

Komparative Visualisierung von Segmentierungen der Beckenorgane

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Medizinische Informatik

eingereicht von

Oliver Reiter, BSc

Matrikelnummer 01226942

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Mitwirkung: Univ.Ass. Renata Raidou, MSc PhD

Wien, 4. März 2018

Oliver Reiter

Eduard Gröller

Comparative Visualization of Pelvic Organ Segmentations

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Medical Informatics

by

Oliver Reiter, BSc

Registration Number 01226942

to the Faculty of Informatics

at the TU Wien

Advisor: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Assistance: Univ.Ass. Renata Raidou, MSc PhD

Vienna, 4th March, 2018

Oliver Reiter

Eduard Gröller

Erklärung zur Verfassung der Arbeit

Oliver Reiter, BSc
Leharstraße 20
4400 Steyr

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 4. März 2018

Oliver Reiter

Acknowledgements

First and foremost, I want to thank everyone who has supported me throughout the time of my studies in some way. I want to thank the Research Division of Computer Graphics, which made this thesis possible and especially my supervisor Renata Raidou, who has supported me a lot with constructive feedback and ideas.

Also many thanks to my family for providing financial and emotional support throughout my studies, without which I would not have been able to pursue my academic interests.

In the end, I want to thank my girlfriend Katharina Zirknitzer for always listening to my concerns when needed and my good friend Daniel Luegmayr, who has made my studies in Vienna a time to remember.

Kurzfassung

Automatische Segmentierung der Beckenorgane spielt eine große Rolle bei der Behandlung von Prostatakrebs und hat hohe Ansprüche an die Genauigkeit der Ergebnisse. Segmentierungsexperten arbeiten ständig daran, ihre Algorithmen zu verbessern, jedoch ist die anatomische Variabilität der Organe und Strukturen oft ein Grund, warum diese Algorithmen schlechte Ergebnisse produzieren. Deswegen ist ein Verständnis dafür, warum dies in einem bestimmten Fall passiert, für die Verbesserung sehr wichtig. Segmentierungsexperten vermuten, dass die Form und Größe von Organen eine große Auswirkung auf das Ergebnis der Algorithmen haben, jedoch fehlen ihnen derzeit noch Möglichkeiten, dies zu untersuchen und zu analysieren, da derzeitige Lösungen noch zu wenig Einblick bieten.

Diese Arbeit beschäftigt sich mit dem Design und der Implementierung einer webbasierten Anwendung, die eine einfache Untersuchung und Analyse der Variabilität von Formen der Organe ermöglicht, um Hypothesen über das Verhältnis von Formen zu den entsprechenden Segmentierungsergebnissen aufzustellen. Außerdem wird eine neue Art der komparativen Visualisierung von mehreren Organen mehrerer Patienten vorgestellt, die einen genauen Vergleich von Formen ermöglicht. Die Anwendung wurde mit einem Datensatz von Segmentierungen einer Kohorte von 17 Patienten getestet, wobei jedes Segmentierungsnetz aus vier Organen und zwei Kontaktflächen besteht, deren Dreiecke jeweils übereinstimmen. Die entwickelten Tools ermöglichen, fehlerhafte Segmentierungen schnell zu erkennen und Hypothesen über die Beziehung der Variabilität von Organen zu anatomischen Gegebenheiten und auch der Qualität der Segmentierung zu erstellen. Obwohl der Ansatz der Arbeit nur auf Segmentierungen der Beckenorgane angewandt wurde, kann er auch für andere Anwendungen erweitert werden, wie zum Beispiel den Vergleich von Segmentierungsalgorithmen oder die Analyse von anatomischer Variabilität.

Abstract

Automatic segmentation of pelvic organs plays a major role in prostate cancer treatment and has high accuracy requirements. Segmentation experts are continuously working on improving their algorithms. However, natural anatomical variability of organs and structures is a common reason for which segmentation algorithms fail. Understanding why an algorithm fails in a specific case is of major importance. Segmentation experts expect that the shape and size of the organs can play an important role in the performance of their algorithms, but current means of exploration and analysis are limited and do not provide the necessary insight.

This thesis discusses the design and implementation of a web-based application allowing for easy exploration and analysis of shape variability in order to generate hypotheses about the relation between algorithm performance and shape of organs. A new way of comparatively visualizing multiple organs of multiple patients is introduced for a detailed shape comparison. The application was tested with segmentation meshes of a cohort of 17 patients, each consisting of four pelvic organs and two organ-interfaces, which are labeled and have per-triangle correspondence. The proposed tools already allow users to quickly identify mis-segmented organs and hypothesize about the relation of variability to anatomical features as well as segmentation quality. The approach was applied on pelvic organ segmentations, but it can be extended to other applications like comparison of segmentation algorithms or analysis of anatomical variability in general.

Contents

Kurzfassung	ix
Abstract	xi
Contents	xiii
1 Introduction	1
1.1 Motivation and Problem Definition	1
1.2 Aim of the Work	1
1.2.1 Tasks and Challenges	2
1.2.2 Requirements	2
1.3 Methodological Approach & Contributions	3
1.4 Structure	3
2 Clinical Background	5
2.1 Anatomy of the Pelvic Organs	5
2.2 Prostate Cancer	5
2.3 Automatic Segmentation of Pelvic Organs	6
3 Related Work	11
3.1 Shape Spaces & Descriptors	11
3.2 Shape Space Visualization	12
3.3 Comparative Visualization	14
4 Visualization Design	19
4.1 Application Overview	19
4.2 User Interface Description	20
4.3 Task 1: Shape Variability	21
4.3.1 Shape Descriptors	21
4.3.2 Similarity Measure	25
4.3.3 Dimensionality Reduction	25
4.3.4 Visualization Strategy	29
4.4 Task 2: Comparative Visualization	31
4.4.1 Displaying Shapes of Multiple Organs	31
	xiii

4.4.2	Visualizing the Median Shape with Variances	32
5	Implementation	39
5.1	Application Overview	39
5.2	Task 1: Shape Descriptor	41
5.3	Task 1: Dimensionality Reduction and 2D Visualization	42
5.4	Task 2: Comparative Visualization	43
6	Results	47
6.1	Exploration of All Organs	50
6.2	Per-Organ Exploration	50
6.3	Comparative Visualization	55
6.4	Exploration based on Different Properties	59
6.5	Critical Review	63
6.6	Evaluation	67
7	Conclusion	69
7.1	Summary	69
7.2	Limitations and Future Work	70
	List of Figures	73
	Glossary	77
	Acronyms	79
	Bibliography	81

Introduction

1.1 Motivation and Problem Definition

Segmentation of pelvic organs plays a major role in prostate cancer treatment and has high accuracy requirements. This is because it is used as input to radiotherapy, a common carcinoma treatment, which uses a high irradiation dose that could also damage healthy tissues, if the latter are not accurately segmented and preserved during the administration of the radiation.

As manual segmentation is time consuming, automatic approaches are often preferred, which can produce inaccurate results in some cases, due to the high variability in shape and size of the organs [CLR11]. Therefore segmentation experts are continuously working on improving their algorithms. Although evaluation of the outcome is often simple, understanding why an algorithm fails in a specific case is a major problem [Rai+16]. Researchers working on the development of segmentation algorithms expect that the shape and size of the organs may play an important role in the performance of their algorithms. Therefore, it is interesting to explore the relation between the shape and size of the pelvic organs with respect to the algorithm performance between patients.

The exploration of this relation is currently a challenge, as the existing means of analysis are limited. Yet, the research area of *Visual Analytics*, which is “dedicated to the development of tools that combine automated analysis techniques with interactive visualizations for an effective analysis of heterogeneous data sets” [Lan+13], can be of great aid.

1.2 Aim of the Work

The aim of this thesis is to design and implement a tool for the visual exploration and analysis of shape variability and its impact on segmentation results. By this, the following research questions should be answered:

1. How can shape variability between organs be measured and visualized at a single- and multi-organ basis? How does this variability affect the segmentation outcome quality?
2. How can multiple segmentations of multiple organs be visualized comparatively, with respect to each other?

1.2.1 Tasks and Challenges

In order to answer the previously formulated questions, the work can be split in the following tasks, each having different challenges:

Shape variability of (more than) one organ

The scope of the thesis is at first, to provide a way to quantify the variability in the shape of the different involved organs, i.e. prostate, bladder, rectum and seminal vesicles, with adequate metrics. Then this variability has to be visualized. The challenge in the visualization design is to effectively show the variability in the segmentation outcomes. In this application we are dealing with different shapes of different organs, e.g. bladder and rectum, as well as a high variability for each organ. Finally, functionality to analyze shape variability at a single- and multi-organ-level has to be provided.

Comparative visualization of shapes of multiple organs as well as multiple patients

Comparative visualization, even between two results, already poses some challenges in design and requirements [Lan+13]. There are some approaches, which try to compare more than two results, but they are all limited to a single organ, as it is described in Chapter 3. The data set of this thesis covers the four mentioned organs and their interfaces in a cohort of 17 patients. The main challenge here is dealing with multiple organs across multiple patients and performing a comparison of the segmentation outcome regarding also the interfaces between them, where avoiding occlusion is a challenging topic. The task itself is to provide a visualization technique, which is able to display multiple organ classes while still retaining the ability to analyze shape variability.

1.2.2 Requirements

As segmentation experts wish to compare a great number of patients, the tool supports this by being compact and understandable and scalable to a larger number of organs and patients. Due to the complexity and large amount of data in Visual Analytics, the general structure of the application follows Schneiderman's mantra "*overview first, zoom and filter, then details on demand*" [Shn96]. To be able to explore and filter the data, the user is provided with a common strategy in Visual Analytics, which is brushing and highlighting in linked views. This further raises the need for high responsiveness and interactive frame-rates of the application, at least during the selection.

1.3 Methodological Approach & Contributions

To fulfill the first task and in order to be able to compare different shapes, they are represented by a shape descriptor, which extracts a number of features usable for further analysis. For the easier exploration of the high-dimensional space of shape descriptors, the distribution of shapes is visualized by using dimensionality reduction for the single- and for the multi-organ variability exploration.

The second task is accomplished by showing an exploded view of all medians of organs. This allows the user to have a summarized visual context of what the multitude of selected organs looks like, as well as the interfaces. A detailed view highlighting local variability of organs is needed to be able to visually compare different shapes. This is done by encoding the local relative extent to the other shapes and showing the local variance in the data.

The *contribution* of this thesis is the design and implementation of a web-based application, which allows segmentation experts to explore and analyze the variability of shapes of pelvic organs. This is achieved by showing the distribution in shape space and local features. Our tools allow segmentation experts to generate new hypotheses about the segmentation algorithms used. The approach of the exploded view, in conjunction with the detailed encoding of features per organ, contributes to the area of comparative visualization and shows a way to handle multiple objects at the same time, preserving their global relation to each other.

1.4 Structure

In the following Chapter 2, a short introduction into the medical background of the discussed organs is given, as well as some fundamentals about automatic segmentation. Related work and other approaches of comparative visualization are discussed in Chapter 3. In addition, different types of shape descriptors are compared with each other.

Chapter 4 deals with the design and implementation of the proposed tool. All techniques used are discussed in detail and compared with other approaches, which were discarded at the implementation phase.

The implemented parts of the visualization design are discussed in Chapter 5, along with the software architecture and other implementation details. The resulting application is then described from the view of segmentation experts in Chapter 6. Therefore, multiple usage scenarios for each research question are proposed. In the last Chapter 7, advantages and disadvantages of the implemented tool are discussed as well as improvements for future work are proposed.

Clinical Background

2.1 Anatomy of the Pelvic Organs

Our application scenario is the segmentation of pelvic organs involved in radiotherapy treatment of prostate cancer. There are three pelvic organs other than the prostate, which are of interest, because of their adjacency to it: the bladder, the seminal vesicles, and the rectum. The anatomy and relation of these organs to each other is shown in Figure 2.1.

The bladder is the largest organ and has high variability, as it is a flexible organ, the shape of which depends on its filling state. The prostate itself normally is not of high shape variability, but has to be segmented best in this scope, because it is the main target in radiotherapy processes. Posterior and between the prostate and the bladder are the pair of seminal vesicles. They have a highly variable shape and their usually small size makes them difficult to segment robustly. Posterior to the three mentioned organs is the rectum, which is the last part of the alimentary canal and because it can also be filled, it is of high shape variability, as shown in Figure 2.2.

2.2 Prostate Cancer

Prostate cancer is the second most common type of cancer in males worldwide, accounting for 15% of all newly diagnosed cancers [Fer+15] and the most common cancer in more developed regions. The age is the dominant risk factor for developing invasive prostate cancer, with a risk of 0.3% from birth to 49 years, but a risk of 12.9% from birth to death [SMJ17]. Screening and detection is both based on the prostate-specific antigen (PSA), which is the best predictor, as well as digital rectal examination (DRE), followed by a diagnosis usually based on biopsy, but which is avoided if possible [Mot+17]. The therapy of the scope, which is radiation therapy, is one of the three major curative




Figure 2.1: Male pelvic organs seen from right side [Gra18].

treatment options, besides radical prostatectomy and brachytherapy, and can be used at all risk groups [Mot+17]. Radiotherapy is a therapeutic procedure, where tumors are irradiated with a high dose, while the surrounding healthy tissue is preserved. Prior to the treatment, a planning is performed to decide the radiation at each location of the body. This is illustrated in Figure 2.3 and requires segmentation.

2.3 Automatic Segmentation of Pelvic Organs

In radiotherapy, segmentation of pelvic organs is used to determine boundaries within which the radiation dose should be administered without harming healthy organ tissue. Segmenting anatomic structures of the pelvis automatically is still challenging [CLR11]




Figure 2.2: Comparison between a filled (left) and an empty rectum

because boundaries between soft tissues are hard to detect. Solutions to this problem are generally model-based, because these approaches are very robust to noise and organ deformations. Some examples are from Chen et al. [CLR11] or Pekar et al. [PMK04], which all use a previously trained model, that is annealed to the optimal segmentation of the data by an energy or cost function. The result is either landmark data or an organ mesh, of which the latter is used in this work. An example is shown in Figure 2.4. The segmentation data assumed in this thesis result from the algorithm of Schadewaldt et al. [Sch+13] for segmentation of pelvic structures in CT images, which is based on the algorithm of Ecabert et al. [Eca+11]. In this algorithm, illustrated in Figure 2.5, for each structure to be segmented, an initial low resolution model is fitted and then adapted to fulfill some boundary criteria, based on selected features. Then, the high resolution model is initialized and further adapted. To speed this up, well segmented parts are frozen.

Statistical shape models in general are built from aligned and annotated ground truth data, of which the variability is modeled. In this context, Principal Component Analysis




Figure 2.3: Radiotherapy planning: The top figure shows the field arrangement of the radiation and the bottom one shows the isodose distribution in the tissue [Pal+08].

(PCA) is a tool to generate principal modes of deviation to the mean shape. The parameters of the shape model then describe a shape by encoding the deviation to the mean shape using the previously obtained modes [Coo+92].

Low contrast and bad image quality, e.g. noise induced by the motion of the organs, can lead to far-from-perfect segmentation outcomes, as the algorithm generally relies too much on the model in this case. In contrast to segmenting bones for example, where model-based approaches are fast and accurate, this is a major issue in segmenting the soft tissues discussed here.


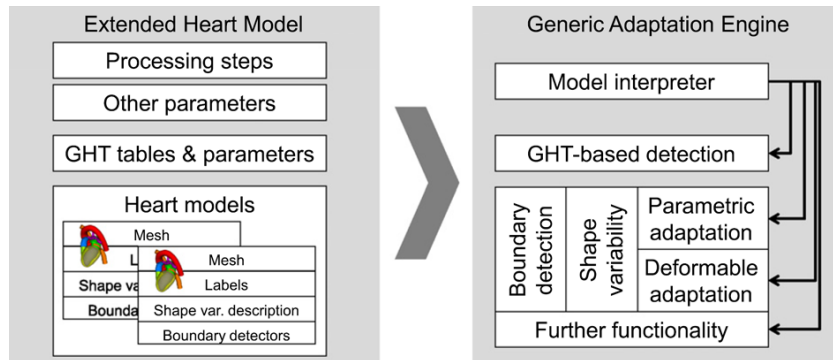



Figure 2.4: Segmentation mesh of the relevant pelvic organs, labeled by color.



(a) Architecture of the segmentation framework. The engine is controlled by the model and provides generic techniques for the automatic adaption of a mesh model to images.



(b) Improvement of the segmentation accuracy over different steps in the pipeline, in which different structures are adapted successively.

Figure 2.5: Architecture (a) and model adaption (b) of the algorithm of Ecabert et al. [Eca+11].

Related Work

3.1 Shape Spaces & Descriptors

A shape is defined by a number of parameters, for example landmark data. Other representations may be “implicit surfaces, CSG (constructive solid geometry), BSP (binary space partitioning) trees, octrees, B-reps (boundary representations), free form surfaces, etc.” [TV04] To be able to compare different shapes, they have to be converted into a suitable vector space, a so-called shape space, which is mostly of reduced dimension. The representation of a shape in this new space is determined by the shape descriptor [Tai12]. Different shape descriptors can be classified by their invariance to transformations of the original shape or by their actual representation. Tangelder et al. [TV04] describe three broad categories of shape matching methods: feature-based, graph-based, and geometry-based. On the other hand, Kazmi et al. [KYZ13] provide a classification for 3D shape descriptors into view-based, histogram-based, transform-based, graph-based and, hybrid ones.

There is always a trade-off between shape descriptor complexity and efficiency, but Peura et al. [PI97] showed that a combination of simple feature-based shape descriptors can already produce good results for shape contours and should be preferred, if it is sufficient, because of its understandability. The conclusion is, that the descriptor has to be suitable for the data anyway, but should be as simple as possible. Zhang and Chen [ZC01] developed an algorithm to extract global features such as volume, moments, and Fourier transform coefficients efficiently.

In general, descriptors based on spatial maps distinguish very well, if details are not taken into account [TV04], e.g. methods using spherical harmonics [KFR03] or 3D Zernike moments [NK03]. The first approach is used in this thesis due to the high discriminative power. An extension for higher robustness is provided by Liu et al. [Liu+04].

Shape descriptors are used for statistical shape analysis. Kendall’s theory of shape space [Ken77] and Generalized Procrustes Analysis (GPA) as a tool for aligning shapes for further multivariate analysis [Ken84; Goo91] have been used for years as a standard for landmark data [ARS04]. To describe shape variability or mean shapes of landmark data, PCA can be performed on the Procrustes coordinates [DM16; Ken94]. The goal is to describe a mean shape and principal components, which are orthogonal to each other and describe the decomposition of the highest variability in the data. It is also used to develop point distribution models [Coo+92], which can be used in Active Shape Models for segmentation [CT92]. PCA is also a useful tool to perform dimensionality reduction, because a limitation to the first principal components always represents the selection of maximum variability.

To compare similarity of shapes, a measure has to be defined. If the properties identity, positivity, symmetry, and triangle inequality are fulfilled, it is a metric [TV04]. The similarity of two feature vectors is quantified by their distance in shape space. Some examples of metrics used in shape analysis are the Minkowski distance (especially the Euclidean distance [ETA02]), Hausdorff distance [ABB95; HKR93] and Procrustes distance [DM16].

3.2 Shape Space Visualization

There are several techniques for visualizing shape space and comparing shapes based on their similarity. One of them is provided by Busking et al. [BBP10] with the shape space scatterplot, shown in Figure 3.1, which is also used and extended by Kirschner et al. [KBW11] for visualizing the underlying distribution of shapes, which is often not Gaussian. This technique was then used by von Landesberger et al. [Lan+13] in a tool for model selection and is basically a 2D projection of higher dimensional data. This projection takes multiple dimensions in shape space into account, which are transformable interactively and allows the user to explore global information about the population. Busking et al. additionally provide a view in object space to visualize the reconstruction error of a shape model.

Besides the global view and to analyze the relation of a single dimension in shape space to the object space, Kirschner et al. [KBW11] show arrow glyphs and Busking et al. [BBP10] provide the shape evolution view. The latter displays multiple projections of reconstructed objects, which only differ in one dimension in shape space. While both techniques are good for estimating properties of dimensions in shape space, they are not suitable in the context of this work, as the main interest is to derive hypotheses from shape similarity and not to understand shape space. They also have the limitation of only dealing with one dimension at a time, which may not have a meaning, e.g. from the output of PCA.

To overcome this limitation and show semantically useful projections of shape space, Hermann et al. [HSK11] compute trait vectors to find the respective best separating axes for a PCA plot, e.g. shown in Figure 3.2. They do this for example to differentiate




Figure 3.1: Shape space scatter plot by Busking et al. [BBP10].

between two subfamilies in a dataset of rodent mandibles. The local impact of a trait vector is visualized by mapping the distance to the mean shape with color on the resulting shape in 3D. This approach could be used in this thesis for visualizing properties like the filling state of the bladder or the rectum, but it requires prior knowledge, which is not available in the dataset. Another approach for quantification and understanding shape variation and covariances in shape data was later described by Hermann et al. [Her+14], where the user can change the shape in specific areas and see other areas respond based on the anatomic shape model. Other shape space navigation and visualization techniques are discussed by Hermann and Klein in [HK15]. Hermann et al. also provide an interactive technique for visualization of large deformations and variability, tackling artifacts and low visual quality introduced by linear interpolation in the tangent space [Her+16].

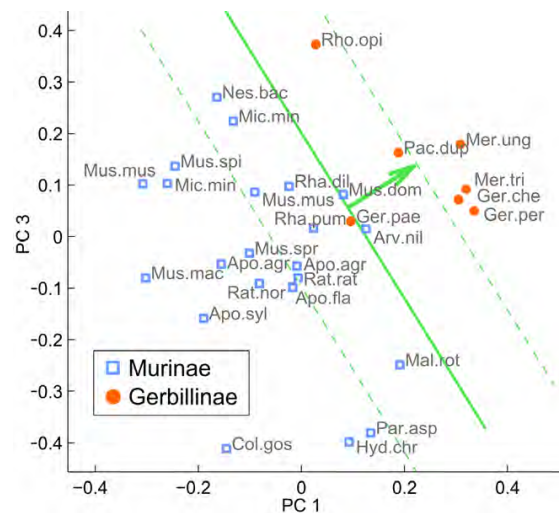


Figure 3.2: Best separating axes for a differentiation based on a specific trait vector. [HSK11]

Klemm et al. [Kle+13] allow for exploration of lumbar spine shapes based on clustering with multiple distance measures. For each cluster, a representative band is chosen for a simplification reason, visible in Figure 3.3. They also extracted a lumbar spine canal centerline in advance to avoid using the whole tetrahedral mesh of the input data.

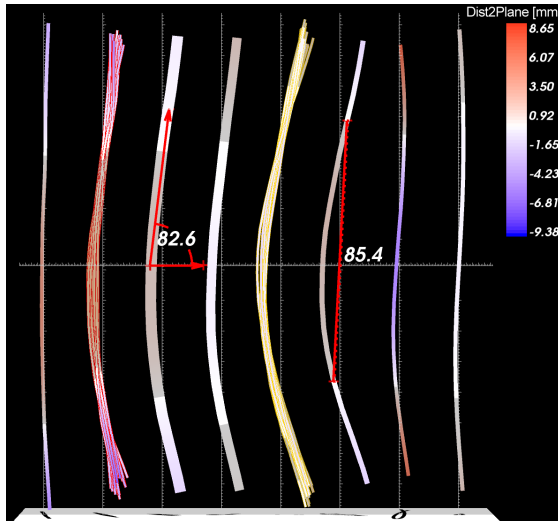


Figure 3.3: Bands representing clusters of lumbar spine canal centerlines used by Klemm et al. [Kle+13].

Torsney-Weir et al. [Tor+11] and Fröhler et al. [FMH16] provide methods to describe the impact of parameters used in segmentation algorithms to the segmentation outcomes, but do not offer comparative visualization of the results. Bergner et al. [Ber+13] propose a tool for the visual exploration of the parameter space by interactive partitioning. In addition to the parameter space analysis, Weissenböck et al. [Wei+16] also analyze the segmentation outcomes as part of their segmentation pipeline. Raidou et al. [Rai+16] proposed a Visual Analytics tool for the exploration of segmentation errors in a cohort of pelvic organs, but they focus on the quality assessment of the segmentation outcomes, without extending to the impact of the organ shape variability.

3.3 Comparative Visualization

Comparative visualization of local shape variation is currently one of the biggest challenges in the research area of *Visual Analytics*, as well as visualizing uncertainty, in the form of variability, of the data [Lan+13]. Uncertainty visualization regarding probability was tackled by Reh et al. [Reh+13]. Even comparison of two shapes provides a lot of visual complexity [Bus+11] and new methods for visualizing and comparing multiple subjects are required.

Kim et al. [KCK17] describe four fundamental types of comparative visualization, i.e. Juxtaposition, Superimposition, Interchangeable and Explicit Encoding, used in the




Figure 3.4: The four fundamental approaches for comparative visualization of 3D data sets described and illustrated by Kim et al. [KCK17]. This simplified illustration shows the case of comparing just two data instances.

literature, as shown in Figure 3.4. They see hybrid approaches, which combine multiple of these or switch between them, as the most promising for future work.

They state that Juxtaposition designs (Figure 3.4a), i.e. side-by-side views, are often used as overview along other techniques for detailed comparison, like Keefe et al. did [Kee+09]. Because the data instances are displayed in different coordinate spaces, a common context is missing by default. This is often provided by interactive techniques like linked camera views or linked highlighting, of which examples are from Sauber et al. [STS06], Schmidt et al. [Sch+14] and Haidacher et al. [HBG11].

Superimposition (Figure 3.4b), i.e. overlay, is the approach of displaying multiple data instances in the same coordinate space, which induces a problem of occlusion. To overcome this, the number of data is often reduced by filtering in advance and visual parameters (e.g. color, opacity) are adjusted. Literature using Superimposition techniques mention the limitation of data number [Bus+11; Ala+12] or show a limited number of volumes by design [Smi+14]. Figure 3.5 shows an example of a Superimposition technique using Weigle’s curvature glyphs [Ala+12; WT05].

In contrast to the two mentioned ones, the Interchangeable approach (Figure 3.4c) displays different instances at different times in the same coordinate space, with the exception of a semi-transparent transition between them. It can be considered as Superimposition over time and hereby solves the occlusion problem, but is often not suitable, as a simultaneous comparison is preferred. Applications range from displaying multiple data instances at the same time using a visibility function, e.g. Phadke et al. [Pha+12] change size or opacity, to showing uncertainty by animation, as proposed by Lundström et al. [Lun+07].

The Explicit Encoding approach (Figure 3.4d) is a solution for visualizing multiple data instances simultaneously in the same coordinate space without occlusion, but does not preserve all data. It derives and displays just one composite from the instances, e.g. just the intersection or the difference. Busking et al. [Bus+11] for example, show distances between surfaces directly. Other approaches show one object with additional information coded in color for local quality assessment. For example, Geurts et al. [Geu+15] show the locally best segmentation algorithm by coloring a region in the respective color. Schmidt et al. [Sch+14] project a heat map according to error variances in the data to



Figure 3.5: Superimposition technique used by Alabi et al. in [Ala+12] based on Weigle’s technique [WT05].

highlight areas of bad reconstruction quality from point-clouds, shown in figure 3.6. Von Landesberger et al. [LBB16] display clusters according to segmentation quality as well as over- and under-segmentation based on color. An approach for comparative visualization of 2D images is provided by Schmidt et al. [SGB13], in which the user is able to explore the space of differences.

A suitable glyph-based approach for comparison of different stent configurations in cerebral aneurysms was found by van Pelt et al. [Pel+14], using different types of glyphs representing blood flow and safety indicators, also based on different zoom levels.

Image-centric cohort studies are a recent development in epidemiology, using MRI and ultrasound imaging in its statistical analysis. Due to the growing complexity of the data, tools of Visual Analytics are proposed to support exploration of the cohort data [Pre+15]. In prior to hypothesis generation, semantics have to be derived from the image data, e.g. with segmentation, and analyzed. Different subgroups are created by brushing and variable selection, which are typical tools of Visual Analytics. Klemm et al. [Kle+14] presented a framework for cohort exploration and visualization using a lumbar spine data set. They visualize different subgroups by Juxtaposition with linked camera views and encode the distance from the group to the global mean shape with color. A dedicated tool for visual subpopulation discovery and validation using subspace clustering was proposed by Alemzadeh et al. [Ale+17]. Their approach also uses coordinated views with




Figure 3.6: Explicit Encoding approach of Schmidt et al. [Sch+14] projecting a heat map of error variances.

an overview of all subpopulations in Juxtaposition.

Visualization Design

In this chapter, all parts of the application are presented and linked to the tasks described in Chapter 1, which were accomplished. Different visualization approaches and alternatives are discussed in detail.

4.1 Application Overview

The application fulfills the two tasks mentioned in Section 1.2.1, dealing with shape variability and comparative visualization. In order to quantify and analyze shape variability, the required steps are shown in Figure 4.1. As the triangle mesh of the input data includes no semantics about the shape by itself, some features have to be extracted to be able to compare different shapes in the first step. This is done by using a shape descriptor. To be able to explore the distribution of the organs and interfaces regarding the features used, the resulting shape vectors are compared using a metric, as well as prepared for visualization, in the second step. The shape vectors are presented to the user in the third one, utilizing different two-dimensional visualization techniques. For exploration and analysis, the user is able to select shapes interactively within linked views and have the resulting selection comparatively visualized, which provides a link to the second task. Regardless of dealing with one or more organs, the steps are the same, but metrics, visualization, and other techniques are likely to be only suitable for one case, as it will be discussed in Section 4.3.4.

Before the actual implementation of Task 2, we need to answer: *What is actually visualized and how?* To reach the goal of effectively showing what the multiple organs and interfaces of multiple patients look like and how they vary, there has to be some kind of shape at first, which represents a cluster of segmentations. In addition, some indication of what other shapes in this cluster may look like is needed. The visualization also has to show variability, which means that some measures have to be defined to be displayed. Based on these requirements, the steps of this task are created and shown in Figure 4.2. The

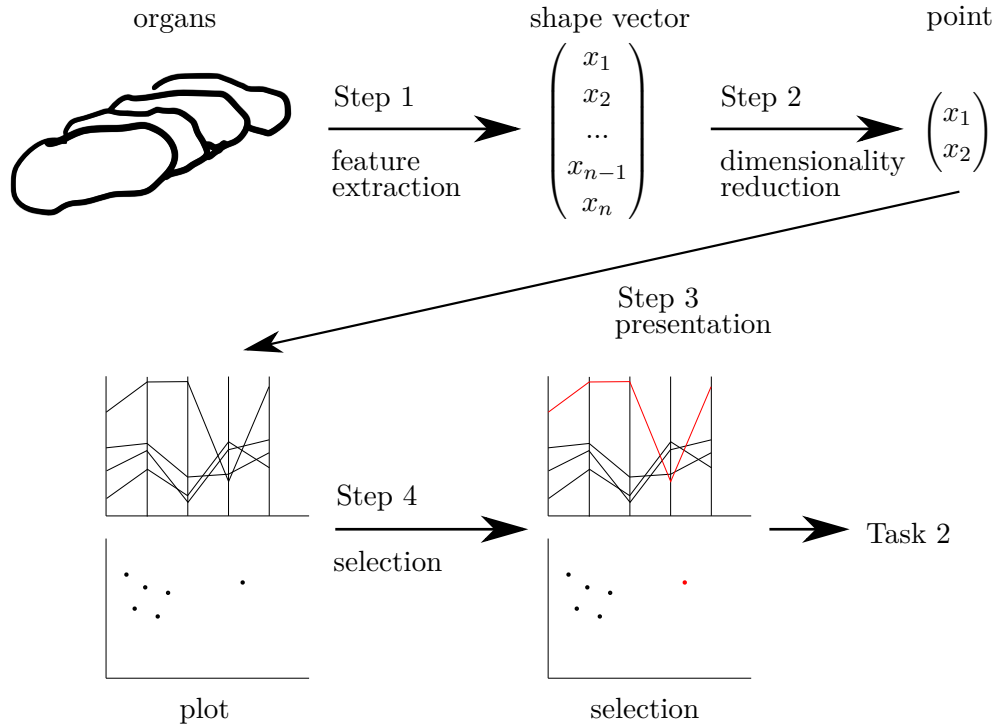


Figure 4.1: Schematic overview of Task 1




Figure 4.2: Schematic overview of task 2

first step is to cluster all shapes based on which organ or interface they are of, find a representative for each cluster, and show a rough overview of all clusters. The second step is a more in-depth view of a cluster, coloring the representative based on if its extent is lower or higher than the others at a specific point. Variance is encoded with glyphs throughout the surface.

4.2 User Interface Description

The entire application is web-based. The main area of the User Interface is divided into three sections, pictured in Figure 4.3. The first area on the top left shows all shapes of the cohort data in a global distribution. On the bottom left, there is a section, which

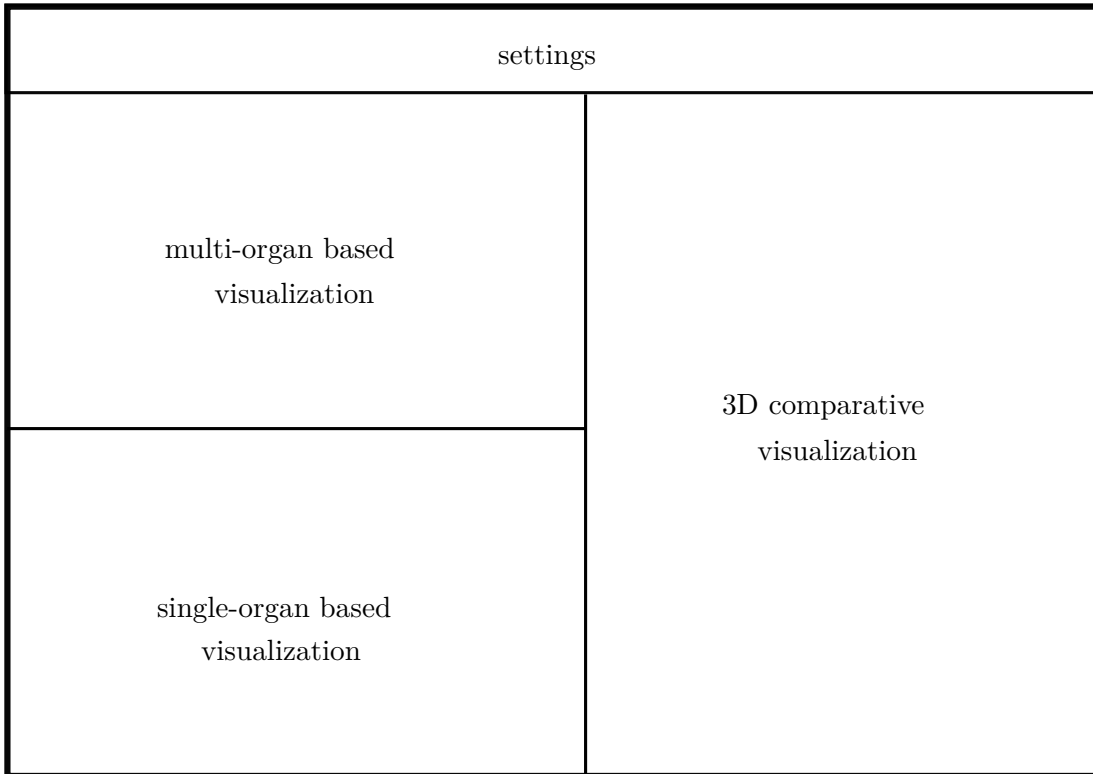


Figure 4.3: The three visualization sections of the UI of our proposed tool

shows the distribution and shape vectors per organ. Together they depict the approach to solve the first task. A 3D-view on the right gives a representation of the current selection of organs and visualizes them comparatively, which is the outcome of the second task.

Linking the different views and providing a common context was of high priority, thus brushing in a view highlights the selection in all other views accordingly and updates the 3D-view. Throughout all views, the colors for each type of organ are the same for intuitive orientation within the view. The color palette is taken from Color Brewer [Bre17].

4.3 Task 1: Shape Variability

As described in Section 4.1, this section deals with finding a suitable shape descriptor and a metric for comparison. Dimensionality reduction techniques for dealing with one and multiple organs as well as visualization approaches are discussed.

4.3.1 Shape Descriptors

The shapes of the input data in this work are defined by a triangle mesh in \mathbb{R}^3 with point-to-point correspondences. A descriptor is suitable if it has invariance regarding

translation and rotation, but not scale, as the size of the organs is also expected to influence segmentation outcome. A fast yet precise method was using a spherical harmonics based descriptor to gain rotation invariance. This is generally proven to be more accurate than performing PCA-based alignment in advance, as described by Kazhdan et al. [KFR03].

Other shape descriptors taken into account were simple global features such as volume or area, and shape distributions, but as they discriminate poorly between mostly similar shapes [TV04], they were rejected. Local feature and graph based descriptors were too inefficient to use and would have required preprocessing. An alternative regarding discriminative performance to descriptors using spherical harmonics are 3D Zernike Descriptors [NK03], but due to simplicity and the direct applicability to the mesh, the former technique is used in this thesis.

Spherical harmonics are orthogonal functions, which are defined on a sphere, defined by a frequency, or degree, l and an order m . Using a shape descriptor based on spherical harmonics, a spherical function has to be defined in advance. The simplest approach is to use the distance to the centroid, which initially arose questions, if it is suitable for non-round organs, like the rectum. This is further discussed in Chapter 6. This spherical function, defined by the angular coordinates θ and ϕ , can then be decomposed into its complex harmonics $Y_l^m(\theta, \phi)$ as described in Equation 4.1.

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{lm} Y_l^m(\theta, \phi) \quad (4.1)$$

By restricting this decomposition to a specific frequency l , a rotation invariant representation can be obtained. The same properties are also valid when using real spherical harmonics Y_{lm} (Equations 4.2 and 4.3) instead of the complex ones [Mor+05], where l is again the frequency, m the order and $P_l^{|m|}$ the associated Legendre polynomials.

$$Y_{lm}(\theta, \phi) = \sqrt{(2l+1) \frac{(l-|m|)!}{(l+|m|)!}} P_l^{|m|}(\cos \theta) N_m(\phi) \quad (4.2)$$

where

$$N_m(\phi) = \begin{cases} \sqrt{2} \cos(m\phi) & (m > 0) \\ 1 & (m = 0) \\ \sqrt{2} \sin(|m|\phi) & (m < 0) \end{cases} \quad (4.3)$$

The descriptor itself consists of the L_2 -norm of the coefficients a_{lm} for each frequency, which does not change when rotating the spherical function. Mathematical proof is provided by Kazhdan et al. [KFR03]. Figure 4.4 shows the computation sequence of the coefficients.

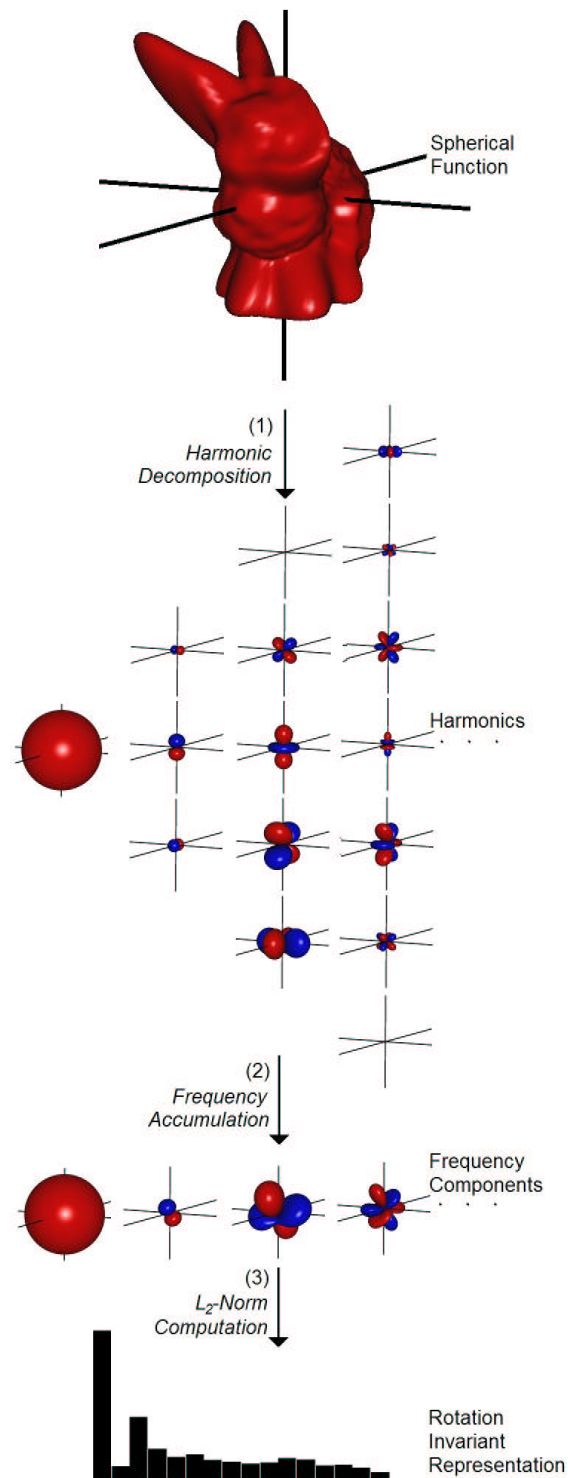


Figure 4.4: Sequence of computing a rotation invariant descriptor based on spherical harmonics [KFR03]. The decomposition into the harmonics and the frequency components are shown for $l = 0$ on the left to $l = 4$ on the right.

Performance

As computing the Legendre polynomials is computationally expensive at higher degrees, this section covers the performance of different band limitations with respect to computation time and differentiation between shapes. Using the L_2 -distance as a measure, the latter is improving when increasing the maximum frequency up to $l = 7$. It then starts to decrease and some similar shapes are having high dissimilarity in the high frequencies, which can be due to the sampling of the organ shapes. Figure 4.5 shows the computation time for maximum frequencies 0-13, using the data-set of 17 patients, each consisting of four organs and two interfaces, which results in 102 shape vectors.

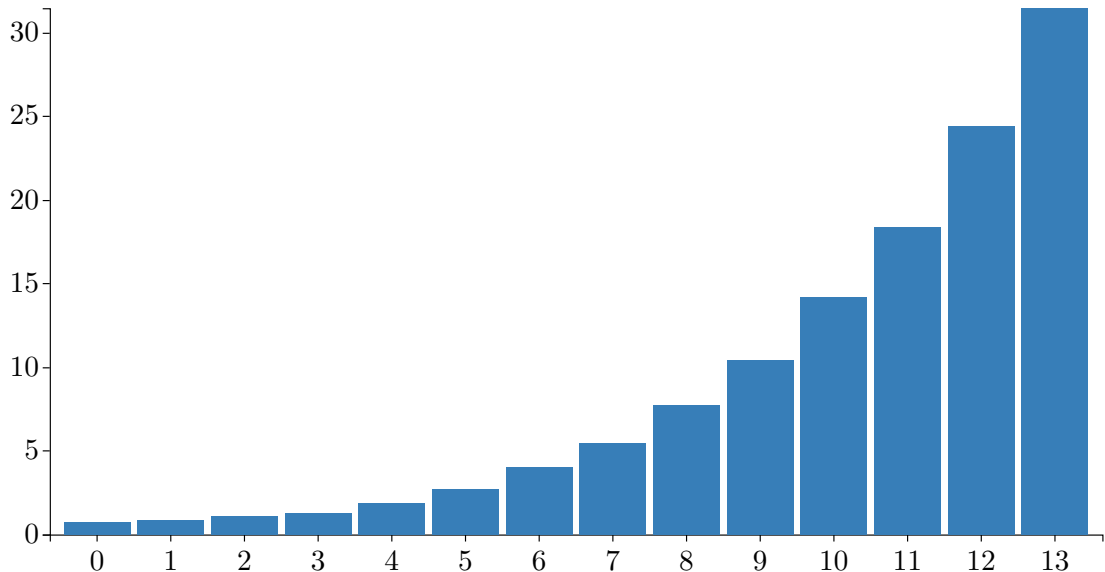


Figure 4.5: Computation time in seconds for the spherical harmonics based descriptor with different maximum frequencies l ranging from 0-13. The performance was measured within Google Chrome version 63 on a system with an Intel Core i7-2600k CPU and a NVIDIA GTX 970 GPU.

Computing the Spherical Coefficients

In theory the spherical coefficients a_{lm} can be obtained by multiplying the spherical function with the complex conjugate of a spherical harmonic and integrating over the solid angle Ω , shown in Equation 4.4.

$$a_{lm} = \int_{\Omega} f(\theta, \varphi) Y_l^{m*}(\theta, \varphi) d\Omega = \int_0^{2\pi} d\varphi \int_0^{\pi} d\theta \sin \theta f(\theta, \varphi) Y_l^{m*}(\theta, \varphi). \quad (4.4)$$

Using the average over the discrete data points as a solution of this integral provides sufficient results in most cases, because the points are, for round organs, adequately sampled on the unit sphere. Results of this computation distinguish well between organs

as well as between features like size and roundness. If the points are unevenly distributed, as in case of the rectum, some errors are not detected, because the dense areas outweigh the sparse ones. In this case, a solution using the method of least squares provides better results in finding outliers, but is not as good in differentiating between organs. In Figure 4.6, a segmentation error is not detected by the average and located near normal shapes, but is clearly visible when using the method of least squares, where it is displayed far away from all the other shapes.

4.3.2 Similarity Measure

To compare different shapes, a metric has to be selected. Some of them are described in Section 3.1. Kazhdan et al. [KFR03] propose to use the most common L_2 -difference with the spherical harmonics, as one property thereof is, that the L_2 -difference of the harmonic representation is a lower bound of the L_2 -difference over all orientations of the spherical function. Both dimensionality reduction techniques described in the next section are based on this metric. Using the L_1 -difference would result in more robustness regarding outliers, but this is not actually a goal of this work.

4.3.3 Dimensionality Reduction

The result of the previous step is a high dimensional shape description vector, which cannot be visualized directly and has to be converted into lower dimensional space. When displaying the distribution of multidimensional data, the problem arises, how it is projected into two- or three-dimensional space to get a meaningful representation. This section deals with finding a suitable solution for reducing multiple classes at once and separate. We investigated two commonly employed dimensionality reduction techniques, because they are proven to produce good results in general for the respective tasks and are fast in computation time. We choose also PCA as a representative of linear dimensionality reduction techniques [MPH07] and t-Distributed Stochastic Neighbor Embedding (t-SNE) as a representative of non-linear techniques [MH08].

Principal Component Analysis

PCA is a technique to represent data with orthogonal variables, i.e. principal components. The application implements PCA based on Singular Value Decomposition (SVD) [Sh14]. The input data \mathbf{X} is centered and standardized in advance. \mathbf{Y} is the new representation in PCA coordinates, transformed by a matrix \mathbf{P} , which consists of the eigenvectors of the covariance matrix $\mathbf{C} = \mathbf{X}^\top \mathbf{X}$.


$$\mathbf{Y} = \mathbf{XP} \tag{4.5}$$

For computation of the eigenvectors, the symmetric matrix \mathbf{C} has to be diagonalized.


$$\mathbf{C} = \mathbf{PDP}^\top \tag{4.6}$$



(a) Segmentation error of a rectum, where only the bottom part seems to be segmented correctly



(b) Average



(c) Method of least squares

Figure 4.6: Comparison of the distributions using the average and the method of least squares to compute the spherical coefficients. Dimensionality reduction was done using PCA with projection on the first two principal components. The segmentation error is highlighted to show the ability to detect outliers.

Now \mathbf{P} contains the eigenvectors and the diagonal matrix \mathbf{D} contains the eigenvalues. The following Equation 4.7 shows the matrix form of SVD.

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \quad (4.7)$$

Constructing the covariance matrix from this, we get

$$\mathbf{X}^\top \mathbf{X} = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^\top \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top \quad (4.8)$$

because

$$\mathbf{U}^\top \mathbf{U} = \mathbf{I} \quad (4.9)$$

due to it being an orthogonal matrix. This shows that the right singular vectors, which form the orthogonal matrix \mathbf{V} , are the same as the eigenvectors, i.e. principal axes, \mathbf{P} and the eigenvalues \mathbf{D} are related to the singular values $\mathbf{\Sigma}$. Therefore the PCA decomposition can be written as

$$\mathbf{Y} = \mathbf{X}\mathbf{V} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \mathbf{V} = \mathbf{U}\mathbf{\Sigma} \quad (4.10)$$

Because the principal components are ordered by variance, the dimensionality reduction, which retains the maximum possible variance, can be obtained by only using the first L principal axes in the transformation, which results in

$$\mathbf{Y}_L = \mathbf{X}\mathbf{V}_L \quad (4.11)$$

PCA is only applicable to one organ at a time, as it does not discriminate well between classes. For the task of reducing multiple classes in the dimensions simultaneously, a dedicated technique such as t-SNE is more suitable.

t-Distributed Stochastic Neighbor Embedding

t-SNE is a dimensionality reduction technique described by van der Maaten [MH08], which iteratively creates a matrix of pairwise similarities and maps similar shapes with small distances and dissimilar shapes with large distances in two- or three-dimensional space. The actual position has, in contrast to PCA, no meaning and can differ at different runs. It outperforms other commonly used dimensionality reduction techniques on the MNIST dataset, the Olivetti faces dataset and others described in the supplementary material of van der Maaten and Hinton [MH08]. Figure 4.7 shows a comparison of PCA and t-SNE on the MNIST dataset with t-SNE producing much better results [Ola14]. The benefit in contrast to PCA is the non-linearity, so clusters can be separated much better, but the absolute position in the reduced space is abstract, so no properties can be deduced from it.

```

3 6 8 1 7 9 6 6 9 1
6 7 5 7 8 6 3 4 8 5
2 1 7 9 7 1 2 8 4 5
4 8 1 9 0 1 8 8 9 4
7 6 1 8 6 4 1 5 6 0
7 5 9 2 6 5 8 1 9 7
2 2 2 2 2 3 4 4 8 0
0 2 3 8 0 7 3 8 5 7
0 1 4 6 4 6 0 2 4 3
7 1 2 8 9 6 9 8 6 1
    
```

(a) Size-normalized examples from the MNIST database. [Lec+98]




Figure 4.7: Comparison of PCA and t-SNE performing on the MNIST dataset. Different colors represent different digits.

4.3.4 Visualization Strategy

This section deals with the visualization step of Task 1. While both mentioned dimensionality reduction techniques support reducing to three dimensions, the usage of two for visualization was preferred. As the last step is to enable selection, using three dimensional visualization could have caused problems with interaction.

Visualization of Shape Vectors with Parallel Coordinate Plots

An intuitive way of displaying all shape vectors at once is the Parallel Coordinate Plot (PCP) and provides insight into the actual values. In this context, it enables the user to see which frequencies a shape is composed of and to compare them to frequencies of other segmentations of this specific organ. The PCP in general shows axes for each dimension of a vector space. Its elements are represented as continuous lines, crossing the axes at the value in the particular dimension. What this means for this thesis is, that all shape vectors of an organ or interface class are displayed as lines for comparison. Selection is implemented by brushing on individual axes. Figure 4.8 shows a PCP for prostate shape vectors, where a polyline represents a shape and the axes represent the frequency.




Figure 4.8: PCP of all prostate shape vectors. Small values, e.g. of frequency one, are abbreviated with SI prefixes.

Visualization of Organ Distributions with PCA Plots

Not only the values of shape vectors are of interest in this step of the first task, but also the position within the distribution, especially to find outliers and to see clusters, which are not easy to see in high dimensional space. It actually allows the user to quantify and

compare distances with each other and relate position to shape features. Showing the distribution of a single organ is realized by performing PCA, projecting the data to the first two modes, and visualizing the resulting vectors in a two-dimensional scatter plot. This shows the highest variance within the data when reducing the dimensionality, an example is given in Figure 4.9.

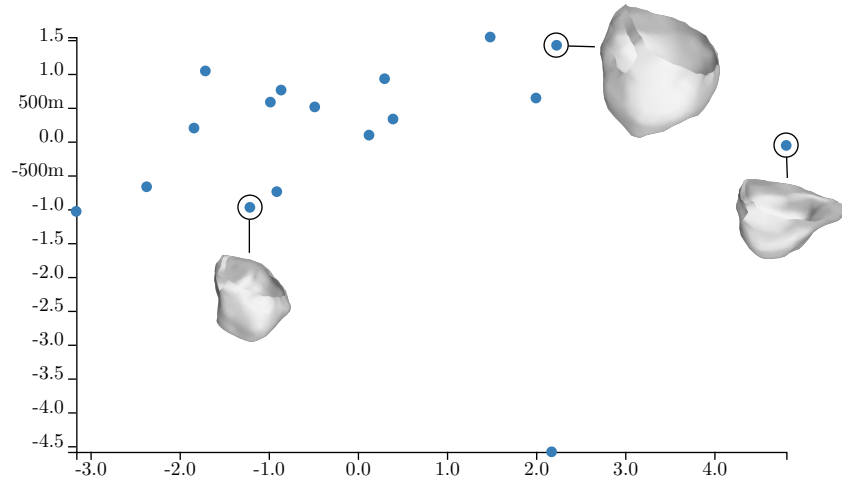


Figure 4.9: PCA plot of the prostate shape vectors shown in Figure 4.8, projected on the first two principal components.

Visualization of the Whole Cohort with a t-SNE Plot

The main challenge in this task was to find a way to deal with multiple organs at once. As discussed in Section 4.3.3, t-SNE produces good results when performing on multiple classes of high-dimensional data and clustering them, especially as it was designed to be used for visualization. The resulting vectors are again displayed in a two-dimensional scatter plot, shown in Figure 4.10, and colored based on which organ or interface class they belong to.

So far, Task 1 is achieved by computing shape vectors of the input data using a shape descriptor based on spherical harmonics. The L_2 -difference is used as a similarity measure in this shape space. The high-dimensional shape vectors are then reduced in dimensionality to be used in the visualization strategies, using PCA on shapes per organ type and t-SNE to process all shapes at once. The former is expected to reveal mis-segmented organs well dealing with one class and the latter is able to deal with multiple classes while still preserving local differences. To visualize the whole cohort, the two-dimensional output of t-SNE is displayed in a plot. For each organ, the shape vectors are directly visualized in a PCP as well as projected on the first two principal components in a plot.

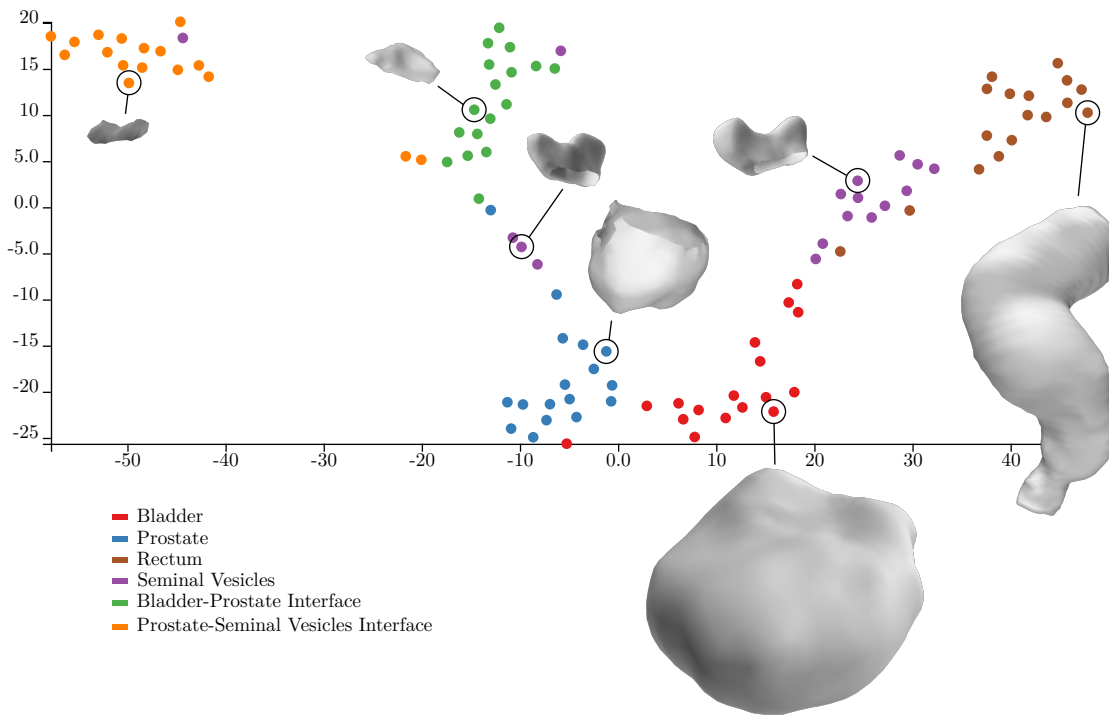


Figure 4.10: t-SNE plot of the whole data-set. The colors show the organ or interface class the shapes belong to. The relative size of the displayed organ shapes to each other is preserved to see the impact of shape and size on the position in the plot.

4.4 Task 2: Comparative Visualization

4.4.1 Displaying Shapes of Multiple Organs

Reviewing the common comparative visualization strategies of Section 3.3, Juxtaposition is the only one suited for an overview of the organs. Because of the number of organs and the different shapes, Encodings and Superimposition would have too much visual information to be useful and an Interchangeable approach is by definition not usable to display multiple instances at a time. Superimposition would be very cluttering as the data have only point-to-point correspondence and no registration. Interchangeable means animations, which have been proven to cause fatigue and also would require a lot of time for a high number of patients. Explicit Encoding would require to find a design that the user would be able to understand, decode, and link back to the initial organs, which can be difficult and probably confusing. In order to compare these approaches to the visualization design described subsequently, Juxtaposition, Superimposition, and encoding of variance in color on the model were implemented in the application and further discussed in Chapter 5.

The organs are displayed separately, but to preserve anatomical context, the positions relative to each other are not changed (see Figure 4.11). They are just displayed with

more spacing, hence the term *exploded view* is used subsequently. To further enhance the perception of connected organs, they are visually connected by a line by the following strategy: if two organs have a direct connection in the mesh through an interface, a line is drawn, shown in Figure 4.12. If the respective interface is also selected, the connection is shown between the organs and the interface. Disconnected components with respect to the mesh are connected to a white dot, like the ones in Figure 4.13. The colors of the circles surrounding each organ is based on the colors used in other parts of the application to visually link the different views.

This exploded view is a simple visual representation of what the current selection looks like, with the median shape of each organ being shown. To give an example, Figure 4.13 shows a different selection of rectums, bladders and seminal vesicles, than the one in Figure 4.11. To compare local differences for each organ, a detail-on-demand view can be obtained by clicking on the desired organ, which shows the same median shape with two additional attributes described in the following section.

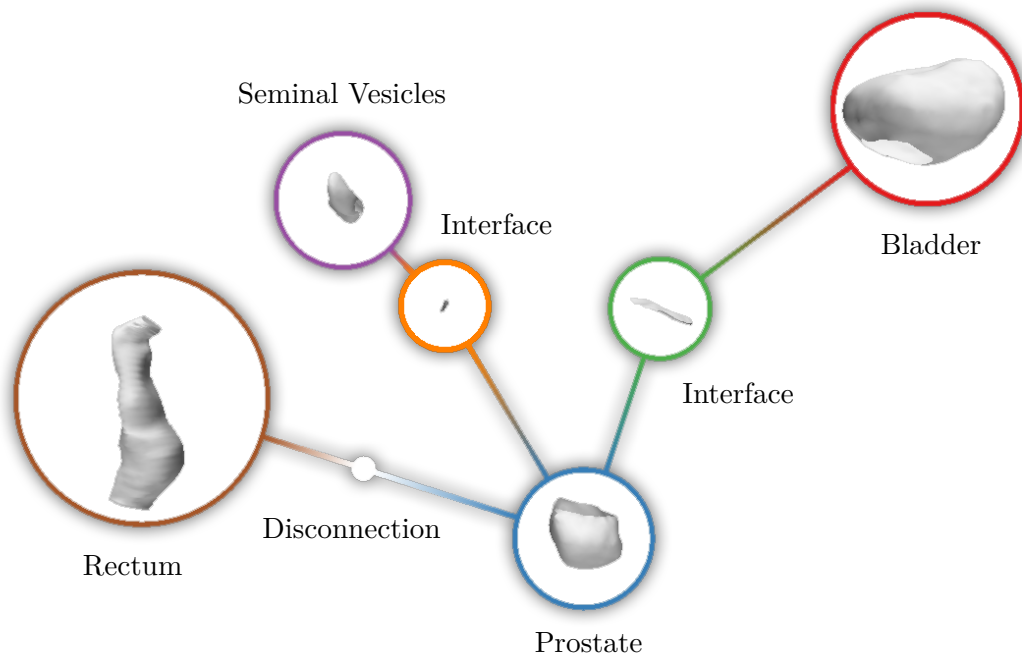


Figure 4.11: Exploded view approach for displaying multiple organs and their interfaces, preserving the connections in the mesh through the interfaces.

4.4.2 Visualizing the Median Shape with Variances

Visualizing the median shape with variances provides insight into a cohort of shapes by showing the global median shape, and therefore an actually existing representative instance. In addition, the per-triangle variances are visualized to show, how locally

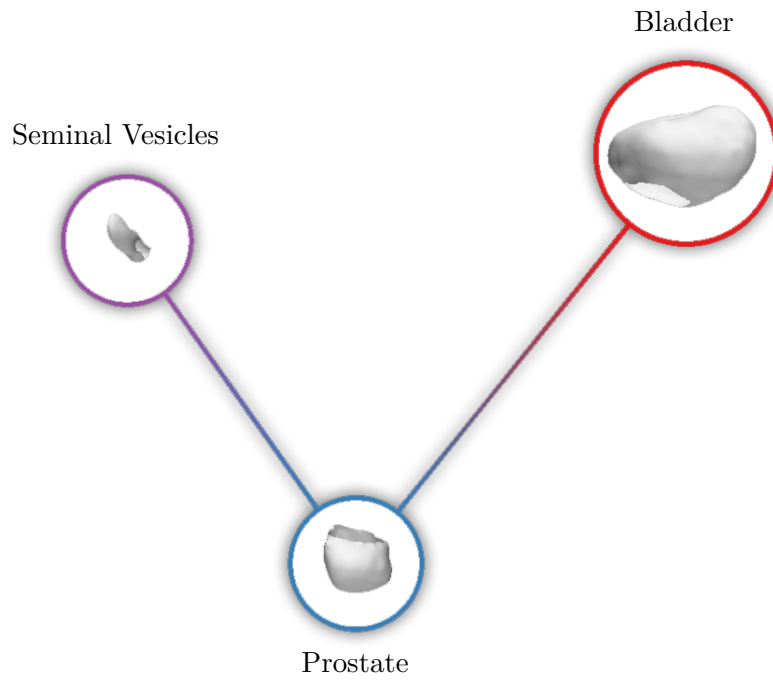


Figure 4.12: Exploded view: connections between organs without showing interfaces.

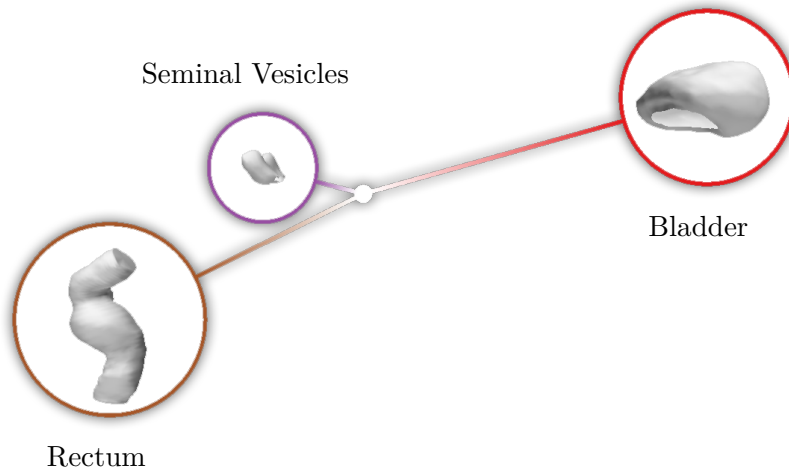


Figure 4.13: Exploded view: lines between disconnected organs.

accurate the displayed mesh is and where areas of high shape variation are. The goal is to understand, how shapes of a cohort actually look like, by pointing out areas of similarity and spots of high variation. This provides, for example, a basis for comparing different subsets of organ shapes.

In contrast to the previous section, the rotation of the organs is now relevant, as we compare local differences and use properties, which are not rotation invariant. Therefore, in order to calculate or show anything, the organs have to be rotated to match each other in the best possible way. A common technique to do this is minimizing the weighted sum of squared deviations, which is solved by the Kabsch algorithm [Kab76; Kab78]. This method is only possible because of the point-to-point correspondence and can be easily implemented using SVD. Both shapes need to be free of translation first, then the optimal rotation is obtained based on SVD of the cross-covariance matrix of the point sets.

The variance in the position of a specific triangle is then calculated based on the L_2 distance to the mean, with \mathbf{x}_i being the centers of a triangle over all patients and $\bar{\mathbf{x}}$ being the mean of all \mathbf{x}_i :

$$\sigma_{\mathbf{x}}^2 = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 \quad (4.12)$$

An addition for providing more information about the distribution and what the selected shapes may look like, is the visualization of outliers, which have to be detected in advance. In this case, two approaches have to be distinguished. The first one is to take a global distance measure like the one used in the t-SNE plot in Section 4.3.4 and the second one is to take local features into account. In our case, the latter is preferred, as the variances are also calculated locally and a global measure may favor other features, which are of no interest in this view. An example would be an area of high frequencies affecting the shape vector, but which is close to the median shape and therefore represents no outlier regarding the general shape. Other features like size are recognized by both approaches.

The described technique is best suitable for a high number of mostly similar shapes. The major drawback is, that the median shape can also be a local outlier and therefore is not a good representative in this region. To take this into account, a solution was found, which shows, how far the global median shape is from the local median at a specific area: encoding the relative position to the other shapes in color. Although this is a limitation, the median is still preferred over the mean shape, because it is, as described above, an actually existing representative and the mean shape could be any shape, maybe even one that would not be anatomically possible. Because the variance is based on the mean, we also chose to add the possibility to calculate it based on the median instead. Another limitation is, that the distribution of shapes is not clearly visible, as well as that local details of only one shape are displayed.

To visualize the local position relative to the other shapes in addition to the local variance, a two-dimensional color space was introduced. The measure for the local position is the

difference in the numbers of shapes, which are inside and outside of the median shape. A possible solution was using a bivariate complementary color scale [Rhe00], but this can interfere with the brightness of shading in the 3D scene. Therefore another way of introducing the second dimension was considered. The local position relative to the other shapes is chosen to be encoded in color and the variance is encoded using glyphs. If the median shape is in the middle of all other shapes, it is colored gray. Two complementary colors are used to show, if it is rather in- or outside relative to all other shapes. The further apart it is, the more saturated the color is. A schematic example of which colors are used to encode which relative position is shown in Figure 4.14.

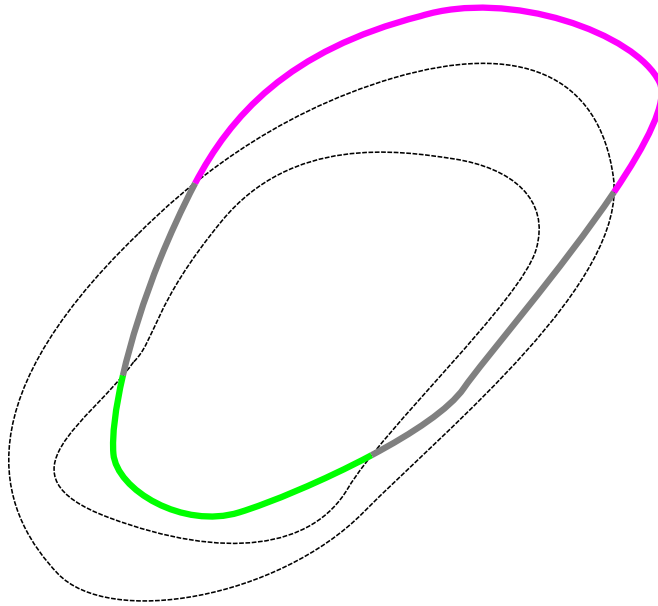


Figure 4.14: Illustration of encoding the relative position with complementary colors. The further inside the shape is relative to the others, the more it changes to the color green, the further outside, the more it changes to magenta. If it is in the center of all shapes, it is gray. The dashed black shapes are not visible in the final visualization and only show the relative position in this figure.

Glyphs are a way to add one or more data attributes to a visual representation, which have to meet certain design criteria to be comprehensive [Bor+13]. The first idea was to use line-based glyphs, which are perpendicular to the surface, and encode the variance with length. Clearly visible in Figure 4.15, there was a perceptive problem with lines parallel to the viewer direction. To overcome this, points on the surface are used. The size was no option for encoding, because of the different triangle sizes in the data. Due to scaling, encoding the size may result in too small points over the whole shape, if the highest variance is in a small triangle. Suitable visual variables are color value, saturation, fuzziness and transparency, with the latter being used in the final approach shown in Figure 4.16.




Figure 4.15: Line-based glyphs encoding the variance of the shapes. The colors show the relative position to the other shapes and the semi-transparent shape is the largest in the cohort.




Figure 4.16: Point-based glyphs encoding the variance of the shapes. The colors show the relative position to the other shapes and the semi-transparent shape is the largest in the cohort.

Implementation

This chapter deals with the implementation of the visualization design described in the previous chapter. The software architecture, used technologies as well as other implemented parts are discussed.

5.1 Application Overview

The web-based application is implemented in Javascript, using the libraries D3.js [Bos17] for the 2D visualizations and three.js [Cab17], which uses WebGL, for the 3D comparative visualization. It is based on modular components, which interact through an event handler, shown in Figure 5.1. After the selection of data, it is stored in a data-container. Starting from here, the organs are extracted by an organ-selector for further processing. The shape descriptor then computes a shape vector from the organ data, which is used in the 2D visualizations, both single- and multi-organ-based. A so-called plot-renderer manages the layout of the plots and delegates the visualization of the shape vector to D3-renderers. The 3D visualization, namely a three-renderer, takes the organ mesh as input. An example of the communication between these components is provided as follows:

In each 2D visualization component, the user has the ability to make a selection, which triggers an event a selection-handler in the data-container is subscribed to. This handler computes an overall selection by taking all current selections into account and propagates an event, that the selection has changed, to which all visualization components are subscribed. The current implementation uses the union for combining selections over components, because they can contain multiple organs, where an intersection of per-organ selections would be empty. We considered using the intersection for combining selections from the PCP and PCA plot per organ or for combining the selection from the t-SNE plot with the combination of all others. We chose the union, as there is currently only one selection possible in the scatterplots and this would restrict the user too much when

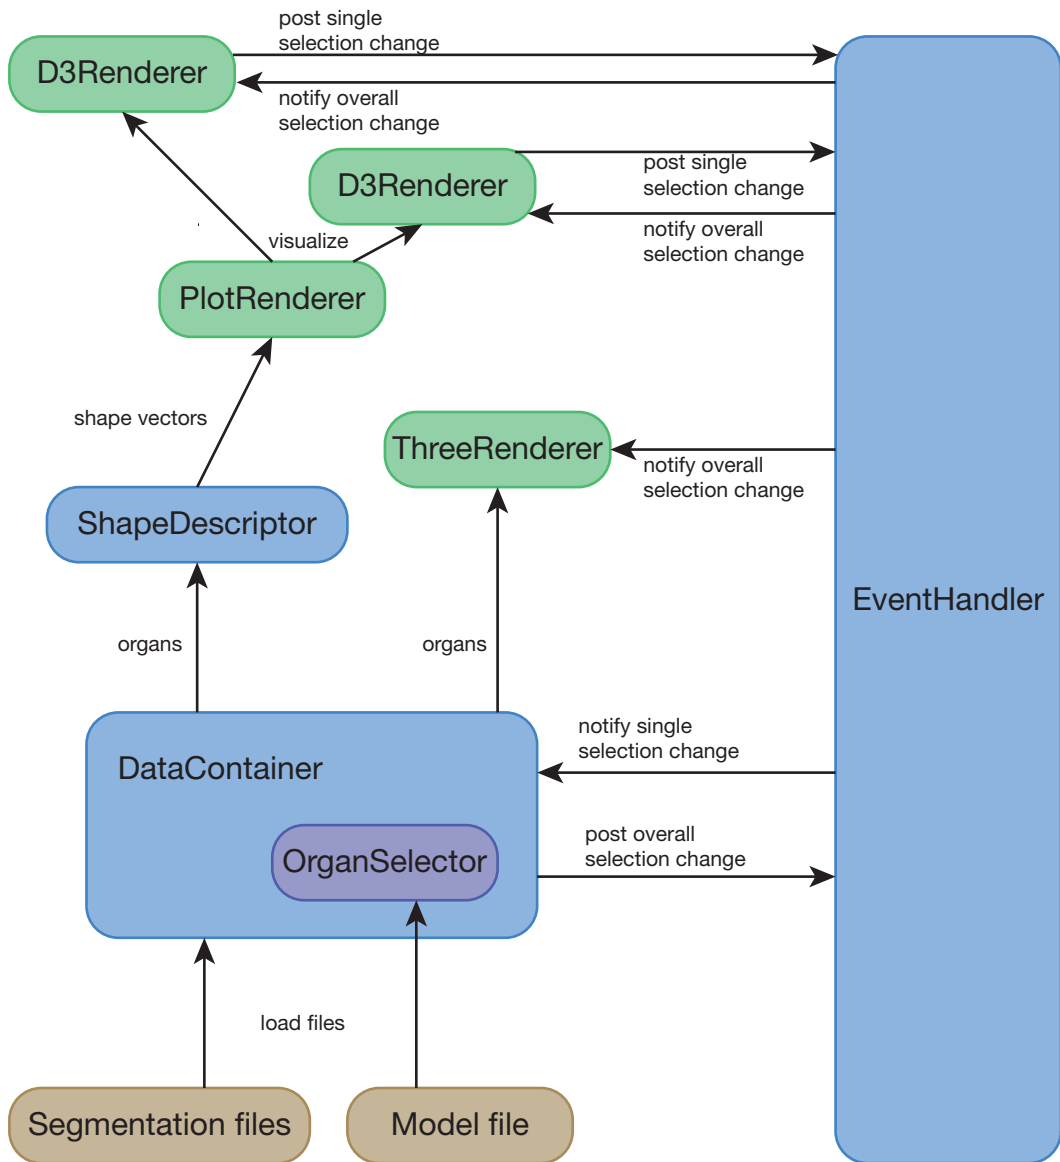


Figure 5.1: Diagram of the main components of the application and their interactions with each other. Green components implement visualization techniques and interact with the Document Object Model (DOM). Blue components are responsible of the application logic and communication.

using the intersection. Selection in the PCP is the only component, where multiple selections are possible and combined by intersection. The overall selection of all axes is then used in the selection-handler of the application. The current selection approach already allows the user to be able to select most interesting subsets of shapes, so custom combination rules, which would allow for every possible selection, were not implemented.

As cohort data can be complex and large, the application follows Schneiderman’s mantra “*overview first, zoom and filter, then details on demand*” [Shn96]. The first view, the user is presented to, is a complete overview of the distribution of all shapes as well as the relative positions of organs to each other. Selections can be made in each 2D view, both in the whole cohort and per organ. The selected shapes can then be examined in detail in the 3D view, still showing the selected organs relative to each other at first, with the ability to inspect local features on demand.

Decoupling and abstraction of components allows for easy extensibility and interchangeability, e.g. with a different shape descriptor implementation or new type of visualization. Although Javascript is not completely object-oriented, the implementation is based on some principles like inheritance and polymorphism, as well as some object-oriented design patterns, e.g. a factory for creating visualization components based on the user settings.

5.2 Task 1: Shape Descriptor

The shape descriptor interface of the application only has one function with the shape as parameter, namely `getShapeVector(shape)`. The implementation, which is based on real spherical harmonics and described in Section 4.3.1, translates the shape at first to have its centroid at the origin of the coordinate system. The spherical coefficients are then computed by a library from Archontis Politis [Pol16] based on `numeric.js` [Loi12] and then stored in separate vectors for each frequency, i.e. the vector of frequency 0 has 1 coefficient and the vector of frequency l has $2l + 1$ coefficients. The L_2 -norm of each of these vectors is then stored in the resulting shape vector to have a rotation invariant representation. The values basically show how much of a specific frequency the spherical function consists of.

To give an example, Figure 5.2 shows the shape vector representations of a sphere, a cube, and a tetrahedron, with the sphere only having one frequency and the cube having values for every even l , because the corresponding harmonics are, like the cube, centrally symmetric, which does not apply to harmonics, where l is uneven. The coefficients of $l = 2$ are also zero, again because of the symmetry properties of a cube. It is also visible, that the coefficients for $l = 3$ are high in the tetrahedron and the cube consists mainly of harmonics of the fourth degree, which shows how values for different frequency components represent specific geometrical shapes. When changing the size of the shapes, the values scale accordingly, but they do not change when the shape is rotated, due to the rotation invariance. For further understanding of the symmetry properties and to see, why the mentioned geometrical shapes consist of specific spherical harmonics, Figure 5.3 is provided, showing the spherical harmonics up to frequency $l = 4$.

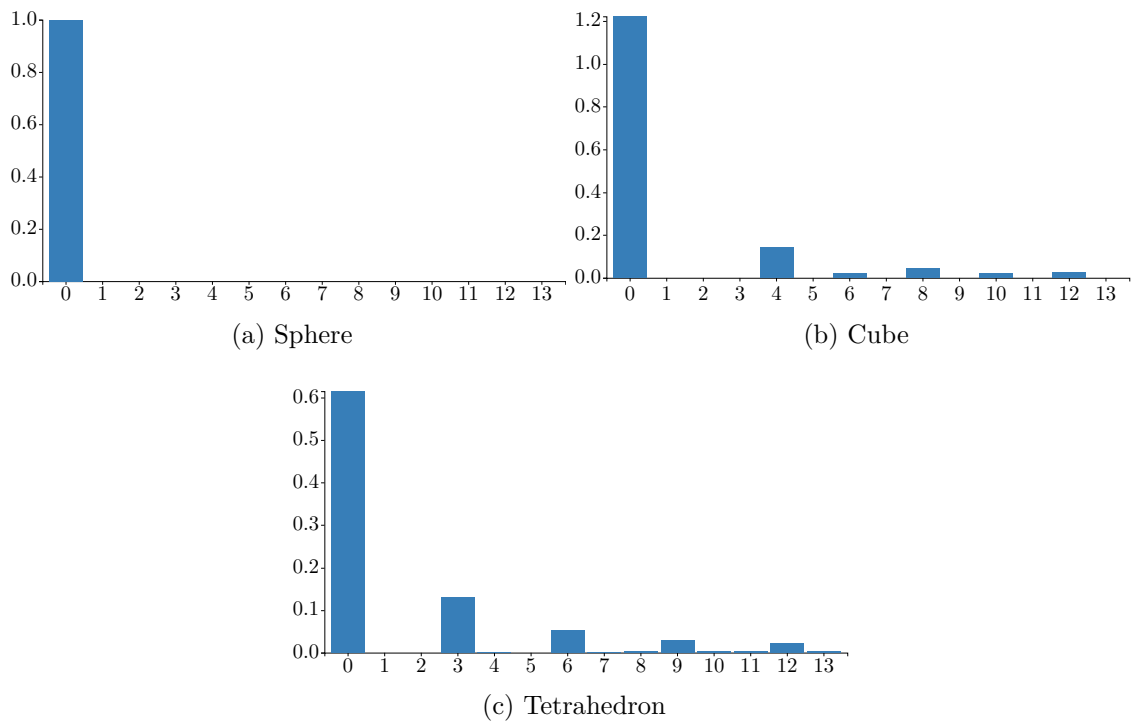


Figure 5.2: Values of the shape vectors of a sphere, a cube and a tetrahedron. It is easily recognizable how different frequencies describe different geometries and symmetry properties.

5.3 Task 1: Dimensionality Reduction and 2D Visualization

The designed 2D visualizations are based on D3.js. Each of these visualization components, despite showing very different plots, is abstracted to the same steps to allow for interchangeability:

1. Plot initialization: the visual elements of the component are initialized by first setting the correct `<div>` element, where they should be placed and then calling the dedicated function.
2. Data initialization: set the raw data, i.e. shape vectors, and perform dimensionality reduction or any other data conversion
3. Drawing: determine colors and visualize the resulting data as points or lines, for example. This step also implements the selection of elements.

In case of t-SNE, the raw data is not processed at once, since the technique is based on approximation and requires some steps to be repeated often, which would take some




Figure 5.3: The real spherical harmonics up to frequency four, green is showing regions where they are positive and red where they are negative. The extent indicates the value in this direction. These values are also encoded in color on the sphere. [Dif17]

time. To still be responsive and show progress, the current positions of the shapes in the distribution are continuously refreshed by performing the drawing step every few runs. The computation is done using the library tSNEJS [Kar16]. PCA was implemented based on numeric.js.

5.4 Task 2: Comparative Visualization

An abstract `ThreeRenderer` provides basic functionality of visualization approach, all done in `Three.js`, and implements a pipeline, in which organs are visualized. This is a small framework for the implementation of comparative visualization approaches and is structured as follows:

1. `initRenderer()` - Creates all DOM elements and sets the dimensions. It also creates a basic scene with some lights and provides functionality for resizing and click listeners. The renderer-specific user-settings are also initialized and displayed here in this step. They are created in a function called `createSettings()`, which is called at this point and allows for different implementations or extension of settings in subtypes.

2. `loadOrgans(organs)` - Loads all organs, which have to be visualized in an internal map with the organ label as key.
3. `alignOrgans()` - Aligns all organs, so that their centroids are at the origin of the scene. Based on this, the position of organs can be adjusted easily afterwards.
4. `onOrgansLoaded()` - This is a function for subtypes of `ThreeRenderer` to provide any functionality, which needs to be implemented at this step and requires the organs to be loaded, e.g. creation and alignment of meshes including computation of their colors.
5. `initCameraSettings()` - This function adds a camera with controls to the scene. To be customizable, the camera is created in a separate function, called `createCamera()`, which returns a basic `THREE.PerspectiveCamera`, but can be overridden to change the camera type without losing camera controls.
6. `onPreRender()` - In this function, any functionality, which is still missing and needs the camera to be initialized, can be implemented by subtypes before the scene is rendered in the next function. An example would be a comparison of organ contours, which was considered in the visualization design, but then discarded.
7. `render()` - The final step, where the scene is rendered.

There are four comparative visualization approaches implemented: Juxtaposition, Superimposition, encoding of variance in color on an organ of the reference model, and the approach described in Section 4.4.1. These are compared in Figure 5.4. Juxtaposition is implemented by just providing different positions for each organ, aligned on a grid. This is technically not exactly juxtaposition, as the organs are in the same coordinate space, but for comparison to the other techniques, it is sufficient, if there is no interaction with the view. Superimposition is just provided in a basic approach, displaying all organs semi-transparently at the same position. So the corresponding `ThreeRenderer` is tweaking with the material settings and also applying rotation with the Kabsch algorithm [Kab76; Kab78] to all organs to be of minimum overall variance. This algorithm is also used in the proposed visualization design of this thesis and described in Section 4.4.2. It was implemented using `numeric.js` [Loi12]. Because WebGL produces artifacts when using the back-facing sides of polygons in transparent objects, only the front-facing sides are used in the superimposition approach.

The last two approaches do not directly visualize all shapes at once, so different meshes were created in the function `onOrgansLoaded()`. The additional calculations of variance and relative position, and the creation of glyphs are added in the provided function before the scene is rendered. Using the framework provided with our application, it is easily extensible with future visualization approaches.

The major limitation of the current implementation is that the comparative visualization has no support for selection within the pipeline and so it is completely reset when the

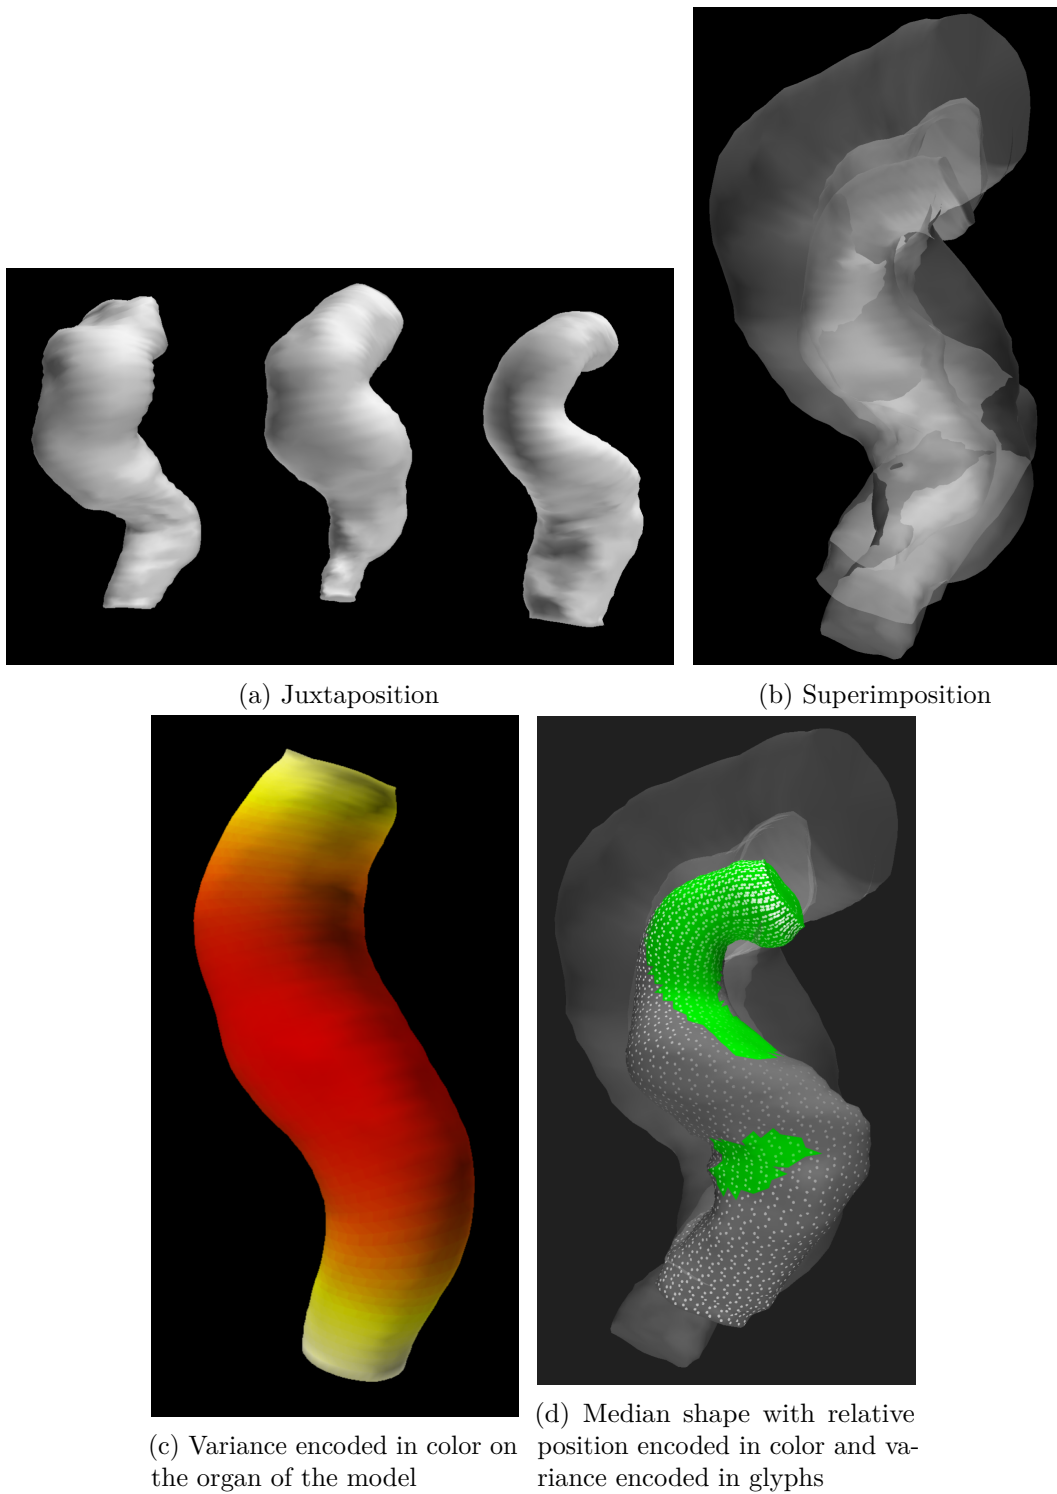


Figure 5.4: Comparison of three rectums visualized by the four implemented approaches of comparative visualization based on the described sequence in Section 5.4.

selection changes. While this is interactive with the data set described in this thesis, using a higher number of patients may slow down the application during selection. Nevertheless, this can still be implemented in the abstract `ThreeRenderer` class without having to change all approaches, as in this case there are just parts of the pipeline, which are called again.

Results

This chapter covers usage scenarios, which are enabled by implementing the two tasks described previously. All of them are performed with a data set of segmentations of pelvic organs, represented as triangle-meshes, from 17 different patients. The segmentations are produced using the algorithm by Schadewaldt et al. [Sch+13] and they include the following organs: prostate, bladder, rectum, and seminal vesicles. All meshes have point-to-point correspondence as well as organs and interfaces between bladder and prostate and between prostate and seminal vesicles labeled. To analyze the data, the application supports a number of settings, which are provided for each scenario and explained as follows:

- 1. Organ/Interface Selection:** With this setting, the user is able to select the organs and the interfaces to process. Because we wanted to take all six classes of the data into account and see how they look like, they are all selected in each scenario.
- 2. Maximum Frequency of the Spherical Harmonics Descriptor:** Determines, to which frequency l the spherical harmonics are computed for the spherical function and therefore determines the length of the shape vectors. As described in Section 4.3.1, the analysis works best with a maximum l of seven and because the computation time is not very high for the amount of data, it is also used in this chapter.
- 3. Spherical Coefficient Computation Strategy:** As described in Section 4.3.1, the spherical coefficients can be computed by using the average of coefficients of all sampled points of the spherical function or by using the method of least squares. The former is better for evenly spread out data points with respect to a sphere and the latter is better for elongated objects as well as describing errors in the data.

6. RESULTS

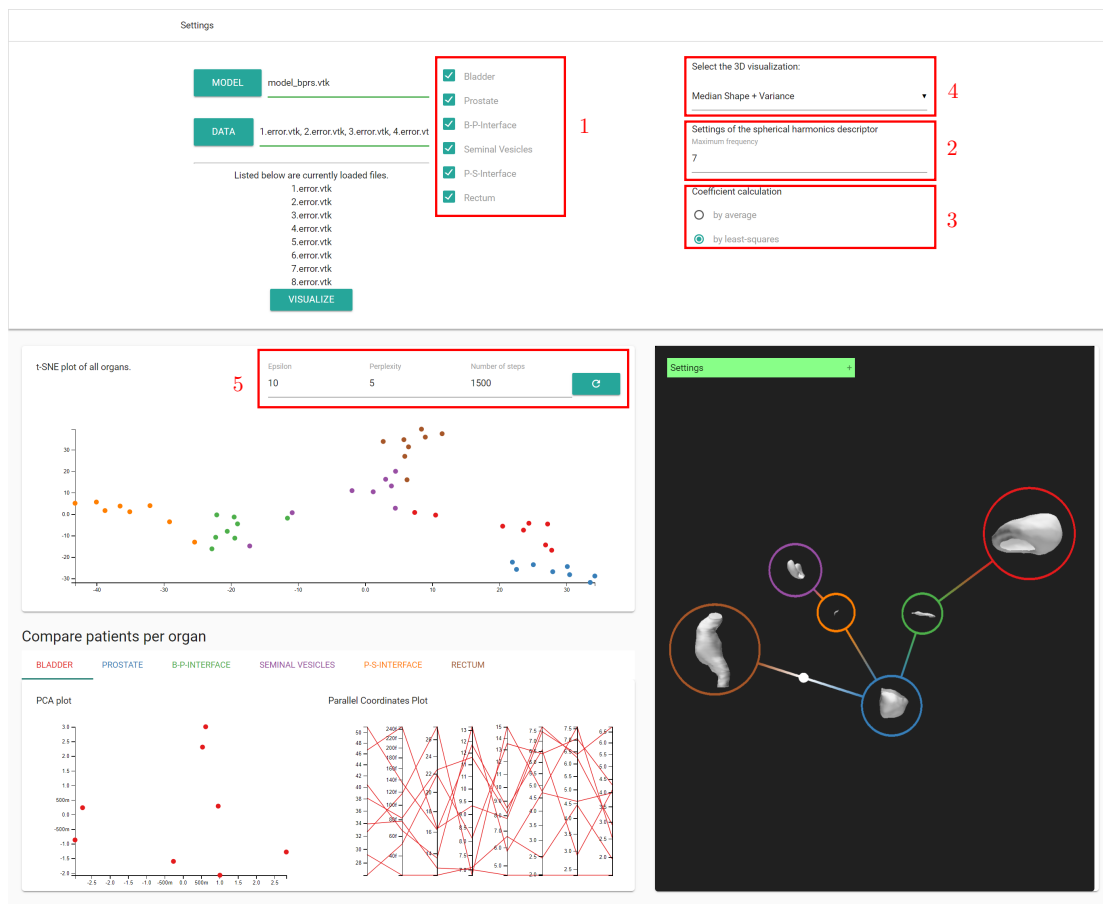


Figure 6.1: Screenshot of the final application with annotations referring to the settings.

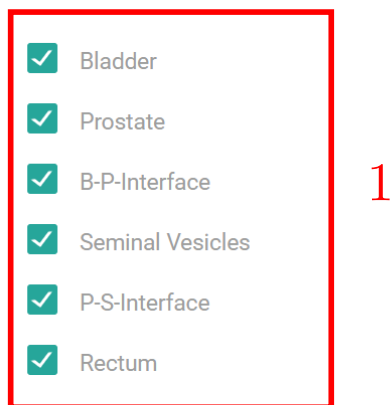


Figure 6.2: Close-up of the settings regarding organ/interface selection.

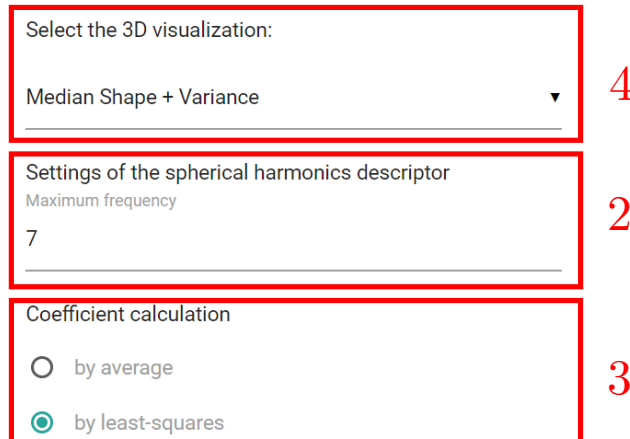


Figure 6.3: Close-up of the settings regarding spherical harmonics and comparative visualization type.



Figure 6.4: Close-up of the t-SNE settings.

4. Visualization Technique: The application has implemented the comparative visualization technique described in Section 4.4.1, which is used in all scenarios of this chapter. For just displaying the mesh of one segmentation or using another view, some other techniques, like Juxtaposition and Superimposition, were also implemented and are accessible through this setting. Each technique also has specific settings. The technique used in these scenarios, for example, allows the user to change the glyph type or using the median shape for computation of the local per-triangle variances, which was described in Section 4.4.2. In this chapter, the settings of the mentioned technique are set to display the point-based glyphs and to not compute the variances based on the median shape, but on the mean.

5. t-SNE Settings: As t-SNE is a parametrized dimensionality reduction technique, which has to be optimized for each data set, its parameters are changeable through the application. Perplexity is about the number of close neighbors of each point in the data, epsilon is the learning rate and the number of steps is needed to limit the approximation. Further information and insight about the parameters is provided by Wattenberg et al. [WVJ16].

6.1 Exploration of All Organs

Figure 6.5 shows all organs of the data set. The selected ones are grouped with the wrong clusters. At further examination in the exploded view, the reason becomes apparent, as they look more like their surrounding organs, than the other ones of their type. Starting from left to right, for example, there is a rectum within the cluster of bladders, probably because of the size. The bladder in the cluster of prostates has an unusual shape, as bladders are typically smoother. It fits in this cluster because of similar size and the slightly increased higher frequencies due to the bumps, which are also present in some prostates. The next step would be to look for a reason, why this bladder has this specific shape. Moving to the right, there are two bladders within the cluster of rectums, which are the only two with an elongated shape and therefore they are more similar to rectums. The prostate, which is located at the bottom, is of normal shape and still misclassified. This is where the shape descriptor can be improved. There is probably a problem with interfaces, which amplifies higher frequencies. There are also two prostates located close to the seminal vesicles. The left one partially has the shape of the segmented seminal vesicles looking from the bottom. Both of them are probably positioned here because of the flat shape and the inaccuracies caused by the interfaces. On the right are two rectums where the segmentation algorithm failed and a pair of seminal vesicles, which are slightly off the others but still similar.

6.2 Per-Organ Exploration

More detailed analysis than just global similarity of shapes can be done in the per-organ section of the application. If not stated otherwise, all results in this section are obtained by limiting the frequency in the calculation of the spherical harmonics to seven and using the method of least squares for computing the spherical coefficients. Figure 6.6 and Figure 6.7 show the exploration of all rectums in the data set. The vertical direction in the PCA plot corresponds to the size. There are two outliers visible, in the PCA plot as well as in the PCP, which are due to segmentation errors. The other shapes are close to each other. After a selection of similar shapes, like in Figure 6.8, local shape analysis can be done in the comparative visualization part of the application.

The axes in the PCA plot of bladders in Figure 6.9 roughly correspond to size and roundness. There are no segmentations, where the algorithm failed completely, like with the rectums, so the distribution depends solely on the variability in shape. The PCP in Figure 6.10 shows two vectors with higher values at high frequencies, which correspond to very flat shapes. Selection shows, that these two vectors are also outliers in the PCA plot as well as the t-SNE plot, where they are grouped with vectors of rectums due to similar size and roundness.

According to the data, segmentations of seminal vesicles have low variability in shape and mainly differ in size. The PCA plot in Figure 6.11 shows this by a very dense cluster and the PCP in Figure 6.12 shows a bundle of similar values across most vectors. The

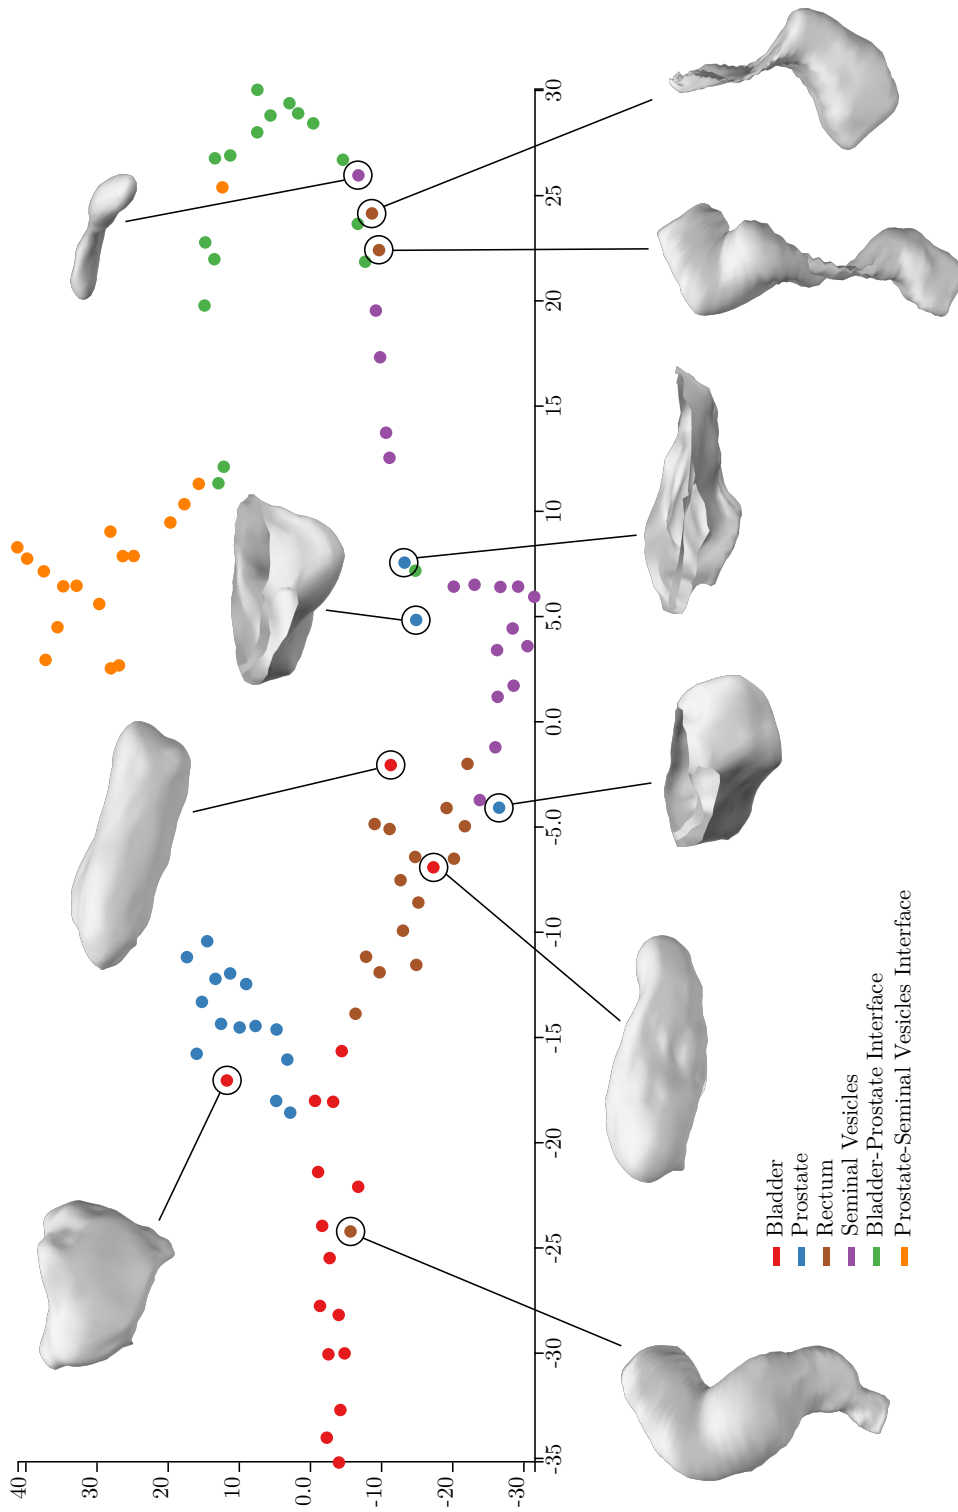


Figure 6.5: t-SNE plot of all organs highlighting the ones, which are grouped with the wrong clusters. The colors show the organ or interface class the shapes belong to and the sizes are not true to scale.

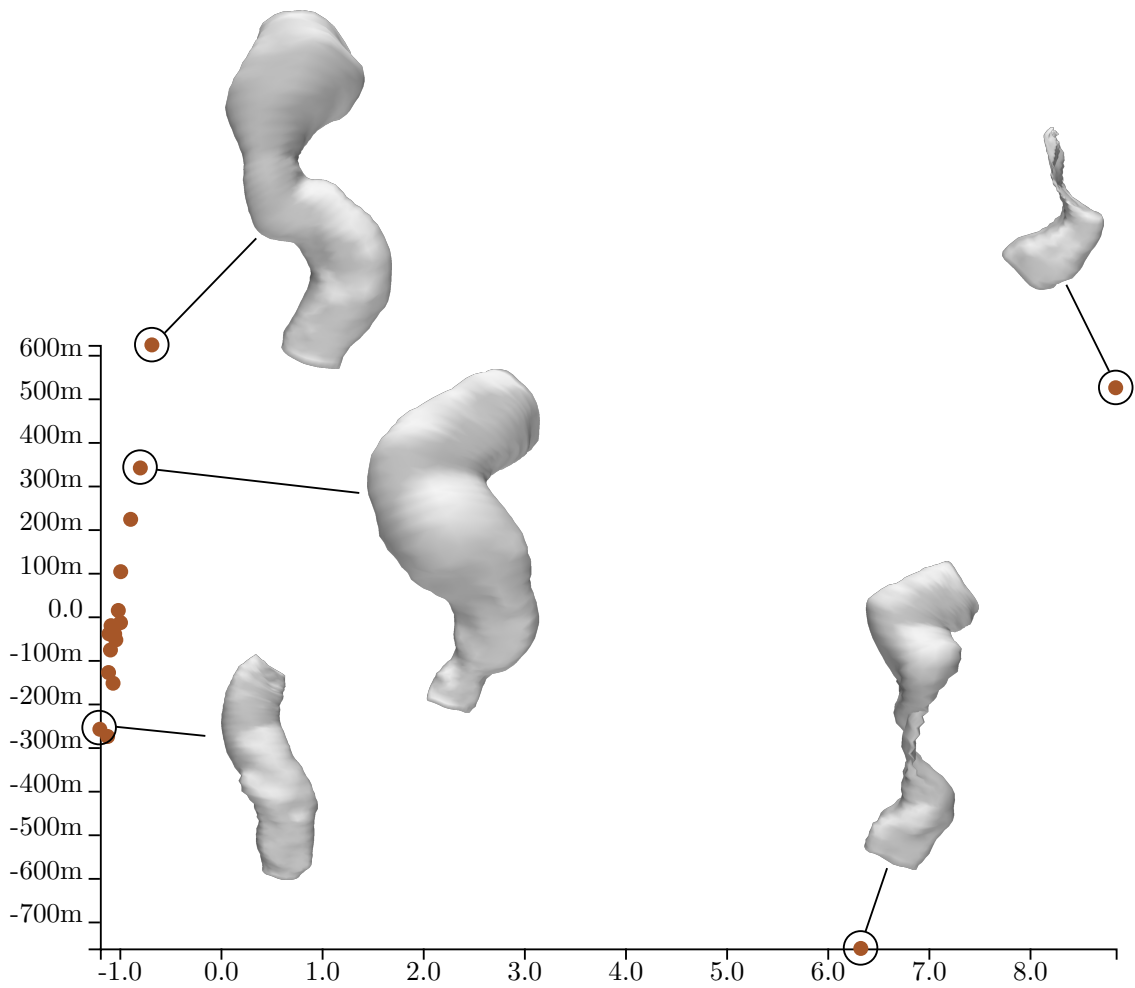


Figure 6.6: PCA plot of all rectums in the data set: Visible are two outliers, which correspond to segmentation errors.

outliers correspond partially to obvious segmentation errors, which become visible in the 3D-view, but one of them seems to be accurately segmented. Their values in the high-frequency coefficients are also similar to the ones of the segmentation errors in the rectums. Also noticeable was, that within these outliers, the zero order coefficients do not correctly encode the size of the shape, which raises the question, if the descriptor is actually suitable for this type of shape.

The segmentation algorithm should perform best on the prostate, as it is prioritized for accuracy for the radiotherapy process. There is only a small difference in size and more in the high frequencies, although it may look different in the PCP of Figure 6.14 due to scaling to the outliers. One of them is completely deformed and probably results from a bad segmentation and the other one looks flattened at the side of the bladder. Looking




Figure 6.7: Parallel Coordinate Plot of all rectums in the data set: Visible are two outliers, which correspond to segmentation errors.

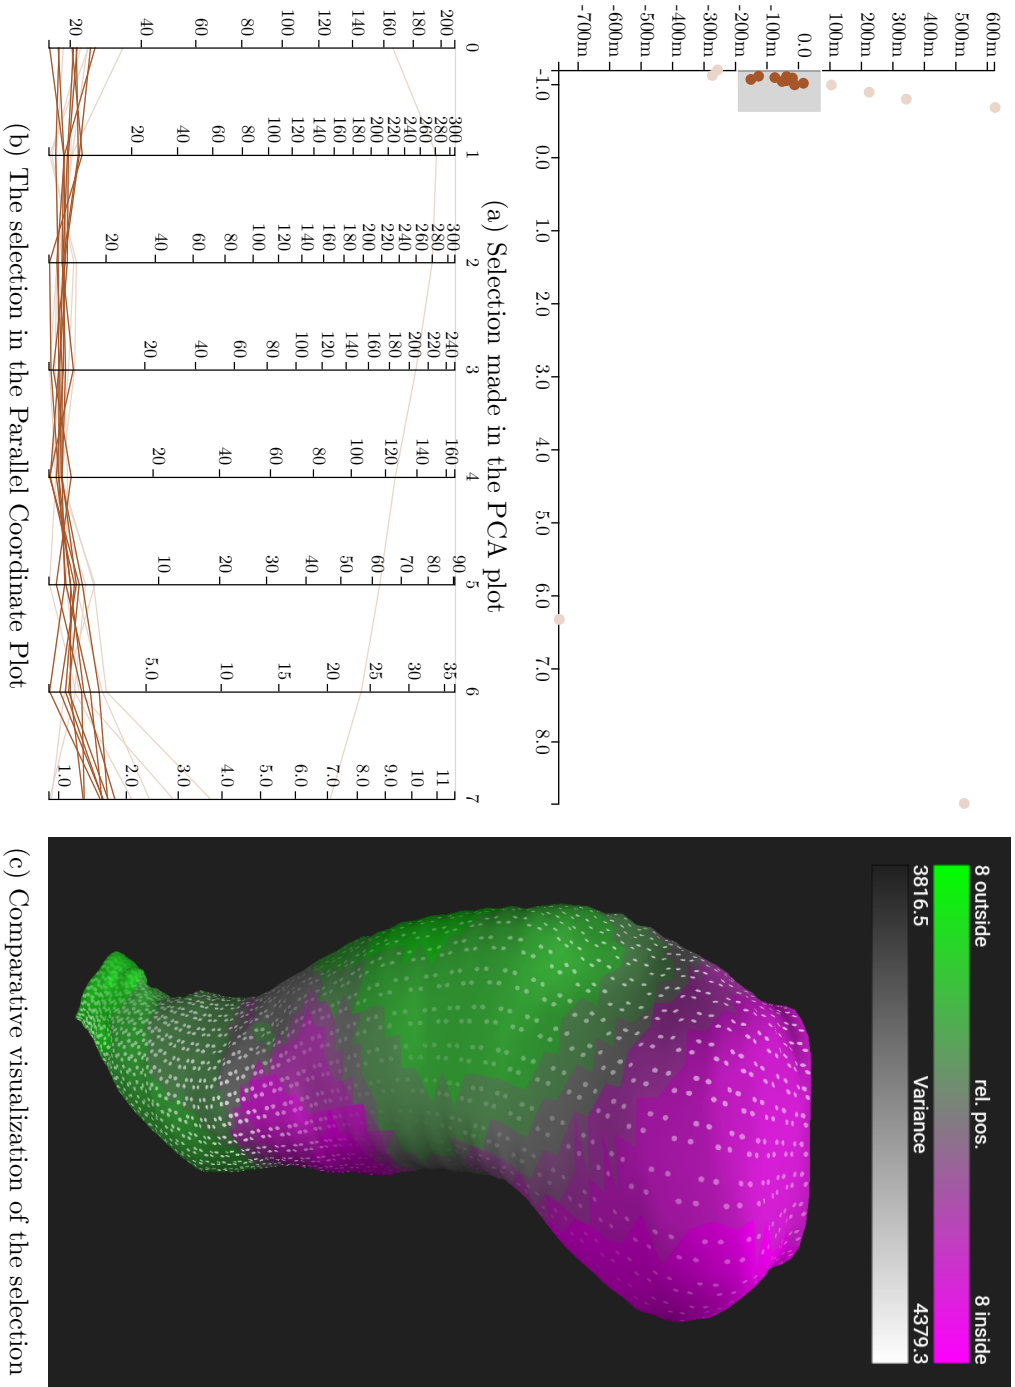


Figure 6.8: Exploration of all rectums in the data set: Similar shapes are selected for further analysis in the comparative visualization to also explore the local variability of rectums.




Figure 6.9: PCA plot of all badders in the data set: The plot shows the distribution well, based on the features size and roundness.

at the whole segmentation mesh of the patient, this is probably caused by the bladder, which is filled and pushes against the prostate. The mis-segmented prostates are also easily seen in the PCA plot in Figure 6.13.

While exploring interfaces, the PCP in Figure 6.15 clearly shows, that the spherical harmonics descriptor of this thesis is not applicable here, as there are no features of the shapes, which correspond to specific values in the shape vectors. For example, the values of frequency zero range from 10 to 27,000 while the shapes are not showing remarkable differences in size. Also for some shapes, the values are abnormally high at each degree, even by a factor of 10^3 compared to regular organs. For comparing interfaces with each other, a different descriptor is required and subject for future work.

6.3 Comparative Visualization

As already mentioned, the exploded view is useful for comparing different organs and their relative positions to each other. The detailed view is for comparing different segmentations of an organ type to create hypotheses about their shapes and the algorithm performance. This section discusses some hypotheses, which were derived from looking at the data in the comparative visualization part of the application. All images and hypotheses in this section are based on selections of representative segmentations, which do not result from a failure of the algorithm.




Figure 6.10: Parallel Coordinate Plot of all bidders in the data set: The axis on the left, which corresponds to frequency zero, mostly describes the size and the other frequencies describe the elongation and features like bumps or edges. Small values, e.g. of the highest frequencies, are abbreviated with SI prefixes.




Figure 6.11: PCA plot of all seminal vesicles in the data set: Most of the shapes are similar, with some outliers being caused by the segmentation algorithm as well as by the shape descriptor. The bottom right shape is probably not suitable for the shape descriptor, as it would correctly belong somewhere closer to the dense cluster.

Looking at the colors on the median shape as well as at the biggest shape of the rectums in Figure 6.16, the high variability in bending and size of the rectums becomes obvious. The triangle positions on the top have a higher variance than the ones in the middle and bottom, which may result from the top being affected more by the filling state. The area of lowest variance is on the back side, probably because of the spine, which is located right there and restricts enlargement in this direction.

As shown in Figure 6.17, the per-triangle variance of the bladders progresses from the tip having the lowest value to the interface with the prostate having the highest value. A hypothesis on this is, that the variance is related to the segmentation quality. The findings of Raidou et al. [Rai+16] show, that the distance of the segmentation to the ground truth data increases from the interface to the tip. Combining these findings with the observed variance would mean, that an area, where the segmentation is accurate, correctly reflects the variability of the bladders and that an area of high segmentation error is of lower per-triangle variance. An explanation could be, that areas of high error are more similar to the reference model than the actual organ and therefore do not reflect the variability.

All areas of the prostates are of similar variance, which results from the size differences. The high frequencies mentioned in the previous section are not noticeable in the variance,




Figure 6.12: Parallel Coordinate Plot of all seminal vesicles in the data set: Most of the shapes are similar, with some outliers being caused by the segmentation algorithm as well as by the shape descriptor.




Figure 6.13: PCA plot of all prostates in the data set: The cluster of similar prostates shows that one dimension is encoding the size. Two outliers are visible. This is an example of where the shape descriptor performs very well.

if enough shapes are selected. The variance increases with proximity to the interface with the seminal vesicles, i.e. from front to back, but is still under 10 %.

6.4 Exploration based on Different Properties

Limiting the frequency in the calculation of the spherical harmonics to lower orders, e.g. zero or one, the variability in terms of size or elongation is expected to be better visible in the PCA and t-SNE plots, but shows bad results. Looking at the bladders in the PCA plot in Figure 6.19, the axes have far less meaning than in Figure 6.9, although the PCP in Figure 6.10 shows, that the size is encoded well with frequency zero. With maximum frequency one, segmentation errors can sometimes still be distinguished, shown in Figure 6.20, but, as expected, far worse than the previous results using higher frequencies.




Figure 6.14: Parallel Coordinate Plot of all prostates in the data set: Frequency zero mostly depicts the size of the shapes and two outliers are visible. This is an example of where the shape descriptor performs very well. Small values, e.g. of the highest frequencies, are abbreviated with SI prefixes.




Figure 6.15: Parallel Coordinate Plot showing the shape vectors of interfaces between bladder and prostate, which have extremely high values for each frequency. To abbreviate large numbers, SI prefixes are used.

Setting the coefficient computation strategy to average, the descriptor depicts features better on the prostates and bladders and is not effective on the other organs and interfaces, as already mentioned in Section 4.3.1. The t-SNE plot in Figure 4.10, which uses the average, generally shows better separated clusters than the one in Figure 6.5, which uses the method of least squares. The axes of PCA plots of the bladders and prostates group similar shapes better, but do not show outliers as well, which is similar to the comparison of PCA plots of rectums in Figure 4.6. The difference is, that the latter does not show an improvement in grouping similar segmentations together due to the different shape of the organ. In the t-SNE visualization there are also normal rectums near to the cluster of seminal vesicles using the average, but the erroneous segmentations are not separated. This is a further example of the average performing badly on shapes, which are not uniformly sampled with respect to a sphere.

Using the median shape in the variance computation instead of the mean, the minimum variance is lower and the maximum is higher, but the areas of low and high variance are not significantly different. Due to the extended range, it is probably better to see some differences in specific parts of the organ, but when the median has an unusual shape, using the mean has higher robustness.




Figure 6.16: Comparative visualization of 15 rectums with low variance at the back side, which is the green area in the middle, and high variance on the top due to the filling state of some rectums. The semi-transparent shape is the biggest rectum of the selection to show the possible extent the organ can have.




Figure 6.17: Comparative visualization of 17 bladders showing a difference in variance between the tip on the top right and the interface on the bottom left, which is probably related to segmentation accuracy. The semi-transparent shape is the biggest bladder of the selection to show the possible extent the organ can have.

6.5 Critical Review

As already mentioned, the shape descriptor cannot handle flat meshes like the interfaces, as the spherical function is based on the distance to the centroid. A solution to this problem could be using another spherical function to describe the mesh or use another shape descriptor in general. For interfaces, a combination of simple global features [PI97; ZC01] could be sufficient.

The approach of comparative visualization fulfills Task 2 completely, yet there are some refinements and additions, which could be useful, e.g. a small axis-indicator to see the camera orientation in the detailed view. It is now only possible to select one organ for a detailed analysis. Sometimes a selection of two or more organs would make sense in terms of comparing variances. The user could quickly see which organs are of higher

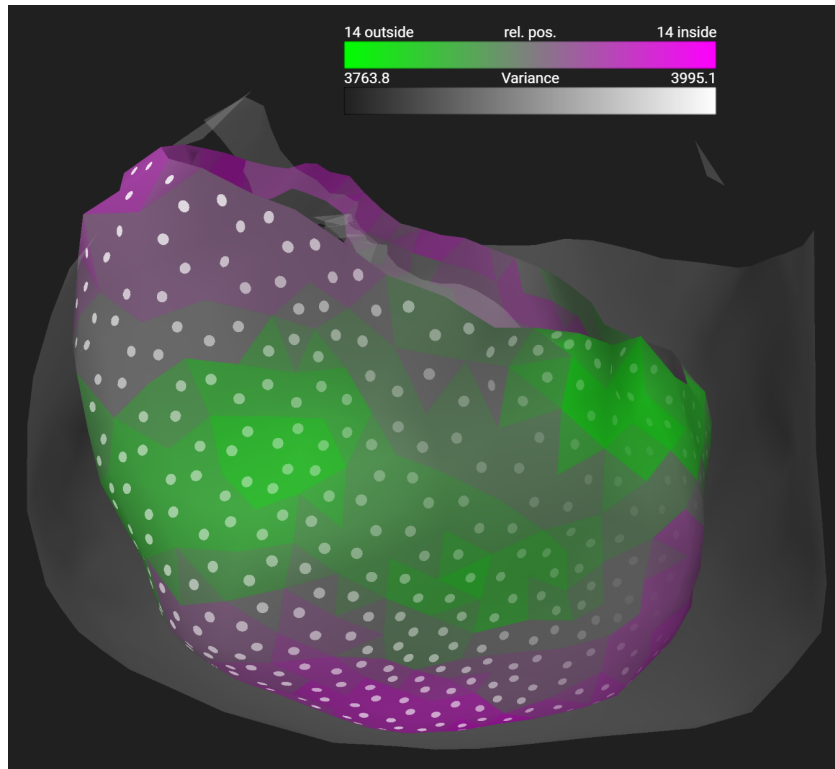


Figure 6.18: Comparative visualization of 15 prostates. The semi-transparent shape is the biggest prostate of the selection to show the possible extent the organ can have.

variability in some areas than others instead of comparing the values after looking them up on the scales provided with the views. To allow for better comparison in the exploded view, a way to rotate individual organs or all at once could be introduced, as currently they are all seen only from one direction. Other interaction methods with the organs could also be included to improve organ exploration.

A drawback of the variance computation is, that the variance is different for specific areas, if two shapes, which are not spheres, have the same shape and only differ in size. To account for this, a setting could be introduced, which changes the variance calculation to take the shape into account. If d is the Euclidean distance from a triangle $\bar{\mathbf{x}}$ of the mean mesh to the center, then the variance of the data at this triangle can be normalized to be

$$\sigma_{\mathbf{x}_{norm}}^2 = \frac{1}{n} \sum_{i=1}^n \left\| \frac{\mathbf{x}_i}{d} - \frac{\bar{\mathbf{x}}}{d} \right\|^2 = \frac{1}{d^2 n} \sum_{i=1}^n \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 = \frac{1}{d^2} \sigma_{\mathbf{x}}^2 \quad (6.1)$$

The point-based glyph design solves some perceptive issues of the line-based one, but it is sometimes hard to see small differences in the variance. Areas with small triangles




Figure 6.19: PCA plot of all bladders in the data set with the shape vector limited to the first two frequencies: The axes do not correspond to the features size and elongation.

also have more points, which means it is brighter overall and may seem to have higher variance compared to areas with large triangles. The color scale to show the relative position is suitable for its purpose and works well in combination with the glyphs, but an extensive evaluation would be needed to quantify this.

The results of the dimensionality reduction done before the two dimensional visualization are good, considering that the detection of outliers is easy and shape features are encoded well in the positions. PCA is also known to perform well on natural datasets [MPH07], besides Sammon mapping [Sam69] and multilayer autoencoders. While t-SNE produces good results, finding suitable parameters requires experience and some trial and error. To overcome this problem, an implementation of Sammon mapping, Locally Linear Embedding (LLE) [RS00], or other non-linear dimensionality reduction techniques should be considered in the context of visualizing all organs simultaneously.

Because segmentation experts also wish to compare a higher number of patients, the scalability of the approach may be important. The current implementation provides interactive frame-rates using all 17 patients of the dataset. Each segmentation mesh, consisting of all organs, has 11606 triangles. The calculation of relative position and variance in the comparative visualization has to be optimized when using a higher number of patients, as it would slow down the selection process. The calculation of shape vectors, which is computationally demanding, as well as the dimensionality reduction




Figure 6.20: PCA plot of all rectums in the data set with the shape vector limited to the first two frequencies: The two outliers are on the outside of the cluster, but still not as easy to distinguish than with the use of higher frequencies, because these show the highest difference to normal shapes.

are performed before the visualization and therefore do not affect interaction. The high number of runs in t-SNE blocks the usability at the beginning, as it still takes some time to finish while all other components are ready. Table 6.1 shows the performance of visualization techniques and other parts of the application with respect to the number of patients within Google Chrome version 63 on a system with an Intel Core i7-2600k CPU and a NVIDIA GTX 970 GPU.

Table 6.1: Computation time for specific tasks in [ms] measured within Google Chrome version 63 on a system with an Intel Core i7-2600k CPU and a NVIDIA GTX 970 GPU.

patients	PCA	PCP	t-SNE	3D overview	3D detail	shape vectors
17	13	14	2958	462	603	3257
80	14	16	34111	1197	971	19495
160	16	20	117963	4412	1504	41233

6.6 Evaluation

An informal evaluation was conducted with a senior scientist at Philips Healthcare Best and professor at the Medical Image Analysis Group at TU Eindhoven, The Netherlands, Prof. dr. ir. Marcel Breeuwer. He has experience in the field of segmentation of more than 25 years and already knew the data set. The purpose was to illustrate the functionality and some initial results, which can be achieved using the proposed visual tool. The usage scenarios described in this thesis were presented remotely to the evaluator, who mentioned that the tool is easy to understand and easy to learn. It provides valuable insight into shape variability and segmentation errors. However, for relating segmentation errors to image quality, e.g. scan artifacts, the original image data should also be accessible through the application, with contours showing the segmentation. To see the actual size of the displayed shapes, a scale should also be implemented. The feedback also includes, that the glyph design encoding the variance, is a discrete representation of continuous values and should therefore be randomly sampled. In the current approach, the model has too much influence on the visual perception of the variance, as one dot per triangle is shown. As for the 2D visualization, the colors of unselected shape representations are not always clearly visible and sometimes, it is hard to distinguish between different organ types. All in all, the framework was regarded very positively, but a more detailed evaluation would need to be conducted in the future.

Conclusion

7.1 Summary

This thesis deals with a visual analytics approach for comparing segmentations of multiple organs over a cohort of multiple patients. This is done in order to develop hypotheses about how algorithms used for automatic segmentation of pelvic organs work in specific cases and what the corresponding segmentation results look like. The automatic segmentation is part of the planning phase in radiotherapy, a common prostate cancer treatment using a high irradiation dose to control malignant cells. As the variability in shape and size of the organs is expected to influence the segmentation outcome quality, its exploration is of high interest.

The first task addressed in this work was to measure and visualize shape variability at a single- and multi-organ basis, which was done by using a shape descriptor based on spherical harmonics and the L_2 distance as similarity measure. This ensured a rotation invariant comparison of shapes while still taking the size into account. The distribution of shapes was further visualized to explore similarities and find mis-segmented shapes, by performing dimensionality reduction and showing the results in plots. The techniques employed in this thesis are PCA - performed on each organ separately - and t-SNE - to be used on all organs at once - , each producing good results. In order to be useful, the latter still needs experience in finding the right parameters, so other techniques like Sammon mapping [Sam69] or LLE [RS00] could be considered as replacement. In addition to the shape space visualization, the shape vectors resulting from the shape descriptor were displayed in a PCP to be able to compare the actual amount of frequencies the shapes consist of. The combination of these plots with interactive selection enables the user to explore the whole cohort based on general shape similarity and different features like size.

To be able to visualize the selection in 3D and actually show what the shapes look like, the problem of visualizing multiple segmentations of multiple organs comparatively,

with respect to each other, was dealt with in the second task of this thesis. Current comparative visualization techniques are limited to only one class of shape, which is the main difference to our approach. It starts by showing an exploded view of representative organs, visually preserving the relative positions and connections. Local shape analysis is further provided by showing shape variances and distribution on demand, encoded in point based glyphs and color.

The tools implemented in this thesis allowed us to quickly identify mis-segmented shapes and explore the data-set based on different properties. Using the comparative visualization technique, some hypotheses about the relation of local segmentation variance and segmentation quality could be made, as well as an exploration of how anatomical relations affect shape variability in different areas.

The *contribution* of this thesis is the design and implementation of a web-based application, which implements the described tasks and allows segmentation experts to understand the results of their algorithms. It is different from previous work in terms of dealing with multiple organs as well as combining shape space exploration with a shape variability analysis of a cohort in 3D. We also propose a way to comparatively visualize multiple organs of multiple patients with the implemented visualization design.

7.2 Limitations and Future Work

The presented work has a lot of advantages, namely the interactive visual exploration of shape variability with simple detection of outliers as well as the intuitive way of analyzing local per-triangle variances in the comparative view. However, a number of limitations and directions for future work has been revealed.

The most significant weakness of the implementation is the shape descriptor, which deals best with sphere-like organs, but produces poor results with flat shapes as well as interfaces, when it comes to comparing different shape features. It is still sufficient for detection of outliers regardless of the organ shape. An improvement of the current application would be finding a suitable shape descriptor, depending on the organ class, or improving the implemented descriptor based on spherical harmonics. We propose the approach of Liu et al. [Liu+04] as an improvement in the robustness of spherical harmonics based descriptors, which could enhance the behavior in areas of the interfaces, where the mesh has holes, as well as in flat organs.

The visualization design is lacking possibilities for interacting with the organs in the exploded view, as it shows them from only one direction and allows the user to select only a single organ for detailed analysis. Selecting multiple organs at once would allow the user to see a comparison of general shape variability between different organs, for example by using the same color or brightness scale for all of them. Implementing an interaction type to be able to rotate organs in the exploded view would allow the user to see more areas and views of an organ before selection, which may be useful for picking organs of interest.

Scalability of the approach is only partially fulfilled, as the interaction would suffer from the slow computation of variances and relative positions in the comparative visualization. This can be solved by a more efficient implementation, because the actual calculation is not computationally demanding.

The next step of understanding segmentation outcomes is to relate the shape variability, which is discussed in this thesis, to actual segmentation outcome quality, which e.g. is discussed by Raidou in [Rai+16]. The examined datasets also include error metrics, for example the distance to the ground truth data, to quantify the segmentation quality. Future work could find a way to analyze the reason behind good or bad quality based on shape variability or features, or even relate per-triangle error variances to shape variances. These relations could provide further insight in how segmentation algorithms work and how to improve them.

The approach of this work could also be used in other applications, where an exploration and analysis of shape features and shape variability is needed, for example the exploration of anatomical variability not only with respect to pelvic organs. It can also be used or extended for comparison of the variability in the outcomes from different parametrizations of segmentation algorithms or different segmentation algorithms in general and would even allow this on multiple patients at once.

List of Figures

2.1	Male pelvic organs seen from right side [Gra18].	6
2.2	Comparison between a filled (left) and an empty rectum	7
2.3	Radiotherapy planning: The top figure shows the field arrangement of the radiation and the bottom one shows the isodose distribution in the tissue [Pal+08].	8
2.4	Segmentation mesh of the relevant pelvic organs, labeled by color.	9
2.5	Architecture (a) and model adaption (b) of the algorithm of Ecabert et al. [Eca+11].	10
3.1	Shape space scatter plot by Busking et al. [BBP10].	13
3.2	Best separating axes for a differentiation based on a specific trait vector. [HSK11]	13
3.3	Bands representing clusters of lumbar spine canal centerlines used by Klemm et al. [Kle+13].	14
3.4	The four fundamental approaches for comparative visualization of 3D data sets described and illustrated by Kim et al. [KCK17]. This simplified illustration shows the case of comparing just two data instances.	15
3.5	Superimposition technique used by Alabi et al. in [Ala+12] based on Weigle's technique [WT05].	16
3.6	Explicit Encoding approach of Schmidt et al. [Sch+14] projecting a heat map of error variances.	17
4.1	Schematic overview of Task 1	20
4.2	Schematic overview of task 2	20
4.3	The three visualization sections of the UI of our proposed tool	21
4.4	Sequence of computing a rotation invariant descriptor based on spherical harmonics [KFR03]. The decomposition into the harmonics and the frequency components are shown for $l = 0$ on the left to $l = 4$ on the right.	23
4.5	Computation time in seconds for the spherical harmonics based descriptor with different maximum frequencies l ranging from 0-13. The performance was measured within Google Chrome version 63 on a system with an Intel Core i7-2600k CPU and a NVIDIA GTX 970 GPU.	24
		73

4.6	Comparison of the distributions using the average and the method of least squares to compute the spherical coefficients. Dimensionality reduction was done using PCA with projection on the first two principal components. The segmentation error is highlighted to show the ability to detect outliers. . .	26
4.7	Comparison of PCA and t-SNE performing on the MNIST dataset. Different colors represent different digits.	28
4.8	PCP of all prostate shape vectors. Small values, e.g. of frequency one, are abbreviated with SI prefixes.	29
4.9	PCA plot of the prostate shape vectors shown in Figure 4.8, projected on the first two principal components.	30
4.10	t-SNE plot of the whole data-set. The colors show the organ or interface class the shapes belong to. The relative size of the displayed organ shapes to each other is preserved to see the impact of shape and size on the position in the plot.	31
4.11	Exploded view approach for displaying multiple organs and their interfaces, preserving the connections in the mesh through the interfaces.	32
4.12	Exploded view: connections between organs without showing interfaces. .	33
4.13	Exploded view: lines between disconnected organs.	33
4.14	Illustration of encoding the relative position with complementary colors. The further inside the shape is relative to the others, the more it changes to the color green, the further outside, the more it changes to magenta. If it is in the center of all shapes, it is gray. The dashed black shapes are not visible in the final visualization and only show the relative position in this figure. .	35
4.15	Line-based glyphs encoding the variance of the shapes. The colors show the relative position to the other shapes and the semi-transparent shape is the largest in the cohort.	36
4.16	Point-based glyphs encoding the variance of the shapes. The colors show the relative position to the other shapes and the semi-transparent shape is the largest in the cohort.	37
5.1	Diagram of the main components of the application and their interactions with each other. Green components implement visualization techniques and interact with the DOM. Blue components are responsible of the application logic and communication.	40
5.2	Values of the shape vectors of a sphere, a cube and a tetrahedron. It is easily recognizable how different frequencies describe different geometries and symmetry properties.	42
5.3	The real spherical harmonics up to frequency four, green is showing regions where they are positive and red where they are negative. The extent indicates the value in this direction. These values are also encoded in color on the sphere. [Dif17]	43
5.4	Comparison of three rectums visualized by the four implemented approaches of comparative visualization based on the described sequence in Section 5.4.	45

6.1	Screenshot of the final application with annotations referring to the settings.	48
6.2	Close-up of the settings regarding organ/interface selection.	48
6.3	Close-up of the settings regarding spherical harmonics and comparative visualization type.	49
6.4	Close-up of the t-SNE settings.	49
6.5	t-SNE plot of all organs highlighting the ones, which are grouped with the wrong clusters. The colors show the organ or interface class the shapes belong to and the sizes are not true to scale.	51
6.6	PCA plot of all rectums in the data set: Visible are two outliers, which correspond to segmentation errors.	52
6.7	Parallel Coordinate Plot of all rectums in the data set: Visible are two outliers, which correspond to segmentation errors.	53
6.8	Exploration of all rectums in the data set: Similar shapes are selected for further analysis in the comparative visualization to also explore the local variability of rectums.	54
6.9	PCA plot of all badders in the data set: The plot shows the distribution well, based on the features size and roundness.	55
6.10	Parallel Coordinate Plot of all badders in the data set: The axis on the left, which corresponds to frequency zero, mostly describes the size and the other frequencies describe the elongation and features like bumps or edges. Small values, e.g. of the highest frequencies, are abbreviated with SI prefixes.	56
6.11	PCA plot of all seminal vesicles in the data set: Most of the shapes are similar, with some outliers being caused by the segmentation algorithm as well as by the shape descriptor. The bottom right shape is probably not suitable for the shape descriptor, as it would correctly belong somewhere closer to the dense cluster.	57
6.12	Parallel Coordinate Plot of all seminal vesicles in the data set: Most of the shapes are similar, with some outliers being caused by the segmentation algorithm as well as by the shape descriptor.	58
6.13	PCA plot of all prostates in the data set: The cluster of similar prostates shows that one dimension is encoding the size. Two outliers are visible. This is an example of where the shape descriptor performs very well.	59
6.14	Parallel Coordinate Plot of all prostates in the data set: Frequency zero mostly depicts the size of the shapes and two outliers are visible. This is an example of where the shape descriptor performs very well. Small values, e.g. of the highest frequencies, are abbreviated with SI prefixes.	60
6.15	Parallel Coordinate Plot showing the shape vectors of interfaces between bladder and prostate, which have extremely high values for each frequency. To abbreviate large numbers, SI prefixes are used.	61
6.16	Comparative visualization of 15 rectums with low variance at the back side, which is the green area in the middle, and high variance on the top due to the filling state of some rectums. The semi-transparent shape is the biggest rectum of the selection to show the possible extent the organ can have.	62

6.17	Comparative visualization of 17 bladders showing a difference in variance between the tip on the top right and the interface on the bottom left, which is probably related to segmentation accuracy. The semi-transparent shape is the biggest bladder of the selection to show the possible extent the organ can have.	63
6.18	Comparative visualization of 15 prostates. The semi-transparent shape is the biggest prostate of the selection to show the possible extent the organ can have.	64
6.19	PCA plot of all bladders in the data set with the shape vector limited to the first two frequencies: The axes do not correspond to the features size and elongation.	65
6.20	PCA plot of all rectums in the data set with the shape vector limited to the first two frequencies: The two outliers are on the outside of the cluster, but still not as easy to distinguish than with the use of higher frequencies, because these show the highest difference to normal shapes.	66

Glossary

brachytherapy is a part of cancer treatment, in which the malignant tissues are irradiated by implanting permanent radioactive seeds into the prostate. It is an option for patients with low-risk disease and selected cases with intermediate-risk disease [Mot+17]. 6

digital rectal examination in men, is a type of physical examination of the prostate, where the doctor uses his finger, which he inserts into the rectum, to check pelvic structures. 5

prostate-specific antigen is a protein, which is produced in the prostate by normal and malignant cells. For screening purposes regarding prostate cancer, a blood test is made, which determines the level of the prostate specific antigen. 5

radical prostatectomy is an option of prostate cancer treatment, where the prostate itself and some surrounding tissue is removed. It is usually performed on younger patients with life expectancy of more than 10 years. 6

radiotherapy is a part of cancer treatment, which uses ionizing radiation to control malignant cells. The patient is exposed to external beams from multiple angles, which intersect at the tumor location to preserve healthy tissues and concentrate on the malignant ones. 1, 5, 6, 52, 63

segmentation in this context, is the process of dividing a medical image into meaningful parts (e.g. organs). xi, xiii, 1–3, 5–10, 12, 15, 16, 20, 21, 24, 25, 47–49, 51, 52, 54, 56, 58–61, 63–65, 68, 69

Acronyms

DOM Document Object Model. 40, 43, 68

DRE digital rectal examination. 5, *Glossary*: digital rectal examination

GPA Generalized Procrustes Analysis. 12

LLE Locally Linear Embedding. 61, 63

PCA Principal Component Analysis. 7, 12, 21, 25–30, 39, 43, 49, 51–55, 59–61, 63, 68, 69

PCP Parallel Coordinate Plot. 27, 29, 30, 39, 41, 49, 51–56, 63, 68, 69

PSA prostate-specific antigen. 5, *Glossary*: prostate-specific antigen

SVD Singular Value Decomposition. 26, 32, 34

t-SNE t-Distributed Stochastic Neighbor Embedding. 26–31, 34, 39, 43, 48–50, 59–61, 63, 68, 69

Bibliography

- [ABB95] Helmut Alt, Bernd Behrends, and Johannes Blömer. “Approximate matching of polygonal shapes”. In: *Annals of Mathematics and Artificial Intelligence* 13.3 (Sept. 1995), pp. 251–265. ISSN: 1573-7470. DOI: 10.1007/BF01530830.
- [Ala+12] Oluwafemi S. Alabi, Xunlei Wu, Jonathan M. Harter, Madhura Phadke, Lifford Pinto, Hannah Petersen, Steffen Bass, Michael Keifer, Sharon Zhong, Chris Healey, et al. “Comparative Visualization of Ensembles Using Ensemble Surface Slicing”. In: *Proceedings of SPIE*. Vol. 8294. NIH Public Access. Jan. 2012. DOI: 10.1117/12.908288.
- [Ale+17] Shiva Alemzadeh, Tommy Hielscher, Uli Niemann, Lena Cibulski, Till Itermann, Henry Völzke, Myra Spiliopoulou, and Bernhard Preim. “Visual Subpopulation Discovery and Validation in Cohort Study Data”. In: *CoRR* abs/1711.09377 (2017). arXiv: 1711.09377.
- [ARS04] Dean C. Adams, F. James Rohlf, and Dennis E. Slice. “Geometric morphometrics: Ten years of progress following the ‘revolution’”. In: *Italian Journal of Zoology* 71.1 (2004), pp. 5–16. DOI: 10.1080/11250000409356545.
- [BBP10] Stef Busking, Charl P. Botha, and Frits H. Post. “Dynamic Multi-View Exploration of Shape Spaces”. In: *Computer Graphics Forum* 29.3 (2010), pp. 973–982. ISSN: 1467-8659. DOI: 10.1111/j.1467-8659.2009.01668.x. URL: <http://dx.doi.org/10.1111/j.1467-8659.2009.01668.x>.
- [Ber+13] Steven Bergner, Michael Sedlmair, Torsten Möller, Sareh Nabi Abdolyousefi, and Ahmed Saad. “ParaGlide: Interactive Parameter Space Partitioning for Computer Simulations”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.9 (Sept. 2013), pp. 1499–1512. ISSN: 1077-2626. DOI: 10.1109/TVCG.2013.61.
- [Bor+13] Rita Borgo, Johannes Kehler, David H.S. Chung, Eamonn Maguire, Robert S. Laramée, Helwig Hauser, Matthew Ward, and Min Chen. “Glyph-based Visualization: Foundations, Design Guidelines, Techniques and Applications”. In: *Eurographics State of the Art Reports*. EG STARS (May 2013). <http://diglib.org/EG/DL/conf/EG2013/stars/039-063.pdf>, pp. 39–63. URL:

<https://www.cg.tuwien.ac.at/research/publications/2013/borgo-2013-gly/>.

- [Bos17] Mike Bostock. *D3.js - Data-Driven Documents*. 2017. URL: <https://d3js.org/> (visited on 02/01/2018).
- [Bre17] Cynthia A. Brewer. *Color Brewer*. 2017. URL: <http://www.colorbrewer.org> (visited on 12/27/2017).
- [Bus+11] Stef Busking, Charl P. Botha, Luca Ferrarini, Julien Milles, and Frits H. Post. “Image-based rendering of intersecting surfaces for dynamic comparative visualization”. In: *The Visual Computer* 27.5 (2011), pp. 347–363. ISSN: 1432-2315. DOI: 10.1007/s00371-010-0541-z.
- [Cab17] Ricardo Cabello. *three.js - Javascript 3D library*. 2017. URL: <https://threejs.org/> (visited on 02/01/2018).
- [CLR11] Siqi Chen, D. Michael Lovelock, and Richard J. Radke. “Segmenting the prostate and rectum in CT imagery using anatomical constraints”. In: *Medical Image Analysis* 15.1 (2011), pp. 1–11. ISSN: 1361-8415. DOI: 10.1016/j.media.2010.06.004. URL: <http://www.sciencedirect.com/science/article/pii/S1361841510000678>.
- [Coo+92] Tim F. Cootes, Chris J. Taylor, D. H. Cooper, and James Graham. “Training Models of Shape from Sets of Examples”. In: *BMVC92: Proceedings of the British Machine Vision Conference, organised by the British Machine Vision Association 22–24 September 1992 Leeds*. Ed. by David Hogg and Roger Boyle. London: Springer London, 1992, pp. 9–18. ISBN: 978-1-4471-3201-1. DOI: 10.1007/978-1-4471-3201-1_2.
- [CT92] Tim F. Cootes and Chris J. Taylor. “Active Shape Models — ‘Smart Snakes’”. In: *BMVC92: Proceedings of the British Machine Vision Conference, organised by the British Machine Vision Association 22–24 September 1992 Leeds*. Ed. by David Hogg and Roger Boyle. London: Springer London, 1992, pp. 266–275. ISBN: 978-1-4471-3201-1.
- [Dif17] Dario Differt. “Holistic methods for visual navigation of mobile robots in outdoor environments”. PhD thesis. Universität Bielefeld, 2017.
- [DM16] Ian L. Dryden and Kanti V. Mardia. “Procrustes analysis”. In: John Wiley & Sons, Ltd, 2016, pp. 125–173. ISBN: 9781119072492. DOI: 10.1002/9781119072492.ch7. URL: <http://dx.doi.org/10.1002/9781119072492.ch7>.
- [Eca+11] Olivier Ecabert, Jochen Peters, Matthew J. Walker, Thomas Ivanc, Cristian Lorenz, Jens von Berg, Jonathan Lessick, Mani Vembar, and Jürgen Weese. “Segmentation of the heart and great vessels in CT images using a model-based adaptation framework”. In: *Medical Image Analysis* 15.6 (2011), pp. 863–876. ISSN: 1361-8415. DOI: 10.1016/j.media.2011.06.004. URL: <http://www.sciencedirect.com/science/article/pii/S1361841511000910>.

- [ETA02] Michael Elad, Ayellet Tal, and Sigal Ar. “Content Based Retrieval of VRML Objects: An Iterative and Interactive Approach”. In: *Proceedings of the Sixth Eurographics Workshop on Multimedia 2001*. Manchester, UK: Springer-Verlag New York, Inc., 2002, pp. 107–118. ISBN: 3-211-83769-8. URL: <http://dl.acm.org/citation.cfm?id=778665.778678>.
- [Fer+15] Jacques Ferlay, Isabelle Soerjomataram, Rajesh Dikshit, Sultan Eser, Colin Mathers, Marise Rebelo, Donald Maxwell Parkin, David Forman, and Freddie Bray. “Cancer incidence and mortality worldwide: Sources, methods and major patterns in GLOBOCAN 2012”. In: *International Journal of Cancer* 136.5 (2015), E359–E386. ISSN: 1097-0215. DOI: 10.1002/ijc.29210. URL: <http://dx.doi.org/10.1002/ijc.29210>.
- [FMH16] Bernhard Fröhler, Torsten Möller, and Christoph Heinzl. “GEMSe: Visualization-Guided Exploration of Multi-channel Segmentation Algorithms”. In: *Computer Graphics Forum* 35.3 (2016), pp. 191–200. ISSN: 1467-8659. DOI: 10.1111/cgf.12895. URL: <http://dx.doi.org/10.1111/cgf.12895>.
- [Geu+15] Alexander Geurts, Georgios Sakas, Arjan Kuijper, Meike Becker, and Tatiana von Landesberger. “Visual Comparison of 3D Medical Image Segmentation Algorithms Based on Statistical Shape Models”. In: *Digital Human Modeling. Applications in Health, Safety, Ergonomics and Risk Management: Ergonomics and Health: 6th International Conference, DHM 2015, Held as Part of HCI International 2015, Los Angeles, CA, USA, August 2-7, 2015, Proceedings, Part II*. Ed. by Vincent G. Duffy. Cham: Springer International Publishing, 2015, pp. 336–344. ISBN: 978-3-319-21070-4. DOI: 10.1007/978-3-319-21070-4_34.
- [Goo91] Colin Goodall. “Procrustes Methods in the Statistical Analysis of Shape”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 53.2 (1991), pp. 285–339. ISSN: 00359246. URL: <http://www.jstor.org/stable/2345744>.
- [Gra18] Henry Gray. *Anatomy of the human body*. Philadelphia: Lea & Febiger, 1918.
- [HBG11] Martin Haidacher, Stefan Bruckner, and Eduard Gröller. “Volume Analysis Using Multimodal Surface Similarity”. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (Dec. 2011), pp. 1969–1978. ISSN: 1077-2626. DOI: 10.1109/TVCG.2011.258.
- [Her+14] Max Hermann, Anja C. Schunke, Thomas Schultz, and Reinhard Klein. “A Visual Analytics Approach to Study Anatomic Covariation”. In: *2014 IEEE Pacific Visualization Symposium*. Mar. 2014, pp. 161–168.
- [Her+16] Max Hermann, Anja C. Schunke, Thomas Schultz, and Reinhard Klein. “Accurate Interactive Visualization of Large Deformations and Variability in Biomedical Image Ensembles”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (Jan. 2016), pp. 708–717. ISSN: 1077-2626. DOI: 10.1109/TVCG.2015.2467198.

- [HK15] Max Hermann and Reinhard Klein. “A visual analytics perspective on shape analysis: State of the art and future prospects”. In: *Computers & Graphics* 53.Part A (2015). 40 years of Computer Graphics in Darmstadt, pp. 63–71. ISSN: 0097-8493. DOI: 10.1016/j.cag.2015.08.008. URL: <http://www.sciencedirect.com/science/article/pii/S0097849315001429>.
- [HKR93] Daniel P. Huttenlocher, Gregory A. Klanderman, and William J. Rucklidge. “Comparing Images Using the Hausdorff Distance”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15.9 (Sept. 1993), pp. 850–863. ISSN: 0162-8828. DOI: 10.1109/34.232073.
- [HSK11] Max Hermann, Anja C. Schunke, and Reinhard Klein. “Semantically steered visual analysis of highly detailed morphometric shape spaces”. In: *2011 IEEE Symposium on Biological Data Visualization (BioVis)*. Oct. 2011, pp. 151–158. DOI: 10.1109/BioVis.2011.6094060.
- [Kab76] Wolfgang Kabsch. “A solution for the best rotation to relate two sets of vectors”. In: *Acta Crystallographica Section A* 32.5 (Sept. 1976), pp. 922–923. DOI: 10.1107/S0567739476001873.
- [Kab78] Wolfgang Kabsch. “A discussion of the solution for the best rotation to relate two sets of vectors”. In: *Acta Crystallographica Section A* 34.5 (Sept. 1978), pp. 827–828. DOI: 10.1107/S0567739478001680.
- [Kar16] Andrej Karpathy. *tSNEJS*. 2016. URL: <https://github.com/karpathy/tsnejs> (visited on 02/17/2018).
- [KBW11] Matthias Kirschner, Meike Becker, and Stefan Wesarg. “3D Active Shape Model Segmentation with Nonlinear Shape Priors”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011: 14th International Conference, Toronto, Canada, September 18-22, 2011, Proceedings, Part II*. Ed. by Gabor Fichtinger, Anne Martel, and Terry Peters. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 492–499. ISBN: 978-3-642-23629-7. DOI: 10.1007/978-3-642-23629-7_60.
- [KCK17] Kyungyoon Kim, John V. Carlis, and Daniel F. Keefe. “Comparison techniques utilized in spatial 3D and 4D data visualizations: A survey and future directions”. In: *Computers & Graphics* 67.Supplement C (2017), pp. 138–147. ISSN: 0097-8493. DOI: <https://doi.org/10.1016/j.cag.2017.05.005>. URL: <http://www.sciencedirect.com/science/article/pii/S0097849317300481>.
- [Kee+09] Daniel F. Keefe, Marcus Ewert, William Ribarsky, and Remco Chang. “Interactive Coordinated Multiple-View Visualization of Biomechanical Motion Data”. In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (Nov. 2009), pp. 1383–1390. ISSN: 1077-2626. DOI: 10.1109/TVCG.2009.152.

- [Ken77] David G. Kendall. “The Diffusion of Shape”. In: *Advances in Applied Probability* 9.3 (1977), pp. 428–430. ISSN: 00018678. URL: <http://www.jstor.org/stable/1426091>.
- [Ken84] David G. Kendall. “Shape Manifolds, Procrustean Metrics, and Complex Projective Spaces”. In: *Bulletin of the London Mathematical Society* 16.2 (1984), pp. 81–121. ISSN: 1469-2120. DOI: 10.1112/blms/16.2.81.
- [Ken94] John T. Kent. “The Complex Bingham Distribution and Shape Analysis”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 56.2 (1994), pp. 285–299. ISSN: 00359246. URL: <http://www.jstor.org/stable/2345900>.
- [KFR03] Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz. “Rotation Invariant Spherical Harmonic Representation of 3D Shape Descriptors”. In: *Symposium on geometry processing*. Vol. 6. 2003, pp. 156–164.
- [Kle+13] Paul Klemm, Kai Lawonn, Marko Rak, Bernhard Preim, Klaus D Tönnies, Katrin Hegenscheid, Henry Völzke, and Steffen Oeltze. “Visualization and Analysis of Lumbar Spine Canal Variability in Cohort Study Data.” In: *VMV*. 2013, pp. 121–128.
- [Kle+14] Paul Klemm, Steffen Oeltze-Jafra, Kai Lawonn, Katrin Hegenscheid, Henry Völzke, and Bernhard Preim. “Interactive Visual Analysis of Image-Centric Cohort Study Data”. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (Dec. 2014), pp. 1673–1682. ISSN: 1077-2626. DOI: 10.1109/TVCG.2014.2346591.
- [KYZ13] Ismail Khalid Kazmi, Lihua You, and Jian Jun Zhang. “A Survey of 2D and 3D Shape Descriptors”. In: *Proceedings of the 2013 10th International Conference Computer Graphics, Imaging and Visualization (Cgiv 2013)*. CGIV '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 1–10. ISBN: 978-0-7695-5051-0. DOI: 10.1109/CGIV.2013.11.
- [Lan+13] Tatiana von Landesberger, Sebastian Bremm, Matthias Kirschner, Stefan Wesarg, and Arjan Kuijper. “Visual Analytics for model-based medical image segmentation: Opportunities and challenges”. In: *Expert Systems with Applications* 40.12 (2013), pp. 4934–4943. ISSN: 0957-4174. DOI: <https://doi.org/10.1016/j.eswa.2013.03.006>. URL: <http://www.sciencedirect.com/science/article/pii/S0957417413001565>.
- [LBB16] Tatiana von Landesberger, Dennis Basgier, and Meike Becker. “Comparative Local Quality Assessment of 3D Medical Image Segmentations with Focus on Statistical Shape Model-Based Algorithms”. In: *IEEE Transactions on Visualization and Computer Graphics* 22.12 (Dec. 2016), pp. 2537–2549. ISSN: 1077-2626. DOI: 10.1109/TVCG.2015.2501813.

- [Lec+98] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11 (Nov. 1998), pp. 2278–2324. ISSN: 0018-9219. DOI: 10.1109/5.726791.
- [Liu+04] Yi Liu, Jiantao Pu, Guyu Xin, Hongbin Zha, Weibin Liu, and Yusuke Uehara. “A robust method for shape-based 3D model retrieval”. In: *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings*. Oct. 2004, pp. 3–9. DOI: 10.1109/PCCGA.2004.1348328.
- [Loi12] Sébastien Loisel. *Numeric Javascript*. 2012. URL: <http://www.numericjs.com/> (visited on 02/17/2018).
- [Lun+07] Claes Lundström, Patric Ljung, Anders Persson, and Anders Ynnerman. “Uncertainty Visualization in Medical Volume Rendering Using Probabilistic Animation”. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (Nov. 2007), pp. 1648–1655. ISSN: 1077-2626. DOI: 10.1109/TVCG.2007.70518.
- [MH08] Laurens van der Maaten and Geoffrey Hinton. “Visualizing Data using t-SNE”. In: *Journal of Machine Learning Research* 9.Nov (2008), pp. 2579–2605.
- [Mor+05] Richard J. Morris, Rafael J. Najmanovich, Abdullah Kahraman, and Janet M. Thornton. “Real spherical harmonic expansion coefficients as 3D shape descriptors for protein binding pocket and ligand comparisons”. In: *Bioinformatics* 21.10 (2005), pp. 2347–2355. DOI: 10.1093/bioinformatics/bti337.
- [Mot+17] Nicolas Mottet, Joaquim Bellmunt, Michel Bolla, Erik Briers, Marcus G. Cumberbatch, Maria De Santis, Nicola Fossati, Tobias Gross, Ann M. Henry, Steven Joniau, Thomas B. Lam, Malcolm D. Mason, Vsevolod B. Matveev, Paul C. Moldovan, Roderick C.N. van den Bergh, Thomas Van den Broeck, Henk G. van der Poel, Theo H. van der Kwast, Olivier Rouvière, Ivo G. Schoots, Thomas Wiegel, and Philip Cornford. “EAU-ESTRO-SIOG Guidelines on Prostate Cancer. Part 1: Screening, Diagnosis, and Local Treatment with Curative Intent”. In: *European Urology* 71.4 (2017), pp. 618–629. ISSN: 0302-2838. DOI: <https://doi.org/10.1016/j.eururo.2016.08.003>. URL: <http://www.sciencedirect.com/science/article/pii/S0302283816304705>.
- [MPH07] Laurens van der Maaten, Eric Postma, and H Herik. “Dimensionality Reduction: A Comparative Review”. In: *Journal of Machine Learning Research - JMLR* 10 (Jan. 2007).
- [NK03] Marcin Novotni and Reinhard Klein. “3D Zernike Descriptors for Content Based Shape Retrieval”. In: *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications*. SM ’03. Seattle, Washington, USA: ACM, 2003, pp. 216–225. ISBN: 1-58113-706-0. DOI: 10.1145/781606.781639.

- [Ola14] Christopher Olah. *Visualizing MNIST: An Exploration of Dimensionality Reduction*. 2014. URL: <https://colah.github.io/posts/2014-10-Visualizing-MNIST/> (visited on 12/30/2017).
- [Pal+08] David Palma, Emily Vollans, Kerry James, Sandy Nakano, Vitali Moiseenko, Richard Shaffer, Michael McKenzie, James Morris, and Karl Otto. “Volumetric Modulated Arc Therapy for Delivery of Prostate Radiotherapy: Comparison With Intensity-Modulated Radiotherapy and Three-Dimensional Conformal Radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics* 72.4 (2008), pp. 996–1001. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2008.02.047>. URL: <http://www.sciencedirect.com/science/article/pii/S0360301608003428>.
- [Pel+14] Roy van Pelt, Rocco Gasteiger, Kai Lawonn, Monique Meuschke, and Bernhard Preim. “Comparative Blood Flow Visualization for Cerebral Aneurysm Treatment Assessment”. In: *Computer Graphics Forum* 33.3 (2014), pp. 131–140. ISSN: 1467-8659. DOI: 10.1111/cgf.12369.
- [Pha+12] Madhura N. Phadke, Lifford Pinto, Femi Alabi, Jonathan Harter, Russell M. Taylor, Xunlei Wu, Hannah Petersen, Steffen A. Bass, and Christopher G. Healey. “Exploring Ensemble Visualization”. In: *Proceedings of SPIE—the International Society for Optical Engineering* 8294.82940B (Feb. 2012). ISSN: 0277-786X. DOI: 10.1117/12.912419. URL: <http://europemc.org/articles/PMC3278305>.
- [PI97] Markus Peura and Jukka Iivarinen. “Efficiency of simple shape descriptors”. In: *Proceedings of the Third International Workshop on Visual Form*. Capri, Italy, May 1997, pp. 443–451.
- [PMK04] Vladimir Pekar, Todd R. McNutt, and Michael R. Kaus. “Automated model-based organ delineation for radiotherapy planning in prostatic region”. In: *International Journal of Radiation Oncology*Biophysics* 60.3 (2004), pp. 973–980. ISSN: 0360-3016. DOI: <https://doi.org/10.1016/j.ijrobp.2004.06.004>. URL: <http://www.sciencedirect.com/science/article/pii/S0360301604009642>.
- [Pol16] Archontis Politis. *Spherical-Harmonic-Transform-JS*. 2016. URL: <https://github.com/polarch/Spherical-Harmonic-Transform-JS> (visited on 02/14/2018).
- [Pre+15] Bernhard Preim, Paul Klemm, Helwig Hauser, Katrin Hegenscheid, Steffen Oeltze, Klaus D. Tönnies, and Henry Völzke. “Visual Analytics of Image-Centric Cohort Studies in Epidemiology”. In: *CoRR* abs/1501.04009 (2015). arXiv: 1501.04009.
- [Rai+16] Renata Raidou, Freek Marcelis, Marcel Breeuwer, Eduard Gröller, Anna Vilanova, and Huub van de Wetering. “Visual Analytics for the Exploration and Assessment of Segmentation Errors”. In: *Eurographics Workshop on Visual Computing for Biology and Medicine* (Sept. 2016), pp. 193–202. URL:

https://www.cg.tuwien.ac.at/research/publications/2016/Groeller_2016_P4/.

- [Reh+13] Andreas Reh, Christian Gusenbauer, Johann Kastner, Eduard Gröller, and Christoph Heinzl. “MObjects—A Novel Method for the Visualization and Interactive Exploration of Defects in Industrial XCT Data”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (Dec. 2013), pp. 2906–2915. ISSN: 1077-2626. DOI: 10.1109/TVCG.2013.177.
- [Rhe00] Penny Rheingans. “Task-based color scale design”. In: *Proceedings of Applied Image and Pattern Recognition '99*. Vol. 3905. SPIE, 2000, pp. 35–43.
- [RS00] Sam T. Roweis and Lawrence K. Saul. “Nonlinear Dimensionality Reduction by Locally Linear Embedding”. In: *Science* 290.5500 (2000), pp. 2323–2326. ISSN: 0036-8075. DOI: 10.1126/science.290.5500.2323.
- [Sam69] John W. Sammon. “A Nonlinear Mapping for Data Structure Analysis”. In: *IEEE Transactions on Computers* C-18.5 (May 1969), pp. 401–409. ISSN: 0018-9340. DOI: 10.1109/T-C.1969.222678.
- [Sch+13] Nicole Schadewaldt, Daniel Bystrov, Torbjørn Vik, Heinrich Schulz, Jochen Peters, Astrid Ruth Franz, C Buerger, and K Bzdusek. “Robust initialization of multi-organ shape models”. In: *MICCAI Challenge Workshop on Segmentation: Algorithms, Theory and Applications*. Vol. 1. 2. 2013, p. 3.
- [Sch+14] Johanna Schmidt, Reinhold Preiner, Thomas Auzinger, Michael Wimmer, Eduard Gröller, and Stefan Bruckner. “YMCA - Your Mesh Comparison Application”. In: *IEEE Visual Analytics Science and Technology*. VAST. Paris, France: IEEE Computer Society, Nov. 2014. DOI: 10.1109/VAST.2014.7042491. URL: <https://www.cg.tuwien.ac.at/research/publications/2014/ymca/>.
- [SGB13] Johanna Schmidt, Eduard Gröller, and Stefan Bruckner. “VAICo: Visual Analysis for Image Comparison”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (Dec. 2013), pp. 2090–2099. ISSN: 1077-2626. DOI: 10.1109/TVCG.2013.213.
- [Shl14] Jonathon Shlens. “A Tutorial on Principal Component Analysis”. In: *CoRR* abs/1404.1100 (2014). arXiv: 1404.1100.
- [Shn96] Ben Shneiderman. “The eyes have it: a task by data type taxonomy for information visualizations”. In: *Proceedings 1996 IEEE Symposium on Visual Languages*. Sept. 1996, pp. 336–343. DOI: 10.1109/VL.1996.545307.
- [Smi+14] Noeska N. Smit, Berend Klein Haneveld, Marius Staring, Elmar Eisemann, Charl P. Botha, and Anna Vilanova. “RegistrationShop: An Interactive 3D Medical Volume Registration System”. In: *Eurographics Workshop on Visual Computing for Biology and Medicine*. Ed. by Ivan Viola, Katja Bühler, and Timo Ropinski. The Eurographics Association, 2014. ISBN: 978-3-905674-62-0. DOI: 10.2312/vcbm.20141193.

- [SMJ17] Rebecca L. Siegel, Kimberly D. Miller, and Ahmedin Jemal. “Cancer statistics, 2017”. In: *CA: A Cancer Journal for Clinicians* 67.1 (2017), pp. 7–30. ISSN: 1542-4863. DOI: 10.3322/caac.21387.
- [STS06] Natascha Sauber, Holger Theisel, and Hans-Peter Seidel. “Multifield-Graphs: An Approach to Visualizing Correlations in Multifield Scalar Data”. In: *IEEE Transactions on Visualization and Computer Graphics* 12.5 (Sept. 2006), pp. 917–924. ISSN: 1077-2626. DOI: 10.1109/TVCG.2006.165.
- [Tai12] Vahid Taimouri. “Shape Analysis in Shape Space”. PhD thesis. Detroit, MI, USA, 2012. ISBN: 978-1-267-11647-5.
- [Tor+11] Thomas Torsney-Weir, Ahmed Saad, Torsten Möller, Hans-Christian Hege, Britta Weber, Jean-Marc Verbavatz, and Steven Bergner. “Tuner: Principled Parameter Finding for Image Segmentation Algorithms Using Visual Response Surface Exploration”. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (Dec. 2011), pp. 1892–1901. ISSN: 1077-2626. DOI: 10.1109/TVCG.2011.248.
- [TV04] Hans Tangelder and Remco Veltkamp. “A survey of content based 3D shape retrieval methods”. In: *Proceedings Shape Modeling Applications, 2004*. June 2004, pp. 145–156. DOI: 10.1109/SMI.2004.1314502.
- [Wei+16] Johannes Weissenböck, Artem Amirkhanov, Eduard Gröller, Johann Kastner, and Christoph Heinzl. “PorosityAnalyzer: Visual analysis and evaluation of segmentation pipelines to determine the porosity in fiber-reinforced polymers”. In: *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. Oct. 2016, pp. 101–110. DOI: 10.1109/VAST.2016.7883516.
- [WT05] Chris Weigle and Russel M. Taylor. “Visualizing intersecting surfaces with nested-surface techniques”. In: *VIS 05. IEEE Visualization, 2005*. Oct. 2005, pp. 503–510. DOI: 10.1109/VISUAL.2005.1532835.
- [WVJ16] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. “How to Use t-SNE Effectively”. In: *Distill* (2016). DOI: 10.23915/distill.00002. URL: <http://distill.pub/2016/misread-tsne>.
- [ZC01] Cha Zhang and Tsuhan Chen. “Efficient feature extraction for 2D/3D objects in mesh representation”. In: *Proceedings 2001 International Conference on Image Processing (Cat. No.01CH37205)*. Vol. 3. 2001, 935–938 vol.3. DOI: 10.1109/ICIP.2001.958278.