. In case multiple data points remain, i.e., occlusion by data occurs, additional user-specified filter mechanisms may be applied. For example, the final value for the encoding in each pixel can be chosen via a maximum filter for quantitative data or via a majority filter for qualitative data.

6.4.2 Presenting occluded data

Additional views that show otherwise hidden information have been proven to be very useful [54]. Our approach is based on additional views as well. We design four coordinated visibility widgets to communicate the detected occlusion of data: horizontal and vertical bands, a

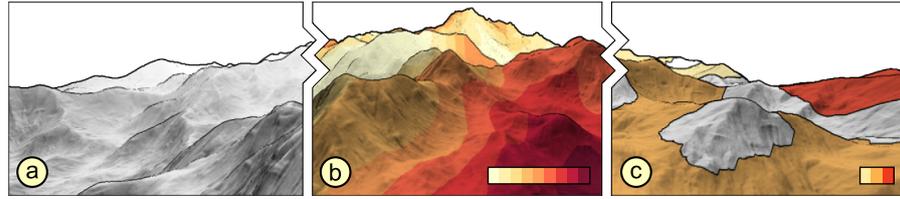


Figure 6.7: Visualizing terrain and data. The geometry of the terrain is encoded via directional lighting, ambient aperture lighting, and enhanced edges (a). The data are directly mapped onto the terrain’s surface via color-coding (b and c). For quantitative data, low values are shown in light yellow and high values in dark red (b). Qualitative data are represented by distinct colors, e.g., three types of weather zones via yellow, orange, and red (c).

panorama, a histogram, and a map. Each widget incorporates a different occlusion management strategy and consists of a dedicated auxiliary view, which can be activated on demand. Figure 6.6 shows an overview of our visual design for data embedded in 3D terrain. Next, we describe the basic terrain visualization and then we explain the visibility widgets in detail.

6.4.2.1 Basic encoding of terrain and data

ENCODING TERRAIN: We apply edge enhancement and various shading techniques for visualizing terrain. Our primary shading aims at illustrating surface properties. For this purpose, we consider two illumination methods: a simple local technique based on a directional lighting model and a more sophisticated global approach based on ambient aperture lighting [130]. Directional lighting depicts the general shape of the terrain’s surface, whereas ambient aperture lighting highlights details and relationships between neighboring parts. Additional stylization allows us to further emphasize certain terrain properties. On the one hand, we utilize a customized toon shading technique (cf. [131]) to create various view-dependent effects, including levels-of-abstraction, aerial perspective, or depth-of-field. For instance, modifying the influence of illumination depending on the distance to the viewpoint results in a continuous abstraction that also reinforces perception of depth. On the other hand, enhancing edges makes it easier to distinguish individual parts. For example, different edge styles allow for tracing ridges or providing additional shape detail. The stylizations can be interactively adjusted or completely deactivated. The user can switch between effects, set the amount of abstraction being applied, and regulate which edges are to be shown. Combining the different encoding techniques allows for effectively communicating inherent geometric features of terrain. Figure 6.7 shows an example of our visual design.

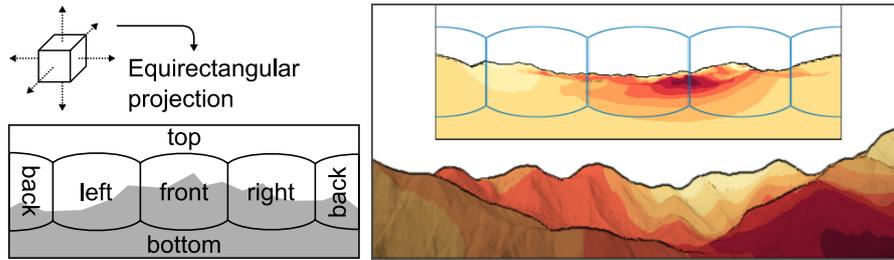


Figure 6.8: Visibility panorama. The diagram illustrates the generation of the 360° overview via environment mapping and equirectangular projection (left). Activating the widget allows for inspecting relevant information in the immediate vicinity of the viewer (right).

ENCODING DATA: We select color for mapping data directly onto the terrain's surface. For encoding visible values ($C_{p_{ij}}$), we apply quantitative or qualitative color palettes from ColorBrewer [132]. Data values are assigned to a limited set of colors per palette to support their identification in an associated legend (Fig. 6.7). By stylizing the terrain, the communication of data is further facilitated. For example, attenuating distant shading allows for interpreting faraway data more easily, while providing nearby surface details enables to relate data and geometric details. Moreover, depicting the terrain solely via gray scales ensures sufficient visual contrast to the color-based encoding of the data. If required, the default color-coding can be interactively adjusted.

6.4.2.2 Design of visibility widgets

VISIBILITY PANORAMA: To manage occlusion by viewpoint, we integrate a multi-perspective panorama (cf. [63]) as a widget into the main 3D visualization. Generating the panorama is based on a customized environment mapping technique. With regard to the main virtual camera, we render the scene in all six viewing directions, i.e., front, back, left, right, up, and down. All rendered images are then joined into a panoramic map via equirectangular projection. This projection method allows for a compact depiction with only low horizontal and varying vertical distortion. To support the interpretation of the panorama, borders of each viewing direction can be superimposed. Figure 6.8 illustrates the resulting display. Activating this widget provides a 360° overview of the terrain. Information outside the main view can be identified without changing the camera's orientation. At the same time, information inside the main view can be related to the surroundings. Focusing on the data is facilitated by abstracting from terrain details in the panorama, e.g., by showing only main slope lines. Moreover, the widget can be set to show only visible, only occluded, or all information regardless of occlusion. We employ rendering filters for showing all information or selected value ranges, which, however, come at the cost of reduced depth perception (Sect. 6.4.1). This way, direct access to

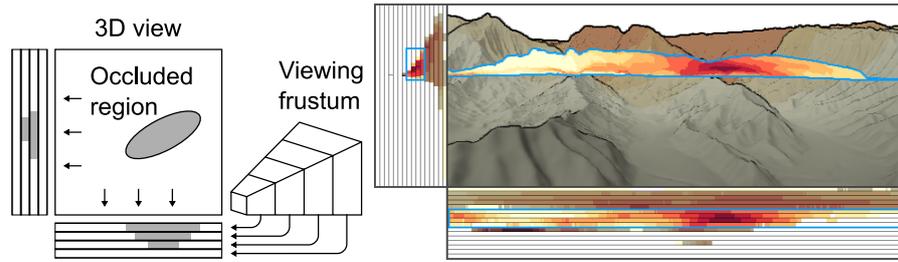


Figure 6.9: Visibility bands. The diagram illustrates how occluded information inside the camera's viewing frustum is successively projected onto both image axes (left). The widget visualizes the projected information via horizontal and vertical bands, each with 16 layers ordered by distance to the viewer (right). When selecting one or multiple layers related occluded data are superimposed in the 3D view.

visible and occluded information in the immediate vicinity of the user is granted.

VISIBILITY BANDS: To address occlusion by terrain, we introduce horizontal and vertical visibility bands with an indirect reference to space. First, we determine all occluded data values inside the viewing frustum of the camera (C_{in}^o). These values are then projected into 2D image space and finally horizontally and vertically onto two 1D value sets. Both value sets are visualized as visibility bands that depict the distribution of occluded information for each image axis. Aligning the bands along the border of the main view, allows users to overview the encoded values and to relate them to the rendered image of the 3D scene. To give the user a better idea at what distance occluded data are located, we additionally subdivide the viewing frustum into equal depth ranges. The occluded values in each depth range are assigned to respective layers in each band. The layer with the closest depth range is placed next to the image border and the remaining layers are arranged by distance to the viewpoint. The number of layers and the total covered depth range can be set by the user. This way, the 3D location can be approximated, with only a marginal increase in required display space. [Figure 6.9](#) shows our design of visibility bands.

VISIBILITY HISTOGRAM: With the histograms we primarily target occlusion by data. The fundamental idea is to abstract from space and take a more data-oriented approach for analyzing hidden information. We achieve this by setting values that are visible in the main view ($C_{p_{ij}}$) in relation to all values located in the camera's viewing frustum (C_{in}). Binning these values allows us to focus on the data distribution. In addition, we compute the ratio of the visible parts and the total amount of values for each bin. These ratios represent the percentage of occlusion for the respective value intervals. For quantitative data, the number

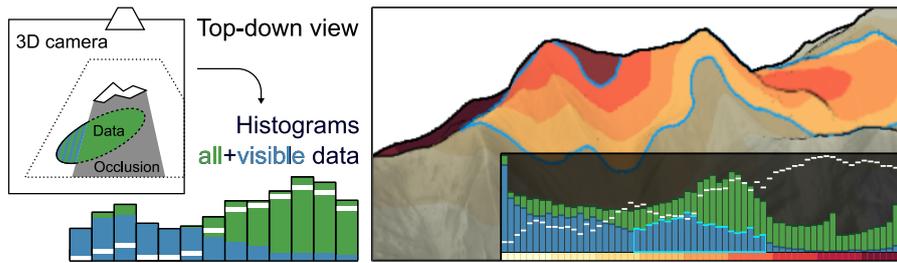


Figure 6.10: Visibility histogram. The diagram summarizes the computation of the histograms (left). The widget shows the relationship between distributions of visible values (blue) and all values (green) inside the current view (right). A white marker per bin indicates the percentage of occlusion in the respective value intervals. A narrow color legend helps associating bins and data values (bottom). When selecting one or multiple bins the respective data values are highlighted in the 3D view.

of bins can be adjusted by the user, whereas for qualitative data, the number of bins matches the count of distinct values by default. Figure 6.10 shows the computation of the histograms and their combined visualization with the ratios as a widget. The main advantage of this widget is the compact representation of all information in the current viewing direction. Looking at the histograms allows users to identify, if all values of interest are actually visible in a chosen view of the scene. In case a specific value interval exhibits occlusion beyond a tolerable threshold, more suitable views can be determined in concert with the other widgets. Utilizing the visibility histograms this way allows for verifying hypotheses and compensates for the inherent abstraction of space.

VISIBILITY MAP: This widget has been designed to address occlusion by viewpoint as well as occlusion by terrain. Integrating a secondary 2D map view can be an effective support for exploring 3D terrain visualizations. The map shows the terrain and the related data from a top-down perspective via an orthogonal projection. This generates an overview of terrain and data, aids navigation, and thus, handles both occlusion problems. Illumination of terrain and encoding of data are applied similar to the main 3D view. Yet, the direction of the lighting is fixed to a north-west incidence angle by default. This design follows general shading guidelines for cartographic maps and helps to differentiate craters and mountains [133]. Furthermore, the coloring can be switched to encode elevation of the terrain. In this case, an alternative encoding for data is applied, e.g., isolines. To further increase the maps utility, we augment the visualization with spatial occlusion information from the main view. The hidden parts of the 3D scene are determined and transferred via a customized shadow mapping technique. Activating this functionality, slightly desaturates and darkens respective parts

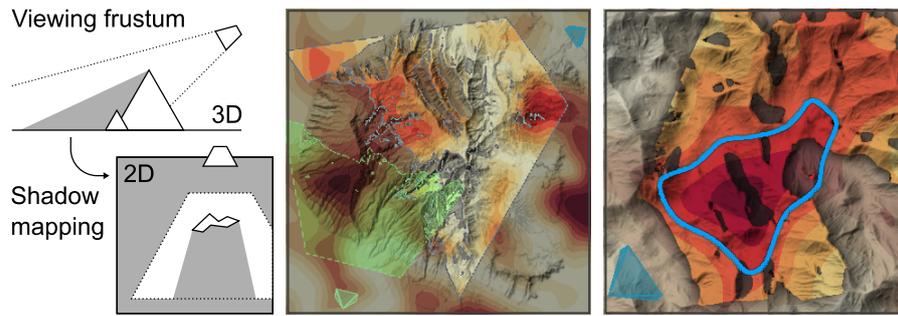


Figure 6.11: Visibility map. The diagram illustrates how occlusion in the main 3D view is transferred to the 2D map view via customized shadow mapping (left). The visibility information is visualized by highlighting parts in sight and by desaturating and darkening parts that are obscured or outside of the field of view (middle). Miniaturized viewing frustums indicate the current viewpoint of the main 3D camera (blue), and an interactively specified destination (green). When selecting spatial regions respective parts of the map are outlined, which allows users to check their visibility (right).

of the map. Additionally, a miniaturized viewing frustum depicts the location and orientation of the main camera. [Figure 6.11](#) shows the final design as a widget.

Individually, the four auxiliary visualizations implement powerful occlusion management strategies. Their real benefit comes into effect when being applied in concert. For example, the map does not directly allow users to analyze occluded information with respect to the terrain's geometry. But, activating the panorama or bands supports locating that information in the main 3D view. Moreover, analyzing multiple data aspects is possible by assigning different data tables to each widget. This helps to focus on one particular aspect in the main view, e.g., visibility ranges, while simultaneously keeping track of other data aspects in the widgets, e.g., no-fly zones, during exploration. As explained next, coordinated interaction further assists in relating encoded information in each widget and in the main 3D visualization.

6.4.3 *Interacting with occluded data*

To support a comprehensive visual exploration, we provide a rich set of tailored interaction techniques. While we focus on communicating occlusion with our widgets, coordinated interaction is essential for discovery, access, and spatial relation. Besides widget-related modifications, e.g., activating and resizing, we particularly support navigation as well as selecting and highlighting certain data. Inspecting details on demand and enhancing the readability of selected regions of interest is further assisted via an interactive lens.

NAVIGATION: To freely navigate through the 3D terrain, the virtual camera of the main view can be interactively controlled to take different points of view on the visualized data. Additionally, predefined camera settings can be applied, e.g., bird's eye view or landmark views. Every movement of the camera is smoothly animated in order to prevent sudden changes in the rendered image. A compass can be superimposed as an additional visual cue for maintaining orientation (Fig. 6.6 (a)). In the 2D map view, we support navigation either independent of or dependent on the main camera. The first option consists of interactive zooming and panning to adjust the visible section of the terrain in the map. For the second option, the map is coupled with the 3D view to keep it centered at the current position of the main camera. All other visualizations are linked with the main 3D view and are automatically updated according to user input.

Target-oriented navigation through the 3D terrain is provided via two additional features in the panorama widget and in the map widget. First, clicking on a point of interest in the panorama centers the 3D view at the selected location via a continuous rotation of the main camera. This particularly eases inspecting information outside of the field of view, e.g., next to or below the viewer, without having to search for suitable viewpoints manually. Second, the map allows users to interactively specify a destination, i.e., a target position and orientation, for the camera of the main 3D view. Activating this feature shows a second miniaturized viewing frustum and highlights all parts of the terrain that are in sight at the currently selected destination (Fig. 6.11). This ensures that the selected parts are actually visible once arrived. Confirming the destination automatically moves the main camera along a collision-free path. This way, even remote locations can be conveniently explored without having to navigate through the 3D terrain by hand and repeatedly checking the map if the correct viewpoint has already been reached.

For a visually assisted navigation, users can switch the coloring of the map to a visibility history. The idea behind the history is to support users in identifying regions which have been visited already as well as regions which are still unexplored due to occlusion. Switching to this coloring mode initially only marks regions in the map that are currently visible in the main view. As other regions become visible during navigation, they are marked in the map as well until all data has been displayed at least once (Fig. 6.14). By utilizing this map feature, users are additionally supported during exploration, whether they use free navigation or target-oriented navigation.

Altogether, our navigation techniques aid discovering occluded data, and thus are particularly helpful in managing occlusion by viewpoint and by terrain.

SELECTING AND HIGHLIGHTING: We enhance the visualization of the interlinked widgets with flexible selection and highlighting facilities. We support selecting individual points in all views and show their assigned values via tooltips. To help associating selected points of the 3D view with the 2D map view, we transform the current selection from one view to the other and duplicate the cursor plus tooltip at the respective location (Fig. 6.6 (b)).

Besides individual points, spatial regions can be selected as well. In the main view and the map view, we support this directly by: (i) polygon selection and (ii) value range selection. The polygon selection allows users to interactively set multiple corners in the terrain to define a spatial region of interest. The value range selection requires only a single seed point as input and automatically forms the respective region by connecting all neighboring parts within a certain threshold. Selecting spatial regions this way, highlights related parts in the 3D terrain and in the 2D map (Fig. 6.11). The utility of these features is twofold: (i) selections of the map can be related to the 3D geometry of the terrain and (ii) selections of the 3D terrain can be inspected with regard to their surroundings.

Further specialized approaches for inspecting hidden information are instantiated by the bands widget and the histogram widget. The bands allow for selecting one or multiple layers to highlight the associated parts in the terrain. This helps users to judge which regions are covered by the layers. Picking individual sections of the bands is enabled either by clicking on colored segments or by using a selection rectangle. Starting this selection in the horizontal band automatically determines and highlights all associated sections in the vertical band and vice versa for clarifying their correspondence. Figure 6.9 illustrates selecting and highlighting occluded information via the bands widget.

The histogram widget allows for selecting single or multiple bins either from the visible distribution, the occluded distribution, or both distributions in combination. All data values that contribute to the chosen bins are highlighted in the terrain. If visible and occluded bins are selected, a user-specified filter is applied to decide which value to encode in each pixel of the display (Sect. 6.4.1). In addition, a modified color-coding ensures that occluded parts can be distinguished from visible parts. Figure 6.10 shows this functionality.

All selections can be expanded or reduced using binary operations. While selecting, the 3D view and the 2D map view are updated to outline selected parts and dim non-selected parts. Essentially, this functionality allows users to relate the abstracted information encoded in the widgets back to space and to reconstruct its spatial context to a certain extent. In sum, our selection and highlighting techniques enable users to access and associate data shown in the interlinked views as well as to spatially relate them with the geometry of the terrain. Hence, these

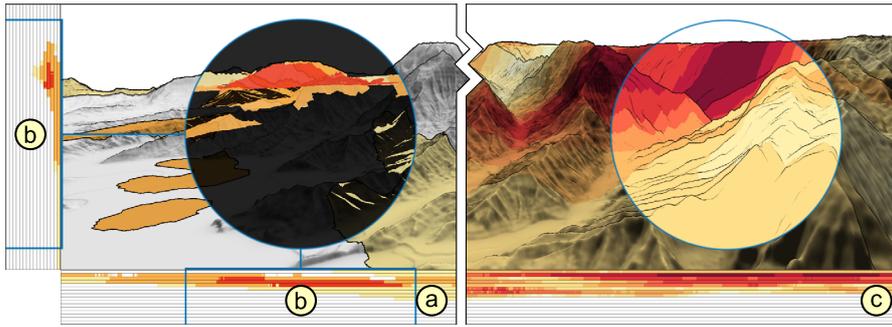


Figure 6.12: Interactive lens. The X-ray mode of the lens reveals qualitative data occluded by surface parts of the terrain (a). Connecting lines help to relate the lens content to patterns in the bands widget (b). The terrain stylization mode of the lens enhances the readability of quantitative data (c). Locally depicting the terrain with emphasized edges only allows to easily read individual data values in an otherwise shadowed terrain region.

features particularly address the inspection of values and occlusion by data.

DETAILS ON DEMAND: Inspecting specific image regions is supported by an *interactive lens* [115]. The lens offers two alternative visualization modes in the main view: (i) an X-ray representation to reveal occlusion by terrain and (ii) a stylized terrain representation to enhance the readability of visible data. In both modes, the position and size of the lens can be interactively adjusted to investigate selected parts in the rendered images.

Activating the lens in X-ray mode allows users to locally investigate hidden information by superimposing and highlighting occluded surface parts of the terrain. If the bands widget is active, the relationship between the bands and lens is illustrated by connecting lines. Optionally, the lens can be configured to show occlusion of one layer, of multiple layers, or aggregated across all layers of the bands. This permits users to examine detected patterns in the bands as well as in the main view.

Activating the lens in terrain stylization mode locally switches the encoding of terrain either to attenuate the influence of illumination or to show enhanced surface edges only. This allows users to clearly identify individual data values as well as transitions between them even in areas otherwise covered by dark shadows. At the same time, the stylized depiction of the terrain provides enough spatial detail to relate data values to the terrain's surface inside the lens as well as to surrounding image areas outside the lens. Figure 6.12 illustrates both lens modes in action and Figure 6.6 (e) shows the lens integrated in our prototype.

In sum, the interactive lens supports the exploration of visible as well as occluded data by locally adapting the visual representation. On

the one hand, users can inspect hidden data even in case of multiple consecutive layers of occlusion. On the other hand, users are able to resolve potential conflicts between the encoding of visible values and shadows induced by the terrain encoding. Hence, the two lens modes particularly assist in managing occlusion by terrain and enhancing the readability of data.

Augmenting our detection and presentation of occluded data with this rich set of interaction techniques completes our visual-interactive design. Next, we give details on the implementation of our visibility widgets and discuss their utility and generalizability for exploring data embedded in 3D terrain.

6.5 IMPLEMENTATION

Our visualization tool consists of a flexible software system for detecting, presenting, and interacting with occluded data. During the participatory development with experts from avionics we conducted consecutive performance tests to enhance the utility of our design.

ARCHITECTURE: We implemented our tool with the requirements for real-time visualization in mind. This requires a sophisticated software architecture. The back end features a dedicated resource management that ensures efficient data processing and rendering. The front end supports stylizing the visualization and interactive exploration of terrain and data.

The resource management handles different data formats. Regarding terrain, our tool supports high-resolution heightmaps, which can cover large spatial regions. For instance, we used terrain models of the Hawaiian islands as well as of central Europe covering areas of up to 300 km² at a resolution of 1/3'' for testing. Regarding embedded data, we mainly focus on two common data types: gridded quantitative data and region-based qualitative data. In the first case, we consider quantitative meteorological data given on a 2D grid, e.g., to represent precipitation, temperature, or wind speed. In the second case, we incorporate qualitative weather zones given as polygonal shapes, e.g., to represent safe zones, potentially safe zones, hazardous zones, and no-go zones. In addition, our resource management also includes the distinction between visible and occluded parts of the data.

The visual encoding is based on a scripting engine. Scripting enables us to include various visualization techniques and maintain associated parameters. Each script specifies how a resource is to be encoded visually. To ensure real-time display of terrain, visible data, and hidden information, the backing algorithms are implemented as a mix of OpenGL shaders and OpenCL kernels for hardware-accelerated execution on the GPU and CPU. Especially, processing and rendering large terrain models can be computational expensive. Hence, our tool includes so-

phisticated level-of-detail techniques to strike a balance between image quality and achieved frame rate. Two techniques are supported: a direct method based on adaptive subdivision via OpenGL tessellation and control shaders, and a customized chunked level-of-detail method (cf. [134]) that ensures sub-pixel precision but requires preprocessing of the terrain data.

Encapsulating and managing the memory, the computational resources, and the visual resources as described has two main advantages. First, it enables us to put each resources state under version control. This allows users to interactively adapt the visualization by modifying the color-coding of the data, choosing from different stylizations, or activating and arranging the individual widgets. Second, it lets us handle resource updates automatically. For example, if new data are available at an external source, the visualization is refreshed to reflect the incoming changes. With regard to the four widgets, occlusion computations are only triggered if currently required, e.g., in case a respective widget is activated. Altogether, our software architecture enables us to rapidly prototype various visualizations for occluded data and to test them with different terrain models and data.

OCCCLUSION COMPUTATION: Specialized and efficient hardware-assisted algorithms for detecting visibility and scalable rendering have been introduced in the past (e.g., [135, 136]). Here, we outline our approach for computational occlusion detection in 3D terrain. Instead of introducing new algorithms, we integrate a combination of common rendering techniques in the back end of our prototypical software architecture. Combining these techniques and applying just few customizations enables us to compute different types of occlusion (cf. Sect. 6.4.1), while being general enough to support a broad range of computing devices and respective programming interface, e.g., OpenGL and OpenCL.

For multiple data tables, we encode all data values per considered table in individual textures. This allows us to switch between them instantaneously. If possible, we reduce overall computational costs by performing calculations in image space, following a deferred rendering strategy. Data points inside (C_{in}) and outside (C_{out}) of the field of view are processed via dynamic environment mapping. To determine visible (C_{in}^v) and occluded (C_{in}^o) data points for each pixel in an image, we consider single-pass techniques, e.g., K-buffer implementations or layered rendering, on newer graphics hardware as well as multi-pass approaches, e.g., depth peeling, on older graphics hardware. Generating specialized mipmaps and integrating respective sampling techniques allows us to realize different types of filters for managing occlusion by data, i.e., in case of multiple data points per pixel. If required, more sophisticated aggregations, e.g., histograms, and summary statistics can be computed in parallel on the GPU (e.g., using scattering-based approaches [137] or, if supported, using atomic counters and compute

Device	FPS	Terrain	Panorama	Bands	Histogram	Map
Desktop (GTX 980 Ti)	170 (300)	0.2 ms	0.11 ms	1.37 ms	0.85 ms	0.22 ms
Laptop (GTX 670M)	24 (35)	8.6 ms	1.08 ms	18.21 ms	2.5 ms	0.91 ms

Table 6.2: Performance of our prototypical visualization tool. The table summarizes the tools refresh rates in frames per second (FPS), the times required for visualizing the terrain and data, and the computation times for each widget for a desktop device and a laptop device. For comparison, FPS are shown for all widgets activated and without the widgets in brackets.

shaders). Additionally, calculated occlusion from different viewpoints, e.g., 3D views and 2D views, is transferred via customized shadow mapping. The synergy of these diverse techniques allows us to perform occlusion detection in real-time. That is, if necessary, all calculations are executed each frame to facilitate interactive exploration.

PERFORMANCE With our prototypical implementation we are able to visualize digital terrain models, e.g., 16 000 px × 16 000 px heightmaps, at interactive frame rates. For occlusion detection, we employ deferred rendering strategies wherever applicable. Hence, the associated computational effort mainly depends on the resolution of the final rendered image and on the size of intermediate buffers used for the detection. [Table 6.2](#) summarizes several performance measurements for a desktop device and a laptop device with different graphics hardware. The shown computation times were taken while navigating through a 3D terrain (4000 px × 4000 px heightmap, cf. [Fig. 6.6](#)), displayed with full HD resolution (1920 px × 1080 px), and averaged over 10 000 frames. On the desktop device we achieved around 170 frames per second (FPS) with all widgets activated (without 300 FPS) and on the laptop device around 24 FPS (without 35 FPS). The time required for visualizing just terrain and data was 0.2 ms (desktop) and 8.6 ms (laptop) respectively. For the individual widgets the timing ranges from 0.2 ms to 1.4 ms for the desktop and from 1.0 ms to 18 ms on the laptop. In all measurements the bands widget required the most computational effort. In summary, our performance testing shows that our occlusion detection and interactive visualization via common rendering techniques can be utilized even on older mobile graphics hardware, e.g., NVIDIA GTX 670M. To improve performance, the occlusion detection can also be performed on downscaled intermediate buffers.

6.6 USER FEEDBACK

We base our approach on well-established and well-evaluated concepts in the field of terrain visualization. Hence, we adapt these working concepts without any further evaluation. Instead, we developed our tool in

the context of a collaborative work. Our industrial project partners are operating in the field of avionics. The cooperation with primarily two domain experts was of mutual benefit. Following a participatory design, we jointly specified suitable strategies for occlusion management and devised corresponding visualizations. While some design decisions were driven by the data and basic prerequisites (e.g., 3D visualization of terrain and data), others were inspired by the domain experts (e.g., presenting information outside the field of view). During the development process, we built upon their domain expertise, identified challenges, and gathered informal feedback.

The achieved results were then assessed in a post-development demonstration and feedback session together with our domain experts. For this purpose, we considered different scenarios. All scenarios more or less required the communication of occluded information. During the discussions, the experts stated that they liked our widget-based design because the basic 3D visualization of terrain and data remains mostly unaltered in the main view. That is, the original shapes of terrain are faithfully represented in the rendered images, except in areas transformed by the interactive lens. Also, they reassured us that activating one or more widgets on demand is indeed helpful for observing occluded parts of the embedded data. Yet, some of the widgets were deemed more effective than others in certain scenarios. For example, the panorama widget and the map widget were selected as candidates for supporting helicopter operations in complex surroundings. Particularly, the abilities to overview data located inside and outside the field of view, and to quickly navigate to certain locations just by clicking in one of the widgets were considered useful for such scenarios. In contrast, the bands widget and histogram widget were judged to be more time-consuming to interpret because of the indirect reference to space. Accordingly, they were taken into account for scenarios with less strict time constraints, e.g., flight route planning. In this regard, the experts approved the provided selection and highlighting techniques, which helped them to relate the information displayed in the widgets to spatial locations depicted in the main view. Applying the lens to inspect parts of the rendered images in detail was considered beneficial as well. But, the experts mentioned that interacting with the lens required slightly more effort to examine an entire image. Hence, at times they relied on the bands widget to locate interesting image parts first, which they then inspected in detail using the lens.

Besides those preferences, the experts appreciated to have one tool at their disposal that integrates several occlusion management strategies. In fact, they confirmed that no single strategy alone would have been sufficient to address all the various problems related to occlusion in different aviation scenarios. They concluded that the abilities to address the different scenarios with our tool, to interactively visualize occluded data even in extensive terrains, and to reduce the need for 3D navigation

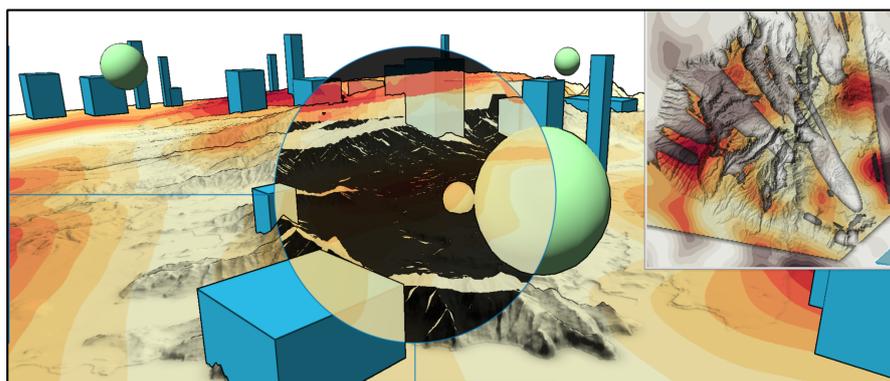


Figure 6.13: Managing occlusion by different types of occluders. To illustrate, stationary pillars (blue) and moving spheres (green) are placed into a 3D terrain. The additional occlusion of data caused by these objects is visualized via the interactive lens and the map widget.

in general are major advantages. As a result, our tool will be used as a basis for further extension, fine-tuning, and evaluation together with our domain experts.

6.7 GENERALIZATION

While we focused on managing occlusion of data embedded in 3D terrain, our solutions can be adopted for handling other geo-visualizations as well. For example, the spatial frame of reference may be directly changed to virtual city models or underwater landscapes. In such environments the geometry can be comparatively more complex. In this regard, our detection and visualization are applicable to occlusion by multiple objects, e.g., added obstacles or landmarks. Our approach also allows for incorporating 3D flight route visualizations (e.g., [138]) and managing related occlusions. Depending on the current task, such objects can be dynamically flagged as occluders and treated in a similar manner as the terrain for the computation of occluded data (Sect. 6.4.1). Figure 6.13 shows an exemplary 3D scene with different types of occluders. Our tool allows to compute and visualize the added occlusion without preprocessing.

Besides geo-spatial data, we successfully applied our visibility widgets to 3D visualizations of volumetric data sets. Particularly, occlusion of data that are mapped onto isosurfaces can be computed with our detection strategies. Our visualizations and interactions help to reveal and analyze hidden information in such 3D scenes. As an example, Figure 6.14 shows the four widgets applied to a visualization of the inner limiting membrane of the human retina. The depicted isosurface is part of a volumetric data set acquired via optical coherence tomography. However, with isosurfaces other types of spatial interactions may occur, e.g., enclosure or intersection (cf. [54]). While our widgets

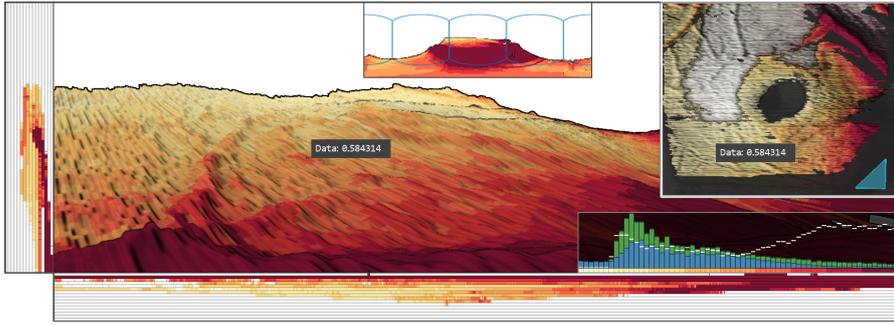


Figure 6.14: Applying visibility widgets to visualization of volumetric data sets. Instead of terrain, an isosurface of the inner limiting membrane of the human retina is shown. The four widgets help to explore visible as well as occluded parts of the color-coded data representing the thickness of the subjacent retinal nerve fiber layer. The map widget shows a visibility history with currently unexplored parts in dark gray.

already can show such occluded data, we are currently improving the visualization to communicate the differences between such types of spatial interaction as well.

Finally, our widget design allows for combinations with other occlusion management strategies that alter the main 3D view. Figure 6.15 illustrates a panorama map [64] and Figure 6.6 shows it integrated in our tool. This way, occlusions can be discovered via our widgets and accessed via the panorama map.

6.8 DISCUSSION AND CONCLUSION

We presented novel visual-interactive support for managing occlusion in 3D terrain visualization. The combination of different strategies opens up new possibilities for detecting, presenting, and interacting with occluded terrain-related data in the aviation domain. Our systematic detection and efficient computation enables processing various types of occlusion in real-time. With the help of four complementary widgets, users can overview visible and occluded data inside and outside of the field of view, and locate and analyze that data in the main 3D visualization. Coordinated interaction facilitates selecting and highlighting in the interlinked views, investigating details on demand, and eases navigation. Integrating the widgets in our flexible visualization tool allows for a full-fledged management of *all* types of occlusion of embedded data in 3D terrain, for which to the best of our knowledge no equivalent approach exists. We conclude that our approach can be a useful aid for exploring any 3D terrain.

A limitation of our approach is that only a single data point per spatial location can be shown in the main view and per widget at a time. While different data tables may be assigned to each activated

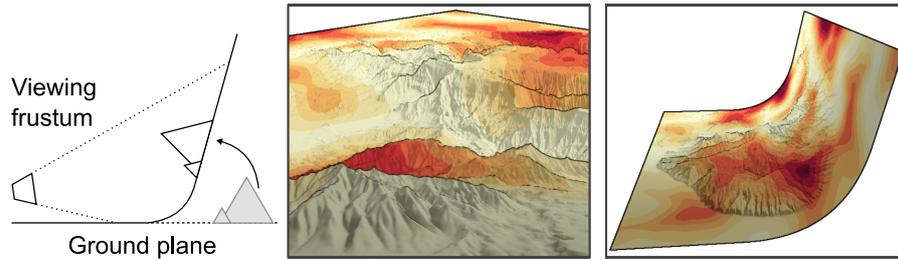


Figure 6.15: Additional occlusion management strategies. The diagram outlines the generation of a 3D panorama by deforming the terrain geometry with regard to a parametric curve [64] (left). The resulting visualization creates a focus zone, a transition zone, and a context zone with increasing distance to the viewer (middle). To exemplify the deformation, the reference point is set to a fixed location, independent of the virtual camera (right).

widget, relating the encoded information can be challenging. So far, we handle overlapping data points from the same data table. More work is required for supporting multiple data tables at once. Likewise, other types of data, e.g., vector field data of wind velocities, have to be considered in future work. Detecting occlusion of such data is possible by computing the contribution of the individual values to each pixel on screen. Visually dissolving occlusion could then be addressed via specific encodings, such as bivariate color-coding, a modified lens, or enhanced histograms.

The discussed examples demonstrate that our approach is suited for analyzing regions of interest, e.g., areas around airports, in moderate-sized environments. We also successfully applied our tool to data covering larger terrain models, e.g., central Europe. Occluded information could be revealed. Yet, interactively exploring large and complex environments can be time-consuming. It remains to be studied how discovering, accessing, and relating relevant information in such scenarios can be further supported. Guiding navigation by computing alternative viewpoints (e.g., [139]) that show interesting phenomena with low occlusion ratios may be an option to investigate in the future.

Our solution is based on multiple coordinated views. This is a well-established visualization approach, which has been proven to be useful for exploring complex data in general [140] and occluded information in particular [54]. Yet, one aspect that has to be considered is the cognitive load imposed on the user for matching separate visualizations. This is why we rely on interaction techniques to help to relate the shown information. Selecting and highlighting indicates associated parts in main view and the widgets, whereas navigation facilitates aligning the visualized data in the views. In addition, users can arrange and resize the widgets freely to focus on specific aspects in combination. An interesting open question is how we could support an automated placement of the widgets to further reduce the cognitive load. This

requires assessing how well a specified arrangement suits a certain scenario.

To further improve our approach, we plan to incorporate automated cues for hinting at occluded data. For example, the widgets could first be deactivated by default. While exploring 3D terrain with our tool, simplified visual indicators can then be shown to make users aware of hidden information. Clicking on an indicator will suggest an appropriate widget based on the given situation. This allows users to initially check and analyze the data, and, if required, retrieve further details on demand. Moreover, it will help users to focus on relevant information by showing only those additional views that are currently needed. In this context, studies for evaluating which occlusion management strategy matches which situation and analysis task will become necessary.

FEATURE-DRIVEN VISUAL ANALYTICS OF CHAOTIC PARAMETER-DEPENDENT MOVEMENT

SCIENTIFIC CONTRIBUTION: In this chapter, we contribute novel **VA** solutions to the complexity problem in visualizations of 3D spatial data. Beyond existing approaches, our concepts support the systematic specification of relevant information as well as the automated extraction and interactive visualization of matching features. This allows to concentrate on key data characteristics, even when comparing large amounts of complex movements. The developed concepts formed the basis for reducing the complex 3D retinal layers in **OCT** data to diagnostically meaningful information. The journal article [26] included here received the best paper award at the EuroVis conference in 2015.

AUTHOR CONTRIBUTION: The author of this thesis co-developed substantial parts of the presented **VA** approach, including joint conceptualization, implementation of essential software parts, data analysis, and evaluation with domain experts. The author further wrote and reviewed major parts of the original publication.

ABSTRACT: Analyzing movements in their spatial and temporal context is a complex task. We are additionally interested in understanding the movements' dependency on parameters that govern the processes behind the movement. We propose a **VA** approach combining analytic, visual, and interactive means to deal with the added complexity. The key idea is to perform an analytical extraction of features that capture distinct movement characteristics. Different parameter configurations and extracted features are then visualized in a compact fashion to facilitate an overview of the data. Interaction enables the user to access details about features, to compare features, and to relate features back to the original movement. We instantiate our approach with a repository of more than twenty accepted and novel features to help analysts in gaining insight into simulations of chaotic behavior of thousands of entities over thousands of data points. Domain experts applied our solution successfully to study dynamic groups in such movements in relation to thousands of parameter configurations.

ORIGINAL PUBLICATION: [26] — M. Luboschik, **M. Röhlig**, A. T. Bittig, N. Andrienko, H. Schumann, and C. Tominski. "Feature-Driven Visual Analytics of Chaotic Parameter-Dependent Movement." In: *Computer Graphics Forum* 34.3 (2015). Best paper award, pp. 421–430. doi: [10.1111/cgf.12654](https://doi.org/10.1111/cgf.12654)

FURTHER READING: Besides movement data, we adapted our interactive visualization design in a new **VA** approach for parameter dependencies of segmented time-series data:

- ▷ **M. Röhlig**, M. Luboschik, M. Bögl, F. Krüger, B. Alsallakh, T. Kirste, S. Miksch, and H. Schumann. “Supporting Activity Recognition by Visual Analytics.” In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 41–48. DOI: [10.1109/VAST.2015.7347629](https://doi.org/10.1109/VAST.2015.7347629)

Based on this, we introduced several extensions for the sequencing of categorical time series and the segmentation and labeling of multivariate time series:

- ▷ C. Richter, M. Luboschik, **M. Röhlig**, and H. Schumann. “Sequencing of Categorical Time Series.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 213–214. DOI: [10.1109/VAST.2015.7347684](https://doi.org/10.1109/VAST.2015.7347684)
- ▷ T. Gschwandtner, H. Schumann, J. Bernard, T. May, M. Bögl, S. Miksch, J. Kohlhammer, **M. Röhlig**, and B. Alsallakh. “Enhancing Time Series Segmentation and Labeling Through the Knowledge Generation Model.” In: *Poster at Eurographics Conference on Visualization (EuroVis 2015)*. 2015
- ▷ J. Bernard, E. Dobermann, M. Bögl, **M. Röhlig**, A. Vögele, and J. Kohlhammer. “Visual-Interactive Segmentation of Multivariate Time Series.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by N. Andrienko and M. Sedlmair. The Eurographics Association, 2016. DOI: [10.2312/eurova.20161121](https://doi.org/10.2312/eurova.20161121)

7.1 INTRODUCTION

Visual Analytics (**VA**) has become an indispensable means to help us understand the characteristics of movements in space and time [68]. Here, we address movements that were synthesized in an effort to simulate processes that are difficult to observe otherwise. Such simulations are typically controlled by parameters whose influence on the simulation outcome is not clear upfront. So, in addition to investigating movement in space and time, there is also the need to understand the movement’s dependency on the parameter configuration.

The analysis of parameter dependencies is a rather challenging issue [144], particularly for simulations of complex movements. A reason is that we have to integrate the visual representation of parameter configurations with the corresponding movement in a comprehensible way. On top of that, there might be thousands of different configurations, each resulting in thousands of unconstrained or even chaotic movements. In such cases, severe clutter and over-plotting will make it hard

to discern even basic movement characteristics from the data, not to mention gaining insight into the influence of parameters.

The related work reviewed in [Section 7.2](#) indicates that extracting and visualizing high-level features can be more appropriate than showing the raw data. For example, time-evolving features have been used successfully to explore and compare single group movements [69]. However, we are still lacking approaches to analyze *all* movements belonging to a specific parameter configuration and to explore these in regard to *all* alternative configurations. To close this gap, we tightly integrate analytic, visual, and interactive means in a novel VA approach for studying chaotic movement data in relation to parameter dependencies. The abstract outline of our approach is as follows:

ANALYTIC: We extract high-level features to capture the characteristics of *all* movements belonging to a specific parameter configuration. We consider basic features, group features, and region features. Advanced features further increase the level of abstraction.

VISUAL: We visualize the features via a novel visual design that integrates (i) an overview of *all* movements conjoining feature and parameter distributions and (ii) detail views reflecting certain aspects of the high-level features back onto the low-level raw data.

INTERACTIVE: We provide interaction techniques facilitating the analysis of parameter dependencies of movement features, the exploration of feature evolution with regard to individual parameter configurations, and the comparison of features corresponding to different configurations.

The analytic feature extraction and the interactive feature visualization are detailed in [Section 7.3](#) and [Section 7.4](#), respectively. In [Section 7.5](#), we demonstrate our solution by applying it to analyze chaotic movements simulated for thousands of different parameter configurations over thousands of time steps.

7.2 MOTIVATION AND RELATED WORK

Next, we outline the motivation for our research and discuss the related work in VA of movements and parameter dependencies.

7.2.1 Motivation

Our work is motivated and driven by recent advances in systems biology. In particular, research on spatial simulation has gained momentum as it expands our ability to understand biological phenomena [145–147]. The key idea is to abstract from nature’s details and create generic models of biological processes. Some of the details abstracted away

during the modeling are captured in parameters to be experimented with when simulating the models. Therefore, multiple simulation runs with different parameter configurations are necessary. The simulation generates large data sets containing parameter-dependent spatial and temporal information about the entities and their movements.

As a concrete example, we consider the investigation of dynamic interactions between *receptor proteins* and *lipid rafts* on the surface of human cells [77]. These interactions play an important role in cellular signaling, for instance, in the cancer-related *Wnt* pathway [148].

Studying such dynamic interactions is a task that is typically difficult to carry out. There are several reasons for that. First, the spatial simulation is based on stochastic Brownian motion. The resulting movement trajectories are in a sense *chaotic*, because they are entangled to a large degree. Second, the entities may pick up, take along, and drop other entities during the simulation. This way, they form *dynamic groups*, which are of high interest, but difficult to grasp. Third, the simulated interactions *depend on the parameter configuration*, where the impact of individual parameters or combinations of specific parameter values is largely unknown.

Our objective is to develop a solution that helps analysts to unveil the influence of parameters on the movement dynamics so that they can evaluate the simulation approach in itself and confirm or reject hypotheses about the underlying biological model. Although we address chaotic movements from systems biology, our approach is generic enough to be applicable to other problems as well.

7.2.2 Related work

Our research is related to [VA](#) of movement and [VA](#) of parameter dependencies.

VISUAL ANALYTICS OF MOVEMENT: concentrates on (1) visualizing spatial and temporal aspects of individual trajectories and sets of trajectories, (2) visualizing movement attributes along trajectories, (3) detecting stops, interactions between trajectories, and other kinds of events, (4) aggregating movement data in space and time and visualizing the resulting aggregates, and (5) revealing relationships between movement and the environment. A profound overview and systematization can be found in [68]. We review existing work with regard to chaotic movements and movements of dynamic groups, which are key aspects of our research.

A general problem when visualizing chaotic movements is the severe over-plotting. Typical approaches to tackle over-plotting include clustering [149], aggregation with density kernels [150], or flow maps [151], which provide summaries of the underlying data. However, with the

chaotic movements that we address in our work, these approaches are likely to fail.

Another widely accepted approach to deal with complex data is feature visualization [71]. The basic idea is to visualize derived features, rather than the raw data. Feature-based approaches have already been used successfully for analyzing movements of groups [69]. Group movements are also studied in [75], yet without following a feature-based approach. In both works, the groups are static, that is, group membership is not allowed to change. Dynamic groups are addressed in [76], but only for rather few simple movements.

VISUAL ANALYTICS OF PARAMETER DEPENDENCIES: deals with visualizing given input parameters of configurable processes along with the corresponding output. A recent survey can be found in [74].

Related to our work are global-to-local approaches, which provide an overview in the beginning and allow the user to drill down into details. An example is the overview-driven approach for parameter dependencies of large time series data [152]. However, the specifics of chaotic parameter-dependent movement have not been addressed yet.

To be able to generate overviews it is often necessary to reduce the number of parameter configurations and output size. A typical way to do so is to use surrogate models, which predict or interpolate the corresponding output [153–155]. Undirected optimization [156, 157] and parameter space partitioning [158] are approaches that make use of reduced result previews or clustering. A problem with these approaches is that they cannot be applied directly to chaotic movements, which can be difficult to predict, interpolate, reduce, or cluster, if this is possible at all.

With regard to movement analysis and parameter dependencies, we also found related work in the flow visualization literature [159–161]. These approaches basically superimpose results of different flows, which are each filtered to a reduced set of streamlines, which in turn are comparable to our raw data trajectories. Given a certain degree of spatial similarity among the resulting trajectories, direct visual comparison of different parameter inputs becomes possible.

Recently, the idea of using features has also been applied to analyze parameter dependencies of complex simulations [162]. In contrast to the single-valued scalar features used there, we propose using time-evolving features. In doing so, we extend previous feature-based approaches for movement analysis [69], as indicated earlier.

In summary, we see several individual solutions, but none that suits our needs directly. Therefore, our goal is to develop a feature-based approach that works with chaotic movements and dynamic groups, and that also supports the analysis of parameter dependencies. To achieve this goal, we (i) introduce tailored features to capture key characteristics of chaotic movements, including dynamic groups, (ii) visualize the

features in association with parameter configurations to enable users to analyze their dependencies, and (iii) integrate appropriate interaction to allow users to look into details and compare different aspects of the data.

7.3 FEATURE EXTRACTION

We consider data of the following form. A data set $D = \{R_1, \dots, R_r\}$ consists of r simulation runs. Each run $R_i = (P_i, M_i)$ with $1 \leq i \leq r$ is a pair of a parameter configuration P_i and the corresponding movement M_i . A movement $M_i = \{T_1, \dots, T_m\}$ consists of the trajectories of m moving entities. Each trajectory T_j with $1 \leq j \leq m$ is sampled at uniform intervals so that we obtain a set of n points $T_j = \{\mathbf{t}_1, \dots, \mathbf{t}_n\}$. A point \mathbf{t}_k with $1 \leq k \leq n$ stores information about an entity at a particular time step. This includes information that is readily available such as the entity's position or type, but also derived attributes such as speed, acceleration, or the distance to particular other entities.

The analytic part of our approach is to condense the complex data down to information that is manageable and relevant. Our method of choice is feature extraction. The feature extraction is based on two principal steps, which are carried out for each movement M_i . First, the data points are enriched with derived measures. This is to inject into the data meaningful information that can help to characterize the movement. The second step is aggregation. As illustrated in [Figure 7.2 \(a\)](#), the goal is to reduce M_i with its m trajectories consisting of n points to a single aggregated feature time series of length n . To this end, we consider all points \mathbf{t}_k of all m trajectories and aggregate them to a single feature value f_k . As we do this for all $1 \leq k \leq n$ time steps, we get n feature values f_1, \dots, f_n that characterize the movement over time.

The net effect is that we replace the complex movement M_i by a feature $F_i = \{f_1, \dots, f_n\}$. The difficult problem of visualizing (P_i, M_i) for $1 \leq i \leq r$ is thus reduced to the problem of visualizing (P_i, F_i) , which is much easier to solve as we will see in [Section 7.4](#).

Apparently, a single feature alone will not suffice to capture the richness of movement data. Therefore, we consider a repository of feature definitions in four categories: (1) basic features, (2) group features, (3) region features, and (4) advanced features. Basic features capture general movement characteristics. Group features address characteristics of dynamic groups, such as fluctuations in memberships and retention periods. Region features perform more complex spatio-temporal aggregations to characterize regions of interest and their evolution over time. Advanced features are a way to extract features over features.

Next, we describe exemplary features from all four categories. For the sake of brevity, we restrict ourselves to brief informal explanations.

For a complete list of features, including design rationales developed with domain experts, we refer to the supplemental material¹.

BASIC FEATURES: General characteristics of movement trajectories can be captured by aggregating basic properties, such as speed, direction, and distance of movements. In addition to considering such features across all moving entities, our approach can compute them also with respect to entities of different types (e.g., receptor proteins or lipid rafts). This allows us to investigate the entire movement as well as specific subsets of moving entities. Some aspects of the movement behavior can for instance be characterized by averaging the distances of all entities of one type to the closest entity of another type.

GROUP FEATURES: As indicated earlier, previous work on analyzing groups mostly considered *static* groups. We are interested in *dynamic* groups, that is groups that emerge, continue to exist with changing members, and decay.

We build upon previous work on non-spatial groups, for example, tracking changes in group sizes [163] or changes of structural properties [164]. To capture the dynamic behavior of spatial groups and to analyze how the behaviors of group members and non-members differ, we specify several group features:

GROUP COUNT: basically captures the number of existing groups per time step.

GROUP AFFILIATION RATIO: describes the overall ratio of group members and entities not being contained in any group.

GROUP LOAD: relates the actual group sizes to the maximum allowed capacity of groups.

GROUP RETENTION PERIOD: captures the time period between entities joining and leaving a group accumulated for all current group members. This measure can further be aggregated for all groups to describe the temporal fluctuation of group memberships.

REGION FEATURES: The feature definitions described so far focus on basic movement characteristics derived directly from the movement trajectories. A limitation of these features is that high-level spatial or temporal characteristics can not be captured well. However, looking at spatial patterns and their temporal evolution is often necessary to fully understand movements and parameter influence.

To better capture spatial aspects, we integrate 2D density maps [165] when computing features. Density maps are computed for all time steps, effectively creating a 3D space-time density volume. This allows us to

¹ Supplemental material of [26] is available at doi:[10.1111/cgf.12654](https://doi.org/10.1111/cgf.12654)

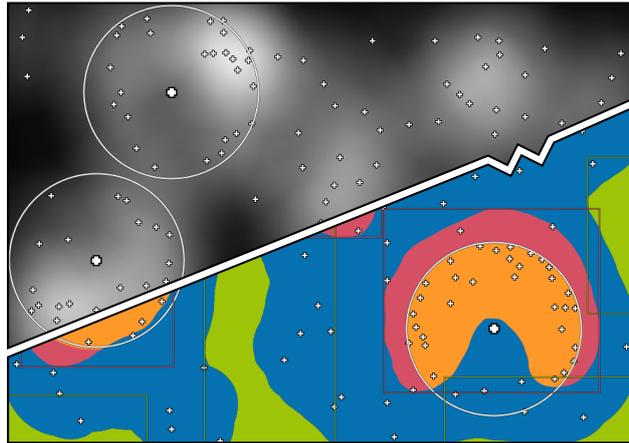


Figure 7.1: Feature extraction via density maps. Entities (white marks) and groups (circles) superimposed on a density map (top) and extracted regions (bottom) with high (red) and low (green) density and overlaps with groups (orange).

capture generic patterns independently of data size and specific data properties, such as extents of entities or topologies of groups [166].

The 2D density maps are further analyzed to extract regions of interest with low and high density regarding suitable thresholds. Figure 7.1 illustrates entities, groups, density map, and extracted regions. Further tracking and aggregating properties of these regions of interest over time enables us to extract spatio-temporal features:

REGION COUNT: is the number of disjoint regions of interest (e.g., high or low density regions) per density map.

REGION SIZE: corresponds to the aggregated size of all regions of interest per density map.

REGION RATIO: is generally applied to relate regions of interest with respect to their density (e.g., low and high density regions) and with regard to certain types of entities (e.g., high density regions for one entity and low density regions of another entity).

ADVANCED FEATURES: With the features introduced so far it is possible to study a variety of movement characteristics. To be able to combine features and to generate even higher level abstractions, we introduce the notion of advanced features, i.e., features over features.

Advanced features can be derived by further analytical processing of previously extracted features. For instance, by applying temporal clustering of feature values it is possible to investigate temporal patterns across multiple simulation runs, such as common feature evolution or time periods of specific behavioral variation. We generate features over features via a Self-Organizing Map (SOM), by which we obtain clusters with similar feature characteristics over time. An exemplary use for such

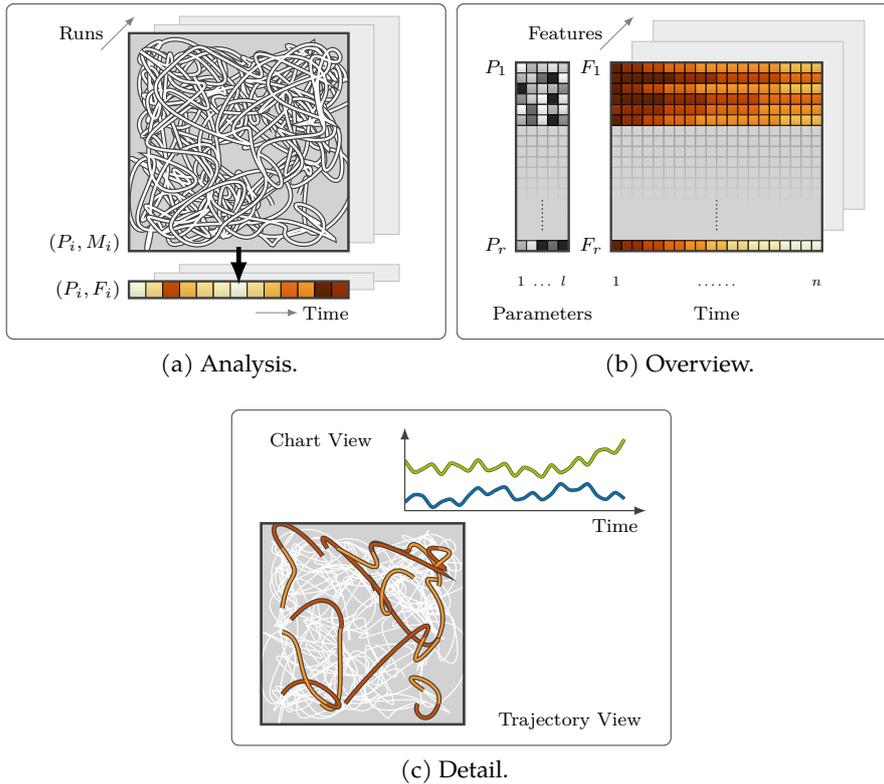


Figure 7.2: Analytically extracted high-level features are visualized via interactive overview and detail representations.

features is to verify that stochastic simulations indeed do not exhibit periodic temporal patterns.

In summary, the feature extraction computes analytic abstractions to capture key characteristics of the movement. We consider a wide variety of feature definitions as collected in our feature repository, which is available as supplemental material². Next we describe how the features are visualized in relation to parameter configurations.

7.4 FEATURE VISUALIZATION AND INTERACTION

We study parameter configurations and associated movements (P_i, M_i) for multiple simulation runs $1 \leq i \leq r$. In terms of parameters, we define a parameter configuration $P_i = \{p_1, \dots, p_l\}$ as a set of l parameter values. The number of parameters l is constant for all simulation runs. As illustrated in Figure 7.2 (a), the analytic feature extraction already reduced the complex movements M_i in space to time series of feature values of the form $F_i = \{f_1, \dots, f_n\}$. To give a rough measure of the size of our data, the number of parameters l is around ten, simulation runs r can be in the thousands, and time steps n can be in the thousands

² Supplemental material of [26] is available at doi:10.1111/cgf.12654

as well. [Section 7.5](#) provides more precise numbers for a use case in systems biology.

Our primary objective is to support the exploration and analysis of the aforementioned data. This involves several visualization tasks, which can be differentiated into overview tasks and detail tasks.

OVERVIEW: At the overview level, users explore the data with respect to temporal evolution of features in relation to all parameter configurations. The goal is to identify interesting patterns and to analyze them with regard to the underlying movements.

DETAIL: For a more detailed investigation, the analysis is focused on selected simulation runs with their parameter configurations and corresponding feature values. Focusing on selected runs enables users to compare interesting patterns in detail and to gain a better understanding of the influence of parameters. Since our features sacrifice spatial information for the sake of analytic abstraction, we also have to support linking back the analysis to the spatial domain, at least partially.

With these data definitions and visualization tasks in mind, we developed a dedicated visualization design based on linked overview and detail views. [Figure 7.2](#) summarizes the overall strategy of our VA solution.

7.4.1 Overview visualization

The overview task focuses on exploring movement characteristics and parameter dependencies across all simulation runs. Therefore, we have to visualize all parameter configurations P_i and associated feature time series F_i for all $1 \leq i \leq r$. To this end, the overview presents the data in a compact matrix-like fashion.

VISUAL ENCODING: In the matrix, the i -th row represents the i -th simulation run (P_i, F_i) . The first part of a row visualizes the parameter configuration $\{p_1, \dots, p_l\}$ and, separated by a small gap, the remainder of the row visualizes the feature values $\{f_1, \dots, f_n\}$. This arrangement is illustrated in [Figure 7.2](#) (b). Note that the matrix shows only one feature definition. Yet, switching between different feature definitions from the feature repository is possible at any time.

The cells of the matrix are color-coded using distinct color maps for parameter values and feature values. For the quantitative feature values we apply color maps from ColorBrewer [132]. Parameter values are color-coded with hue-less shades of gray. This clearly separates feature values (colors w/ hue) from parameters (colors w/o hue). Darker shades of gray represent lower parameter values and brighter shades stand for higher values. If required, the default color-coding can be interactively adjusted.

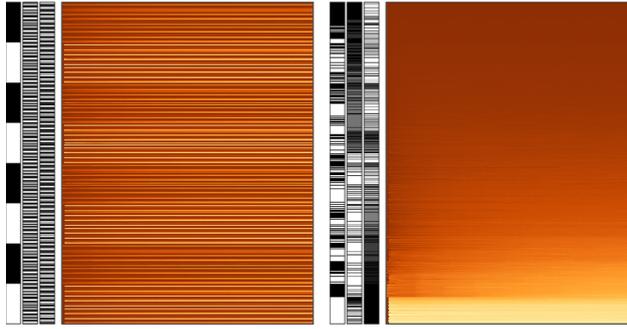


Figure 7.3: Overviews sorted row-wise by parameter configurations (left) and according to feature behavior (right).

SORTING: When displaying the data of thousands of simulation runs in a row-wise fashion the order of rows is vital for discovering patterns in the data. Because fully manual sorting is impractical, we provide two ways for automatic sorting: (i) parameter-based sorting and (ii) feature-based sorting. Sorting based on parameter values facilitates the interpretation of parameter influence on the data, e.g., for hypothesis testing regarding the parameters. On the other hand, sorting based on feature values helps to identify simulation runs with similar behavior, e.g., to build hypotheses when inspecting the related parameters.

While sorting individual values is trivial, sorting sets of values (our $\{p_1, \dots, p_l\}$ and $\{f_1, \dots, f_n\}$) is rather challenging. We provide various metrics and algorithms for sorting, including Euclidean, Hausdorff, Fréchet, and Levenshtein distance, average squared error [167], and dynamic time warping [168], as well as SOMs and gradient decent. We achieved good results with Euclidean distance combined with SOM for feature-based sorting and gradient decent for parameter-based sorting. Figure 7.3 shows the visual effect of sorting on the emergence of patterns.

As sorting according to all parameters or all time steps might not lead to the desired insight, the user can choose to restrict the sorting to subsets of parameters or time steps. To support further data exploration, it is possible to experiment with the different sorting methods and apply them to different parameter subsets or time intervals. If the automatic sorting is still not satisfactory, there is always the option to reorder individual rows or groups of rows manually.

INTERACTIVE EXPLORATION: Given the size of the data, showing all observed time steps and all parameter values for all simulation runs can easily exceed the available screen space. Therefore, the matrix resides in a zoomable space allowing independent scaling along the axis of simulation runs (rows) and along the time axis (columns). By incorporating interactive zooming and panning, users are enabled to steer the visual analysis process according to their task-specific needs.

To help users in maintaining orientation during exploration, we integrate additional visual cues. Miniature scroll bars indicate where the current view is located with respect to the entire data representation. Further, we use two over-plotting indicators. They tell the user whether simulation runs (rows) and/or time steps (columns) are affected by over-plotting. Red indicators warn the user that perceived patterns could be artifacts due to over-plotting. To resolve such ambiguities quickly, the over-plotting indicators can be clicked to smoothly animate the view to a zoom level where no over-plotting occurs.

Equipped with these interaction facilities, our compact matrix-like visualization provides an overview and supports the identification of basic value distributions and temporal patterns. For example, constant feature values at specific time intervals or uniform feature evolution over time are reflected by rows with ranges of constant or gradually changing colors, respectively. Furthermore, dependencies between parameters and features (and hence the underlying movements) can be discerned by looking vertically for patterns under different orderings. Absence of vertical patterns may also indicate weak or no parameter influence. Although the overview can already lead to specific insights its real value is to initiate more targeted follow-up investigations.

7.4.2 *Detail visualization*

To go beyond overview visualization and basic exploration, we integrate techniques for detailed comparison and re-establishing the spatial context, as illustrated in [Figure 7.2](#) (c). These techniques are key to exploring movement data in changing analysis scenarios.

DETAILED COMPARISON: More targeted investigation typically means comparing selected subsets of the data in detail. However, color-coded visual representations are less suited for analyzing and comparing numerical values in detail [169]. Therefore, we integrate a separate chart view.

As illustrated in [Figure 7.4](#) (top), the chart view shows time series plots representing feature values for simulation runs selected from the overview. Selections and plots are associated with unique highlighting colors to make them easier to distinguish and track across the visualization. In [Figure 7.4](#), the time series plot in blue corresponds to the selected simulation run indicated by a blue bar across the main matrix. The chart view is positioned above the overview and is aligned horizontally to maintain the temporal context. Additionally, zoom and pan operations in time are linked to preserve temporal alignments between the overview and the line plots.

By showing selected time series as line charts we facilitate a more precise analysis and direct comparison of simulation runs. This allows the user to focus on specific parameter combinations (e.g., similar pa-

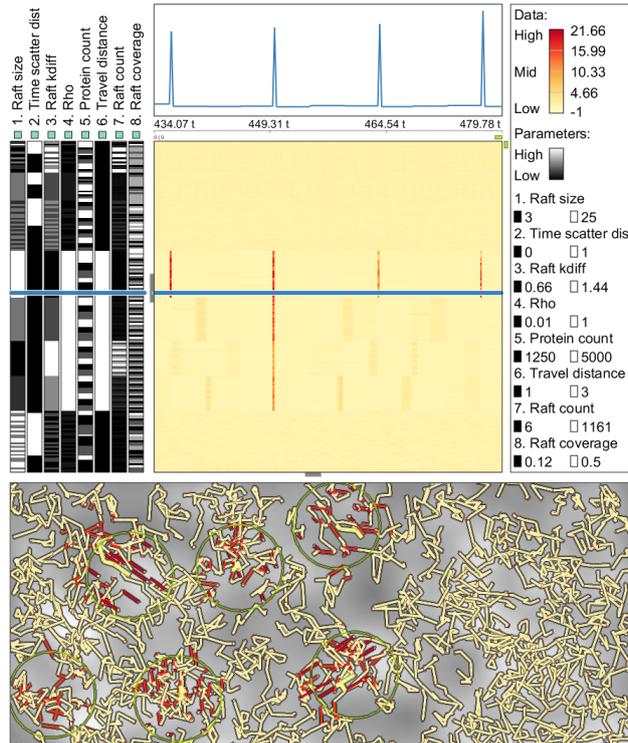


Figure 7.4: The overview of parameters and features (center) in conjunction with the chart view (top), the trajectory view (bottom), and a legend (right) facilitate spatio-temporal data exploration.

parameter values) and to compare related feature characteristics. The other way around, the user can also start with interesting patterns of feature values and inspect their relation to the associated parameters.

RELATING BACK TO SPACE: Our approach is based on analytical abstractions of the rather complex and even chaotic movements. These abstractions make it possible to reduce the amount of data to be displayed at a time. Yet, this comes at the cost that the spatial context and the influence of individual movement trajectories is lost to some extent.

To compensate for this, we incorporate an additional trajectory view that relates features back to the raw movement data $M_i = \{T_1, \dots, T_m\}$, but for selected simulation runs only, as indicated in Figure 7.4 (bottom). The movements are shown as trajectories $T_j = \{\mathbf{t}_1, \dots, \mathbf{t}_n\}$. Spatial aspects and feature characteristics are married by combining a spatial layout based on the trajectory points $\mathbf{t}_1, \dots, \mathbf{t}_n$ with a color-coding based on the feature values f_1, \dots, f_n . Moreover, the trajectory view can be blended with a selected 2D density map generated during the feature extraction. Figure 7.4 (bottom) shows a gray-scale density map in the background.

As color-coding and density maps establish a connection between feature values and locations where the raw movements took place, the spatial context is partially restored. Linked zooming and panning in time further strengthens the connection to temporal aspects of the

raw movement trajectories. Focusing on a selected temporal interval of interest also significantly reduces visual clutter. This helps to investigate relationships between certain patterns of feature values and observed movement behaviors. For example, the user might be able to relate feature values representing low movement speeds to spatial conglomerations of trajectories.

The chart view for detailed comparison and the trajectory view for linking to spatial aspects complete our feature-driven VA approach. In the next section, we apply this approach to a problem from systems biology.

7.5 APPLICATION TO SYSTEMS BIOLOGY

The approach presented so far is a general solution applicable to different types of parameter-dependent movements. Yet our work has been largely motivated and driven by applications in systems biology. In the following, we present a use case where domain experts apply our solution to study dynamic interactions between receptor proteins and lipid rafts on the surface of human cells.

The data were generated using an ML-Space simulator in combination with movement synthesis based on Brownian motion [170]. Several properties describe the lipid rafts and proteins, including position and size. Lipid rafts and proteins move according to a diffusion parameter k_D . The Brownian motion is simulated by individually calculating displacement vectors with a random direction and a normally distributed average length depending on the smallest entity. During the movement, dynamic groups are formed by proteins docking to lipid rafts. Proteins enclosed in lipid rafts move along with them depending on their fluidity factor ρ , which also controls the probability of proteins leaving the lipid raft.

Movement updates also include collision detection. Overlaps between entities of the same type (i.e., protein with protein and lipid raft with lipid raft) are prohibited and are resolved stochastically. Collisions of a protein with a lipid raft are handled by pushing the protein a little further so that it is either fully inside or outside the lipid raft depending on if the protein is entering or leaving the lipid raft.

In summary, eight parameters control the simulation, including fluidity, entity size, entity counts, and traveled distances. The domain experts determined 1,968 meaningful parameter configurations for which simulations were run. Each simulation run describes the chaotic movement of up to 1,161 lipid rafts and 5,000 proteins depending on the parameter configuration. The individual simulations covered 4,000 time steps.

The domain experts applied our solution to analyze the simulation outcome. In a pre-process, all features of the repository were computed to allow the experts to quickly switch between different movement characteristics. Among others, the following results could be obtained.

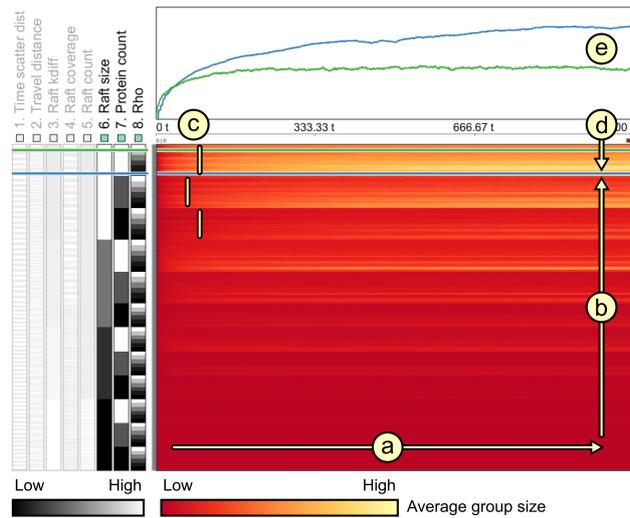


Figure 7.5: The visualization shows that average group size depends on the size of lipid rafts and the fluidity ρ .

INSIGHTS RELATED TO GROUPS: Figure 7.5 shows the parameter dependency of the *average group size* feature, which captures the average number of proteins inside lipid rafts over time. Although this feature is rather simple, it nicely illustrates parameter dependencies in our data. For this purpose, a hierarchical sorting has been applied based on the parameters *raft size*, *protein count*, and ρ . At first glance, Figure 7.5 shows an overall temporal trend of low to high group sizes from left to right (a) and also a trend across simulation runs from bottom to top (b). The temporal trend (a) reflects the fact that lipid rafts start empty and collect proteins incrementally. The trend (b) represents the dependency of group sizes on the parameter *raft size*.

A second observation can be made by looking at the row-wise bands (c), which show different shades aligned with the parameter *protein count*. This dependency is due to the fact that large groups can only emerge if the number of potential group members is sufficiently high.

A third pattern (d) is visible within the bands. It appears to be related to the parameter ρ . To study this pattern further, two simulation runs with potentially large groups (large raft size and high protein count) were selected and are shown in the chart view (e) in greater detail. The line chart reveals that low values of ρ lead to constantly increasing group sizes (blue line), whereas high values of ρ result in stagnation of group sizes below their potential maximum (green line).

Switching to the *average retention period* feature, as shown in Figure 7.6, while maintaining the selection and order of the simulation runs leads to further findings. Similar intra-band gradients (a) show the logical influence of ρ on the duration for which proteins remain inside lipid rafts (b). It becomes clear that the stagnating group sizes (green line (e) in Figure 7.5) are caused by low retention periods (green line (b) in Figure 7.6). In other words, because of the high fluctuation, the lipid

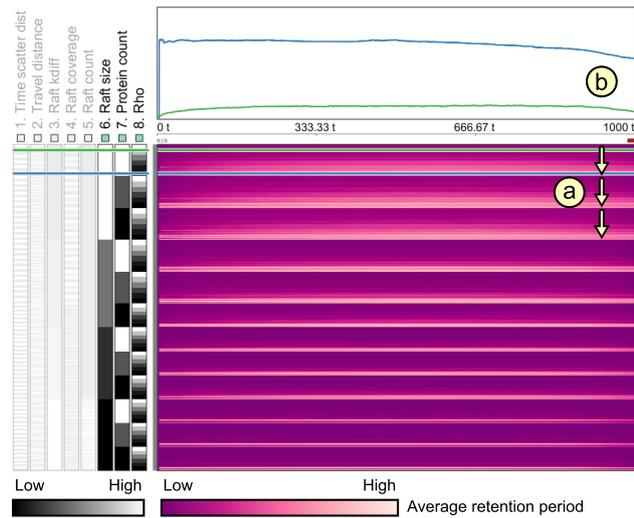


Figure 7.6: Visualizing average group retention confirms the influence of parameters *rho* and *raft size*.

rafts drop the same number of proteins as they pick up, which inhibits further growth of groups.

CONFIRMING THE SWEEPING EFFECT: One particular pattern our domain experts were anticipating is the so-called *sweeping effect*. The sweeping effect relates to the fact that the space around lipid rafts is only sparsely populated. This phenomenon is due to the lipid rafts' random movement, which causes them to collect nearby proteins, effectively emptying the space around them. Visualizing the raw data trajectories of lipid rafts and proteins helps to identify this effect visually, yet only for a limited number of moving entities [171].

To investigate this effect, the domain experts set up the hypothesis that the empty space slowly emerges over time and thus the sweeping effect should become apparent by a steady increase of the average distance of non-member proteins to the nearest lipid raft. A corresponding feature was specified and extracted from the data. Figure 7.7 shows a SOM-sorted visualization of the feature. For several simulation runs, the hypothesized steady increase is particularly visible (a), making the sweeping effect quantitatively graspable for the first time. The result of the effect can be emphasized further by displaying the feature values of selected simulations as line charts (b) or by showing the 2D density maps for selected time points (c). Regarding parameter dependencies of the effect, the experts identified the parameters *raft size* and *rho* to have major influence. For example, (d) shows that only large lipid rafts with a low fluidity *rho* are capable of gathering and holding surrounding proteins in a way that establishes a noticeable effect.

The previous paragraphs illustrated how our approach can be applied to gain insight into chaotic movement simulations from systems biology.

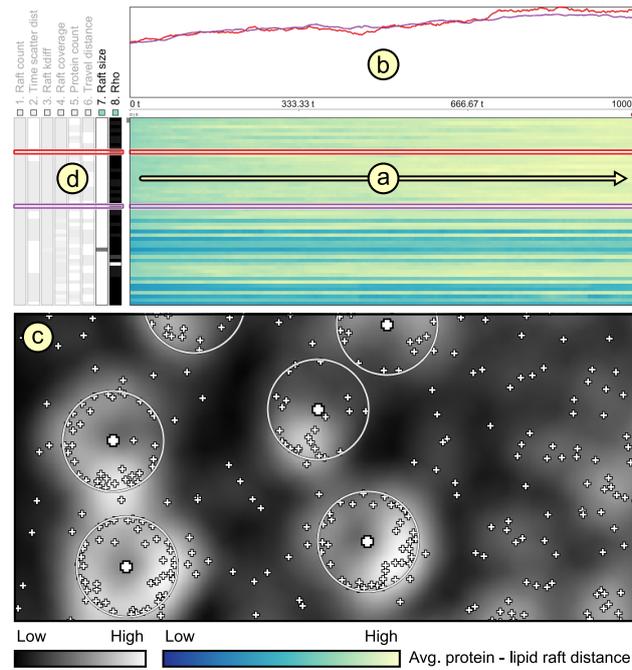


Figure 7.7: Visualizing the *average distance of non-member proteins to the nearest lipid raft feature* in conjunction with the detail line chart and a selected 2D density map helps in studying the *sweeping effect*.

Next we briefly outline how we designed our solution together with simulation experts.

USER PARTICIPATION: Our solution is the result of a participatory design process starting from prior work [171] and [152]. We cooperated with a group of five domain experts. The cooperation was of mutual benefit. We could build upon their domain expertise and devise and specify interesting movement behaviors as features. Collaborative data analysis sessions and informal user feedback helped us to design the visualization and the associated interaction so that they are indeed helpful to the analysts. While some design decisions were driven by the addressed data and tasks (e.g., compact color-coded matrix representation), others were inspired by the domain experts (e.g., chart view for comparison and trajectory view for linking back to space).

In turn, the experts benefited from our solution as it provided them with valuable *new* insight into their simulations, of which we could describe only a few here. The ability to explore and even compare different aspects of the chaotic movements of up to 25,000 entities across multiple parameter configurations was identified to be a major advantage.

GENERALIZATION: While our use case focuses on data from systems biology, we envision applications in other fields as well. Particularly promising are parameter-dependent simulations of crowd behav-

ior in mass events, which certainly involve dynamic groups and rather chaotic movement trajectories. But also actual real-world data, even without dependencies on parameters, could be interesting to analyze. Examples are dynamic groups in sports (e.g., breakaway groups in cycling) or animal behavior in the wild (e.g., wolves patrolling their territory).

The visualization and interaction part of our solution will be one-to-one applicable to such data. In scenarios where no parameters are involved, the visualization could instead show the conditions or influential factors under which the data have been recorded. The described feature repository will be useful for a broad range of analytic questions. It is also possible to adjust existing features or develop new ones to better address the particularities in alternative application backgrounds.

Finally, it is obvious that our feature-driven approach is applicable to non-chaotic and also constrained movements, with or without the consideration of dynamic groups.

7.6 DISCUSSION AND CONCLUSION

We presented a VA approach for parameter-dependent movements. The analytic extraction of features of different kinds opens up new possibilities for exploring unconstrained, crowded, and chaotic movements where the moving entities group dynamically. With the help of an overview visualization of parameters and features, users can spot interesting patterns. Selecting individual simulation runs allows the user to conduct in-depth inspections using a chart view and a trajectory view. Coordinated interaction facilitates data exploration. We can conclude that through combining analytic, visual, and interactive means, our approach is a useful aid for analyzing complex movements.

A limitation of our solution as well as any other feature-driven approach is that the expressiveness of the visualization is limited by the expressiveness of the features. This is why we rely on a feature repository to capture many data characteristics. Basic features support basic analytic tasks, whereas more complex features can provide more high-level insights. However, finding meaningful feature definitions that match specific data sets or analysis tasks is challenging in general. An interesting question is how we could support the user in designing feature definitions on the fly and in steering the feature extraction process while it is running. This requires assistance in evaluating how well a feature captures certain characteristics and to which extent individual trajectories influence the outcome of the feature extraction.

The examples we described here demonstrated that our approach is suitable for around ten of parameters. The limited number of parameters allowed us to simply unroll the parameter space to a linear order of parameter configurations and show the features with regard to them. Basic parameter dependencies could be revealed. Yet, com-

plex influences of a larger number of parameters and high-dimensional correlations may be difficult to grasp. It remains to be studied how such scenarios can be handled. An interactive aggregation of multiple parameters or parameter configurations into clusters may be one option to investigate in the future.

We further plan to improve our approach based on the feedback from domain experts. One particular issue is that our visualization shows only one feature at a time. Comparing features by switching between different visualizations is not the best solution, because users have to memorize considerable amounts of information in their short-term memory, which makes the comparison error-prone. More work is needed to be able to show multiple features simultaneously. Cognitive constraints and screen space limitations will be challenging factors to deal with. Hence, it also makes sense to extend our advanced features to be able to capture characteristics of multiple source features. Such higher-level features would accumulate much more information, but would remain straight-forward to visualize. Additionally, we could make use of large, high-resolution displays to physically extend the space for showing multiple features at a time.

VISUAL ANALYSIS OF RETINAL CHANGES WITH OPTICAL COHERENCE TOMOGRAPHY

SCIENTIFIC CONTRIBUTION: In this chapter, we contribute novel **VA** solutions to improve the visibility of diagnostically relevant information in retinal **OCT** data. With our concepts, the complex 3D retinal layers are reduced to 2D attribute maps. Using a unique combination of linked 2D and 3D views, the maps are visualized in their 3D spatial context with interesting parts highlighted. Beyond previous approaches, this ensures that the main characteristics of all layers in a dataset are actually visible and the user can browse through them. The journal article [27] included here is an extended version of [79], which received a best paper award.

AUTHOR CONTRIBUTION: The author of this thesis developed major parts of the presented **VA** approach, including conceptualization, implementation, and evaluation with domain experts, and wrote the main body of both publications [27, 79].

ABSTRACT: **OCT** enables noninvasive high-resolution 3D imaging of the human retina, and thus plays a fundamental role in detecting a wide range of ocular diseases. Despite the diagnostic value of **OCT**, managing and analyzing resulting data is challenging. We apply two **VA** strategies for supporting retinal assessment in practice. First, we provide an interface for unifying and structuring data from different sources into a common basis. Fusing that basis with medical records and augmenting it with analytically derived information facilitates thorough investigations. Second, we present a tailored **VA** tool for presenting, emphasizing, selecting, and comparing different aspects of the attributed data. This enables free exploration, reducing the data to relevant subsets, and focusing on details. By applying both strategies, we effectively enhance the management and the analysis of retinal **OCT** data for assisting medical diagnoses. Domain experts applied our solution successfully to study early retinal changes in patients suffering from type 1 diabetes mellitus.

ORIGINAL PUBLICATION: [27] — **M. Röhlig**, C. Schmidt, R. K. Prakasam, P. Rosenthal, H. Schumann, and O. Stachs. “Visual Analysis of Retinal Changes with Optical Coherence Tomography.” In: *The Visual Computer* 34.9 (2018), pp. 1209–1224. DOI: [10.1007/s00371-018-1486-x](https://doi.org/10.1007/s00371-018-1486-x)

FURTHER READING: Beyond the presented VA approach, we introduced two extensions to reveal the influence of acquisition parameters on the quality of retinal OCT data and to ensure the visibility of all retinal layers in an OCT dataset at once using a new compound map:

- ▷ **M. Röhlig**, M. Luboschik, R. K. Prakasam, O. Stachs, and H. Schumann. “Visually Analyzing Parameter Influence on Optical Coherence Tomography Data in Ophthalmology.” In: *EuroVis 2017 - Posters*. Ed. by A. P. Puig and T. Isenberg. The Eurographics Association, 2017. DOI: [10.2312/eurp.20171176](https://doi.org/10.2312/eurp.20171176)
- ▷ J. Stüwe, **M. Röhlig**, H. Schumann, R. K. Prakasam, and O. Stachs. “Visual Analysis of Abnormal Thickness of Intraretinal Layers.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2018

In addition, in the master’s thesis of A. Röhm (supervised by the author of this thesis), the interactive visualization and editing of transfer functions to emphasize interesting parts of the layer attributes and the OCT images were investigated:

- ▷ A. Röhm. “Visualization and editing of transfer functions for visual exploration of 3D retinal data.” MA thesis. University of Rostock, 2017

8.1 INTRODUCTION

In the context of ophthalmology, Optical Coherence Tomography (OCT) is a widely applied method to support the medical diagnosis of various ocular diseases. Based on high-resolution 3D imaging of the human retina, many medical conditions can be detected. Examples are diabetic retinopathy, age-related macular degeneration, and glaucoma induced retinal changes [16]. The diagnostic procedures involve searching for subtle retinal changes, analyzing multiple OCT datasets, and correlating findings with other clinical information. Especially for identifying early signs or investigating progressions, all of these steps are mandatory. Yet, due to the complexity of the data, already exploring a single volumetric OCT dataset can be difficult. This becomes even more complicated if details in multiple datasets have to be compared or different proprietary data formats have to be dealt with, e.g., in case of datasets from different OCT devices. Hence, managing and analyzing OCT data and relating them to other information are challenging and time-consuming tasks.

We present an interactive visual approach to address the peculiarities of OCT data: (i) data originating from different acquisition modalities and (ii) volumetric data of high spatial resolution combined with extracted information. Our contributions are:

UNIFIED DATA MANAGEMENT: We convert and structure OCT data from different sources into a common basis, enabling unified data access and management. We fuse that basis with other medical records and augment it with analytically derived information.

INTERACTIVE VISUAL ANALYTICS: We propose a novel visual design for presenting, emphasizing, and comparing different aspects of the data. Coordinated interaction facilitates exploration, selecting relevant subsets, and inspecting details on demand.

This paper covers an extended version of our initial work on Visual Analytics (VA) of OCT data [79]. Here, we introduce a new compact representation of multiple OCT datasets that enables visual comparison and eases the identification of subtle retinal changes. We integrate the visualization, interaction, and computation techniques in a VA workflow and describe how it helps to tackle the problem at hand. As another extension, we outline our collaboration with medical experts, which allowed us to understand their current analysis procedure, define design requirements, and test our solutions. The combined work steered the development and helped to address the experts' needs. Lastly, we broaden the discussion of our results, point out limitations, and indicate future research directions.

The remainder of the paper is structured as follows. The background and current analysis procedure are described in Section 8.2. Related work is reviewed in Section 8.3 and design requirements are listed in Section 8.4. The unified data management and the interactive visualization are detailed in Section 8.5 and Section 8.6. In Section 8.7 and Section 8.8, application results and user feedback are presented. Finally, limitations and future work are discussed in Section 8.9.

8.2 BACKGROUND

The structure of the multi-layered retina in the posterior segment of the eye cannot be examined with conventional ophthalmic methods. Starting with ultrahigh-resolution OCT-based retinal imaging [13], it is now possible to display different layers of the retina and provide unmatched detail and contrast images in an extraordinary way. This improves the diagnosis of pathologies, even at an early stage.

8.2.1 *Optical coherence tomography data*

State-of-the-art OCT devices are based on frequency domain optical coherence technology (also referred to as swept source OCT or spectral domain OCT). The OCT procedure is noninvasive and completely safe without light hazard. A laser beam scans the retina using dedicated scanning patterns in combination with active eye tracking. In this process, consecutive 1D reflectivity profiles (A-scans) are acquired. The

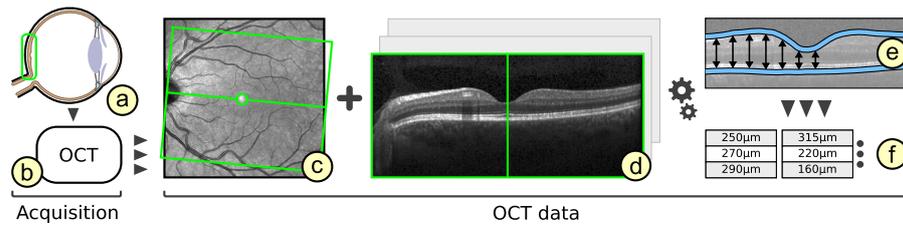


Figure 8.1: Common data components of volumetric OCT datasets. During acquisition, primarily the area around the macula and the optic disk (a) are captured with an OCT device (b). The resulting data consists of a fundus image of the interior surface of the eye (c) and multiple 2D image slices that form a 3D tomogram (d). Further processing allows to extract several retinal layers (e) and computed measurements (f), e.g., the layer thickness.

profiles contain information about the spatial dimensions and location of inner structures. Cross-sectional 2D image slices (*B-scans*) are obtained by laterally combining a series of these axial depth scans (*A-scans*). Multiple 2D image slices are then merged into a 3D tomogram. Commercially available systems currently have an approximate axial and lateral resolution of 3–7 $\mu\text{m}/\text{px}$ and 5–10 $\mu\text{m}/\text{px}$ respectively, covering an area of about 9 mm \times 9 mm at 2 mm scan depth. The actual image resolution depends on acquisition settings and differs between manufacturers. As an example, datasets from the Spectralis OCT (Heidelberg Engineering GmbH) can have a maximal resolution of 1536 equally spaced *B-scans* with 1536 px \times 512 px.

The raw volumetric image data are further processed using various image analysis techniques, including image registration and noise reduction [173]. Complementary segmentation algorithms are applied on the datasets to extract up to 11 retinal layers [85]. Optionally, measurements are computed to additionally characterize the data, e.g., thickness maps of the layers. Figure 8.1 illustrates common components of OCT data. Altogether, a typical OCT dataset contains one 3D tomogram composed of multiple 2D image slices, several extracted layers, and various computed measurements. In addition, one fundus image of the interior surface of the eye around the OCT acquisition area is usually recorded, e.g., via scanning laser ophthalmoscopy.

8.2.2 Current analysis procedure

Given the complexity of OCT data, it is challenging to identify subtle and localized changes of various abnormal conditions of the retina. Currently common analysis procedures involve both analyzing individual datasets and comparing multiple datasets.

INDIVIDUAL OCT DATASETS: are analyzed in search for retinal changes related to signs, symptoms, and other clinical information of

ocular diseases. This is usually done to diagnose subtle variations that cannot be captured with other examination methods or to obtain additional findings for investigating indications of a disease. To detect localized retinal changes, the ophthalmologists have to explore three main aspects of OCT data: the raw tomogram, the extracted layers, and associated measurements. This can entail high manual effort. For example, the experts are often required to inspect individual image slices or extracted layers one by one. This causes strain, particularly in case of high-resolution datasets that comprise hundreds of image slices. However, a high density of image slices is necessary for an exact assessment of certain retinal areas, e.g., to detect small macular holes [174] and to reduce errors of computed measurements [175].

Alternatively, current analysis procedures are either based on selective measurements or on excessive global data reduction with regard to both, space and information. In the first case, measurements are computed for specific points in the data. Depending on the analysis goal, the data are sampled at a known location, e.g., where changes are visible in the fundus image. As another example, several sample points may be aligned to form custom sample patterns, e.g., to obtain approximate layer profiles from the center of the fovea in a certain direction. In the second case, the tomogram or the extracted layers are subdivided into coarse sectors and for each sector aggregated measurements are retrieved. The locations and areas of these sectors are, for example, defined according to grids from the Early Treatment Diabetic Retinopathy Study (ETDRS) [89]. These grids consist of nine regions that subdivide central, pericentral, and peripheral concentric rings around the fovea in nasal, temporal, inferior, and superior direction. As a result, entire OCT datasets are represented by few values. While this reduces the amount of information that has to be examined, it also renders the analysis results spatially unspecific.

MULTIPLE OCT DATASETS: are analyzed to compare datasets either with each other, e.g., a baseline to a follow-up dataset, or to reference data, e.g., datasets of patients to data of a control group. Such analyses are more involved and go beyond the steps necessary for analyzing individual datasets.

In the first case, ophthalmologists typically have to open separate views for each dataset. The views are then individually adjusted to keep them consistent and to ensure that matching parts per dataset are shown. Searching for differences is realized by juxtaposing the views or by switching back and forth between them. Yet, manually coordinating and inspecting the separate views is perceived cumbersome.

In the second case, ophthalmologists have to rely on data reduction, e.g., based on ETDRS grids as described above. To start, aggregated measurements for all considered datasets are computed. These values are then exported for further statistical analysis. Yet, the available export

functionality is usually limited and so, dozens of values per dataset have to be copied by hand from the acquisition platform to the analysis software. Particularly in case of larger studies with hundreds of OCT datasets, this is time-consuming and error-prone. Moreover, the statistical outcome often only shows comparisons of mean values from larger areas, e.g., sectors of ETDRS grids. However, it does not disclose subtle local variations or the actual morphological changes.

In addition to analyzing OCT datasets, ophthalmologists have to relate their findings to other clinical data, e.g., information on data quality, patient-related records, and results of other examination methods. For this purpose, they typically have to rely on third-party software to show necessary information on demand. This further complicates the data analysis.

Our new interactive visual approach shows such information directly and simplifies identifying and processing abnormal retinal characteristics.

8.3 RELATED WORK

Current analysis procedures are based on a combination of commercial and non-commercial OCT software and general-purpose analysis software. Besides visualizing OCT data, statistical evaluation of pointwise measurements or aggregated values of certain regions is a common approach in ophthalmic research. Particularly, the thickness of retinal layers has been continuously used as a basis for investigations of early structural changes of the retina for a variety of diseases, e.g., diabetes mellitus [176–178] or glaucoma [91]. The statistical results are reported via basic displays, e.g., tables, box-plots, or color-coded ETDRS grid diagrams. Commercial and non-commercial software offer additional methods for visually analyzing OCT data.

8.3.1 *Commercial analysis software*

Commercial OCT software is predominantly distributed by OCT device manufactures. Currently, several major platforms are available, including software from Nidek, Optovue, Zeiss, Topcon, Heidelberg Engineering and others. Modern 3D retinal imaging has led to advances regarding the display and analysis functionality [179, 180]. Yet, commercial software typically matches the respective device’s capabilities and hence, software features often differ between tools. In general, they all support managing, analyzing, and presenting OCT data.

For managing OCT data, users have to rely on device-specific file formats and databases. This prohibits the exchange of OCT data. Moreover, it complicates analyzing and comparing datasets from diverse manufactures, e.g., due to deviations in resolution, numerical precision, registration, noise reduction, and image segmentation [174].

For analyzing OCT data, users take measurements based on OCT tomograms directly or based on prior extracted retinal layers. Yet, supporting algorithms and associated parameters are proprietary, and thus deviations between measurements may occur [181]. This makes comparisons of analysis results from different tools error-prone. Moreover, while some tools allow to relate measurements to a proprietary normative database, available methods generally do not support comparative analysis of data from different individuals or groups. Thus, such investigations have to be conducted based on exported data and external software.

For displaying OCT data, three types of presentations are common. First, the acquired 2D image slices are shown individually. This allows to view details but flipping through the images is time-consuming. Second, a fundus image is shown together with superimposed retinal layers. This helps to link the layers to the fundus but the layers can only be examined one at a time. Third, the OCT tomogram is shown in 3D. This provides an overview of the data but adjusting the visual representation, e.g., via navigation, is often limited. Also, combined 3D visualizations of the tomogram and the layers are typically not available, and thus spatial relationships might go unnoticed. Other drawbacks include inappropriate hard-wired color-coding, lacking consideration of data quality, and limited support for additional information.

Given those limitations, ophthalmologists are currently advised to use the same OCT device and associated analysis software for all examinations of their patients to ensure comparability of results [174].

8.3.2 *Non-commercial analysis software*

Besides commercial software, few approaches for visually analyzing retinal OCT data exist. The open-source software ImageJ can be used to analyze OCT images [182]. The Iowa Reference Algorithms are a research-oriented software that supports segmentation, visualization, and measurement of retinal layers [80]. Instead of extracted retinal layers, reflectivity profiles allow to characterize retinal conditions [82]. 3D visualization based on ray-tracing and artificial shadows shows subtle structures more distinctly but images can take multiple seconds to render [81]. Likewise, virtual reality can be employed to enhance spatial perception and facilitate an immersive data access [183, 184]. However, selecting and comparing parts of multiple datasets or relating them to other information is often not considered. Real-time 3D rendering has also been studied to enable online display of OCT tomograms during acquisition and to preselect relevant subsets for reduced storage costs [185, 186]. Yet, in-depth analysis of details still has to be done in a post-acquisition stage.

Aside from ophthalmology, visualization of OCT data is an inherent part of other fields as well. Examples include different biomedical ap-

plications, e.g., OCT-based exploration of cerebral vessel walls [187], or material science, e.g., detection of subsurface defect [188] or investigation of internal structures of bast fibers [189]. Although visualization proved to be useful in those domains, respective solutions are not directly applicable to retinal OCT data.

In summary, existing works offer different approaches for managing, analyzing, and presenting OCT data. Yet, each solution covers only a certain aspect. Our goal is to develop an integrated approach that (i) incorporates previous efforts regarding a common data basis, (ii) extends that basis with supplementary information, (iii) visualizes and emphasizes different aspects of the attributed data, and (iv) allows selecting and comparing relevant subsets. With the resulting flexibility, we are able to address the peculiarities of OCT data and support the retinal assessment.

8.4 REQUIREMENTS

Discussing with domain experts and analyzing existing approaches, we derived two lists of requirements related to: (i) managing OCT data and (ii) visualizing OCT data.

DATA-RELATED REQUIREMENTS: The data-related requirements reflect the experts' needs with regard to managing and processing OCT data. We derived the following list by talking with the experts about current limitations and by analyzing the computation and storage capabilities of OCT devices.

MANAGE DATA FROM DIFFERENT MANUFACTURERS (DR₁): The experts require support for OCT data from different manufacturers. Particularly, they ask for a unified access for datasets acquired using different OCT devices. So, providing a common data basis that incorporates all main aspects of OCT data, i.e., the raw tomogram, extracted retinal layers, and computed measurements, is a fundamental requirement.

COMMON AND OPEN DATA FORMAT (DR₂): Based on the common data basis, the experts want to make the data openly available. This is to allow easy exchange, storage, and further processing of datasets. Hence, the second requirement is to provide a common and open data format.

DERIVE AND INCORPORATE ADDITIONAL INFORMATION (DR₃): Incorporating and retrieving additional information is another request from the experts. This is to provide them with further details to help to characterize the condition of the retina. Especially, measurements based on extracted layers are necessary, as they allow to capture subtle changes which would be hard to identify by

looking at the raw OCT data alone. Therefore, augmenting OCT datasets with respective information has to be supported.

COHERENT ANALYSIS PIPELINE (DR₄): Altogether, the experts wish for a consistent set of data processing and computation methods. This is to ensure comparability of analysis results across datasets. These results also have to be reflected in the common data basis. Accordingly, the attributed data should provide a starting point for a coherent analysis and visualization pipeline.

VISUALIZATION-RELATED REQUIREMENTS: The visualization-related requirements correspond to the experts' goals when visually analyzing OCT data. We derived the requirements following a participatory design process. Together with the experts we made general design decisions and devised suitable visualization concepts:

EXPLORE AND RELATE DIFFERENT DATA ASPECTS (VR₁): To detect localized retinal changes, the experts have to explore the raw tomogram, the extracted retinal layers, and computed measurements. In this regard, established display concepts have to be considered (cf. Sect. 8.3.1). At the same time, the context of the three data aspects has to be maintained. Therefore, visualizing the aspects and interactively relating them needs to be facilitated.

SHOW RELATED INFORMATION ON DEMAND (VR₂): For assessing the condition of the retina, the experts have to relate findings to other information. So, besides incorporating supplementary information in the common data basis (cf. DR₃), showing that information on demand together with the OCT data is required.

DISCOVER AND INSPECT REGIONS OF INTEREST (VR₃): The experts require quick access to different parts of the data and to be able to inspect these parts in greater detail. This is particularly important in case of high-resolution datasets and subtle retinal changes. Hence, identifying, highlighting, and selecting regions of interest is another basic requirement.

COMPARE MULTIPLE DATASETS (VR₄): The experts need to compare findings in multiple OCT datasets. This involves both comparing individual datasets to each other and analyzing datasets in relation to reference data. Thus, we consider visually comparing multiple datasets as a necessary design requirement.

Based on the experts' demands, we incorporate a unified data management that addresses requirements DR₁, ..., DR₄ and develop an associated visualization tool to meet requirements VR₁, ..., VR₄.

8.5 UNIFY, STRUCTURE & ENHANCE DATA

Currently, no standardized description of OCT data exists. This hinders data processing and analysis. Moreover, the raw data alone may not provide enough information for definite diagnoses. In an initial effort to address these issues, we integrate a common data basis and enrich that basis with derived layer attributes.

8.5.1 Common data basis

No matter if clinical or research use is intended, it is often beneficial or at times even needed to take data from different OCT devices into consideration. The most problematic part with this is the fact that manufacturers only provide software for the data of their own OCT scanners. Consequently, there is no approved way of viewing, analyzing, or comparing data from different devices. Furthermore, there are only limited possibilities to implement own prototypes to perform such tasks, since software libraries are provided with exclusive licenses and incomplete data specifications.

To develop VA solutions for OCT data, it is essential to find a remedy for that issue. On this account, we adapt and extend previous work [181]. Particularly, we integrate a common data format and a library (Unified OCT Explorer (UOCTE)) that is capable of parsing data from the majority of current devices, e.g., Nidek, Eyeteq, Topcon, and Heidelberg Engineering (DR₁). Based on the UOCTE library, we map the parsed data into one common domain. The result is independent from constrained manufacturer-provided functionality and it opens up new possibilities for data unification.

To structure the data, we utilize a common data modeling language (Unified OCT Markup Language (UOCTML)), which accompanies the library (DR₂). UOCTML consists of an easily amendable Extensible Markup Language (XML) file and a set of binary files for fundus, tomogram, and layers. The resulting data format promotes data exchange. Moreover, we are able to enrich the data with derived information and other medical records. To this end, we provide an associated interface with dedicated processing capabilities. It allows either to retrieve pre-segmented layer data or to apply basic layer segmentation methods directly (e.g., [190]). In addition, it enables extraction of OCT metadata and computation of several layer attributes.

Together, the UOCTE library, the UOCTML data format, and the interface represent initial steps towards the development of common data standards in the domain of OCT data. The software and associated specifications are freely available [191]. Further technical details regarding our efforts to unify distinct OCT data sources are described in [181].

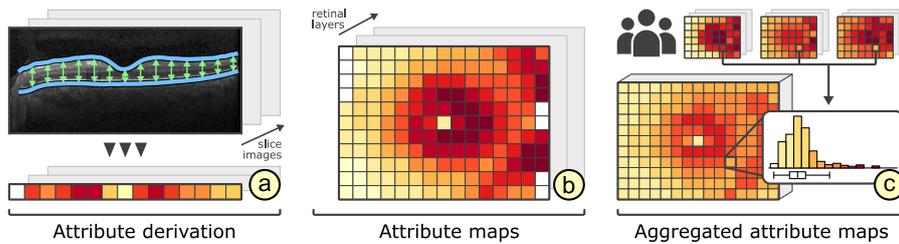


Figure 8.2: Derivation of layer attributes. An attribute is computed by analyzing the extracted retinal layers in each acquired image slice of a dataset (a). One value is assigned to every pixel along the horizontal image axes, e.g., by measuring the vertical distance between the upper and the lower layer boundaries. The values are combined into 2D attribute maps for each extracted layer and each attribute definition (b). Aggregated attribute maps are compiled by applying summary statistics to all associated pixels in the individual maps from multiple datasets (c).

8.5.2 Derivation of layer attributes

To facilitate the analysis of OCT data, we incorporate a set of derived attributes (DR_3). This is to inject meaningful information into the data that can help to characterize the condition of the retina. It allows us to condense complex OCT data down to information that is manageable and relevant. Moreover, the derived attributes capture even subtle and localized retinal changes. This way, they represent an effective starting point for the subsequent visual data analysis (cf. Sect. 8.6).

The attribute derivation comprises two steps that are carried out for each extracted layer. First, all points that define a layer's shape are enriched with derived attribute values. Figure 8.2 illustrates this process. This is to obtain attribute maps that characterize different layer properties, e.g., thickness maps. Second, for each attribute, a ranking of the layers is computed. A layer's rank is determined by analyzing the distribution of attribute values and their spatial locations in the attribute maps. The exact way how a ranking is computed can be controlled via different parameters. That is, the ranks are either defined by the total ratio of attribute values that are considered abnormal in a map or by the amount and size of connected abnormal regions. In addition, weights can be assigned to regard certain areas on a layer with abnormal characteristics more important than others. For instance, a targeted investigation may adapt the ranking to favor localized changes only in selected sectors of the ETDRS grid. This helps users to focus on layers with specific patterns.

In general, a single attribute alone will not suffice to capture all abnormal conditions of the retina. Therefore, we consider a collection of layer properties to target symptoms of different ocular diseases. In previous works mainly the thickness of the retinal layers has been studied to distinguish abnormal eyes from healthy eyes (e.g., [91, 177]). Based on

discussion with domain experts, we additionally incorporate several other properties. Examples are the curvature of the layers' shapes or the homogeneity of OCT values enclosed by layer boundaries. This may help to detect localized deformations, e.g., drusen induced in an early stage of age-related macular degeneration, and variations in reflectivity, e.g., small macular holes or fluid deposits caused by diabetic retinopathy, respectively. Deriving corresponding attribute maps enables a more sophisticated investigation of the different layer properties.

Based on the individual attribute maps, we also support the computation of reference data (Fig. 8.2 (c)). The user can select any group of OCT datasets and compute aggregated attribute maps for single or multiple attributes per layer. As the datasets may cover different retinal areas, the individual attribute maps are projected into a common spatial frame of reference, e.g., using the center of the fovea and of the optic disk as reference points. Each point of an aggregated map then stores summary statistics of the individual attribute maps, e.g., mean, range, standard deviation, and histograms. This provides a compact representation of multiple OCT datasets. Moreover, it facilitates overview and comparison, while being spatially more precise than analysis approaches based on global data reduction according to ETDRS grids (cf. Sect. 8.2). If necessary, such currently common abstractions of attribute values are easily derived from individual or aggregated attribute maps.

On top of the common data basis and derived layer attributes, we develop a VA tool with augmented visualization techniques for the attributed OCT data (DR_4), which will be specified in the next section.

8.6 VISUALIZE, EMPHASIZE, SELECT & COMPARE DATA

We aim at supporting users in visually analyzing OCT data and related information. For this purpose, we design a flexible visualization tool based on multiple coordinated views. Figure 8.3 shows an overview of the user interface. Our tool supports: (i) visualizing the data, (ii) emphasizing details, (iii) selecting subsets, and (iv) comparing multiple aspects.

8.6.1 Visualizing data

In order to facilitate a comprehensive analysis of attributed OCT data, we visualize different types of information. Our solution shows raw data together with extracted layers, considers derived information, takes the quality of the data into account, and facilitates the exploration of relationships (VR_1 and VR_2). To this end, we support four types of views: (i) a 2D top-down view, (ii) a 3D view, (iii) a 2D cross-sectional view, and (iv) a view for related information.

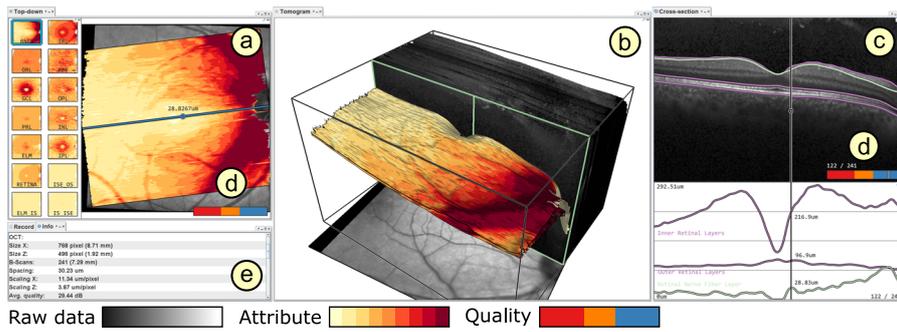


Figure 8.3: Overview of our prototypical VA tool. The user interface allows to add and arrange multiple linked views for visualizing different aspects of OCT data. Depicted are (a) a top-down view of the fundus overlaid with a derived attribute map, (b) a 3D view showing a volume visualization together with retinal layers as surfaces, (c) a cross-sectional view with extracted layers and a detail chart, and (e) a view for supplementary information. The signal strength of image slices is color-coded and shown together with associated legends (d).

THE 2D TOP-DOWN VIEW: provides an overview of derived attribute maps with regard to the interior surface of the eye (Fig. 8.3 (a)). The OCT acquisition area is visualized on top of a fundus image. Each acquired image slice is represented as a line. All extracted retinal layers are shown as thumbnails on the side, ordered according to the computed rankings. Optionally, the thumbnails are enlarged and superimposed over the fundus image. Derived attribute maps are color-coded using suitable and adjustable palettes [132]. The opacity of the overlay can be adjusted using a slider, e.g., to help to relate attribute values in the maps to noticeable structures in the subjacent fundus image. If required, colored ETDRS grids are shown in addition to the maps. This design extends existing displays, in that it allows to view attributes for all layers in one image without having to flip through them manually. Hence, layers with abnormal characteristics can be easily discovered.

THE 3D VIEW: shows the tomogram together with the extracted layers at a glance (Fig. 8.3 (b)). The raw data are visualized via direct volume rendering and the layers are displayed as surfaces. Blending both 3D presentations allows to relate the tomogram to the layers. On demand, the layer surfaces are color-coded based on derived attributes. This illustrates interrelations between attribute values and the layer shapes. For spatial reference, a fundus image can be vertically moved through the display. Showing the data in 3D allows for a faithful representation of internal retinal structures. This way, regions of interest can be precisely localized.

THE 2D CROSS-SECTIONAL VIEW: supports the investigation of details and the comparison of several layers at once (Fig. 8.3 (c)). The

acquired 2D image slices are depicted individually. The extracted layers are displayed as superimposed lines along the horizontal image axes. This allows to identify the exact layer profiles, to visually relate them, and to check for segmentation errors. A detail chart shows plots of derived attributes of a single layer or of multiple layers. The chart is positioned below the image slices and aligned horizontally to maintain the spatial context. Layers and plots are associated with unique colors. Visualizing the layer profiles and the attribute plots together facilitates a more precise analysis and direct comparison in addition to the color-coded visual representations in the other views.

Data quality is an important characteristic of OCT data [175, 192]. It has to be considered to prevent wrong conclusions from the data analysis. Therefore, in the first three views, missing values are either mapped to a special highlighting color to bring them to the users attention or to a background color to focus on certain parts of the data instead. Moreover, quality measures, e.g., the signal strength for 2D image slices, and associated legends are displayed in each view (Fig. 8.3 (d)). Indicating the data quality helps to differentiate between artifacts and actual retinal changes, and thus aids the assessment of findings.

THE INFORMATION VIEW: displays general properties of datasets, logs about selected values, or other patient-related records (Fig. 8.3 (e)). Depending on the type of information, different basic visualizations are available, e.g., tabular presentations or document viewers. In case multiple datasets are loaded, the view can show respective listings and help to assign datasets to the other views. This allows users to directly check the additional information together with the different perspectives of the shown OCT data, without having to rely on external software tools.

8.6.2 *Emphasizing data*

Showing important information in the context of entire datasets is essential for visually analyzing complex OCT data [193]. We support users in comprehending the data by emphasizing spatial relationships and characteristic data values (VR_3). For this purpose, we provide methods for both spatial and data-oriented accentuation. Figure 8.4 illustrates how these methods can be applied.

SPATIAL EMPHASIS: conveys the spatial properties of the visualized data. This is crucial in the 3D view (Fig. 8.4 (a)). We consider two illumination methods: a local technique based on a directional lighting model and a global approach based on ambient aperture lighting [130]. Directional lighting illustrates the general shape of the volumetric data and enhances the perception of small-scale spatial structures. Ambient aperture lighting highlights details and relationships between neighbor-

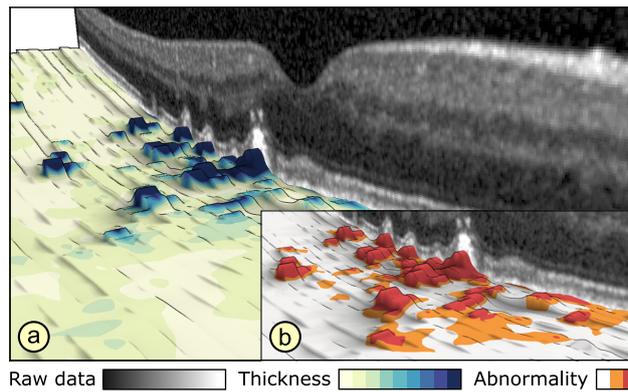


Figure 8.4: Selecting and emphasizing data. The 3D view shows a selected layer, colored by its thickness (thin: green, thick: blue). Illumination and enhanced edges facilitate the spatial perception of the layer's shape (a and b). Abnormal thickness according to predefined thresholds is emphasized via spatial and data-driven selections (b). The small scattered elevations (orange, red) indicate the presence of drusen, a sign for age-related macular degeneration.

ing parts on the layer surfaces. Stylization further emphasizes certain spatial properties. On the one hand, enhancing edges makes it easier to distinguish spatial features of the tomogram and the layers. On the other hand, customized toon-shading generates a continuous abstraction of the layers that also reinforces the perception of depth [131]. Optionally, stereoscopic rendering can be activated to further facilitate 3D spatial perception.

DATA-ORIENTED EMPHASIS: highlights values of interest and steers the user's attention. For this purpose, we utilize three visual variables: color, transparency, and blur [194]. In the 2D and 3D views, visual prominence is controlled via an interactive transfer function editor. The editor modifies the visual encoding to either strengthen or attenuate the influence of values in the rendered images. For instance, invalid values may be filtered out by lowering their opacity. Other value ranges can be assigned to special color palettes to compare them in the different views. The views are linked with the editor and automatically update according to user input. Figure 8.4 (b) shows an example of modified colors for identifying abnormal attribute values. To further steer the user's attention, data-oriented emphasis is generated via customized depth-of-field rendering [195]. Regions with values of interest are depicted sharply, whereas their surroundings are blurred in the visualization. This helps to focus on details and to maintain the context at once.

Spatial emphasis and data-oriented emphasis can also be combined. This is valuable for relating certain data values to spatial properties. For example, in the 3D view, abnormal attribute values are highlighted using color. Applying enhanced edges illustrates their relationship to the

shape of the layers, while still being able to precisely read off different color shades. In addition, toon-shading or depth-of-field rendering can be activated to support the identification of affected regions.

8.6.3 *Selecting data*

To enable the exploration of large OCT datasets, the data have to be reduced to relevant subsets. We support this via coordinated selection techniques (VR₃). Subsets are interactively defined both spatially and data-driven. Figure 8.4 exemplifies the application of both techniques.

SPATIAL SELECTIONS: enable users to specify regions of interest. We integrate various selection methods based on points and geometric shapes. Individual points can be selected in all views to show their assigned values via tooltips. In the 3D view, tomogram and layer selections are realized via interactive clipping geometry, including planes, spheres, or layer surfaces. For example, clipping the tomogram via adjustable planes helps to relate selected parts to the layers (Fig. 8.4 (a)). In the 2D views, polygon selection allows users to interactively set multiple corners to define spatial regions of interest. Users can choose to apply such selections to single image slices and layers, or to groups of them. This way, inspecting the same regions in different parts of the data is possible. All spatial selections can be expanded or reduced using binary operations. The provided selection methods aid the specification of precise boundaries of spatial regions, and hence are an effective way to reduce the data directly via their visual representation.

DATA-DRIVEN SELECTIONS: permit users to specify value ranges of interest. For raw OCT data, the transfer function editor facilitates selecting one or multiple value ranges. To support such selections, a histogram is shown for reference. For derived attributes, users can choose which attributes are to be mapped onto the layer representations (Fig. 8.4 (b)). This helps to focus on those data characteristics and to relate them to their spatial context. In addition, the detail chart in the 2D cross-sectional view enables selecting value ranges by brushing parts of the attribute plots. Data-driven selections allow to capture all regions with specific data characteristics across entire datasets, which would be hard to dissect via direct spatial selections in the views.

Specified selections can be applied in two ways: locally, i.e., solely in one view, or globally, i.e., by automatically propagating them to interlinked views. Moreover, selections may be applied in combination with the emphasizing methods to adapt the visual representations. This allows to show subsets either exclusively or highlighted in context of other data. For example, global selections support linking and brushing. Selecting a point in one view shows respective cursors in all other views. To browse through the data, the cursor in the top-down view is

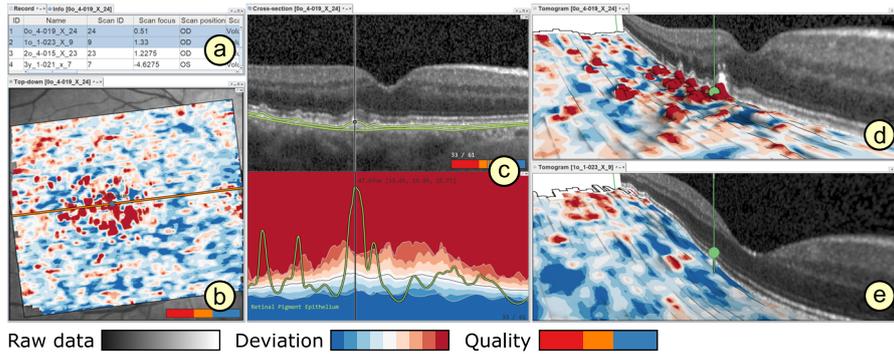


Figure 8.5: Comparison of multiple datasets. Two datasets are selected in the information view (a). The first dataset is assigned to views (b, c, d) and the second dataset to view (e). Deviations of layer thickness from a confidence interval of reference data are color-coded using lighter colors for values close to the reference mean and darker colors for values outside the interval. The upper and the lower 3D views (e and d) are linked to show matched parts of each dataset.

moved. The 3D view and cross-sectional view are updated accordingly and continuously display matched image slices and clippings of the tomogram. This allows to quickly access and relate different parts of the data.

8.6.4 Comparing data

Subtle retinal changes can often only be revealed by relating two or more datasets to each other. Hence, besides displaying one dataset of one individual at a time, our tool also allows to visually compare multiple intraindividual datasets, e.g., from follow-up examinations, and interindividual datasets, e.g., from cohort studies (VR₄). We support this via a combination of superposition and explicit encoding of deviations, and juxtaposition of interlinked views (cf. [196]). Figure 8.5 shows an overview of the user interface for visual comparisons.

ENCODING DEVIATIONS: enables comparing one dataset to reference data. This is possible by adapting the visual representation of extracted layers in the 2D views and 3D views (Fig. 8.5 (b) and (d)). On demand, the coloring of the layers is switched to show derived attributes in relation to summary statistics stored in aggregated attribute maps (cf. Sect. 8.5.2). Different properties of the attribute maps can be chosen for the encoding. Examples are deviations to the mean or to a confidence interval of the reference data. The resulting deviation maps are encoded using diverging color palettes [132]. Optionally, properties of the reference data can be inspected on their own or in relation to other reference data.

The chart in the 2D cross-sectional view shows deviations for selected image slices in detail (Fig. 8.5 (c)). Attribute profiles are depicted as superimposed line plots on top of the mean, the minimum, the maximum, and percentiles of the reference data. The chart's background is either colored according to a confidence interval to match the encoding in the other views, or based on a histogram to show the distribution of the reference data.

Switching to this encoding allows to countercheck findings made while analyzing a dataset on its own. Moreover, subtle and localized changes can be easily detected and compared to raw data and extracted layers with high spatial specificity.

JUXTAPOSING VIEWS: supports the simultaneous comparison of several individual datasets (Fig. 8.5 (d) and (e)). The user interface allows to dynamically add and freely arrange instances of the four types of views. Using the information view, listed datasets can be selected and assigned to any other view instance simply via drag and drop (Fig. 8.5 (a)). Those view instances then automatically update and show the respective OCT data directly. This allows to quickly switch between datasets. At the same time, specified selections and emphases are maintained to facilitate comparison of data aspects. To compare multiple datasets at once, one view instance is created per dataset. Juxtaposing and linking the view instances ensures that matching parts of the data are shown during navigation. All navigation is smoothly animated to prevent sudden changes in the rendered images. For example, different parts of the data can be explored in the 3D views via an interactive virtual camera and in the 2D views via zooming and panning. Relating the shown data back to the selected datasets is supported via special highlighting colors.

Applying such a view layout helps to analyze several data aspects in multiple datasets. Based on the common data basis this is even possible for datasets acquired via different OCT devices. This provides a distinct advantage compared to existing software tools that are based on fixed layouts, and thus only allow to visualize one aspect of one dataset at a time.

Figure 8.5 shows an example of comparing datasets of two patients to data of a control group. In the 2D top-down view and the 3D views the layers are colored by setting the patients' attribute values in contrast to the 95% interval of the control data. Light colors mark areas on the layers with thickness values close to the reference mean, whereas dark colors indicate thinning or thickening outside the reference interval. In addition, one 3D view is added for each dataset. The upper 3D view is adjusted to show a region of interest. These view settings are dragged onto the second 3D view below. The second view automatically adapts according to the applied changes and shows the same region and encoding. This way, abnormal retinal changes can be easily detected in each

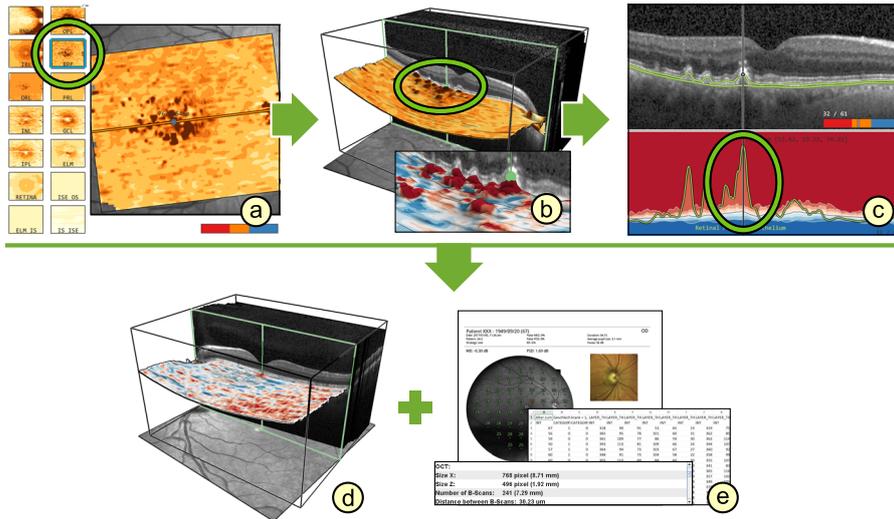


Figure 8.6: Exemplary analysis workflow. An OCT dataset is gradually reduced to relevant information by selecting a layer in the top-down view (a), defining a region of interest in the 3D view (b), and inspecting details in the cross-sectional view (c). The obtained findings are then compared to further datasets (d) or to results of other examination methods (e).

dataset and subsequently compared to each other via the juxtaposed views.

8.6.5 Analysis example

Our prototypical visualization tool can be applied in various ways. [Figure 8.6](#) illustrates an exemplary analysis workflow. First, the overview of derived attribute maps in the top-down view helps to identify and select a retinal layer with abnormal characteristics ([Fig. 8.6 \(a\)](#)). Enlarging the attribute map over the fundus image reveals centralized abnormal changes in the selected layer. Second, visualizing the layer in the 3D view allows to compare the attribute values to the layer's shape and to the volume rendering of the raw tomogram ([Fig. 8.6 \(b\)](#)). This way, a region of interest can be localized and selected. Third, to analyze the region in detail, individual image slices and respective attribute plots are shown in the cross-sectional view for the specified region ([Fig. 8.6 \(c\)](#)). Comparing the region to reference data is possible by loading an aggregated attribute map. Deviations of attribute values are emphasized by switching the color-coding and applying further stylization methods. Finally, the findings can be directly related to additional OCT datasets ([Fig. 8.6 \(d\)](#)), e.g., a follow-up dataset, or to results of other examination methods ([Fig. 8.6 \(e\)](#)). Applying such an analysis procedure promotes an informed interactive reduction of complex OCT data to clinical relevant information.

8.7 APPLICATION

In this use case, we applied our solution in cooperation with domain experts to study if retinal changes in pediatric patients suffering from Type 1 Diabetes Mellitus (T1DM) can be visualized in comparison with healthy control subjects of the same age. Interestingly, none of the patients showed signs of diabetic retinopathy in results of other examination methods. In fact, the principal goal of this study was to identify changes of the retina that occur in a very early stage of a primary disease, e.g., diabetes mellitus. These changes are precursors of secondary complications that may lead to severe structural and functional defects. Secondary complications, e.g., diabetic retinopathy, are often only diagnosed in an advanced state once symptoms are obvious. However, the high sensitivity of OCT and the VA of OCT data support detecting subtle and localized changes, and hence may contribute to an early diagnosis and immediate intervention.

The study data were acquired via OCT examination of two groups of subjects. The first group consists of 19 children with T1DM (average 14.2 years of age and 5.25 years duration of T1DM) and the second group of 30 healthy controls with matched sex and age (average 14.3 years of age). For each subject, one volumetric OCT dataset of the macula was selected for further investigation (49 datasets in total). In a pre-process, aggregated attribute maps were computed for both groups. Individual attribute maps for all layers in the datasets were derived on the fly during visualization.

We jointly applied our tool to analyze the study data. In addition, a currently common analysis procedure based on statistical evaluation of ETDRS grids was performed (cf. Sect. 8.2). Figure 8.7 summarizes a part of the obtained results.

RELATING PATIENTS TO CONTROLS: We started by visually comparing the aggregated attribute map of the patient group to the map of the control group (Fig. 8.7 (a)). Deviations of the mean values in both maps were color-coded. The resulting visualization demonstrated that patients exhibit thinning primarily in pericentral ETDRS sectors around the fovea for several retinal layers, e.g., retinal nerve fiber layer, ganglion cell layer, and inner plexiform layer. These first insights are in accordance with the statistical analysis results and findings of previous work [197]. The experts stated that our approach eases the detection of such patterns and that it provides more spatial detail. Hence, hypotheses can be generated quickly without having to rely on external software.

Beyond group versus group analyses, also individual patients were compared to the reference data of the control group. This helped to explain the visualization of aggregated data and the results of the statistical evaluation. Apart from common patterns, different thickness

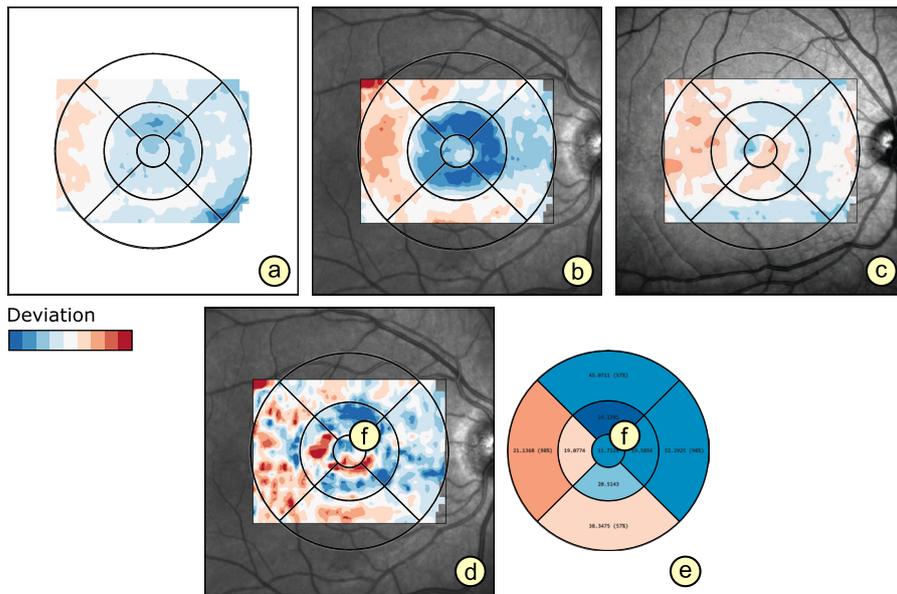


Figure 8.7: Selected results from analyzing data of a clinical study in our use case. The layer thickness of patients with T_1DM is compared to the 95 % confidence interval of a healthy control group. The deviation maps show thinning (blue) and thickening (red) of the overall retinal thickness for the mean of all patients (a) and for two individual patients (b and c). Borders of an $ETDRS$ grid are superimposed for spatial reference. The first map reveals pericentral (middle $ETDRS$ ring) thinning across all patients (a). This is reflected by one of the patients (b), whereas the other patient exhibits slight thickening in the same area (c). Comparing a deviation map (d) to an $ETDRS$ grid (e) shows that local pericentral thickening (f) of the retinal nerve fiber layer can be accurately captured by the map but not with the grid due to the applied data reduction.

patterns between patients were also identified. Figure 8.7 (b) illustrates the data of one selected patient. The visualization clearly shows pericentral areas with reduced layer thickness. This confirmed prior assumptions and is in line with the group versus group analysis. In fact, by consulting related clinical information, the patient could be assigned to a subgroup of severe cases (long duration of T_1DM and high blood sugar values). Observing the actual patterns of spatial changes and relating them to other aspects in the different views was new to the experts. For instance, some patients showed asymmetric patterns, e.g., temporal thickening and nasal thinning, in the peripheral ring of the $ETDRS$ grids. Moreover, the experts discovered cases with less thinning or even marginal local thickening in pericentral areas (Fig. 8.7 (c)). Most of these patients were subsequently categorized as mild cases (short duration of T_1DM and low blood sugar values). This implies that even in mild cases, noticeable patterns can be detected with our tool. Thus, our VA approach is effective in visualizing early retinal changes in pediatric patients with T_1DM and no diabetic retinopathy.

DATA SCREENING: The attribute maps allow to capture more information compared to [ETDRS](#) grids. This supports the data analysis in several ways. On the one hand, the experts can consider local variations during retinal assessment. Such information is often lost when looking at [ETDRS](#) diagrams alone ([Fig. 8.7](#) (d) and (e)). On the other hand, errors in the data can be detected easily. For example, at first the experts had identified local variations in attribute values for some datasets. However, these changes could later be traced back to segmentation errors. Visualizing the data quality of respective image slices and inspecting the layer boundaries for the affected areas helped the experts to come to this conclusion. Accordingly, an additional data screening pass was conducted with our tool. Some datasets were replaced by substitutes and others were marked for correction of segmentation errors. This way, only datasets of sufficient quality were included in the study.

The described results illustrate how our approach can be applied to gain insights into [OCT](#) data. Following, we outline how we designed and tested our solution together with domain experts.

8.8 USER FEEDBACK

Our solution is the result of a participatory design process starting from prior work [[79](#)] and [[181](#)]. We cooperated with a group of domain experts, including ophthalmic research scientists and ophthalmologists who deal with the treatment of retinal diseases. Throughout the development, we had close contact with primarily two experts. We jointly identified challenges and specified respective requirements (cf. [Sect. 8.4](#)). Their informal feedback helped us to design suitable visualizations and associated interaction techniques. While some design decisions were driven by the addressed data and task, e.g., the common data format and derivation of attributes, others were inspired by the domain experts, e.g., the comparison of datasets to reference data. The cooperation was of mutual benefit. We could build upon the experts domain expertise and in turn, they benefited from our solution as it provided them with new insights into the data of their studies (cf. [Sect. 8.7](#)).

To ensure maximal practical relevance of the achieved results, we assessed our solution in repeated demonstration and feedback sessions together with the experts. The sessions were run using a pair analytics approach [[29](#)]. They were scheduled after each major development cycle and held at least once a month to review new functionality. One session typically lasted between one and two hours depending on the complexity of the added methods and associated analysis tasks. Mainly one visualization expert and at least one or two medical experts participated. The visualization expert played the role of the driver of the [VA](#) tool, whereas the medical experts played the role of the navigator based on their contextual knowledge in the application domain. The data and

analytical tasks were consistent between sessions and correspond to our use case. Selecting the use case, i.e., currently relevant tasks and familiar datasets, to test our solutions created a naturalistic setting for observations and reasoning. The tasks involved both free exploration, e.g., discovering abnormal characteristics in patients, and targeted investigations, e.g., retrieving values from specific deviation maps in group versus group comparison. Session protocols were recorded and exchanged after each session. In addition to the results outlined in the previous section, we obtained feedback regarding the general design and functionality provided by our tool.

During the discussions, the experts stated, that they liked the VA tool because it allowed them to explore and relate different aspects of the data. They appreciated the consideration of established display concepts, e.g., the top-down view and the cross-sectional view. At the same time, the integration of our extensions was pointed out as meaningful, e.g., the overview of all extracted layers or the view for supplementary information. The experts also reassured us that the provided emphasis and selection methods are indeed helpful for inspecting different subsets in detail.

In earlier tests with initial prototypes of our tool, they had mentioned that the provided controls required some time to get accustomed with. Based on this feedback, we selected several interface components for a revision, including the controls of the virtual camera in the 3D view and control panels with associated display options. To come up with alternative solutions, we referred to available design guidelines (e.g., [198]) and prototyped new interaction concepts. For example, with respect to the 3D camera, we devised a new control mode that is solely based on mouse input, whereas before mouse and keyboard input was required. In addition, smooth animated transitions between camera movements were added to help users maintain their mental model during navigation. Regarding the interface, indirections were minimized by using direct manipulation where possible and by showing control panels with advanced settings only on demand. The reworked design was presented and discussed in subsequent feedback sessions together with the experts and again further refined. This iterative process improved the user experience substantially.

As a major advantage, the experts identified the ability to compare multiple and possibly differing datasets. The flexible interface and the interlinked views helped them to quickly switch between datasets and to show related information when necessary. For the analysis of larger studies, they considered the visualization of aggregated attribute maps to be particularly useful. Utilizing these maps as reference data to compare individual datasets to groups of datasets helped them to detect even subtle and localized changes. In this regard, they approved to have both overviews, e.g., thumbnails of deviation maps in the top-down view, and detail views, e.g., the detail chart showing confidence intervals and

histograms for image slices, at their disposal. Overall, they concluded that reducing the manual analysis effort and being able to obtain results with higher accuracy compared to the current analysis procedure are great benefits. In fact, for the first time, they were able to perform such analyses for volumetric macular OCT datasets. This is an important step towards finding characteristic retinal changes that indicate early symptoms of certain diseases.

Applying pair analytics to test and reflect on the introduced enhancements has been most fruitful so far, as it combines the expertise of visualization experts and domain experts. Findings in the data could be directly explained and new views were added or adapted on the fly to investigate further details. Applying our tool in this way triggered discussions and sparked new ideas for further improvements, which we will continue to share with the ophthalmic community (e.g., [98]). Besides, the experts recently asked us if we could provide them with a dedicated standalone version of our tool, so they are able to conduct further tests and apply it in clinical practice on their own. Accordingly, our tool will be used as a basis for further extension, fine-tuning, and evaluation together with them.

8.9 DISCUSSION AND CONCLUSION

We presented an interactive visual approach for managing, analyzing, and presenting OCT data. A unified data basis incorporates data from various devices and derived information. A VA tool supports exploration and emphasis of different data aspects, and selections of relevant subsets in interlinked views. Multiple datasets can be compared to each other and to reference data. Our approach constitutes a systematic enhancement of existing work, and hence can be a useful aid for retinal assessment using OCT.

DIAGNOSTIC SUPPORT: From an ophthalmic perspective, VA of retinal OCT data may be utilized in various ways. With regard to clinical practice, it contributes in confirming retinal changes in relation to other clinical information and in getting a more complete idea of the retinal condition of individual patients. Therefore, the actual diagnosis and treatment decisions are often made on an individual basis only after detailed clinical and systemic investigation in addition to OCT analysis. With regard to ophthalmic research, VA eases the evaluation of cohort studies and provides results with higher spatial specificity. This may lead to new insights into how certain pathologies affect the structures of the retina. The goal is to identify meaningful characteristics that can later be used to define biomarkers. In this regard, experimental and prospective studies, as presented in our use case, play an inherent role of ophthalmic research. We highly recommend additional prospective longitudinal studies to observe the structural changes over a period of

time and to relate these changes with the functional aspects of the patients. To support such investigations, both qualitative and quantitative analyses of OCT data have to be integrated. Hence, we plan to extend our approach to support time-oriented data analysis and targeted statistical tests.

DATA SUPPORT: Our data management is capable of parsing and converting data from different devices. By making the software sources freely available, we ensure its dissemination to an interested audience. This is a first step towards developing device-independent data standards. Although it remains difficult to cover all possible data sources, we will continue our efforts to support file formats of major manufacturers and retrieve respective metadata. In addition, we plan to directly relate the data to further information for characterizing the retina. This includes results from other examination methods, e.g., angiography or microperimetry, and diagnosis-related annotations. For this purpose, the common data basis has to be extended to adequately register, store, and exchange such information. Regarding the extraction of retinal layers, our data interface supports two methods. On the one hand, we retrieve pre-segmented layer information from proprietary data files of the OCT device manufacturers. On the other hand, if such information is not available, we provide basic segmentation methods [190] or use free software (e.g., [199]). Yet, this does not emulate the functionality of commercial software to its full extent. More work is required to integrate state-of-the-art segmentation algorithms (e.g., [85]). Directly incorporating advanced methods into our tool and disclosing related parameters, will help to reduce deviations between measurements and facilitate the comparison of analysis results from different datasets. In this regard, it remains to be studied to what extent existing algorithms can be adapted to address manufacturer-related differences in the data, e.g., varying numerical precision or noise characteristics.

ANALYTIC SUPPORT: We aid the VA of OCT data via various computational methods in our tool. This includes the derivation of layer attributes, the ranking of the layers, and the automated alignment of interlinked views. We plan to enhance our approach based on further feedback from domain experts. One particular request was to extend the support for analyzing multiple datasets. Interesting patterns could be automatically detected and subsequently suggested to the user. For example, similarly to the computation of layer rankings, the experts want to be able to search for matching subsets across datasets. One possible approach would be to extract and compare features based on attribute maps and reference data. An interesting open question is how we can find meaningful feature definitions that address specific analysis tasks, e.g., the identification of certain symptoms. This requires assistance in evaluating how well a feature captures certain characteristics

and how the feature extraction process can be steered to obtain robust results.

VISUAL SUPPORT: The described examples demonstrate that our approach is suitable for analyzing multiple datasets at once. Visual comparison is supported either via explicit encoding of deviations to reference data or by assigning different datasets to juxtaposed views. However, adding multiple view instances only makes sense for up to two or three datasets at a time. Comparing numerous datasets by switching between multiple visualizations may not be the best solution. Hence, we plan to extend our tool to provide dedicated overview visualizations. An effective design would allow to directly compare the data of an entire study. Cognitive constraints and screen space limitations will be challenging factors to deal with. In a first prototype, we tested a small multiple design, which shows synchronized top-down views for several datasets in one image. This can be a feasible approach for tens of datasets. In the future, we will investigate alternative encodings that are scalable enough to handle even more datasets.

USER SUPPORT: We ascertained the general utility of our solutions in first tests with domain experts (two being authors of this paper). To improve our design, we plan to integrate guidance for different diagnostic tasks. In this context, specifications of dedicated workflows and further evaluations of our tool will become necessary.

ENHANCED GRID-BASED VISUAL ANALYSIS OF RETINAL LAYER THICKNESS WITH OPTICAL COHERENCE TOMOGRAPHY

SCIENTIFIC CONTRIBUTION: In this chapter, we contribute novel **VA** solutions to reduce the complexity of multiple retinal **OCT** datasets. In particular, we introduce concepts for the grid-based extraction of high-level features over local regions of the retinal layers. A new **VA** tool and a complementary **VA** procedure enable the comparison of these features between patients and controls with higher spatial specificity and lower analysis effort than previous analysis approaches. The journal article [28] included here is an extended version of [88], which received a best student paper award.

AUTHOR CONTRIBUTION: The author of this thesis developed major parts of the presented **VA** approach, including conceptualization, implementation, and evaluation with domain experts, and wrote the main body of both publications [28, 88]. The individual contributions of all authors are listed in [28]. The implementation of the grid representations is partly based on the bachelor's thesis "Attribute-driven visualization of retinal OCT data" by J. Stüwe [200], which was substantially supervised by the author of this thesis.

ABSTRACT: **OCT** enables high-resolution 3D imaging of retinal layers in the human eye. The thickness of the layers is commonly assessed to understand a variety of retinal and systemic disorders. Yet, the thickness data are complex and currently need to be considerably reduced prior to further processing and analysis. This leads to a loss of information on localized variations in thickness, which are important for an early detection of certain retinal diseases. We propose an enhanced grid-based reduction and exploration of retinal thickness data. Alternative grids are computed, their representation quality is rated, and best fitting grids for given thickness data are suggested. Selected grids are then visualized, adapted, and compared at different levels of granularity. A **VA** tool bundles all computational, visual, and interactive means in a flexible user interface. We demonstrate the utility of our tool in a complementary analysis procedure, which eases the evaluation of ophthalmic study data. Ophthalmologists successfully applied our solution to study localized variations in thickness of retinal layers in patients with diabetes mellitus.

ORIGINAL PUBLICATION: [28] — M. Röhlig, R. K. Prakasam, J. Stüwe, C. Schmidt, O. Stachs, and H. Schumann. “Enhanced Grid-Based Visual Analysis of Retinal Layer Thickness with Optical Coherence Tomography.” In: *Information* 10.9 (2019), p. 266. DOI: [10.3390/info10090266](https://doi.org/10.3390/info10090266)

FURTHER READING: The initial ideas for the representation of retinal layers by different grids were developed as part of the bachelor’s thesis of J. Stüwe under the supervision of the author of this thesis:

- ▷ J. Stüwe. “Attribute-driven visualization of retinal OCT data.” BA thesis. University of Rostock, 2018

Based on the presented VA approach, in the master’s thesis of J. Stüwe (supervised by the author of this thesis), an interactive similarity search for abnormal regions of the retinal layers in the OCT data of numerous patients was investigated:

- ▷ J. Stüwe. “Interactive similarity search for attribute maps of intraretinal layers.” MA thesis. University of Rostock, 2019

9.1 INTRODUCTION

Optical Coherence Tomography (OCT) [1] is a widely-used noninvasive technique to capture high-resolution 3D images of retinal substructures. Ophthalmologists analyze the resulting data to understand a variety of retinal and systemic disorders [16]. This includes common eye diseases, e.g., age-related macular degeneration [4], diabetic retinopathy [5], or glaucoma [6], other pathologies with ocular signs, e.g., multiple sclerosis [7], and retinal vascular disorders [202]. Particularly, the thickness of segmented retinal layers is used to assess the condition of the retina. However, derived thickness data are complex, as one thickness value is computed for every single point of each retinal layer. On the one hand, this enables a spatially precise judgment of the layers. On the other hand, the large amounts of thickness values are difficult to deal with, and a summary of thickness changes in larger retinal areas is missing. Hence, ophthalmologists rely on considerable data reduction prior to further processing and analysis of the data.

Established reduction methods for retinal thickness data are based on retinal grids. These grids are used to spatially divide the retina into few large regions and to derive aggregated thickness measures for each region. This helps to get a quick overview of the layers’ thickness in anatomically predefined areas. Moreover, it drastically reduces the amount of information to be analyzed, particularly in larger studies with dozens or hundreds of OCT datasets. Yet, due to the spatial aggregation, information loss may occur. This is because small and localized

variations in thickness are not always reasonably represented via aggregated values of large regions. Capturing such information is however mandatory for detecting early signs of certain diseases or investigating progressions.

We aim at supporting ophthalmologists in their grid-based analysis of retinal layer thickness. We propose an enhanced data reduction scheme together with a Visual Analytics (VA) tool for the exploration of alternative grids. Our approach helps to strike a balance between obtaining a compact grid representation of thickness data and being able to capture more relevant information of retinal layers. Our contributions are:

ENHANCED GRID DESIGN: We propose an adaptive grid-based reduction of retinal thickness data. New grid layouts are derived based on radial and sector-wise subdivision of well-established grids. The representation quality of alternative grids is rated and best options are suggested to the user.

GRID-BASED VA: We develop a flexible VA tool for grid-based data exploration. Grids are interactively refined, compared to other grids, and cell-related details are shown on demand. A complementary procedure eases the analysis of ophthalmic study data.

EVALUATION OF OPHTHALMIC STUDIES: We apply our tool to investigate localized variations in retinal layer thickness in two cross-sectional studies with patients suffering from diabetes mellitus. The main findings are summarized and compared to results of current analysis procedures.

This paper constitutes an extension of our previous work [88]. We continue our collaboration with ophthalmologists, which allows us to define design requirements and to test our solutions in comparison with their current analysis practices. We integrate our enhanced grid design and exploration in a flexible VA tool and describe how an associated analysis procedure helped ophthalmologists to evaluate data of two ophthalmic studies. We summarize gathered feedback, broaden the discussion of results, and exemplify other application scenarios in ophthalmology.

After describing the medical background, related work, and current analysis procedures in [Section 9.2](#), we present our enhanced grid design in [Section 9.3](#). Based on new grid layouts, we introduce our interactive exploration of retinal thickness data in [Section 9.4](#). We then propose a matched VA tool and analytic process in [Section 9.5](#), and apply them to ophthalmic study data in [Section 9.6](#). Lastly, we summarize our work and discuss future research directions in [Section 9.7](#).

9.2 BACKGROUND

Our work is motivated and driven by advances in the detection of retinal diseases [17]. Particularly, the dynamic development of OCT technology with respect to image quality offers a unique possibility of differentiating and precisely measuring retinal substructures [2, 13, 14], close to an in-vivo ‘optical biopsy’ of the retina [15]. Modern OCT devices enable high-resolution imaging of biological tissue in the micrometer range with an approximate axial and lateral resolution of $2\ \mu\text{m}$ [203] or even almost submicron lateral resolution [204]. The commercially available system used in our work (Spectralis® HRA + OCT, OCT2 Module, Heidelberg Engineering GmbH, Heidelberg, Germany) is based on frequency domain optical coherence technology (also called spectral-domain or Fourier-domain OCT). It captures Cross-sectional Tomographic Depth Scans (B-scans) of the retina at high speed (axial scan rate of 85 kHz) with an axial and lateral image resolution of $3.5\ \mu\text{m}/\text{px}$ and $6\text{--}11\ \mu\text{m}/\text{px}$, respectively. Typically, multiple B-scans are combined to obtain volumetric 3D images and further processed, including noise reduction [173] and segmentation of retinal layers [85]. This provides a straightforward method of assessing even subtle retinal changes and accurately monitoring the progression of a disease. Consequently, OCT examinations are now a standard of care in clinical practice and an essential tool in ophthalmic research [8]. Based on the data, ophthalmologists aim at performing both:

- *Patient-specific* assessments of the retinal condition of individuals
- *Group-specific* evaluations of experimental and prospective studies

In this context, they often need to compare multiple intraindividual datasets, e.g., follow-up examinations of a single patient, and interindividual datasets, e.g., examinations of patients in relation to normative data of controls. Yet, the data analysis can be complex and the available analysis methods differ between existing software tools. In this regard, our work is related to the VA of retinal OCT data in general, and to the representation of retinal layer thickness via grids in particular.

9.2.1 Visual analysis of OCT data

Current analysis procedures are based on a combination of commercial OCT software, non-commercial OCT software, and general-purpose analysis software. Segmentation of retinal layers and measurement of layer thickness are supported by both commercial software and non-commercial software [80, 199, 205]. Commercial OCT software is commonly distributed by OCT device manufactures. Besides commercial software, few approaches for visually analyzing retinal OCT data exist. Examples are the research-oriented Iowa Reference Algorithms [80], the

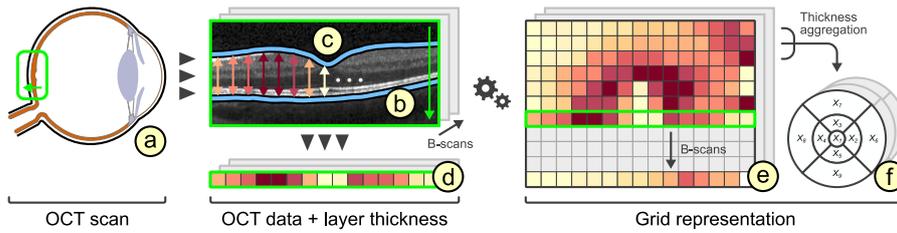


Figure 9.1: Grid representation of retinal layer thickness. An OCT scan captures the area around the macula and the optic disk (a). Multiple 2D depth images (B-scans) are acquired (b). Retinal layers are segmented per B-scan (c) and thickness values are computed for every point along the horizontal image axes (d). The thickness values are combined per layer (e) and aggregated into thickness grids (f).

open-source software ImageJ [182] and its application to OCT data [82], a ray-traced 3D visualization of OCT data [81], machine learning-based systems for visual diagnosis and referral support [93], and a recent VA software based on multiple coordinated views [27]. These software tools are typically applied in ophthalmic research rather than in clinical practice. All available software packages support at least one of three fundamental presentation methods: cross-sectional views, 3D views, and top-down views.

Cross-sectional views show individual 2D depth images (B-scans) of volumetric OCT data together with overlaid profiles of segmented retinal layers. This allows to view details but flipping through the images is time-consuming, as OCT datasets can consist of hundreds of depth images. *3D views* show an entire OCT dataset as a 3D volume rendered tomogram. This provides an overview of the data but combined 3D visualizations of the tomogram and the layers are only provided by few tools, e.g., [27]. *Top-down views* show a fundus image of the interior surface of the eye around the OCT acquisition area together with superimposed retinal layers. This facilitates a layer-centric analysis of the data and helps to link the layers to retinal areas in the fundus image. In general, cross-sectional views and 3D views are mainly used for the analysis of individual datasets, whereas top-down views are also applied to anatomically localize and compare areas under investigation in multiple OCT datasets. In this regard, top-down views are most relevant to our work.

Instead of showing the raw OCT data, top-down views typically represent retinal layers via their derived layer thickness. The layer thickness is displayed either via thickness maps or via spatially aggregated thickness grids. This helps to reveal retinal changes, which may be difficult to identify by visualizing the raw OCT data alone. While recently there has been progress in the comparative visualization of thickness maps in OCT data [27, 30], thickness grids are still predominantly used in most ophthalmic applications. One reason for this is that thickness maps, albeit being spatially precise, can be unnecessarily complex compared

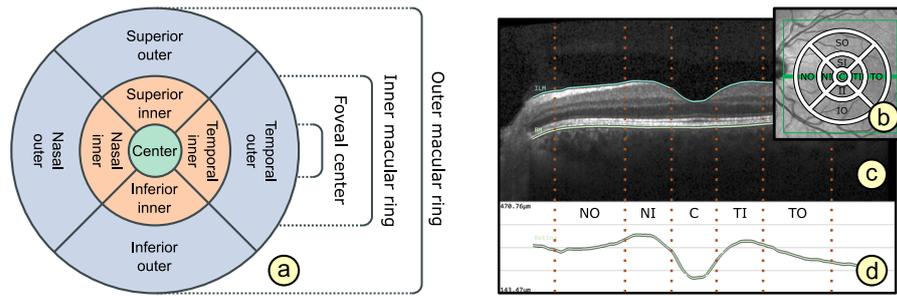


Figure 9.2: The layout of the standard **ETDRS** grid. The **ETDRS** grid cells divide the retina into nine regions defined by three rings and four quadrants (a). An overview shows the grid on top of a fundus image of a healthy left eye (b). The location of an **OCT** depth image (**B-scan**) is marked (central green line) (c), segmented layer boundaries of total retina, and the layer thickness as a line chart (d). The dotted lines denote anatomically distinct areas along the image axis.

to simpler grids in certain situations and they tend to be more difficult to analyze with respect to anatomically distinct retinal areas. On the other hand, the design of thickness grids has been hardly investigated in earlier work, despite the continuous development of **OCT** technology and analysis methods. In fact, in ophthalmic research, only few standardized grids have been employed for data analysis and study results are mostly shown as statistical plots or numbers in a table without spatial context. On this account, we focus on the grid-based exploration of retinal thickness data. Our main objective is to support an information-aware data reduction that preserves relevant details while summarizing remaining parts to a manageable amount.

9.2.2 Representation of retinal layer thickness via **ETDRS** grids

The most common grid type for retinal thickness data was established by the Early Treatment Diabetic Retinopathy Study (**ETDRS**) [89]. The **ETDRS** grid cells divide the retina into nine large regions defined by three rings, i.e., a central foveal ring with 1mm diameter, an inner macular ring (pericentral) with 3mm diameter, and an outer macular ring (peripheral) with 6mm diameter. The inner and outer rings are divided into four quadrants, namely nasal, temporal, superior, and inferior. For each grid cell and retinal layer, one aggregated thickness measurement, e.g., the arithmetic mean, is stored. Figure 9.1 illustrates the generation of **ETDRS** grids and Figure 9.2 shows their layout. **ETDRS** grids have been widely applied for various purposes in ophthalmic research, including investigations of early retinal changes for a variety of diseases, e.g., diabetes mellitus [90] and glaucoma [91]. Although other grid types exist, they have been mostly designed for special applications, e.g., rect-

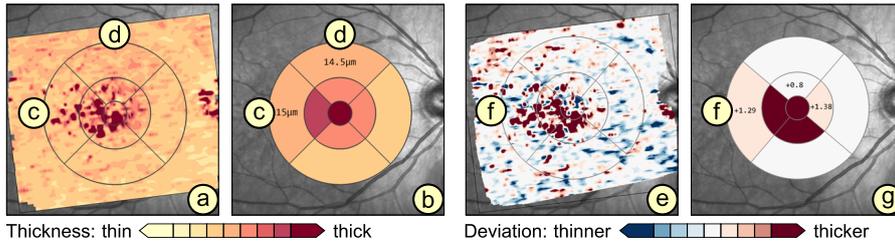


Figure 9.3: Problems with an [ETDRS](#) grid-based data representation. The thickness of a retinal layer is shown via a thickness map (a) and an [ETDRS](#) grid (b). A cell with localized regions of high thickness, visible in the map (dark red; c), has almost the same aggregated value in the grid as a cell without such regions (d). Localized regions of positive and negative deviations in thickness in the map (dark red and dark blue; e) are nearly nullified due to data aggregation in cells (f) of the grid (g).

angular grids for asymmetry analysis of retinal thickness in glaucoma diagnosis [92].

An advantage of [ETDRS](#) grids is their compact representation of complex thickness data. This allows a quick overview of thickness changes in predefined retinal areas. The layout of [ETDRS](#) grid cells also accounts for natural differences in the shape of the retina. It enables the localization of important areas of the macula near the center of the retina and the assessment of thickness data with respect to each anatomically distinct location. Moreover, the fixed data reduction eases the evaluation and comparison of multiple datasets, even from different studies.

The main problem with [ETDRS](#) grids is, however, that they represent aggregated thickness data (Fig. 9.3). Due to the considerable spatial data reduction, localized variations in thickness are not accurately captured via a single aggregated thickness value per grid cell. This is the case for both small variations in thickness within a grid cell and variations divided by the boundary of two or more grid cells (Fig. 9.3 (c)). Moreover, when evaluating deviations between thickness data of several datasets, aggregation artifacts may lead to false negative findings. The reason for this is that localized positive and negative deviations within a grid cell can get nullified during data reduction (Fig. 9.3 (f)). Detecting such information is vital for the identification of early retinal changes of certain diseases.

9.3 GRID DESIGN

[ETDRS](#) grids are well-established and widely-used. We aim at designing alternative grids that combine the advantages of [ETDRS](#) grids with the possibility to capture more relevant information. To this end, we identified grid-related design requirements, devised a new subdivision scheme that allows to compare grid layouts, and developed a method to rate the representation quality of grids.

9.3.1 *Grid-related requirements*

The grid-related design requirements reflect the ophthalmologists' needs with regard to processing and analyzing retinal thickness data. We derived the following list by talking with ophthalmologists about current limitations and the way they utilize existing grids to analyze retinal thickness data.

LAYOUT BASED ON **ETDRS** GRID (GR_1): The basic layout of alternative grids should correspond to the **ETDRS** grid. This is to maintain the ability to localize anatomically important areas of the macula near the center of the retina.

COMPARABILITY OF GRID LAYOUTS (GR_2): Alternative grid layouts should be comparable to both the basic **ETDRS** grid and other alternative grids. This is to ensure that analysis results from multiple datasets with different grids are relatable.

COMPACT DATA REPRESENTATION (GR_3): The number of grid cells should be small and the content of a grid cell should be represented by mainly one descriptive value. Nevertheless, an appropriate representation of thickness data has to be facilitated.

9.3.2 *Subdivision of grids*

Based on the ophthalmologists' demands, we design alternative grids that meet requirements GR_1 and GR_3 . Taking the **ETDRS** grid as a basis, we employ a radial or sector-wise partitioning strategy for **ETDRS** grid cells. This allows us to obtain alternative grids that represent the underlying thickness data at different levels of granularity. [Figure 9.4](#) (a) exemplifies both partitioning strategies.

Radial partitions add rings to a grid. The resulting grid cells enable a more fine-grained analysis of areas with respect to the distance from the foveal center. For example, retinal layer thickness of macular rings can be investigated in-depth via secondary inner or outer ring cells. *Sector-wise partitions* add separating lines at certain angles to a grid. The generated grid cells facilitate a direction-centric analysis of areas with respect to the foveal center. For example, by adding further directional cells, the thickness of areas facing nasally may be evaluated in greater detail. Radial and sector-wise partitions can be combined to obtain fine grids that support analyses with focus on both properties.

To restrict the number of all possible combinations of subdivisions, we start by equally dividing grid cells and increase the number of radii or sectors in power of two steps. By dividing grid cells in half in each subdivision pass, the amount of information stored in a subdivided grid is increased in constant steps. The resulting set of grids can then be refined interactively. Just like standard **ETDRS** grids, we represent the

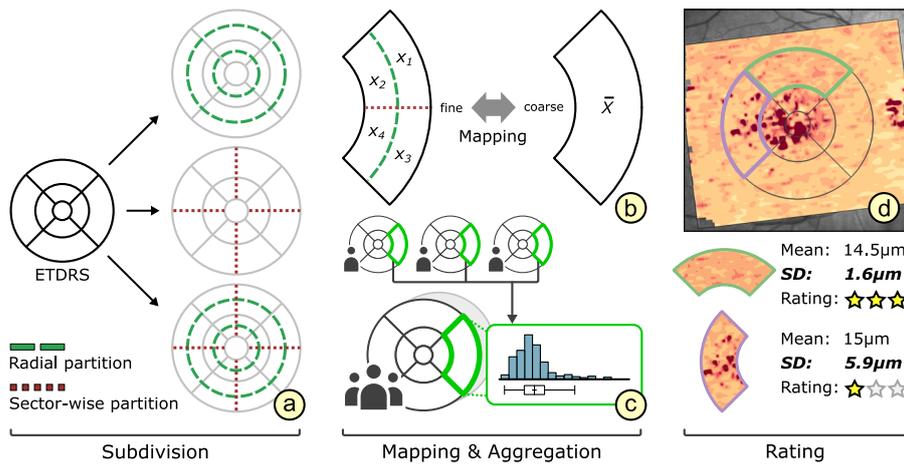


Figure 9.4: Subdivision, mapping, aggregation, and rating of grids. The standard **ETDRS** grid layout is subdivided via radial (dashed line) or sector-wise (dotted line) partitions (a). A coarse grid is mapped to a fine grid or vice versa by subdividing or merging corresponding grid cells (b). Multiple source grids are compiled into a single aggregated grid (c). The representation quality of grids is rated by measuring the standard deviation of thickness values within each grid cell (d).

content of each grid cell via one aggregated thickness measurement. Optionally, summary statistics, e.g., mean, percentiles, and standard deviation, are stored per cell to provide further information on the distribution of the underlying thickness data.

Subdividing the standard **ETDRS** grid helps ophthalmologists to familiarize themselves with the layout of derived grids. Added cells are easily associated with anatomically distinct retinal areas. This increases the acceptance of alternative grids.

9.3.3 Comparability of grids

Ophthalmologists are often interested in relating grid-based analysis results from multiple datasets (cf. GR_2). Each dataset of individual patients may be best represented by another grid with a different layout. To ensure the comparability of the grids, we support mapping a fine grid to a coarse grid and vice versa. This is possible as in our design a fine grid is basically a subdivided version of a coarse grid. **Figure 9.4** (b) illustrates both mapping strategies. In addition, grids of a patient group may have to be compared to grids of another patient group or of a control group. On this account, we allow to compile a set of multiple source grids of a group into a single aggregated grid. The aggregated grid cells can then be mapped and assigned values can be directly related. **Figure 9.4** (c) shows the aggregation of grids.

Mapping a fine grid to a coarse grid entails that subdivided cells have to be merged together to the granularity level of coarser cells. The values

of merged cells are determined by aggregating the values of subdivided cells, e.g., by computing and storing their arithmetic mean. A prominent example for this mapping strategy is the backtracking of alternative grids to the standard [ETDRS](#) grid. This allows to compare new analysis results obtained via our grid design to results based on the [ETDRS](#) grid in previous ophthalmic studies. Another example is to pick the coarsest subdivision grid from a set of alternatives and to map the other grids to that reference. This is necessary if the thickness data are no longer available and thus, merging finer cells is the only option.

Mapping a coarse grid to a fine grid involves that coarse cells are subdivided to the granularity level of fine cells. The values of the subdivided cells are then assigned either by recomputing them based on the underlying thickness data or simply by copying the values of corresponding coarser cells. A precondition for the second case is that the coarse grid is a good data representation and consequently, a subdivided version of that grid represents the data equally well, i.e., it just consists of more grid cells. An advantage of this mapping strategy is that no details stored via finer grid cells are lost during a comparison. This is important if a fine grid representing a patient with abnormal localized variations in thickness has to be compared to a coarser grid representing normative data of healthy controls, which commonly show less variations in thickness.

Compiling an aggregated grid implies that multiple source grids are transformed into a single grid representation. First, a layout for the aggregated grid is determined and then the source grids are mapped to that common layout. The aggregated grid represents the source grids via one descriptive value per cell, e.g., the arithmetic mean of corresponding source cell values. In addition, summary statistics of source values may be stored per aggregated grid cell. To analyze two or more aggregated grids, they are mapped using one of the above strategies and related based on corresponding cell values. This way, a grid of a single patient or an aggregated grid of a patient group can be compared to an aggregated grid of a control group.

9.3.4 *Rating of grids*

Generally, there is no single grid layout that represents all spatial distributions of thickness data equally well. Based on our partitioning strategies, we instead get a set of readily available alternative grids. In order to suggest which of these grids actually corresponds to the given data, we developed a rating procedure that measures the representation quality of individual grids. Our rating of grids addresses the second aspect of grid-related requirement GR_3 . [Figure 9.4](#) (d) exemplifies the grid rating.

We compute the representation quality of a grid by measuring the homogeneity of thickness values within each cell. This is based on

the assumption that a cell with high homogeneity covers values that are more or less equal, whereas a cell with low homogeneity encloses strongly varying values. Thus, a cell value with high homogeneity matches the data, whereas a cell value with low homogeneity indicated information loss. Consequently, a grid consisting of homogeneous cells corresponds to a good representation of the thickness data and vice versa.

One possible measure to quantify the amount of variation or homogeneity of a set of thickness values is the standard deviation. With this simple measure, we obtain the overall rating of a grid via the weighted arithmetic mean of the standard deviations of all grid cells.

$$\bar{x} = \frac{\sum_{i=1}^n w_i x_i}{\sum_{i=1}^n w_i}$$

In this equation, \bar{x} represents the final rating of the grid, n denotes the number of grid cells, x_i is the standard deviation of a cell, and w_i is the weight for that cell given by the normalized amount of enclosed data points (equal to the size of the cell). Based on the final ratings, we compute a ranking for a set of alternative grids and suggest a best fit for thickness data of one retinal layer while considering secondary constraints by ophthalmologists. Such constraints are a specified maximum number of allowed grid cells or a preference for either radial or sector-wise subdivisions. This promotes a more patient-specific analysis in contrast to generalizing all given thickness data to just the [ETDRS](#) grid representation. Likewise, the rating and ranking can be adapted to support a group-specific assessment of grids. The procedures allow to determine one best fitting grid for multiple layers of one dataset or even for one or several layers of multiple datasets, e.g., one grid to represent the data of all patients in a study. The resulting rankings are used to steer the grid-based visual exploration of thickness data.

In summary, our grid design allows to obtain compact representations of retinal thickness data comparable to the [ETDRS](#) grids, while also being able to capture more relevant information. Thus, our solution is a first step towards helping ophthalmologists to use grids that correspond to the thickness data they examine. In addition, we develop interactive visualization techniques that support the exploration of different grids and the adaption of their visual representation.

9.4 GRID EXPLORATION

We aim at supporting users in their grid-based analysis of retinal layer thickness. For this purpose, we identify visualization-related requirements and design matched visualization techniques. [Figure 9.5](#) shows an overview of the user interface. Our design supports: (i) visualizing grids, (ii) showing grid details, (iii) interactively adapting grids, and (iv) comparing different grids.

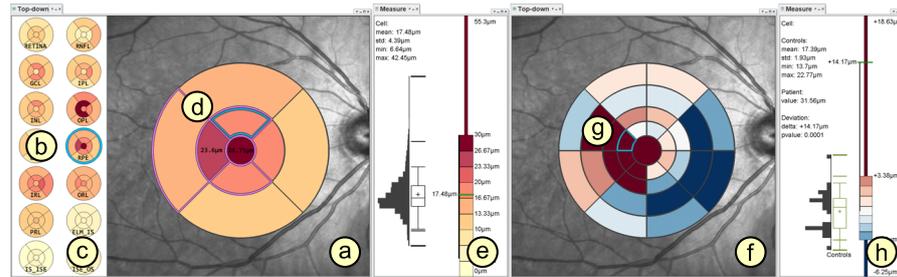


Figure 9.5: Overview of our visualization design. In the left top-down view (a), a retinal layer (b) is selected in the layer overview (c) and an associated thickness grid is shown on top of a fundus image. Cell color encodes aggregated thickness values, labels indicate cell values close to specified thresholds, and borders of cells with low ratings are highlighted (purple). Details of a selected cell (d) are shown in a linked measurement view (e). In the right top-down view (f), a grid is mapped and compared to a control grid and deviations are color-coded. Details of a selected cell (g) are shown in a second measurement view (h) in relation to the control distribution.

9.4.1 Visualization-related requirements

The visualization-related requirements correspond to the ophthalmologists' goals when visually analyzing retinal thickness data. We devised suitable visualization concepts following a participatory design process. Together with ophthalmologists, we derived a list of design requirements and made general design decisions:

VISUALIZATION IN SPATIAL CONTEXT (VR_1): The thickness values of grids have to be presented in their underlying spatial frame of reference. This ensures that they are relatable to respective regions within the retina and allows to understand their spatial distribution.

GRID DETAILS ON DEMAND (VR_2): Details of grids with respect to space, e.g., finer subdivision of cells, and encoded information, e.g., the distribution of underlying thickness values, have to be made visually available upon request. This way, interactive investigations of findings at different levels of granularity are possible.

RELATION OF MULTIPLE GRIDS (VR_3): The thickness differences between multiple grids have to be graphically presented. In connection with the comparability of grid layouts (cf. GR_2), the differences need to be visualized in a common space. In case of ophthalmic studies, statistical quantification on top of a pure comparative display is required.

9.4.2 Presentation of grids

The visualization design is based on multiple coordinated views. It includes a top-down view for coloring and labeling of grids and a measurement view for showing details of cells. With the linked first and second view, we address the visualization-related requirements VR_1 and VR_2 .

The *top-down view* provides an overview of different thickness grids with regard to the interior surface of the eye (Fig. 9.5 (a)). A fundus image depicts the OCT acquisition area. On top, colored grids of selected retinal layers are visualized. The opacity of the grid overlay can be adjusted using a slider to relate cell values to noticeable structures in the fundus image. All other retinal layers are shown as grid thumbnails on the side, ordered according to their anatomical location within the retina (Fig. 9.5 (c)). This view design presents grids for all layers in one image without having to flip through them manually. Thus, layers with abnormal characteristics can be easily discovered.

The *coloring of grids* is based on suitable and adjustable palettes [132]. Sequential palettes encode the actual thickness of individual grids or the averaged thickness of aggregated grids. The cells are colored by evaluating the stored values with respect to clinically established thresholds. Two boundary values are given for all cells of each retinal layer. Low thickness values are assigned to light colors, high thickness values to dark colors, and thickness values outside the specified ranges to distinctively lighter or darker colors. Figure 9.5 (a) shows an example. This allows to judge the thickness data globally in relation to the given boundaries. The coloring based on common thresholds also helps to relate grid presentations across different datasets.

The *labeling of grids* enriches the colored grids with additional text labels and highlighted cell borders. Text labels are added to show cell values in detail. Optionally, only text labels of selected cells or of cells with values outside of specified thresholds are displayed. This is to prevent visual clutter in the image, particularly in fine grids with a lot of small cells. Instead of showing numeric cell values, the labeling can be switched to encode location-oriented cell names (Fig. 9.6 (a)). To this end, existing naming conventions of the ETDRS grid (cf. Fig. 9.2) can be adapted to derived grids. One alternative naming scheme is based on notations of points on a compass for directions and recursively partitioned macular rings for distances. For instance, a location is denoted by *SSN:OO*, which encodes a cell in the direction of *superior superior-nasal* (analogous to north northwest) at distance *outer outer* (outer half of outer macular ring) between the foveal center and the macular border. Next to text labels of cells, cell borders are outlined to illustrate the structure of grids and to highlight specific cells (Fig. 9.6 (b)). By default, all cell borders are outlined. The presentation can be adjusted to match selectively shown text labels of abnormal cells or to mark

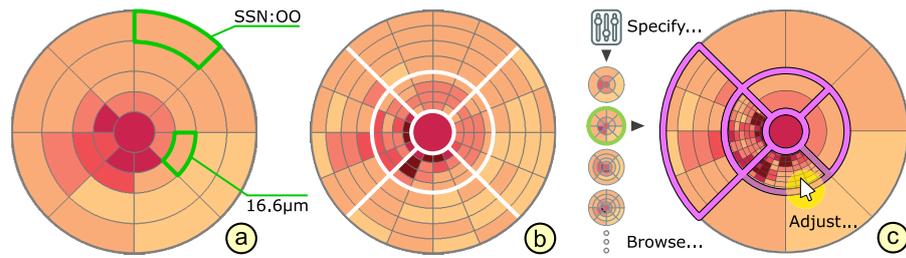


Figure 9.6: Interactive labeling and adaption of grids. The labeling of grids shows either numerical cell values or location-oriented cell names (a). White cell borders mark the layout of the standard [ETDRS](#) grid on top of a subdivided grid (b). The adaption of grids involves specifying initial grids, browsing through them, and adjusting cells of selected grids (purple borders) on demand (c).

the [ETDRS](#) grid layout in subdivided grids. This further facilitates the localization of cells. Finally, cell borders may be emphasized to reflect the computed grid ratings, e.g., to indicate information loss in cells with low homogeneity ([Fig. 9.6](#) (c)).

The *measurement view* goes beyond the basic grid presentation by showing further details of the thickness data ([Fig. 9.5](#) (e)). Thickness values of selected cells in the top-down view are visualized as lines and numerical text labels on top of a color legend in the measurement view. The legend reflects the applied coloring and specified thresholds of the grid presentation. Next to the legend, summary statistics of selected cells are optionally shown as statistical plots. For individual grids, the distribution of the contained thickness data is visualized via box-and-whisker plots and histograms together with numerical labels of descriptive values. This helps to judge the grid representation and to understand the rating of cells, e.g., by showing the distribution of highlighted cells with low homogeneity. For aggregated grids, the statistical plots encode the distribution of thickness values of all corresponding source cells. This allows to assess the variability of thickness values of a group of patients.

9.4.3 Adaptation of grids

To promote an in-depth data analysis, the thickness data have to be explored at an appropriate level of granularity. We support users in specifying initial grids to start the exploration, browsing through grids at different levels of granularity, and adjusting grids on demand. [Figure 9.6](#) (c) exemplifies the adaptation of grids. These interactive means address visualization-related requirement VR_2 .

Specifying initial grids is possible in two ways: manually or automatically. The first option is based on interactive grid design, which gives ophthalmologists the greatest freedom and allows them to obtain grids that meet predefined needs. To manually specify a grid, the user has to

set layout parameters, i.e., the number of radii and sectors. The second option is based on computed grid suggestions, which provide grids with the highest ratings of a number of alternative grid layouts. Alternative layouts may consist of increasingly finer subdivisions of the [ETDRS](#) grid using combinations of both radial and sector-wise partitioning strategies. Automated grid suggestions do not require user input. However, an ophthalmologist may steer the suggestions by setting additional constraints, e.g., the maximum number of allowed grid cells. This helps to obtain grids that adapt to variations in thickness but do not exceed a predefined level of granularity.

Browsing through grids is supported by interactive selections from a set of predefined grid layouts with different subdivisions. Selected grids are immediately shown in the top-down view with cells colored and labeled. By flipping through grids, the data are explored at different levels of granularity. For instance, selecting increasingly finer subdivisions of the [ETDRS](#) grid allows to analyze the data from overview, i.e., the coarsest grid, to detail, i.e., the finest grid. This helps to understand the data and to localize areas of interest.

Interactively adjusting grids allows to refine or extend a predefined subset of possible grid layouts. On demand, one or multiple grid cells are selected and then merged together or subdivided by means of the provided partitioning strategies (cf. [Sect. 9.3.2](#)). This helps to fine-tune grids for given data based on the expertise of ophthalmologists. For example, while browsing through alternative grids, a specific grid is selected, cells are visualized, and information about the thickness data is shown. Individual grid cells with low ratings are further subdivided to investigate respective areas in greater detail and the remaining cells are merged together to provide context information of less relevant areas.

9.4.4 Comparison of grids

Both a patient-specific and a group-specific analysis of retinal layer thickness requires the comparison of grids from multiple datasets. We enable such comparisons by juxtaposition of views, explicit encoding of deviations, and application of statistical tests. With the combination of these techniques, we address visualization-related requirement VR_3 .

The juxtaposition of views supports the comparison of several individual grids ([Fig. 9.5](#) (a) and (f)). In the user interface, multiple instances of the top-down view and measurement view can be dynamically added and freely arranged. For a patient-specific analysis, different grids of the same patient, e.g., follow-up examinations, or of related patients, e.g., similar medical cases, can be assigned to these view instances. Linking the view instances ensures that matching parts of the data are shown. For a group-specific analysis, individual grids of all members of

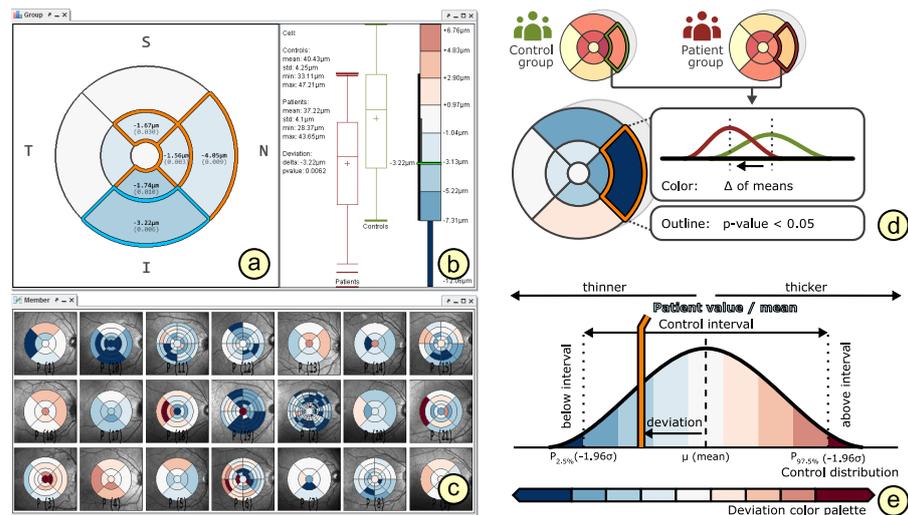


Figure 9.7: Comparison of grids. The top-down view (a) and linked measurement view (b) show a comparison between patients and controls via aggregated grids of both groups. The juxtaposed small multiple views (c) show an overview of grids of all individual patients in relation to controls. Patient values or group means are compared to the distribution of controls (d). A diverging color palette encodes differences per cell (e) and borders of cells with significant differences are highlighted (orange).

a group are shown as small multiple views (Fig. 9.7 (c)). This provides an overview of differences between group members.

The explicit encoding of deviations facilitates a direct comparison of a grid to a reference grid (Fig. 9.7 (d)). For this purpose, the coloring of grids is switched. Diverging color palettes are used to encode cell value with respect to threshold intervals of an aggregated reference grid (Fig. 9.7 (e)). The boundary values vary cell by cell for each layer and are based on confidence intervals or percentile boundaries stored in reference cells. Small deviations are depicted via a light neutral color and larger negative or positive deviations via respective darker colors. Displaying deviations helps to evaluate a grid locally in relation to a reference grid. Figure 9.5 (f) and Figure 9.7 (a) exemplify comparisons of a grid of a single patient and of an aggregated grid of an entire patient group to intervals ranging between the 2.5th and 97.5th percentile boundaries of an aggregated grid of a control group, respectively.

The application of statistical tests allows to quantify differences between multiple aggregated grids. Different statistical tests are supported, e.g., Student's t-test for quantifying differences between patients and controls or Tukey HSD test for multiple comparison between several patient subgroups and controls. The resulting measurements of statistical significance (p -values) and effect size are shown as numerical labels. Borders of cells with significant differences are highlighted with respect to a user-specified significance level, e.g., $p < 0.05$. In addition, statistical plots of cells for each grid under consideration are shown in the

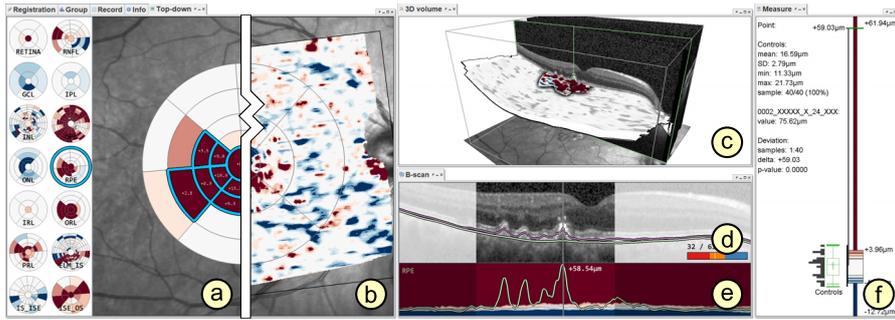


Figure 9.8: A VA tool for retinal layer thickness. A top-down view shows either grids (a) or maps (b). A 3D view presents a volume visualization of raw OCT data (c), and a 2D cross-sectional view displays individual depth images (d) and thickness profiles of selected layers (e). All views are linked, selections in grids or maps are highlighted, and details are shown in a measurement view (f).

measurement view (Fig. 9.7 (b)). This illustrates respective thickness distributions and provides details of the color-coded deviations in the comparative grid presentation.

In summary, our grid exploration enables the analysis and comparison of the spatial distribution of retinal layer thickness at different levels of granularity. Our design helps ophthalmologists to perform in-depth data analyses and to interactively adapt them to their needs. On top of this, we develop a VA tool that bundles our solutions in a flexible user interface and a complementary analysis procedure for evaluating ophthalmic study data.

9.5 GRID-BASED VISUAL ANALYTICS OF RETINAL LAYER THICKNESS

We integrate our grid-based computation, visualization, and interaction techniques into a prototypical VA tool for retinal OCT data. We demonstrate the utility of its main components in an analysis procedure for the evaluation of ophthalmic study data.

9.5.1 A visual analytics tool

Offering our solutions in a VA tool for retinal OCT data enables a comprehensive data analysis beyond the current use of thickness grids. The tool is a result of our previous work [27], which aimed at providing a unified interface for OCT data from different sources and an interactive visualization design for different aspects of OCT data, particularly with respect to thickness maps. We extend the available functionality and complement it with our grid design and exploration techniques. Figure 9.8 shows an overview of the user interface. The extended tool consists of three functional components for computation, visualization, and interaction:

The *computation component* combines diverse automated means for processing and analyzing retinal OCT data. Regarding our grid design, the computation of alternative grids, the rating of their representation quality, and the mapping and aggregation of grids are part of this component. It also provides a common data basis and parsing options for OCT data from different sources. Based on the converted data, layer thickness is derived together with other layer attributes, e.g., curvature, to characterize the condition of the retina. The thickness data are stored as maps or grids. Semi-automated registration methods allow to transform imported datasets into a common spatial frame of reference and to remove basic errors in the segmentation and thickness data. This enables precise comparisons of multiple grids or maps. Various statistical methods support the measurement, summarization, and quantification of grid cells or regions on maps.

The *visualization component* covers the necessary designs to present and visually relate different parts of OCT data, i.e., raw OCT data, segmentation data, and thickness data. Together with our grid exploration, three types of views are supported: a top-down view, a 3D view, and a cross-sectional view (Fig. 9.8). The top-down view depicts thickness grids and maps with respect to their location within the retina. The 3D view displays OCT data and surfaces of segmented layers in a blended volume visualization. The cross-sectional view shows individual depth images (B-scans) together with thickness profiles of selected layers. An extra measurement view provides additional information on cells, points, or regions selected in the other views. Supplementary methods for emphasizing spatial relationships and characteristic data values allow to show important information in the context of entire datasets. This includes information on the data quality, e.g., the signal strength of B-scans or invalid and missing values in the OCT data, segmentation data, and thickness data.

The *interaction component* supplies techniques that enable users to interact with the visualized data. This includes the adaptation and comparison of different grids. Comparisons of maps from individuals or groups are equally possible. Coordinated spatial and data-driven selections allow users to specify points, regions, cells, or value ranges of interest and to highlight them in the different views. Selections made in one view are automatically propagated to other linked views. For instance, selecting a grid cell allows to quickly access and relate respective locations on the surfaces of layers in the 3D view and associated depth images in the cross-sectional view. To browse through the data, layers are switched or a cursor is moved in one view. The other views are updated accordingly and continuously display matched parts of the data. All navigation techniques are based on smooth animated transitions between camera movements to help users maintain their mental model of the visualized data.

The three functional components of our tool are implemented in a modular software architecture and a flexible user interface. The back end features a dedicated resource management that ensures efficient data processing and presentation. While most computations and renderings are performed on the fly (e.g., the adaptation of grids and automated measurements), some data preparations are preferably done offline (e.g., the initial registration and aggregation of many datasets). The front end supports stylizing the visualization and interactive exploration of retinal thickness data. This allows us to consider alternative display types, e.g., stereoscopic rendering to enhance the 3D perception of the spatial data. Altogether, our tool enables an interactive switch between different views, different levels of granularity, and two different tasks of analyzing individual patients and groups of patients.

9.5.2 *A visual analytics procedure*

To demonstrate our prototypical visualization tool, we developed an exemplary procedure for evaluating data from cross-sectional ophthalmic studies. In this connection, we discuss advantages and disadvantages of a Current Analysis (CA) approach with ETDRS grids in relation to our enhanced grid-based VA approach. An overview of both methods is shown in Figure 9.9 and respective outcomes of two studies are described in Section 9.6.2. Together with ophthalmologists, we compiled the steps involved in CA and VA by observing and discussing data analysis practices. Four common analysis stages were identified: (1) data preparation, (2) data transformation, (3) comparison of patients and controls, and (4) quantification of findings.

In the first stage, the study data are prepared. Volumetric macular OCT images are acquired for every single subject and retinal layers are segmented using device-specific software. The image and segmentation data are screened to identify datasets with insufficient signal quality and erroneous or invalid layer boundaries. One or more best datasets per subject are selected and layer thickness data are computed for subsequent analyses.

In both CA and VA, data quality checks are crucial to prevent acquisition artifacts from impairing the data analysis. In CA, every dataset of every subject has to be validated one by one. This entails considerable manual effort, particularly in case of datasets with a high number of depth images (B-scans). A high density of B-scans is necessary for accurate retinal thickness measurements, prevention of interpolation errors [175], and capturing small and localized areas of retinal changes that occur in an early stage of a disease [174].

In VA, our tool provides additional support for displaying multiple datasets at once, indicating missing values and variations in signal strength, and highlighting and correcting segmentation errors [27]. Ensuring adequate segmentation quality is important in VA, as localized

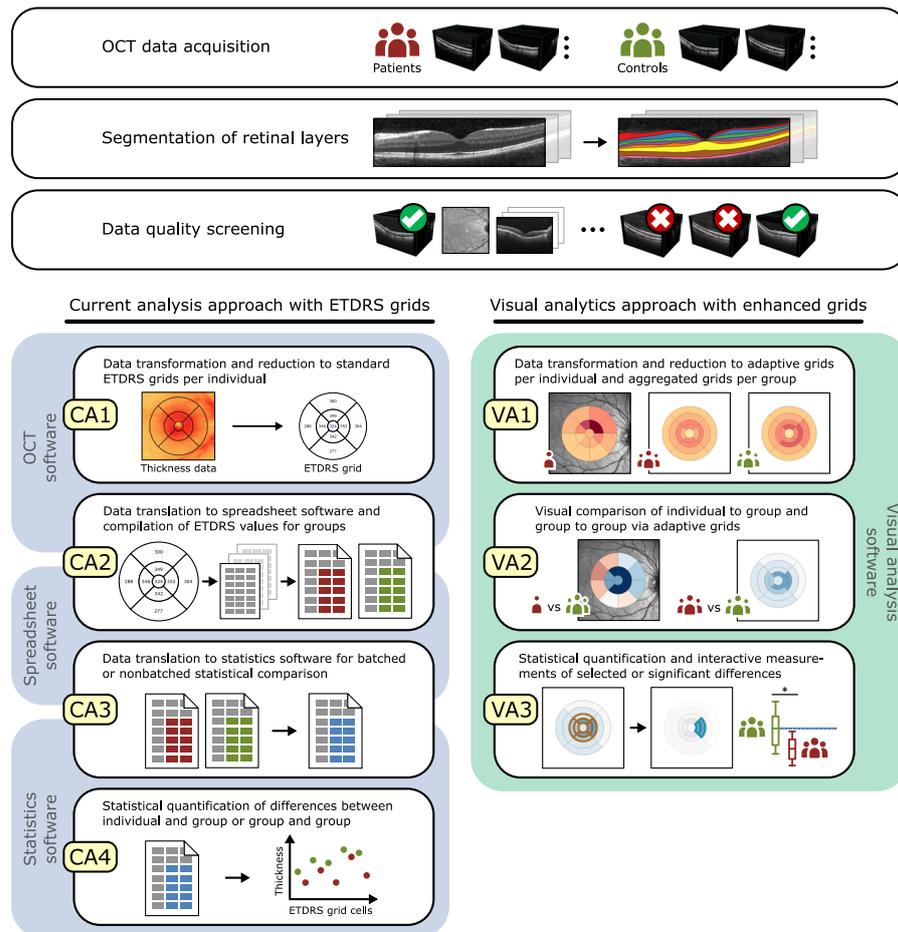


Figure 9.9: Exemplary procedures for evaluating cross-sectional ophthalmic study data using the current **ETDRS** grid-based analysis (**CA**) approach and our enhanced grid-based **VA** approach. After a common data preparation stage, **CA** comprises four analysis steps (**CA1** to **CA4**) with three software tools and **VA** involves three analysis steps (**VA1** to **VA3**) with our **VA** tool.

errors may become directly visible in smaller cells of subdivided grids. This is not necessarily the case in **CA** due to data aggregation in larger cells of **ETDRS** grids.

In the second stage, the study data are transformed. All datasets have to be registered and the thickness data are projected into a common spatial frame of reference. The registration is necessary to account for peculiarities of individual datasets, e.g., opposite laterality, tilted or off-center scan location, or different data resolution. Afterwards, the thickness data are reduced via given grid layouts.

In **CA**, data registration is performed by **OCT** device software. However, inter-individual positioning and laterality matching are often done by hand. Even though dedicated software options exist, their accessibility is limited because they are not supplied and installed as standard software features. The device software is then used for automated re-

duction of thickness data to standard [ETDRS](#) grids per layer and dataset ([Fig. 9.9 \(CA1\)](#)). This provides a manageable amount of information per layer, which is easier to handle in the following data translation and statistical comparison of [CA](#).

In [VA](#), our tool offers semi-automated registration of datasets using the centers of the fovea and the optic disk as reference points [27]. Differences in laterality or resolution are compensated. Appropriate grid layouts are either generated based on the computed ratings or specified by the user. The thickness data are aggregated for each layer in datasets of individual patients, of entire patient and control groups, and of optional subgroups of patients ([Fig. 9.9 \(VA1\)](#)). The adapted grids of individual patients provide a balanced representation of the thickness data. The aggregated grids of study groups are equally compact and precise and promote a cell-by-cell analysis with high spatial specificity. This aspect differs from [CA](#) in which data reduction and analysis based on [ETDRS](#) grids are applied.

In the third stage, the study data of patients and controls are compiled and compared. According to the study design, all data of investigated layers from individual subjects, subgroups, and groups are selected and assorted for the desired comparisons. After the comparisons, identified patterns are marked for further quantification and detail analysis.

In [CA](#), this stage is divided into two steps. First, the study data are translated and exported from [OCT](#) software to spreadsheet software ([Fig. 9.9 \(CA2\)](#)). The spreadsheet software is used for data fusion and compilation of groups or subgroups of patients and controls. Yet, the available export functionality can be limited. This is because, in our experience, [OCT](#) software from different manufacturers does not always provide an option to batch process multiple datasets and to export [ETDRS](#) grid values to external files, or such functionality is not activated by default. Hence, dozens of values per dataset may have to be copied by hand. This is time-consuming and error-prone, particularly in the context of larger studies with a lot of subjects. For example, in one case, more than 4000 data values were read from display of one software and manually copied to another software. Second, the compiled study data are again transferred to statistics software to perform batched or non-batched statistical comparison ([Fig. 9.9 \(CA3\)](#)). Typical comparisons are made either with respect to a fixed set of predefined study hypotheses or by computing and evaluating all potentially relevant combinations. The output of statistics software is summarized and reviewed in form of lists, tables, or statistical plots.

In [VA](#), adapted grids of individual patients, subgroups of patients, and the patient and control groups are mapped and visually compared ([Fig. 9.9 \(VA2\)](#)). The visualizations help to identify where exactly and to what degree retinal changes occur. This is a distinct advantage compared to [CA](#), since changes in layer thickness of patients can be examined in their spatial context. First insights into the data are gained early in

the overall analysis process and without the need for further data translation and export. The exploratory nature of **VA** does not necessarily rely on predefined study questions but rather assists in discovering patterns and investigating only involved layers and cells in details. Typically, **VA** of study data starts by comparing the patient group to the control group to get a general idea of thickness deviation in patients. Following, other possible comparisons are between patient subgroups or between patient subgroups and the control group, which are helpful in identifying confounding factors that influence the study variables. Individual patients can also be compared to the control group, which is useful in understanding patient-specific retinal changes and tagging mild or severe cases. Optionally, our tool supports automated export of all or selected grid values of the study data to external files for further processing using external software.

In the last stage, findings are quantified and further investigated. Suitable statistical tests are applied to assess the observed differences in the comparisons of patients and controls. The most interesting findings are selected, graphically summarized, and reported.

In **CA**, statistics software provides various tests on top of the computed differences. Significant differences are highlighted in numerical lists or tables. Relevant results are then summarized via different types of charts, e.g., plots of mean deviation + standard error, that illustrate patient values in relation to controls for each **ETDRS** grid cell or ring (Fig. 9.9 (CA4)). These tables and charts are useful to prove predefined study hypotheses and to draw study conclusions based on statistical numbers. The standardized **ETDRS** grid cells facilitate direct comparisons of statistical results from other studies. However, they do not disclose the presence or absence of localized variations in thickness, which were left unmeasured due to the applied data reduction in the earlier steps of **CA**. This may render the analysis results spatially unspecific.

In **VA**, statistical tests are applied to every grid cell and significant cells are highlighted in the compared grids. This way, the shape and extent of areas with significant differences can be accurately detected. The location and structure of significant cells in subdivided grids illustrate the spatial distribution of thickness changes. These unique patterns cannot be directly compared to statistical results of **CA**. Generally, it is only possible to compare the existence of smaller significant cells in relation to values of **ETDRS** grid cells. Although, based on our mapping of grids, **VA** results may be generalized and traced back to **ETDRS** results of **CA** if necessary. Our tool also supports interactive measurements of single or multiple cells to show statistical plots together with numerical labels of descriptive values. The measurements of cells in one grid are visually related to cells in another grid, e.g., cells of an individual patient to cells of a patient group. Further, the cells are shown together with the raw **OCT** data in the linked views of our tool. This interactive back and

forth between comparisons, measurements, and adaptation of grids exemplifies the flexibility of VA. In contrast, if, after some analysis steps in CA, it turns out that a previous step has to be adjusted, the entire procedure must be reset to make the desired change and then repeated from that point on.

All in all, VA extends CA in several ways. Most importantly, it allows to reveal study-specific localized thickness differences between patients and controls. It also simplifies the overall analysis procedure by combining the computation and visualization functionality required for most analysis steps within a single tool. This is in contrast to CA in which applied OCT software is mainly targeted at clinical practice and thus, ophthalmologist currently have to rely on a collection of diverse special and general purpose tools to evaluate study data. Moreover, the interactivity of our tool enables a flexible analysis process, which can be adapted on demand. Although, an extended support for data preprocessing in the first steps of the analysis procedure is required.

9.6 APPLICATION

We evaluated our research prototype in cooperation with ophthalmologists. Its utility was assessed in the context of one experimental and two cross-sectional studies. Following, we describe three use cases, present exemplary results, and reflect on benefits and limitations of our approach.

9.6.1 *Experimental evaluation of patients with age-related macular degeneration*

In this experimental use case, our principal goal was to demonstrate the feasibility of our approach. We applied our solutions to retinal thickness data of a small group of 8 adult patients suffering from Age-related Macular Degeneration (AMD) and a group of 20 control subjects. We wanted to find out if early retinal changes related to AMD can be captured via different grid representations and how they relate to grid representation of healthy eyes. To reason about the utility of our approach, we assessed the obtained results from a data-centric perspective rather than conducting a comprehensive ophthalmic evaluation. Figure 9.10 illustrates the overall study setup and summarizes the findings.

An early sign of AMD is the presence of drusen in the macula. Drusen are small accumulations of extracellular material that build up between lower retinal layers of the eye [16] (Fig. 9.10 (a)). Particularly, in the thickness data of the Retinal Pigment Epithelium (RPE), they become apparent as small and localized increases in thickness. We hypothesized that standard ETDRS grids are not always the best choice for representing the thickness data in such a setting (cf. Sect. 9.2.2).

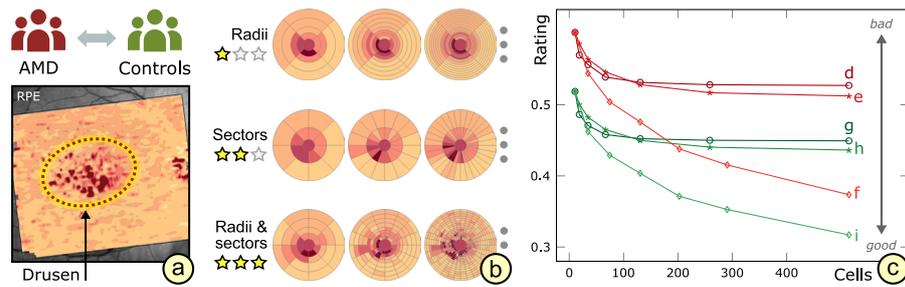


Figure 9.10: The experimental study setup and results in AMD patients and matched controls. Drusen are indicated by localized thickening in lower retinal layers of patient eyes (a). Examples of evaluated grids show increasingly finer subdivision of cells (b). The line plot illustrates ratings of grids with respect to cell counts for thickness data of patients (red) and controls (green) in relation to three partitioning strategies: radial (d, g), sector-wise (e, h), and radial and sector-wise (f, i).

To test this hypothesis, we specified 21 alternative grids and rated them with respect to thickness data of the RPE from patients and controls. For each partitioning strategy, 7 layouts with incrementally added radii, sectors, or radii plus sectors were included (Fig. 9.10 (b)). The numbers of radii and sectors were increased in constant steps up to a maximum of 129 radii, 256 sectors, and 17 radii plus 32 sectors, respectively. The finest grid layouts had 513 cells in each strategy. The ratings of the ETDRS grid and the alternative grids were first computed per subject and then summarized per layout to judge the overall representation quality for each group. In this connection, we assumed that a better rating is equivalent to a superior data representation based on the measures introduced in Section 9.3.4, i.e., a low mean standard deviation implies a good representation. Figure 9.10 (c) shows the obtained results.

Generally, all alternative grids showed better ratings than the ETDRS grid. In patients, grids based on sector-wise partitions performed slightly better than grids based on radial partitions (Fig. 9.10 (d) and (e)). For both partitioning strategies, an increase in cell count only up to a certain point led to better ratings, suggesting that further added cells did not improve the representation quality of grids. In contrast, combined radial and sector-wise partitions showed continuously improved ratings with an increase in cell count (Fig. 9.10 (f)). We suspect that cells generated by this partitioning strategy more closely match the small, roughly circular-shaped regions of high thickness, which exist as a result of early drusen. In controls, the general trends in grid ratings per partitioning strategy were similar to those observed in patients (Fig. 9.10 (g), (h), (i)). This similarity can be explained by the consideration of patients with early AMD, i.e., patients with changes in layer thickness which are just noticeable but not significantly different from the control data. A remarkable difference, however, was that all control grids required fewer cells than the respective patient grids in order to obtain equally

good ratings. We assume that this is due to the fact that thickness data of healthy eyes contain less localized variations in thickness and can be appropriately represented with fewer and coarser cells.

In summary, the alternative grid layouts provide a better representation of the studied thickness data of AMD patients. The results show that our approach allows us to determine suitable grids specifically tailored to each group. This indicates that a spatially-aware data subdivision helps to accurately characterize the condition of retinal layers.

9.6.2 *Grid-based analysis of retinal layer thickness in patients with diabetes mellitus*

Motivated by the experimental results in AMD patients, we applied our solutions to analyze data of two other cross-sectional ophthalmic studies. The first study was focused on early retinal changes in pediatric patients suffering from Type 1 Diabetes Mellitus (T1DM) and the second one on retinal changes in adult patients suffering from Type 2 Diabetes Mellitus (T2DM). In both studies, the ophthalmologists were interested in evaluating the effectiveness of our enhanced grid-based exploration in comparison to a current practice of using only ETDRS grids for analyzing retinal layer thickness. In contrast to the experimental study, which focused on alternative grids to represent the actual thickness data per study group, the ophthalmologists specifically wanted to study how well abnormal thickness deviations of diabetic patients in relation to age-matched healthy controls are detected by each analysis method.

Previous ophthalmic studies have demonstrated deviations in the thickness of retinal layers in pediatric and adult patients with T1DM [90] or T2DM [21]. These deviations in the thickness profiles of the retina serve as biomarkers or surrogate measures that aid an early detection and monitoring the progression of the disease. In an advanced stage of diabetes, the disease-induced alteration of layer thickness are expected to affect large retinal areas and exhibit a clear deviation from control data. In an early stage of diabetes, however, more subtle, localized, and scattered areas of thickness deviation are suspected that may go undetected without a spatially-aware subdivision and aggregation of the thickness data. Hence, we again hypothesized that such changes are not always adequately captured by a current analysis approach based on ETDRS grids. In fact, we were recently able to confirm the existence of localized areas of abnormal thickness in patients with early T1DM using a point-by-point data analysis based on deviation maps [30]. In the present studies, we aimed at obtaining similar results with our enhanced grid-based analysis, while reducing the data to a meaningful extent. The ophthalmologists also wanted to compare their findings in both studies to see if any general similarities between pediatric and adult patients with T1DM and T2DM exist.

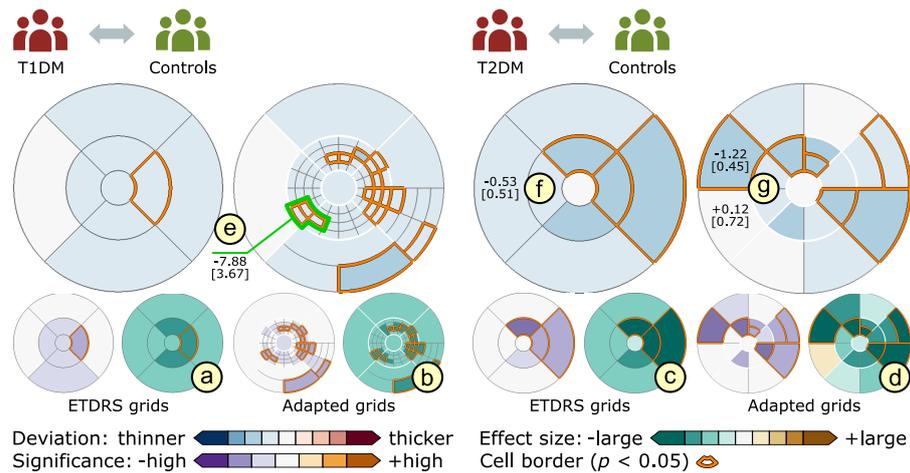


Figure 9.11: Exemplary results of two cross-sectional studies. The grids represent TR in pediatric T1DM patients (a, b) and the thickness of RNFL in a subgroup of adult T2DM patients (c, d) compared to matched controls. Cell color encodes thickness deviation, statistical significance, or effect size. Highlighted cell borders mark significant differences ($p < 0.05$). In T1DM patients, the ETDRS grids (a) depict significant thinning limited to a single cell. The subdivided grids (b) show additional cells of significant localized thinning and measurements of multiple small cells (e). In T2DM patients, the ETDRS grids (c) present a general overview of thickness deviation. The subdivided grids (d) show more details of the spatial distribution and degree of thinning (f).

The data of both studies were acquired via OCT examinations of each subject per group followed by retinal layer segmentation. In the first study, 26 pediatric T1DM patients and 29 age-matched controls participated and in the second study, 33 adult T2DM patients and 40 age-matched controls were involved. For each subject, the Total Thickness of the Retina (TR) and the thickness of inner retinal layers, i.e., the Retinal Nerve Fiber Layer (RNFL), the Ganglion Cell Layer (GCL), the Inner Plexiform Layer (IPL), and the Inner Nuclear Layer (INL), were selected for further investigation. Per selected layer, the ETDRS grid layout and one best fitting alternative grid layout were determined. One aggregated grid was computed per layout for every group in both studies. The resulting grids were then mapped and directly compared using explicit encoding of deviations together with statistical tests to quantify the differences between patients and controls (cf. Sect. 9.4.4). Lastly, based on the statistical results and computed grid ratings, the alternative grids were further adapted to investigate cells with significant differences in greater detail. Figure 9.11 shows exemplary results of the two studies.

Regarding the similarity of results in both studies, the ophthalmologists observed some common patterns. In general, they noted a significant amount of abnormal thinning in several investigated layers of both patient groups. In T1DM patients, cells with significant thinning were

found primarily in the inner macular ring of **TR** and **RNFL**, and selectively in nasal direction of the outer macular ring. In **T2DM** patients, similar findings were made for **TR** and **RNFL** but also for **GCL** and **IPL**. The least amount of significant cells was detected for **INL** in both studies. These results are in accordance with results of other studies that reported thickness reduction of inner retinal layers in patients with **T1DM** (e.g., [22, 96] or [90]) or **T2DM** (e.g., [21, 97]) as an early event of neurodegenerative change.

Regarding the difference of findings between both studies, the ophthalmologists noticed a greater amount, spatial extent, and higher degree of significant thinning in adult **T2DM** patients than in pediatric **T1DM** patients. Particularly, they noted larger cells and darker cell colors (higher deviation) in subdivided grids of most layers in the **T2DM** study (e.g., grids of **RNFL** in Fig. 9.11 (d)). In contrast, subdivided grids from the **T1DM** group rather consisted of smaller cells due to a higher subdivision rate, hinting at the presence of more localized areas of abnormal thickness (e.g., grids of **TR** in Fig. 9.11 (b)). This is in line with our previous map-based findings of localized thickness changes in **T1DM** patients [30]. Possible reasons for the differences in the present study findings are the longer mean duration of disease in **T2DM** patients (**T1DM**: 5.6 years; **T2DM**: 13.3 years) and the inclusion of **T2DM** patients with mild to moderate diabetic neuropathy. A longer duration of diabetes together with an early diabetic neuropathy may have an effect on neuronal cells, which could explain the relatively greater amount and extent of thinning in retinal layers of **T2DM** patients. In this regard, further studies and analyses of correlations between thickness deviations and clinical parameters are necessary.

Regarding the comparison of analysis methods, our enhanced grid-based analysis and the standard **ETDRS** grid analysis both demonstrated general thickness changes of patients in relation to controls. However, the ophthalmologists pointed out a key difference between both methods: the enhanced grids revealed additional details of the distribution and degree of abnormal deviations, whereas the **ETDRS** grids failed to detect significant changes in certain areas. They also noticed a tendency of the **ETDRS** grids to underestimate changes (Fig. 9.11 (f)) and to produce false negative results in the studies. As an example, the neighboring fine cells of significant thinning crossing the inferior-temporal quadrants of the inner macular ring were only detected by the enhanced grid-based analysis (Fig. 9.11 (b)) but not as significant cells by the analysis based on the **ETDRS** grid (Fig. 9.11 (a)). In such a case, our design enables the measurement of multiple subdivided cells involving two or more quadrants within a macular ring (Fig. 9.11 (e)). Detecting and analyzing localized changes that span across boundaries of standard **ETDRS** grid cells was not possible with previous analysis methods. The measurement results allow to quantify the visualized data and help to get a greater understanding of a particular disease condition.

In summary, the enhanced grids are effective in detecting early and localized deviations in retinal layer thickness of pediatric and adult patients with [T1DM](#) and [T2DM](#). The study results emphasize the importance of selecting an appropriate analysis method that supports a precise assessment of abnormal retinal alterations.

9.7 DISCUSSION AND CONCLUSIONS

We presented an enhanced grid-based reduction and exploration approach for retinal thickness data. Our solutions are a continuation of our previous work, in which we introduced the general design of our research prototype [[27](#)] and evaluated first studies using deviation maps [[30](#), [102](#)]. Here, we extended this basis with a new grid design for precise detection of localized thickness variations and bundled the developed techniques in a flexible [VA](#) tool. We also proposed a complementary analysis procedure to demonstrate the tool's utility in the context of two cross-sectional studies. Collaborating ophthalmologists noted the added value of our solutions with respect to more findings and reduced manual analysis effort. We conclude that our approach offers a systematic enhancement of existing work and helps ophthalmologists in their grid-based analysis of retinal layer thickness.

So far, we implemented a data-driven approach to design grids and judge their representation quality. On the one hand, this offers high flexibility in balancing the granularity of grids and the amount of encoded information. On the other hand, automated grid suggestions are dependent on how well the employed quality measures, e.g., the standard deviation, quantify variations in the underlying thickness data. For some applications, very fine subdivided grids with a large number of cells may be practically unfeasible. On this account, our experimental evaluation method (cf. [Sect. 9.6.1](#)) helps to determine how much variance is acceptable in a grid with a certain amount of cells. This is particularly easy in case of grids that show no or only minimally improved ratings in further subdivisions with a certain partitioning strategy (e.g., [Fig. 9.10](#) (d) and (g)). In contrast, suggesting best fitting grids in case of steadily improving ratings with an increase in cell count is more challenging and currently requires input by the user. This also applies to the grids generated via combined radial and sector-wise partitions in our experimental study of [AMD](#) patients, which showed better ratings compared to grids of other partitioning strategies (cf. [Fig. 9.10](#) (f) and (i)). We will evaluate our grid-based methods in further studies to assess their reliability in relation to different patients and retinal conditions, and to share the collected medical insights in full detail with the ophthalmic community.

In this connection, an interesting extension is to establish a diagnosis-driven design and rating of grids. This requires custom strategies for measuring the quality of grids to target different ophthalmic analysis

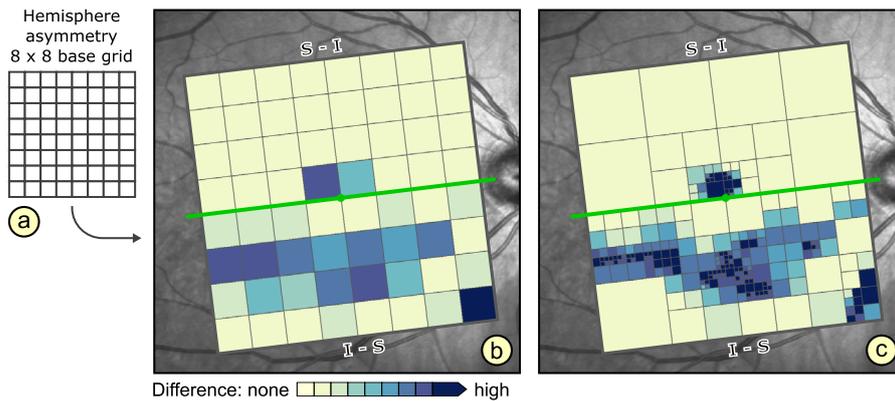


Figure 9.12: Asymmetry analysis of retinal layer thickness using rectangular grids. A standard 8 x 8 base grid layout (a) shows the mean difference between corresponding cells of opposite hemispheres (b). The fovea-to-disc axis (green line) marks the symmetry line and cell color encodes the degree of asymmetry (darker colors represent higher differences). In the upper grid half, negative differences between superior and inferior hemisphere are displayed and vice versa in the lower grid half. The subdivided grid (c) shows additional information on the spatial distribution and degree of differences.

tasks. Moreover, default settings of layout parameters, e.g., the allowed maximum number of radii and sectors, and suitable rating cutoffs for the selection of best fitting grids have to be provided. Our ophthalmologists reported that they are trying to find new grid types that adequately represent retinal changes of specific diseases. A diagnosis-driven approach to rating and ranking of grids may support such efforts by evaluating newly designed grids and sorting out existing grid types.

We are also adapting the main ideas of our grid design and exploration to other basic grid types in more specific ophthalmic applications. One example is the asymmetry analysis of retinal thickness data for glaucoma diagnosis using rectangular grids [92]. Currently, a standard 8 x 8 grid layout is used to analyze cell-by-cell differences in thickness between a superior and an inferior hemisphere. Figure 9.12 illustrates this application. Similar to the standard ETDRS grid, the predefined rectangular grid does not support an information-aware reduction of thickness data. Hence, we are actively extending and testing our grid design with respect to this grid type to support a fine-grained data analysis.

We further plan to augment our analysis procedure with other software tools. This will provide support for evaluating different types of ophthalmic studies. For example, in addition to retinal layer thickness, often other clinical parameters of patients are investigated. Typically, a diverse set of tools besides our solutions is required for preparation, analysis, and presentation of combined OCT data and clinical data. Fully integrating all tools into our software or replacing all their

functionality with custom solutions is hardly possible. Instead, we follow a lightweight approach of loosely coupling individual tools into toolchains [206]. This way, ophthalmologists will be relieved from having to start, adjust, and transfer data to each required tool individually. In an initial effort, we are investigating the definition and execution of temporal tool sequences with an interactive editor.

DEVIATION MAPS FOR UNDERSTANDING THICKNESS CHANGES OF INNER RETINAL LAYERS IN CHILDREN WITH TYPE 1 DIABETES MELLITUS

SCIENTIFIC CONTRIBUTION: In this chapter, we evaluate our novel VA solutions in a cross-sectional study with pediatric patients suffering from type 1 diabetes mellitus. We show that our methods improve the identification of early retinal defects compared to the current analysis practice. In particular, we point out cases of small and localized regions with abnormal variations in retinal thickness that could not be detected before. Since we gained new medical insights with our solutions, we published them in a biomedical journal [30].

AUTHOR CONTRIBUTION: The author of this thesis co-developed major parts of the presented study, especially with regard to the VA approach. The author also contributed significantly to the data preparation, the data analysis, and the evaluation of the results and wrote and reviewed substantial parts of the original publication.

ABSTRACT: *Purpose:* To analyze the use of deviation maps to understand thickness changes of inner retinal layers in OCT data. To test a new VA method with reduced complexity of OCT data analysis by comparing the layer thickness of children with type 1 diabetes mellitus to matched controls.

Methods: OCT was performed on unilateral eyes of 26 children with type 1 diabetes mellitus without diabetic retinopathy and 29 healthy children to obtain macular volume scans. Subsequently, segmented inner retinal layers were analyzed using VA. Deviation maps were generated to readily visualize thickness differences between both groups, and to investigate thickness changes of individual patients in relation to the control group.

Results: In deviation maps of the patient group, the total thickness of the retina demonstrated localized, irregular areas of thinning (mean deviation and standard deviation) involving foveal center, inner macula, and inferior-nasal outer macula ($-9.31 \mu\text{m} \pm 1.73 \mu\text{m}$; $p < 0.05$). Similarly, retinal nerve fiber layer showed continuous and localized areas of thinning in both inner and outer macula, extending nasally ($-5.45 \mu\text{m} \pm 4.31 \mu\text{m}$; $p < 0.05$). In deviation maps of individual patients, the total thickness of the retina and the thickness of the inner retinal layers revealed remarkable changes that were present between patients at both late and early stages of diabetes.

Conclusion: The VA method simplifies the in-depth analysis of OCT volume data from different groups and is effective in detecting retinal thickness changes in children with diabetes. It can be easily adopted in a clinical set-up and intuitively used in complex multidisciplinary studies.

ORIGINAL PUBLICATION: [30] — R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Götze, A. Jünemann, H. Schumann, and O. Stachs. “Deviation maps for understanding thickness changes of inner retinal layers in children with type 1 diabetes mellitus.” In: *Current Eye Research* 44.7 (2019). DOI: [10.1080/02713683.2019.1591463](https://doi.org/10.1080/02713683.2019.1591463)

FURTHER READING: Prior to the evaluation of the study data with our VA approach, we used the current approach based on the statistical analysis of ETDRS grids to compare the T1DM patient group with the control group. We also used the current analysis approach to study variations in retinal thickness in pediatric patients with mild to moderate chronic kidney disease:

- ▷ A. Götze, S. von Keyserlingk, S. Peschel, U. Jacoby, C. Schreiber, B. Köhler, S. Allgeier, K. Winter, **M. Röhlig**, A. Jünemann, R. Guthoff, O. Stachs, and D.-C. Fischer. “The corneal subbasal nerve plexus and thickness of the retinal layers in pediatric type 1 diabetes and matched controls.” In: *Scientific Reports* 8.1 (2018). DOI: [10.1038/s41598-017-18284-z](https://doi.org/10.1038/s41598-017-18284-z)
- ▷ R. K. Prakasam, A. Götze, S. von Keyserlingk, A. Jünemann, **M. Röhlig**, O. Stachs, and D.-C. Fischer. “Spectral-Domain Optical Coherence Tomography for Determination of Retinal Thickness in Pediatric Patients with Mild-To-Moderate Chronic Kidney Disease: A Cross-Sectional Study.” In: *Current Eye Research* 44.2 (2018), pp. 206–211. DOI: [10.1080/02713683.2018.1522649](https://doi.org/10.1080/02713683.2018.1522649)

Based on our observations in these two studies, we first introduced the basic idea of using VA to evaluate retinal OCT data in cross-sectional studies. We then presented our initial VA results of the re-evaluated study data at ophthalmic conferences, including the largest ophthalmic research meeting, ARVO 2009:

- ▷ **M. Röhlig**, A. Jünemann, D.-C. Fischer, R. K. Prakasam, O. Stachs, and H. Schumann. “Visual Analysis of Retinal OCT Data.” In: *Klinische Monatsblätter für Augenheilkunde* 234.12 (2017), pp. 1463–1471. DOI: [10.1055/s-0043-121705](https://doi.org/10.1055/s-0043-121705)
- ▷ O. Stachs, A. Jünemann, H. Schumann, and **M. Röhlig**. “Visual Analytics in der Augenheilkunde.” In: *Spitzenforschung in der Ophthalmologie* 2.1 (2017), pp. 160–164

- ▷ R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Jünemann, H. Schumann, and O. Stachs. “Visual Analysis: A new method to analyze OCT thickness data of retinal layers.” In: *Investigative Ophthalmology & Visual Science*. Vol. 59. Proceedings of the ARVO annual meeting 9. ARVO, 2018

10.1 INTRODUCTION

Optical Coherence Tomography (OCT) [17], a noninvasive imaging technology, attained an important place in the field of diagnostics, especially in detecting and managing a wide range of retinal and nerve diseases. The latter is relevant even in pediatric patients with Diabetes Mellitus (DM) since subclinical neuropathy with the inherent risk of progression to severe diabetic neuropathy is common among children and adolescents with Type 1 Diabetes Mellitus (T1DM) [207]. It has been hypothesized that the neurodegenerative changes, i.e., the apoptotic loss of neural cells in the retina, occur well before clinically visible vascular damage or Diabetic Retinopathy (DR) [208–210]. Thus, OCT has been suggested as a powerful tool for monitoring neurodegenerative changes in T1DM patients. In fact, we [90] and others [177, 178] have successfully used OCT to demonstrate significant thinning of distinct neuronal layers within the retina in pediatric and adult DM patients with no or minimal DR compared to healthy controls.

On the one hand, improved OCT image quality with a high axial and lateral resolution offers a unique possibility of differentiating and presenting substructures of the retina. On the other hand, analysis of OCT data is still challenging, because of both, the limited features of available analysis software and the associated increase in complexity when multiple three-dimensional volume datasets have to be studied. Commercial OCT software enables segmentation of intraretinal layers and production of thickness maps by measuring every single data point. However, for a comparative evaluation of multiple datasets, the Early Treatment Diabetic Retinopathy Study (ETDRS) [211] grid is commonly utilized. The ETDRS grid reduces the data by extracting only few spatially averaged thickness measurements from a given thickness map relative to 9 grid sectors, i.e., a central foveal ring with 1 mm diameter, an inner macula ring (pericentral) with 3 mm diameter, and an outer macula ring (peripheral) with 6 mm diameter. The inner and outer rings are divided into four quadrants (nasal, temporal, superior, and inferior). Thus, an entire OCT dataset is represented by only one arithmetic mean thickness value for each grid sector and retinal layer. While this Current Analysis (CA) approach allows a quick overview on thickness changes in predefined anatomical areas, it consequently reduces the amount of information and even bears the risk of information loss secondary to false negative findings due to averaging artifacts.

Recently, we reported [27, 98] on the development of a new software tool, based on Visual Analytics (VA) methods, i.e., a special field of research in computer science. VA supports in simplifying the analysis of large OCT volume datasets by utilizing both, automated algorithms and human abilities to visually detect trends and patterns in the data. Here, we present new software features for a Map-based Analysis (MA) of cross-sectional study data with high spatial accuracy, and discuss their advantages and disadvantages in relation to the CA approach. We hypothesize that our VA-driven software tool is a meaningful extension to manual or semi-manual conventional approaches for data analysis and interpretation, especially with regard to subtle retinal changes. Therefore, we re-investigated previously obtained data from a cross-sectional study in pediatric T1DM patients and healthy controls [90].

10.2 MATERIALS AND METHODS

In our previously reported cross-sectional study [90], the participated pediatric patients were suffering from T1DM for at least 12 months and had an inclusion criteria of 6–18 years of age at time of enrollment, C-peptide below 0.3 ng/ml, stable therapeutic regimen with insulin for at least three months. Exclusion criteria included limited ability to cooperate, pregnancy, tumoral diseases, history of any other systemic disease including neurological diseases and idiopathic small fibre neuropathy, clinical evidence of diabetic peripheral neuropathy, or ophthalmological diseases especially myopia of more than 6 diopters, corneal and retinal disorders. Demographic and clinical data for characterization of the disease (duration of the disease, mode of therapy, mean daily insulin dosages, and mean HbA_{1c} during the last year) were gained by interview and/or chart review. The participated pediatric T1DM patients were being treated at the university between October 2012 and December 2013 and healthy age-matched controls were recruited from different schools in Rostock. The participants of the control group were further investigated in our out-patient clinic, and the children who presented without any history of systemic or ocular diseases were included in the study. The study was conducted in accordance with the Declaration of Helsinki after receiving appropriate ethics committee approval from the institutional review board (Rostock University Medical Centre Ethics committee). The nature of the study was explained to all participants and/or to their parents prior to obtaining informed written consent for participating in the study.

All participants underwent an ophthalmological investigation, i.e., assessments of visual acuity, intraocular pressure, slit lamp biomicroscopy, fundus examination, and finally unilateral spectral-domain OCT (SPECTRALIS HRA + OCT, Heidelberg Engineering, Heidelberg, Germany). In this process, multiple cross-sectional B-scans of the macular regions of the central retina were recorded. The instrument was equipped with

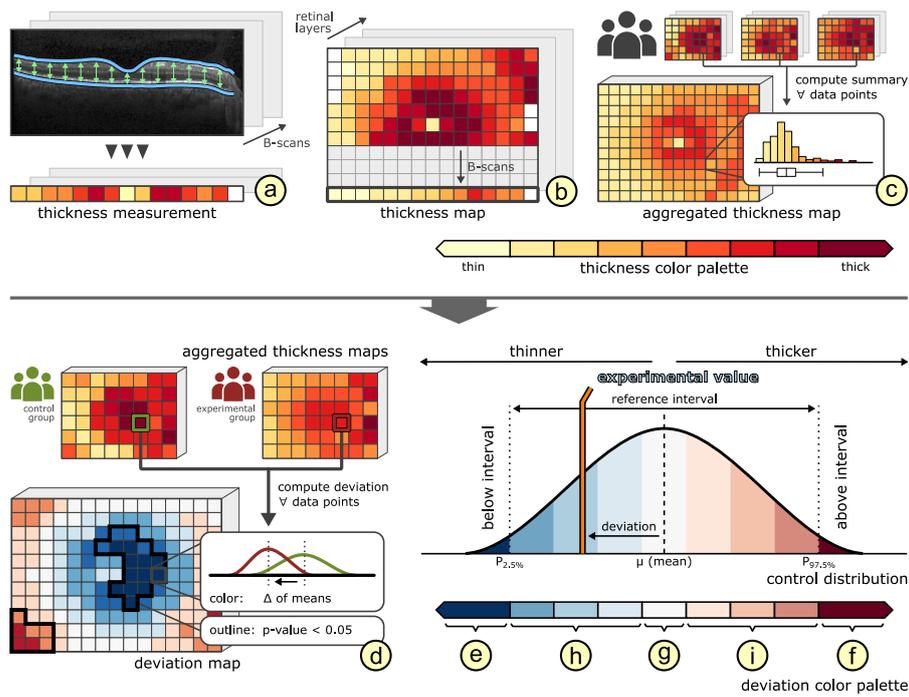


Figure 10.1: Generation of TMs, RMs, and DevMs. TMs were computed by measuring the distance between layer boundaries in every B-scan of a dataset (a, b). RMs were compiled by aggregating a set of user-specified base TMs (c). DevMs were derived by comparing either a TM to a RM or two RMs (d). The color scale explains the thickness deviation of retinal layers compared via DevMs. Dark colors indicate values outside the reference interval (e, f), white color indicate values close to the reference mean (g), and intermediate colors (blue or red) denote either thinning (h) or thickening (i).

an active eye tracking system (TruTrack™) to ensure precise registration of OCT scans to fundus cSLO image. The predetermined scanning pattern of a rectangular area (volume scan) was measuring $20^\circ \times 15^\circ$ ($5.8 \text{ mm} \times 4.4 \text{ mm}$), centered on the fovea to generate 19 consecutive macular B-scans at an interval of $243 \mu\text{m}$ per participant. Automatic real-time tracking was used for averaging of 9 frames per B-scan. All OCT images were acquired at high-speed scanning modality with an axial and lateral resolution of $3.9 \mu\text{m} \times 11 \mu\text{m}$, and with a signal quality of at least 27 dB. The volume datasets from all participants were further processed using the device-specific software (Heidelberg Eye Explorer®, HEYEX version 6.5) to perform retinal layer segmentation. Finally, per participant and per retinal layer three averaged thickness values, reflecting the central foveal ring, the inner macula ring, and the outer macula ring of the ETDRS grid were obtained [90].

VISUAL ANALYTICS APPROACH OF ANALYZING OCT THICKNESS DATA OF INTRARETINAL LAYERS: The previously obtained segmented OCT volume scans were transferred to VA software for systematic

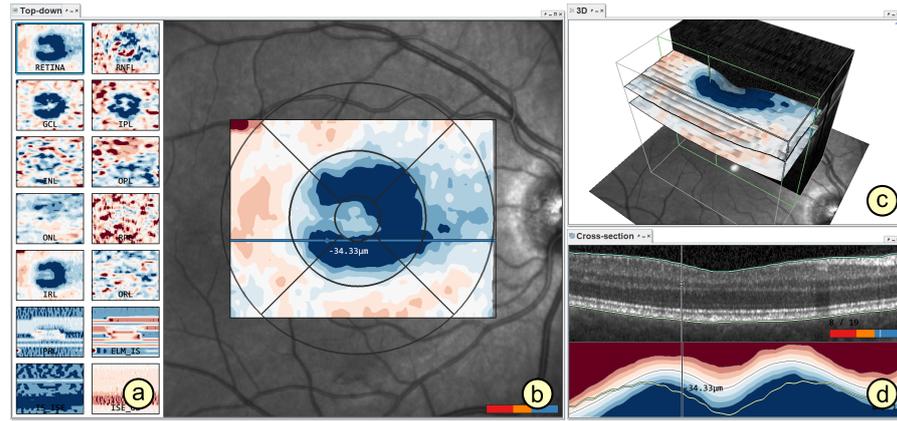


Figure 10.2: Overview of the graphical user interface of our VA software for analysis of individuals. The visualizations show the data of a patient in relation to the control group. Depicted were (a) an overview of Deviation Maps (DevMs) of all segmented retinal layers, (b) a top-down view showing a fundus image overlaid with a DevM of a selected layer, (c) a 3D view showing a volume visualization of OCT data together with boundary surfaces of a selected layer colored by the DevM, and (d) a cross-sectional view for inspecting individual B-scans and associated thickness profiles extracted from the DevMs.

quality assessment and subsequent analysis of measurements of the Total Thickness of the Retina (TR) and the thickness of the inner retinal layers, including the Retinal Nerve Fiber Layer (RNFL), the Ganglion Cell Layer (GCL), the Inner Plexiform Layer (IPL), and Inner Nuclear Layer (INL). Scans with considerable image artifacts, e.g., very noisy or distorted B-scans, and incomplete or failed retinal layer segmentation were excluded from the analysis. A detailed description of the technical features, design, and functionality of the VA software has been published recently [88].

To support a spatially precise analysis of cross-sectional study data, we designed new comparative visualizations based on three types of maps: Thickness Maps (TMs), aggregated Reference Maps (RMs), and Deviation Maps (DevMs). Figure 10.1 (a, b) illustrates the generation of TMs by measuring the pointwise distance between boundaries of each segmented retinal layer in every B-scan of a given dataset. Based on the computed TMs, aggregated RMs were compiled. Each point of these combined maps encloses summary statistics, e.g., mean, percentiles, and standard deviation, of a set of user-supplied base TMs (Fig. 10.1 (c)). Base TMs from datasets with varying properties of different laterality, scan size, or resolution were transformed to correctly summarize thickness values at corresponding spatial locations. To support the comparison of multiple datasets, DevMs were derived from selected TMs and RMs (Fig. 10.1 (d)). DevMs enable a point by point comparison of thickness values by encoding deviations between TMs and RMs using

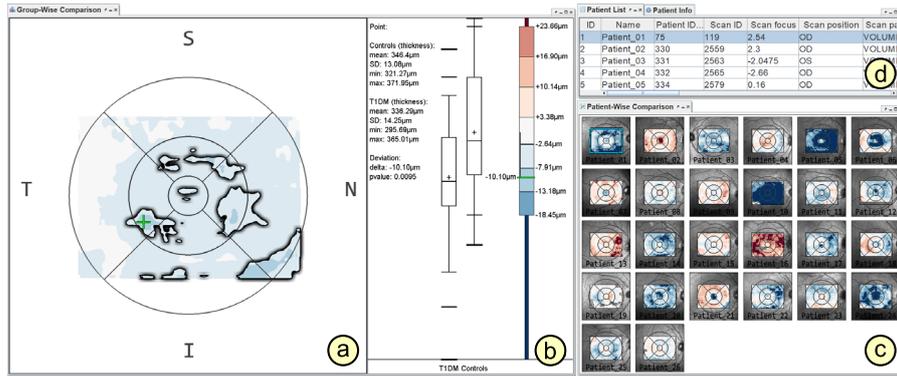


Figure 10.3: Overview of the graphical user interface of our VA software for analysis of groups. The visualizations show the TR of the patient group in relation to the control group. Depicted were (a) a view showing a DevM for comparing the mean thickness of both groups in general, and (b) a linked measurement view, depicting automated statistical summaries of both groups for an interactively selected point on the DevM (green cross). The box-and-whisker plots represent the difference between both groups graphically and the numerical labels provide descriptive values of the comparison results. In addition, an overview of DevMs from all individual members of the patient group compared to the control group were shown (c) together with a list of all members of the patient group and associated metadata (d).

diverging color palettes. Thus, they basically demonstrate the thickness difference in microns of a retinal layer either for an individual relative to the averaged data of a reference group (TM compared to RM) or for two different groups (RM compared to RM).

The color scale defined to our DevMs explains the deviation of the thickness of the retinal layers in comparison to reference data obtained from the control group, i.e., the color scale reflects the interval between 2.5th and 97.5th percentiles. The darkest colors indicate thinner (dark blue) or thicker (dark red) areas below and above the interval boundaries, respectively (Fig. 10.1 (e) and (f)). White areas indicate no or only minimal difference in thickness compared to the mean of the reference group (Fig. 10.1 (g)). For the distinction of areas with intermediate deviations, the color scale was gradually extended from white towards both sides of the scale to the interval minimum and maximum (Fig. 10.1 (h) and (i)). Further, as another new feature of our approach, the thickness changes observed on DevMs are tested for statistical significance ($p < 0.05$) by applying Student's t-test at every single point in the maps. Neighboring points with significant differences are combined to form connected significant areas highlighted within the maps.

We utilize these maps to support the analysis of retinal thickness data from individuals and groups. Overviews of respective graphical user interfaces are given in Figures 10.2 and 10.3. TMs, RMs, and DevMs were integrated in multiple coordinated views that are interlinked and

	Patients (10f/16m)	Controls (15f/14m)
Age [years]	14.50 ± 2.23	14.50 ± 2.00
Height [SDS]	0.15 ± 0.89	0.60 ± 1.20
Weight [SDS]	0.31 ± 0.97	0.60 ± 0.90
BMI [SDS]	0.30 ± 0.96	0.40 ± 0.70
BP _{sys} [SDS]	1.95 ± 1.20	1.45 ± 1.15
BP _{dias} [SDS]	0.90 ± 0.76	0.48 ± 0.99
Duration of disease [years]	5.60 ± 4.20	NA
Mean daily insulin dosage [IU/kg]	0.39 ± 0.14	NA
Actual HbA _{1c} [%]	8.96 ± 1.86	NA
Mean HbA _{1c} [%]	8.73 ± 1.51	NA

Table 10.1: Anthropometric and clinical characteristics of patients and controls. Results are given as mean and standard deviation for all parameters. SDS: Standard deviation score, BMI: Body mass index, BP: blood pressure, HbA_{1c}: Hemoglobin A_{1c}, Mean HbA_{1c}: two different time points.

can be freely arranged to associate different aspects. Our VA software also supports automated measurements to evaluate the visualized data. Points, user-specified regions, and highlighted significant areas on the maps can be interactively selected to show respective statistical summaries (Fig. 10.3 (b)). The distributions of the thickness measurements are visualized as box-and-whisker plots next to the applied color scales with numerical labels of descriptive values. Graphically representing both the patients and the controls as statistical plots side by side helps to understand their differences with respect to the selections made in the maps. Optionally, our software supports relating the MA results to an enhanced ETDRS grid-based analysis of retinal thickness data [88].

10.3 RESULTS

Previously, we applied OCT and conventional CA [90] to investigate pediatric T1DM patients (n=26) and healthy controls (n=30), which showed significant thinning of the TR and the thickness of GCL and IPL in the inner macula ring of the ETDRS grid¹. Now, we re-evaluated these results with our new VA software. The descriptive results from both anthropometric and clinical characteristics of patients [90] and controls are summarized in Table 10.1. The previously acquired multiple volume scans from every single participant have been processed for a quality

¹ All results of [90] are available as supplemental material at doi:10.1038/s41598-017-18284-z

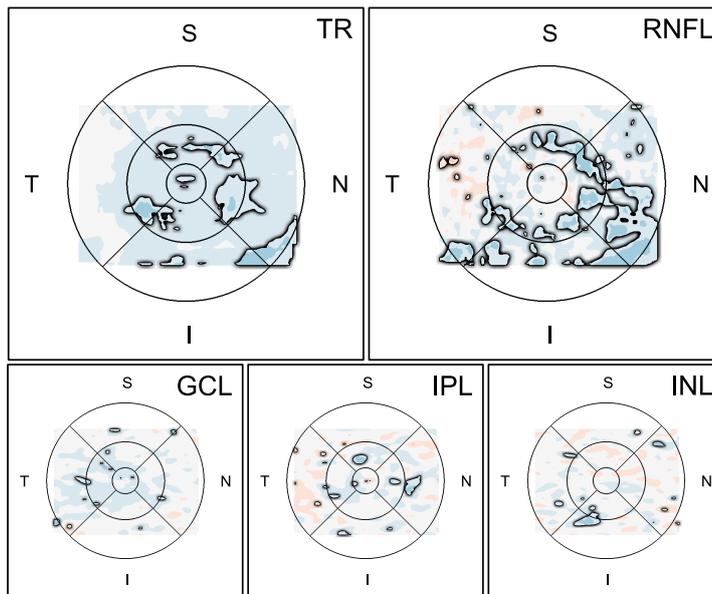


Figure 10.4: DevMs of TR and inner retinal layers (RNFL, GCL, IPL, and INL) for group-wise comparison between patients and controls. Color denotes macular areas of thinning (blue) and thickening (red), and black outlines highlight areas of statistically significant differences ($p < 0.05$). Mean Deviation and Standard Deviation ($MD \pm SD$) values were calculated by measuring highlighted areas on DevMs. The DevM of TR demonstrated localized, irregular areas of significant thinning ($MD \pm SD$; $-9.31 \mu\text{m} \pm 1.73 \mu\text{m}$) involving foveal center, inner macula regions, and inferior-nasal quadrant of the outer macula ring. The DevM of RNFL demonstrated both continuous and localized areas of significant thinning ($MD \pm SD$; $-5.45 \mu\text{m} \pm 4.31 \mu\text{m}$) in the inner and outer macula rings, extending nasally. Other retinal layers (GCL, IPL, and INL) showed only minimal thickness changes.

assessment using VA, which resulted in a rejection of several scans with various quality issues. Failed retinal layer segmentation has been the primary reason for the rejection. The resulted one best of scans for 26 children suffering from T1DM for $5.6 \text{ years} \pm 4.2 \text{ years}$ and for 29 out of 30 healthy children were selected for further analysis. All selected scans were aligned to enable a precise point by point comparison, and minor segmentation errors if any were corrected.

Using our VA software, DevMs for all inner retinal layers were generated for (i) group-wise comparisons between patients and controls and (ii) patient-wise comparisons between individual patients and the control group.

10.3.1 Group-wise comparison between patients and controls

The results from the group-wise comparison are demonstrated in Figure 10.4. The TR showed localized and irregular areas of significant

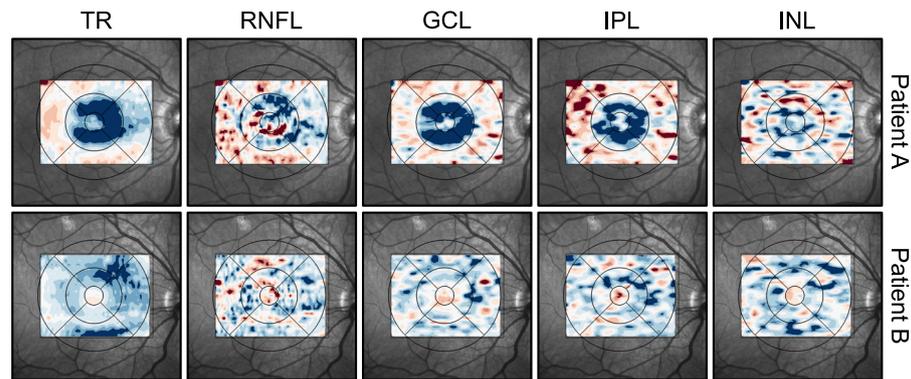


Figure 10.5: DevMs of TR and inner retinal layers (RNFL, GCL, IPL, and INL) for patient-wise comparison between two example patients and the control group. Patient A had a longer duration of disease and a high HbA_{1c} value, whereas patient B had a shorter duration of disease with good control of blood glyceemic levels. MD ± SD values for considerable thinning (dark blue) or thickening (dark red) were calculated by measuring areas with thickness values outside the control interval on DevMs. Patient A clearly demonstrated a ‘bullseye’ like pattern of normal foveal center surrounded by thinning of inner macula ring, readily visible in TR and the layers GCL and IPL (MD ± SD; $-33.36 \mu\text{m} \pm 6.19 \mu\text{m}$, $13.55 \mu\text{m} \pm 3.91 \mu\text{m}$ and $-10.16 \mu\text{m} \pm 2.08 \mu\text{m}$, respectively). Adjacent localized areas of thickening were noted in all layers, primarily evolving at the outer macula ring. Mean thickening of TR and other layers (mean of mean RNFL, GCL, IPL, and INL) was $25.60 \mu\text{m} \pm 4.04 \mu\text{m}$ and $10.13 \mu\text{m} \pm 0.73 \mu\text{m}$, respectively. DevMs of patient B demonstrated thinning of TR ($-21.63 \mu\text{m} \pm 3.27 \mu\text{m}$) involving both inner and outer macula rings, while all other layers exhibited only minimal thickness changes compared to the layers of patient A.

thinning in the foveal center, all four quadrants of inner macula ring, and in the inferior-nasal outer macula regions. The RNFL demonstrated both continuous and localized areas of significant thinning in both the inner and outer macula rings, with a greater extent towards the nasal quadrant. Other retinal layers, including GCL, IPL, and INL, showed only minimal and localized areas of thickness changes.

10.3.2 Patient-wise comparison between individual patients and the control group

For every patient, 5 DevMs of the TR and inner retinal layers (RNFL, GCL, IPL, and INL) were generated and subjectively analyzed for the presence or absence of thickness changes in relation to the individual characteristics of the disease, i.e., duration of T1DM and glyceemic control. Eleven out of 26 patients (42 %) suffered from T1DM for more than 5 years (range, 5.0–15.5 years) and presented with poor glyceemic control, i.e., a mean HbA_{1c} above 7.5 % (range, 7.8–11.3 %). In 6 patients (23 %),

duration of disease was less than five years (range, 1.2–4.0 years) and good glycemic control (mean HbA_{1c} range, 6.7–7.4 %) was noted. The remaining 9 patients (35 %) presented either with a duration of disease for more than 5 years and good glycemic control or less than five years of T₁DM and poor glycemic control. Subjective analysis of DevMs of individual patients (patient by patient analysis) in both the higher (duration of T₁DM > 5 years and poor glycemic control) and lower (duration of T₁DM < 5 years and good glycemic control) categories revealed remarkable thickness changes with varying degree and extent between subjects for all inner retinal layers. As an example, DevMs of one selected patient from each of these categories are given in Figure 10.5. Patient A suffered from diabetes for 13 years and presented with a mean HbA_{1c} value of 11.3 %, whereas patient B suffered from T₁DM for 1.8 years and showed an overall good glycemic control (mean HbA_{1c} of 6.7 %). In patient A, a normal foveal center was surrounded by areas of considerable thinning (thickness values outside the control interval) at the inner macula ring, a classic ‘bullseye’ like appearance more obvious in the TR and layers GCL and IPL. In contrast, DevMs of patient B revealed only minimal deviations from mean thickness values of controls. However, a diffuse area of thinning involving superior, inferior, and nasal quadrants of the inner and outer macula rings was noted in TR.

10.4 DISCUSSION

Ocular biomarkers or surrogate measures of a disease are important and useful for the early diagnosis of ocular [212, 213] and neurodegenerative diseases [197, 214–216]. OCT is considered an important tool to identify biomarkers and is widely used in the fields of ophthalmology and neuro-medicine to quantify and study neurodegenerative changes in the eye [197, 212, 216]. Neurodegenerative changes of the retina that occur at very early stages in children with T₁DM serve as a biomarker that needs to be closely assessed, to prevent or control the progression of the disease by implementing an early intervention.

In the present study, MA supports detecting and measuring thickness changes in the macular region as distinct patterns on comparative DevMs, varying between different layers of the retina in children with T₁DM without clinical evidence of DR. Although the overall results from the group-wise comparison, i.e., the presence of statistically significant thinning in the TR and layers RNFL and GCL in T₁DM patients compared to healthy controls, are similar for both methods of CA [90] and MA, the latter allows the unique visualization and measurement of topographical patterns of subtle changes on DevMs. Among the investigated inner layers, RNFL had demonstrated a greater amount of thinning in the nasal, superior, and inferior quadrants of both the inner and outer macula rings, similar to the results found in recent OCT studies which reported selective thinning of RNFL [22] or both RNFL and GCL [96].

Evaluation of patient-wise comparison results, i.e individual patients against the control group, revealed thickness changes in *DevMs* of all patients. However, these changes were apparently not related to glycemic control (good/poor) and duration of the disease (long/short). This is well in line with published results obtained by means of the *CA* approach (e.g., [178, 217] or [90]). In contrast, Gundogan et al. [96] and Tekin et al. [22] have shown significant negative correlations between either duration of *T1DM* or mean HbA_{1c} and thickness of *GCL* and/or *RNFL*, respectively. We hypothesize that these contrasting results mainly reflect differences in the methods used for *OCT* data analysis. The absence of any clinical signs of *DR* together with the missing association between retinal thickness changes and duration of disease or glycemic control point to an abnormal response of these structures to diabetes-related metabolic changes (e.g., [218] or [90]). Thus, a major concern is to test meticulously the clinical importance of the thickness changes observed in the present study to avoid any mislead on early diagnosis of neurodegenerative changes and consecutive overtreatment. Therefore, we highly recommend carefully designed longitudinal cohort studies together with *MA* approach of data analysis to precisely assess neurodegenerative alterations of the retinal layers to the time course of the disease and to relate morphological and functional aspects to long-term glycemic control and therapeutic regimen. Further, we strongly support a secure transition of *VA* methodologies into the clinical setting for an efficient analysis and investigation of retinal and other systemic pathologies. This will promote an intuitive determination of surrogate measures of a disease in complex multidisciplinary studies.

In the past, retinal layer thickness measurement studies [176–178, 217, 219, 220] have used different *OCT* devices from various manufacturers along with device-specific software to analyze *OCT* data by either manual or semi-manual *CA* approaches, which are generally based on data reduction via *ETDRS* grids. Similar to our approach, *DevMs* have been used to compare thickness data of an individual patient to predefined normative data and to study follow-up scans of an individual [176, 178]. However, existing approaches did not consider the comparison of multiple datasets from different study groups and in-depth analysis of subtle and localized retinal changes so far. In this context, our new approach provides a unique feature of applying patient-wise and group-wise *DevMs* for a direct comparison of two different groups (e.g., patient group and control group). To our knowledge, no equivalent approach exists of *VA*-driven interactive visualization and measurement of thickness data of intraretinal layers. Further, the *VA* software is highly flexible in defining groups and in the interactive generation of *DevMs*. Therefore, for matched study groups the *DevM*-based comparison is automatically age-adjusted and it facilitates a targeted and thorough analysis of subgroups, e.g., comparison of retinal thickness changes in patients with different treatment strategies. The current study, however,

has not presented subgroup analysis results due to smaller sample sizes of subgroups derived from the **T1DM** group.

Regarding the advantages and disadvantages **MA** in relation to **CA**, we argue that **CA** is often based on considerable data reduction, where aggregated mean thickness values for each coarse sector of **ETDRS** grids were calculated and then used to statistically analyze differences in retinal layer thickness between layers and/or between study subjects. This grid-based analysis approach provides an overview of area-based thickness changes in the macular region and the standardized **ETDRS** sectors facilitate comparison of analysis results between different studies. Yet, subtle information on local thickness changes and their topographical extent or irregular patterns, e.g., the sector-overlapping localized and irregularly shaped areas of thinning noticed on **DevM** of **TR** in group-wise comparison between **T1DM** patients and controls, reported in the present study may go undetected. On the other hand, our **MA** approach [27] extends the **CA** in that it provides an advanced fully automated computation together with 2D and 3D visualizations of **OCT** data. **VA** is useful for studying individual patients for diagnosis and management of the disease as well as for the comparison of multiple large **OCT** datasets from different groups in cross-sectional studies. Most importantly, thickness information is derived for every single point in the data to generate color-coded **DevMs**. This helps to understand the distribution of local thickness changes with greater spatial specificity and without information loss, in addition to the spatial summaries provided by **CA**. Generally, it is hardly possible to directly compare the point-to-point thickness changes from **DevMs** to statistical results of **ETDRS** defined macular regions from **CA**. On this account, the application of statistical tests on every single point on **DevMs** for group comparisons and automatically highlighting and measuring respective significant areas are effective features that have been incorporated in our **MA** approach to advance **OCT** analysis. Consequently, the dedicated **VA** software allows to generate results faster than a manual or semi-manual **CA** and without having to rely on external software for statistical analysis.

In conclusion, our **VA** methods are a meaningful extension to **CA**. The **MA** eases the in-depth analysis of **OCT** data and is effective in detecting and demonstrating subtle thickness changes of inner retinal layers in children with **T1DM** even at the early stages of the disease, which might provide useful insights in early diagnosis and treatment management of **DR**.

THICKNESS OF INTRARETINAL LAYERS IN PATIENTS WITH TYPE 2 DIABETES MELLITUS DEPENDING ON A CONCOMITANT DIABETIC NEUROPATHY: RESULTS OF A CROSS-SECTIONAL STUDY USING DEVIATION MAPS FOR OCT DATA ANALYSIS

SCIENTIFIC CONTRIBUTION: In this chapter, we evaluate our novel [VA](#) solutions in a cross-sectional study with adult patients suffering from type 2 diabetes mellitus. Especially, we introduce an advanced measurement method for retinal defects made visible with our visualization techniques. Compared to the current analysis approach, this improves the quantification of irregular-shaped regions with abnormal variations in retinal layer thickness. The new medical insights were published in a biomedical journal [[31](#)].

AUTHOR CONTRIBUTION: The author of this thesis co-developed major parts of the presented study, especially with regard to the [VA](#) approach. The author also contributed significantly to the data preparation, the data analysis, and the evaluation of the results and wrote and reviewed substantial parts of the original publication. The individual contributions of all authors are listed in [[31](#)].

ABSTRACT: [OCT](#) supports the detection of thickness changes in intraretinal layers at an early stage of diabetes mellitus. However, the analysis of [OCT](#) data in cross-sectional studies is complex and time-consuming. We introduce an enhanced map-based analysis and demonstrate its effectiveness in detecting early changes in intraretinal layer thickness in adults with type 2 diabetes mellitus compared to the current analysis with the common [ETDRS](#) grid. To this end, we obtained [OCT](#) scans of unilateral eyes from 33 patients with type 2 diabetes mellitus without diabetic retinopathy and 40 healthy controls. The patients were categorized according to concomitant diabetic peripheral neuropathy. The results of the map-based and current analysis demonstrated statistically significant differences in retinal thickness between patients and controls. Thinning was most pronounced in total retinal thickness and the thickness of the inner retinal layers in areas of the inner macular ring, selectively extending into areas of the outer macular ring and foveal center. Patients with clinically proven diabetic peripheral neuropathy showed the strongest thinning of the inner retinal layers. The map-based analysis showed additional areas of thinning whereas the current analysis tended to underestimate thickness changes, especially

in areas with localized thinning. We conclude that the map-based analysis enables a precise analysis of retinal thickness data and contributes to the understanding of localized changes in intraretinal layers in adults with type 2 diabetes mellitus.

ORIGINAL PUBLICATION: [31] — R. K. Prakasam, A. Matuszewska-Iwanicka, D.-C. Fischer, H. Schumann, D. Tschöpe, B. Stratmann, H.-J. Hettlich, R. F. Guthoff, O. Stachs, and **M. Röhlig**. “Thickness of intraretinal layers in patients with type 2 diabetes mellitus depending on a concomitant diabetic neuropathy: results of a cross-sectional study using deviation maps for OCT data analysis.” In: *Biomedicines* 8.7 (2020). DOI: [10.3390/biomedicines8070190](https://doi.org/10.3390/biomedicines8070190)

FURTHER READING: Prior to the presented evaluation of the study data, we introduced our advanced measurement method for local variations of retinal layer thickness:

- ▷ **M. Röhlig**, C. Schmidt, R. K. Prakasam, O. Stachs, and H. Schumann. “Towards Accurate Visualization and Measurement of Localized Changes in Intraretinal Layer Thickness.” In: *Poster at the IEEE Workshop on Visual Analytics in Healthcare (VAHC)*. 2019

On this basis, we presented our first VA-based study results at two national and international ophthalmic conferences:

- ▷ O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, H. J. Hettlich, and **M. Röhlig**. “Visual Analytics of OCT data: Utility of deviation maps in describing retinal layer thickness changes.” In: *Investigative Ophthalmology & Visual Science*. Proceedings of the ARVO annual meeting. ARVO, 2019
- ▷ O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, B. Stratmann, H. J. Hettlich, and **M. Röhlig**. “Visual analysis of retinal OCT data in patients with type 2 diabetes mellitus.” In: *Ophthalmologe*. Proceedings of the DOG congress. Springer Medicine, 2019

11.1 INTRODUCTION

Optical Coherence Tomography (OCT) [17] is an established method for detecting early signs of neurodegenerative changes in patients with Diabetes Mellitus (DM). Previous retinal OCT studies by us [30] and others [21, 22, 220, 221] have clearly demonstrated changes in the thickness profiles of distinct neuronal layers, primarily involving the Retinal Nerve Fiber Layer (RNFL), Ganglion Cell Layer (GCL) and Inner Plexiform Layer (IPL) in pediatric and adult DM patients at an early stage

or even before any manifestation of clinical signs of microangiopathy (e.g., Diabetic Retinopathy (DR)). Such changes are thought to serve as biomarkers for the identification of patients at risk of developing DR. In this regard, software algorithms for the analysis of OCT data, in particular those for image segmentation [222], are a prerequisite for detailed analyses of individual retinal layers and thus are helpful for studying distinct ocular or systemic conditions. In general, the process of boundary identification and segmentation is followed by the generation of topographical color-coded thickness maps by measuring every single point of OCT data. While these color-coded thickness maps are spatially informative, it is hard to relate such data to the widely used Early Treatment Diabetic Retinopathy Study (ETDRS) grid [89]. In fact, the established analysis approach of thickness profiles per retinal layer is currently based on averaged thickness values with respect to the ETDRS grid cells. The grid cells are defined by a central foveal ring with a 1 mm diameter, an inner macula ring (pericentral) with a 3 mm diameter and an outer macula ring (peripheral) with a 6 mm diameter. The inner and outer rings are divided into four quadrants (nasal, temporal, superior and inferior). To date, this grid-based analysis has been frequently used for the comparison of OCT findings between groups [221, 223]. While this approach allows a quick overview of thickness changes in predefined macular regions, it consequently reduces the available amount of information. Hence, subtle and localized thickness differences within one cell are likely to be overlooked due to data averaging.

To address these issues, we developed an analysis software based on Visual Analytics (VA) methods [27, 28]. VA is a special field in computer science that offers means for simplifying the analysis of multiple volumetric OCT datasets by combining automated algorithms and human abilities to visually detect trends and patterns in the data. Our VA software has been successfully tested on children with Type 1 Diabetes Mellitus (T1DM) and age-matched controls [30]. Most importantly, we demonstrated a new Map-based Analysis (MA) using comparative color-coded Deviation Maps (DevMs), which helped us to understand the distribution of local thickness changes between two different study groups with greater spatial specificity. Statistical tests were directly applied to every single point on the DevMs to automatically highlight macular areas with significant aberration from the control.

In the present study, we extend our design with a novel method for precisely measuring irregular shaped areas of differences between study groups on DevMs (cf. Sect. 11.2.3) and discuss its advantages in detecting early neurodegenerative changes in the inner retina. We demonstrate our approach by comparing results from the Current Analysis (CA) approach with ETDRS grids to results from our enhanced MA of thickness changes in the inner retinal layers in patients with Type 2 Diabetes Mellitus (T2DM) and controls of comparable age. Moreover, we

show how our VA software can be used to flexibly define and compare subgroups of patients with Diabetic Peripheral Neuropathy (DN).

11.2 MATERIALS AND METHODS

11.2.1 *Inclusion and exclusion criteria for patients and controls*

All T2DM patients were recruited from the Diabeteszentrum at Herz- und Diabeteszentrum Nordrhein-Westfalen, Germany, and the control group was assembled from colleagues at University Medical Center Rostock, Germany, on a voluntary basis. The study protocol was approved by the Institutional Review Board (Rostock University Medical Center Ethics Committee, Registration number: A 2015-0111 (dated 04/09/2015) with approved amendment (dated 26/02/2019)). All study procedures followed the tenets of the Declaration of Helsinki, and written informed consent was obtained from all subjects before enrolment.

Patients with a minimum age of 18 years, a clinical diagnosis of T2DM without DR, and an HbA_{1c} ranging between 6.5% and 9.5% were eligible. Individuals suffering from T1DM, diabetic macular edema, proliferative DR, mental disabilities or a history of serious systemic illness (e.g., tumoral disease or neurological disorders except DN) were not eligible. Likewise, individuals with uveitis, optic nerve pathology, vitreous hemorrhage, contact lens wear (soft and rigid), symptomatic dry eyes, conjunctivitis at the time of examination, severe hyperopia, previous corneal or retinal surgery, media opacities such as dense cataracts, corneal opacities and epiretinal membranes were excluded.

Study examinations included the gathering of demographic and clinical data for the characterization of the disease (duration of the disease, mode of therapy, mean daily insulin dosages and mean HbA_{1c}) by interview and/or medical chart review. All participants underwent comprehensive ophthalmic investigation, including unilateral retinal imaging (see below) and indirect ophthalmoscopy, to confirm the absence of DR. The Neuropathy Symptom Score (NSS) and Neuropathy Deficient Score (NDS) were used to test for DN. The patients with an NSS and/or NDS score ≤ 2 were classified as Diabetic Peripheral Neuropathy Negative (DN_{neg}) and the remaining patients as Diabetic Peripheral Neuropathy Positive (DN_{pos}).

11.2.2 *Spectral domain OCT*

Randomly selected unilateral eyes of all participants underwent retinal imaging using spectral-domain OCT (SPECTRALIS HRA + OCT, Heidelberg Engineering GmbH, Heidelberg, Germany). The non-invasive procedures of OCT imaging were explained prior to examination and all participants were instructed to steadily fixate on an objective system of

the device during the examination. The fixation was monitored by an active eye tracking system (TruTrack™) to ensure the precise registration of the OCT scans in the fundus image acquired with a Confocal Scanning Laser Ophthalmoscope (cSLO). Using the volume scan protocol, at least one volume scan was recorded for each participant. The scan presets were selected to closely match the area covered by the ETDRS grid. The volume scan size was defined to cover a rectangular area centered on the fovea and measuring at least $20^\circ \times 20^\circ$ (5.9 mm \times 5.9 mm). Per volume scan, at least 49 consecutive cross-sectional macular Cross-sectional Tomographic Depth Scans (B-scans), each consisting of 16 averaged images (ART mean), with an interval minimum of 120 μm , were obtained.

All OCT images were acquired at a high-speed scanning modality with an axial and lateral resolution of $3.9 \mu\text{m} \times 11 \mu\text{m}$ and with a signal quality of at least 27 dB. The device-specific software (Heidelberg Eye Explorer®, HEYEX version 6.5, Heidelberg Engineering GmbH, Heidelberg, Germany) was used for the retinal layer segmentation throughout the OCT datasets. Subsequently, the datasets were transferred to our VA software. The data transfer was established by means of a raw data export with the data format specification provided by the device manufacturer. Using the VA software, each dataset was screened to ensure data quality. In particular, data were checked for significant motion artifacts as well as for the accuracy and completeness of the automatically segmented retinal layers. The automated segmentations were verified by checking the generated thickness maps for artifacts (e.g., holes) and examining all individual B-scans for faulty layer boundary detection. Where possible, minor segmentation errors were corrected using the device software. Datasets which did not fulfill the quality criteria (e.g., very noisy or distorted B-scans and incomplete or failed retinal layer segmentation) were excluded. In this way, the best OCT scan per participant was selected. All selected scans were precisely aligned by determining the center of the fovea and the center of the optic disc as registration points. Finally, the center of the ETDRS grid was placed on the center of each fovea to allow accurate comparison during the interactive evaluation with the VA software.

11.2.3 Visual analysis

The technical design of our VA software and a detailed description of its main features regarding the CA and MA of retinal OCT thickness data have been published previously [27, 28, 30]. Here, we utilized the VA software to evaluate the previously recorded OCT data of T₂DM patients with and without DN relative to controls and to compare the results obtained with both methods.

OCT data were analyzed with ETDRS-based deviation grids [28] in CA and via DevMs [30] in MA. Basically, both methods demonstrate the absolute thickness difference of a retinal layer between groups, although

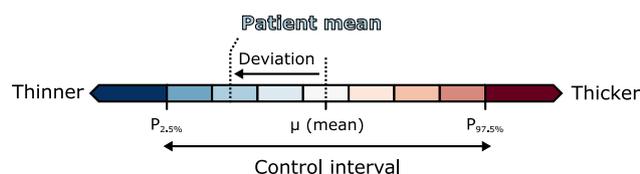


Figure 11.1: The color scale used to encode the deviation in mean thickness per grid cell or per map point. White indicates minimal patient deviation from the controls, and blue and red indicate negative (thinning) and positive (thickening) deviations, respectively. The degree of deviation is encoded by the color lightness, with the darkest colors representing values below or above the control interval.

with distinct differences regarding the presentation of the results. The colored grids offer a cell by cell overview on averaged thickness changes in predefined anatomical areas, whereas colored *DevMs* provide a point by point detail view on the actual shape and extent of thickness changes across different macular areas. The diverging color scale defined by the grids and maps explains the deviation in thickness of a retinal layer in comparison to reference data per grid cell or per map point (Fig. 11.1). In the present study, the reference data were obtained from the thickness distribution of the control group, i.e., color reflects the thickness deviation with respect to the interval between the 2.5th and 97.5th percentiles of the control data. White indicates minimal deviation from the mean of the control group, and blue and red denote negative (thinning) or positive (thickening) deviations, respectively. The degree of deviation is encoded by the color lightness, with the darkest colors representing values below or above the 2.5th and 97.5th percentiles. The color-coded deviations were further tested for statistical significance ($p < 0.05$) at every grid cell and map point. We selected Student's independent t-test for quantifying the comparison between *T2DM* patients and controls and Tukey's Honestly Significant Difference (*HSD*) test for multiple comparisons between the subgroups of DN_{neg} and DN_{pos} patients and controls. The resulting areas of significant differences were highlighted within the grids (orange cell borders and black cell labels) and within the maps (black outlines).

In addition, we used the *VA* software to automatically quantify the visualized differences in grids and maps. Based on the statistical test results, measurements were obtained of those grid cells in the *CA* and map areas in the *MA* that showed significant differences. Regarding the measurements in *MA*, a new method for the precise assessment of *DevMs* was integrated. The main idea was to use the *ETDRS* grid as a spatial frame of reference and to split the continuous significant areas on *DevMs* according to the anatomically distinct macular areas defined by its grid cells. Subsequently, the method involved measuring the significant fraction of the partitioned map areas within each superimposed *ETDRS* grid cell. This process included four automated steps: (i) point-based de-

	All T ₂ DM (n = 33)	DN _{pos} T ₂ DM (n = 21)	DN _{neg} T ₂ DM (n = 12)
Age (years)	60.4 ± 12.0	64.0 ± 9.0	54.0 ± 14.0
Gender (male/female)	21/12	15/6	6/6
Height (cm)	174 ± 10	176 ± 10	171 ± 9
Weight (kg)	99.3 ± 21.4	101.2 ± 21.4	95.9 ± 21.9
BMI (kg/m ²)	32.7 ± 6.3	32.5 ± 5.2	33.1 ± 8.0
Systolic blood pressure (mmHg)	133 ± 16	134 ± 18	130 ± 10
Diastolic blood pressure (mmHg)	79 ± 11	79 ± 12	77 ± 6
Duration of DM (years)	13.3 ± 8.8	16.0 ± 9.6	8.7 ± 4.5
Mean daily insulin dosage (IU/kg)	0.515 ± 0.323 (n = 30)	0.516 ± 0.335 (n = 19)	0.512 ± 0.317 (n = 11)
HbA _{1c} (%)	8.2 ± 0.8	8.2 ± 0.8	8.2 ± 0.9
NDS (pts)	1.7 ± 1.9	2.7 ± 1.8	0
NSS (pts)	2.7 ± 1.9	3.9 ± 0.9	0.5 ± 0.9

Table 11.1: Clinical and demographic data of all T₂DM patients and subgroups of DN_{pos} and DN_{neg} T₂DM patients.

tection of statistically significant areas on DevMs, (ii) anatomy-oriented partitioning of the areas based on the ETDRS grid, (iii) measurement of only significant fractions of the areas within each grid cell and (iv) checking the fractions of areas against a coverage threshold. The coverage threshold was defined as the minimum percentage of grid cell area required to be covered by significant points on the DevMs. It was introduced to avoid the misinterpretation of particularly small areas in the maps. We selected a threshold of at least 10 % coverage per ETDRS grid cell and the corresponding absolute Mean Deviation and Standard Error ($MD \pm SE$) values in the MA were compared to the respective values of the CA.

11.3 RESULTS

To test our VA software for OCT data analysis, we recruited a group of T₂DM patients and assembled a control group of comparable age. The mean age of the T₂DM patients (n = 33) and healthy controls (n = 40) was 60.4 years ± 12.0 years and 56.0 years ± 10.5 years, respectively. The anthropometric and clinical characteristics of the patients are summarized in Table 11.1.

Using the VA software, CA and MA results for thickness deviation between patients and controls were generated. The concepts and algorithms involved in this process have been described previously [27, 28]. The Total Thickness of the Retina (TR) and the thickness of the RNFL,

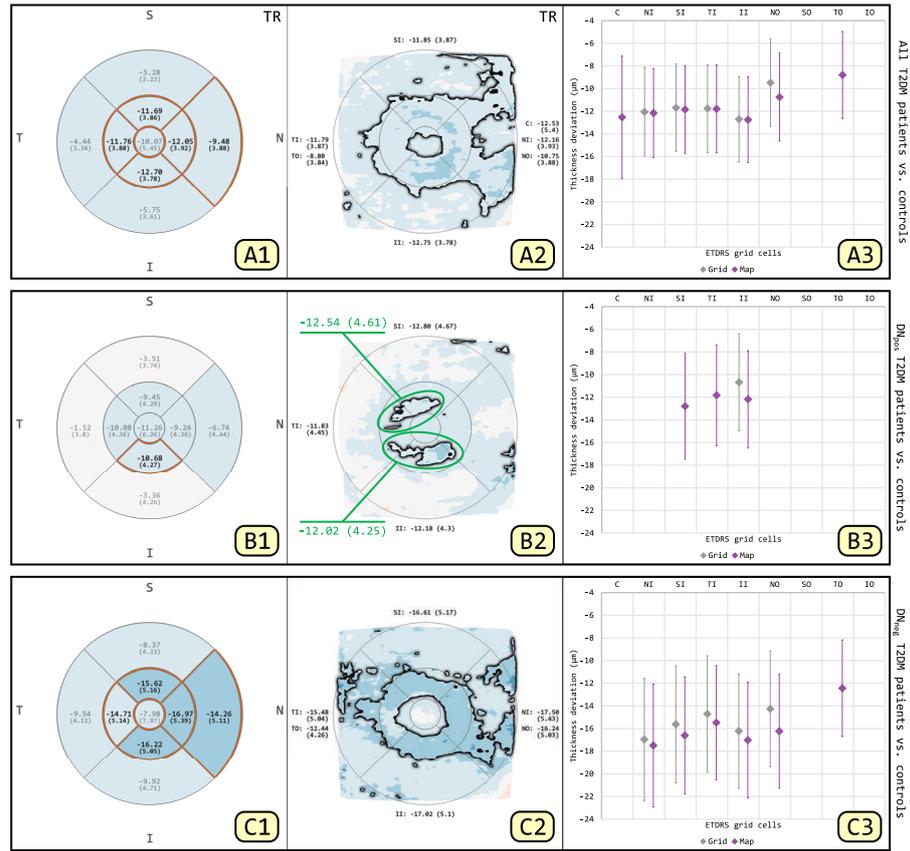


Figure 11.2: Comparisons for TR between all T₂DM patients and controls (A1 to A3), between subgroup of DN_{pos} T₂DM patients and controls (B1 to B3) and between subgroup of DN_{neg} T₂DM patients and controls (C1 to C3). The image columns correspond to results from CA (A1, B1 and C1), results from MA (A2, B2 and C2) and their comparison using MD ± SE plots (A3, B3 and C3). Color encodes thickness deviation (thinning: blue; thickening: red) and statistically significant differences ($p < 0.05$) are highlighted (grid cells: orange borders and black labels; maps: black outline). Two irregular shaped areas of interest were selected in one DevM and measured as individual defects (B2; green circles).

the GCL, the IPL and the Inner Nuclear Layer (INL) were analyzed. Subsequently, patients were stratified according to the presence or absence of DN and the results of either subgroup were compared to the controls. For all comparisons, measurements only of statistically significant areas ($p < 0.05$) per ETDRS grid cell in CA and per fraction of those grid cells in MA were retrieved and reported as MD ± SE. Generally, both analysis methods revealed the presence of a considerable amount and extent of thinning primarily affecting the Inner Macular Ring (IMR) and selectively extending into quadrants of the Outer Macular Ring (OMR) in all investigated retinal layers, except for INL.

11.3.1 Comparison of T2DM patients and patient subgroups to controls for TR

The results of CA and MA for TR are used to illustrate the differences between both methods (Fig. 11.2). Comparison of the T2DM patients and controls for CA and MA revealed thinning of the TR primarily involving IMR, extending towards the nasal OMR. Within the subgroup of DN_{pos} T2DM patients, the thinning of the TR appears to be limited to the inferior IMR with CA, while MA detected an additional irregular shaped area of thinning involving the superior and temporal IMRs. By contrast, in DN_{neg} T2DM patients, CA and MA revealed thinning of the TR primarily involving the IMR, extending only towards the nasal OMR in CA but towards the nasal and temporal OMRs in MA. To highlight the differences between CA and MA, the respective results were compared using MD ± SE plots (Fig. 11.2 (A3, B3, C3)). Overall, CA tended to underestimate thickness changes in ETDRS grid-defined areas, particularly in cells with localized thinning detected by MA. In addition, MA allowed the measuring of selected regions of interest as a single defect, such as the irregular shaped areas of localized thinning located at the inferior, temporal and superior IMRs on DevMs (Fig. 11.2 (B2), green circles).

11.3.2 Comparison of all T2DM patients to controls for RNFL, GCL, and IPL

The results of the comparison between T2DM patients and controls with respect to the thickness of RNFL, GCL and IPL are depicted in Figure 11.3. For all of these layers, both CA and MA demonstrated areas of thinning primarily involving the IMR and selected quadrants of the OMR. In MA, the DevMs revealed a topographical overview of continuous or discrete thickness changes across different macular quadrants. The RNFL showed a greater amount and extent of thinning, extending towards the nasal OMR. The thinning in the GCL was more confined to the IMR, which exhibited a typical pattern of the normal Foveal Center (FC) surrounded by a circular band of thinning. The IPL exposed multiple discrete areas of thinning across the macular quadrants. Moreover, the FC of the IPL showed areas of thinning, which were not present in the RNFL or the GCL.

11.3.3 Comparison between DN_{pos} T2DM patients and controls for RNFL, GCL, and IPL

The results of the comparison between DN_{pos} T2DM patients and controls for RNFL, GCL and IPL are shown in Figure 11.4. For all of these layers, the differences observed were similar to those obtained for the comparison between all T2DM patients and controls (Fig. 11.3). Both CA and MA demonstrated areas of thinning involving the IMR and selected

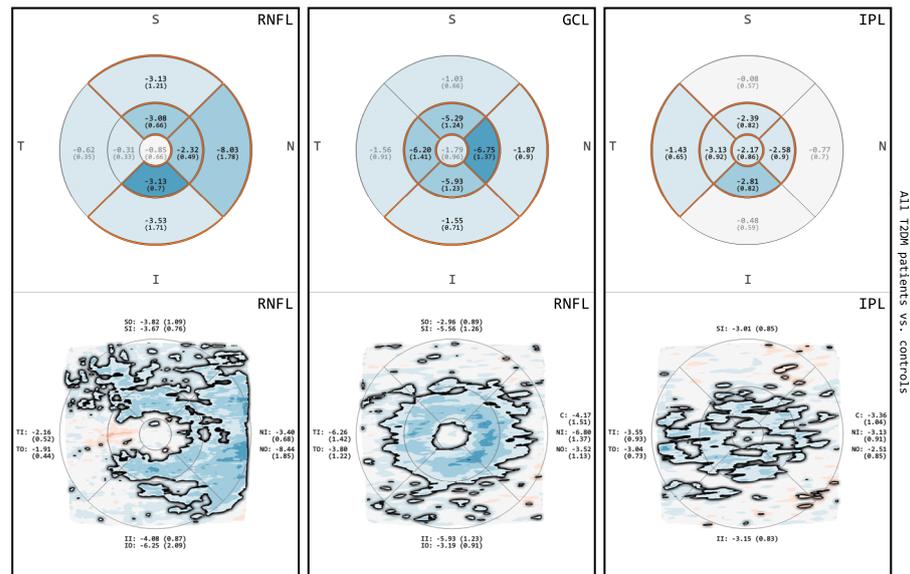


Figure 11.3: Comparisons for RNFL, GCL and IPL between all T2DM patients and controls. The image rows correspond to results from CA (top) and results from MA (bottom). Color encodes thickness deviation (thinning: blue; thickening: red) and statistically significant differences ($p < 0.05$) are highlighted (grid cells: orange borders and black labels; maps: black outline).

quadrants of the OMR. In MA, the RNFL showed a greater amount and extent of thinning extending towards the nasal OMR, the GCL exhibited a normal FC surrounded by a circular band of thinning, and the IPL exposed multiple discrete areas of thinning across the macular quadrants. The FC of the IPL again showed areas of significant thinning, which were not present in the RNFL or the GCL.

11.3.4 Comparison between DN_{neg} T2DM patients and controls for RNFL, GCL, and IPL

The results of the comparison between DN_{neg} T2DM patients and controls for RNFL, GCL and IPL are demonstrated in Figure 11.5. In contrast to the findings described above (Fig. 11.3 and Fig. 11.4), both CA and MA showed areas of thinning involving only parts of the IMR and OMR. In CA, the RNFL exhibited thinning involving the nasal and superior IMRs extending to the nasal OMR, the GCL showed a relatively greater amount and extent of thinning confined to the IMR, and the IPL exposed thinning of the inferior IMR. In MA, two areas of localized thinning in the RNFL (Fig. 11.5; at 10 to 11 o'clock position in the OMR) and in the IPL (Fig. 11.5; at temporal OMR) were noted in addition to those described with CA. In particular, the clustered multiple defects (small scattered areas of localized thinning) in the RNFL (Fig. 11.5; bottom left; green circle)

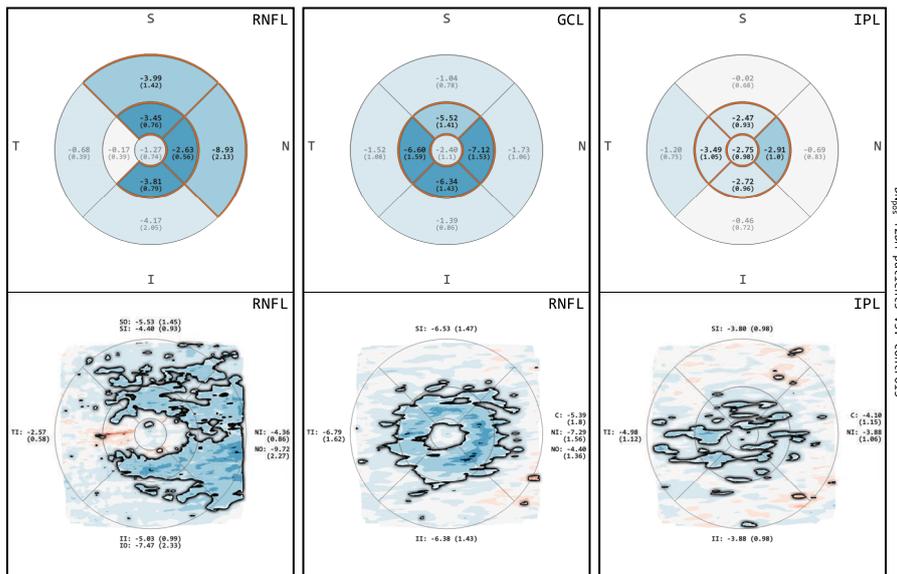


Figure 11.4: Comparisons for RNFL, GCL and IPL between subgroup of DN_{pos} T2DM patients and controls. The image rows correspond to results from CA (top) and results from MA (bottom). Color encodes thickness deviation (thinning: blue; thickening: red) and statistically significant differences ($p < 0.05$) are highlighted (grid cells: orange borders and black labels; maps: black outline).

spanning both the superior and temporal OMRs were not identified in the CA (Fig. 11.5; top left).

Furthermore, a few small and isolated areas of statistically significant thickening were observed only in MA in some DevMs of the RNFL, GCL and IPL, while in CA, only areas of thinning were demonstrated in all comparisons.

11.4 DISCUSSION

In the present study, we demonstrated and compared results from both CA and MA to investigate early neurodegenerative changes in the retinas of patients with diabetes.

11.4.1 Comparison of grid-based and map-based analyses

In general, the areas defined by the ETDRS grid cells are well suited to describing the thickness profile in four quadrants and three anatomically different rings (FC, IMR and OMR) of the central retina. Taking the natural shape of the retina into account, it is often acceptable to summarize and compare thickness measurements of quadrants within the IMR or OMR. However, considering the structural differences between the IMR and OMR, it is generally inappropriate to summarize and compare thickness measurements across macular rings.

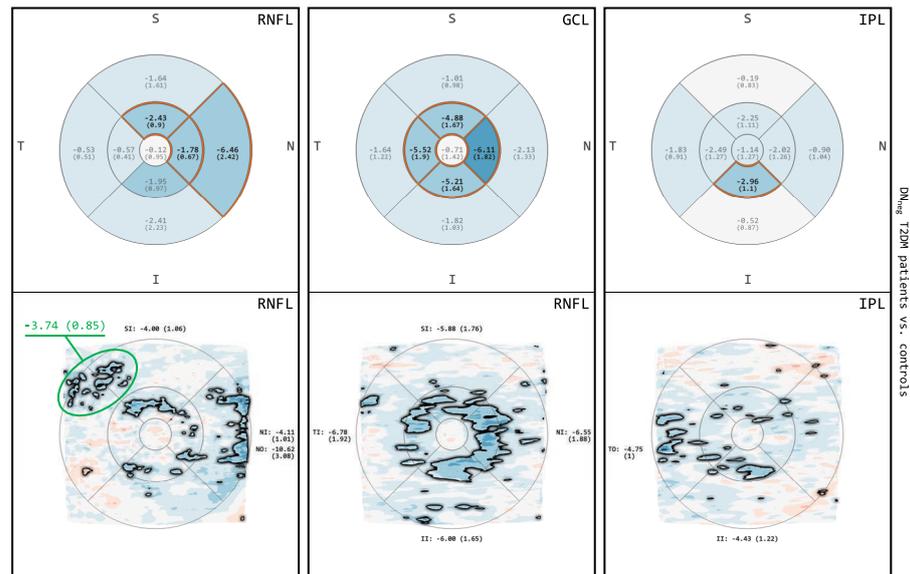


Figure 11.5: Comparisons for RNFL, GCL and IPL between subgroup of DN_{neg} T2DM patients and controls. The image rows correspond to results from CA (top) and results from MA (bottom). Color encodes thickness deviation (thinning: blue; thickening: red) and statistically significant differences ($p < 0.05$) are highlighted (grid cells: orange borders and black labels; maps: black outline). An interesting cluster of small areas of thinning spanning two grid cells was selected in one DevM and measured as a single defect (bottom left; green circle).

Previous retinal OCT studies [97, 176, 220, 221, 224] have adopted this grid-based analysis as a common way of presenting the CA results of group comparisons, either by deriving a mean thickness value for each individual grid cell or by averaging the mean thickness measures of four cells within one ring. This conventional approach grants an overview of thickness changes, but it bears the risk of losing information on subtle and localized changes within these generalized macular areas. The spatially precise DevMs have the ability to preserve such information. However, they have been solely applied to analyzing the thickness data of individuals and to measure single points on the maps in the past. Prior studies [21, 176, 178] include investigations of follow-up scans of an individual or relating the data of a single patient to predefined, often proprietary, normative data. Only recently, we have demonstrated a novel DevM design that, for the first time, supports point-based comparisons of layer thickness between different study groups and helps to analyze subtle and localized changes in the intraretinal layers in children with T1DM and matched controls [30]. Here, we extended our software with an additional feature that allows the accurate analysis of areas with statistically significant differences on DevMs, while also taking natural variations in retinal thickness in anatomically distinct macular areas into account. Compared to CA, which ends up with a

mean thickness value for each **ETDRS** grid cell by averaging all normal and abnormal points, our enhanced **MA** provides more precise measurements of only abnormal points within each grid cell.

To avoid misinterpretation of the automated measurements, we introduced a coverage threshold as the minimum percentage of the grid cell area required to be covered by abnormal points on the **DevMs** (cf. **Sect. 11.2.3**). For example, it might be inappropriate to generalize a measurement of small and scattered areas of significant thinning (coverage threshold $< 10\%$, $p < 0.05$) within a grid cell to the entire area of that cell (e.g., **Fig. 11.5**, nasal and superior quadrants on **DevM** of **IPL**). At the same time, clusters of small areas of significant thinning, even with a coverage threshold of $< 10\%$, may appear relevant (e.g., **Fig. 11.5**, superior-temporal **OMR** at around 10 o'clock on **DevM** of **RNFL**). We argue that these clusters need to be closely followed up in order to assess the thickness changes with regard to clinical findings and disease progression. For this purpose, our **MA** method enables the selection and measurement of any localized changes spanning two or more quadrants within a macular ring, such as the two irregular shaped areas of thinning on the **DevM** of **TR** (**Fig. 11.2** (B2); green circles), measuring $12.54 \mu\text{m} \pm 4.61 \mu\text{m}$ and $-12.02 \mu\text{m} \pm 4.25 \mu\text{m}$, respectively. The measurement of such patterns might help to improve the understanding of a particular disease condition. As an example, in patients with diabetes or glaucoma, it might be interesting to compare the changes between the superior and inferior macular regions in order to investigate correlations related to the anatomical pattern of the **RNFL** bundle.

The **MA** method also enables direct comparisons between **DevM** measurements and corresponding grid measurements in **CA**. While both methods are similarly effective in representing general thickness changes in intraretinal layers, we expect that our new method will provide additional and probably meaningful details of the exact amount and extent of retinal changes. Furthermore, we evidently inferred a tendency of **CA** to produce false negative results due to the averaging of artifacts in the present study. For example, the irregular shaped area of significant thinning crossing the superior-temporal quadrants of the **IMR** was only detected in **MA** (**Fig. 11.2** (B2)) but not as a significant cell in **CA** (**Fig. 11.2** (B1)). Likewise, partial changes of the **FC** could be revealed in **MA** (**Fig. 11.2** (A2)) but not in **CA** (**Fig. 11.2** (A1)). Such foveal and macular involvements are a direct measure of central visual function. It is therefore essential to use an appropriate data analysis method for the accurate detection of changes in order to support early intervention.

11.4.2 Neurodegenerative changes in patients with T2DM

In general, the results from both **CA** and **MA** were comparable, and each method demonstrated significant thinning primarily involving the **IMR** in all except one of the investigated retinal layers, i.e., the **INL**.

The thinning of the IMR also seemed to extend into the nasal OMR in TR and RNFL, which was particularly marked in RNFL. Our results are in line with those reported by others [21, 97, 221, 225–228] in that the thickness of the inner retinal layers is reduced in patients with T2DM and this thinning is thought to be a marker of the beginning of neurodegenerative changes. While the majority of researchers evaluated the inner retinal layers, a number of studies [229, 230] demonstrated morphological alterations in the choroidal and outer retinal layers in diabetes and DR. Therefore, it would be interesting to apply the VA software to study these parts of the retina in children and adults with diabetes. In the present study, we focused on the detection of spatial thickness changes relative to the presence or absence of DN. Our findings demonstrate the flexibility of our VA software regarding the utilization of interactive DevMs for comparison between controls and even subgroups of patients, i.e., those with and without DN. This approach revealed significant amounts of thinning in the RNFL and GCL, but there were no significant thickness changes in between subgroups (DN_{pos} T2DM patients compared to DN_{neg} T2DM patients) in any of the investigated retinal layers. However, the amount and extend of thinning detected in the DevMs of the RNFL and GCL was most prominent in the DN_{pos} T2DM patients. First of all, and in line with a recent study by Srinivasan, this indicates that DN is not restricted to the periphery but also affects the thickness of intraretinal layers [231, 232]. Secondly, this strengthens the utilization of OCT as a measure for the sensitive monitoring of DM patients at risk of developing DN.

11.5 CONCLUSION

In conclusion, the enhanced MA method is highly spatially specific and effective in detecting subtle and early neurodegenerative changes in the intraretinal layers. The importance and advantages of selecting an appropriate data analysis method to accurately measure retinal thickness changes is emphasized. As part of our continued research efforts, we plan to make the developed VA software available on a software sharing website in the future.

Part III

THE OUTRO

This third part summarizes the key insights of this work and presents the overall conclusion. The thesis is completed with future research topics on the visual analytics of optical coherence tomography data of the retina in ophthalmology.

CONCLUSION AND FUTURE WORK

In the preceding two parts of this thesis, we have presented both a coherent summary and the details of our individual approaches. This last chapter contains an overall conclusion and outlook. First, in [Section 12.1](#), we derive and discuss insights about the bigger picture behind this work. Then, in [Section 12.2](#), we propose research topics for future work on the [VA](#) of retinal [OCT](#) data in ophthalmology.

12.1 CONCLUDING REMARKS

In this work, we dealt with the development of novel [VA](#) methods for retinal data acquired via [OCT](#). We started with the need of ophthalmologists to translate the growing number of volumetric images into insights. In particular, they called for improvements in the identification and measurement of small and local defects of the retinal layers in the early stages of diseases. These defects, however, are buried in the complex 3D image data, whose spatial resolution is constantly increasing due to advances in [OCT](#) scanning devices. Looking at the current analysis practice of ophthalmologists, we found that they either inspected individual cross-sections to analyze details in one dataset or applied a strong automated reduction based on the [ETDRS](#) grid to get an overview of multiple datasets. Both methods were associated with a high manual analysis effort and the risk of losing valuable information.

From the [VA](#) perspective, the research to support the analysis of [OCT](#) data of the multi-layered retina had been limited. For this reason, we first determined two fundamental challenges: ensuring the visibility of relevant information and reducing the data complexity. The investigation of these well-known challenges was our main research focus throughout this thesis. First, we addressed them with regard to complex 3D spatial data in general and then, with regard to retinal [OCT](#) data in particular. This helped us to develop new [VA](#) methods to cope with both challenges. In general, we found that they cannot be solved with a single technique. This is why we have developed several innovative concepts for partial problems and consolidated them into new comprehensive solutions. In particular, the combination of two strategies – multiple coordinated 2D and 3D views and the extraction of high-level features – proved to be effective in not only reducing the data complexity but also improving the visibility of diagnostically relevant information. In our tests, we were able to show that our methods based on these strategies actually support ophthalmologists in the early detection of small and local defects of the retinal layers.

Nowadays, OCT examinations of the human eye are an integral part of clinical practice and ophthalmic research. With our research, we have shown that VA is an important technology to help ophthalmologists understand the complex data, make decisions, and generate knowledge. Our results suggest that our solutions can contribute to the handling of today's volumes and diversity of retinal OCT data.

12.2 TOPICS FOR FUTURE WORK

This thesis covered various aspects related to the visibility and complexity of retinal OCT data. In addition to the focus of our work, there are other open questions and opportunities for future VA research.

FURTHER DATA ASPECTS: In this work, we focused on the visual analysis of attributes of retinal layers in their 3D spatial context. Besides these two data components (attributes and spatial context), there are other aspects of retinal OCT data that require consideration in future work. One of these aspects is time, which is important when monitoring the progression of a disease in follow-up examinations of individual patients or patient groups. Relevant questions are how we can compute, visually communicate, and interact with temporal changes of local regions in the attribute space, i.e., the development of abnormal values over time, and of local regions in the spatial context, i.e., the dynamics in the location and shape of the grid cells and map regions. In this respect, an extension of our solutions to track features [67] over local regions of retinal layers may be worth exploring.

Another relevant aspect of retinal OCT data are uncertainties. Besides our initial considerations about missing values and insufficient signal strength [27], there are other influences on the data quality that need to be accounted for. These include quality differences depending on the acquisition parameters [87] and patient compliance, the segmentation algorithms applied, or the different OCT devices used. The general challenge is that these influences must be additionally displayed in order to take them into account in the visual analysis. This in turn drastically increases the complexity of the visualization. Hence, more work is needed on how to effectively encode them together with the other data aspects in 2D and 3D. One possible approach could be to adapt our ideas of a prioritized visualization for attributes, uncertainties, and 3D spatial context [57] to select and visually highlight those aspects that currently require the most attention.

HIGHER-ORDER CORRELATIONS: We presented solutions to analyze the relationship between the layer attributes and the 3D spatial context. When discussing future research directions with ophthalmologists and industry experts, we were asked whether it would be possible to extend the analysis of these relationships to higher-order correla-

tions with respect to the data of other examination methods. Among the data considered were measurements of retinal blood circulation by retinal angiography [202], visual field sensitivities by perimetry [233], or corneal nerve cells by confocal scanning laser microscopy [90]. The inclusion of such measurements could help to link structural and functional changes of the eye with influencing factors and thus to detect eye diseases more accurately. In future work, it would therefore be interesting to investigate how the correlations between retinal OCT data and these measurements in the attribute space and the spatial context can be computed and visually communicated. To present them in the spatial context, one way could be to extend our feature extraction and visualization to make local regions with relevant correlation visible in the grids and maps. To present them in the attribute space, our VA approach for the identification of ophthalmic markers in multivariate clinical data [234, 235] might be a useful starting point.

EXTENDED ANALYSIS SUPPORT: As discussed in Section 4.1, our VA tool provides support for numerous steps required in the evaluation of single-subject and multi-subject data. Taking other data aspects and the analysis of higher-order correlations into account will, however, further complicate the entire analysis process. From the perspective of VA research, it makes sense to develop new solutions for some parts of this process, but to use existing tools for others, e.g., the administration of datasets or the segmentation of retinal layers. This leads to the problem that the ophthalmologists still have to use different software packages to fulfill their analysis tasks. A future research question is therefore how we can integrate our VA methods and other existing tools to offer further analysis support. Instead of developing the tools from scratch or trying to bring everything together in a single monolithic framework, one idea is to loosely couple the separate tools and provide access to them through a unified interface. In this regard, a lightweight coordination for the VA of retinal OCT data could be established by considering our initial work on the visualization of the analysis steps together with the visual output of the tools [236] and linking of the tools based on our layered toolchaining approach [206, 237].

OTHER ANALYSIS SCENARIOS: We demonstrated our solutions using data from individual patients and patient groups in relation to control data in the context of cross-sectional studies. The evaluation of such study data plays an important role in ophthalmic research. Nonetheless, it makes sense to consider additional analysis scenarios in the future. One particularly interesting scenario is the exploratory evaluation of retinal OCT data in large-scale prospective and population-based studies, such as the Gutenberg Health Study with 15 000 participants [238]. The scope of these studies presents several new challenges. For example, here the analysis was focused on the early detection of retinal layer de-

fects by comparing predefined patient and control groups. In contrast, in population-based studies a challenge would be to find out which retinal data of individuals or subgroups differ from the rest of the cohort, how they differ, and what makes them special. The development of respective VA approaches might contribute to the documentation of the study participants' health development and the progression of occurred eye diseases. While our solutions could be a starting point for VA research in this direction, more work is needed to translate the data volumes of such complex multidisciplinary studies into insights.

The ideas and questions described here contain some pointers and possibilities for future work. The range of topics suggests that there is still work to be done in connection with the VA of retinal OCT data. With regard to the results presented in this thesis, we conclude that it is worthwhile to address the associated challenges in future VA research.

BIBLIOGRAPHY

- [1] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, and J. G. Fujimoto. "Optical Coherence Tomography." In: *Science* 254.5035 (1991), pp. 1178–1181. doi: [10.1126/science.1957169](https://doi.org/10.1126/science.1957169).
- [2] M. R. Hee, J. A. Izatt, E. A. Swanson, D. Huang, J. S. Schuman, C. P. Lin, C. A. Puliafito, and J. G. Fujimoto. "Optical Coherence Tomography of the Human Retina." In: *JAMA Ophthalmology* 113.3 (1995), pp. 325–332. doi: [10.1001/archophth.1995.01100030081025](https://doi.org/10.1001/archophth.1995.01100030081025).
- [3] C. A. Puliafito, M. R. Hee, C. P. Lin, E. Reichel, J. S. Schuman, J. S. Duker, J. A. Izatt, E. A. Swanson, and J. G. Fujimoto. "Imaging of Macular Diseases with Optical Coherence Tomography." In: *Ophthalmology* 102.2 (1995), pp. 217–229. doi: [10.1016/S0161-6420\(95\)31032-9](https://doi.org/10.1016/S0161-6420(95)31032-9).
- [4] P. A. Keane, P. J. Patel, S. Liakopoulos, F. M. Heussen, S. R. Sadda, and A. Tufail. "Evaluation of Age-related Macular Degeneration With Optical Coherence Tomography." In: *Survey of Ophthalmology* 57.5 (2012), pp. 389–414. doi: [10.1016/j.survophthal.2012.01.006](https://doi.org/10.1016/j.survophthal.2012.01.006).
- [5] G. Virgili, F. Menchini, G. Casazza, R. Hogg, R. R. Das, X. Wang, and M. Michelessi. "Optical coherence tomography (OCT) for detection of macular oedema in patients with diabetic retinopathy." In: *The Cochrane database of systematic reviews* 1 (2015), p. CD008081. doi: [10.1002/14651858.CD008081.pub3](https://doi.org/10.1002/14651858.CD008081.pub3).
- [6] D. S. Grewal and A. P. Tanna. "Diagnosis of glaucoma and detection of glaucoma progression using spectral domain optical coherence tomography." In: *Current Opinion in Ophthalmology* 24.2 (2013), pp. 150–161. doi: [10.1097/ICU.0b013e32835d9e27](https://doi.org/10.1097/ICU.0b013e32835d9e27).
- [7] J. Dörr, K. D. Wernecke, M. Bock, G. Gaede, J. T. Wuerfel, C. F. Pfueller, J. Bellmann-Strobl, A. Freing, A. U. Brandt, and P. Friedemann. "Association of Retinal and Macular Damage with Brain Atrophy in Multiple Sclerosis." In: *PLOS ONE* 6.4 (2011), pp. 1–6. doi: [10.1371/journal.pone.0018132](https://doi.org/10.1371/journal.pone.0018132).
- [8] J. Fujimoto and E. Swanson. "The Development, Commercialization, and Impact of Optical Coherence Tomography." In: *Investigative Ophthalmology & Visual Science* 57.9 (2016), OCT1–OCT13. doi: [10.1167/iovs.16-19963](https://doi.org/10.1167/iovs.16-19963).

- [9] Centers for Medicare and Medicaid Services. *Medicare Provider Utilization and Payment Data: Physician and Other Supplier*. Retrieved from Physician and Other Supplier Data CY 2017 at cms.gov. July 2019. URL: <https://www.cms.gov/research-statistics-data-and-systems/statistics-trends-and-reports/medicare-provider-charge-data/physician-and-other-supplier.html>.
- [10] US National Library of Medicine and National Institutes of Health. *PubMed database: Citations corresponding to optical coherence tomography and ophthalmology*. Retrieved from PubMed search at pubmed.gov. Jan. 2020. URL: <https://www.ncbi.nlm.nih.gov/pubmed/?term=optical+coherence+tomography+ophthalmology>.
- [11] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE, 2005. ISBN: 9780769523231.
- [12] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. "Visual Analytics: Definition, Process, and Challenges." In: *Information Visualization: Human-Centered Issues and Perspectives*. Ed. by A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North. Springer Berlin Heidelberg, 2008, pp. 154–175. DOI: [10.1007/978-3-540-70956-5_7](https://doi.org/10.1007/978-3-540-70956-5_7).
- [13] W. Drexler, U. Morgner, R. K. Ghanta, F. X. Kärtner, J. S. Schuman, and J. G. Fujimoto. "Ultrahigh-resolution ophthalmic optical coherence tomography." In: *Nature Medicine* 7.4 (2001), pp. 502–507. DOI: [10.1038/86589](https://doi.org/10.1038/86589).
- [14] M. A. Choma, M. V. Sarunic, C. Yang, and J. A. Izatt. "Sensitivity advantage of swept source and Fourier domain optical coherence tomography." In: *Optics Express* 11.18 (2003), pp. 2183–2189. DOI: [10.1364/OE.11.002183](https://doi.org/10.1364/OE.11.002183).
- [15] M. Adhi and J. S. Duker. "Optical coherence tomography – current and future applications." In: *Current opinion in ophthalmology* 24.3 (2013), pp. 213–221. DOI: [10.1097/ICU.0b013e32835f8bf8](https://doi.org/10.1097/ICU.0b013e32835f8bf8).
- [16] N. Yoshimura and M. Hangai. *OCT Atlas*. Springer, 2014. ISBN: 978-3-642-38624-4.
- [17] R. A. Costa, M. Skaf, L. A. Melo, D. Calucci, J. A. Cardillo, J. C. Castro, D. Huang, and M. Wojtkowski. "Retinal assessment using optical coherence tomography." In: *Progress in Retinal and Eye Research* 25.3 (2006), pp. 325–353. DOI: [10.1016/j.preteyeres.2006.03.001](https://doi.org/10.1016/j.preteyeres.2006.03.001).

- [18] J. D. Stein, B. W. Hanrahan, G. M. Comer, and F. A. Sloan. "Diffusion of Technologies for the Care of Older Adults With Exudative Age-Related Macular Degeneration." In: *American Journal of Ophthalmology* 155.4 (2013), 688–696.e2. doi: [10.1016/j.ajo.2012.10.003](https://doi.org/10.1016/j.ajo.2012.10.003).
- [19] J. C. Buchan, W. Amoaku, B. Barnes, A. Cassels-Brown, B. Y. Chang, J. Harcourt, D. Shickle, A. F. Spencer, S. A. Vernon, and C. MacEwen. "How to defuse a demographic time bomb: the way forward?" In: *Eye* 31 (2017), pp. 1519–1522. doi: [10.1038/eye.2017.114](https://doi.org/10.1038/eye.2017.114).
- [20] T. Ilginis, J. Clarke, and P. J. Patel. "Ophthalmic imaging." In: *British medical bulletin* 111.1 (2014), pp. 77–88. doi: [10.1093/bmb/ldu022](https://doi.org/10.1093/bmb/ldu022).
- [21] J. Chhablani, A. Sharma, A. Goud, H. K. Peguda, H. L. Rao, V. U. Begum, and G. Barteselli. "Neurodegeneration in Type 2 Diabetes: Evidence From Spectral-Domain Optical Coherence Tomography." In: *Investigative Ophthalmology & Visual Science* 56.11 (2015), pp. 6333–6338. doi: [10.1167/iovs.15-17334](https://doi.org/10.1167/iovs.15-17334).
- [22] K. Tekin, M. Inanc, E. Kurnaz, E. Bayramoglu, E. Aydemir, M. Koc, H. Kiziltoprak, and Z. Aycan. "Quantitative evaluation of early retinal changes in children with type 1 diabetes mellitus without retinopathy." In: *Clinical and Experimental Optometry* 101.5 (2018), pp. 680–685. doi: [10.1111/cxo.12667](https://doi.org/10.1111/cxo.12667).
- [23] B. Preim and C. Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. Morgan Kaufmann Publishers Inc., 2014. ISBN: 978-0-12-415873-3.
- [24] N. Andrienko and G. Andrienko. *Exploratory Analysis of Spatial and Temporal Data*. Springer, 2006. doi: [10.1007/3-540-31190-4](https://doi.org/10.1007/3-540-31190-4).
- [25] **M. Röhlig**, M. Luboschik, and H. Schumann. "Visibility Widgets for Unveiling Occluded Data in 3D Terrain Visualization." In: *Journal of Visual Languages & Computing* 42.Supplement C (2017), pp. 86–98. doi: [10.1016/j.jvlc.2017.08.008](https://doi.org/10.1016/j.jvlc.2017.08.008).
- [26] M. Luboschik, **M. Röhlig**, A. T. Bittig, N. Andrienko, H. Schumann, and C. Tominski. "Feature-Driven Visual Analytics of Chaotic Parameter-Dependent Movement." In: *Computer Graphics Forum* 34.3 (2015). Best paper award, pp. 421–430. doi: [10.1111/cgf.12654](https://doi.org/10.1111/cgf.12654).
- [27] **M. Röhlig**, C. Schmidt, R. K. Prakasam, P. Rosenthal, H. Schumann, and O. Stachs. "Visual Analysis of Retinal Changes with Optical Coherence Tomography." In: *The Visual Computer* 34.9 (2018), pp. 1209–1224. doi: [10.1007/s00371-018-1486-x](https://doi.org/10.1007/s00371-018-1486-x).

- [28] **M. Röhlig**, R. K. Prakasam, J. Stüwe, C. Schmidt, O. Stachs, and H. Schumann. "Enhanced Grid-Based Visual Analysis of Retinal Layer Thickness with Optical Coherence Tomography." In: *Information* 10.9 (2019), p. 266. DOI: [10.3390/info10090266](https://doi.org/10.3390/info10090266).
- [29] R. Arias-Hernandez, L. T. Kaastra, T. M. Green, and B. D. Fisher. "Pair Analytics: Capturing Reasoning Processes in Collaborative Visual Analytics." In: *Proceedings of the Hawaii International Conference on System Sciences*. IEEE Computer Society, 2011, pp. 1–10. DOI: [10.1109/HICSS.2011.339](https://doi.org/10.1109/HICSS.2011.339).
- [30] R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Götze, A. Jünnemann, H. Schumann, and O. Stachs. "Deviation maps for understanding thickness changes of inner retinal layers in children with type 1 diabetes mellitus." In: *Current Eye Research* 44.7 (2019). DOI: [10.1080/02713683.2019.1591463](https://doi.org/10.1080/02713683.2019.1591463).
- [31] R. K. Prakasam, A. Matuszewska-Iwanicka, D.-C. Fischer, H. Schumann, D. Tschöpe, B. Stratmann, H.-J. Hettlich, R. F. Guthoff, O. Stachs, and **M. Röhlig**. "Thickness of intraretinal layers in patients with type 2 diabetes mellitus depending on a concomitant diabetic neuropathy: results of a cross-sectional study using deviation maps for OCT data analysis." In: *Biomedicines* 8.7 (2020). DOI: [10.3390/biomedicines8070190](https://doi.org/10.3390/biomedicines8070190).
- [32] S. Dübel, **M. Röhlig**, H. Schumann, and M. Trapp. "2D and 3D Presentation of Spatial Data: A Systematic Review." In: *Proceedings of the IEEE VIS International Workshop on 3DVis: Does 3D really make sense for Data Visualization?* 2014. DOI: [10.1109/3DVis.2014.7160094](https://doi.org/10.1109/3DVis.2014.7160094).
- [33] P. R. Keller and M. M. Keller. *Visual Cues: Practical Data Visualization*. IEEE Computer Society Press, 1994. ISBN: 0818631023.
- [34] D. A. Keim. "Information visualization and visual data mining." In: *IEEE Transactions on Visualization and Computer Graphics* 8.1 (2002), pp. 1–8. DOI: [10.1109/2945.981847](https://doi.org/10.1109/2945.981847).
- [35] M. D. Plumlee and C. Ware. "Zooming versus Multiple Window Interfaces: Cognitive Costs of Visual Comparisons." In: *ACM Transactions on Computer-Human Interaction* 13.2 (2006), pp. 179–209. DOI: [10.1145/1165734.1165736](https://doi.org/10.1145/1165734.1165736).
- [36] J. Wood, P. Isenberg, T. Isenberg, J. Dykes, N. Boukhelifa, and A. Slingsby. "Sketchy Rendering for Information Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2749–2758. DOI: [10.1109/TVCG.2012.262](https://doi.org/10.1109/TVCG.2012.262).
- [37] B. B. Bederson and A. Boltman. "Does animation help users build mental maps of spatial information?" In: *Proceedings of the IEEE Symposium on Information Visualization*. 1999, pp. 28–35. DOI: [10.1109/INFVIS.1999.801854](https://doi.org/10.1109/INFVIS.1999.801854).

- [38] J. Vion-Dury and M. Santana. "Virtual Images: Interactive Visualization of Distributed Object-Oriented Systems." In: *Proceedings of the Ninth Annual Conference on Object-Oriented Programming Systems, Languages, and Applications*. Ed. by J. McKenna, J. E. B. Moss, and R. L. Wexelblat. ACM, 1994, pp. 65–84. DOI: [10.1145/191080.191095](https://doi.org/10.1145/191080.191095).
- [39] A. Cockburn and B. McKenzie. "Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery, 2002, pp. 203–210. DOI: [10.1145/503376.503413](https://doi.org/10.1145/503376.503413).
- [40] A. Cockburn. "Revisiting 2D vs 3D Implications on Spatial Memory." In: *Proceedings of the Fifth Conference on Australasian User Interface (AUIC)*. Australian Computer Society, 2004, pp. 25–31.
- [41] M. Tory, A. E. Kirkpatrick, M. S. Atkins, and T. Moller. "Visualization task performance with 2D, 3D, and combination displays." In: *IEEE Transactions on Visualization and Computer Graphics* 12.1 (2006), pp. 2–13. DOI: [10.1109/TVCG.2006.17](https://doi.org/10.1109/TVCG.2006.17).
- [42] G. Fuchs and H. Schumann. "Visualizing abstract data on maps." In: *Proceedings of the International Conference on Information Visualisation*. 2004, pp. 139–144. DOI: [10.1109/IV.2004.1320136](https://doi.org/10.1109/IV.2004.1320136).
- [43] M. Kreuseler. "Visualization of geographically related multi-dimensional data in virtual 3D scenes." In: *Computers & Geosciences* 26.1 (2000), pp. 101–108. DOI: [10.1016/S0098-3004\(99\)00036-9](https://doi.org/10.1016/S0098-3004(99)00036-9).
- [44] C. Tominski, H. Schumann, G. Andrienko, and N. Andrienko. "Stacking-Based Visualization of Trajectory Attribute Data." In: *IEEE Transactions on Visualization and Computer Graphics* 18.12 (2012), pp. 2565–2574. DOI: [10.1109/TVCG.2012.265](https://doi.org/10.1109/TVCG.2012.265).
- [45] C. Ware and M. Plumlee. "Designing a better weather display." In: *Visualization and Data Analysis*. Ed. by P. C. Wong, D. L. Kao, M. C. Hao, C. Chen, R. Kosara, M. A. Livingston, J. Park, and I. Roberts. Vol. 8294. SPIE, 2012, pp. 75–87. DOI: [10.1117/12.906213](https://doi.org/10.1117/12.906213).
- [46] R. M. Pickett and G. G. Grinstein. "Iconographic Displays For Visualizing Multidimensional Data." In: *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*. Vol. 1. 1988, pp. 514–519. DOI: [10.1109/ICSMC.1988.754351](https://doi.org/10.1109/ICSMC.1988.754351).
- [47] G. Kindlmann and C. Westin. "Diffusion Tensor Visualization with Glyph Packing." In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (2006), pp. 1329–1336. DOI: [10.1109/TVCG.2006.134](https://doi.org/10.1109/TVCG.2006.134).

- [48] G. Ellis and A. Dix. "A Taxonomy of Clutter Reduction for Information Visualisation." In: *IEEE Transactions on Visualization & Computer Graphics* 13.06 (2007), pp. 1216–1223. DOI: [10.1109/TVCG.2007.70535](https://doi.org/10.1109/TVCG.2007.70535).
- [49] S. Few. *Solutions to the Problem of Over-Plotting in Graphs*. Perceptual Edge - Visual Business Intelligence Newsletter. Sept. 2008.
- [50] N. Andrienko and G. Andrienko. "Spatial Generalization and Aggregation of Massive Movement Data." In: *IEEE Transactions on Visualization and Computer Graphics* 17.2 (2011), pp. 205–219. DOI: [10.1109/TVCG.2010.44](https://doi.org/10.1109/TVCG.2010.44).
- [51] T. Hägerstrand. "What about people in regional science?" In: *Papers in Regional Science* 24.1 (1970), pp. 7–24. DOI: [10.1111/j.1435-5597.1970.tb01464.x](https://doi.org/10.1111/j.1435-5597.1970.tb01464.x).
- [52] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale. "A Review of Temporal Data Visualizations Based on Space-Time Cube Operations." In: *EuroVis - STARs*. Ed. by R. Borgo, R. Maciejewski, and I. Viola. The Eurographics Association, 2014. DOI: [10.2312/eurovisstar.20141171](https://doi.org/10.2312/eurovisstar.20141171).
- [53] M. Jobst and J. Döllner. "Better Perception of 3D-Spatial Relations by Viewport Variations." In: *Visual Information Systems. Web-Based Visual Information Search and Management*. Ed. by M. Sebillio, G. Vitiello, and G. Schaefer. Springer Berlin Heidelberg, 2008, pp. 7–18. DOI: [10.1007/978-3-540-85891-1_4](https://doi.org/10.1007/978-3-540-85891-1_4).
- [54] N. Elmqvist and P. Tsigas. "A Taxonomy of 3D Occlusion Management for Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 14.5 (2008), pp. 1095–1109. DOI: [10.1109/TVCG.2008.59](https://doi.org/10.1109/TVCG.2008.59).
- [55] J. Engel, A. Semmo, M. Trapp, and J. Döllner. "Evaluating the Perceptual Impact of Rendering Techniques on Thematic Color Mappings in 3D Virtual Environments." In: *Vision, Modeling & Visualization*. Ed. by M. Bronstein, J. Favre, and K. Hormann. The Eurographics Association, 2013. DOI: [10.2312/PE.VMV.VMV13.025-032](https://doi.org/10.2312/PE.VMV.VMV13.025-032).
- [56] D. A. Szafir, A. Sarikaya, and M. Gleicher. "Lightness Constancy in Surface Visualization." In: *IEEE Transactions on Visualization and Computer Graphics* 22.9 (2016), pp. 2107–2121. DOI: [10.1109/TVCG.2015.2500240](https://doi.org/10.1109/TVCG.2015.2500240).
- [57] S. Dübel, M. Röhlig, C. Tominksi, and H. Schumann. "Visualizing 3D Terrain, Geo-spatial Data, and Uncertainty." In: *InformatICS* 4.1 (2017). DOI: [10.3390/informatICS4010006](https://doi.org/10.3390/informatICS4010006).

- [58] **M. Röhlig** and H. Schumann. “Visibility Widgets: Managing Occlusion of Quantitative Data in 3D Terrain Visualization.” In: *Proceedings of the 9th International Symposium on Visual Information Communication and Interaction (VINCI)*. Best paper award. ACM, 2016, pp. 51–58. DOI: [10.1145/2968220.2968230](https://doi.org/10.1145/2968220.2968230).
- [59] R. Stoakley, M. J. Conway, and R. Pausch. “Virtual Reality on a WIM: Interactive Worlds in Miniature.” In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 1995, pp. 265–272. DOI: [10.1145/223904.223938](https://doi.org/10.1145/223904.223938).
- [60] M. Plumlee and C. Ware. “Integrating multiple 3D views through frame-of-reference interaction.” In: *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*. 2003, pp. 34–43. DOI: [10.1109/CMV.2003.1215001](https://doi.org/10.1109/CMV.2003.1215001).
- [61] S. K. Feiner and D. D. Seligmann. “Cutaways and ghosting: Satisfying visibility constraints in dynamic 3D illustrations.” In: *The Visual Computer* 8.5 (1992), pp. 292–302. DOI: [10.1007/BF01897116](https://doi.org/10.1007/BF01897116).
- [62] N. Elmqvist, U. Assarsson, and P. Tsigas. “Employing Dynamic Transparency for 3D Occlusion Management: Design Issues and Evaluation.” In: *Proceedings of the IFIP International Conference on Human-Computer Interaction (INTERACT)*. Springer, 2007, pp. 532–545. DOI: [10.1007/978-3-540-74796-3_54](https://doi.org/10.1007/978-3-540-74796-3_54).
- [63] D. N. Wood, A. Finkelstein, J. F. Hughes, C. E. Thayer, and D. H. Salesin. “Multiperspective Panoramas for Cel Animation.” In: *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '97. ACM Press/Addison-Wesley Publishing Co., 1997, pp. 243–250. DOI: [10.1145/258734.258859](https://doi.org/10.1145/258734.258859).
- [64] S. Möser, P. Degener, R. Wahl, and R. Klein. “Context Aware Terrain Visualization for Wayfinding and Navigation.” In: *Computer Graphics Forum* 27.7 (2008), pp. 1853–1860. DOI: [10.1111/j.1467-8659.2008.01332.x](https://doi.org/10.1111/j.1467-8659.2008.01332.x).
- [65] W. Straßer. “Schnelle Kurven- und Flächendarstellung auf grafischen Sichtgeräten.” PhD thesis. Technische Universität Berlin, 1974.
- [66] E. E. Catmull. “A Subdivision Algorithm for Computer Display of Curved Surfaces.” PhD thesis. The University of Utah, 1974.
- [67] F. H. Post, B. Vrolijk, H. Hauser, R. S. Laramee, and H. Doleisch. “The State of the Art in Flow Visualisation: Feature Extraction and Tracking.” In: *Computer Graphics Forum* 22.4 (2003), pp. 775–792. DOI: [10.1111/j.1467-8659.2003.00723.x](https://doi.org/10.1111/j.1467-8659.2003.00723.x).

- [68] G. Andrienko, N. Andrienko, P. Bak, D. Keim, and S. Wrobel. *Visual Analytics of Movement*. Springer, 2013. doi: [10.1007/978-3-642-37583-5](https://doi.org/10.1007/978-3-642-37583-5).
- [69] T. von Landesberger, S. Bremm, T. Schreck, and D. W. Fellner. “Feature-Based Automatic Identification of Interesting Data Segments in Group Movement Data.” In: *Information Visualization* 13.3 (2014), pp. 190–212. doi: [10.1177/1473871613477851](https://doi.org/10.1177/1473871613477851).
- [70] J. Kehrer and H. Hauser. “Visualization and Visual Analysis of Multifaceted Scientific Data: A Survey.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.3 (2013), pp. 495–513. doi: [10.1109/TVCG.2012.110](https://doi.org/10.1109/TVCG.2012.110).
- [71] F. Reinders, F. H. Post, and H. J. W. Spoelder. “Visualization of Time-Dependent Data with Feature Tracking and Event Detection.” In: *The Visual Computer* 17.1 (2001), pp. 55–71. doi: [10.1007/PL00013399](https://doi.org/10.1007/PL00013399).
- [72] R. Zeimer, M. Shahidi, M. Mori, S. Zou, and S. Asrani. “A new method for rapid mapping of the retinal thickness at the posterior pole.” In: *Investigative Ophthalmology & Visual Science* 37.10 (1996), pp. 1994–2001.
- [73] D. E. Lederer, J. S. Schuman, E. Hertzmark, J. Heltzer, L. J. Velazques, J. G. Fujimoto, and C. Mattox. “Analysis of macular volume in normal and glaucomatous eyes using optical coherence tomography.” In: *American Journal of Ophthalmology* 135.6 (2003), pp. 838–843. doi: [10.1016/S0002-9394\(02\)02277-8](https://doi.org/10.1016/S0002-9394(02)02277-8).
- [74] M. Sedlmair, C. Heinzl, S. Bruckner, H. Piringer, and T. Möller. “Visual Parameter Space Analysis: A Conceptual Framework.” In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2161–2170. doi: [10.1109/TVCG.2014.2346321](https://doi.org/10.1109/TVCG.2014.2346321).
- [75] N. Andrienko, G. Andrienko, L. Barrett, M. Dostie, and P. Henzi. “Space Transformation for Understanding Group Movement.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2169–2178. doi: [10.1109/TVCG.2013.193](https://doi.org/10.1109/TVCG.2013.193).
- [76] K. Reda, C. Tantipathananandh, T. Berger-Wolf, J. Leigh, and A. Johnson. “SocioScape – a Tool for Interactive Exploration of Spatio-Temporal Group Dynamics in Social Networks.” In: *Poster at the IEEE Conference on Information Visualization (InfoVis)*. 2009.
- [77] D. V. Nicolau, K. Burrage, R. G. Parton, and J. F. Hancock. “Identifying Optimal Lipid Raft Characteristics Required To Promote Nanoscale Protein-Protein Interactions on the Plasma Membrane.” In: *Molecular and Cellular Biology* 26.1 (2006), pp. 313–323. doi: [10.1128/MCB.26.1.313-323.2006](https://doi.org/10.1128/MCB.26.1.313-323.2006).

- [78] **M. Röhlig**, M. Luboschik, M. Bögl, F. Krüger, B. Alsallakh, T. Kirste, S. Miksch, and H. Schumann. “Supporting Activity Recognition by Visual Analytics.” In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 41–48. DOI: [10.1109/VAST.2015.7347629](https://doi.org/10.1109/VAST.2015.7347629).
- [79] **M. Röhlig**, P. Rosenthal, C. Schmidt, H. Schumann, and O. Stachs. “Visual Analysis of Optical Coherence Tomography Data in Ophthalmology.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by M. Sedlmair and C. Tominski. Best paper award. The Eurographics Association, 2017. DOI: [10.2312/eurova.20171117](https://doi.org/10.2312/eurova.20171117).
- [80] M. K. Garvin, M. D. Abramoff, X. Wu, S. R. Russell, T. L. Burns, and M. Sonka. “Automated 3-D Intraretinal Layer Segmentation of Macular Spectral-Domain Optical Coherence Tomography Images.” In: *IEEE Transactions on Medical Imaging* 28.9 (2009), pp. 1436–1447. DOI: [10.1109/TMI.2009.2016958](https://doi.org/10.1109/TMI.2009.2016958).
- [81] C. Glittenberg, I. Krebs, C. Falkner-Radler, F. Zeiler, P. Haas, S. Hagen, and S. Binder. “Advantages of Using a Ray-traced, Three-Dimensional Rendering System for Spectral Domain Cirrus HD-OCT to Visualize Subtle Structures of the Vitreoretinal Interface.” In: *Ophthalmic Surgery Lasers and Imaging* 40.2 (2009), pp. 127–134. DOI: [10.3928/15428877-20090301-08](https://doi.org/10.3928/15428877-20090301-08).
- [82] M. G. Garrido, S. C. Beck, R. Mühlfriedel, S. Julien, U. Schraermeyer, and M. W. Seeliger. “Towards a Quantitative OCT Image Analysis.” In: *PLOS ONE* 9.6 (2014), pp. 1–10. DOI: [10.1371/journal.pone.0100080](https://doi.org/10.1371/journal.pone.0100080).
- [83] J. Weiss, U. Eck, M. A. Nasser, M. Maier, A. Eslami, and N. Navab. “Layer-Aware iOCT Volume Rendering for Retinal Surgery.” In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. Ed. by K. Lawonn and R. G. Raidou. The Eurographics Association, 2019. DOI: [10.2312/vcbm.20191239](https://doi.org/10.2312/vcbm.20191239).
- [84] R. Kafieh, H. Rabbani, and S. Kermani. “A review of algorithms for segmentation of optical coherence tomography from retina.” In: *Journal of medical signals and sensors* 3.1 (2013), pp. 45–60.
- [85] A. Ehnes, Y. Wenner, C. Friedburg, M. N. Preising, W. Bowl, W. Sekundo, E. M. zu Bexten, K. Stieger, and B. Lorenz. “Optical Coherence Tomography (OCT) Device Independent Intraretinal Layer Segmentation.” In: *Translational Vision Science & Technology* 3.1 (2014). DOI: [10.1167/tvst.3.1.1](https://doi.org/10.1167/tvst.3.1.1).
- [86] J. Stüwe, **M. Röhlig**, H. Schumann, R. K. Prakasam, and O. Stachs. “Visual Analysis of Abnormal Thickness of Intraretinal Layers.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2018.

- [87] **M. Röhlig**, M. Luboschik, R. K. Prakasam, O. Stachs, and H. Schumann. “Visually Analyzing Parameter Influence on Optical Coherence Tomography Data in Ophthalmology.” In: *EuroVis 2017 - Posters*. Ed. by A. P. Puig and T. Isenberg. The Eurographics Association, 2017. doi: [10.2312/eurp.20171176](https://doi.org/10.2312/eurp.20171176).
- [88] **M. Röhlig**, J. Stüwe, C. Schmidt, R. K. Prakasam, O. Stachs, and H. Schumann. “Grid-Based Exploration of OCT Thickness Data of Intraretinal Layers.” In: *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*. Ed. by A. Kerren, C. Hurter, and J. Braz. Vol. 3. International Conference on Information Visualization Theory and Applications (IVAPP). Best student paper award. SciTePress, 2019, pp. 129–140. doi: [10.5220/0007580001290140](https://doi.org/10.5220/0007580001290140).
- [89] Early Treatment Diabetic Retinopathy Study Research Group. “Grading Diabetic Retinopathy from Stereoscopic Color Fundus Photographs – An Extension of the Modified Airlie House Classification: ETDRS Report Number 10.” In: *Ophthalmology* 98.5 (1991), pp. 786–806. doi: [10.1016/S0161-6420\(13\)38012-9](https://doi.org/10.1016/S0161-6420(13)38012-9).
- [90] A. Götze, S. von Keyserlingk, S. Peschel, U. Jacoby, C. Schreiber, B. Köhler, S. Allgeier, K. Winter, **M. Röhlig**, A. Jünemann, R. Guthoff, O. Stachs, and D.-C. Fischer. “The corneal subbasal nerve plexus and thickness of the retinal layers in pediatric type 1 diabetes and matched controls.” In: *Scientific Reports* 8.1 (2018). doi: [10.1038/s41598-017-18284-z](https://doi.org/10.1038/s41598-017-18284-z).
- [91] Q. Chen, S. Huang, Q. Ma, H. Lin, M. Pan, X. Liu, F. Lu, and M. Shen. “Ultra-high resolution profiles of macular intra-retinal layer thicknesses and associations with visual field defects in primary open angle glaucoma.” In: *Scientific Reports* 7 (2017), p. 41100. doi: [10.1038/srep41100](https://doi.org/10.1038/srep41100).
- [92] S. Asrani, J. Rosdahl, and R. Allingham. “Novel software strategy for glaucoma diagnosis: Asymmetry analysis of retinal thickness.” In: *Archives of Ophthalmology* 129.9 (2011), pp. 1205–1211. doi: [10.1001/archophthamol.2011.242](https://doi.org/10.1001/archophthamol.2011.242).
- [93] J. De Fauw, J. R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O’Donoghue, D. Visentin, G. van den Driessche, B. Lakshminarayanan, C. Meyer, F. Mackinder, S. Bouton, K. Ayoub, R. Chopra, D. King, A. Karthikesalingam, C. O. Hughes, R. Raine, J. Hughes, D. A. Sim, C. Egan, A. Tufail, H. Montgomery, D. Hassabis, G. Rees, T. Back, P. T. Khaw, M. Suleyman, J. Cornebise, P. A. Keane, and O. Ronneberger. “Clinically applicable deep learning for diagnosis and referral in retinal disease.” In: *Nature Medicine* 24 (2018), pp. 1342–1350. doi: [10.1038/s41591-018-0107-6](https://doi.org/10.1038/s41591-018-0107-6).

- [94] **M. Röhlig**, C. Schmidt, R. K. Prakasam, O. Stachs, and H. Schumann. "Towards Accurate Visualization and Measurement of Localized Changes in Intraretinal Layer Thickness." In: *Poster at the IEEE Workshop on Visual Analytics in Healthcare (VAHC)*. 2019.
- [95] R. K. Prakasam, A. Götze, S. von Keyserlingk, A. Jünemann, **M. Röhlig**, O. Stachs, and D.-C. Fischer. "Spectral-Domain Optical Coherence Tomography for Determination of Retinal Thickness in Pediatric Patients with Mild-To-Moderate Chronic Kidney Disease: A Cross-Sectional Study." In: *Current Eye Research* 44.2 (2018), pp. 206–211. DOI: [10.1080/02713683.2018.1522649](https://doi.org/10.1080/02713683.2018.1522649).
- [96] F. C. Gundogan, F. Akay, S. Uzun, U. Yolcu, E. Çağiltay, and S. Toyran. "Early Neurodegeneration of the Inner Retinal Layers in Type 1 Diabetes Mellitus." In: *Ophthalmologica* 235 (2016), pp. 125–132. DOI: [10.1159/000442826](https://doi.org/10.1159/000442826).
- [97] U. Frydkjaer-Olsen, R. S. Hansen, T. Peto, and J. Grauslund. "Structural neurodegeneration correlates with early diabetic retinopathy." In: *International Ophthalmology* 38.4 (2018), pp. 1621–1626. DOI: [10.1007/s10792-017-0632-1](https://doi.org/10.1007/s10792-017-0632-1).
- [98] **M. Röhlig**, A. Jünemann, D.-C. Fischer, R. K. Prakasam, O. Stachs, and H. Schumann. "Visual Analysis of Retinal OCT Data." In: *Klinische Monatsblätter für Augenheilkunde* 234.12 (2017), pp. 1463–1471. DOI: [10.1055/s-0043-121705](https://doi.org/10.1055/s-0043-121705).
- [99] O. Stachs, A. Jünemann, H. Schumann, and **M. Röhlig**. "Visual Analytics in der Augenheilkunde." In: *Spitzenforschung in der Ophthalmologie* 2.1 (2017), pp. 160–164.
- [100] R. K. Prakasam, O. Stachs, C. Schmidt, M. Tafel, D.-C. Fischer, R. Guthoff, A. Jünemann, H. Schumann, and **M. Röhlig**. "Visual Analysis of Optical Coherence Tomography Data." In: *Ophthalmologe*. Vol. 114. Proceedings of the DOG congress Supplement 2. Springer Medicine, 2017, p. 39. DOI: [10.1007/s00347-017-0551-6](https://doi.org/10.1007/s00347-017-0551-6).
- [101] R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Jünemann, H. Schumann, and O. Stachs. "Visual Analysis: A new method to analyze OCT thickness data of retinal layers." In: *Investigative Ophthalmology & Visual Science*. Vol. 59. Proceedings of the ARVO annual meeting 9. ARVO, 2018.
- [102] O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, H. J. Hettlich, and **M. Röhlig**. "Visual Analytics of OCT data: Utility of deviation maps in describing retinal layer thickness changes." In: *Investigative Ophthalmology & Visual Science*. Proceedings of the ARVO annual meeting. ARVO, 2019.

- [103] O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, B. Stratmann, H. J. Hettlich, and M. Röhlig. "Visual analysis of retinal OCT data in patients with type 2 diabetes mellitus." In: *Ophthalmologe*. Proceedings of the DOG congress. Springer Medicine, 2019.
- [104] S. Takahashi, K. Yoshida, T. Nishita, and K. Shimada. "Occlusion-Free Animation of Driving Routes for Car Navigation Systems." In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (2006), pp. 1141–1148. DOI: [10.1109/TVCG.2006.167](https://doi.org/10.1109/TVCG.2006.167).
- [105] H. Lorenz and J. Döllner. "3D feature surface properties and their application in geovisualization." In: *Computers, Environment and Urban Systems* 34.6 (2010), pp. 476–483. DOI: [10.1016/j.compenvurbsys.2010.04.003](https://doi.org/10.1016/j.compenvurbsys.2010.04.003).
- [106] H. Deng, L. Zhang, X. Mao, and H. Qu. "Interactive Urban Context-Aware Visualization via Multiple Disocclusion Operators." In: *IEEE Transactions on Visualization and Computer Graphics* 22.7 (2015), pp. 1862–1874. DOI: [10.1109/TVCG.2015.2469661](https://doi.org/10.1109/TVCG.2015.2469661).
- [107] S. Pasewaldt, A. Semmo, M. Trapp, and J. Döllner. "Multi-perspective 3D panoramas." In: *International Journal of Geographical Information Science* 28.10 (2014), pp. 2030–2051. DOI: [10.1080/13658816.2014.922686](https://doi.org/10.1080/13658816.2014.922686).
- [108] S. Brooks and J. L. Whalley. "Multilayer hybrid visualizations to support 3D GIS." In: *Computers, Environment and Urban Systems* 32.4 (2008), pp. 278–292. DOI: [10.1016/j.compenvurbsys.2007.11.001](https://doi.org/10.1016/j.compenvurbsys.2007.11.001).
- [109] N. F. Polys and D. A. Bowman. "Design and display of enhancing information in desktop information-rich virtual environments: Challenges and techniques." In: *Virtual Reality* 8.1 (2004), pp. 41–54. DOI: [10.1007/s10055-004-0134-0](https://doi.org/10.1007/s10055-004-0134-0).
- [110] T. T. Elvins, D. R. Nadeau, and D. Kirsh. "Worldlets: 3D Thumbnails for Wayfinding in Virtual Environments." In: *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'97)*. ACM, 1997, pp. 21–30. DOI: [10.1145/263407.263504](https://doi.org/10.1145/263407.263504).
- [111] S. Fukatsu, Y. Kitamura, T. Masaki, and F. Kishino. "Intuitive Control of "Bird's Eye" Overview Images for Navigation in an Enormous Virtual Environment." In: *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST)*. ACM, 1998, pp. 67–76. DOI: [10.1145/293701.293710](https://doi.org/10.1145/293701.293710).
- [112] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1983. ISBN: 9781930824133.

- [113] I. Viola, M. Feixas, M. Sbert, and M. E. Gröller. “Importance-Driven Focus of Attention.” In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (2006), pp. 933–940. DOI: [10.1109/TVCG.2006.152](https://doi.org/10.1109/TVCG.2006.152).
- [114] S. Discher, R. Richter, and J. Döllner. “Interactive and view-dependent see-through lenses for massive 3D point clouds.” In: ed. by A. Abdul-Rahman. *Lecture Notes in Geoinformation and Cartography*. Springer, 2016. Chap. Advances in 3D Geoinformation, pp. 49–62. DOI: [10.1007/978-3-319-25691-7_3](https://doi.org/10.1007/978-3-319-25691-7_3).
- [115] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. “Interactive Lenses for Visualization: An Extended Survey.” In: *Computer Graphics Forum* 36.6 (2016), pp. 173–200. DOI: [10.1111/cgf.12871](https://doi.org/10.1111/cgf.12871).
- [116] M. Agrawala, D. Zorin, and T. Munzner. “Artistic Multiprojection Rendering.” In: *Proceedings of the Eurographics Workshop on Rendering Techniques*. Springer, 2000, pp. 125–136.
- [117] S. Takahashi, N. Ohta, H. Nakamura, Y. Takeshima, and I. Fujishiro. “Modeling Surperspective Projection of Landscapes for Geographical Guide-Map Generation.” In: *Computer Graphics Forum* 21.3 (2002), pp. 259–268. DOI: [10.1111/1467-8659.t01-1-00585](https://doi.org/10.1111/1467-8659.t01-1-00585).
- [118] P. Degener and R. Klein. “A Variational Approach for Automatic Generation of Panoramic Maps.” In: *ACM Transactions on Graphics* 28.1 (2009). DOI: [10.1145/1477926.1477928](https://doi.org/10.1145/1477926.1477928).
- [119] J. Cui, P. Rosen, V. Popescu, and C. Hoffmann. “A Curved Ray Camera for Handling Occlusions through Continuous Multiperspective Visualization.” In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1235–1242. DOI: [10.1109/TVCG.2010.127](https://doi.org/10.1109/TVCG.2010.127).
- [120] F. Grabler, M. Agrawala, R. W. Sumner, and M. Pauly. “Automatic Generation of Tourist Maps.” In: *ACM Transactions on Graphics* 27.3 (2008). DOI: [10.1145/1360612.1360699](https://doi.org/10.1145/1360612.1360699).
- [121] H. Qu, H. Wang, W. Cui, Y. Wu, and M. Y. Chan. “Focus+Context Route Zooming and Information Overlay in 3D Urban Environments.” In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 1547–1554. DOI: [10.1109/TVCG.2009.144](https://doi.org/10.1109/TVCG.2009.144).
- [122] D. Hirono, H.-Y. Wu, M. Arikawa, and S. Takahashi. “Constrained optimization for disoccluding geographic landmarks in 3D urban maps.” In: *Proceedings of the IEEE Pacific Visualization Symposium*. 2013, pp. 17–24. DOI: [10.1109/PacificVis.2013.6596123](https://doi.org/10.1109/PacificVis.2013.6596123).

- [123] H. Deng, L. Zhang, J. Ma, and Z. Kang. "Interactive panoramic map-like views for 3D mountain navigation." In: *Computers & Geosciences* 37.11 (2011), pp. 1816–1824. DOI: [10.1016/j.cageo.2011.04.002](https://doi.org/10.1016/j.cageo.2011.04.002).
- [124] H. Deng, L. Zhang, C. Han, Y. Ren, L. Zhang, and J. Li. "Efficient occlusion-free visualization for navigation in mountainous areas." In: *Computers & Geosciences* 52 (2013), pp. 389–397. DOI: [10.1016/j.cageo.2012.09.016](https://doi.org/10.1016/j.cageo.2012.09.016).
- [125] S. Pasewaldt, M. Trapp, and J. Döllner. "Multi-Perspective Detail+Overview Visualization for 3D Building Exploration." In: *Theory and Practice of Computer Graphics*. 2013, pp. 57–64. DOI: [10.2312/LocalChapterEvents.TPCG.TPCG13.057-064](https://doi.org/10.2312/LocalChapterEvents.TPCG.TPCG13.057-064).
- [126] M. Trapp, L. Schneider, C. Lehmann, N. Holz, and J. Döllner. "Strategies for visualising 3D points-of-interest on mobile devices." In: *Journal of Location Based Services* 5.2 (2011), pp. 79–99. DOI: [10.1080/17489725.2011.579579](https://doi.org/10.1080/17489725.2011.579579).
- [127] T. Glander, M. Trapp, and J. Döllner. "A Concept of Effective Landmark Depiction in Geovirtual 3D Environments by View-Dependent Deformation." In: *Proceedings of the International Symposium on LBS and Telecartography*. 2007.
- [128] A. Semmo, M. Trapp, J. E. Kyprianidis, and J. Döllner. "Interactive Visualization of Generalized Virtual 3D City Models using Level-of-Abstraction Transitions." In: *Computer Graphics Forum* 31.3 (2012), pp. 885–894. DOI: [10.1111/j.1467-8659.2012.03081.x](https://doi.org/10.1111/j.1467-8659.2012.03081.x).
- [129] T. Ortner, J. Sorger, H. Steinlechner, G. Hesina, H. Piringer, and E. Gröller. "Vis-A-Ware: Integrating Spatial and Non-Spatial Visualization for Visibility-Aware Urban Planning." In: *IEEE Transactions on Visualization and Computer Graphics* 23.2 (2017), pp. 1139–1151. DOI: [10.1109/TVCG.2016.2520920](https://doi.org/10.1109/TVCG.2016.2520920).
- [130] C. Oat and P. V. Sander. "Ambient aperture lighting." In: *Proceedings of the Symposium on Interactive 3D Graphics and Games*. ACM, 2007, pp. 61–64. DOI: [10.1145/1230100.1230111](https://doi.org/10.1145/1230100.1230111).
- [131] P. Barla, J. Thollot, and L. Markosian. "X-toon: An Extended Toon Shader." In: *Proceedings of the International Symposium on Non-photorealistic Animation and Rendering*. Ed. by D. DeCarlo and L. Markosian. ACM, 2006, pp. 127–132. DOI: [10.1145/1124728.1124749](https://doi.org/10.1145/1124728.1124749).
- [132] M. A. Harrower and C. A. Brewer. "ColorBrewer.org: An Online Tool for Selecting Color Schemes for Maps." In: *The Cartographic Journal* 40.1 (2003), pp. 27–37. DOI: [10.1179/000870403235002042](https://doi.org/10.1179/000870403235002042).
- [133] E. Imhof. *Cartographic relief presentation*. ESRI Press, 2007. ISBN: 1589480260.

- [134] T. Ulrich. "Rendering massive terrains using chunked level of detail control." In: *SIGGRAPH Course Notes*. Vol. 3. 5. 2002.
- [135] C. Dick, J. Krüger, and R. Westermann. "GPU Ray-Casting for Scalable Terrain Rendering." In: *Proceedings of Eurographics - Areas Papers*. 2009, pp. 43–50. DOI: [10.2312/ega.20091007](https://doi.org/10.2312/ega.20091007).
- [136] D. Strnad. "Parallel terrain visibility calculation on the graphics processing unit." In: *Concurrency and Computation: Practice and Experience* 23.18 (2011), pp. 2452–2462. DOI: [10.1002/cpe.1808](https://doi.org/10.1002/cpe.1808).
- [137] T. Scheuermann and J. Hensley. "Efficient Histogram Generation Using Scattering on GPUs." In: *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D)*. ACM, 2007, pp. 33–37. DOI: [10.1145/1230100.1230105](https://doi.org/10.1145/1230100.1230105).
- [138] S. Buschmann, M. Trapp, and J. Döllner. "Animated visualization of spatial-temporal trajectory data for air-traffic analysis." In: *The Visual Computer* 32.3 (2016), pp. 371–381. DOI: [10.1007/s00371-015-1185-9](https://doi.org/10.1007/s00371-015-1185-9).
- [139] P. Kohlmann, S. Bruckner, A. Kanitsar, and E. Gröller. "LiveSync: Deformed Viewing Spheres for Knowledge-Based Navigation." In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1544–1551. DOI: [10.1109/TVCG.2007.70576](https://doi.org/10.1109/TVCG.2007.70576).
- [140] J. C. Roberts. "State of the Art: Coordinated Multiple Views in Exploratory Visualization." In: *Proceedings of International Conference on Coordinated and Multiple Views in Exploratory Visualization*. 2007, pp. 61–71. DOI: [10.1109/CMV.2007.20](https://doi.org/10.1109/CMV.2007.20).
- [141] C. Richter, M. Luboschik, **M. Röhlig**, and H. Schumann. "Sequencing of Categorical Time Series." In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 213–214. DOI: [10.1109/VAST.2015.7347684](https://doi.org/10.1109/VAST.2015.7347684).
- [142] T. Gschwandtner, H. Schumann, J. Bernard, T. May, M. Bögl, S. Miksch, J. Kohlhammer, **M. Röhlig**, and B. Alsallakh. "Enhancing Time Series Segmentation and Labeling Through the Knowledge Generation Model." In: *Poster at Eurographics Conference on Visualization (EuroVis 2015)*. 2015.
- [143] J. Bernard, E. Dobermann, M. Bögl, **M. Röhlig**, A. Vögele, and J. Kohlhammer. "Visual-Interactive Segmentation of Multivariate Time Series." In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by N. Andrienko and M. Sedlmair. The Eurographics Association, 2016. DOI: [10.2312/eurova.20161121](https://doi.org/10.2312/eurova.20161121).
- [144] H. Obermaier and K. I. Joy. "Future Challenges for Ensemble Visualization." In: *IEEE Computer Graphics and Applications* 34.3 (2014), pp. 8–11. DOI: [10.1109/MCG.2014.52](https://doi.org/10.1109/MCG.2014.52).

- [145] B. N. Kholodenko. "Cell-Signalling Dynamics in Time and Space." In: *Nature Reviews Molecular Cell Biology* 7.3 (2006), pp. 165–176. DOI: [10.1038/nrm1838](https://doi.org/10.1038/nrm1838).
- [146] K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. "Spatio-temporal Correlations Can Drastically Change the Response of a MAPK Pathway." In: *Proceedings of the National Academy of Sciences* 107.6 (2010), pp. 2473–2478. DOI: [10.1073/pnas.0906885107](https://doi.org/10.1073/pnas.0906885107).
- [147] F. Haack, K. Burrage, R. Redmer, and A. M. Uhrmacher. "Studying the Role of Lipid Rafts on Protein Receptor Bindings with Cellular Automata." In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 10.3 (2013), pp. 760–770. DOI: [10.1109/TCBB.2013.40](https://doi.org/10.1109/TCBB.2013.40).
- [148] A. Kikuchi, H. Yamamoto, and A. Sato. "Selective Activation Mechanisms of Wnt Signaling Pathways." In: *Trends in Cell Biology* 19.3 (2009), pp. 119–129. DOI: [10.1016/j.tcb.2009.01.003](https://doi.org/10.1016/j.tcb.2009.01.003).
- [149] S. Rinzivillo, D. Pedreschi, M. Nanni, F. Giannotti, N. Andrienko, and G. Andrienko. "Visually Driven Analysis of Movement Data by Progressive Clustering." In: *Information Visualization* 7.3–4 (2008), pp. 225–239. DOI: [10.1057/palgrave.ivs.9500183](https://doi.org/10.1057/palgrave.ivs.9500183).
- [150] N. Willems, H. van de Wetering, and J. J. van Wijk. "Visualization of Vessel Movements." In: *Computer Graphics Forum* 28.3 (2009), pp. 959–966. DOI: [10.1111/j.1467-8659.2009.01440.x](https://doi.org/10.1111/j.1467-8659.2009.01440.x).
- [151] J. Wood, A. Slingsby, and J. Dykes. "Visualizing the Dynamics of London's Bicycle-Hire Scheme." In: *Cartographica* 46.4 (2011), pp. 239–251. DOI: [10.3138/carto.46.4.239](https://doi.org/10.3138/carto.46.4.239).
- [152] M. Luboschik, S. Rybacki, F. Haack, and H.-J. Schulz. "Supporting the Integrated Visual Analysis of Input Parameters and Simulation Trajectories." In: *Computers & Graphics* 39 (2014), pp. 37–47. DOI: [10.1016/j.cag.2013.09.004](https://doi.org/10.1016/j.cag.2013.09.004).
- [153] H. Piringer, W. Berger, and J. Krasser. "HyperMoVal: Interactive Visual Validation of Regression Models for Real-Time Simulation." In: *Computer Graphics Forum* 29.3 (2010), pp. 983–992. DOI: [10.1111/j.1467-8659.2009.01684.x](https://doi.org/10.1111/j.1467-8659.2009.01684.x).
- [154] T. Torsney-Weir, A. Saad, T. Möller, H.-C. Hege, B. Weber, and J.-M. Verbavatz. "Tuner: Principled Parameter Finding for Image Segmentation Algorithms Using Visual Response Surface Exploration." In: *IEEE Transactions on Visualization and Computer Graphics* 17.10 (2011), pp. 1892–1901. DOI: [10.1109/TVCG.2011.248](https://doi.org/10.1109/TVCG.2011.248).

- [155] W. Berger, H. Piringer, P. Filzmoser, and E. Gröller. “Uncertainty-Aware Exploration of Continuous Parameter Spaces Using Multivariate Prediction.” In: *Computer Graphics Forum* 30.3 (2011), pp. 911–920. DOI: [10.1111/j.1467-8659.2011.01940.x](https://doi.org/10.1111/j.1467-8659.2011.01940.x).
- [156] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. “Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation.” In: *Proceedings of the 24th annual conference on computer graphics and interactive techniques (SIGGRAPH)*. ACM, 1997, pp. 389–400. DOI: [10.1145/258734.258887](https://doi.org/10.1145/258734.258887).
- [157] S. Bruckner and T. Möller. “Result-Driven Exploration of Simulation Parameter Spaces for Visual Effects Design.” In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 1468–1476. DOI: [10.1109/TVCG.2010.190](https://doi.org/10.1109/TVCG.2010.190).
- [158] S. Bergner, M. Sedlmair, T. Möller, S. N. Abdolyousefi, and A. Saad. “ParaGlide: Interactive Parameter Space Partitioning for Computer Simulations.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.9 (2013), pp. 1499–1512. DOI: [10.1109/TVCG.2013.61](https://doi.org/10.1109/TVCG.2013.61).
- [159] V. Verma and A. Pang. “Comparative Flow Visualization.” In: *IEEE Transactions on Visualization and Computer Graphics* 10.6 (2004), pp. 609–624. DOI: [10.1109/TVCG.2004.39](https://doi.org/10.1109/TVCG.2004.39).
- [160] H. Guo, X. Yuan, J. Huang, and X. Zhu. “Coupled Ensemble Flow Line Advection and Analysis.” In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2733–2742. DOI: [10.1109/TVCG.2013.144](https://doi.org/10.1109/TVCG.2013.144).
- [161] R. van Pelt, R. Gasteiger, K. Lawonn, M. Meuschke, and B. Preim. “Comparative Blood Flow Visualization for Cerebral Aneurysm Treatment Assessment.” In: *Computer Graphics Forum* 33.3 (2014), pp. 131–140. DOI: [10.1111/cgf.12369](https://doi.org/10.1111/cgf.12369).
- [162] K. Matković, D. Gračanin, R. Splechtna, M. Jelović, B. Stehno, H. Hauser, and W. Purgathofer. “Visual Analytics for Complex Engineering Systems: Hybrid Visual Steering of Simulation Ensembles.” In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1803–1812. DOI: [10.1109/TVCG.2014.2346744](https://doi.org/10.1109/TVCG.2014.2346744).
- [163] S. Bremm, T. von Landesberger, G. Andrienko, N. Andrienko, and T. Schreck. “Interactive Analysis of Object Group Changes over Time.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. 2011, pp. 41–44. DOI: [10.2312/PE/EuroVAST/EuroVA11/041-044](https://doi.org/10.2312/PE/EuroVAST/EuroVA11/041-044).

- [164] C. Turkay, J. Parulek, N. Reuter, and H. Hauser. "Interactive Visual Analysis of Temporal Cluster Structures." In: *Computer Graphics Forum* 30.3 (2011), pp. 711–720. DOI: [10.1111/j.1467-8659.2011.01920.x](https://doi.org/10.1111/j.1467-8659.2011.01920.x).
- [165] U. Demšar and K. Verrantaus. "Space-Time Density of Trajectories: Exploring Spatio-Temporal Patterns in Movement Data." In: *International Journal of Geographical Information Science* 24.10 (2010), pp. 1527–1542. DOI: [10.1080/13658816.2010.511223](https://doi.org/10.1080/13658816.2010.511223).
- [166] H. Jeung, M. L. Yiu, and C. S. Jensen. "Trajectory Pattern Mining." In: *Computing with Spatial Trajectories*. Springer, 2011, pp. 143–177. DOI: [10.1007/978-1-4614-1629-6_5](https://doi.org/10.1007/978-1-4614-1629-6_5).
- [167] M. Garland and P. S. Heckbert. "Surface Simplification Using Quadric Error Metrics." In: *Proceedings of the 24th annual conference on computer graphics and interactive techniques (SIGGRAPH)*. ACM, 1997, pp. 209–216. DOI: [10.1145/258734.258849](https://doi.org/10.1145/258734.258849).
- [168] S. Salvador and P. Chan. "Toward Accurate Dynamic Time Warping in Linear Time and Space." In: *Intelligent Data Analysis* 11.5 (2007), pp. 561–580.
- [169] H. Lam, T. Munzner, and R. Kincaid. "Overview Use in Multiple Visual Information Resolution Interfaces." In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1278–1285. DOI: [10.1109/TVCG.2007.70583](https://doi.org/10.1109/TVCG.2007.70583).
- [170] A. T. Bittig, F. Haack, C. Maus, and A. M. Uhrmacher. "Adapting Rule-based Model Descriptions for Simulating in Continuous and Hybrid Space." In: *Proceedings of the 9th International Conference on Computational Methods in Systems Biology*. ACM, 2011, pp. 161–170. DOI: [10.1145/2037509.2037533](https://doi.org/10.1145/2037509.2037533).
- [171] M. Luboschik, C. Tominski, A. T. Bittig, A. M. Uhrmacher, and H. Schumann. "Towards Interactive Visual Analysis of Microscopic-Level Simulation Data." In: *Proceedings of SIGRAD 2012, Interactive Visual Analysis of Data*. Linköping University Electronic Press, 2012, pp. 91–94.
- [172] A. Röhm. "Visualization and editing of transfer functions for visual exploration of 3D retinal data." MA thesis. University of Rostock, 2017.
- [173] A. Baghaie, Z. Yu, and R. M. D'Souza. "State-of-the-art in retinal optical coherence tomography image analysis." In: *Quantitative imaging in medicine and surgery* 5.4 (2015), pp. 603–617. DOI: [10.3978/j.issn.2223-4292.2015.07.02](https://doi.org/10.3978/j.issn.2223-4292.2015.07.02).
- [174] Berufsverband der Augenärzte Deutschlands e. V., Deutsche Ophthalmologische Gesellschaft, and Retinologische Gesellschaft e. V. "Quality assurance of optical coherence tomography for diagnostics of the fundus: Positional statement of the BVA,

- DOG and RG." In: *Der Ophthalmologe* 114.7 (2017), pp. 617–624. doi: [10.1007/s00347-017-0508-9](https://doi.org/10.1007/s00347-017-0508-9).
- [175] D. Odell, A. M. Dubis, J. F. Lever, K. E. Stepien, and J. Carroll. "Assessing errors inherent in OCT-derived macular thickness maps." In: *Journal of Ophthalmology* 2011 (2011). doi: [10.1155/2011/692574](https://doi.org/10.1155/2011/692574).
- [176] Y. Chen, J. Li, Y. Yan, and X. Shen. "Diabetic macular morphology changes may occur in the early stage of diabetes." In: *BMC Ophthalmology* 16.1 (2016), p. 12. doi: [10.1186/s12886-016-0186-4](https://doi.org/10.1186/s12886-016-0186-4).
- [177] H. W. van Dijk, P. H. B. Kok, M. Garvin, M. Sonka, J. H. DeVries, R. P. J. Michels, M. E. J. van Velthoven, R. O. Schlingemann, F. D. Verbraak, and M. D. Abramoff. "Selective Loss of Inner Retinal Layer Thickness in Type 1 Diabetic Patients with Minimal Diabetic Retinopathy." In: *Investigative Ophthalmology & Visual Science* 50.7 (2009), pp. 3404–3409. doi: [10.1167/iovs.08-3143](https://doi.org/10.1167/iovs.08-3143).
- [178] D. El-Fayoumi, N. M. Badr Eldine, A. F. Esmael, D. Ghalwash, and H. M. Soliman. "Retinal Nerve Fiber Layer and Ganglion Cell Complex Thicknesses Are Reduced in Children With Type 1 Diabetes With No Evidence of Vascular Retinopathy." In: *Investigative Ophthalmology & Visual Science* 57.13 (2016), pp. 5355–5360. doi: [10.1167/iovs.16-19988](https://doi.org/10.1167/iovs.16-19988).
- [179] E. Moisseiev, S. Park, G. Yiu, J. S. Werner, and R. J. Zawadzki. "The Third Dimension: Advantages of 3D-OCT in Retina – Unprecedented detail of perfusion and other structures." In: *Retinal Physician* 13 (2016), pp. 24–33.
- [180] M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, T. Ko, J. S. Schuman, A. Kowalczyk, and J. S. Duker. "Three-dimensional Retinal Imaging with High-Speed Ultrahigh-Resolution Optical Coherence Tomography." In: *Ophthalmology* 112.10 (2005), pp. 1734–1746. doi: [10.1016/j.opthta.2005.05.023](https://doi.org/10.1016/j.opthta.2005.05.023).
- [181] P. Rosenthal, M. Ritter, D. Kowerko, and C. Heine. "OphthalVis – Making Data Analytics of Optical Coherence Tomography Reproducible." In: *Proceedings of EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization (EuroRV3)*. Ed. by K. Lawonn, M. Hlawitschka, and P. Rosenthal. The Eurographics Association, 2016. doi: [10.2312/eurorv3.20161109](https://doi.org/10.2312/eurorv3.20161109).
- [182] J. Schindelin, C. T. Rueden, M. C. Hiner, and K. W. Eliceiri. "The ImageJ ecosystem: An open platform for biomedical image analysis." In: *Molecular Reproduction and Development* 82.7-8 (2015), pp. 518–529. doi: [10.1002/mrd.22489](https://doi.org/10.1002/mrd.22489).

- [183] G. D. Aaker, L. Gracia, J. S. Myung, V. Borchering, J. R. Banfelder, D. J. D'Amico, and S. Kiss. "Three-dimensional reconstruction and analysis of vitreomacular traction: Quantification of cyst volume and vitreoretinal interface area." In: *Archives of Ophthalmology* 129.6 (2011), pp. 805–820. DOI: [10.1001/archophthalmol.2011.123](https://doi.org/10.1001/archophthalmol.2011.123).
- [184] J. P. Schulze, C. Schulze-Döbold, A. Erginay, and R. Tadayoni. "Visualization of three-dimensional ultra-high resolution OCT in virtual reality." In: *Studies in health technology and informatics* 184 (2013). Ed. by J. D. Westwood, S. W. Westwood, L. Felländer-Tsai, R. S. Haluck, R. A. Robb, S. Senger, and K. G. Vosburgh, pp. 387–391. DOI: [10.3233/978-1-61499-209-7-387](https://doi.org/10.3233/978-1-61499-209-7-387).
- [185] J. Probst, P. Koch, and G. Hüttmann. "Real-time 3D rendering of optical coherence tomography volumetric data." In: *Proceedings of SPIE Optical Coherence Tomography and Coherence Techniques IV*. Ed. by P. E. Andersen and B. E. Bouma. SPIE, 2009, 73720Q–73731Q. DOI: [10.1117/12.831785](https://doi.org/10.1117/12.831785).
- [186] M. Sylwestrzak, D. Szlag, M. Szkulmowski, and P. Targowski. "Real-time massively parallel processing of Spectral Optical Coherence Tomography data on Graphics Processing Units." In: *Proceedings of SPIE Optical Coherence Tomography and Coherence Techniques V*. Ed. by R. A. Leitgeb and B. E. Bouma. SPIE, 2011, pp. 80910V–80917V. DOI: [10.1117/12.889805](https://doi.org/10.1117/12.889805).
- [187] S. Glaßer, T. Hoffmann, A. Boese, S. Voß, T. Kalinski, M. Skalej, and B. Preim. "Virtual Inflation of the Cerebral Artery Wall for the Integrated Exploration of OCT and Histology Data." In: *Computer Graphics Forum* (2016). DOI: [10.1111/cgf.12994](https://doi.org/10.1111/cgf.12994).
- [188] M. D. Duncan, M. Bashkansky, and J. Reintjes. "Subsurface defect detection in materials using optical coherence tomography." In: *Optics Express* 2.13 (1998), pp. 540–545. DOI: [10.1364/OE.2.000540](https://doi.org/10.1364/OE.2.000540).
- [189] V. Placet, J. Méteau, L. Froehly, R. Salut, and M. L. Boubakar. "Investigation of the internal structure of hemp fibres using optical coherence tomography and Focused Ion Beam transverse cutting." In: *Journal of Materials Science* 49.24 (2014), pp. 8317–8327. DOI: [10.1007/s10853-014-8540-5](https://doi.org/10.1007/s10853-014-8540-5).
- [190] S. Kahl, M. Ritter, and P. Rosenthal. "Automated assessment of the injury situation in patients with age-related macular degeneration (AMD)." In: *Proceedings of Forum Bildverarbeitung*. Ed. by F. P. Leon and M. Heizmann. KIT Scientific Publishing, 2014, pp. 179–190. DOI: [10.5445/KSP/1000043608](https://doi.org/10.5445/KSP/1000043608).
- [191] P. Rosenthal, M. Ritter, D. Kowerko, and C. Heine. *Unified OCT Explorer*. <http://bitbucket.org/uocte/>. 2016.

- [192] D. N. Koleva-Georgieva. "Optical Coherence Tomography – Segmentation Performance and Retinal Thickness Measurement Errors." In: *European Ophthalmic Review* 6.2 (2012), pp. 78–82. doi: [10.17925/EOR.2012.06.02.78](https://doi.org/10.17925/EOR.2012.06.02.78).
- [193] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. "Visual Analytics: Scope and Challenges." In: *Visual Data Mining: Theory, Techniques and Tools for Visual Analytics*. Ed. by S. J. Simoff, M. H. Böhlen, and A. Mazeika. Springer, 2008, pp. 76–90. doi: [10.1007/978-3-540-71080-6_6](https://doi.org/10.1007/978-3-540-71080-6_6).
- [194] K. W. Hall, C. Perin, P. G. Kusalik, C. Gutwin, and S. Carpendale. "Formalizing Emphasis in Information Visualization." In: *Computer Graphics Forum* 35.3 (2016), pp. 717–737. doi: [10.1111/cgf.12936](https://doi.org/10.1111/cgf.12936).
- [195] R. Kosara, S. Miksch, and H. Hauser. "Semantic Depth of Field." In: *Proceedings of the IEEE Symposium on Information Visualization*. IEEE Computer Society, 2001, pp. 97–104. doi: [10.1109/INFVIS.2001.963286](https://doi.org/10.1109/INFVIS.2001.963286).
- [196] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. "Visual comparison for information visualization." In: *Information Visualization* 10.4 (2011), pp. 289–309. doi: [10.1177/1473871611416549](https://doi.org/10.1177/1473871611416549).
- [197] E. E. B. D. Clerck, J. S. A. G. Schouten, T. T. J. M. Berendschot, A. G. H. Kessels, R. M. M. A. Nuijts, H. J. M. Beckers, M. T. Schram, C. D. A. Stehouwer, and C. A. B. Webers. "New ophthalmologic imaging techniques for detection and monitoring of neurodegenerative changes in diabetes: a systematic review." In: *The Lancet Diabetes & Endocrinology* 3.8 (2015), pp. 653–663. doi: [10.1016/S2213-8587\(15\)00136-9](https://doi.org/10.1016/S2213-8587(15)00136-9).
- [198] N. Elmqvist, A. Vande Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. J. Jankun-Kelly. "Fluid Interaction for Information Visualization." In: *Information Visualization* 10.4 (2011), pp. 327–340. doi: [10.1177/1473871611413180](https://doi.org/10.1177/1473871611413180).
- [199] M. A. Mayer, J. Hornegger, C. Y. Mardin, and R. P. Tornow. "Retinal Nerve Fiber Layer Segmentation on FD-OCT Scans of Normal Subjects and Glaucoma Patients." In: *Biomedical Optics Express* 1.5 (2010), pp. 1358–1383. doi: [10.1364/BOE.1.001358](https://doi.org/10.1364/BOE.1.001358).
- [200] J. Stüwe. "Attribute-driven visualization of retinal OCT data." BA thesis. University of Rostock, 2018.
- [201] J. Stüwe. "Interactive similarity search for attribute maps of intraretinal layers." MA thesis. University of Rostock, 2019.

- [202] A. H. Kashani, C.-L. Chen, J. K. Gahm, F. Zheng, G. M. Richter, P. J. Rosenfeld, Y. Shi, and R. K. Wang. "Optical coherence tomography angiography: A comprehensive review of current methods and clinical applications." In: *Progress in Retinal and Eye Research* 60 (2017), pp. 66–100. DOI: [10.1016/j.preteyeres.2017.07.002](https://doi.org/10.1016/j.preteyeres.2017.07.002).
- [203] A. Cogliati, C. Canavesi, A. Hayes, P. Tankam, V.-F. Duma, A. Santhanam, K. P. Thompson, and J. P. Rolland. "MEMS-based handheld scanning probe with pre-shaped input signals for distortion-free images in Gabor-domain optical coherence microscopy." In: *Optics Express* 24.12 (2016), pp. 13365–13374. DOI: [10.1364/OE.24.013365](https://doi.org/10.1364/OE.24.013365).
- [204] W. Drexler, U. Morgner, F. X. Kärtner, C. Pitris, S. A. Boppart, X. D. Li, E. P. Ippen, and J. G. Fujimoto. "In vivo ultrahigh-resolution optical coherence tomography." In: *Optics Letters* 24.17 (1999), pp. 1221–1223. DOI: [10.1364/OL.24.001221](https://doi.org/10.1364/OL.24.001221).
- [205] J. Mazzaferri, L. Beaton, G. Hounye, D. N. Sayah, and S. Costantino. "Open-source algorithm for automatic choroid segmentation of OCT volume reconstructions." In: *Scientific Reports* 7 (2017). DOI: [10.1038/srep42112](https://doi.org/10.1038/srep42112).
- [206] H.-J. Schulz, **M. Röhlig**, L. Nonnemann, M. Aehnelt, H. Diener, B. Urban, and H. Schumann. "Lightweight Coordination of Multiple Independent Visual Analytics Tools." In: *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*. Ed. by A. Kerren, C. Hurter, and J. Braz. Vol. 3. International Conference on Information Visualization Theory and Applications (IVAPP). SciTePress, 2019, pp. 106–117. DOI: [10.5220/0007571101060117](https://doi.org/10.5220/0007571101060117).
- [207] D. Nelson, J. K. Mah, C. Adams, S. Hui, S. Crawford, H. Darwish, D. Stephure, and D. Pacaud. "Comparison of conventional and non-invasive techniques for the early identification of diabetic neuropathy in children and adolescents with type 1 diabetes." In: *Pediatric Diabetes* 7.6 (2006), pp. 305–310. DOI: [10.1111/j.1399-5448.2006.00208.x](https://doi.org/10.1111/j.1399-5448.2006.00208.x).
- [208] A. J. Barber. "A new view of diabetic retinopathy: a neurodegenerative disease of the eye." In: *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 27.2 (2003). Apoptosis and Human Neurodegenerative Diseases, pp. 283–290. DOI: [10.1016/S0278-5846\(03\)00023-X](https://doi.org/10.1016/S0278-5846(03)00023-X).
- [209] M. Villarroel, A. Ciudin, C. Hernández, and R. Simó. "Neurodegeneration: an early event of diabetic retinopathy." In: *World journal of diabetes* 1.2 (2010), pp. 57–64. DOI: [10.4239/wjd.v1.i2.57](https://doi.org/10.4239/wjd.v1.i2.57).

- [210] A. J. Barber, E. Lieth, S. A. Khin, D. A. Antonetti, A. G. Buchanan, and T. W. Gardner. "Neural apoptosis in the retina during experimental and human diabetes. Early onset and effect of insulin." In: *The Journal of Clinical Investigation* 102.4 (1998), pp. 783–791. DOI: [10.1172/JCI2425](https://doi.org/10.1172/JCI2425).
- [211] E. Y. Chew, M. L. Klein, F. L. Ferris III, N. A. Remaley, R. P. Murphy, K. Chantry, B. J. Hoogwerf, and D. Miller. "Association of Elevated Serum Lipid Levels With Retinal Hard Exudate in Diabetic Retinopathy: Early Treatment Diabetic Retinopathy Study (ETDRS) Report 22." In: *Archives of Ophthalmology* 114.9 (1996), pp. 1079–1084. DOI: [10.1001/archophth.1996.01100140281004](https://doi.org/10.1001/archophth.1996.01100140281004).
- [212] G. Wollstein, J. S. Schuman, L. L. Price, A. Aydin, P. C. Stark, E. Hertzmark, E. Lai, H. Ishikawa, C. Mattox, J. G. Fujimoto, and L. A. Paunescu. "Optical Coherence Tomography Longitudinal Evaluation of Retinal Nerve Fiber Layer Thickness in Glaucoma." In: *Archives of Ophthalmology* 123.4 (2005), pp. 464–470. DOI: [10.1001/archophth.123.4.464](https://doi.org/10.1001/archophth.123.4.464).
- [213] B. Mrak, T. Čaljkusić-Mance, B. Mrak, B. Cerovski, and Z. Trobonjača. "The role of the spectral domain ocular coherence tomography in detection of age-related macular degeneration." In: *Collegium antropologicum* 35.2 (2011), pp. 157–160.
- [214] A. Petzold, J. F. de Boer, S. Schippling, P. Vermersch, R. Kardon, A. Green, P. A. Calabresi, and C. Polman. "Optical coherence tomography in multiple sclerosis: a systematic review and meta-analysis." In: *The Lancet Neurology* 9.9 (2010), pp. 921–932. DOI: [10.1016/S1474-4422\(10\)70168-X](https://doi.org/10.1016/S1474-4422(10)70168-X).
- [215] J.-g. Yu, Y.-f. Feng, Y. Xiang, J.-h. Huang, G. Savini, V. Parisi, W.-j. Yang, and X.-a. Fu. "Retinal Nerve Fiber Layer Thickness Changes in Parkinson Disease: A Meta-Analysis." In: *PLOS ONE* 9.1 (2014), pp. 1–8. DOI: [10.1371/journal.pone.0085718](https://doi.org/10.1371/journal.pone.0085718).
- [216] J. den Haan, F. D. Verbraak, P. J. Visser, and F. H. Bouwman. "Retinal thickness in Alzheimer's disease: a systematic review and meta-analysis." In: *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring* 6 (2017), pp. 162–170. DOI: [10.1016/j.dadm.2016.12.014](https://doi.org/10.1016/j.dadm.2016.12.014).
- [217] C. Biallostowski, M. E. J. van Velthoven, R. P. J. Michels, R. O. Schlingemann, J. H. DeVries, and F. D. Verbraak. "Decreased optical coherence tomography-measured pericentral retinal thickness in patients with diabetes mellitus type 1 with minimal diabetic retinopathy." In: *British Journal of Ophthalmology* 91.9 (2007), pp. 1135–1138. DOI: [10.1136/bjo.2006.111534](https://doi.org/10.1136/bjo.2006.111534).

- [218] D. A. Antonetti, A. J. Barber, S. K. Bronson, W. M. Freeman, T. W. Gardner, L. S. Jefferson, M. Kester, S. R. Kimball, J. K. Krady, K. F. LaNoue, C. C. Norbury, P. G. Quinn, L. Sandirasegarane, and I. A. Simpson. "Diabetic Retinopathy — Seeing Beyond Glucose-Induced Microvascular Disease." In: *Diabetes* 55.9 (2006), pp. 2401–2411. DOI: [10.2337/db05-1635](https://doi.org/10.2337/db05-1635).
- [219] A. Chan, J. S. Duker, T. H. Ko, J. G. Fujimoto, and J. S. Schuman. "Normal Macular Thickness Measurements in Healthy Eyes Using Stratus Optical Coherence Tomography." In: *Archives of Ophthalmology* 124.2 (2006), pp. 193–198. DOI: [10.1001/archophth.124.2.193](https://doi.org/10.1001/archophth.124.2.193).
- [220] N. Mendez, S. S. Kommana, B. Szirth, and A. S. Khouri. "Structural Changes by Spectral Domain Optical Coherence Tomography in Patients With Type 1 Diabetes Mellitus." In: *Journal of Diabetes Science and Technology* 10.2 (2016), pp. 271–276. DOI: [10.1177/1932296815603371](https://doi.org/10.1177/1932296815603371).
- [221] D. C. Debuc and G. M. Somfai. "Early detection of retinal thickness changes in diabetes using optical coherence tomography." In: *Medical science monitor* 16.3 (2010), pp. 15–21.
- [222] A. M. Bagci, M. Shahidi, R. Ansari, M. Blair, N. P. Blair, and R. Zelkha. "Thickness Profiles of Retinal Layers by Optical Coherence Tomography Image Segmentation." In: *American Journal of Ophthalmology* 146.5 (2008), 679–687.e1. DOI: [10.1016/j.ajo.2008.06.010](https://doi.org/10.1016/j.ajo.2008.06.010).
- [223] W. Goebel and T. Kretzchmar-Gross. "Retinal Thickness in Diabetic Retinopathy: A Study Using Optical Coherence Tomography (OCT)." In: *Retina* 22.6 (2002), pp. 759–767. DOI: [10.1097/00006982-200212000-00012](https://doi.org/10.1097/00006982-200212000-00012).
- [224] X. Liu, M. Shen, Y. Yuan, S. Huang, D. Zhu, Q. Ma, X. Ye, and F. Lu. "Macular Thickness Profiles of Intraretinal Layers in Myopia Evaluated by Ultrahigh-Resolution Optical Coherence Tomography." In: *American Journal of Ophthalmology* 160.1 (2015), 53–61.e2. DOI: [10.1016/j.ajo.2015.03.012](https://doi.org/10.1016/j.ajo.2015.03.012).
- [225] J. A. Choi, H. W. Kim, J.-W. Kwon, Y.-s. Shim, D. H. Jee, J.-S. Yun, Y.-B. Ahn, C. K. Park, and S.-H. Ko. "Early inner retinal thinning and cardiovascular autonomic dysfunction in type 2 diabetes." In: *PLOS ONE* 12.3 (2017). Ed. by P. E. Fort, e0174377. DOI: [10.1371/journal.pone.0174377](https://doi.org/10.1371/journal.pone.0174377).
- [226] R. Dhasmana, S. Sah, and N. Gupta. "Study of Retinal Nerve Fibre Layer Thickness in Patients with Diabetes Mellitus Using Fourier Domain Optical Coherence Tomography." In: *Journal of Clinical & Diagnostic Research* 10.7 (2016), NC05–NC09. DOI: [10.7860/JCDR/2016/19097.8107](https://doi.org/10.7860/JCDR/2016/19097.8107).

- [227] S.-t. Li, X.-n. Wang, X.-h. Du, and Q. Wu. "Comparison of spectral-domain optical coherence tomography for intra-retinal layers thickness measurements between healthy and diabetic eyes among Chinese adults." In: *PLOS ONE* 12.5 (2017). Ed. by D. G. Vavvas, e0177515. DOI: [10.1371/journal.pone.0177515](https://doi.org/10.1371/journal.pone.0177515).
- [228] T. Oshitari, K. Hanawa, and E. Adachi-Usami. "Changes of macular and RNFL thicknesses measured by Stratus OCT in patients with early stage diabetes." In: *Eye* 23.4 (2009), pp. 884–889. DOI: [10.1038/eye.2008.119](https://doi.org/10.1038/eye.2008.119).
- [229] P. Gupta, S. G. Thakku, C. Sabanayagam, G. Tan, R. Agrawal, C. M. G. Cheung, E. L. Lamoureux, T.-Y. Wong, and C.-Y. Cheng. "Characterisation of choroidal morphological and vascular features in diabetes and diabetic retinopathy." In: *British Journal of Ophthalmology* 101.8 (2017), pp. 1038–1044. DOI: [10.1136/bjophthalmol-2016-309366](https://doi.org/10.1136/bjophthalmol-2016-309366).
- [230] X.-n. Wang, S.-t. Li, W. Li, Y.-J. Hua, and Q. Wu. "The thickness and volume of the choroid, outer retinal layers and retinal pigment epithelium layer changes in patients with diabetic retinopathy." In: *International Journal of Ophthalmology* 11.12 (2018), pp. 1957–1962. DOI: [10.18240/ijo.2018.12.14](https://doi.org/10.18240/ijo.2018.12.14).
- [231] S. Srinivasan, N. Pritchard, D. Vagenas, K. Edwards, G. P. Sampson, A. W. Russell, R. A. Malik, and N. Efron. "Retinal Tissue Thickness is Reduced in Diabetic Peripheral Neuropathy." In: *Current Eye Research* 41.10 (2016), pp. 1359–1366. DOI: [10.3109/02713683.2015.1119855](https://doi.org/10.3109/02713683.2015.1119855).
- [232] S. Srinivasan, N. Pritchard, G. P. Sampson, K. Edwards, D. Vagenas, A. W. Russell, R. A. Malik, and N. Efron. "Diagnostic capability of retinal thickness measures in diabetic peripheral neuropathy." In: *Journal of Optometry* 10.4 (2017), pp. 215–225. DOI: [10.1016/j.optom.2016.05.003](https://doi.org/10.1016/j.optom.2016.05.003).
- [233] P. Cirafici, G. Maiello, C. Ancona, A. Masala, C. E. Traverso, and M. Iester. "Point-wise correlations between 10-2 Humphrey visual field and OCT data in open angle glaucoma." In: *Eye* (2020). DOI: [10.1038/s41433-020-0989-7](https://doi.org/10.1038/s41433-020-0989-7).
- [234] M. Luboschik, **M. Röhlig**, G. Kundt, O. Stachs, S. Peschel, A. Zhivov, R. F. Guthoff, K. Winter, and H. Schumann. "Supporting an Early Detection of Diabetic Neuropathy by Visual Analytics." In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by M. Pohl and J. Roberts. Eurographics Association, 2014. DOI: [10.2312/eurova.20141145](https://doi.org/10.2312/eurova.20141145).
- [235] **M. Röhlig**, M. Luboschik, G. Kundt, O. Stachs, S. Peschel, A. Zhivov, R. F. Guthoff, K. Winter, and H. Schumann. "Visual Analysis for an Early Detection of Diabetic Neuropathy."

- In: *Klinische Monatsblätter für Augenheilkunde* 231.12 (2014), pp. 1162–1169. DOI: [10.1055/s-0034-1383332](https://doi.org/10.1055/s-0034-1383332).
- [236] **M. Röhlig**, O. Stachs, and H. Schumann. “Detection of Diabetic Neuropathy - Can Visual Analytics Methods Really Help in Practice?” In: *Proceedings of the EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization (EuroRV3)*. Ed. by K. Lawonn, M. Hlawitschka, and P. Rosenthal. The Eurographics Association, 2016. DOI: [10.2312/eurorv3.20161111](https://doi.org/10.2312/eurorv3.20161111).
- [237] H.-J. Schulz, **M. Röhlig**, L. Nonnemann, M. Hogräfer, M. Aehnelt, B. Urban, and H. Schumann. “A Layered Approach to Lightweight Toolchaining in Visual Analytics.” In: *Computer Vision, Imaging and Computer Graphics Theory and Applications*. Ed. by Cláudio, A. P. Bouatouch, K. Chessa, M. Paljic, A. Kerren, A. Hurter, C. Tremeau, A. Farinella, and G. Maria. Communications in Computer and Information Science. Springer, 2019, pp. 313–337. DOI: [10.1007/978-3-030-41590-7_13](https://doi.org/10.1007/978-3-030-41590-7_13).
- [238] P. Wild, T. Zeller, M. Beutel, M. Blettner, K. Dugi, K. Lackner, N. Pfeiffer, T. Münzel, and S. Blankenberg. “Die Gutenberg Gesundheitsstudie.” In: *Bundesgesundheitsblatt - Gesundheitsforschung - Gesundheitsschutz* 55.6 (2012), pp. 824–830. DOI: [10.1007/s00103-012-1502-7](https://doi.org/10.1007/s00103-012-1502-7).

LIST OF SCIENTIFIC CONTRIBUTIONS

SUMMARY

BOOK CHAPTERS:	1
JOURNAL ARTICLES:	12 (3 invited)
CONFERENCE AND WORKSHOP PAPERS:	12 (4 best paper awards)
INVITED CONFERENCE TALKS:	1
POSTERS AND DEMOS:	10

BOOK CHAPTERS

- [165] H.-J. Schulz, **M. Röhlig**, L. Nonnemann, M. Hogräfer, M. Aehnelt, B. Urban, and H. Schumann. "A Layered Approach to Lightweight Toolchaining in Visual Analytics." In: *Computer Vision, Imaging and Computer Graphics Theory and Applications*. Ed. by Cláudio, A. P. Bouatouch, K. Chessa, M. Paljic, A. Kerren, A. Hurter, C. Tremeau, A. Farinella, and G. Maria. Communications in Computer and Information Science. Springer, 2019, pp. 313–337. DOI: [10.1007/978-3-030-41590-7_13](https://doi.org/10.1007/978-3-030-41590-7_13).

JOURNAL ARTICLES

- [34] O. Stachs, A. Jünemann, H. Schumann, and **M. Röhlig**. "Visual Analytics in der Augenheilkunde." In: *Spitzenforschung in der Ophthalmologie* 2.1 (2017), pp. 160–164.
- [67] M. Luboschik, **M. Röhlig**, A. T. Bittig, N. Andrienko, H. Schumann, and C. Tominski. "Feature-Driven Visual Analytics of Chaotic Parameter-Dependent Movement." In: *Computer Graphics Forum* 34.3 (2015). Best paper award, pp. 421–430. DOI: [10.1111/cgf.12654](https://doi.org/10.1111/cgf.12654).
- [73] R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Götze, A. Jünemann, H. Schumann, and O. Stachs. "Deviation maps for understanding thickness changes of inner retinal layers in children with type 1 diabetes mellitus." In: *Current Eye Research* 44.7 (2019). DOI: [10.1080/02713683.2019.1591463](https://doi.org/10.1080/02713683.2019.1591463).
- [79] **M. Röhlig**, A. Jünemann, D.-C. Fischer, R. K. Prakasam, O. Stachs, and H. Schumann. "Visual Analysis of Retinal OCT Data." In: *Klinische Monatsblätter für Augenheilkunde* 234.12 (2017), pp. 1463–1471. DOI: [10.1055/s-0043-121705](https://doi.org/10.1055/s-0043-121705).

- [83] **M. Röhlig**, C. Schmidt, R. K. Prakasam, P. Rosenthal, H. Schumann, and O. Stachs. “Visual Analysis of Retinal Changes with Optical Coherence Tomography.” In: *The Visual Computer* 34.9 (2018), pp. 1209–1224. DOI: [10.1007/s00371-018-1486-x](https://doi.org/10.1007/s00371-018-1486-x).
- [89] **M. Röhlig**, M. Luboschik, and H. Schumann. “Visibility Widgets for Unveiling Occluded Data in 3D Terrain Visualization.” In: *Journal of Visual Languages & Computing* 42.Supplement C (2017), pp. 86–98. DOI: [10.1016/j.jvlc.2017.08.008](https://doi.org/10.1016/j.jvlc.2017.08.008).
- [103] R. K. Prakasam, A. Götze, S. von Keyserlingk, A. Jünemann, **M. Röhlig**, O. Stachs, and D.-C. Fischer. “Spectral-Domain Optical Coherence Tomography for Determination of Retinal Thickness in Pediatric Patients with Mild-To-Moderate Chronic Kidney Disease: A Cross-Sectional Study.” In: *Current Eye Research* 44.2 (2018), pp. 206–211. DOI: [10.1080/02713683.2018.1522649](https://doi.org/10.1080/02713683.2018.1522649).
- [119] S. Dübel, **M. Röhlig**, C. Tominksi, and H. Schumann. “Visualizing 3D Terrain, Geo-spatial Data, and Uncertainty.” In: *Informatics* 4.1 (2017). DOI: [10.3390/informatics4010006](https://doi.org/10.3390/informatics4010006).
- [145] A. Götze, S. von Keyserlingk, S. Peschel, U. Jacoby, C. Schreiber, B. Köhler, S. Allgeier, K. Winter, **M. Röhlig**, A. Jünemann, R. Guthoff, O. Stachs, and D.-C. Fischer. “The corneal subbasal nerve plexus and thickness of the retinal layers in pediatric type 1 diabetes and matched controls.” In: *Scientific Reports* 8.1 (2018). DOI: [10.1038/s41598-017-18284-z](https://doi.org/10.1038/s41598-017-18284-z).
- [158] **M. Röhlig**, M. Luboschik, G. Kundt, O. Stachs, S. Peschel, A. Zhivov, R. F. Guthoff, K. Winter, and H. Schumann. “Visual Analysis for an Early Detection of Diabetic Neuropathy.” In: *Klinische Monatsblätter für Augenheilkunde* 231.12 (2014), pp. 1162–1169. DOI: [10.1055/s-0034-1383332](https://doi.org/10.1055/s-0034-1383332).
- [169] **M. Röhlig**, R. K. Prakasam, J. Stüwe, C. Schmidt, O. Stachs, and H. Schumann. “Enhanced Grid-Based Visual Analysis of Retinal Layer Thickness with Optical Coherence Tomography.” In: *Information* 10.9 (2019), p. 266. DOI: [10.3390/info10090266](https://doi.org/10.3390/info10090266).
- [193] R. K. Prakasam, A. Matuszewska-Iwanicka, D.-C. Fischer, H. Schumann, D. Tschöpe, B. Stratmann, H.-J. Hettlich, R. F. Guthoff, O. Stachs, and **M. Röhlig**. “Thickness of intraretinal layers in patients with type 2 diabetes mellitus depending on a concomitant diabetic neuropathy: results of a cross-sectional study using deviation maps for OCT data analysis.” In: *Biomedicines* 8.7 (2020). DOI: [10.3390/biomedicines8070190](https://doi.org/10.3390/biomedicines8070190).

CONFERENCE AND WORKSHOP PAPERS

- [28] **M. Röhlig** and H. Schumann. “Visibility Widgets: Managing Occlusion of Quantitative Data in 3D Terrain Visualization.” In:

- Proceedings of the 9th International Symposium on Visual Information Communication and Interaction (VINCI)*. Best paper award. ACM, 2016, pp. 51–58. DOI: [10.1145/2968220.2968230](https://doi.org/10.1145/2968220.2968230).
- [33] M. Luboschik, **M. Röhlig**, G. Kundt, O. Stachs, S. Peschel, A. Zhivov, R. F. Guthoff, K. Winter, and H. Schumann. “Supporting an Early Detection of Diabetic Neuropathy by Visual Analytics.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by M. Pohl and J. Roberts. Eurographics Association, 2014. DOI: [10.2312/eurova.20141145](https://doi.org/10.2312/eurova.20141145).
- [42] **M. Röhlig**, P. Rosenthal, C. Schmidt, H. Schumann, and O. Stachs. “Visual Analysis of Optical Coherence Tomography Data in Ophthalmology.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by M. Sedlmair and C. Tominski. Best paper award. The Eurographics Association, 2017. DOI: [10.2312/eurova.20171117](https://doi.org/10.2312/eurova.20171117).
- [51] **M. Röhlig**, M. Luboschik, M. Bögl, F. Krüger, B. Alsallakh, T. Kirste, S. Miksch, and H. Schumann. “Supporting Activity Recognition by Visual Analytics.” In: *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 41–48. DOI: [10.1109/VAST.2015.7347629](https://doi.org/10.1109/VAST.2015.7347629).
- [56] S. Dübel, **M. Röhlig**, H. Schumann, and M. Trapp. “2D and 3D Presentation of Spatial Data: A Systematic Review.” In: *Proceedings of the IEEE VIS International Workshop on 3DVis: Does 3D really make sense for Data Visualization?* 2014. DOI: [10.1109/3DVis.2014.7160094](https://doi.org/10.1109/3DVis.2014.7160094).
- [60] D.-C. Fischer, A. Götze, S. Peschel, N. Redepenning, **M. Röhlig**, and O. Stachs. “Die optische Kohärenztomografie zur Früherkennung der retinalen Neurodegeneration: eine Querschnittsstudie an pädiatrischen Patienten mit eingeschränkter Nierenfunktion oder einem Typ 1 Diabetes mellitus.” In: *Nieren- und Hochdruckkrankheiten*. Vol. 47. Proceedings of the annual meeting of the Society of Pediatric Nephrology (GPN) 2. Dustri-Verlag, 2018.
- [69] **M. Röhlig**, O. Stachs, and H. Schumann. “Detection of Diabetic Neuropathy - Can Visual Analytics Methods Really Help in Practice?” In: *Proceedings of the EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization (EuroRV3)*. Ed. by K. Lawonn, M. Hlawitschka, and P. Rosenthal. The Eurographics Association, 2016. DOI: [10.2312/eurorv3.20161111](https://doi.org/10.2312/eurorv3.20161111).
- [96] **M. Röhlig**, J. Stüwe, C. Schmidt, R. K. Prakasam, O. Stachs, and H. Schumann. “Grid-Based Exploration of OCT Thickness Data of Intraretinal Layers.” In: *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*. Ed. by A. Ker-

- ren, C. Hurter, and J. Braz. Vol. 3. International Conference on Information Visualization Theory and Applications (IVAPP). Best student paper award. SciTePress, 2019, pp. 129–140. DOI: [10.5220/0007580001290140](https://doi.org/10.5220/0007580001290140).
- [123] J. Bernard, E. Dobermann, M. Bögl, **M. Röhlig**, A. Vögele, and J. Kohlhammer. “Visual-Interactive Segmentation of Multivariate Time Series.” In: *Proceedings of the EuroVis Workshop on Visual Analytics (EuroVA)*. Ed. by N. Andrienko and M. Sedlmair. The Eurographics Association, 2016. DOI: [10.2312/eurova.20161121](https://doi.org/10.2312/eurova.20161121).
- [152] R. K. Prakasam, O. Stachs, C. Schmidt, M. Tafel, D.-C. Fischer, R. Guthoff, A. Jünemann, H. Schumann, and **M. Röhlig**. “Visual Analysis of Optical Coherence Tomography Data.” In: *Ophthalmologie*. Vol. 114. Proceedings of the DOG congress Supplement 2. Springer Medicine, 2017, p. 39. DOI: [10.1007/s00347-017-0551-6](https://doi.org/10.1007/s00347-017-0551-6).
- [166] H.-J. Schulz, **M. Röhlig**, L. Nonnemann, M. Aehnelt, H. Diener, B. Urban, and H. Schumann. “Lightweight Coordination of Multiple Independent Visual Analytics Tools.” In: *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*. Ed. by A. Kerren, C. Hurter, and J. Braz. Vol. 3. International Conference on Information Visualization Theory and Applications (IVAPP). SciTePress, 2019, pp. 106–117. DOI: [10.5220/0007571101060117](https://doi.org/10.5220/0007571101060117).
- [168] C. Schmidt, **M. Röhlig**, B. Grundel, P. Daumke, M. Ritter, A. Stahl, P. Rosenthal, and H. Schumann. “Combining Visual Cleansing and Exploration for Clinical Data.” In: *Proceedings of the IEEE Workshop on Visual Analytics in Healthcare (VAHC)*. 2019, pp. 25–32. DOI: [10.1109/VAHC47919.2019.8945034](https://doi.org/10.1109/VAHC47919.2019.8945034).

INVITED CONFERENCE TALKS:

- [171] **M. Röhlig**, C. Schmidt, R. K. Prakasam, H. Schumann, and O. Stachs. “Visual Analysis of Retinal Changes with Optical Coherence Tomography.” Invited talk at Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM). 2019.

POSTERS AND DEMOS:

- [10] C. Richter, M. Luboschik, **M. Röhlig**, and H. Schumann. “Sequencing of Categorical Time Series.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2015, pp. 213–214. DOI: [10.1109/VAST.2015.7347684](https://doi.org/10.1109/VAST.2015.7347684).

- [30] T. Gschwandtner, H. Schumann, J. Bernard, T. May, M. Bögl, S. Miksch, J. Kohlhammer, **M. Röhlig**, and B. Alsallakh. “Enhancing Time Series Segmentation and Labeling Through the Knowledge Generation Model.” In: *Poster at Eurographics Conference on Visualization (EuroVis 2015)*. 2015.
- [53] **M. Röhlig**, M. Luboschik, M. Tavakoli, and H. Schumann. “Assessing Corneal Nerve Morphology with Visual Analytics.” In: *Poster at the 25th NEURODIAB Conference*. 2015.
- [64] J. Stüwe, **M. Röhlig**, H. Schumann, R. K. Prakasam, and O. Stachs. “Visual Analysis of Abnormal Thickness of Intraretinal Layers.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2018.
- [101] **M. Röhlig**, M. Luboschik, H. Schumann, M. Bögl, B. Alsallakh, and S. Miksch. “Analyzing Parameter Influence on Time-Series Segmentation and Labeling.” In: *Poster at the IEEE Conference on Visual Analytics Science and Technology (VAST)*. 2014. DOI: [10.1109/VAST.2014.7042524](https://doi.org/10.1109/VAST.2014.7042524).
- [111] O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, H. J. Hettlich, and **M. Röhlig**. “Visual Analytics of OCT data: Utility of deviation maps in describing retinal layer thickness changes.” In: *Investigative Ophthalmology & Visual Science*. Proceedings of the ARVO annual meeting. ARVO, 2019.
- [116] R. K. Prakasam, **M. Röhlig**, D.-C. Fischer, A. Jünemann, H. Schumann, and O. Stachs. “Visual Analysis: A new method to analyze OCT thickness data of retinal layers.” In: *Investigative Ophthalmology & Visual Science*. Vol. 59. Proceedings of the ARVO annual meeting 9. ARVO, 2018.
- [138] **M. Röhlig**, M. Luboschik, R. K. Prakasam, O. Stachs, and H. Schumann. “Visually Analyzing Parameter Influence on Optical Coherence Tomography Data in Ophthalmology.” In: *EuroVis 2017 - Posters*. Ed. by A. P. Puig and T. Isenberg. The Eurographics Association, 2017. DOI: [10.2312/eurp.20171176](https://doi.org/10.2312/eurp.20171176).
- [167] **M. Röhlig**, C. Schmidt, R. K. Prakasam, O. Stachs, and H. Schumann. “Towards Accurate Visualization and Measurement of Localized Changes in Intraretinal Layer Thickness.” In: *Poster at the IEEE Workshop on Visual Analytics in Healthcare (VAHC)*. 2019.
- [170] O. Stachs, R. K. Prakasam, D.-C. Fischer, H. Schumann, A. Matuszewska, D. Tschöpe, B. Stratmann, H. J. Hettlich, and **M. Röhlig**. “Visual analysis of retinal OCT data in patients with type 2 diabetes mellitus.” In: *Ophthalmologe*. Proceedings of the DOG congress. Springer Medicine, 2019.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede and Ivo Pletikosić. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both \LaTeX and LyX :

<https://bitbucket.org/amiede/classicthesis/>