

**The Virtual Surgical Pelvis
Anatomy Visualization for Education and Surgical Planning**

Smit, Noeska

DOI

[10.4233/uuid:b065ea24-0fb8-4cab-b427-9612ae6a2113](https://doi.org/10.4233/uuid:b065ea24-0fb8-4cab-b427-9612ae6a2113)

Publication date

2016

Document Version

Final published version

Citation (APA)

Smit, N. (2016). *The Virtual Surgical Pelvis: Anatomy Visualization for Education and Surgical Planning*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:b065ea24-0fb8-4cab-b427-9612ae6a2113>

Important note

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

Copyright


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

Takedown policy

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

The **Virtual Surgical Pelvis**

Anatomy Visualization for Education and Surgical Planning



Noeska Smit

THE VIRTUAL SURGICAL PELVIS:

ANATOMY VISUALIZATION FOR
EDUCATION AND SURGICAL PLANNING

THE VIRTUAL SURGICAL PELVIS:

ANATOMY VISUALIZATION FOR
EDUCATION AND SURGICAL PLANNING

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 31 oktober 2016 om 12:30 uur

door

Noeska Natasja SMIT

Master of Science, Faculteit Elektrotechniek, Wiskunde en Informatica,
Technische Universiteit Delft, Nederland,
geboren te Heenvliet, Nederland.

Dit proefschrift is goedgekeurd door de

promotors: prof. dr. E. Eisemann en prof. dr. M.C. de Ruiter
copromotor: dr. A. Vilanova Bartroli

Samenstelling promotiecommissie:


Rector Magnificus,	voorzitter
Prof. dr. E. Eisemann,	Technische Universiteit Delft
Prof. dr. M.C. de Ruiter,	Leiden Universitair Medisch Centrum
Dr. A. Vilanova Bartroli,	Technische Universiteit Delft

Onafhankelijke leden:

Prof. dr. B.P.F. Lelieveldt,	Technische Universiteit Delft
Prof. dr. C.J.H. van de Velde,	Leids Universitair Medisch Centrum
Prof. dr. A. Ynnerman,	Linköping University
Dr. I. Viola,	Vienna University of Technology
Prof. dr. F.W. Jansen,	Technische Universiteit Delft, reservelid



This work was carried out in the TU Delft Graduate School.



This work was carried out in the ASCI graduate school. ASCI dissertation series number 347.



This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO) and partly funded by the Ministry of Economic Affairs (project number 10903: “High-definition Atlas-based surgical planning for Pelvic Surgery”)

Keywords: Medical visualization, anatomy, education, surgical planning

Printed by: Gildeprint

Front & Back: Cover design by Gerrit Rijken - Os Coxae rendered in VolumeShop

Copyright © 2016 by N.N. Smit

ISBN 978-94-6186-719-3

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

Endure. In enduring, grow strong.

Dak'kon

CONTENTS

Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 Outline of thesis.	3
References	6
I Representation & Construction	9
2 Model-based Visualization of Heterogeneous Anatomy Data	13
2.1 Introduction	14
2.2 Medical Background	15
2.2.1 Computed Tomography (CT)	15
2.2.2 Magnetic Resonance Imaging (MRI)	15
2.2.3 Ultrasound.	16
2.2.4 Modalities from Nuclear Medicine	16
2.3 Related Work	17
2.4 Method	18
2.5 Implementation	20
2.6 Results	20
2.7 Conclusion	23
References	24
3 The Virtual Surgical Pelvis: Atlas Construction	27
3.1 Introduction	28
3.2 Material & Methods	28
3.2.1 Material	28
3.2.2 Methods	29
3.3 Results	30
3.4 Discussion	33
3.5 Conclusion	34
References	36
II Education	39
4 Web-based Anatomy Education	43
4.1 Introduction	44
4.2 Background.	45
4.3 Related Work	45

4.4	The Online Anatomical Human	48
4.4.1	Data	49
4.4.2	Features	49
4.4.3	Interaction	51
4.4.4	Course Setup.	51
4.5	Implementation	52
4.6	Conclusion and Future Work	54
	References	54
5	Visualizing Anatomical Variation in Branching Structures	57
5.1	Introduction	58
5.2	Related Work	59
5.3	VarVis	61
5.3.1	Representations	61
5.3.2	Visualization	66
5.4	Implementation	69
5.5	Results	69
5.6	Evaluation	70
5.6.1	Individual tree view	71
5.6.2	Summary tree	72
5.6.3	Topology similarity graph and matrix	72
5.6.4	Applicability	72
5.6.5	Evaluation conclusion	73
5.7	Conclusion and Future Work	73
	References	74
III	Surgical Planning	77
6	A Survey on Multimodal Medical Data Visualization	81
6.1	Introduction	82
6.2	Medical Background	84
6.3	Clinical Workflow and Requirements	86
6.3.1	Clinical Workflow	86
6.3.2	Requirement Analysis	87
6.4	Visualization Techniques	89
6.4.1	Basic Representation.	89
6.4.2	Smart Visibility.	90
6.4.3	Focus-and-context visualization	93
6.4.4	Summary	95
6.5	Multimodal Rendering and Interaction Techniques	95
6.5.1	Rendering Techniques	96
6.5.2	Interaction Techniques	98

6.6	Applications	99
6.6.1	Medical Research	99
6.6.2	Diagnosis	100
6.6.3	Treatment Planning	102
6.6.4	Discussion	106
6.7	Conclusion and Future Challenges	107
6.8	Summary	109
	References	110
7	RegistrationShop: An Interactive 3D Medical Volume Registration System	123
7.1	Introduction	124
7.2	Related Work	125
7.3	RegistrationShop	126
7.3.1	Visualizations	126
7.3.2	Transformation tools.	129
7.4	Implementation	131
7.5	Evaluation	131
7.5.1	Case study design	132
7.5.2	Case study results	133
7.6	Conclusion and Future Work	135
	References	135
8	PelVis: Atlas-based Surgical Planning for Oncological Pelvic Surgery	139
8.1	Introduction	140
8.2	Related Work	140
8.3	Medical Background	142
8.4	Requirements Analysis	143
8.5	PelVis	144
8.5.1	Preprocessing	145
8.5.2	Visualization	145
8.5.3	Interaction	148
8.6	Implementation	149
8.7	Results	150
8.8	Evaluation	152
8.8.1	Evaluation Setup.	152
8.8.2	Evaluation Results	153
8.8.3	Evaluation Conclusion	155
8.9	Discussion	156
8.10	Conclusion	157
	References	157
9	Conclusion	163
9.1	Contributions	163
9.2	Outlook	165
	References	166

Acknowledgements	167
Curriculum Vitæ	169
List of Publications	171

SUMMARY

THIS thesis deals with visualizing anatomical data for medical education and surgical planning purposes. To this end, we have developed a detailed virtual atlas, the Virtual Surgical Pelvis (VSP), which unifies surgically relevant knowledge on pelvic anatomy. We provide methods to share the knowledge contained in the VSP for educational purposes, and to visualize the VSP in the context of individual patients for pre-operative planning purposes.

Chapter 2 deals with a representation developed to unify spatial and non-spatial anatomical knowledge. Via this representation, it is possible to store, access and visualize these heterogeneous datasets through a shared coordinate system. This allows us to construct the VSP atlas, a process which we describe in detail in Chapter 3, where we also detail the application potential of the VSP. We present several examples of the VSP mapped to clinical pre-operative MRI scans, as examples of how the VSP can be used to enrich clinical data with surgically relevant information that is not available from the scans themselves.

To share the VSP for educational purposes, we present an online tool, the Online Anatomical Human (OAH) in Chapter 4. OAH runs directly in the browser and can be used to explore the complex relation between 2D and 3D anatomy. Furthermore, annotations can be added directly on the 3D structures for quizzing purposes, or to enrich the VSP further with annotations performed by experts. The OAH was successfully deployed in a Massive Open Online Course (MOOC), where thousands of students worldwide used the application to study pelvic anatomy.

While the VSP is based on multiple datasets, it does not include all potential topological anatomical variations in branching structures such as vessels and nerves. Illustrations and text are traditionally used by medical specialists to study these variations, but it is difficult to compare complex variations in such illustrations. Therefore, in Chapter 5 we present an interactive visualization application for anatomical variations, which allows the user to compare and explore variations of branching structures interactively for educational purposes. With methods inspired by graph theory, users can intuitively select groups of variations, based on a similarity measure, and compare local differences.

In Chapter 6, we present a state-of-the-art report on multimodal medical visualization. We describe the basics of medical image acquisition, and the clinical workflow for dealing with such data. We discuss suitable rendering and visualization techniques appropriate for rendering multiple modalities. The core contribution of this work is a taxonomy based on the multimodal medical visualization applications so far, the visualization techniques they employ, and the medical domain context. Additionally, we provide an outlook on open problems and potential future research directions.

To make the VSP patient-specific and to enrich the VSP with more datasets, registration is needed. Unfortunately, current registration software is often difficult to use for non-medical-imaging-experts. In Chapter 7 we present a new registration applica-

tion, RegistrationShop, that allows user to register 3D medical image datasets based on 3D visualizations and simple interactive transformation tools. Based on real-time visual feedback via comparative visualization techniques, users can inspect the current registration result and iteratively improve the alignment. Besides basic interactive transformation tools, we propose a novel way of placing corresponding landmark-pairs in 3D volumes.

After combining the VSP atlas with patient-specific pre-operative MRI scans, we visualize the results in an interactive application for surgical planning aimed at pelvic oncological procedures, entitled PelVis, which is described in Chapter 8. We present visualization methods to represent context, target, and risk structures for surgical planning of the Total Mesorectal Excision (TME) procedure. We employ distance-based and occlusion management techniques to represent the patient-specific pathology and anatomy. Furthermore, we visualize the confidence in the registration outcome in relation to the distance of the target structure to the risk zones.

The research described in this thesis was supported by the Dutch Technology Foundation STW via project 10903: “High-definition Atlas-based surgical planning for Pelvic Surgery”.

SAMENVATTING

DEZE thesis beschrijft de visualisatie van anatomische data voor medische educatie en chirurgische planning. Hiertoe hebben we een virtuele atlas ontwikkeld, de Virtual Surgical Pelvis, die alle chirurgisch relevante kennis over de anatomie van het bekken verenigd. We stellen methodes voor om de kennis in de VSP te delen voor educatieve doeleinden en om de VSP te visualiseren in een patient-specifieke context voor pre-operatieve planning.

Hoofdstuk 2 beschrijft een representatie die ontwikkeld is om ruimtelijke en niet-ruimtelijke anatomische kennis samen te brengen. Met deze representatie is het mogelijk om deze heterogene datasets op te slaan, te laden en te visualiseren. Dit stelt ons in staat om de VSP atlas te creëren. Dit proces is beschreven in hoofdstuk 3. We laten verschillende voorbeelden zien waarin de VSP geregistreerd wordt met klinische data om deze te verrijken met chirurgisch relevante informatie die niet in de scans zelf zichtbaar is.

Om de VSP te kunnen delen voor educatieve doeleinden, hebben we een online applicatie ontwikkeld, de Online Anatomical Human (OAH), beschreven in hoofdstuk 4. OAH kan direct vanuit de browser gebruikt worden om de complexe relatie tussen 2D en 3D anatomie te bestuderen. Hiernaast kunnen ook annotaties worden toegevoegd op de 3D structuren. Dit stelt de gebruikers in staat om vragen over specifieke structuren te stellen of de VSP verder te verrijken met extra informatie. De OAH is succesvol ingezet in een zogenaamde Massive Open Online Course (MOOC), waar duizenden studenten wereldwijd de applicatie gebruiken om bekken anatomie te bestuderen.

Ondanks dat de VSP is gebaseerd op meerdere datasets, omvat deze niet alle mogelijke topologische anatomische variaties in vertakkende structuren. Vaak worden illustraties, al dan niet gecombineerd met tekstuele beschrijvingen, ingezet om deze variaties te bestuderen. Het is in deze illustraties lastig om complexe variaties met elkaar te vergelijken. Om deze situatie te verbeteren, presenteren we in hoofdstuk 5 een interactieve visualisatie applicatie om anatomische variaties te bestuderen. Hierin kunnen gebruikers interactief verschillende variaties in vertakkende structuren verkennen en vergelijken. De methodes die we hiervoor ontwikkelden zijn geïnspireerd door grafentheorie en stellen de gebruikers in staat om groepen variaties te selecteren, gebaseerd op een similarity measure, alsmede lokale verschillen te vergelijken.

In hoofdstuk 6 presenteren we een state-of-the-art report over multimodale medische visualisatie. We beschrijven hoe medische beelden verkregen worden en hoe deze data in de kliniek wordt bekeken. Hiernaast beschrijven we geschikte rendering en visualisatie technieken om deze multimodale datasets weer te geven. De hoofdbijdrage van dit werk is een taxonomie van de multimodale medische visualisatie applicaties tot dusver, de gebruikte visualisatie technieken en in welk medisch domein ze worden toegepast.

Om de VSP patient-specifiek te maken en om de VSP te verrijken met meer datasets, is registratie nodig. Helaas is de huidige registratiesoftware gericht op registratie experts en minder geschikt voor andere gebruikers. In hoofdstuk 7 beschrijven we een nieuwe registratie applicatie, RegistrationShop, die gebruikers ertoe in staat stelt 3D medische datasets te registreren door middel van 3D visualisaties en simpele transformatie tools. Gebruikers kunnen hierin gemakkelijk het huidige registratie-resultaat beoordelen door real-time visuele feedback en deze iteratief verbeteren. Naast simpele interactieve transformatie tools, bieden we een nieuwe manier om paarsgewijs corresponderende landmarks te plaatsen in 3D volumes.

Nadat de VSP gecombineerd is met patient-specifieke pre-operatieve MRI scans, visualiseren we de resultaten in een interactieve applicatie voor chirurgische planning, gericht op oncologische procedures in de pelvis, genaamd PelVis, beschreven in hoofdstuk 8. We presenteren visualisatie technieken voor context, doel en risico structuren voor chirurgische planning van de Total Mesorectal Excision (TME) procedure. We passen verschillende visualisatie-technieken toe, zoals occlusion management en weergave van afstanden, om patient-specifieke pathologie en anatomie te presenteren. Hiernaast visualiseren we het vertrouwen in het registratie resultaat in relatie tot de invloed op de afstand van de doel structuur tot de risico-gebieden.

Het onderzoek beschreven in deze thesis is ondersteund door de Technologiestichting STW door middel van project 10903: “High-definition Atlas-based surgical planning for Pelvic Surgery”.

1

INTRODUCTION

C OLORECTAL cancer is the third most frequent cancer diagnosed in men (after lung and prostate cancer) and the second most frequent in women (after breast cancer). Rectal cancer treatment is a multi-disciplinary treatment with standardized surgical, pathological and radiotherapeutical procedures. Improvement of the surgical techniques has been a key factor in the progression of rectal cancer treatment over the last 100 years. The current incidence rate is increasing mainly due to increasing age of the population. Surgical resection of the rectum with a Total Mesorectal Excision procedure, or TME, is the predominant surgical treatment option for rectal cancer. TME is the complete removal of the rectum including the surrounding mesorectal fat and connective tissue. Introduction of the TME technique and a combination with radiotherapy has reduced local recurrence rates from +/-20 to 4 percent. However, the anatomy of the bony pelvis and its position within the body hamper a good view on the internal organs (e.g., rectum, uterus, vagina, prostate and bladder) and their very closely related but difficult to detect nerves, lymphatics and vessels during surgery. Although world-wide TME training programs pay special attention to nerve sparing techniques, clinical studies still report a high incidence of pelvic organ dysfunction (urinary incontinence in 34%, fecal incontinence in 39% and sexual dysfunction in up to 79% of the cases) [1]. From the Dutch TME trial we learned that poor functional outcome after surgery is mainly caused by surgical damage to the nervous system [2]. These nerves are often damaged, because they cannot be seen by the surgeon in pre-operative MRI scans or during surgery due to their small size. Furthermore, there was a lack of knowledge and consensus on the exact location of these nerves due to the complex 3D anatomy of the region [3]. Therefore, knowledge of the complicated courses and positions of the crucial nerves, i.e., superior and inferior hypogastric nerves, pelvic splanchnic nerves, fascia sheets [4, 5], and the recently rediscovered levator ani nerve, with regard to the surgical approach are essential in order to increase the quality of life after surgery.

The work presented in this thesis is of an interdisciplinary nature, resulting from close collaboration between anatomists, surgeons and visualization researchers. The first goal of this thesis is to address the lack of knowledge about these nerves by unifying existing and newly discovered surgically relevant anatomical knowledge. To this end, we developed an integrated histology-based model of the pelvis, the Virtual Surgical Pelvis (VSP), which features a larger amount of surgically relevant detail than previous models until now [6]. Our basic atlas is based on high resolution histological sections and can be enriched with additional metadata, e.g., hyperlinks to relevant articles, histology data, and relational information. Furthermore, the VSP can be combined with MRI data, enabling us to include all soft tissue structures, including smaller nerves and vessels, that are difficult or impossible to detect in MRI data alone. It is then possible to perform pre-operative planning of the Total Mesorectal Excision surgical procedure, based on the knowledge regarding the nerves available from the atlas.

A secondary goal is share the knowledge gathered in the VSP for educational purposes in an intuitive way. Previous educational applications provide an artistic general impression of anatomy and lack a link to 2D anatomy. We opt for sharing the information via the web, so that it is available to students and medical experts worldwide. Additionally, by allowing medical experts to share their knowledge we can enrich our database

faster.

The final goal of this thesis is to improve pre-operative planning and surgical training for the Total Mesorectal Excision surgical procedure. In order to do this, we developed the necessary techniques to fit the VSP to a specific patient by making use of pre-operative MRI data. For this, registration of the atlas to patient-specific MRI scans is needed. Combining the registered atlas and MRI scan in a surgical planning application will enable the surgeon to increase their awareness of the patient-specific anatomy and location of the nerves, thus potentially improving post-operative outcome.

1.1. OUTLINE OF THESIS

This thesis is divided in three main parts, dealing with gathering and representing anatomical knowledge, education, and surgical planning respectively (see Figure 1.1 for a visual overview). Due to the fact that the topics presented in this thesis are multifaceted, there is no global related work chapter for the whole thesis. Instead, for the sake of readability, each of the core chapters is based on one or more publications and provides their own related work section.

Part 1 focuses on the representation and construction of the Virtual Surgical Pelvis.

First of all, to acquire the anatomical level of detail required to be useful in surgical planning, multiple heterogeneous datasets need to be combined together in a shared coordinate system. This allows multiple modalities, each providing unique anatomical details, to be merged in a single 3D model. Hereby, available data can be stored, queried and visualized in a single unified system. To make this feasible, we developed the **Unified Anatomical Human (UAH) representation** [7, 8], described in Chapter 2, that forms the foundations for storing, querying and visualizing heterogeneous anatomical data in a unified way. It allows us to construct a detailed model and to map it to patient-specific scans, providing the required anatomical detail for surgical planning. 3D models of the pelvis so far lacked surgically relevant anatomical detail [6]. Therefore, we built a new atlas model that does contain the desired level of detail. The acquisition and construction of our **Virtual Surgical Pelvis** [9] atlas is described in Chapter 3.

In **Part 2** we focus on the educational contributions of this work. Even without making the VSP patient-specific, the atlas itself already has potential to help achieve educational goals aimed at medical students and surgeons.

In order to be of real clinical use, the 3D model needs to include large amounts of detailed and correct anatomical information. This anatomical information needs to be added to our UAH database by medical domain experts. Since this process is important, but very time-consuming, we would like to allow the medical community to share their knowledge via the model by using a web-interface. This software will allow users to add valuable information by using 'crowd-sourcing' and interaction techniques to enrich the existing database of the UAH project via the Internet. Preliminary work on visualizing and interacting with the 3D model from the database for **web-based anatomy education** [10] is presented in Chapter 4. The system presented was successfully deployed in a Massive Open Online Course (MOOC) on pelvic anatomy to educate students worldwide in the complex 3D pelvic anatomy and the link to 2D images.

Another important aspect of this project is modeling and visualizing anatomical variability. No two people are exactly anatomically alike and therefore also in patients anatom-

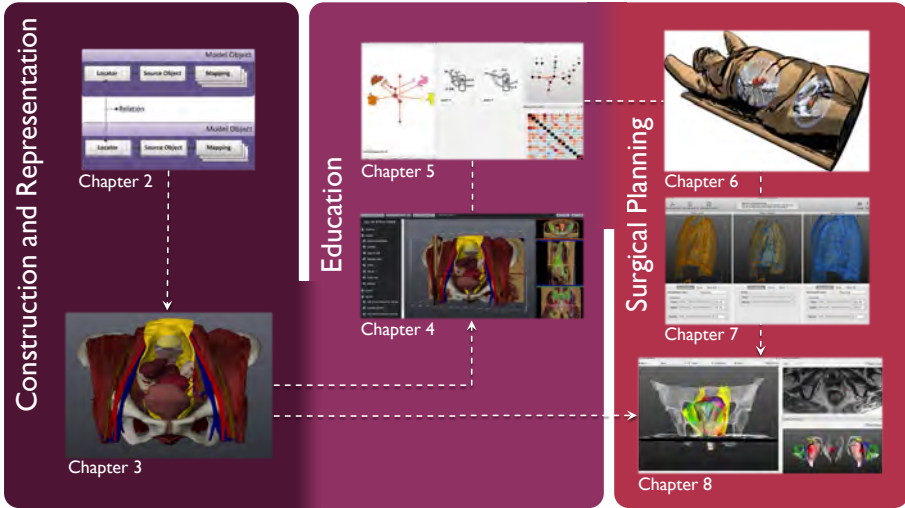


Figure 1.1: A visual overview of the thesis organization into parts and chapters.

ical variations are expected. These variations are not just limited to size and shape, but also include different nerve and vessel layouts. The exact location of critical anatomical structures such as nerves are highly important, but unfortunately impossible to detect in standard pre-operative MRI-scans. Therefore our detailed 3D model needs to include these potential variations and visualize them in such a way that the types of variations and their probabilities are immediately apparent. As an initial step, we focused on **visualizing anatomical variability in branching structures** [11] presented in Chapter 5. This technique builds upon graph matching techniques, which are used to provide an overview of similarities and dissimilarities between variations, presented in an interactive network visualization.

Part 3 presents the contributions related to surgical planning. For the surgical planning component of this thesis, we need to visualize both VSP atlas and MRI information together. This generates visualization challenges similar to those that arise when visualizing multiple modalities together. Therefore, Part 3 starts with a survey of current **State-of-the-Art on multimodal medical visualization** [12] in Chapter 6.

The first challenge in surgical planning lies in the registration process. To map the model information to patient-specific scans, registration needs to be performed. This registration can be done by registering the VSP atlas, which is based on a manual segmentation of a Visible Human dataset by an anatomical expert, with a patient-specific scan. By using the transformation that is obtained by registering the atlas with the patient-specific scan, model information from the atlas can be mapped to the scan. This model information can make structures that are not visible in the original scan, such as nerves, visible in the context of the individual patient. Ideally, an end-user such as a radiographer could easily complete the registration process by applying a mixture of existing techniques such as rigid, deformable, hybrid or articulated registration through direct vi-

sual interaction with the volumes. To this end, **RegistrationShop** [13] was designed and implemented as an interactive system for registration. RegistrationShop is presented in Chapter 7. The focus here is not on creating new registration techniques, but on combining existing techniques to improve user interaction and visual feedback. RegistrationShop is not only an indispensable tool for the VSP project, but has potential to be used in any application that requires registration, such as combining different imaging modalities of the same patient or comparing scans before and after treatment. RegistrationShop was made freely available as an Open Source project.

The final part of the project involves taking all the individual, but closely related, subprojects and bringing them together in the first prototype of a **model-based surgical planning** [14] tool presented in Chapter 8. By registering the VSP model atlas to a pre-operative patient-specific scan using RegistrationShop, a patient-specific model of the pelvis can be built. Once this registration is complete, the surgeon can prepare for a procedure before going into the operating room using our surgical planning prototype application, PelVis. In PelVis, the anatomical context, target and risk structures are visualized using illustrative techniques to convey distance information, and provide spatial context, while limiting occlusion. The software also has potential to be used in surgical training, by allowing for different pathologies to be represented in several scenarios. The focus of this tool is on the TME, but it is easily extendable to other procedures of interest.

Finally, Chapter 9 presents conclusions and an outlook on the future.

The research described in this thesis was funded by the STW: a technology foundation that realizes the transfer of knowledge between technical sciences and users, by bringing researchers and users together. This project was carried out as a collaboration between the department of Anatomy and Embryology and Surgery at the LUMC and the Computer Graphics and Visualization group at the TU Delft.

REFERENCES

- [1] M. Lange, C. Marijnen, C. Maas, H. Putter, H. Rutten, A. Stiggelbout, E. M.-K. Krannenbarg, C. van de Velde, C. C. I. of the Dutch, *et al.*, *Risk factors for sexual dysfunction after rectal cancer treatment*, *European Journal of Cancer* **45**, 1578 (2009).
- [2] C. Wallner, M. M. Lange, B. A. Bonsing, C. P. Maas, C. N. Wallace, N. F. Dabhoiwala, H. J. Rutten, W. H. Lamers, M. C. DeRuiter, and C. J. van de Velde, *Causes of fecal and urinary incontinence after total mesorectal excision for rectal cancer based on cadaveric surgery: a study from the cooperative clinical investigators of the dutch total mesorectal excision trial*, *Journal of clinical oncology* **26**, 4466 (2008).
- [3] A. C. Kraima *et al.*, *New insights in pelvic anatomy and implications for radical nerve-sparing pelvic surgery. From microscopy to a virtual 3D anatomical atlas*, Ph.D. thesis, Department of Anatomy and Embryology, Faculty of Medicine, Leiden University Medical Center (LUMC), Leiden University (2015).
- [4] A. Kraima, N. West, D. Treanor, D. Magee, R. Bleys, H. Rutten, C. van de Velde, P. Quirke, and M. DeRuiter, *Understanding the surgical pitfalls in total mesorectal excision: Investigating the histology of the perirectal fascia and the pelvic autonomic nerves*, *European Journal of Surgical Oncology (EJSO)* **41**, 1621 (2015).
- [5] A. Kraima, N. West, D. Treanor, D. Magee, H. Rutten, P. Quirke, M. DeRuiter, and C. Van De Velde, *Whole mount microscopic sections reveal that denonvilliers' fascia is one entity and adherent to the mesorectal fascia; implications for the anterior plane in total mesorectal excision?* *European Journal of Surgical Oncology (EJSO)* **41**, 738 (2015).
- [6] A. C. Kraima, N. N. Smit, D. Jansma, C. Wallner, R. Bleys, C. Van De Velde, C. P. Botha, and M. C. DeRuiter, *Toward a highly-detailed 3d pelvic model: Approaching an ultra-specific level for surgical simulation and anatomical education*, *Clinical Anatomy* **26**, 333 (2013).
- [7] N. Smit, A. Kraima, D. Jansma, M. deRuiter, and C. Botha, *A unified representation for the model-based visualization of heterogeneous anatomy data*, in *EuroVis - Short Papers* (Eurographics Association, 2012) pp. 85–89.
- [8] N. Smit, A. Kraima, D. Jansma, M. DeRuiter, and C. Botha, *The unified anatomical human (beta): Model-based representation of heterogeneous anatomical data*, in *Workshop 3D Physiological Human (3DPH), CASA* (2012).
- [9] A. C. Kraima, N. N. Smit, D. Jansma, E. Eisemann, M. Chung, N. West, P. Quirke, H. Rutten, C. Van De Velde, A. Vilanova, and M. C. DeRuiter, *A novel 3d virtual anatomical atlas: taking the next step in enhancing surgical anatomical education and patient-specific clinical guidance*, *Anatomical Sciences Education* (2016), in Submission.
- [10] N. N. Smit, C.-W. Hofstede, A. C. Kraima, D. Jansma, M. C. Deruiter, and E. Eisemann, *The online anatomical human: Web-based anatomy education*, in *Proceedings of Eurographics - Education papers* (2016) accepted, to appear.

- [11] N. N. Smit, A. C. Kraima, D. Jansma, M. C. Deruiter, E. Eisemann, and A. Vilanova, *Varvis: Visualizing anatomical variation in branching structures*, in *Proceedings of EuroVis - Short Papers* (2016) accepted, to appear.
- [12] K. Lawonn, N. N. Smit, K. Bühler, and B. Preim, *A survey on multimodal medical data visualization*, *Computer Graphics Forum* (2016), in Submission.
- [13] N. N. Smit, B. K. Haneveld, M. Staring, E. Eisemann, C. P. Botha, and A. Vilanova, *Registrationshop: An interactive 3d medical volume registration system*, *EG Workshop on Visual Computing in Biology and Medicine (VCBM)*, (2014).
- [14] N. Smit, K. Lawonn, A. Kraima, M. DeRuiter, H. Sokooti, S. Bruckner, E. Eisemann, and A. Vilanova, *Pelvis: Atlas-based surgical planning for oncological pelvic surgery*, *IEEE Transactions on Visualization and Computer Graphics* (Proceedings of Scientific Visualization 2016) **23** (2017), accepted to appear.

I

REPRESENTATION & CONSTRUCTION

THIS part of the thesis deals with the first goal of the work: unifying existing surgically relevant anatomical knowledge. To this end, we have developed a representation that allows multiple heterogeneous data sources to be combined in a single atlas, which we describe in Chapter 2. Using this representation, we are able to visualize an anatomical level of detail revealing information that is not available from clinical medical imaging acquisition, such as the region containing the autonomic nerves.

Based on this representation, we constructed our pelvic atlas, the Virtual Surgical Pelvis (VSP). The construction details and data sources for the VSP are described in Chapter 3. The VSP is applicable both for medical education as well as surgical planning. In this chapter, we also describe results from registering the VSP atlas to five patient-specific pre-operative MRI scans.

2

MODEL-BASED VISUALIZATION OF HETEROGENEOUS ANATOMY DATA

In the course of anatomical research, anatomists acquire and attempt to organize a great deal of heterogeneous data from different sources, such as MRI and CT data, cryosections, immunohistochemistry, manual and automatic segmentations of various structures, related literature, the relations between all of these items, and so forth. Anatomical variation between subjects further complicates this organization. Currently, there is no way of storing, accessing and visualizing these heterogeneous datasets in an integrated fashion. Such capabilities would have great potential to empower anatomy research. In this chapter, we present methods for the integration of heterogeneous spatial and non-spatial data from different sources, as well as the complex relations between them, into a single model, all using standardized anatomical coordinates. All captured data can then be interactively visualized in various ways, depending on the anatomical question. Furthermore, our model enables data to be queried both structurally, i.e., relative to existing anatomical structures, and spatially, i.e., with anatomical coordinates. When new patient-specific medical scans are added to the model, all available model information can be mapped to them. Using this mapping, model information can be transferred back to the new scans, thus enabling the creation of visualizations enriched with information not available in the scans themselves. While the presented methods in this chapter are generic, we present several examples of results focused on the pelvic area.

This chapter is based on three publications: **Noeska Smit**, Annelot Kraima, Marco DeRuiter, Daniel Jansma and Charl Botha, "A Unified Representation for the Model-based Visualization of Heterogeneous Anatomy Data", *Proceedings of EuroVis 2012 Short Papers* [1] and **Noeska Smit**, Annelot Kraima, Daniel Jansma, Marco DeRuiter, and Charl Botha, "The unified anatomical human (beta): Model-based representation of heterogeneous anatomical data", *Proceedings of the 3D Physiological Human (3DPH) Workshop, CASA. 2012* [2]. Kai Lawonn, **Noeska Smit**, Katja Bühler, and Bernhard Preim, "A Survey on Multimodal Medical Data Visualization", In Submission, *Computer Graphics Forum*, 2016 [3]

2.1. INTRODUCTION

Anatomical knowledge is not only an important part of medical education, but also of great relevance in daily practice. In surgery, for instance, knowing the exact location of nerves and arteries can be of paramount importance in improving surgical outcome. Anatomists currently have no comprehensive and intuitive way of storing and sharing the knowledge, both spatial and non-spatial, they possess with medical professionals. Examples of spatial volumetric data they work with are cryosectional slices, CT scans, MRI scans and histological slices. This however can be enriched by data that is inherently non-spatial, such as the related literature, knowledge about anatomical topologies, anatomical systems and the relationships between structures. There was clearly a need for a system that can integrate all of these different heterogeneous data types into a unified model. Such a model by itself already has potential for anatomy education, which we describe further in Chapter 4 of this thesis. Once the information is in a shared coordinate system, model information can then also be mapped to patient-specific scans. All available knowledge can then be visualized in the anatomical context of a specific patient for surgical planning purposes, which we describe further in Chapter 8.

With this work, we present a model-based representation for the storage, flexible querying and visualization of heterogeneous anatomical data. Our approach is based on a standardized coordinate system of the human body to which arbitrary anatomical datasets, both spatial and non-spatial, can be associated. One of the unique aspects of our method, is that we perform lazy normalization. In other words, datasets are stored in their raw form, enriched with a locator and a number of mappings. The locator enables us to perform spatial indexing, whilst each mapping describes a different task-specific transformation from the raw dataset space to the standardized anatomical space. In addition, our approach enables an arbitrary number of overlapping, differently sampled, multi-modal datasets. Besides storage, querying and visualization, our pipeline enables the task-specific mapping of model-based information onto new patient-specific datasets, also creating possibilities for surgical planning and guidance.

For the system we envision, the technical requirements are the following:

1. The system needs to be able to **store** arbitrary types of heterogeneous anatomical information and should be able to handle spatial data in arbitrary resolutions and spacings.
2. The system needs to allow the user to **query** stored information per topic of interest and all information in the system should be queryable spatially in a shared coordinate system.
3. The system needs to be able to **visualize** all available information relevant to the user's interest in anatomical spatial context, and the relation to the other available representations should be clearly defined.

The system that we designed provides a way of integrating all anatomical knowledge in one model. With this, our contributions in this Chapter are the following:

- We present a novel generic model to store all anatomical information (both spatial and non-spatial). The use of standardized anatomical coordinates enables

domain-specific queries with spatial/visual querying due to the use of a schema-less database and kd-tree.

- Furthermore we describe freeform relations that can be defined to represent any type of relationship between model objects, capturing not only the data itself, but also the connectivity of the data.

2.2. MEDICAL BACKGROUND

In this section, we provide a brief overview of the most commonly used medical imaging modalities that provide spatial data. The characteristics of each of the modalities involved are summarized, and their advantages and disadvantages and associated visualization challenges are discussed.

2.2.1. COMPUTED TOMOGRAPHY (CT)

CT is an x-ray based tomographic imaging technique that creates stacks of 2D cross sectional images that can be combined to 3D volumes. It is especially suitable to distinguish tissues such as bone, water, fat, and the air in the lungs. A contrast agent can be applied to enhance vascular structures. Recently, hybrid scanners such as dual source, or dual energy CT scanners are available delivering a final image, which fuses information from high and low voltage image acquisition performed at the same time [4]. Depending on the chosen imaging protocol this technique allows for a differentiation of structures like bone and contrast-enhanced blood vessels.

In 2007 approximately 70 million scans were made in the US alone [5], illustrating the relevance of this imaging method in clinical practice. CT data is especially suited for high quality direct volume rendering, due to its high resolution (in general 512×512 in slice resolution and 0.3-2 mm slice distance), high signal-to-noise ratio and standardized intensity values (Hounsfield Units), allowing the definition of re-usable and task specific transfer functions.

2.2.2. MAGNETIC RESONANCE IMAGING (MRI)

In Magnetic Resonance Imaging (MRI), a scan is made using a powerful magnetic field. In contrast to CT, MRI scanners are highly configurable and provide a large variety of imaging protocols allowing to capture structural as well as functional information. In general, several different MRI sequences, such as T1 and T2-weighted scans, are acquired at the same time, leading to (more or less) co-registered images. Since signal strength relies on the presence of Hydrogen atoms, MRI sequences are emphasizing image contrast in tissue containing water. Intensities in MR images are not standardized. MRI data often exhibit an inhomogeneous gray level distribution, requiring careful preprocessing of the data, and intensity values vary depending on scanner vendor and clinic. Therefore MRI data is challenging to visualize.

MRI data can not be easily classified, i.e. a simple 1-D transfer function does often not deliver a clear separation of structures due to the inhomogeneities in the magnetic field. Furthermore, due to the unpredictability of the intensity values, transfer functions are not directly applicable across several datasets without dynamic adaptation [6]. Finally, MR images have generally a lower resolution and lower signal-to-noise ratio than

CT images.

Besides, the standard MRI scanning protocols, there are specific MRI sequences and protocols, such as MRSI, DCE-MRI and DTI. In MRSI, spatially localized metabolites in body tissues are measured. DCE-MRI (Dynamic Contrast-Enhanced MRI) is a perfusion imaging technique that measures the perfusion of tissues by blood indicating regions damaged by stroke or infarction as well as characterizing the vascularization of tumors, helping to assess whether they are benign [7]. DTI is an extension of Diffusion Weighted Imaging (DWI), that detects the direction of white matter tracts in the brain, which represent connectivity between different areas of gray matter. DTI is used in clinical practice to assess the deformation of white matter by tumors, neurosurgical planning and for (early) diagnosis of brain pathologies such as Alzheimer disease, schizophrenia and multiple sclerosis [8]. DTI data is often visualized as a scalar field consisting of the Fractional Anisotropy (FA), using glyphs or fiber tracking [9].

Functional MRI (fMRI): Besides using MRI to acquire structural information, it is possible to detect neural activity. fMRI records subtle changes in blood flow in response to stimuli or actions and uses this information to visualize cortical activity. The most frequently employed technique is blood oxygenation level dependent (BOLD) fMRI. By having the subject perform tasks categorized into visual, motor, speech or memory tasks, different functional areas of the brain 'light up' and can be associated to the tasks performed. Additionally, fMRI is used in a research context to improve the understanding of neural networks in the brain even when the user has no task, as is the case in resting state fMRI [10].

2.2.3. ULTRASOUND

In medical ultrasound, high-frequency sound waves are employed to characterize tissue. Ultrasound can be used both for diagnosis and to guide interventional therapeutic procedures. Due to the nature of the modality, ultrasound is suitable to image soft tissues, such as tendons, vessels and organs, but cannot visualize bone and air, or structures lying underneath these tissue types. Based on the Doppler effect, blood flow in the heart and blood vessels can be detected. Advantages of ultrasound compared to other modalities are that it is cheap, safe, portable and real-time. However, ultrasound is difficult to interpret, due to the low signal-to-noise ratio, artifacts, and the limited field of view. Recent advances in ultrasound technology include 3D ultrasound [11], elastography and contrast-enhanced ultrasound using microbubbles.

2.2.4. MODALITIES FROM NUCLEAR MEDICINE

Positron Emission Tomography (PET) relies on the indirect detection of gamma rays emitted by the patient after a positron-emitting radionuclide (tracer) is introduced. Using this technique, metabolic uptake processes can be visualized. While CT and MRI scans can provide detailed anatomical data, PET scans are able to reveal functional information, such as metabolism. A common application of PET scans is to search for metastases and for this, the radioactive substance fluorodeoxyglucose (FDG) tracer is used, a substance similar to glucose. The metastases have higher glucose uptake than normal and specific abnormal metabolic activity can be captured in this way. Besides oncological applications, PET is also used for neurological and cardiological diagnos-

tic purposes. While commonly used, FDG is not the only available tracer for PET and different tracers may be better suited for specific applications [12].

The PET data needs to be attenuation-corrected before visualization. Visualizing PET data in 3D is challenging, since normal metabolic information is also contained in PET, the highest activity measures are not always the most interesting.

Single-photon emission computed tomography (SPECT) is a nuclear medicine tomographic imaging technique that uses radioactive tracer material to detect gamma rays. In this way, it is similar to PET, however, in contrast to PET, the gamma radiation is measured directly. SPECT can be used for oncological diagnosis, but also for infection, thyroid or bone imaging. Besides these applications, SPECT can also provide localized function within organs for functional cardiac or brain imaging. SPECT suffers from a lower spatial resolution and contrast than PET [13]. Similar to PET, SPECT data is not straight-forward to render in 3D.

2.3. RELATED WORK

In this section, we first discuss the body of work produced by the VOXEL-MAN project over the past decades. The VOXEL-MAN research focused on the storage, query and visualization of anatomical data and has made a significant impact on the field. We also compactly discuss the BrainGazer project, as it also shows a number of similarities with the research presented in this paper.

The VOXEL-MAN project, started in 1985 in Germany by the research group led by Professor Karl Heinz Höhne, has made great progress in combining spatial models with symbolic descriptions. Using the Visible Human dataset combined with segmentation and visualization techniques, the group created an anatomical atlas, combining anatomy, function and radiological appearance.

In 1993 Tiede et al. created a 3D anatomical atlas of the human skull and brain [14]. In the next year Pommert et al. defined several concepts for structuring anatomical information in a semantic network model [15]. By assigning an anatomical structure to every voxel in a 3D volume and connecting the anatomical knowledge base, the group created a medical education tool [16]. In 1995, the group presented the intelligent volume approach, by combining attribute volumes as big as the image volume, for instance one for blood supply, one for function volume, and one for morphology volume [17]. In the same paper, they also demonstrated that it was possible to derive X-ray projections from available CT-data.

The group also did some work on high quality rendering techniques for attributed volume data at subvoxel precision [18]. The group created a high-resolution spatial/symbolic model of the inner organs based on the Visible Human Project and presented a segmentation tool in color space [19, 20]. Furthermore, an interactive atlas of the hand was presented where the group came across a number of limitations of the Visible Human dataset: cryosection cutting artifacts, poor CT quality and lack of visibility of small vessels and nerves [21]. In 2006 Pommert et al. presented their work on using the VOXEL-MAN model for simulation of surgical procedures [22].

Our approach extends the work done by the VOXEL-MAN Group in the following ways. First of all, our model allows for multiple anatomical structures to be defined at any point in model space using information from various different data sources. Sec-

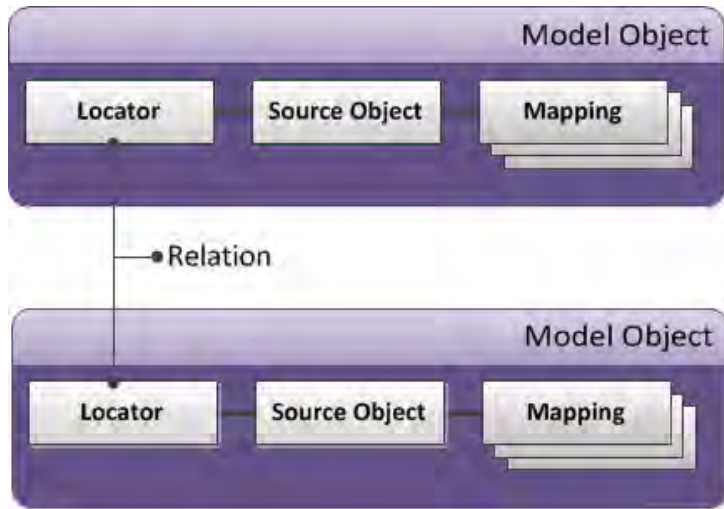


Figure 2.1: A raw dataset (a source object) can be added to the anatomical model by enriching it with a locator, used in spatial indexing, and any number of mappings to the standardized coordinate system. Arbitrary relations between datasets can be described as well.

only, the VOXEL-MAN model represents a single general anatomy, while our model is enriched by the anatomical information from multiple datasets. For this reason, the VOXEL-MAN group had no way to describe interindividual variations, age variations or “fuzzy” anatomical object boundaries [17].

In the BrainGazer project by Brucker et al., visual queries for neurobiological research are introduced [23]. The BrainGazer system uses large databases of transgenic specimens and the acquisition of confocal microscope images of fruit fly brains in which distinct neuronal types are highlighted together with annotated anatomical structures to enable neurobiologists to query this data both visually and through the database interface. The research presented in this chapter differs from that of the BrainGazer research, in that our model needs to support a number of different modalities with significantly differing sampling resolutions and strategies, and that it also needs to cope with the storage of pristine data sources, each packaged with a number of different task-specific spatial transformations. This last characteristic is in fact one of the main factors differentiating our work from similar research.

2.4. METHOD

The method we propose is a model-based representation for storage, querying and visualization of heterogeneous anatomy data. Using an anatomical standardized coordinate system, this system enables users to integrate arbitrary anatomical data into a single unified model. Figure 2.1 shows the primary concepts used in our system. The model is designed in such a way that it forms a solid foundation for further development of anatomical and surgical applications.

Source objects represent original unprocessed information that the user would like

to add to the model. Relevant anatomical knowledge can occur in various forms, such as cryosectional slices, CT scans, MRI scans, histological slices, anatomical structure names and related scientific literature. All these different types of knowledge can be divided in two categories. The first category of source objects are those that have an inherent geometry. The spatial data types include acquisitions from medical imaging devices that can be acquired in vivo, such as MRI-scans, CT-scans and PET-scans. Other examples of spatial data include cryosectional slices and histological images. The second category of source objects is those that do not have an inherent geometry. Examples of this include anatomical terms, literature that is deemed relevant to a certain anatomical structure, statistics, and bio-mechanical tissue characteristics. A special feature of these types of data is that even though they do not have an inherent spatial component themselves, they can be spatially embedded in model space through their relations with other model objects.

Source objects are added to the model by augmenting them with a locator and one or more task-specific mappings to model space. The combination of the source object, its locator and its mappings is then called a **model object**. Once a dataset becomes a model object, it becomes a part of the standardized anatomical coordinate system. This means that it can be queried and visualized in the same space as all other model objects. Once a source object becomes part of model space, when queries are executed at a certain point or even a region, this added source object will show up in the query results. This is possible because the locator describes the spatial extent of a model object. Using one of the mappings that was added, the source object can be transformed to model space and visualized together with other model objects of interest in the standardized coordinate system. The standardized coordinate system has its origin in the sacral promontory. This is a bony anatomical landmark that can easily be found in any patient scan that includes the pelvis. Another benefit of choosing this point is that it is independent of patient pose and central in the human body. The axis are defined in standard anatomical pose. The z-axis points in the cranial or superior direction, the x-axis points to the left hand side of the patient and the y-axis points forward to the anterior or front of the body. Using the sacral promontory as the origin, any point in the human body, arranged in a standard pose, can be intuitively defined with respect to this point.

In order to be able to represent source objects added to the model, **mappings** need to be added. These mappings consist of the transformations that map a source object to model space. The transformations required to do this are acquired during a registration process and can be rigid, affine, deformable or hybrid, for example articulated registration [24] (see Table 2.1). Because of the different modalities available, registration of the different source objects is no easy task. Furthermore, inter-patient variability further complicates the process. Initially, we used a mixture of tools such as MITK, Elastix and 3DSlicer, to create gold standard mappings interactively. Later, we developed our own registration software, which we describe in Chapter 7.

The **locators** are used to define where in model space the model objects are defined. Depending on the data type, the locator can be defined in several forms (see Table 2.2). Specifically, the locator can represent a point set, a volume or a non-geometric model object. In the case of a volume locator, the origin, extent and spacing are stored in order to be able to check if a volume is available at any given point in model space. The

Mapping	Category	Type
rigid	linear	translation and rotation
affine	linear	translation, rotation, shearing, scaling
nonrigid	non-linear	deformable
hybrid	non-linear	combination of the above

Table 2.1: Currently defined mapping types.

Locator	Category	Information
points	spatial	set of points
volume	spatial	origin, extent, spacing
non-geometric	non-spatial	-

Table 2.2: Supported locator types.

pointsets are used as a spatial index for fast spatial querying, by storing them as kd-trees for instance. For non-geometric model objects, the locators do not store any extra information, but are used in defining relations between model objects.

Relations are always defined between locators. They link the model objects together through their locators. The relationships are free-form, which means they can be one-to-many, many-to-many or one-to-one and have any meaning required. A relation then consists of a type, one or more independent variables and one or more dependent variables. Also an extra parameter can be defined. Examples of relation types are `associated_to`, `lookupvalue`, `defined_by`, `landmark` or `subdivision` (see Table 2.3). An `associated_to` relation can for instance define the link between an anatomical structure and the paper that is relevant to that specific anatomical structure.

2.5. IMPLEMENTATION

The first prototype application to store, explore and query model information was implemented in Python as a DeVIDE module[25], employing the Visualization Toolkit (VTK) for its visualization functionality. Later, a standalone version was made. The proof-of-concept standardized anatomical space was based on a semi-automatic segmentation in Amira[26] of cryosectional images of a Dutch female pelvis. After this initial prototype atlas was made, we built a more detailed model, which is described in Chapter 3. An important component of our system is the underlying database, which was implemented using MongoDB. MongoDB is a *schema-less* document-oriented database technology that is designed to be agile and scalable. Since there is no fixed schema design required in the initial stages of the project, a benefit of using this technology is that it is easy to further extend our model as new data is provided to us by the anatomists, without any changes required to the existing database. Furthermore, MongoDB's GridFS enables us to store large volume data and rapidly retrieve it from the database.

2.6. RESULTS

The prototype application offers several functionalities. The user interface consists of three render windows that represent the surface model along with two linked slice viewers, that can display arbitrary volume data. The user is not only able to query anatomical

<i>Relation</i>	<i>Dependent</i>	<i>Independent</i>	<i>Parameter</i>
associated_to	literature	structure	-
lookupvalue	value	volume	-
defined_by	structure	structure	operation
landmark	structure	points	type
subdivision	structure	structure	planes

Table 2.3: Types of relations that are currently supported by our system. This can easily be extended with new types at any time.

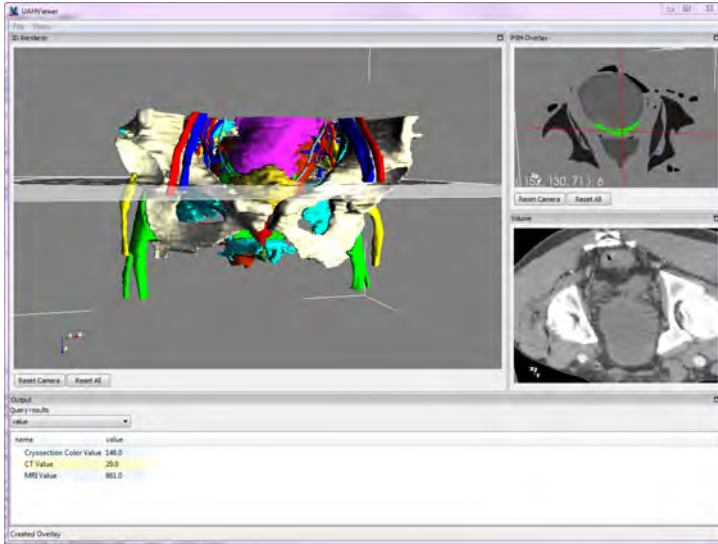


Figure 2.2: Value querying in arbitrary views presents query results for all available volumes in that point. The surface render shows an MRI slice while the top right render window displays the linked segmentation label volume slice and the lower right render window displays the mapped CT slice

structures topically, by selecting an anatomical term from a drop-down list, but can also query the model spatially. In this section we demonstrate the utility of our model-based anatomy visualization prototype using three examples.

Generic multi-modal data querying By mapping volumes to anatomical model coordinates, it becomes possible to simultaneously slice through volumes of different modalities and different subjects. The linked volume representations enable the user to compare arbitrary multi-modal volumes side-by-side in an intuitive way.

It is also possible to query a specific voxel value in all available volumes at a selected point. This type of spatial querying is done by either clicking on a structure in the surface rendering or in one of the slice viewers. For example, in Figure 2.2 a point is queried in the segmentation label volume. This query returns the voxel values for all volumes available in that point. The linked representations make it easy to locate points of interest in any view preferred by the user. The CT value can then be used for instance to calculate the Young's modulus of certain tissues.

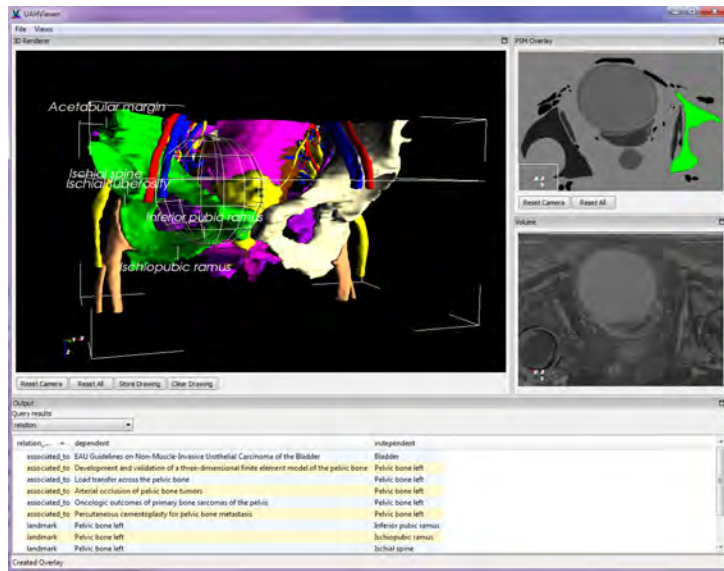


Figure 2.3: Spatial querying through sphere selection. A distance query reveals the literature and anatomical landmarks associated to the structures within the selection sphere.

Distance querying It is also possible to query a specific volume of the model space by using a selection sphere. In this way, users can select an area of interest and perform a distance query in this area. The application then returns the query results of the point that was selected as well as all locators in the region that was selected by the user. Figure 2.3 shows how the sphere selection returns query results based on the locators found inside the sphere. The query results contain all structures within the sphere of interest along with their relations to other model objects. In this case, the user has added related literature to several structures and annotated anatomical landmarks on the structure in 3D. By following the relations that model objects have, it is also possible to query the related model objects and to chain these related query results up to arbitrary distances in the connectivity graph.

Mapping model information to patient data A third example of a prototype functionality is enriching unseen patient scan data with the information provided by model objects. By registering a patient-specific scan from an arbitrary modality and providing a mapping to model space, it becomes possible to reveal structures that can not be denoted in that modality by default. For instance, because there is not enough contrast between nerves and surrounding soft tissues in CT-scans, the location of nerves is not visible in a CT-scan. In histological and cryosectional images, however, the nerves can be seen and segmented. In Figure 2.4, a CT-scan added to model space is enriched with an isosurface render of the nerves segmented out of the cryosectional images that form the basis of the model.

Using this same principle, it is also possible to map the segmentation label volume




Figure 2.4: Using our approach, the nerves that were segmented in the cryosectional images can be mapped to unseen patient specific CT-scan data. The location of these nerves is not visible in the original CT-scan.

<i>System</i>	<i>Modalities</i>	<i>Resolution</i>	<i>Variations</i>	<i>Relations</i>
VM	multi	single	single general anatomy	structural/functional/abstraction
BG	single	single	multiple specimens	semantic/spatial
Our	multi	multi	multiple specimens	free-form

Table 2.4: Comparison between VOXEL-MAN (VM), Braingazer (BG) and our system.

to patient-specific data. If the registration is correct, this creates an automatic segmentation of the patient-specific scan. This mapped segmentation can be used to make a 3D patient model by rendering isosurfaces.

Table 2.4 shows a comparison between our proposed system and existing systems such as VOXEL-MAN and Braingazer. While there are many similarities, our proposed system excels in allowing multi-modal, multi-resolution datasets to be stored, queried and visualized. Considering the ability to handle variations, we observe that the VOXEL-MAN-project represents only a single general anatomy, while our system allows the user to store and visualize datasets with age, interindividual and anatomical variations. When we examine the relational capabilities, we see that in the VOXEL-MAN-project, relations are used to describe anatomical links that are structural, functional or represent abstraction in a semantic network [17]. For the Braingazer project, semantic and spatial relations can be defined. Our relations can be used to define all of these links, but extend this idea by allowing relations to be defined between locators of arbitrary types of model objects.

2.7. CONCLUSION

In this chapter, we have presented a method to integrate heterogeneous spatial and non-spatial data from different sources, as well as the complex relations between them, into a single model. Using a standardized coordinate system, all available anatomical knowledge can be queried interactively in the prototype application, both topically and spatially. The model data can be visualized as is, compared in linked views, or be used to

enrich patient-specific scans with additional information from the model. Due to the flexibility of the model, anatomists will be able to store arbitrary heterogeneous data in the model in a generic way. As such the system forms an excellent foundation for future pre-operative surgical planning and intra-operative guidance applications.

In the Chapter 3, based on the method presented here, we describe the construction of an atlas of the human pelvis, based on a manual segmentation of the cryosectional slices of the Visible Korean female dataset and histological studies.

REFERENCES

- [1] N. Smit, A. Kraima, D. Jansma, M. deRuiter, and C. Botha, *A unified representation for the model-based visualization of heterogeneous anatomy data*, in *EuroVis - Short Papers* (Eurographics Association, 2012) pp. 85–89.
- [2] N. Smit, A. Kraima, D. Jansma, M. DeRuiter, and C. Botha, *The unified anatomical human (beta): Model-based representation of heterogeneous anatomical data*, in *Workshop 3D Physiological Human (3DPH), CASA* (2012).
- [3] K. Lawonn, N. N. Smit, K. Bühler, and B. Preim, *A survey on multimodal medical data visualization*, *Computer Graphics Forum* (2016), in Submission.
- [4] B. Krauss, B. Schmidt, and T. G. Flohr, *Dual source CT*, in *Dual energy CT in clinical practice* (Springer, 2011) pp. 11–20.
- [5] A. Berrington de González, M. Mahesh, K.-P. Kim, M. Bhargavan, R. Lewis, F. Mettler, and C. Land, *Projected cancer risks from computed tomographic scans performed in the United States in 2007*, *Archives of Internal Medicine* **169**, 2071 (2009).
- [6] C. Rezk-Salama, P. Hastreiter, J. Scherer, and G. Greiner, *Automatic adjustment of transfer functions for 3D volume visualization*, in *Proc. of VMV* (2000) pp. 357–364.
- [7] P. S. Tofts, G. Brix, D. L. Buckley, J. L. Evelhoch, E. Henderson, M. V. Knopp, H. B. Larsson, T.-Y. Lee, N. A. Mayr, G. J. Parker, *et al.*, *Estimating kinetic parameters from dynamic contrast-enhanced t_1 -weighted mri of a diffusable tracer: standardized quantities and symbols*, *Journal of Magnetic Resonance Imaging* **10**, 223 (1999).
- [8] D. Le Bihan, J.-F. Mangin, C. Poupon, C. A. Clark, S. Pappata, N. Molko, and H. Chabriat, *Diffusion tensor imaging: concepts and applications*, *Journal of magnetic resonance imaging* **13**, 534 (2001).
- [9] I. Hotz and T. Schultz, *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data* (Springer, 2015).
- [10] M. P. Van Den Heuvel and H. E. H. Pol, *Exploring the brain network: a review on resting-state fMRI functional connectivity*, *European Neuropsychopharmacology* **20**, 519 (2010).
- [11] I. Viola, Å. Birkeland, V. Solteszova, L. Helljesen, H. Hauser, S. Kotopoulos, K. Nylund, D. M. Ulvang, and Øye, *High-Quality 3D Visualization of In-Situ Ultrasonography*,

- in *Proc. of Eurographics 2013 - Dirk Bartz Prize*, edited by H.-C. Hege and A. Vilanova.
- [12] G. Treglia, P. Castaldi, M. F. Villani, G. Perotti, A. Filice, V. Ambrosini, N. Cremonini, A. Versari, S. Fanti, A. Giordano, *et al.*, *Comparison of different positron emission tomography tracers in patients with recurrent medullary thyroid carcinoma: our experience and a review of the literature*, in *Theranostics, Gallium-68, and Other Radionuclides* (2013) pp. 385–393.
- [13] D. L. Bailey and K. P. Willowson, *An evidence-based review of quantitative spect imaging and potential clinical applications*, *Journal of Nuclear Medicine* **54**, 83 (2013).
- [14] U. Tiede, M. Bomans, K. H. Höhne, A. Pommert, M. Riemer, T. Schiemann, R. Schubert, and W. Lierse, *A computerized three-dimensional atlas of the human skull and brain*. *American Journal of Neuroradiology* **14**, 551 (1993).
- [15] A. Pommert, *Symbolic modeling of human anatomy for visualization and simulation*, (SPIE, 1994) pp. 412–423.
- [16] A. Pommert, M. Riemer, T. Schiemann, R. Schubert, U. Tiede, K. H. Hoehne, R. Schubert, U. Tiede, and K. H. Höhne, *Knowledge-Based and 3D imaging systems in medical education*, (1994).
- [17] K. H. Höhne, B. Pflesser, A. Pommert, M. Riemer, T. Schiemann, R. Schubert, and U. Tiede, *A new representation of knowledge concerning human anatomy and function*, *Nature Medicine* **1**, 506 (1995), PMID: 7585108.
- [18] U. Tiede, T. Schiemann, and K. H. Höhne, *High quality rendering of attributed volume data*, in *Visualization Conference, IEEE* (IEEE Computer Society, Los Alamitos, CA, USA, 1998) p. 255.
- [19] A. Pommert, *Creating a high-resolution spatial/symbolic model of the inner organs based on the visible human*, *Medical Image Analysis* **5**, 221 (2001).
- [20] K. H. Höhne, B. Pflesser, A. Pommert, M. Riemer, R. Schubert, T. Schiemann, U. Tiede, and U. Schumacher, *A realistic model of human structure from the visible human data*, *Methods of Information in Medicine* **40**, 83 (2001), PMID: 11424309.
- [21] S. Gehrmann, K. Höhne, W. Linhart, B. Pflesser, A. Pommert, M. Riemer, U. Tiede, J. Windolf, U. Schumacher, and J. Rueger, *A novel interactive anatomic atlas of the hand*, *Clinical Anatomy* **19**, 258–266 (2006).
- [22] A. Pommert, K. H. Höhne, E. Burmester, S. Gehrmann, R. Leuwer, A. Petersik, B. Pflesser, and U. Tiede, *Computer-Based anatomy: A prerequisite for Computer-Assisted radiology and surgery1*, *Academic Radiology* **13**, 104 (2006).
- [23] S. Bruckner, V. Šoltészová, E. Gröller, J. Hlad?vka, K. Bühler, J. Y. Yu, and B. J. Dickson, *BrainGazer - visual queries for neurobiology research*, *IEEE Transactions on Visualization and Computer Graphics* **15**, 1497 (2009).

- [24] M. Baiker, J. Milles, J. Dijkstra, T. D. Henning, A. W. Weber, I. Que, E. L. Kaijzel, C. W. Löwik, J. H. Reiber, and B. P. Lelieveldt, *Atlas-based whole-body segmentation of mice from low-contrast Micro-CT data*, *Medical Image Analysis* **14**, 723–737 (2010).
- [25] C. P. Botha and F. H. Post, *Hybrid scheduling in the DeVIDE dataflow visualisation environment*, in *Proceedings of Simulation and Visualization*, edited by H. Hauser, S. Strassburger, and H. Theisel (SCS Publishing House Erlangen, 2008) p. 309–322, best paper award.
- [26] D. Stalling, M. Westerhoff, and H. christian Hege, *Amira: A highly interactive system for visual data analysis*, in *The Visualization Handbook* (Elsevier, 2005) pp. 749–767.

3

THE VIRTUAL SURGICAL PELVIS: ATLAS CONSTRUCTION

In rectal cancer surgery, excellent knowledge of pelvic anatomy is essential to remove tumors radically without damaging structures. Surgeons are challenged to relate 2D textbook anatomy to 3D in vivo anatomy. Previously, there were no pelvic models available with the surgically relevant detail required for surgical education and planning. Education is the key to optimize anatomical knowledge and surgical training and should be given in light of the current technological advancements that replace tactile feedback for visual feedback. This chapter describes the development and application potentials of the 'Virtual Surgical Pelvis' (VSP): a virtual 3D anatomical atlas, for educational and surgical planning purposes.

The VSP is an anatomical atlas of various synchronously linked 2D anatomical data and a 3D reconstruction, showing important relationships between pelvic organs, blood vessels, fasciae and nerves. As a proof of concept, the autonomic nerves included in the atlas were mapped onto five preoperative MRI scans from rectal cancer patients. The Sørensen-Dice coefficient, reflecting the registration accuracy, was 97-99%.

The VSP and method of developing a virtual 3D anatomical atlas can be of great potential in surgical anatomical education, which we describe in Chapter 4. It can provide the future anatomical context for surgical planning, as is described in Chapter 8, and in the future potentially for robotic surgery, image-guided surgery and radiotherapy planning.

This chapter is based on a publication currently in submission to Anatomical Sciences Education: Annelot Kraima, **Noeska Smit**, Daniel Jansma, Elmar Eisemann, Jin Seo Park, Beom Sun Chung, Nicholas West, Philip Quirke, Harm Rutten, Cornelis van de Velde, Anna Vilanova, Marco DeRuiter, "A novel 3D virtual anatomical atlas: taking the next step in enhancing surgical anatomical education and patient-specific clinical guidance", In Submission, *Anatomical Sciences Education*, 2016 [1].

3.1. INTRODUCTION

Every surgeon relies on anatomical knowledge when operating in the human body. As described in Chapter 1, for regions such as the pelvis this anatomy is more complex than for other regions. Good-quality TME surgery strongly relies on how surgeons perform the actual surgery. To this end, they must have an excellent understanding of the detailed pelvic anatomy. Unfortunately, incomplete surgical specimens and postoperative urogenital and anorectal dysfunction are still reported [2–6], demonstrating the need to further optimise treatment.

In daily clinical practice, surgeons must link two-dimensional (2D) anatomical images to three-dimensional (3D) in vivo surgical anatomy. This may seem common practice to the experienced consultant surgeon, but surgical residents, i.e., surgeons in training, often struggle much more to correctly identify anatomy. Moreover, the progressive usage of laparoscopic and robotic approaches in colorectal surgery replaces tactile feedback for visual feedback, challenging surgeons to recognize detailed anatomy on 2D projection monitors. With respect to TME, precise knowledge of the relationship of perirectal fasciae, forming the dissection planes, and autonomic nerves is essential [7, 8]. This has been often studied, yet various contradicting descriptions exist in the literature, hampering to define uniform dissection planes along which the rectum should be mobilised [9–13]. Thus, there is a need to improve the detailed anatomical knowledge of the pelvis by enhancing the surgical anatomical education, taking into account that TME is increasingly conducted in technologically driven surgical fields. 3D anatomical modeling is a popular method to visualize complex anatomical structures or regions. The few existing 3D pelvic models fail in showing detailed anatomy, such as perirectal fasciae and autonomic nerves. These models are generic models showing average anatomy rather than patient-specific anatomy [14].

As described in Chapter 2, we have established a method to integrate heterogeneous anatomical data in one single atlas [15, 16]. This chapter describes the creation of a novel 3D virtual anatomical atlas of the pelvis, showing important anatomical relationships and surgical risk zones: ‘The Virtual Surgical Pelvis.’ To avoid visualization of anatomical misconceptions, pelvic cadaveric specimens have been histologically analyzed. Additionally, the first results of patient-specific 3D anatomical modeling will be presented, which forms a basis for the work presented in Chapter 8 on patient-specific surgical planning.

3.2. MATERIAL & METHODS

3.2.1. MATERIAL

Adult cadaveric specimens Two human adult female cadaveric specimens were obtained through the NHS Leeds GIFT Research Tissue Programme¹ from consented donor bodies. Ethical approval was granted by the Northern and Yorkshire Regional Ethics Committee, Jarrow, UK (unique reference number 11/H0903/6). The specimens were retrieved during tissue donation autopsies performed at the St. James’s University Hospital in Leeds. None of the specimens suffered from pelvic pathology at post mortem examination. The specimens were essentially pelvic exenterations, which involves the re-

¹www.gift.leeds.ac.uk

removal of all organs from the pelvic cavity, and comprised the whole en-bloc soft-tissues kept in the bony pelvis from the pelvic brim to the perineum. These were transversely sectioned at 1 cm after which the slices were dissected to fit in Super Mega Cassettes measuring 74.8 x 52.5 x 16.5 mm (CellPath; Powys; UK). The tissues were rehydrated in graded ethanol and xylene, and embedded in paraffin mega-blocks.

The Visible Korean Female The Visible Korean Female dataset comprising digitalised transverse cryosections of a donor body was used for anatomical segmentation [17–19]. The donor body belonged to a female aged 26 years who died from gastric cancer. We detected a rectocele and a small amount of intra-abdominal fluid, of which the latter could be explained by metastatic disease. There was no other pelvic pathology. The spatial resolution of the images was 5,616 x 2,300 pixels and the cross-sectional interval was 0.2 mm. 910 consecutive images were selected comprising the whole pelvis from the 5th lumbar vertebra to the perineum.

3.2.2. METHODS

Histological staining and image acquisition adult cadaveric specimens The mega-blocks were serially cut in 5 μm transverse sections. Every 50th section was collected and stained with H&E resulting in a cross-sectional interval of 250 μm . Additional sections were collected from all mega-blocks for extra stains with Masson's trichrome and Miller's elastin [20]. The stained glass slides were digitally scanned in Leeds with an Aperio XT slide scanner (Aperio, San Diego, California, USA) at 20x magnification, creating a resolution of 0.46 microns per pixel. The digital images were compressed with JPEG2000 quality 70 and viewed in Aperio ImageScope (version 10.2.2.2.2319).

Anatomical segmentation of the VKF dataset 910 selected images from the VKF were loaded in Amira software version 5.3.3 (Amira, Oregon, USA). Structures of interest were manually and semi-automatically segmented (see Table 3.1). The voxel size of the image stack was 0,1 x 0,1 x 0,2 mm. Our goal was to visualize the relationship of the perirectal fasciae and autonomic nerves. Because it was impossible to identify and segment the pelvic autonomic network in the VKF, risk zones were segmented showing where surgical damage of the autonomic nerves might occur. These risk zones were defined based on the anatomical insights gained through the microscopic analysis. All segmentations were exported as label files.

The Unified Anatomical Human The Unified Anatomical Human (UAH) representation [15, 16], described in Chapter 2, was used as a platform to combine all relevant anatomical knowledge of the pelvis into a single anatomical atlas. The accompanying UAHviewer is a software system facilitating interactive exploration of the anatomical atlas in three linked views: one 3D view and two 2D views.

Image processing of the segmentation labels The segmentation labels of the VKF dataset formed the basis of the atlas. The initial 3D reconstruction was created in DeVIDE [21] using a Marching Cubes algorithm, followed by post-processing in Meshlab [22] to simplify and smooth the automatically generated meshes. All anatomical structures were




Figure 3.1: Schematic overview of the registration process through which MR images can be combined with anatomical data captured in the VSP in order create a patient-specific model.

reconstructed in 3D and stored in the UAH. Since the segmentation labels corresponded directly to the original cryosections of the VKE, the 2D images could be synchronously linked with the 3D reconstruction in the atlas.

Registration onto patient-specific radiologic images Patient-specific MR images were selected to register anatomical data contained in the atlas onto. A schematic of the registration pipeline can be seen in Figure 3.1). Volume registration is defined as the process of establishing image correspondence between a source and a target volume. In this project, the source volume was the dataset we wished to match, i.e. the volume of the segmentation labels of the atlas. The target volume was the dataset we aimed to deform towards, i.e. the patient-specific MR images. We did not register the atlas directly onto the MR images, but segmented in the MR images gross anatomical structures expected to undergo most inter-individual variation, such as the rectum and mesorectum, uterus and vagina, bladder and bony pelvis. Firstly, a manual affine registration was performed, which is a transformation of data and includes translation, scaling and rotations to coarsely align the atlas and scan segmentation labels. Secondly, the atlas organs were fully deformed to fit the patient-specific organs. The same transformations could be applied to anatomical structures not visible in the patient-specific MR images, such as the risk zones, making the atlas information patient-specific. The entire registration process was performed using DeVIDE [21], Slicer [23], and Elastix [24].

3.3. RESULTS

Adult cadaveric specimens To integrate the correct anatomical relationships of the perirectal fasciae and autonomic nerves, cadaveric specimens were histologically analysed. The surgically famous ‘holy plane’ was bordered by the mesorectal and parietal fascia. The mesorectal fascia (also referred to as the visceral fascia or fascia propria recti) was the innermost fascia encircling the mesorectal fat. The parietal fascia (also referred to as the presacral fascia, and sometimes mistakenly as Waldeyer’s fascia) was the outermost fascia covering the presacral space. The pelvic autonomic network included the hypogastric nerves, inferior hypogastric plexus, pelvic splanchnic nerves and neurovascular bundles. This part of the pelvic autonomic nervous system was segmented as the surgical risk zone containing the autonomic nerves. The hypogastric nerves and inferior hypogastric plexus were located directly lateral to the mesorectal fascia, and consequently the risk zones as well (see Figure 3.2). The pelvic splanchnic nerves originated




Figure 3.2: Whole mount microscopic section of a female cadaveric pelvic specimen showing the precise relationship between the autonomic nerves (arrowheads) and mesorectal fascia (arrow). C: cervix; R: rectum; M: mesorectum; MT: Masson's trichrome. Scale bar window a: 6mm, window a.i: 1 mm.

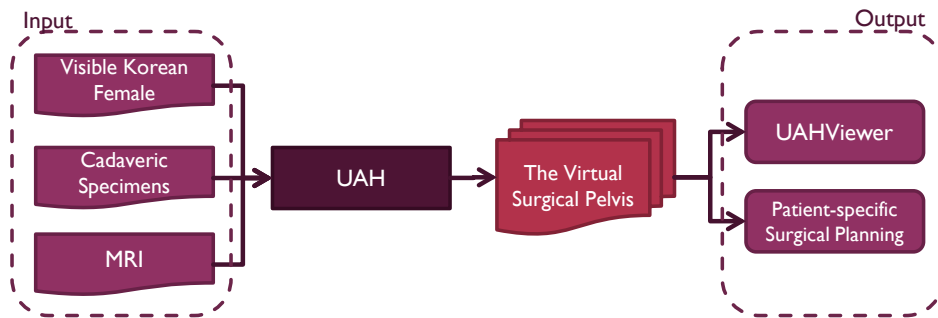


Figure 3.3: Schematic overview of the development of the Virtual Surgical Pelvis (VSP). The input data are the cryosections and label images of the Visible Korean Female, the information gained by histological analysis of cadaveric specimens and patient-specific MR images. Through the Unified Anatomical Human representation all this information can be combined and the VSP can be developed. The anatomical information captured in the VSP can be viewed with the UAHViewer, OAHViewer and used for patient-specific planning.

from the sacral foramina, penetrated the parietal fascia and merged with the hypogastric nerves building up the inferior hypogastric plexus. This part of the pelvic autonomic nervous system was segmented as the risk zone containing the pelvic splanchnic nerves.

The Virtual Surgical Pelvis The VSP is a collection of all available anatomical data, including the original cryosections of the VKF, label images and 3D reconstruction, of a female pelvis integrated in one anatomical atlas, as shown in the overview in Figure 3.3. It contains surgically relevant anatomical structures and relationships, divided into several anatomical groups: the bony pelvis, musculature, organs, vasculature, nervous system, fasciae and risk zones (see Table 3.1). These risk zones are indicated as regions containing the distal part of the superior hypogastric plexus, hypogastric nerves, inferior hypogastric plexus, pelvic splanchnic nerves and neurovascular bundles. The perirectal fasciae are reconstructed in 3D based on the spaces they delineate. The resulting 3D pelvic atlas model can be seen in Figure 3.4.



Figure 3.4: The complete VSP female pelvic atlas model, featuring the structures in Table 3.1 with the autonomic nerve risk zones indicated in bright yellow.

The VSP can be applied in two important ways. Firstly, the VSP can be interactively explored and annotated for educational purposes (see Chapter 4). Secondly, anatomical structures or zones of interest captured in the VSP can be mapped onto patient-specific radiologic images, such as preoperative MR images or CT images for pre-operative planning purposes (see Chapter 8). In this initial study, we registered the VSP to five MRI scans of patients suffering from high-, mid- and low-rectal tumors (see Table 3.2). All tumors were staged as being T3 or T4 tumors, meaning that these tumors grow deep into the mesorectal fat (T3) or penetrate the mesorectal fascia (T4).

In Figure 3.5 the results of the registration process are shown. In the column on the left we show the risk zones containing the autonomic nerves mapped to the patient-specific MRI. The whole 3D volume is registered, but here we visualize the results on a single slice. The column on the right shows patient-specific 3D models we were able to reconstruct based on the registration results. We visualize patient-specific anatomy, as well as an indication of the areas where nerve damage may occur, which are not available from the MRI scans. Note that patient #5 lacks a uterus; this patient had undergone a uterus extirpation.

Category	Structures
Bones/intervertebral discs	Fifth lumbar vertebra Lumbar disc Os coxae Sacral disc Coccygeal disc
Muscles	Iliospoas muscle Piriformis muscle Levator ani muscle Coccygeal muscle Bulbospongiosus muscle Ischiocavernosus muscle External anal sphincter
Organs	Rectum Bladder Ureter Uterus Ovaries Fallopian tubes Vagina Corpora cavernosa
Vasculature	Arteries Veins
Nerves	Somatic nerves Risk zones autonomic nerves Risk zones splanchnic nerves
Fasciae	Parietal fascia Mesorectal fascia

Table 3.1: This table lists the anatomical structures and zones of interest that were segmented in the Visible Korean Female dataset.

3.4. DISCUSSION

An excellent understanding of detailed pelvic anatomy must be part of every surgeon's compendium to operate safely in the pelvis. Good-quality TME relies on a comprehensive understanding of the pelvic anatomy, and specifically the arrangement of perirectal fasciae and nerves. To the best of our knowledge, we are the first to have developed a solid method to build a virtual 3D anatomical atlas of the pelvis by combining heterogeneous anatomical data in one single atlas. The creation of "the Virtual Surgical Pelvis" (VSP) concerns work in progress and to date, the VSP represents most important surgical anatomy of the pelvis and specifically reveals relationships between organs, muscles, blood vessels, fasciae and somatic and autonomic nerves. Histological analysis of cadaveric specimens revealed the course of the autonomic nerves and their location just laterally to the perirectal fasciae.

Anatomy of interest (either structures or zones) can be mapped onto radiologic images. The present work has shown that it is feasible to enrich patient-specific preoperative MR images with anatomical data captured in the VSP. As a proof of concept, risk zones of the autonomic nerves were mapped onto five patient-specific MRI scans. It was possible to highlight the risk zones in the MR scans and single MR images as well as to render a patient-specific 3D volume. To our knowledge, the VSP is the first anatomical

Patient ID	Patient Characteristics	Registration Accuracy
1	Female 71 years old T4 mid-rectal tumor	Sørensen–Dice index: 0.99
2	Female 92 years old T4 high-rectal tumor	Sørensen–Dice index: 0.99
3	Female 92 years old T3 low-rectal tumor	Sørensen–Dice index: 0.97
4	Female 58 years old T3 low-rectal tumor	Sørensen–Dice index: 0.99
5	Female 48 years old T3 high-rectal tumor	Sørensen–Dice index: 0.99

Table 3.2: Patient information for the five selected MRI scans. The Sørensen-Dice index of the registration outcome is listed additionally.

atlas allowing patient-specific 3D modeling rather than generic modeling. Preoperative exploration of patient-specific anatomy in both 2D and 3D views helps surgeons to focus on specific anatomy that is not visible in conventional MR images.

This work has some limitations. Currently, the VSP involves the female pelvis and not suitable to explore (patient-specific) 3D male anatomy. In the future, we aim to segment the Visible Korean Male dataset [17, 25] and enrich the VSP with a male pelvis. Furthermore, the pelvic anatomy as segmented in the VKF is not “complete” in a way that all anatomical structures are included in the VSP. For example, the venous network is not included as post-mortem venous distension hampered to identify the complete venous network in the VKF. The atlas deals with potential anatomical variability in the individual nerves by defining the risk zones in a broad region to account for all possible individual variations. Further improvements in the atlas and a more specific definition of the risk zones via cadaveric studies would have beneficial consequences for our application.

3.5. CONCLUSION

In conclusion, we are the first to have developed a novel method to build a virtual 3D anatomical atlas containing heterogeneous anatomical data. This method is applicable to every anatomical region. The first spin-off is the VSP: an anatomical atlas containing cryosections, segmentation labels and a 3D reconstruction of a female pelvis. The VSP currently shows the intricate relationships between pelvic organs, muscles, blood vessels, fasciae and nerves, but can be constantly enriched with new anatomical data of different image modalities. As a proof of concept, we have shown that anatomical data entailed in the VSP can be mapped onto patient-specific MR images. By visualizing patient-specific anatomy either in 2D by enriching MR images or in 3D by reconstructing a patient-specific model, we provided a basis for next steps in enhancing preoperative surgical planning and decision-making.




Figure 3.5: The row numbers correspond to the patient's ID listed in Table 2. Column a shows the result of registration of the atlas autonomic nerve-risk zones onto patient-specific MR images. Column b shows the patient-specific 3D models. The bony pelvis is depicted in grey, the rectum in green, the vagina and uterus in purple, the bladder in orange and the risk zones of the autonomic nerves in yellow.

REFERENCES

- [1] A. C. Kraima, N. N. Smit, D. Jansma, E. Eisemann, M. Chung, N. West, P. Quirke, H. Rutten, C. Van De Velde, A. Vilanova, and M. C. DeRuiter, *A novel 3d virtual anatomical atlas: taking the next step in enhancing surgical anatomical education and patient-specific clinical guidance*, Anatomical Sciences Education (2016), in Submission.
- [2] M. M. Lange and C. J. van de Velde, *Urinary and sexual dysfunction after rectal cancer treatment*, Nature Reviews Urology **8**, 51 (2011).
- [3] M. Lange, C. Marijnen, C. Maas, H. Putter, H. Rutten, A. Stiggelbout, E. M.-K. Kranenbarg, C. van de Velde, C. C. I. of the Dutch, *et al.*, *Risk factors for sexual dysfunction after rectal cancer treatment*, European Journal of Cancer **45**, 1578 (2009).
- [4] C. Wallner, M. M. Lange, B. A. Bonsing, C. P. Maas, C. N. Wallace, N. F. Dabhoiwala, H. J. Rutten, W. H. Lamers, M. C. DeRuiter, and C. J. van de Velde, *Causes of fecal and urinary incontinence after total mesorectal excision for rectal cancer based on cadaveric surgery: a study from the cooperative clinical investigators of the dutch total mesorectal excision trial*, Journal of clinical oncology **26**, 4466 (2008).
- [5] M. Den Dulk, H. Putter, L. Collette, C. A. Marijnen, J. Folkesson, J.-F. Bosset, C. Rödel, K. Bujko, L. Pählman, and C. J. Van De Velde, *The abdominoperineal resection itself is associated with an adverse outcome: the european experience based on a pooled analysis of five european randomised clinical trials on rectal cancer*, European Journal of Cancer **45**, 1175 (2009).
- [6] N. West, C. Anderin, K. Smith, T. Holm, and P. Quirke, *Multicentre experience with extralevator abdominoperineal excision for low rectal cancer*, British Journal of Surgery **97**, 588 (2010).
- [7] R. Heald, *The'holy plane'of rectal surgery*. Journal of the Royal Society of Medicine **81**, 503 (1988).
- [8] R. Heald and B. Moran, *Embryology and anatomy of the rectum*, in *Seminars in surgical oncology*, Vol. 15 (Wiley Online Library, 1998) pp. 66–71.
- [9] Y. Kinugasa, G. Murakami, D. Suzuki, and K. Sugihara, *Histological identification of fascial structures posterolateral to the rectum*, *British Journal of Surgery* **94**, 620 (2007).
- [10] I. Lindsey, B. F. Warren, and N. J. Mortensen, *Denonvilliers' fascia lies anterior to the fascia propria and rectal dissection plane in total mesorectal excision*, Diseases of the colon & rectum **48**, 37 (2005).
- [11] Y. Kinugasa, G. Murakami, K. Uchimoto, A. Takenaka, T. Yajima, and K. Sugihara, *Operating behind denonvilliers' fascia for reliable preservation of urogenital autonomic nerves in total mesorectal excision: a histologic study using cadaveric specimens, including a surgical experiment using fresh cadaveric models*, Diseases of the colon & rectum **49**, 1024 (2006).

- [12] F. Peschard, D. Moszkowicz, B. Alsaid, T. Bessede, C. Penna, and G. Benoit, *Preservation of genital innervation in women during total mesorectal excision: which anterior plane?* World journal of surgery **36**, 201 (2012).
- [13] N. Clausen, T. Wolloscheck, and M. A. Konerding, *How to optimize autonomic nerve preservation in total mesorectal excision: clinical topography and morphology of pelvic nerves and fasciae*, World journal of surgery **32**, 1768 (2008).
- [14] A. C. Kraima, N. N. Smit, D. Jansma, C. Wallner, R. Bleys, C. Van De Velde, C. P. Botha, and M. C. DeRuiter, *Toward a highly-detailed 3d pelvic model: Approaching an ultra-specific level for surgical simulation and anatomical education*, Clinical Anatomy **26**, 333 (2013).
- [15] N. Smit, A. Kraima, D. Jansma, M. deRuiter, and C. Botha, *A unified representation for the model-based visualization of heterogeneous anatomy data*, in *EuroVis - Short Papers* (Eurographics Association, 2012) pp. 85–89.
- [16] N. Smit, A. Kraima, D. Jansma, M. DeRuiter, and C. Botha, *The unified anatomical human (beta): Model-based representation of heterogeneous anatomical data*, in *Workshop 3D Physiological Human (3DPH)*, CASA (2012).
- [17] J. S. Park, M. S. Chung, S. B. Hwang, Y. S. Lee, D.-H. Har, and H. S. Park, *Visible korean human: improved serially sectioned images of the entire body*, Medical Imaging, IEEE Transactions on **24**, 352 (2005).
- [18] J. S. Park, Y.-W. Jung, J. W. Lee, D. S. Shin, M. S. Chung, M. Riemer, and H. Handels, *Generating useful images for medical applications from the visible korean human*, Computer methods and programs in biomedicine **92**, 257 (2008).
- [19] M. S. Chung, *Visible korean human. improved serially sectioned images of the entire body*, in *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the* (IEEE, 2006) pp. 6563–6566.
- [20] M. M. Everett and W. A. Miller, *The role of phosphotungstic and phosphomolybdic acids in connective tissue staining i. histochemical studies*, The Histochemical Journal **6**, 25 (1974).
- [21] C. P. Botha and F. H. Post, *Hybrid scheduling in the DeVIDE dataflow visualisation environment*, in *Proceedings of Simulation and Visualization*, edited by H. Hauser, S. Strassburger, and H. Theisel (SCS Publishing House Erlangen, 2008) p. 309–322, best paper award.
- [22] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia, *Meshlab: an open-source mesh processing tool*. in *Eurographics Italian Chapter Conference*, Vol. 2008 (2008) pp. 129–136.
- [23] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J.-C. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, *et al.*, *3d slicer as an image computing platform for the quantitative imaging network*, Magnetic resonance imaging **30**, 1323 (2012).

- [24] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, *Elastix: a toolbox for intensity-based medical image registration*, Medical Imaging, IEEE Transactions on **29**, 196 (2010).
- [25] J. S. Park, M. S. Chung, S. B. Hwang, B.-S. Shin, and H. S. Park, *Visible korean human: its techniques and applications*, Clinical Anatomy **19**, 216 (2006).

II

EDUCATION

IN this part of the thesis, we present visualization methods related to the second goal of the work focusing on medical education. By combining the anatomical knowledge on the pelvis using the methods described in Chapter 2, we were able to build a complete atlas of the human pelvis, the Virtual Surgical Pelvis (VSP), as described in Chapter 3. We have developed a way to share the knowledge contained in the VSP with students worldwide via the Internet, which is described in Chapter 4. This educational platform supports 3D annotations, which makes it possible to further enrich the VSP with knowledge, or quiz students.

The VSP does not take topological anatomical variations in branching structures, such as the arteries into account. To visualize anatomical variations such as these for educational purposes, we have developed an interactive application that is described in Chapter 5. The application reveals similarities and dissimilarities in sets of variations, based on graph representations of the structures.

4

WEB-BASED ANATOMY EDUCATION

Human anatomy, the study of the structure of the human body, is complex and as early as the late bronze age, people have been trying to gain insights into the functioning of the human body. Nowadays, resources such as books and software are available to educate medical students. Nonetheless, these media usually have some restrictions; anatomical images in books present static views and do not allow readers to freely explore the information, while software tools often present an idealized average human anatomy.

In this work, we introduce the Online Anatomical Human (OAH), a web-based viewer for studying anatomy. It is based on real human anatomy and incorporates medical image data in linked 2D and 3D views that students can freely interact with. Our application is the only to support both 2D and 3D views based on real medical imaging data. The main goal of this anatomical online resource is to serve as an educational platform available to anyone with access to a modern web browser. Our solution allows users to annotate regions, add comments, and provide hyperlinks to additional media. By making our work accessible to medical experts, we can ensure an increasing amount of information, leading to an ever-growing gain in educational value. The OAH is currently used in an Massive Open Online Course (MOOC) to teach anatomy of the pelvis and made available to students worldwide via the web.

This chapter is based on the following publication: **Noeska Smit**, Cees-Willem Hofstede, Annelot Kraima, Daniel Jansma, Marco Deruiter, Elmar Eisemann, and Anna Vilanova, "The Online Anatomical Human: Web-based Anatomy Education", *Proceedings of EuroGraphics Education Papers*. 2016 [1], which was based on work performed by Cees-Willem Hofstede in his Master thesis [2].

4.1. INTRODUCTION

Human anatomy is the science of shape and structure of the human body. One of the oldest known papers on human anatomy, known as the Edwin Smith Papyrus, dates back to approximately 1700 BC and still today, most education on anatomy is based on books [3]. Even though software exists in addition, the presented anatomical models are often idealized and only an artistic rendition of the real anatomy. There are also cases in which no consensus is reached among anatomical researchers regarding the exact configuration of certain anatomical structures, but current resources have no way of indicating potential alternatives. Even if a consensus is reached, many outdated books remain accessible and it might prove futile trying to update all resources.

The Unified Anatomical Human (UAH) [4, 5] is a software system that was developed to allow anatomical researchers to store, query and visualize heterogeneous anatomical data stored in a schemaless database. In the associated UAHViewer prototype, users can interact with this database using a 3D view linked with 2D slice-viewers. Heterogeneous spatial and non-spatial data from different sources can be represented, as well as relations between this data. Currently, the UAH database is focused on the human female pelvis and contains a 3D atlas of the pelvis built through manual segmentation of the Visible Korean Human dataset [6–8].

In this paper, we present the Online Anatomical Human (OAH), an online viewer and annotation system for anatomical information developed for educational purposes. Similarly to the UAH system, it is based on real human anatomy and incorporates medical image data in linked 2D orthogonal slice-based views and 3D views, but runs entirely in a web browser. This property makes our OAH solution a good choice for educational purposes, as it is basically available to anyone with an internet-enabled device. A screenshot is shown in Figure 4.1.

Within OAH, a user can interact with the 3D model and 2D slices of a real medical data set to gain a better understanding and improve the mental mapping between the two domains. The latter is a critical skill since imagining modalities have been introduced in clinical practice. While interpreting medical scans, such as MRI and CT, used to be done mainly in radiology departments, nowadays scans are available in digital form throughout the hospital and many specialists examine the images directly. Consequently, for medical students, it is important but challenging to mentally reconstruct a 3D model based on 2D images, as well as to relate 3D anatomy to 2D medical imaging data.

Additionally, by allowing medical experts to share their knowledge via annotations and hyperlinks, our representation becomes an online database, able to keep track of a variety of information and the most recent developments. Finally, the 3D annotation system has several additional purposes; it can be used by teachers to indicate anatomical landmarks, lines and regions, or to acquire different expert opinions on, e.g., optimal incision lines for surgical procedures. The annotation can also be used by students, to verify their understanding of the presented material or to ask questions. In consequence, our solution is highly valuable for educational purposes. Finally, our approach can be interacted with by using a Leap Motion. Hereby, the application can be controlled via hand and finger gestures. Besides a general ease of use, this possibility is of advantage when using the OAH during a medical practical.

Specifically, the contributions of this chapter are as follows:

- A prototype tool, the Online Anatomical Human, capable of providing the user with linked 2D and 3D anatomical information for educational purposes via a web browser;
- An open platform to enable users to annotate anatomical structures in 3D, leading to an ever-growing wealth of information;
- 3D interaction techniques via hand and finger gestures.

The remainder of this chapter is structured as follows. In Section 4.2, we provide background information on teaching anatomy to medical students and present related work in Section 4.3. In Section 4.4, we introduce the Online Anatomical Human application and its educational value. Section 4.5 provides implementation details and presents methods developed to support annotation. Finally, we conclude the paper and present an outlook on potential future directions in Section 4.6.

4.2. BACKGROUND

Human anatomy is an important part of the curriculum for future doctors, but also radiographers, dentists, physiotherapists, nurses, paramedics, and many other health professionals. Anatomy education typically consists of several teaching methods varying from lectures to dissection of bodies (gross anatomy).

One way of teaching anatomy is a regional approach, where the human body is divided in physical regions, such as the abdomen and lower extremities. Organs and tissues are also studied in relation to one another [9]. Alternatively, one can use a systemic approach, where the focus is on systems, such as the vascular or respiratory system as a whole.

To support the learning process, a wealth of written and visual material exists (textbooks, diagrams and photographs). Additionally, an important component of the curriculum is the more practical way of learning by dissecting and inspecting human bodies. During a dissection, the students can explore a human body themselves, whereas, during a prosection, the students either observe an experienced anatomist perform a dissection or they examine a previously-dissected specimen themselves.

An important observation is that textbook anatomy is vastly different from what is experienced in dissection, which in turn is different from what is experienced in clinical practice (when examining patients directly and using medical imaging). However, this link between the different modalities is vital in the education of the students. While books can present illustrations of 3D patient anatomy or 2D medical images from scanners, they are static and cannot present a direct and consistent link between both. In dissection, only 3D information is available and the link to corresponding 2D information can be unclear.

4.3. RELATED WORK

In this section, we describe several tools with a similar educational goal, followed by a description of related literature around 3D structure annotation. Finally, we shortly




Figure 4.1: Main window of OAH: The structure list is available on the left for selection and showing or hiding (groups of) structures. In the center, the 3D model is shown. On the right, orthogonal 2D sliceviewers are available.

discuss 3D graphics rendering in a web-browser and alternative input devices.

3D anatomy education Estevez et al. evaluated the use of 3D visualization in teaching neuroanatomy and concluded that 3D physical modeling is an effective method for teaching spatial relationships of brain anatomy [10].

Web interfaces for anatomy atlases have been around for over ten years. In the beginning, these interfaces did not work well enough in terms of interaction speed to replace existing atlases in books [11]. Only recently, using new techniques, such as WebGL, full 3D applications are feasible in the browser. Zygote Body was introduced by Blume et al. [12] in 2011, at the time under the name Google Body. The tool features a full overview of the human body, male and female. It originally contained a feature to add an annotation point, consisting of a position and a note. Single annotations could be shared via a special url. The shown anatomy stems from a model designed by artists and not from medical imaging data, in consequence, it is not always representative of real patient anatomy.

Qualter et al. introduced the Biodigital Human [13]. This tool has since grown to be the most feature-rich WebGL anatomy platform. It again only supports pinpoint annotations (with a label and description). Region annotations or line segments are not available, often making it impossible to give a detailed indication of the relevant part of a structure. The complete view and setup, including annotations, can be shared with other users. The model is again not created from real anatomical data, but a manually-modeled artistic generalization. 2D medical images are therefore not available.

InnerBody [14] is another human body exploration tool. It does not feature a 3D model, but a collection of illustrations. The tool offers various anatomical systems, such as the vascular or lymphatic system, to be explored in detail. For each system, a hierarchical list of structures and sub-structures is available. When a structure is selected, an

interactive 2D illustration is shown.

In the Anatomography project, anatomical polygonal surfaces are made freely available via the dataBodyParts3D website for educational purposes [15]. Acland's video atlas of human anatomy provides videos of dissections online via a subscription service [16]. 4D Anatomy Online is a subscription-based online education platform that provides dissection-simulation resources based on dissection images [17]. Furthermore, apps have been developed aimed at mobile phones and tablets such as Anatomy 4D by DAQRI and the Human Anatomy Atlas by Visible Body.

The main difference with other applications is that our approach supports both 2D and 3D views using real medical imaging data.

Annotation of 3D structures When looking at depictions of anatomy in books, such as the works by Agur and Dalley [18], Platzer et al. [19] and Attene et al. [20], several types of annotations are used. For larger regions, an internal label is shown directly on top of the depicted structure, whereas interesting landmarks can be labeled externally with a connecting line to the point of interest [21]. In the work by Platzer et al., color is used to distinguish different tissue types, whereas Attene et al. use color to label sub-parts in a larger model covering the entire human body. To allow users to make similar annotations, we examine several methods and systems for the annotation of surface meshes.

A digital example of interactive direct painting on a mesh, within a web-browser, is DeathPaint [22]. In principle, this technique could be used to annotate regions on anatomical structures. Due to the restriction to WebGL, Deathpaint uses a color-picking technique to choose a drawing location. Our approach uses a similar technique, but does not require pre-rendered textures and works even for new models without any pre-processing. Other tools, such as Space Pen [23] allow users to annotate 3D models with text comments positioned in 3D space by directly drawing on a model. It is web-based, but requires external plugins and runs only in older versions of Internet Explorer. Gorgan et al. [24] also aimed at medical education and support a pen-based annotation in 2D and 3D, but the software needs to be installed and is not web-based. Furthermore, all these solutions only support freehand annotation.

Drawing lines on 3D surfaces is not straightforward. Bonneau and Hahmann [25] discuss several methods. One way is to let the user draw a line segment on the screen, which is point sampled. Through each sample point, a ray is cast to find the corresponding surface intersection (alternatively, a projection method can be used [26]). However, these surface points need to be connected and might be far apart at occlusions or in high-curvature regions. The main method of Bonneau and Hahmann only considers the two segment endpoints and use an adaption of Dijkstra's algorithm [27] to find a shortest path on the surface. This path is then smoothed to yield the final annotation. Unfortunately, the generated curve will usually deviate from the user-defined line segment and not appear straight.

3D graphics in a web-browser An elaborate survey of 3D web graphics techniques can be found in the survey by Evans et al. [28]. In order to reduce the stress on client computers, Blazona and Mihajlovic [29] proposed a method of in-server rendering. In this approach, the rendering process is performed by the server and only the resulting im-



Figure 4.2: When selecting an element in the original pelvic model (left), the other structures become semi-transparent. Hereby, we optimize visibility, while maintaining an anatomical context (right).

4

agery is sent to the client. This solution increases the load on the server when many clients are present. The release of WebGL made rendering complex 3D medical data within the browser feasible [30]. WebGL quickly gains in popularity and is supported by all major modern browsers. Since our models are relatively lightweight, and we expect many clients in a Massive Open Online Course (MOOC) context, WebGL suits our needs best.

Input devices A thorough overview of several input devices for 3D interaction can be found in the survey by Chris Hand. [31]. Standard interaction is based on a computer mouse and keyboard. The mouse is essentially a 2D device and in case a 3D scene is rendered, interaction can become counter-intuitive. Gallo et al. [32] proposed a method for model interaction with a Wiimote, the wireless Nintendo Wii controller. Since this device can be controlled in the air, the user is not limited to the typical 2D mouse interaction. 3D devices usually map motion more intuitively on 3D interaction. The Leap Motion, which was released in 2013, captures hand gestures in the air. Unlike when using the Wiimote, there is no need to hold anything in your hands. Weichert et al. [33] concluded that the device is robust enough to use in applications requiring a high level of accuracy for in-the-air movement. The average measured accuracy was 0.7 mm, while the accuracy attainable by a human hand is 0.4 mm on average.

OAH was inspired by the research and tools described above, but none of the mentioned tools use real anatomical data. The OAH features a 3D model combined with the 2D medical image data from which it was created. It is also the first such system to allow region and line annotations on anatomical 3D models in a standard web-browser. Besides keyboard and mouse, our application also supports 3D input devices, such as the Leap Motion to enable interaction without touching a device, which is crucial during medical practicals.

4.4. THE ONLINE ANATOMICAL HUMAN

This section describes the application and design of the OAH system. First, we describe the data used by our approach. We then describe the interface and available features,

including the annotation system. Finally, we describe the interaction options.

4.4.1. DATA

OAH uses the VSP atlas, which is described in chapter 3. The final model consist of 28 structures divided into 6 groups. When all structures are loaded, the complete model consists of 666992 faces and 323280 vertices. For the 2D slice-based views, we have 910 transversal, 134 coronal and 246 sagittal images. To reduce loading times, we downsampled the images and added an internal setting to limit the number of images loaded in the browser. The latter can also be automated based on the connection speed of the user. Instead of showing cryosectional slices, it is also possible to show other imaging modalities (e.g., CT and MRI).

4.4.2. FEATURES

The OAH interface can be seen in Figure 4.1. Users are able to interact with the 3D model and browse through the linked 2D cryosectional and segmentation slices on which the 3D model is based. Structures can be selected directly in the 3D view or by choosing them from a nested list of all available elements. Whenever a structure is selected, the non-selected elements become semi-transparent to prevent occlusion, while preserving anatomical context (see Figure 4.2). Further, query results from the UAH database, such as literature associated to the selection, become available to the user. Additionally, the user can decide to hide or unhide a selected structure.

Once a user finds an interesting view of the 3D model, it is possible to store the camera position. Such camera views can also be retrieved and make the camera move from its current view to the selected one via a tweening animation. These interactions help students familiarize themselves with the anatomical structures and their relations. They can learn their names and even find related literature easily.

To ensure that students understand the link between 2D medical imaging and 3D anatomy, the 2D cryosectional images, which were used to create our 3D model, as well as the corresponding segmentation labels, are made available in the OAH in transversal, coronal and sagittal directions. For each of these three directions, a 'sliceviewer' is available (see Figure 4.3). Each view is linked, hence, updates in one view are automatically propagated to all other views. One update relates to the movement of a cross-haired marker shown in each 2D sliceviewer and the 3D view. The axes of these markers are color-coded according to corresponding global directions. Whenever the crosshair is moved, all point representations in all views move accordingly, including the 3D point in 3D space. In each sliceviewer, one axis is invisible and represented by the stack of images associated to it. The user can scroll through the image stack using the mousewheel.

Our system also supports annotations. Teachers can use these to provide details regarding the structures, such as landmark names. Additionally, the annotation system can be used to evaluate student performance, by asking them for instance to indicate certain landmarks on a blank model. The same system is also suitable for the students to indicate relevant areas when asking questions about structure details.

Examples of different annotations made in the OAH are shown in Figure 4.4. We offer three solutions to place annotations directly on the 3D mesh, which are detailed below.




Figure 4.3: The sliceviewer provides three orthogonal slicing planes in the transversal, coronal and saggital directions through the cryosection and segmentation data in a linked 3D (left) and 2D (right) representation. The crosshair lines show the currently selected point.

Landmarks are single points on the surface of an anatomical structure, used to either label an exact point, or to sub-label a structure without a specific region. We illustrate such landmarks with a small sphere at the selected point, so that it is visible from both sides of the surface.

Regions are used to coarsely annotate regions of structures. It can be useful when, for example, two parts of a structure need to be distinguishable, but the actual border is not evident. The method uses a brush, which makes annotations of large areas simple and quick. In our brush metaphor, all vertices within the brush diameter are colored and a fadeout can be defined to indicate the precision of the annotation. When two parts of the model are close together, we ensure that our brushing method does not overflow into the disconnected region of the model.

Line annotations are used when precision is needed. While regions have a fuzzy fade-out, lines are accurate. They can be set to form a region contour by connecting start and endpoints. Line segments can be drawn in two modes. The first mode, only uses a line segment drawn by the user. Our approach finds a path on the surface, whose projection matches the user input, while handling surface discontinuities. The second mode adds such line segments continuously, while the pointer is moving over the model.

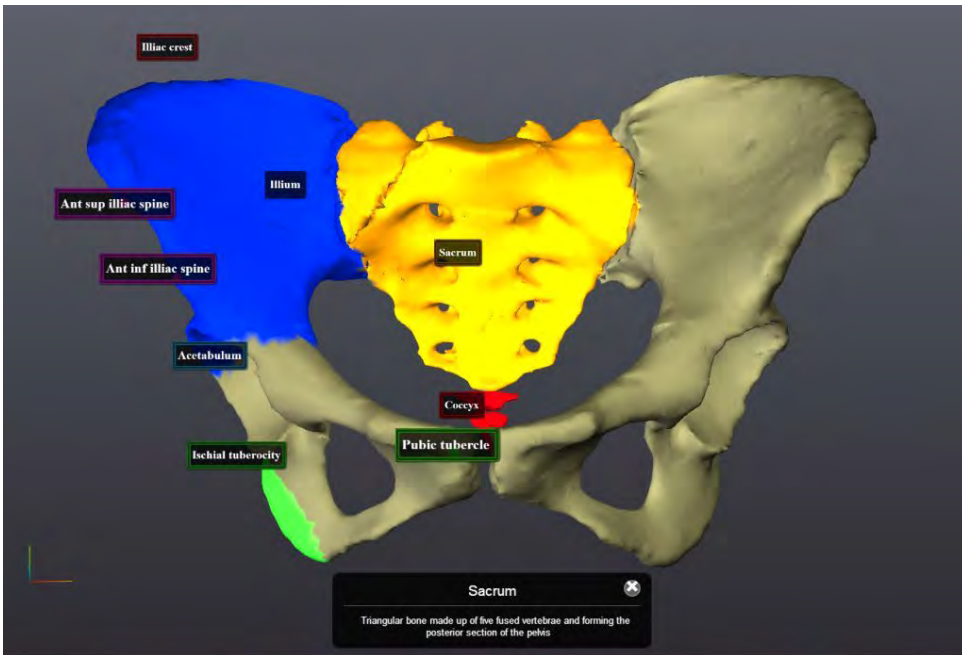


Figure 4.4: In this annotation view, pelvic bones are shown with landmark, region and line annotations.

4.4.3. INTERACTION

Our system can be used entirely with a mouse and the user can click (and drag) to navigate and define annotations. However, one particularity of the OAH is its support of the Leap Motion. It is a sensor capturing input via hand and finger gestures, which can also be used alongside the keyboard and mouse. It registers full 3D positioning, making it particularly suitable for 3D interaction. The fact that the device works touchless, is particularly useful if OAH is used during anatomical dissection classes because students do not need to clean their hands, as they would using standard input devices. Additionally, the Leap Motion provides a playful way for students to interact with the models and can be used for demonstration purposes.

When using the Leap Motion in conjunction with a mouse, the user can rotate and zoom the camera with one hand, while the other hand is used to control the mouse for annotations. This process is similar to holding the model in one hand, while painting with the other. It is also possible to annotate using only the Leap Motion, when involving both hands. To represent a click the user opens the left hand and points with one finger of the right hand to the screen. When the left hand is closed, the pointer will act as a click was performed. Controls can be inverted to support left handed interaction.

4.4.4. COURSE SETUP

The OAH system is used in an upcoming Massive Open Online Course (MOOC) as an interactive part of a Coursera course on human anatomy. This online course is aimed

at large-scale participation and open access via the internet to students worldwide. Instructors provide questions on the anatomy of the pelvis, which students need to answer by interacting with the OAH. Through these questions, which can have different goals, mutual relations between organs and their clinical relevance are interactively explored by students.

A first goal is to focus on important organs in order to better illustrate their shape, extent, and location. For example, 'Show in three cross-sections that the Douglas Pouch is the deepest part of the peritoneal cavity. Which organs border this cavum?'

Another type of questions deals with relations between organs. For example, 'Explain what peritonitis is and what symptoms it causes'. The answer here would be abdominal pain, but also vaginal pain, due to the close proximity of the peritoneum to the vagina.

Questions can also be oriented towards risk investigations, which can lead to awareness of typical pathologies in certain organs. For example, 'A small kidney-stone is able to travel through the ureter to the bladder. Where can it get stuck along the way and what type of pain can this cause?'. Here, the students will find that the most critical point is crossing of the ureter and the Arteria Iliaca.

Further, the students can be trained in communication. For this purpose, they can be asked to first identify all possible illustrations from the data set and present the finding to a colleague or patient. For example, 'Show the potential problem points for a kidney-stone in the 3D model as well as on all cross-sections revealing these points. Then choose the representation that is most effective in illustrating this case and present it to the patient.'

Finally, our system also allows us to verify the quality of the 3D model by comparing it to the segmented medical images. Similarly, students can be trained to find deviations from scanned data, hereby, learning to better interpret the scans.

Through interaction with the OAH, these various types of questions can lead students to gain a deeper insight into anatomy. Hereby, they will not only be familiarized with the organs themselves, but also develop an understanding of their relationships and the symptoms caused by various pathologies.

4.5. IMPLEMENTATION

In this section, we describe the techniques and libraries that were used during the development of OAH. Further, we give insights into implementation details.

Libraries For the rendering of the 3D views, we use WebGL to utilize the power of the graphics card inside of a web-browser environment. Since 2011, WebGL is rapidly gaining in popularity and support by all major browsers, including those in mobile devices. The project relies on Three.js, a javascript framework that simplifies the way of working with WebGL. A challenge is to keep the system as lightweight as possible to ensure a smooth execution even on tablet computers. Consequently, models should be of reasonable size, while maintaining enough detail to be anatomically correct. We converted the models to a JSON format which can be directly imported in Three.js.

To translate movement detected by the Leap Motion, we use Leap.js and threeleap-controls. The latter converts gestures to camera transformations. Leap.js is required to

get information from the device and used it to map hand and finger movement to the pointer control.

jQuery and jQuery UI were used for the interface. jQuery UI offers several widgets to create the menus and contains a widgetfactory to build custom stateful widgets. The latter was used for the 2D sliceviewers, which can be displayed within the 3D view as sliceplanes to illustrate their spatial relation to the model. There is always one 3D intersection point of these slices. Whenever a change is made in one of the sliceviewers, the location of this point updates in all views, which is triggered by a DOM event. As in the system presented in Chapter 2, literature is linked to the anatomical structures and annotations. We maintain this information in a MongoDB database and PHP is used to communicate with this database.

Annotation Implementation At the core of our annotation functionality is the selection of a point on the 3D surface. Raycasting is the most-commonly used technique. Although accurate, a ray test against all primitives is too slow. While octrees or spatial indexing methods [34] are standard, they increase memory usage for larger models and lead to increased preprocessing times, which makes it a suboptimal choice for mobile devices. Instead, we turned to offscreen rendering.

The drawback of Slack's offscreen method [22] is the use of a predefined color texture, which requires texture mapping of the models and can introduce artifacts, especially at texture seams. Our method does not require such a pre-generated texture. We use an offscreen render target into which the model is rendered with a distinct color for each face. The color of the pixel underneath the cursor then directly reveals the selected face. To this extent, we convert the model's face indices to hexadecimal values, which are used to determine the corresponding RGB value [35]. For example, index 8128 is 0x1fc0 in base16, which would translate to the RGB triplet R = 0, G = 31, B = 192. Using 8-bit color depth, $256^3 = 16777216$ faces can be distinguished, which is enough for our purposes. Although this technique requires an extra render pass, the buffer is only a single pixel in size, making it highly efficient. If needed, one can then compute the actual intersection point by testing only the recovered single face against the ray.

For our other two annotation types (regions and lines), we rely on topological information. Unfortunately, in Three.js, the connectivity of models is not directly accessible. Each face has a list of its vertices, but there is no easy way of finding the face connected to a vertex, or to find vertices forming an edge. We used the work of Stemkoski [36] to transform the geometry into a topological structure. Hereby, mesh traversal and path finding become much simpler.

For the region brushing, we need to find all vertices within the diameter of the brush and color them according to their distance. We first detect the central face and then perform a forward search approach on the mesh to select all vertices within the brush radius.

For the line annotation, we first project the two endpoints of the segment on the surface. We then virtually cut the entire model by the plane defined by camera center and the drawn segment. Hereafter, we launch a shortest path algorithm from the first to the second projected endpoint, traveling only on the newly created intersection edges between the mesh and the cutting plane. As all traversed edges do project exactly to the

drawn line segment and the shortest path algorithm only proceeds on connected edges, the resulting path on the surface is continuous, precise, and stable under high curvature.

4.6. CONCLUSION AND FUTURE WORK

In this chapter, we presented the Online Anatomical Human (OAH), an online viewer and annotation system that is of interest to both students and experts. It makes 2D and 3D anatomy based on medical imaging data available to anyone with an internet connection, as it runs entirely inside of a web-browser and avoids any kind of installation procedure. Besides interactive exploration, users can annotate the model directly in 3D by adding landmark, region and/or line annotations. We support the use of the Leap Motion, enabling an intuitive and playful way to interact with the system. The OAH has been used in a Massive Open Online Course (MOOC) to allow students to learn about the relation between 2D and 3D pelvic anatomy in an interactive environment.

In the future, we intend to use our OAH system to perform large-scale user studies with students from the MOOC. Furthermore, since experts might disagree regarding certain annotations, visualizing combined opinions and offering intuitive ways to explore this information is an interesting future challenge.

REFERENCES

- [1] N. N. Smit, C.-W. Hofstede, A. C. Kraima, D. Jansma, M. C. Deruiter, and E. Eise-mann, *The online anatomical human: Web-based anatomy education*, in *Proceedings of Eurographics - Education papers* (2016) accepted, to appear.
- [2] C. Hofstede, *The Online Anatomical Human: An online browser and annotation system for real human anatomy*, Master's thesis, Delft University of Technology, the Netherlands (2014).
- [3] J. F. Nunn, *Ancient egyptian medicine* (University of Oklahoma Press, 2002).
- [4] N. N. Smit, A. C. Kraima, D. Jansma, M. C. Deruiter, and C. P. Botha, *A unified representation for the model-based visualization of heterogeneous anatomy data*, in *EuroVis - Short Papers*, edited by M. Meyer and T. Weinkauff (Eurographics Association, Vienna, Austria, 2012) pp. 85–89.
- [5] N. N. Smit, A. C. Kraima, D. Jansma, M. C. Deruiter, and C. P. Botha, *The unified anatomical human (beta): model-based representation of heterogeneous anatomical data*, in *Proceedings of the 3D Physiological Human (3DPH) Workshop, CASA* (2012).
- [6] J. S. Park, M. S. Chung, S. B. Hwang, Y. S. Lee, D.-H. Har, and H. S. Park, *Visible korean human: improved serially sectioned images of the entire body*, *Medical Imaging, IEEE Transactions on* **24**, 352 (2005).
- [7] J. S. Park, M. S. Chung, S. B. Hwang, B.-S. Shin, and H. S. Park, *Visible korean human: its techniques and applications*, *Clinical Anatomy* **19**, 216 (2006).

- [8] J. S. Park, Y.-W. Jung, J. W. Lee, D. S. Shin, M. S. Chung, M. Riemer, and H. Handels, *Generating useful images for medical applications from the visible korean human*, *Computer methods and programs in biomedicine* **92**, 257 (2008).
- [9] H. Gray, *Anatomy of the human body* (Lea & Febiger, 1918).
- [10] M. E. Estevez, K. A. Lindgren, and P. R. Bergethon, *A novel three-dimensional tool for teaching human neuroanatomy*, *Anatomical sciences education* **3**, 309 (2010).
- [11] S. Bradley, C. Rosse, and J. F. Brinkley, *Web-based access to an online atlas of anatomy: the digital anatomist common gateway interface*. in *Proceedings of the Annual Symposium on Computer Application in Medical Care* (American Medical Informatics Association, 1995) p. 512.
- [12] A. Blume, W. Chun, D. Kogan, V. Kokkevis, N. Weber, R. W. Petterson, and R. Zeiger, *Google body: 3d human anatomy in the browser*, in *ACM SIGGRAPH 2011 Talks* (ACM, 2011) p. 19.
- [13] J. Qualter, F. Sculli, A. Olikier, Z. Napier, S. Lee, J. Garcia, S. Frenkel, V. Harnik, and M. Triola, *The biodigital human: a web-based 3d platform for medical visualization and education*. *Studies in health technology and informatics* **173**, 359 (2011).
- [14] T. Taylor, *InnerBody*, <http://www.innerbody.com> (1997), [Online; accessed 10-March-2016].
- [15] The Database Center for Life Science, *Anatomography - BodyParts3D*, <http://lifesciencedb.jp/bp3d/>, [Online; accessed 10-March-2016].
- [16] Robert Acland, *Acland's Video Atlas of Human Anatomy*, <http://aclandanatomy.com/>, [Online; accessed 10-March-2016].
- [17] 4D Anatomy Inc., *4D Anatomy Online*, <https://www.4danatomy.com/>, [Online; accessed 10-March-2016].
- [18] A. M. Agur and A. F. Dalley, *Grant's atlas of anatomy* (Lippincott Williams & Wilkins, 2009).
- [19] W. Platzer, G. Spitzer, H. Fritsch, W. Kühnel, H. Leonhardt, W. Kahle, and M. Frotscher, *Sesam atlas van de anatomie* (Sesam, 2007).
- [20] M. Attene, F. Robbiano, M. Spagnuolo, and B. Falcidieno, *Characterization of 3d shape parts for semantic annotation*, *Computer-Aided Design* **41**, 756 (2009).
- [21] T. Ropinski, J.-S. Praßni, J. Roters, and K. Hinrichs, *Internal labels as shape cues for medical illustration*. in *VMV*, Vol. 7 (Citeseer, 2007) pp. 203–212.
- [22] J. Slack, *Cartelle deathpaint study*, <http://www.cartelle.nl/deathpaint/study/>, [Online; accessed 10-March-2016].

- [23] T. Jung, M. D. Gross, and E. Y.-L. Do, *Annotating and sketching on 3d web models*, in *Proceedings of the 7th international conference on Intelligent user interfaces* (ACM, 2002) pp. 95–102.
- [24] D. Gorgan, T. Stefanut, and B. Gavrea, *Pen based graphical annotation in medical education*, in *null* (IEEE, 2007) pp. 681–686.
- [25] G.-P. Bonneau and S. Hahmann, *Smooth polylines on polygon meshes*, in *Geometric modeling for scientific visualization* (Springer, 2004) pp. 69–84.
- [26] O. Tolba, J. Dorsey, and L. McMillan, *Sketching with projective 2d strokes*, in *Proceedings of the 12th annual ACM symposium on User interface software and technology* (ACM, 1999) pp. 149–157.
- [27] E. W. Dijkstra, *A note on two problems in connexion with graphs*, *Numerische mathematik* **1**, 269 (1959).
- [28] A. Evans, M. Romeo, A. Bahrehmand, J. Agenjo, and J. Blat, *3d graphics on the web: A survey*, *Computers & Graphics* **41**, 43 (2014).
- [29] Ž. Mihajlović and B. Blazona, *Visualization service based on web services*, *CIT. Journal of Computing and Information Technology* **15**, 339 (2007).
- [30] J. Congote, A. Segura, L. Kabongo, A. Moreno, J. Posada, and O. Ruiz, *Interactive visualization of volumetric data with WebGL in real-time*, in *Proceedings of 3D Web Technology* (ACM, 2011) pp. 137–146.
- [31] C. Hand, *A survey of 3d interaction techniques*, in *Computer graphics forum*, Vol. 16 (Wiley Online Library, 1997) pp. 269–281.
- [32] L. Gallo, G. De Pietro, and I. Marra, *3d interaction with volumetric medical data: experiencing the wiimote*, in *Proceedings of Ambient media and systems (ICST, 2008)* p. 14.
- [33] F. Weichert, D. Bachmann, B. Rudak, and D. Fisseler, *Analysis of the accuracy and robustness of the leap motion controller*, *Sensors* **13**, 6380 (2013).
- [34] D. Meagher, *Geometric modeling using octree encoding*, *Computer graphics and image processing* **19**, 129 (1982).
- [35] E. Eisemann, H. Winnemöller, J. C. Hart, and D. Salesin, *Stylized vector art from 3d models with region support*, *Computer Graphics Forum (Proc. of EGSR)* **27** (2008).
- [36] L. Stemkoski, *Three.js topology*, <https://github.com/stemkoski/stemkoski.github.com/blob/master/Three.js/js/topology.js>, [Online; accessed 10-March-2016].

5

VISUALIZING ANATOMICAL VARIATION IN BRANCHING STRUCTURES

In the Chapter 4, we described how to reach students worldwide to educate them in pelvic anatomy, but this teaching method does not accommodate for topological anatomical variations in branching structures. Anatomical variations are naturally-occurring deviations from typical human anatomy. While these variations are considered normal and non-pathological, they are still of interest in clinical practice for medical specialists such as radiologists and transplantation surgeons. The complex variations in branching structures, for instance in arteries or nerves, are currently visualized side-by-side in illustrations or expressed using plain text in medical publications.

In this work, we present a novel way of visualizing anatomical variations in complex branching structures for educational purposes: VarVis. We propose a novel graph representation to provide an overview of the topological changes. Our solution involves a topological distance metric to deliver insights into the similarity of variations. Further, the graph is linked to associated 3D surface representations of patient-specific CT-data or anatomical illustrations. We also incorporate additional information regarding the probability of the various cases. Our solution has several advantages over traditional approaches, which we demonstrate by an evaluation.

This chapter is based on the following publication: **Noeska Smit**, Annelot Kraima, Daniel Jansma, Marco DeRuiter, Elmar Eisemann, and Anna Vilanova, "VarVis: Visualizing Anatomical Variation in Branching Structures", *Proceedings of EuroVis Short Papers*. 2016 [1]

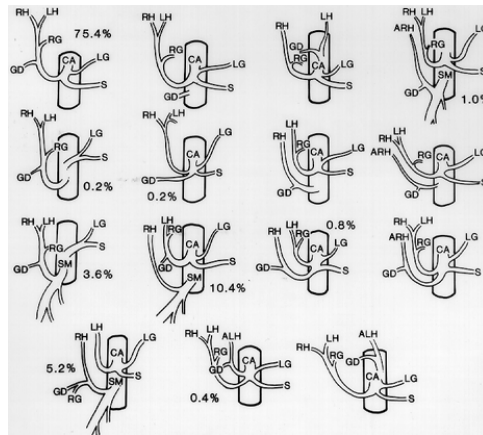


Figure 5.1: Gastroduodenal artery variation illustrations [6]

5

5.1. INTRODUCTION

Anatomical variations are natural variations in anatomy, which are considered normal and non-pathological. The variations of anatomical structures such as bones and organs consist mainly of geometric differences in size and shape. For branching structures such as the vascular and nervous system, the variations can be more complex and include changes in topology. Among these structures, veins vary the most, while nerves are the least variable [2]. Prominent examples of complex variations are missing structures, extra structures, different connectivity patterns, or different relations to surrounding structures. First introduced by Vesalius in *Humani Corporis Fabrica* (1542), the variations are still clinically relevant today [3]. Knowledge of existing variations and their occurrence frequencies are essential in order to correctly diagnose and safely treat patients [4]. Vascular variations specifically are important for many interventional procedures and, thus, for radiologists, vascular surgeons, transplantation surgeons, but also neurosurgeons who perform surgery on aneurysms and malformations in the brain. For example, in hepatic transplant surgery, preoperative knowledge of vascular and biliary anatomical variations is mandatory for surgical planning. It helps reduce postoperative complications for both the organ donor and the recipient [5]. After initial discussions with collaborating domain experts, we elicited the aspects of interest to them in an educational setting. Relevant knowledge includes finding out what types of variations exist, how often they occur, which of the variations are similar, and what the differences between variations are. Especially the quick discovery of the most common variations is relevant to them because the chances of encountering them in clinical practice are the highest.

The current way of presenting and acquiring knowledge on variations is either based on textual information or depiction. For the first case, the reader has to construct a mental model of the exact variations, which can be challenging. More common is the use of simplified illustrations exemplifying the different layouts side by side (See Figure 5.1). In these illustrations, it is not clear how to visualize the complete range of variations entirely when many possibilities exist, leading to a significant visual overload. Further, in

case of complex anatomical structures, illustrations easily become too difficult to interpret. The links between topological changes are often challenging to understand and too manifold to organize easily in an illustration. Further, the additional information regarding the probability of the various cases in form of numeric values is hard to interpret simultaneously. Currently, nothing is done to emphasize similarity and dissimilarity between variations. Comparing two individual variations is doable, but it becomes more challenging as soon as more than two need to be compared. It is also not possible to identify groups of similar variations or to examine probabilities of group occurrences at a single glance.

We propose an interactive visualization application for anatomical variations, which allows the user to compare and explore variations of branching structures interactively for educational purposes. We consider two potential input sources: anatomical illustrations and medical imaging scans. Our visualization technique shows a simplified geometrical surface representation and a topological graph-based representation of the hierarchy. Using our application, users can gain insight into existing variations and their differences and similarities. Specifically, the contributions in this chapter are: 1) Interactive visualization of the variations and the differences between them based on their graph representations combined with a topology similarity metric. 2) A summary tree view, synthesized by a graph-matching algorithm, which allows users to explore local similarities and differences between multiple complex anatomical variations interactively. 3) An interactive prototype application, which offers various ways to explore the graphs, linked illustrations, and anatomical geometric and volumetric representations. In the following, we will discuss related work (Section 5.2), our representation and visualization methods (Section 5.3), and implementation details (Section 5.4). We then present the results (Section 5.5) and an informal evaluation with experts (Section 5.6). Finally, we conclude and give an outlook on future work (Section 5.7).

5.2. RELATED WORK

To the best of our knowledge, there is no prior work in developing interactive tools to explore anatomical variations in branching structures. Morphological variations in organs, however, were represented and visualized in the context of the VOXEL-MAN project [7, 8]. By using a medial representation of an organ population, the authors propose a compact representation of shape variation. This representation is intended to include size and shape variation, but cannot describe topological differences. Graph representations have been proposed for medical applications in previous work. In the field of neuroscience, Stam and Reijneveld [9] used a graph-theoretical analysis on the complex networks in the brain. They state that the brain can be represented as a network with a small-world structure. Rubinov and Sporns [10] presented work on using complex network measurements to characterize brain connectivity. These measurements are not applicable to branching structures, since we cannot assume anatomical structures to adhere to the same characteristics as brain connectivity. Crippa et al. [11] used graph averaging to compare multichannel EEG coherence networks. By computing the distance to an application-specific mean graph, individual trees can be quantitatively compared. Perez et al. [12] measured and quantified geometric and topological properties of vascular trees in fundus images of the human eye. Unfortunately, these quantitative measures




Figure 5.2: Two example variation trees based on illustrations (a,b) and the resulting summary tree (c).

do not help the anatomical understanding of the topological differences and similarities between the graphs.

Further related work has been done in extracting graphs from medical-imaging acquisitions. The prior art closest to our work was presented recently by Al-Awami et al [13]. In their *NeuroLines* work, they successfully apply a subway metaphor to visualize complex neuronal connectivity data. Their work focuses on scalable interactive visualizations for large data sets, but does not attempt to summarize topological variations, nor establish relationships between variations. Zhang et al. [14] propose a method to extract a graph representation automatically for the hepatic arteries in contrast-enhanced CT data. Kruszyński et al. [15–17] applied graph representations in biology in order to quantitatively compare coral colonies. In their work, they propose morphological measures to quantify the variation based on skeletonized CT-scans of corals. The output of their approach is a statistical distribution of the morphological features they measure per graph visualized in histograms. None of these methods provides a way to visually represent variations in a summarized tree-like structure based on topology.

Considerable research has been done to find ways of optimizing the layout of nodes in a graph [18–22], but a full review of these techniques is out of the scope of this paper. Won et al. [23] use a simulated annealing approach for the layout of the abdominal aortic vessel tree. Using articulated rotation, they project the 3D anatomy to a 2D plane and optimize the orientation of the branches to avoid overlap. Barsky et al. [24] propose a biologically-guided graph layout in their *Cerebral* system. Using small multiple views of different experimental conditions, they allow the biologists to analyze molecular interaction graphs. Of all these graph layout algorithms, none is directly suitable for a graph representing anatomical branching structures in an anatomically familiar and consistent way. They form a good basis and we rely on a force-directed solution for our layout algorithm.

In order to create a summary of the variations, we require a special form of graph matching. Exact graph matching is an NP-hard problem [25], but heuristic approaches have been proposed such as inexact graph matching [26]. In his survey on tree matching, Bille [27] describes methods for comparing labeled trees and discusses the concept of a tree edit distance, based on simple local operations of deleting, inserting and relabeling nodes. Closely related to graph matching is graph averaging, or finding the mean of a set of graphs. For this, the concept of a graph edit distance is used, which can be defined as the number of operations needed to turn one graph into another. Bunke and Kandel [28, 29] describe the mean of a pair of graphs as a graph that minimizes the sum of graph edit

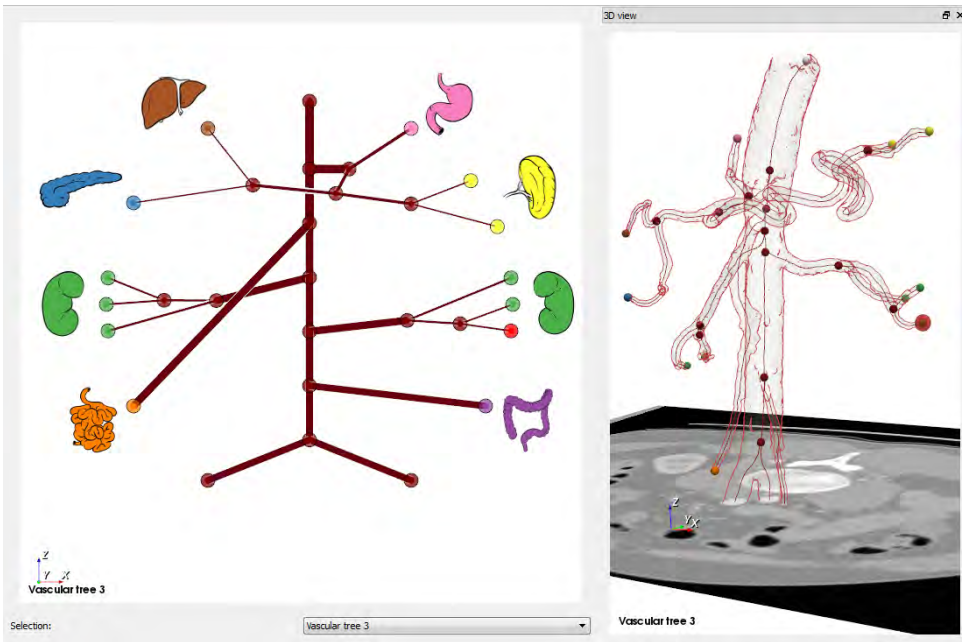


Figure 5.3: An individual variation tree on the left with the associated linked medical imaging data and surface representation on the right.

distances between them. Jiang et al. [30] propose the concept of a median graph among a set of graphs, which minimizes the graph edit distances to the whole set of graphs. For our specific anatomical variation summarization, we are looking for an algorithm that can guarantee anatomical correctness. The previous work in graph matching does not consider anatomical information, and does not preserve its correctness. The concept of a graph edit distance, however, is a good candidate for a similarity metric to define the distance between graphs.

5.3. VARVIS

5.3.1. REPRESENTATIONS

In this section, we discuss the construction of the main components for our VarVis method: the graph representation, the summary-tree generation and the topology similarity graph construction.

GRAPH REPRESENTATION

There are two possible inputs for VarVis. First, we can visualize anatomical variations based on illustrations currently used to convey variation information, such as the ones in Figure 5.1. From these illustrations, we manually create nodes for every endpoint and junction and define the connectivity by placing edges between these nodes (see Figure 5.2). A second potential input for VarVis consists of centerlines extracted from medical scans, such as contrast-enhanced CT-scans. In this way, we can describe anatomical variations directly from medical imaging data (see Figure 5.3). To transform a CT scan




Figure 5.4: The matching algorithm starts at the matched endpoints of the source and summary trees. It traverses to the root in both trees, adding labels and merging nodes along the way to the summary tree.

with extracted centerlines into a tree, we proceed as follows. We extract the centerlines based on the work of Antiga et al. [31], which uses a segmented surface representation of the vessel tree. From these centerlines we encode the topology of the branching structure in form of a graph. In our context, the resulting structure will form a cycle-free tree. Vascular structures always exhibit a source (main vessel) that is branching from there to reach specific regions. We assume that each endpoint of the graph will be feeding one specific organ or region, except for the source, which we make the root node of the tree. In this way, endpoint correspondence between different graphs is established. In case the tree is generated from medical imaging data, and the annotation in a first tree has already been done, we can rely on this information to make an initial guess for the organ membership of any unidentified endpoints in a second tree. We rely on the Euclidean distance of the corresponding 3D locations of the endpoints in a labeled volume, after registering the unlabeled source volume to it. In many cases, this step is successful, but the user can still perform manual adjustments. The latter option is also used in case of illustration-based variations, for which no volumetric data exists. Finally, we make the graph edges directed going from the endpoint nodes towards the root node. The resulting graph encodes vessel connectivity and the topological branching. To facilitate the matching process, in every node we store the labels of the endpoints they feed.

SUMMARY-TREE GENERATION

To be able to identify the exact local differences and similarities between trees, we create a summary tree of the variations for two or more given individual trees. This summary tree will encode a matching between individual tree nodes and can then be used to interactively navigate through the topological differences and similarities. An example of

Algorithm 1 : Graph matching

```

1: \\Summary tree:  $T_s = (N_s, E_s)$ , Source tree:  $T_o = (N_o, E_o)$ , N: nodes, E: edges
2:  $l = \text{MatchEndpoints}(T_s, T_o)$ 
3:  $\text{MatchGraphs}(T_s, T_o, l)$ 
4: for  $(n_s, n_o)$  in  $l$  do
5:    $\text{current}_s = n_s, \text{current}_o = n_o$ 
6:    $\text{current}_s.\text{AddLabel}(\text{current}_o.\text{id})$ 
7:    $\text{mergelist} = []$ 
8:    $\text{seenlist} = [\text{current}_o.\text{id}]$ 
9:   while  $\text{current}_o$  is not root and  $\text{current}_o$  not in  $\text{seenlist}$  do
10:      $\text{next}_s = T_s.\text{GetNext}(\text{current}_s)$ 
11:      $\text{next}_o = T_o.\text{GetNext}(\text{current}_o)$ 
12:     if  $\text{next}_o.\text{id}$  not in  $\text{next}_s.\text{ids}$  then
13:       \\case 1: no match found, merge later
14:        $\text{mergelist}.\text{Append}(\text{next}_s.\text{id})$ 
15:        $\text{current}_s = \text{next}_s$ 
16:     else
17:       \\case 2: match found, merge if needed
18:       if not  $\text{mergelist}.\text{IsEmpty}()$  then
19:          $T_s.\text{Merge}(\text{next}_s, \text{mergelist})$ 
20:          $\text{mergelist} = []$ 
21:       end if
22:        $\text{next}_s.\text{AddLabel}(\text{next}_o.\text{id})$ 
23:        $\text{seenlist}.\text{Append}(\text{current}_o.\text{id})$ 
24:        $\text{current}_s = \text{next}_s, \text{current}_o = \text{next}_o$ 
25:     end if
26:   end while
27: end for

```

two individual trees and their summary tree can be seen in Figure 5.2. The differences between the trees are collapsed into bigger nodes, while similar nodes and edges are preserved. In order to produce a summary tree encoding various variations, each variation is integrated one at a time. To explain the details, we define the following concepts: **Source tree**: the tree to be matched to the summary tree. **Summary tree**: ultimately, the tree that summarizes the variations. Initially, the structure the source tree is matched too. If empty, the full source tree is being used as a summary. **Endpoint**: a leaf node in the tree, assumed to feed a specific region or organ. The endpoints are denoted with ids by letters. For instance, 'A' could indicate the left gastric artery. **Node id (id)**: id identifying a node in a tree, consists of a tree id combined with what endpoints the node feeds, for instance, '2-AB' means that this is a node in the tree with tree id 2 and it feeds endpoints A and B.

The pseudocode for the basic algorithm can be seen in Algorithm 1. The algorithm starts its matching process at matched endpoint pairs (same organ) in source and summary tree. Please note that endpoints from one tree might not match with an endpoint

in another, in case of accessory or missing branches, a case we will ignore for now and treat later. For each matched endpoint pair, both trees are then traversed simultaneously, matching and merging, if possible, nodes into new summary tree nodes along the way to the root. We keep track of the nodes that need to be merged in a *mergelist* and merge them together when we find the first pair of nodes that match. Two nodes match when they feed the same endpoints, as denoted by their labels (please recall from the previous section, that all nodes in both trees are labelled).

We can divide the possible actions of this algorithm in the following two basic cases: **Case 1 - Merge:** When no match is found between an examined source tree node and a summary tree node, add the next summary node to the *mergelist* and only continue to the next node in the summary tree. **Case 2 - Match:** When a match is found: add the *mergelist* elements and the label of the source node to the matched summary tree node. Advance in both trees and empty the *mergelist*.

We illustrate the basic graph matching process with an example in Figure 5.4. Initially, endpoints A and 2-A are matched in the summary and source tree respectively. The same holds for the endpoint pairs B/2-B and C/2-C. As a first step, the label of the source node 2-A is added to node A in the summary tree. The node id of 2-ABC, ABC, is in the node id of ABC, so we are in case 2 and have a match. There are no nodes in the *mergelist*, so we simply add the 2-ABC node label to the ABC node of the summary tree and move on in both trees until we reach the root. For the next matching endpoint pair, B and 2-B, we again add the 2-B node label to the ids of B. We find that the node id for the next summary node, BC, does not contain the ABC we are looking for. We are now in case 1 and add BC to the *mergelist* and move on only in the summary tree. In the next step, we find matching node ids again, ABC in the source and ABC in the summary tree. This means we are in case 2, however, this time we have nodes to merge. The BC node gets merged with the ABC node. For the final endpoint pair, C and 2-C, we once again add the 2-C label and we find that the current source node is already in the summary tree, which means that the rest of the tree is already matched for this endpoint, and we can stop traversing.

After the full matching algorithm is executed, the summary tree has been enriched with the matched nodes of the source tree, thus forming a summary. Where the trees are different and cannot be matched, nodes merge into bigger nodes. These merged nodes can be recognized by the higher amount of labels that they have.

The basic algorithm above does not account for accessory branches, which generate extra endpoints in one of the trees that cannot be matched. Any unmatched endpoints in the summary tree can be left as is, since they will not be traversed in the basic algorithm. For unmatched endpoints in the source tree, however, additional steps are needed to add the variation to the summary tree, leading to a **Case 3: Add:** There is an extra branch in the source tree. It causes an extra junction node, where one adds a node to the summary tree between the current and next and advances in both trees.

Figure 5.5 illustrates the process for matching additional branches. The source tree has an accessory branch to region C, which cannot be matched to the summary tree. While traversing the matching endpoints, node 2-ABC is added as an additional junction point, since the reduced node id AB is encountered twice in the source tree. In the final phase of the algorithm, the new node 2-C is added to the summary tree and connected




Figure 5.5: In the first pass of the algorithm, a junction is found created by an accessory branch in the source tree. This junction is added to the matching tree and in the second pass the new subtree is added and attached to the new node.

by an edge to 2-ABC.

In case of complex variations, there is a risk of merging too many nodes, until only a single node connected to the endpoints remains. To prevent this often undesirable outcome, we can introduce a threshold on the number of merges that are allowed per endpoint pair. If the threshold is exceeded, we treat the corresponding endpoint as an additional branch and re-run the algorithm. This creates two outgoing edges from the endpoint, which means that we need to consider all possible paths during subsequent graph traversal and pick the paths that causes the least merge operations. The threshold can be controlled by the user to define the level of merging that is desirable.

After the tree matching is complete, we can highlight local differences and similarities based on the above matching scheme. There are three possibilities. There can be a one-to-one match between nodes, which represents an exact match. There can also be a mismatch, either because a node is not available in the other tree(s) or because the nodes are different. Matched, extra and merged nodes are highlighted in green, purple and orange respectively (see Figure 5.6).

TOPOLOGY SIMILARITY GRAPH

To help the user in assessing, which anatomical variation trees are similar, and which are different, we create a topology similarity weighted graph, where each node represents a variation instance. The graph provides insight on how variation trees are related

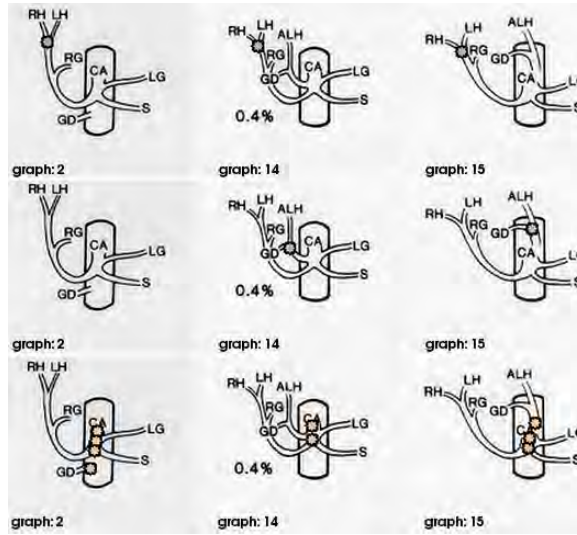


Figure 5.6: Matching nodes are highlighted in green (top), extra nodes in purple (center) and merged nodes in orange (bottom).

via its connectivity (disconnected nodes indicate outliers different from all others - see Figure 5.7, top left). Specifically, an edge in this graph indicates that the similarity of a tree pair is above a user-defined threshold. We have defined this similarity metric based on the number of identical nodes in a set of trees. By identical, we mean that the labels indicating which endpoints are fed are the same. To calculate the distance between a pair of trees t_1 and t_2 , we find the intersection of the sets of all node labels in each tree, s_1 and s_2 , and normalize this measure by dividing it by the maximum number of nodes:

$$\text{similarity}_{t_1, t_2} = \frac{|s_1 \cap s_2|}{\max(|s_1|, |s_2|)} \quad (5.1)$$

This similarity metric is 1 when trees are identical and 0 when the labels of all nodes are different. The similarity metric can be replaced by other metrics if desired, such as the graph edit distance [27].

5.3.2. VISUALIZATION

In this section, we discuss the presentation components of VarVis. We start with the graph layout algorithms, afterwards we describe the design decisions made for the summary tree and the topology similarity graph. Finally, we will explain the anatomical view linking 2D and 3D representations.

GRAPH LAYOUTS

We opt for an easy-to-understand 2D representation based on the individual vessel trees and summary trees to visualize the topological configurations. We derive a planar layout, which should remain intuitive for medical professionals, by associating the actual body configuration with the graph; we fix the root and organ endpoints in a circular layout,

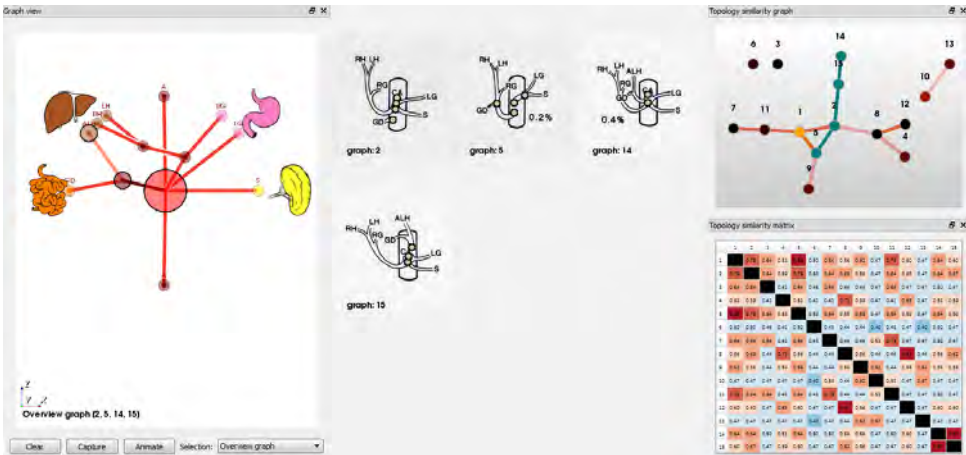


Figure 5.7: The VarVis interface. In the center view, individual and summary trees can be examined. Top left shows the similarity graph with three variations selected in green. The similarity matrix can be explored in the top right window. The bottom right window shows the illustrations of the selected trees. The bottom left shows which nodes of the individual trees are summarized in the summary tree.

which roughly preserves the naturally-occurring organ location, while keeping the graph compact. We discussed several potential layouts with our domain experts. We proposed a subway map metaphor similar to the Neurolines work [13], a rectangular layout with the organs arranged in two parallel lines and a circular layout, with the organ endpoints arranged on circular arcs. The domain experts preferred the circular layout and found the subway map organization too abstract to relate it easily to the anatomy, even though it minimizes intersections.

Projecting a 3D tree structure directly as a planar graph is unlikely to deliver a good illustration and avoid intersections or clutter. Hence, we position junction points by a force-directed mechanism to spread the graph, while the endpoints are fixed to keep them close to the organ arrangement (see Figures 5.2 and 5.3). A zero-crossing solution cannot always be guaranteed due to the inherent 3D nature of the problem. The user can arrange nodes manually, if needed.

An additional layout mechanism is used for the topology variation graph. In this case, we wish to preserve the distance between individual variations, whenever there are edges between them, i.e., if the similarity is above a certain threshold. We map the similarity between the trees to the distance between nodes and the edge length. We achieve this by adding the similarity metric as an edge weight and by using again a force-directed layout algorithm by Fruchterman and Reingold [20], which takes this measure into account (see Figure 5.7, top left).

SUMMARY TREE

In the graph view, individual trees or summary trees can be displayed and interacted with. The endpoints related to the organs are fixed in the circular layout as described in the previous section. To keep the main vessel recognizable in the graph layout, we can constrain it to a vertical line in the center of the circle, with its junctions equally spaced, leaving the other junction nodes to be placed by the force-directed layout as

before. When a node is selected in the tree or 3D representation, the same node is highlighted in the other views to provide the user with an intuitive way of visually linking the two representations together. We make use of halos around the tree edges to show which edge is in front in case of overlap [32]. The visual encoding of the summary tree (see Figure 5.7, center) uses node size and edge coloring to visualize the level of summarization. Bigger nodes indicate that more individual nodes from different trees had to be merged to create the summary tree, while the edge coloring further highlights the differences. When the user selects a node in the summary tree, detailed views of the corresponding variations can be explored. The variation view shows close-ups of the selected nodes in the individual trees. In case a selected node consists of many merged nodes, all of the original nodes are highlighted in the variation view. The variation view can also be used to capture (summary or individual) trees of interest and compare them side by side. Additionally, when illustrations are available, selected nodes are highlighted in green (see Figure 5.7, bottom right). The summary tree is also used to steer an animation between the summarized variation trees. By fixing matched nodes in the same positions and force-directing the merged nodes, initialized at the merged node location, the transitions between individual variations are highlighted.

TOPOLOGY SIMILARITY GRAPH

The topology similarity metric is visualized in two views. First, as a *node-link graph*, which reveals the connections between the variations that are most similar to each other. This graph visually encodes the probabilities of occurrence for each of the individual trees using a node color map. The node color indicates variation probability in a heated body color map. The similarity between the variations is encoded in both the edge color, as well as the position for extra emphasis. Distances between the nodes are preserved in the layout, so that pairs of nodes far away in the topology similarity graph are dissimilar. The graph can be used to discover groups of similar trees, which will be represented as closely positioned interconnected groups of variations in the topology similarity graph.

Second, we include a full *similarity metric matrix*. Since edges are only generated above a similarity threshold, the distances between all possible pairs are lost in the topology similarity graph. To answer questions related to the similarity of one tree to all other trees, we include the full similarity measures in a matrix representation. The similarity is quantified in this view, further emphasized by coloring the cells based on the similarity. In the matrix view, a blue to red color encoding is used for the cells, where blue represents dissimilar and red represents similar pairs. In this way, it can be used to see the similarity of a tree to all other trees at a glance. The matrix can also be re-ordered by different criteria, such as tree id, average similarity and probability. To select variations that are interesting to summarize or explore, the topology similarity graph can be used. For instance, the user can select nodes in the topology similarity graph that form a cluster. By interacting with the matrix one can either select individual variations or generate summary trees by selecting multiple variations. The summary tree representation is generated on the fly and consists of all the selected trees summarized using the algorithm described in Section 3.1.2.

ANATOMICAL VIEW

We also connect the variation trees to the 3D vascular anatomy, whenever medical imaging data is available (see Figure 5.3). The 3D view presents a silhouette rendering of the surface created from the segmentation of the vascular tree. It is rendered in a transparent way to reveal the centerline embedded within. The endpoints and junctions are accentuated with small spheres. The endpoints are colored according to their organ membership to create a visual link to the graph representation. The 3D view also features an interactive CT slice plane that can be adjusted freely in various slicing directions to provide anatomical context for the structures surrounding the arteries. When the variations are based on illustrations, the original illustration is shown in place of the 3D representation.

5.4. IMPLEMENTATION

To evaluate the applicability and potential of the method, we implemented a prototype application, VarVis. The core of VarVis is written in Python and uses PyQt for the user interface. PyQt and Python were chosen to support cross-platform development. We implemented the graph construction and matching algorithm in Python. We rely on the Visualization Toolkit (VTK) for the core visualization. We used the 2D Chart API of VTK for the graph rendering and extended it with a custom graph class in C++ to support our graph layout algorithm and custom representation. The force-directed part of our tree layout algorithm is based on the force layout algorithm in D3 [33] and the layout algorithm for the topology similarity graph is based on Fruchterman and Reingold's algorithm [20]. The segmentations, surfaces and centerlines of the CT data required as input for our system were preprocessed from clinical CT datasets using the Vascular Modeling Toolkit (VMTK).

5.5. RESULTS

In this section, we present results of our method for three cases using illustrations and medical imaging as the input. The prototype VarVis implementation can be seen in Figure 5.7.

The illustration-based variations in Figure 5.7 are constructed based on the variations of the gastroduodenal artery, as can be seen in Figure 5.1. In the original study, 500 cases resulted in 15 variations with various probabilities. To get this result, we first selected three trees of interest using the topology similarity graph (top left), causing them to be highlighted in green. In the center view, we see the summary tree that is generated from these three trees. By inspecting a node of the summary tree, e.g., a node that was merged, we can examine which nodes in the original tree are matched to that node and highlight the differences and similarities between the trees in the individual variations and illustration views, respectively located in the bottom left and right. For these three trees, we find that the junction nodes vary quite extensively among them, as demonstrated by the number of highlighted merged nodes in these views. Furthermore, we find that these trees have an accessory right hepatic artery in common, which causes them to cluster together in the topology similarity graph. In Figure 5.7, the topology similarity graph and matrix are shown in the top left and right window panes. The similarities

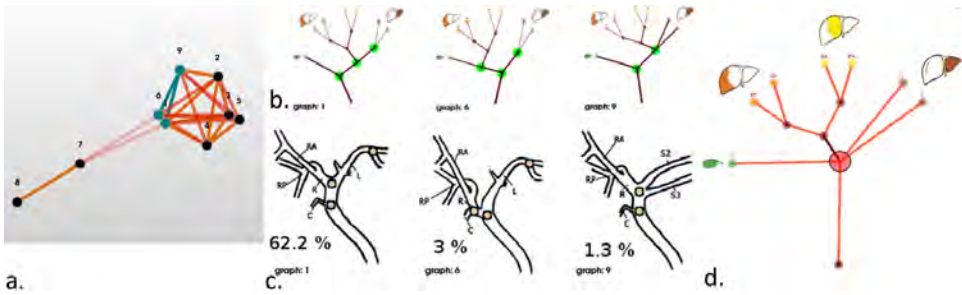


Figure 5.8: Interacting in the topology similarity graph (a) and summary tree (b) highlights differences in trees (c) and their linked illustrations (d).

between the fifteen variations of the gastroduodenal artery can be explored. The matrix view can be used to see which trees are most similar to all others, and which are most distant using the color encoding. To find the exact similarity between pairs of variations, the matrix displays the values inside the corresponding cells. Using the topology similarity graph to select subgroups of trees, we can discover clusters. Trees 4, 8 and 12 for instance, all have an accessory right hepatic artery, as we already noted in the previous summary example in Figure 5.7. Trees 14 and 15 have an accessory left hepatic artery.

When the variations have 3D volumetric data associated to them, we can explore the relation between the 2D graph representation and the 3D volumetric medical imaging data as well as the surface representation in linked views. The medical imaging datasets used in our the prototype implementation are three clinical CT-scans of the abdomen that feature contrast-enhanced arteries from the diaphragm level down to the aortic bifurcation. In order to only visualize anatomical variations and not pathological conditions, we selected scans that were made following a kidney transplant donor protocol. In Figure 5.3, the view on the left shows an individual variation tree, that is presented using our graph layout algorithm. The 3D view on the right displays the anatomical and radiological context of the same variation. The currently selected node is highlighted in both views.

An additional case is presented in Figure 5.8. Here we examine variations of the bile ducts [34]. First, we selected a cluster of three variations of interest in the topology similarity graph (a). By interacting with the summary tree (b), we highlight the differences in both the individual graph views as well as the illustrations.

5.6. EVALUATION

We performed an informal user study with three domain experts in order to evaluate the potential of VarVis. The participants had diverse medical backgrounds. The first expert is trained as an M.D. and has taught medical subjects. Currently, he is involved in developing computer aided learning techniques. The second expert is trained as an M.D. and currently involved in surgical anatomy research. The third expert is a medical biologist that currently works on e-learning programs within the department of Anatomy for medical students and supervises them. These experts have affinity with anatomical education, from the perspective of students as well as teachers. The evaluation consisted

Question	User 1	User 2	User 3
1. (3D) How many branching points does the aorta of tree 1 have?	1' 26"✗	0' 31"✓	0' 40"✗
2. (graph) How many branching points does the aorta of tree 2 have?	0' 21"✓	0' 07"✓	0' 10"✓
3. (all) How many branching points does the aorta of tree 3 have?	0' 05"✓"	0' 06"✓	0' 11"✓
4. (summary) Which branch is only available in one of the trees?	>5'00"✓	0' 15"✓	0' 08"✓
5. What are the differences in the right renal arteries?	0' 28"✓	0' 27"✓	0' 24"✓
6. (image) Which pair of trees are most similar?	4' 12"✗	0' 23"✗	0' 26"✗
7. (topology similarity graph) Which pair of trees are most similar?	0' 31"✓	0' 10.7"✓	0' 09"✓
8. (similarity matrix) Which pair of trees are most similar?	0' 31"✓	0' 10"✗	0' 14"✗
9. (image) Which 3 trees occur the most frequently?	0' 21"✓	0' 20"✓	0' 18"✗
10. (topology similarity graph) Which 3 trees occur most frequently?	0' 39"✓	0' 29"✓	0' 15"✓
11. (image/graph/matrix) Which tree is most similar to all others?	0' 13"✓	0' 20"✓	0' 26"✓
12. (animation) What is the difference between tree 14 and 15?	0' 59"✓	NA	0' 40"✓

Table 5.1: Evaluation questions asked with timing and correctness results. The time it took the three users to answer is shown, followed by a cross or checkmark, indicating the correctness of their answers.

of two sessions. In the first session, we examined the potential of the graph representation and summary tree as well as the link to medical imaging data. In the followup, we examined the utility of the topology similarity graph and matrix, as well as the relation to medical illustrations. In the evaluations, we used the same datasets that were presented in the results section. The evaluation sessions took on average one hour per user and session and consisted of four phases; 1) We provided an explanation of VarVis and its features in a live demo. 2) The experts used VarVis themselves to get used to the controls and took a look at the available visualization options. During this phase, the experts were asked a series of open questions to verify their understanding of the concepts and interaction possibilities. 3) The users were given concrete assignments to perform in one or more of the views. We measured the time it took them to perform these assignments and checked the correctness of their answers for the first session (questions 1 to 5) and the second session (questions 6 to 12) — see Table 5.1). 4) We asked a series of open questions designed to elicit feedback on the potential and future directions for VarVis.

5.6.1. INDIVIDUAL TREE VIEW

With respect to the 2D graph representation of the individual trees, User 1 remarked it might be too abstract for medical professionals, since they are accustomed to making their decisions based on a combination of experience and details. He raised a concern regarding the loss of these important details in the 2D tree. He did state that the 2D tree is well suited for checking structural connectivity of the vessels and as an overview, especially if the nodes were positioned with a closer relation to anatomy, based on their positions in 3D. User 2 felt that the relation between the 2D tree and the anatomy was clear, still she appreciated that the locations of the nodes were interactively adjustable and that she could change the positions into more intuitive locations for her, from an anatomical perspective. She preferred the individual trees over the 3D views for spotting variations between the trees, as the CT and 3D information would take her too much time to interpret properly. User 3 saw the relation of the 2D tree to the 3D anatomy clearly and appreciated the abstraction level that leaves out exact position details in favor of a clearer schematic layout. Questions 1, 2 and 3 in Table 5.1 revealed that the 2D views allowed all users to find the correct answer faster than in the 3D view. When the

source data is an illustration, however, all users pointed out that the individual illustrations are clearer and easier to interpret directly, than the individual graph representations. In this case, they commented that the summary tree could be used to highlight where the variations are in the individual illustrations.

5.6.2. SUMMARY TREE

The summary tree representation received mixed reactions. User 1 found it too abstract and was unable to relate it to actual anatomy. User 2 indicated she found the summary tree clear and understood what it was representing, but had some trouble understanding the exact meaning of the colors. User 3 fully understood the summary tree and its meaning. He saw it as a good way to quickly inspect a set of trees and find regions where anatomical variations occur. He used it to navigate to the detailed views of the variations themselves for closer inspection. Question 4 in Table 5.1 reflected that User 1 did not fully understand the overview representation. It took him over 5 minutes to reach enough understanding of the tree to spot the variation he was looking for. The other two users were able to use it to reach a correct conclusion in under 20 seconds. When asked to use whatever approach they preferred to spot a specific variation, the first two users relied on the individual tree to quickly reach a correct answer and the third user used the summary tree. The animation of the summary tree was found useful, though the level of abstraction prevented User 2 from finding the answer in Question 12. The linked highlighting in the illustrations feature was added after the evaluation phases. When shown to the users afterwards, they commented that it made the summary tree easier to use.

5.6.3. TOPOLOGY SIMILARITY GRAPH AND MATRIX

We asked several questions to examine the usefulness of the similarity graph and matrix compared to looking at the illustrations alone. From Question 6, 7 and 8, we find that none of the users can tell which pair of trees is most similar from only looking at the illustrations. User 2 even commented that it is impossible to do. It takes all users a long time (Question 6) and still could not give a correct answer. Using the topology similarity graph allows all users to reach a correct answer quickly (Question 7), while the matrix was similarly fast to use, it led to incorrect answers for two of the three users, due to just selecting similar pairs, but not the single most similar (Question 8). Questions 9 and 10 were related to detecting how often individual variations occur. The illustrations and graph both worked well for this purpose, although using the graph, all users answered correctly. In Question 11, two users used the similarity matrix to reach a correct answer, while one used the illustration. All users appreciated the topology similarity metric to identify groups. The matrix ordering was considered useful to quickly find variations with the highest probability or similarity to other trees.

5.6.4. APPLICABILITY

We asked the users to present their general impression on the tool using open questions. User 1 would like to make an adjusted version of the interactive graphs freely available via the Internet to provide medical students with an educational tool to be used in courses. Secondly, he sees potential in using the tool to teach surgeons common anatomical variations, but he would prefer to do this from a typical individual

anatomy representation, instead of having them use the summary tree. He commented that knowledge of anatomical variations can help surgeons and interventional radiologists to recognize variations more easily when they encounter them in clinical practice. User 2 would like to see the application used as an educational tool for teaching anatomy of vascular structures. In the future, she indicates it might be useful for surgeons in surgical planning if patient-specific graphs can be easily constructed by semi-automatically segmenting the vessels of interest in a pre-operative CT-scan. User 3 stated that he would like to use this application in education to teach systemic anatomy, a style of anatomy teaching that describes full systems such as the cardiovascular system, to students. He would first teach them 2D anatomy and then advance into the 3D part of the application once they grasp the 2D systemic information. As a teacher, he envisions using it as a quizzing system to assess the students. In the future, he would like to see more systems, such as the digestive and nervous system, represented in such an abstract 2D representation and to integrate or link these together if possible. Furthermore, he indicated there might be additional applications of this tool in research, when simulating vessel growth.

5.6.5. EVALUATION CONCLUSION

Overall, VarVis was found to be an interesting and promising application by the users. The 3D view and illustrations linked with the more abstract individual 2D tree representation provided the users with insight into the topological variations between vascular trees. Individual tree views were not considered added value compared to 2D illustrations, but they were appreciated to get an overview of the vascular connectivity when the initial data came from 3D medical imaging. The abstraction level of the summary tree representation was received with mixed reviews, but all users saw the utility of it in identifying regions that are similar or regions that are different. The timed questions revealed that the users were quicker and more correct in answering a systemic anatomy question in the 2D graph representation than in the 3D view. They did prefer to use the 2D illustrations over the 2D graph representations, when dealing with questions regarding individual graphs. Furthermore, they were able to quickly spot specific variations using either the individual 2D graphs or the summary tree. The topology similarity graph and matrix were found useful to identify similarities and groups of trees. This task was found especially difficult to achieve with the 2D illustration or using the 3D vascular trees. All users saw the potential of using VarVis as an educational tool in the future.

5.7. CONCLUSION AND FUTURE WORK

We presented VarVis, a way of visualizing anatomical variations in branching structures that allows the user to compare and explore these variations in an interactive environment. VarVis features a topological graph-based visualization of the vessel tree combined with illustrations and 3D anatomical surface representations in linked views. The topology similarity graph provides the user with an overview of the existing variations and their similarities at a glance. Individual variations can be summarized and animated to provide insight into the local differences and similarities. We visualize the graph representations using a semi-automatic force-directed layout algorithm specifically designed to present the graph in a clear and compact overview, while maintaining anatomical

structure. We performed an informal evaluation with three domain experts that revealed that VarVis has potential as an educational tool for teaching systemic anatomy. For more conclusive results, a larger follow-up study is needed with medical students. In the future, we would like to extend this work to build patient-specific graph representations for anatomical branching structures that have robust automatic segmentation methods available. In this way, the specific variation the patient has and how it differs from the most common configuration could be emphasized.

REFERENCES

- [1] N. N. Smit, A. C. Kraima, D. Jansma, M. C. Deruiter, E. Eisemann, and A. Vilanova, *Varvis: Visualizing anatomical variation in branching structures*, in *Proceedings of EuroVis - Short Papers* (2016) accepted, to appear.
- [2] K. L. Moore, A. F. Dalley, A. M. Agur, *et al.*, *Clinically oriented anatomy* (Wolters Kluwer Health, 2013).
- [3] J. R. Sañudo, R. Vázquez, and J. Puerta, *Meaning and clinical interest of the anatomical variations in the 21st century*, *European Journal of Anatomy* **7**, 1–4 (2003).
- [4] P. L. Willan and J. R. Humpherson, *Concepts of variation and normality in morphology: Important issues at risk of neglect in modern undergraduate medical courses*, *Clinical Anatomy* **12**, 186–190 (1999).
- [5] O. A. Catalano, A. H. Singh, R. N. Uppot, P. F. Hahn, C. R. Ferrone, and D. V. Sahani, *Vascular and biliary variants in the liver: Implications for liver surgery*, *RadioGraphics* **28**, 359 (2008).
- [6] E. H. Daseler and B. J. Anson, *The cystic artery and constituents of the hepatic pedicle; a study of 500 specimens*. *Surgery, gynecology & obstetrics* **85**, 47 (1947).
- [7] S. Hacker and H. Handels, *Representation and visualization of variability in a 3d anatomical atlas using the kidney as an example*, in *Medical Imaging* (International Society for Optics and Photonics, 2006) pp. 61410B–61410B.
- [8] H. Handels and S. Hacker, *A framework for representation and visualization of 3D shape variability of organs in an interactive anatomical atlas*, *Methods of information in medicine* **48**, 272 (2009).
- [9] C. J. Stam and J. C. Reijneveld, *Graph theoretical analysis of complex networks in the brain*, *Nonlinear biomedical physics* **1**, 3 (2007).
- [10] M. Rubinov and O. Sporns, *Complex network measures of brain connectivity: Uses and interpretations*, *NeuroImage* **52**, 1059 (2010).
- [11] A. Crippa, N. M. Maurits, M. M. Lorist, and J. B. Roerdink, *Graph averaging as a means to compare multichannel EEG coherence networks and its application to the study of mental fatigue and neurodegenerative disease*, *Computers & Graphics* **35**, 265 (2011).

- [12] M. Perez, A. Highes, A. Stanton, S. Thorn, N. Chapman, A. Bharath, and K. Parker, *Retinal vascular tree morphology: a semi-automatic quantification*, IEEE Transactions on Biomedical Engineering **49**, 912 (2002).
- [13] A. Al-Awami, J. Beyer, H. Strobel, N. Kasthuri, J. Lichtman, H. Pfister, and M. Hadwiger, *NeuroLines: A subway map metaphor for visualizing nanoscale neuronal connectivity*, IEEE Transactions on Visualization and Computer Graphics **20**, 2369 (2014).
- [14] X. Zhang, J. Tian, K. Deng, X. Li, and F. Yang, *Graph representation of hepatic vessel based on centerline extraction and junction detection*, (2012) pp. 83143G–83143G–6.
- [15] M. V. Filatov, J. A. Kaandorp, M. Postma, R. v. Liere, K. J. Kruszynski, M. J. A. Vermeij, G. J. Streekstra, and R. P. M. Bak, *A comparison between coral colonies of the genus madracis and simulated forms*, Proceedings of the Royal Society of London B: Biological Sciences **277**, 3555 (2010).
- [16] K. J. Kruszynski, J. A. Kaandorp, and R. Liere, *A computational method for quantifying morphological variation in scleractinian corals*, Coral Reefs **26**, 831 (2007).
- [17] K. J. Kruszynski, R. van Liere, and J. A. Kaandorp, *An interactive visualization system for quantifying coral structures*. in *Proceedings of the Eighth Joint Eurographics/IEEE VGTC conference on Visualization* (Eurographics Association, 2006) pp. 283–290.
- [18] P. Eades, *A heuristics for graph drawing*, Congressus numerantium **42**, 146 (1984).
- [19] T. Kamada and S. Kawai, *An algorithm for drawing general undirected graphs*, Information Processing Letters **31**, 7 (1989).
- [20] T. M. J. Fruchterman and E. M. Reingold, *Graph drawing by force-directed placement*, Softw Pract Exp **21**, 1129 (1991).
- [21] R. Davidson and D. Harel, *Drawing graphs nicely using simulated annealing*, ACM Trans. Graph. **15**, 301 (1996).
- [22] T. Dwyer, Y. Koren, and K. Marriott, *IPSep-CoLa: An incremental procedure for separation constraint layout of graphs*, IEEE TVCG **12**, 821 (2006).
- [23] J.-H. Won, J. Rosenberg, G. D. Rubin, and S. Napel, *Uncluttered single-image visualization of the abdominal aortic vessel tree: Method and evaluation*, Medical Physics **36**, 5245 (2009).
- [24] A. Barsky, T. Munzner, J. Gardy, and R. Kincaid, *Cerebral: Visualizing multiple experimental conditions on a graph with biological context*, IEEE TVCG **14**, 1253 (2008).
- [25] F. Zhou and F. De la Torre, *Deformable graph matching*, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2013) pp. 2922–2929.
- [26] H. Bunke and G. Allermann, *Inexact graph matching for structural pattern recognition*, Pattern Recognition Letters **1**, 245 (1983).

- [27] P. Bille, *A survey on tree edit distance and related problems*, Theoretical computer science **337**, 217 (2005).
- [28] H. Bunke and S. Günter, *Weighted mean of a pair of graphs*, Computing **67**, 209 (2001).
- [29] H. Bunke and A. Kandel, *Mean and maximum common subgraph of two graphs*, Pattern Recognition Letters **21**, 163 (2000).
- [30] X. Jiang, A. Munger, and H. Bunke, *On median graphs: properties, algorithms, and applications*, IEEE Transactions on Pattern Analysis and Machine Intelligence **23**, 1144 (2001).
- [31] L. Antiga, B. Ene-Iordache, and A. Remuzzi, *Centerline computation and geometric analysis of branching tubular surfaces with application to blood vessel modeling*. in WSCG (Citeseer, 2003).
- [32] A. Appel, F. J. Rohlf, and A. J. Stein, *The haloed line effect for hidden line elimination.*, Vol. 13 (ACM, 1979).
- [33] M. Bostock, V. Ogievetsky, and J. Heer, *D³ data-driven documents*, IEEE TVCG **17**, 2301 (2011).
- [34] A. Ragab, R. I. Lopez-Soler, A. Oto, and G. Testa, *Correlation between 3D-MRCP and intra-operative findings in right liver donors*, Hepatobiliary surgery and nutrition **2**, 7 (2012).

III

SURGICAL PLANNING

THE last part of this thesis deals with aspects of the work related to the goal of improving pre-operative planning of the Total Mesorectal Excision (TME) surgical procedure.

Combining atlas and medical imaging data is an example of multimodal visualization. To examine the related prior art in this field, we start this part of the thesis with a survey on multimodal medical data visualization in Chapter 6.

To achieve the goal of improved pre-operative planning for the TME procedure, two components are needed. First, the VSP needs to be registered to patient-specific pre-operative MRI scans, as was already described in Chapter 3. To simplify this registration process, we have developed an application for interactive 3D volume registration, as described in Chapter 7.

Furthermore, the patient-specific anatomy and pathology needs to be visualized to prepare for the surgery. An application aimed at such atlas-based surgical planning for oncological pelvic surgery is described in Chapter 8.

6

A SURVEY ON MULTIMODAL MEDICAL DATA VISUALIZATION

Visualizing atlas and patient-specific MRI data simultaneously, presents similar challenges to those that arise in general in multimodal medical visualization. Therefore, we performed a survey on techniques and applications visualizing multimodal medical data.

Multimodal data of complex human anatomy contains a wealth of information. To visualize and explore such data, techniques for emphasizing important structures and controlling visibility are essential. Such fused overview visualizations guide physicians to suspicious regions to be analyzed in detail, e.g., with slice-based viewing. We give an overview of state of the art in multimodal medical data visualization techniques. Multimodal medical data consists of multiple scans of the same subject using various acquisition methods, often combining multiple complimentary types of information. 3D visualization techniques for multimodal medical data can be used in diagnosis, treatment planning, doctor-patient communication, as well as interdisciplinary communication. Over the years, multiple techniques were developed in order to cope with the various associated challenges and presenting the relevant information from multiple sources in an insightful way. In this chapter, we present an overview of these techniques and analyze the specific challenges that arise in multimodal data visualization and how recent works aimed to solve these, often using smart visibility techniques. We provide a taxonomy of these multimodal visualization applications, that highlights which modalities were used and which visualization techniques were employed. Additionally, we identify unsolved problems as potential future research directions.

This chapter is based on a publication currently in submission to Computer Graphics Forum: Kai Lawonn, **Noeska Smit**, Katja Bühler, and Bernhard Preim, "A Survey on Multimodal Medical Data Visualization", In Submission, *Computer Graphics Forum*, 2016 [1]




Figure 6.1: 2D images with superimposed CT and PET data, represented in gray scale and color respectively.

6.1. INTRODUCTION

Multimodal medical datasets consist of multiple scans of the same subject using various acquisition methods. In this way, it is possible to image different tissue characteristics. Several modalities can also be integrated directly into the hardware of a single scanner. Of these hybrid medical imaging techniques, PET/CT is currently the most widespread. More recently, PET/MR became clinically feasible and additionally SPECT/CT is also used in clinical practice. The common denominator in the aforementioned hybrid imaging combinations is that often the radiologic modality (CT or MRI) depicts the anatomy of the patient in high spatial resolution, whereas the nuclear medicine modality (PET or SPECT) depicts functional processes, such as metabolism, in a lower resolution. Additionally, it is possible to combine structural information from a scan with functional information from the same scanner, such as combining an anatomical MRI acquisition with functional MRI information. Different tissue types cannot easily be distinguished based on metabolism alone. CT and MRI, however, are able to visualize anatomical information at a much higher resolution, but lack metabolic information. Thus, a combination of both imaging modalities allows visualization and localization of metabolic information in anatomical context. The overall visualization goal in this case is to accurately localize regions featuring abnormal functional values based on their relation to structural anatomical information.

Typically, radiologists or specialists in nuclear medicine examine data in a slice-based fashion consisting of superimposed 2D images with a combined visualization of CT and PET image data, as is shown in Figure 6.1. Here, physicians browse through the slices and set window and level parameters to adjust the brightness and contrast, and examine structures of interest. Unfortunately, presenting these 2D images in such a way can hinder full and quick analysis of the data. First, the physician needs to check every slice, which may be time-consuming as the number of slices increases. Furthermore, physicians need to mentally fuse all information from these slices to form a correct diagnosis. Therefore, visualization techniques that provide a 3D overview are helpful to see potential abnormalities at a glance. This enables the experts to navigate to these suspicious regions for a more detailed exploration in the slice views.

In recent years, the amount of work focusing on multimodal medical data visualization increased. In 2010, there was an IEEE VIS contest on multimodal visualization for

neurosurgical planning [2], further highlighting the interest in and importance of multimodal visualization. With this chapter of the thesis, we aim to provide a survey of this literature. A number of medical visualization techniques are described in a survey style in the book by Preim and Botha [3]. Multimodal visualization, however, is only slightly touched upon as an add-on to registration and is also not considered in other survey articles up until now. In general, concepts used for visualizing medical data can be extended to other multimodal visualization application domains, but there are some specific challenges in dealing with medical multimodal data. Thus, the presented works are of interest to any visualization researcher dealing with multimodal data.

The contributions of this survey are the following:

- We provide an overview of multimodal acquisition techniques and relate this to requirements and challenges.
- We propose a taxonomy of multimodal medical data visualization applications.
- We provide an outlook on open problems in multimodal medical visualization and a perspective on future research directions.

Scope. There have been several surveys on multi-field data visualizations [4, 5]. In contrast to general multi-field data, measured medical image data is not as 'clean' as simulation data, due to the acquisition process. Living tissue is imaged with scanning parameters that favor the patients' safety over image quality, unlike for instance applications in material testing. Furthermore, several tasks such as searching for metastases, or assessment of infiltration, are unique and therefore also evaluation of multimodal medical visualization that needs to consider such or similar relevant tasks. While multimodal medical image data are available at both microscopic and macroscopic level, we focus on the macroscopic level, e.g. radiological image data. We consider also interaction with multimodal data. Due to the complexity of these data, interaction is even more important than in the exploration of single data.

Please note that in this survey we mainly focus on combining imaging modes acquired by different scanning techniques, and less on combining imaging modes of a single scanner, e.g. MRI T1 and T2. We consider multiple modes from the same scanner such as these, examples of multi-field data, similar to the definitions stated by Blaas et al. [6]. For the same reason, though it is multi-field data, we also do not consider Diffusion Tensor Imaging (DTI) data to be inherently multimodal when used with normal MRI. DTI can additionally be used in multimodal applications when it is combined with fMRI or additional imaging modalities. Multimodal data is often acquired from separate scanners, which creates a need for software-based registration to align the volumes. Due to the time between the scans and patient pose differences it is not straightforward to register multiple volumes accurately. This in turn generates additional challenges in dealing with uncertainty in the form of processing errors introduced by the registration process. For an overview of medical image registration we refer to the survey by Maintz and Viergever [7] and work on mutual-information-based registration methods by Pluim et al. [8]. In essence, there are many registration techniques that can be employed to align multimodal datasets, such as maximizing an information-theoretic measure, e.g. normalized cross-correlation and normalized mutual information. Furthermore, landmark-based registration can be employed, which requires appropriate crisp

anatomical landmarks, e.g., vessel bifurcation points or brain folds. These optimization strategies can be costly, in terms of processing time required, in which case GPU support is necessary [9].

While blood flow measurements can be considered functional information, a review of these techniques is out of the scope of this survey (see the recent survey on Cardiac 4D PC-MRI [10]).

In our survey we focus on three core applications areas, namely visualizations aimed at research, diagnosis and treatment planning. Within these areas, we distinguish between different application subdomains. We divide the applications aimed at diagnosis in oncologic and cardiologic visualizations, and the applications aimed at treatment planning in neurosurgery and radiotherapy planning.

Paper selection criteria. We searched for papers that were related to multimodal medical visualization using the EG digital library, where we looked for the following conferences and workshops: EuroVIS, VCBM, and VMV in the last ten years. Furthermore, we used the IEEE digital library and the proceedings of the IEEE VIS conference for the last ten years. Additionally, we used Google Scholar for finding papers related to our survey, in order to integrate older papers. Specifically, we looked for the following keywords and combinations of these: combined, concurrent, CT, DTI, dual, fMRI, fused, fusion, hybrid, integrated, intermixing, neurosurgery, MRI, multi-field, multi-variate, multi-volume, multimodal, multiple, PET, planning, simultaneous, SPECT, and visualization. We employed "neurosurgery" as the only application-specific search term, since this area is indeed the key application for many multimodal visualization applications so far.

Organization. Our survey is structured as follows. In SECTION 6.2 we provide an overview of hybrid imaging acquisition and what the characteristics, advantages and disadvantages of each of these modalities are. In SECTION 6.3, we describe the typical workflow of a specialist in radiology and nuclear medicine for exploration and analysis of the data. Based on this workflow, we derive requirements for 3D visualizations. In SECTION 6.4 we provide a brief overview of relevant visualization techniques, such as smart visibility techniques and focus-and-context techniques, that can be employed for multimodal visualization. SECTION 6.5 continues with multimodal rendering (Section 6.5.1) and interaction (Section 6.5.2) techniques. In SECTION 6.6, we focus on applications of multimodal visualization techniques to real-world clinical data. We subdivide this section into the application areas these visualizations are developed for:

- Research (Section 6.6.1),
- Diagnosis (Section 6.6.2),
- Treatment planning and guidance (Section 6.6.3).

SECTION 6.7 concludes the survey and outlines unsolved problems and challenges for future research. Finally, in SECTION 6.8 we conclude with a brief summary.

6.2. MEDICAL BACKGROUND

In this section, we provide an overview of multimodal and hybrid imaging, as well as their applications in clinical practice.



Figure 6.2: A PET/CT scanner developed by Siemens Healthcare. The frontal part of the gantry features a CT scanner, while the rear part contains a PET scanner.

When combining acquisitions from multiple modalities from different scanners, registration problems arise. To simplify the registration process, efforts have been made to integrate multiple modalities into hybrid scanners that combine the best of both worlds in structural and functional imaging. Hybrid imaging scanners combining PET or SPECT with CT are already commonplace in clinical practice, while combining PET, SPECT with MRI is more recent development [11] for which the first prototypes have been made [12]. In a preliminary study comparing clinical impact of PET/CT and PET/MRI, PET/MRI imaging outperformed PET/CT and more frequently affected patient management [13]. Furthermore, integrated whole-body PET/MRI was found feasible in a clinical setting with comparable reliability to PET/CT for this purpose [14]. Since 2014 the very first devices are legally allowed and used in clinical practice, focused on head and neck imaging initially.

PET/CT: While both PET, CT, and MRI provide valuable information by themselves, they can be combined to provide more insight into the exact localization of suspicious metabolic activity. By combining PET and CT/MRI scans, physicians are able to detect (abnormal) metabolic activity using the PET scan and to localize this activity using the CT or MRI scan. This can aid the user in distinguishing which activity is physiologically normal and which is pathological. Since the adaptation of PET/CT is more widespread clinically, we focus on this hybrid imaging techniques in the rest of this section (see Figure 6.2).

The CT data is also used to perform noiseless attenuation correction on the PET data,

thus eliminating the need for an additional PET transmission scan for this purpose. This can reduce the total scanning time by up to 40 percent [15]. A drawback of PET/CT is that the imaging is performed sequentially, which takes more time and eliminates temporal correlation between the modalities [16].

Clinical oncological applications of the combined PET/CT scanner include diagnosing and staging primary tumors as well as localization of metastatic disease in almost any region of the body [17]. Further applications include decision making with respect to surgical operability of the tumor or treatment selection. PET/CT is additionally used to determine if cancer has recurred or to determine the difference between scar tissue and active cancer tissue. Especially when monitoring therapy, PET can detect changes in tumors earlier than CT, because metabolic changes occur sooner than anatomical size changes.

For a more elaborate description of PET/CT acquisition, we refer to Townsend et al. [15].

SPECT/CT: The benefits of combining PET and CT also extend to combining SPECT and CT. As with PET/CT, the CT can provide structural anatomical information for localization, while the SPECT can provide metabolic information. Similarly, the CT can be used to perform attenuation correction for the SPECT. SPECT is less expensive than PET, but suffers from lower contrast and spatial resolution [18]. A general description of the technical developments and future directions of SPECT/CT can be found in the review by Buck et al. [19]. A more elaborate discussion on the clinical uses of SPECT/CT can be found in Mariani et al. [20].

In essence, CT and PET or SPECT provide complimentary information that needs to be integrated and thus fused visualizations are desirable.

6.3. CLINICAL WORKFLOW AND REQUIREMENTS

In this section we examine the clinical workflow for the analysis of multimodal medical imaging data. Furthermore, we provide an overview of visualization requirements that should be considered when developing a novel visualization technique for such data.

6.3.1. CLINICAL WORKFLOW

Multimodal medical imaging acquisitions are interpreted by radiologists and/or nuclear medicine physicians. Since each of the two disciplines have their own specific skill-set, ideally, two imaging specialists (one from each discipline) would interpret multimodal images together. However, due to availability and workload this is often not possible [21]. Additionally, there are now more and more physicians specialized in both radiology and nuclear medicine in joint departments for radiology and nuclear medicine. While specific protocols may vary per hospital, in general the images are analyzed similarly in radiology workstations of major vendors. The modalities can be separately analyzed in the views that are common for the individual modalities, such as slice-based viewing in the axial, coronal and sagittal direction. For nuclear imaging modalities, Maximum Intensity Projections (MIP) are commonplace in clinical practice. Often these MIP renderings are viewed from the front or the side of the patient. Depending on the clinical indication, more elaborate visualizations can be used, such as polar maps for cardiac SPECT

data [22].

To visualize both modalities simultaneously, the typical approach in for instance PET/CT data, is to examine the CT images in grayscale with the PET images superimposed using a colormap in a 2D fused view (Figure 6.1). The combination of a colored overlay of functional information with gray-valued anatomical information is quite effective since brightness and color are different perceptual channels and thus can be processed simultaneously [23]. Interestingly, the choice of colors in the functional color map is not standardized and varies per vendor. While most systems map the highest activity region to red, for instance General Electric (GE) workstations map it to blue. Common clinically used visualization combinations of PET/CT data can be seen in the work by Griffith [24]. The cases presented here are often shown with a frontal MIP of the PET, axial slices of both the PET and CT separately, and finally a fused superimposed slice of the PET and CT combined. The exact configuration of the available views may vary, but typically there will be a central view with supporting linked views or a display featuring linked equally-sized views of the different modalities and slice directions. For oncological diagnosis and treatment planning, also the combination of CT and MRI data may be important [25]. Since these data are both high-resolution data, their combination is more challenging. The simultaneous exploration of CT and MRI, with either PET or SPECT, is currently not feasible with clinically available software.

While more advanced visualization techniques could be employed to visualize multimodal images, these are not yet broadly available in clinical practice. Examples of these techniques can be found in the OsiriX software, which was specifically designed for navigation and visualization of multimodal and multidimensional data [26]. Visualization options include Multiplanar Reformation (MPR), surface rendering and direct volume rendering for fused multimodal datasets.

6.3.2. REQUIREMENT ANALYSIS

From a medical visualization perspective, the goals of multimodal medical data visualization include:

- reducing complexity and thus cognitive load,
- enabling, improving, or accelerating decision making, and
- providing tailored visualizations for specific applications.

In diagnosis for instance, examining all 2D slices individually to examine anatomy, pathology and metabolic uptake can become time-consuming and cumbersome. Therefore, besides the traditional 2D approaches, 3D techniques can additionally be used to give an overview of the full datasets at a glance. Existing methods such as maximum intensity projection (MIP), which was first developed for use in Nuclear Medicine [27], can provide such an overview, but suffer from depth perception issues. In most cases, since the intensities are not comparable, a combined volume MIP can suffer from occlusion by the highest intensity regions.

3D visualization techniques that are suitable for displaying hybrid multimodal data would allow the users to get a quick overview of areas of interest and localize for instance foci of elevated metabolic activity in an anatomical context. When developing a novel

visualization technique or application designed for use with multimodal medical data, there are some general requirements that should be kept in mind. The requirements for designing a 3D multimodal medical visualization technique are the following:

1. Visualization parameters should be easily adjustable to fit the needs of the user.
2. In order to be suitable for clinical use, the technique should be fast and interactive, with minimal or no pre-processing required.

When combining functional with structural information for diagnostic purposes, the following additional requirements should be fulfilled [28]:

1. The technique should show the combination of the two or more modalities in a fused view in which the functional activity of interest is always visible.
2. The technique should relate metabolic activity to nearby anatomical structures for accurate localization.

Such a technique can be used to guide the exploration of the datasets by bringing the attention to regions of interest, after which detailed inspection of these regions can be performed in 2D images. Furthermore, besides aiding diagnosis, 3D visualization techniques could be beneficial for research and treatment planning.

For research purposes, the visualization requirements are similar to the requirements for a diagnostic application, and further specific requirements are application-dependent and should be formulated based on the needs of the domain experts. Surgical treatment planning could benefit from multi-modal 3D patient-specific visualization to support access planning. In oncologic neurosurgery for instance, an access path to the tumor needs to be planned, taking the location of risk structures, such as arteries, into account. In this neurosurgical context, a visualization application should [25]:

- Provide high-quality, interactive and flexible 3D visualization
- Offer multimodal visualization for modalities such as CT, MRI, fMRI, PET or DSA.
- Provide interactive manipulation such as simulated surgical procedures, endoscopic views or virtual cutting planes.

In radiotherapy treatment planning, the pathology, i.e., radiation target, should be visualized in the context of the healthy structures at risk to receive undesired radiation damage. The requirements for such an application include [29]:

- Support for 4D PET and CT data and fusion of these modalities in a 3D view
- Visualization of segmentation data
- Visualization of dose information
- Clipping and/or masking parts of the volume
- No pre-processing required

These requirements demand advanced visualization techniques to fulfill the presented requirements. In the next section, we present several visualization concepts that have the potential to be applied in multimodal medical applications.

6.4. VISUALIZATION TECHNIQUES

Due to the nature of multimodal medical data visualization, there are always two or more volumes that need to be visualized. With overlapping extents, each of these volumes may be occluded by the others, and thus no clear insight can be given into the features of each of the modalities simultaneously. Visualization techniques supporting this visualization problem need to incorporate heuristics for assessing the importance of information as well as emphasis techniques to adjust the visualization to the derived importance. Therefore, smart visibility, focus-and-context techniques and other emphasis techniques are crucial when dealing with multimodal medical data. In this section, we describe basic visualization techniques to cope with challenges such as dealing with occlusion, improving depth perception and presenting relevant information from multiple sources in an insightful way.

6.4.1. BASIC REPRESENTATION

For multi-modal visualization, often pre-processing is required in terms of registration or segmentation. When combining datasets from multiple modalities using separate scanners, registration needs to be performed to align the volumes. An example of this can be found in the work by Wein et al., who proposed a fully automatic registration method to align 3D ultrasound with CT scans [30, 31]. Segmentation is not performed routinely in all clinical applications, but for instance in radiation treatment planning, the target structure (tumor) and all organs at risk are segmented manually. Here, segmentation may be performed as a basis quantification for further analysis, e.g., volume measurements.

Both direct volume rendering (DVR) and indirect volume rendering (IVR) techniques can be applied to visualize multimodal medical imaging data. DVR visualizes the data directly, whereas IVR attempts to reconstruct a triangulated surface mesh representative of the original image data and as such is suitable to define crisp boundaries.

In IVR, to determine the triangulated surface mesh, usually the Marching cubes algorithm is applied to a segmented subset of the image data [32]. The method may lead to ambiguities such that correct results cannot be guaranteed. Using the resulting mesh, various shading techniques as well as different visualization methods can be applied. On such surface meshes, it is easier to perform illustrative visualization techniques such as stippling and hatching, than it would be in DVR, and these techniques have potential to be applied in multimodal medical visualization. To determine the surface, a segmentation of the medical image data needs to be performed to select and delineate structures of interest. Segmentation can be performed by point-based, edge-based, region-based, or model-based approaches. Point-based approaches includes intensity-based methods, e.g., setting a threshold. Thus, there are various algorithms that determine an optimal threshold [33–35]. For edge-based approaches the LiveWire [36, 37] technique can be applied. Region-based approaches include region-growing methods [38, 39] and watershed segmentation [40, 41]. Model-based methods attempt to fit a predefined model to the intensities, gradients, curvature and/or texture of an image. Active contours or snakes are mostly used fit in a closed curve to a certain structure [42]. Since discussing alternative methods for image segmentation is out of the scope of this paper, we refer to the books by Bankman et al. [43] and Setarehdan et al. [44] for additional information on

this topic.

In contrast to IVR, in DVR, we are not susceptible to segmentation pre-processing. In this way, we do not need to pre-determine suitable thresholds and do not risk losing critical aspects of the original data, based on these choices. However, suitable transfer functions still need to be defined.

6.4.2. SMART VISIBILITY

To display multiple structures from different modalities leads to challenging problems. Mostly, the user is interested in one specific structure, e.g., a tumor, or a class consisting of several important structures. These important structures may be surrounded by other objects, resulting in structures that may occlude or overlap each other. In the following, we describe advanced visualization techniques that require the important structures to be segmented. This prerequisite is fulfilled in some clinical applications, e.g., in radiation treatment planning and complex neurosurgery planning. A primary concern is how to resolve the occlusion problem such that the most important structures are clearly visible, while still indicating surrounding structures. Hiding the surrounding structures is not appropriate, since they provide spatial context. Whenever a class of interesting objects are occluded, smart visibility techniques can be applied. For the smart visibility term, we refer to Viola and Gröller [45] who gave the following definition “expressive visualization techniques that smartly uncover the most important information in order to maximize the visual information in the resulting images”. This visualization technique category comprises several sub-categories and different ways to remedy the occlusion problem. Here, we introduce the most prominent ways to resolve visibility problems comprising of cutaways, ghosted views and exploded views and present preliminary guidelines on how they could be applied to multimodal medical data.

CUTAWAY VIEWS

A common way to depict the focus object and to separate it from the surrounding context structures is to use cutaway views. Cutaway views employ, for instance, a conic object placed around the focus object. The conic object is aligned towards the camera in such a way that the camera looks inside the object from the base to the tip. Every non-focus structure that is inside the conic object is cut such that only the focus object inside the cone is visualized. The main advantage is the clear illustration of the main object without occlusion by the surrounding objects. Thus, the user gets a direct view on the focus object. Unfortunately, structures such as vessels around the object, as well as important regions between the focus object and the view point of the camera are lost. In general, cutaways vary from *predefined models* [46] to *focus-oriented objects* [47, 48]. The algorithmic complexity differs strongly based on the object used.

Predefined models: With predefined models, view-dependent and view-independent approaches can be distinguished. Independently, the predefined model is loaded into the framework and subsequently all structures are tested to see if they are inside or outside the focus object. For instance, a lymph node can be approximated by a cylinder with a suitable size. Here, during the rendering an inside/outside test can be applied to the surrounding (context) objects with the predefined model. In case the context objects overlap with the predefined model, the fragments are discarded; otherwise the frag-

ments are drawn. Based on the view-dependence, the predefined model is either aligned along the view direction of the camera or not. Mostly, view-dependent cutaway techniques are used, where the conic objects are oriented along the view-vector. Predefined models can be used if the the focus object has a simple shape. Moreover, it is also only applicable if the model does not vary in time. In case of animated blood flow, predefined models cannot be employed, due to the change of the form and position of the blood flow [49].

Focus-oriented objects: The second class of cutaway techniques uses focus-oriented objects. Here, the conic object is constructed based on the focus object and can be arbitrarily shaped. To achieve real time performance, the determination of the conic object and the associated cutaway needs to be efficient. Previous work, such as the work by Viola et al. [50], uses a Chamfer distance transform approximation [51] to the conic object in order to calculate the cutaway. Thus, a depth image of the conic surface is calculated, which is then used to cut away regions that are closer to the viewer than the conic object itself. If the depth of the surrounded objects is smaller than the depth of the conic object, the corresponding fragment is discarded, otherwise it is shown. Other approaches were developed to speed up this calculation [52]. The key idea is to determine the conic object in the view plane with a jump flooding algorithm. Usually, the conic object is generated by applying the standard flood fill, an iterative process in which in every iteration a pixel passes its value to its direct neighbors. Contrarily, the jump flood algorithm uses the same approach, but the neighbors vary for each iteration. Thus, a pixel passes its value to its neighbors, which have a certain distance. For each iteration the step size is halved. Cutaway generation was used to visualize animated blood flow in such a way that a clear view on the blood flow within a vessel is always guaranteed [53].

Application to multimodal data: Cutaway views can be directly employed in multimodal medical data visualization applications without much adaptation. For instance, one modality could be displayed outside the cutaway region while another is displayed within the cutaway. Furthermore, cutaways can be used on both modalities to reveal structures of interest hidden within a larger anatomical structure.

For instance, a cutaway can be used to display hybrid imaging PET/CT data, where high PET activity needs to be depicted. By using a cutaway view in these regions, the high PET activity is always visible and it provides a clear view inside the anatomical CT data and can help guarantee a correct localization of the occurrence, whereas a simple maximum intensity projection cannot give reasonable depth cues. Such a view might be useful not only for diagnosis, but also for patient-doctor and interdisciplinary communication, as well as for treatment planning.

GHOSTED VIEWS

An easy way to illustrate the focus object and the surrounding objects is to use transparency. The main drawback of employing multiple transparent objects is that they may lead to visual clutter. Furthermore, it is hard to interpret the shape and ordering of overlapping regions. To remedy this, ghosted view techniques attempt to improve the transparency depending on the underlying object. Here, a distinction can be made between smart transparency techniques and interactive approaches. In smart transparency, the strength of transparency depends mostly on curvature or normal vector information.

One approach is to set the transparency very high if the surface is nearly planar and to increase the opacity if the curvature varies. That helps to identify regions where the morphology changes strongly. This was successfully applied for instance to vessel visualization with embedded flow [54, 55]. The transparency is set by $F = 1 - \langle V, N \rangle^r$, where V represents the view vector, N the surface normal, and r represents the shininess constant. The superiority of this approach over standard semi-transparent visualization related to cerebral blood flow was confirmed in the perceptual study by Baer et al. [56].

An additional transparency technique was inspired by suggestive contour lines, a line drawing technique that emphasizes characteristics of the surface. This approach was used to develop a shading approach, based on a suggestive contour scalar field [53]. Here, the scalar field is used to determine zero-crossings after which positive and negative values are used to shade the object in orange or cyan, respectively.

Application to multimodal data: Ghosted views can be applied to multimodal data visualization in order to remove emphasis of structures that are less important from one or more modalities. A ghosted view can be employed to show an anatomical context in order to localize functional data without occluding the view. For instance, a brain surface can be rendered as a ghosted view, while functional data is rendered in an opaque style. This yields a visualization where the brain surface can be recognized, without being occluding, whereas the functional brain data is directly depicted. The main challenge is to provide a transparency technique that is not disturbing, but provides as much morphological context as possible, while not occluding underlying structures. Conventionally, a simple transparency setting is used which does not provide much in the way of shape impression.

6

EXPLODED VIEWS

To illustrate a focus structure and representing the numbers of layers around the focus region, a carefully arranged exploded view, such as the one used in technical illustrations, may be suitable. When this technique is employed by a skilled illustrator, it may result in valuable educational illustrations. To uncover a hidden focus object, either the visual representations can be changed or the spatial relations can be adjusted in order to provide more insight. The exploded view approach makes use of the latter to 'break' the object apart and order it differently in the region surrounding the focus object. This technique was first applied on technical illustrations [57, 58], before it was adapted to visualize medical data sets [59, 60]. When using exploded view techniques, the original spatial relations may be lost. Furthermore, it may happen that important structures around a focus object, e.g., vessels around the tumor, are relocated, or even worse, break apart.

Application to multimodal data: For reasons outlined above, exploded views are not directly suitable for multimodal medical data visualization. One could consider to have one modality rendered in a traditional style with the information from an additional modality arranged around it. However, since the spatial relations, which are crucial in the interpretation of medical data are lost, these techniques are not promising from this perspective.

6.4.3. FOCUS-AND-CONTEXT VISUALIZATION

For focus-and-context visualizations, we assume that at least the focus object is segmented. Even though a lot of techniques exist that provide insight into the focus object, the question is how to represent them in the context of the surrounding structures. A straightforward approach is to use the same shading technique on all structures, which may lead to a hindered perception of the focus object by the context objects. Thus, the main goal of the *focus-and-context* visualization is to illustrate the focus object so that it is perceivable, and simultaneously illustrating the surrounding context structures without distracting from the focus. Techniques vary from using unsaturated colors for shading, to different shading methods per object class, to even varying entire rendering concepts.

Line drawings. One of the first line drawing that appeared in combination with focus objects was presented by Interrante et al. [61]. They used a hatching technique for the surrounding objects which are represented by the isosurface of radiation. In this case, anatomical data is combined with a scalar field resulting from the simulation of dose distribution, a standard approach for radiation treatment planning. The lines were aligned along the principle curvatures directions for an improved shape perception, whereas the inside focus object was simply shaded. Another line drawing technique that appeared in combination with focus objects was presented by Treavett and Chen [62]. They developed line drawings that can be applied on volume data sets and furthermore, they showed how it can be used to illustrate the surrounding context object, e.g., the skin, with the focus object, e.g., the skull. Again in the context of volume rendering, Hauser et al. [63] used combined visualization techniques to illustrate different objects. For instance, bones are rendered with direct volume rendering, surface rendering is applied to the vessels, and contours indicate the skin. Lum and Ma [64] presented a way to render a dataset with non-photorealistic techniques efficiently. Several illustrative techniques could then be applied with interactive framerates. It was shown that line drawing concepts can successfully be used for focus-and-context visualization [65]. The core idea here was to provide the flexibility to combine surface-, volume-data, and line drawing for a sparse visual representation of context objects. Here, not only line drawings, but illustrative visualization techniques in general were used to visualize context objects. The key idea of illustrative visualization is to provide a simplified and expressive depiction of a scene or problem [66]. Here, details are omitted purposefully in order to convey only the important information.

Indeed these different rendering styles may be used for multimodal visualization, e.g. contours representing boundaries of objects in one modality may be overlaid to the visualization of a second dataset, where no contours are used.

Tietjen et al. combined and tested [65] various rendering techniques. While their technique was not shown on multimodal data, it is applicable for feature fusion. They confirmed that line drawings can be a meaningful extension to existing visualization techniques in an evaluation. Later, Tietjen et al. [67] used slice-based visualizations in combination with 3D illustration techniques and employed focus-and-context visualization in the medical context. In a similar way, a slice or a slab of CT/MRI data may be integrated in a SPECT/PET visualization or vice versa. The slice/slab may be even enhanced with contours from a segmented object. For PET/CT, Kim et al. [68] used different ren-

dering techniques to illustrate the PET enrichment and the anatomical context by using the CT data. Various rendering techniques were also used in the context of simultaneously rendering anatomical and functional brain data [69]. For an overview of several focus-and-context approaches, we refer to Cohen [70] and Bruckner et al. [71]. Wan et al. [72] proposed an importance-based accumulated transparency modulation to visualize focus objects. Their approach is a generalization of the maximum intensity difference accumulation technique. Another importance-based direct volume rendering approach was presented by Pinto and Freitas [73]. Furthermore, they combined standard volume rendering approaches with illustrative visualization techniques. For focus-and-context visualization with PET/CT hybrid data, Bramon et al. [74] suggested visualization techniques to emphasize PET activity in anatomical context. Another approach to explore and visualize interesting regions was proposed by Abellán et al. [75]. They showed their approach on MR/PET data and used illustrative visualization techniques to convey information.

Lens-based visualization. Focus-and-context visualization can also be applied interactively. One way to approach this, is to place a region on the surface that allows the user to look inside the data and reveals everything underneath the surface. This idea is called *magic lens* as it can be used to look through objects and to reveal interesting parts [76]. The magic lens can be applied in many situations, ranging from whole body that uses a lens to reveal organs, to vessels where a lens depicts special properties of the blood flow [77]. A special application in medicine is to employ the vertical symmetry of the human body and propagate the movement of the lens in one part of the body to the other part, thus supporting the comparison between a suspicious region and the analogue region on the other part of the human body. Such lenses are employed, e.g., in the exploration of brain perfusion data, where one perfusion parameter is shown in the lens region and the other in the background. For 3D ultrasound, Schulte zu Berge et al. [78] presented a framework where the user can set various predicates such that interesting regions are highlighted. Increasing the slider for the vessels makes it more visible than the surrounding structures. In the context of vascular models and tumor visualization, Lawonn et al. [79] presented a way to depict the vessels with line drawings techniques whereas the tumor is visualized with diffuse shading. Furthermore, lenses can also be employed to control which source is displayed, e.g., in the lens region PET (or PET and CT combined) and outside only CT. Thus different layers may be shown interactively [80].

The distinction between smart visibility and focus-and-context visualization is not that clear. We argued that focus-and-context visualization techniques are a subcategory within the broader category of smart visibility approaches. The main goal of smart visibility is to provide insights into interesting regions or objects and in this way, the same can be said about focus-and-context techniques. However, this goal can be also be achieved by adding 2D image planes in a 3D visualization representation [81, 82]. In this case, there is no focus object involved, although it is still clearly a smart visibility technique. Thus, we consider focus and context visualization a subgroup of smart visibility techniques, which is a broader category including cutaways, ghosted views or exploded views, as well as techniques in which there is no focus object.

Application to multimodal data: Focus-and-context techniques may be applicable to

multimodal medical data and can be employed in two main ways. First, both modalities may be combined and a region of interest within the datasets can serve as a focus area, while the context is rendered in a different style. Secondly, one modality with a higher importance can be rendered as the focus dataset, while the other is rendered to provide context. The rendering style required strongly depends on the application. For instance, if the focus region consists of a tumor and the physician is interested in the distance to nearby vessels, which serves as the context object, then the morphological vessel structure may be not as important as an indicated position of the vessel structures. In this case a simple contour may be sufficient that shows only the outline and allows for a fast recognition of the spatial extent. But in case of a liver tumor, the spatial impression and the shape of the liver is important for the treatment planning. In this case, a more advanced approach needs to be applied to illustrate the tumor in a reasonable way while simultaneously depicting the liver surface with important spatial and morphological features.

6.4.4. SUMMARY

There are different channels, with which we perceive information, e.g., color, or orientation of elongated objects. Thus, we can perceive information simultaneously if it is encoded in different channels. The major theory here is the Feature Integration Theory, developed by Anne Treisman [23]. Multimodal visualization applications may benefit from the use of different visualization techniques that can be easily perceived simultaneously, e.g., color is only used for depicting one dataset, and line drawing techniques for the other. In order to represent multiple modalities in a single rendering, smart visibility techniques are crucial to prevent visual clutter and occlusion problems. Specifically, multimodal medical visualization can benefit from deciding which dataset has priority, i.e., the "focus dataset", and which is the context dataset. In multimodal visualization design, one could consider indicators of importance of a modality, such as the size and resolution, which could for instance be indicative of which dataset is best used for anatomical detail. Then one could think about which techniques are primarily suitable to show the focus dataset, likely revealing more details, and which to show the context dataset. This approach forms a bridge between single data focus-and-context techniques and focus-and-context in the context of multimodal data. There may be also situations where both volumes have the same importance, e.g., CT and MRI, and in these situations, it might be more appropriate to consider a region within both datasets the focus area, while the region outside is rendered as context. In conclusion, visualization techniques that have been developed for single modality visualization can be adapted to multimodal data, but this requires careful consideration of the application area requirements and in some cases extensions of existing techniques.

6.5. MULTIMODAL RENDERING AND INTERACTION TECHNIQUES

In this section, we discuss rendering and interaction techniques for multiple or fused volumes. First, we discuss recent contributions on rendering aiming at fusing multiple volumes together at different points in the rendering pipeline. After this, we discuss interaction techniques that are designed for use with multimodal medical data, including

clipping and specialized transfer functions.

6.5.1. RENDERING TECHNIQUES

In this subsection we provide an overview of rendering techniques suitable for rendering two or more volumes simultaneously. Although, Drebin [83] and Levoy [84] are often cited as the developers of the first volume rendering approach, actually Höhne [85] was the first who rendered a volume based on CT data. As a follow up, Höhne [86] were the first who developed a multimodal visualization technique visualizing MRI combined with MRA. Cai and Sakas presented a seminal work that allows rendering of multiple volumes [87]. They presented a method to achieve a reasonable mix of opacity, color, and illumination. Furthermore, they defined three levels of volume intermixing:

- Illumination model level intermixing
- Accumulation level intermixing
- Image level intermixing

In illumination model level intermixing, the opacity and intensity at each sample point is calculated directly from a multi-volume illumination model, instead of mixing several opacity and intensity values. This type of volume intermixing is the most complex, but also the most realistic. In accumulation level intermixing, the intermixing is performed during ray sample accumulation by mixing opacities and intensities from different volumes. This method requires changes to the rendering pipeline and the time consumption can be high, but provides correct depth cues. The simplest intermixing method is image level intermixing, where the two rendering images are merged on the pixel level. For this, no changes need to be made to the rendering pipeline, but there is a lack of correct depth cueing. This approach was successfully applied in the context of radiotherapy treatment planning, combining CT, dose and segmented object volumes.

Several multi-volume techniques were shown on multimodal medical data specifically. Wilson et al. presented a visualization approach to depict different image modalities together at interactive framerates using different rendering styles [88]. To render the result efficiently, they used the GPU and demonstrated their results with two case studies. One of the case studies featured a multimodal mouse data set combining PET and MRI acquisitions. Ferre et al. discussed strategies to visualize multimodal volume datasets using direct multimodal volume rendering (DMVR) [89]. Furthermore, they stated requirements and proposed five rendering methods, that differ in the step of the rendering pipeline at which the fusion is performed: property -, property and gradient -, material -, shading - and color fusion. They used SPECT and MRI data to present a combined visualization and evaluate the suitability of the five methods for this data. Another technique to render multimodal volumes that additionally solves depth cuing issues was presented by Hong et al. [90]. They conducted experiments and could show that their method can distinguish the depth of overlapping regions and produces rendering results faster than conventional software using High Level Shader Language (HLSL) for efficient volume fusion. They show the results of their method on PET and CT images in a performance evaluation. Rößler et al. described a GPU-based multi-volume

rendering scheme that allows users to visualize an arbitrary number of volumes interactively, focused on functional brain images [91]. Subsequently, they introduced dynamic shader generation for multi-volume ray casting [92]. In an accumulation-level volume rendering context, their system automatically generates different shaders per volume from the configuration of an interactive abstract render graph. Brecheisen et al. employed a framework to render different volumetric datasets together [93]. They used depth peeling to visualize overlapping volumes that can be intersected with an arbitrary number of geometric shapes. Kainz et al. proposed a framework for rendering multiple volumes on the GPU [94]. Theirs was the first framework for multi-volume rendering that still provided interactive frame rates while rendering a considerable amount of overlapping volumes. Recently, Sunden et al. introduced a volume illumination technique specifically designed for use with multimodal data [95]. They proposed a new *light-space-based* volume rendering algorithm, that employs illumination-importance metrics to compress and transform multimodal data into an illumination-aware representation. Their method was applied to CT and MRI data, as well as microscopy and simulation data and evaluated based on the quality and performance. Lindholm et al. [96] presented a rendering algorithm for hybrid volume-geometry data. They combined volume data and geometry-based data, i.e., a mesh. They applied their approach to proteins, vessels with blood flow, and DTI.

Additional related work can be found in the state-of-the-art report on the visualization of multivariate scientific data presented by Fuchs and Hauser [5]. Schubert and Scholl provide a performance and perceptual comparison of GPU-based multi-volume ray casting techniques [97]. Furthermore, they presented an overview of visualization techniques that use data intermixing approaches and direct volume rendering methods that use ray casting. They add Classification Level Intermixing to the three levels defined by Cai and Sakas [87], which mixes volumes by linear combination of the sample values at the beginning of the rendering pipeline. For a general overview of large-scale volume visualization techniques, we refer to the state-of-the-art report by Beyer et al. [98].

In the area of medical image fusion, Gupta et al. [99] presented a measure to evaluate PET/MRI image fusion. They could confirm the usefulness of their approaches in different experiments. For this, they used an entropy measure to combine the interesting parts in PET and MRI, but this approach also was applied to CT combined with MRI. Lindholm et al. introduced fused multi-volume direct volume rendering using a BSP, aimed at medical volume rendering of multimodal data [100]. Lindholm et al. provided a GPU-based ray casting approach to visualize intersecting volumes with regard to an efficient depth sorting of the resulting fragments. Bramon et al. fused image modalities and applied their technique to CT, MRI and PET data [101]. This information-theoretic framework automatically selects the most informative voxels from two volume data sets. They evaluated the potential of their technique with medical experts and concluded that it is potentially useful for planning radiotherapy, treatment monitoring, and planning brain surgery. In their evaluation, they proposed several information maps and fused data sets and had the experts vote on the quality. Their method performed well in differentiating between bone and cerebral tissue, as well as between morphological and functional data. More recently, Kim et al. introduced a slab-based intermixing method for fusion rendering of multiple medical objects [102].

We refer to James and Dasarathy [103] and Galande and Patil [104] for a more extensive discussion.

6.5.2. INTERACTION TECHNIQUES

In this subsection we describe several interaction techniques that were successfully used in the context of multimodal medical data. Interaction techniques primarily serve to adjust which portions of the dataset are visible. Visibility may be adjusted with clipping and cutting on a geometrical basis and with transfer functions on the basis of attribute values, e.g., intensity values, gradient magnitude or curvature.

Clipping and cutting. In the work by Hastreiter et al. clipping planes were employed per volume to cut the separate modalities in different ways [81]. This can be used to look at the slices while simultaneously analyzing the 3D visualization. Furthermore, they combined CT and MR images in a single view. Fairfield et al. proposed a custom clipping solution for co-registered MRI and CT, which they refer to as 'curtaining' [105]. This allows the user to define a clipping region in which the other modality will be shown. Manssour et al. propose a method to visualize inner structures in multimodal volume data employing cutting and data intermixing [106]. They applied it to the skull where cutting planes were used to analyze the brain. More advanced ways of clipping data were developed by Weiskopf et al. via per-fragment operations in texture-based volume visualization [107]. While their method was applied to a single modality, their technique was successfully employed by Rößler et al. to clip the brain based on regions in an atlas of the human brain combined with fMRI data [91]. More recently, a membrane clipping approach was proposed that avoids cutting through features in the data [108]. While it was not directly applied to multimodal datasets, it could be used in such a way. Furthermore, their technique features slab rendering which is a trade-off between slice-based viewing and full 3D rendering.

Transfer functions. To tackle occlusion problems, transfer functions are essential and have been designed for multi-modal data. 2D transfer functions, introduced by Levoy in 1988 [84], are an essential method for multimodal visualization, since they allow the user to emphasize boundaries and details within structures. Kniss et al. applied a multi-dimensional transfer function to multivariate weather data that allows experts to readily identify specific zones for their analysis task [109]. They presented a case study in which they explore the utility of multidimensional transfer functions for the visualization of multivariate fields. Kim et al. introduced a dual-lookup table for PET/CT data such that medical experts can set different transfer functions for every volume in a single view [68]. They provided techniques such that the volumes can be merged properly with different transfer functions. A general information-based approach for transfer functions was introduced by Haidacher et al. [110]. They provided a transfer function space to visualize certain tissues of multimodal data. They applied their method successfully to PET/CT data of the brain. Maciejewski et al. proposed a clustering technique to generate 2D transfer functions [111]. Their approach yields a clustering in the 2D histogram, which can then be adjusted by the user to get appropriate 3D visualization results. While this is an advanced and elegant solution, the necessary type of interaction limits users to visualization experts.

Furthermore, Bramon et al. introduced an information-theoretic observation chan-

nel for volume visualization, which quantifies the information transfer between the source data set and the rendered image [112]. With these quantities the authors provide the possibility to select informative viewpoints, to design transfer function, and to position the light. Joshi et al. introduced new interaction techniques to explore and visualize multimodal data [113]. Their technique allow for a precise control of the shape of the region in the brain that can be used to crop away to gain insight into the data during exploration and surgery. For their approach they used SPECT/MRI and fMRI data. Additionally, Ropinski et al. proposed interaction techniques to show a close-up of an interesting region in the context of the rest of the data [114]. They applied their technique to PET/CT data registered with MRI. More recently, Haidacher et al. proposed a approach for volume analysis based on multimodal surface similarity [115]. With this, the similarity space generated can be used for isosurface selection in applications like Dual Energy Computed Tomography (DECT). The similarity map can be used to set different isovalues for, e.g., DECT, such that inner structures can be set opaque whereas outer structures can be set as transparent. Correa and Ma [116] employed visibility-driven transfer functions to illustrate important structures in comparison to the surrounded regions. Their approach provides the user with graphical cues that inform about the contribution of particular scalar values to the final image. Furthermore, they presented a semi-automatic transfer function design that solves an energy minimization problem, such that the visibility of an initial opacity transfer function that provides the desired importance is maximized.

6.6. APPLICATIONS

In this section, we provide an overview of recent multimodal medical visualization papers that are application-oriented. We further subdivide this section in applications developed in a research-oriented, diagnosis or treatment planning and guidance context. For diagnostic purposes, we distinguish applications developed for cardiology and those aimed at oncology. The treatment planning and guidance applications are further categorized into neurosurgery and radiotherapy planning.

6.6.1. MEDICAL RESEARCH

Multimodal medical data visualization applications were developed in a research context related to vascular pathology development as well as neuroscience.

PET/CT scans are often requested in cases related to clinical oncology for initial cancer staging or a follow-up during or after treatment [21]. However, Ropinski et al. visualized mouse aorta PET/CT scans in a medical research-oriented application related to the formation of plaque [118]. They propose a linked multi-view approach, using a specialized straightened multipath curved planar reformation combined with a multimodal vessel flattening technique. In a follow-up work, Diepenbrock et al. additionally looked into PET/CT visualization of mice arteries [117]. They use the vessel wall extracted from the CT scan to perform a normalized circular projection which allows the user to judge PET signal distribution in relation to the deformed vessel.

Nguyen et al. used an approach to visualize and interact with real-time fMRI data in a neuroscience research context [122]. They treat the fMRI signal as light emission and

Table 6.1: Table of references sorted according to the application type (Type), area (App), and to the acquisition scanners (hybrid scanner, single scanner and multiple scanners). It is stated if any segmentation (Seg) is needed (○ for unclear cases), if the approach was evaluated (qualitatively (Ql), quantitatively (Qt), with domain experts (D) or non-domain experts (N) or the performance (P), and what types of image modalities are used. The visualization techniques that are used are mentioned, subdivided into cutaways (CA), ghosted view (GV), focus-and-context (FC), and illustrative visualization (I-Vis).

References	Type	App	Scanner	Seg	Eval	Image Modality								Vis technique				
						CT	MRI	fMRI	DTI	PET	SPECT	US	CA	GV	FC	I-Vis		
[117]	Res	Vasc	Hybr	✓	QID	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗
[118]	Res	Vasc	Hybr	✗	QIN	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓	✗
[119]	Res	Neuro	Mult	✓	✗	✗	✓	✗	✓	✗	✗	✗	✓	✓	✓	✓	✗	
[120]	Res	Neuro	Sing	✗	QID	✗	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	
[121]	Res	Neuro	Sing	✗	QIN	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	
[122]	Res	Neuro	Sing	✗	QtN	✗	✓	✓	✗	✗	✗	✗	✓	✗	✓	✓	✗	
[91]	Res	Neuro	Sing	✗	QtP	✗	✓	✓	✗	✗	✗	✗	✓	✓	✓	✓	✗	
[123]	Res	Neuro	Sing	✓	QID	✗	✓	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	
[124]	Res	Neuro	Mult	✓	QID	✗	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	
[125]	Diag	Onco	Hybr	✓	QtP	✓	✗	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	
[68]	Diag	Onco	Hybr	✗	✗	✓	✗	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	
[28]	Diag	Onco	Hybr	✗	QID	✓	✗	✗	✗	✓	✗	✗	✓	✓	✓	✓	✓	
[126]	Diag	Onco	Mult	○	QIN	✓	✗	✗	✗	✓	✗	✓	✓	✗	✓	✓	✓	
[127]	Diag	Cardio	Mult	✓	QtD	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	
[128]	Diag	Cardio	Mult	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	
[129]	Diag	Cardio	Mult	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	
[130]	Diag	Cardio	Mult	✓	QtD	✓	✗	✗	✗	✗	✓	✗	✗	✗	✓	✓	✗	
[25]	Trea	Neuro	Mult	○	QID	✓	✓	✓	✗	✓	✗	✗	✓	✗	✓	✓	✗	
[131]	Trea	Neuro	Sing	✓	QID	✗	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓	○	
[132]	Trea	Neuro	Mult	✓	QID	✓	✓	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗	
[133]	Trea	Neuro	Sing	✓	✗	✗	✓	✓	✓	✗	✗	✗	✓	✓	✓	✗	✓	
[2]	Trea	Neuro	Sing	✓	QID	✓	✓	✓	✓	✗	✗	✗	✓	✓	✓	✗	✗	
[69]	Trea	Neuro	Sing	✓	QID	✗	✓	✓	✗	✗	✗	✗	✓	✓	✗	✓	✓	
[134]	Trea	Neuro	Sing	✓	QID	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	
[113]	Trea	Neuro	Mult	✓	QID	✓	✓	✓	✗	✗	✓	✗	✓	✓	✓	✓	✗	
[135]	Trea	Neuro	Mult	✓	QtD	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	
[136]	Trea	Neuro	Mult	✗	QtP	✗	✓	✗	✓	✗	✗	✗	✓	✗	✓	✓	✗	
[137]	Trea	Neuro	Sing	✓	QID	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	
[138]	Trea	Neuro	Mult	✓	QID	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	
[139]	Trea	Neuro	Sing	✓	QID	✗	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗	✗	
[140]	Trea	Neuro	Mult	✗	✗	✗	✓	✗	✗	✗	✗	✗	✓	✗	✓	✓	✗	
[141]	Trea	Neuro	Mult	✓	Ql/QtD	✓	✓	✗	✗	✗	✗	✗	✓	✗	○	✗	✗	
[142]	Trea	Neuro	Sing	✓	QID	✗	✓	✗	✗	✗	✗	✗	✓	✗	✓	✗	✗	
[143]	Trea	Radio	Sing	○	Ql/QtD	✗	✓	✗	✗	✗	✗	✗	✗	✗	○	✗	✗	
[29]	Trea	Radio	Hybr	✓	✗	✓	✗	✗	✗	✓	✗	✗	✓	✓	✓	✓	✗	

render it in the context of a patient-specific high resolution reference MRI scan. Using this technique, the brain glows and emits light from active functional regions with a 2 second delay of the measured fMRI signal. In the work by van Dixhoorn et al., the focus is on examining whole brain functional network connectivity at a voxel-level [120]. They visualize the correlation of the functional activity, where fMRI time-signals at each voxel are correlated with every other voxel in the brain to determine functional connectivity. Therefore they propose an application for the interactive visual analysis of this high resolution brain network data, both in a linked matrix representation as well as in its anatomical context based on MRI in a GPU raycasting framework.

6.6.2. DIAGNOSIS

Visualization techniques aimed at improving diagnostic value of multimodal medical data have been primarily developed in the fields of oncology and cardiology.




Figure 6.3: PET/CT visualization with PET activity presented as a focus area using dynamic cutaways and CT providing anatomical context for accurate localization [28].

ONCOLOGY

Among the oncological applications, Kim et al. introduced a dual-lookup table specifically designed for use with PET/CT data such that medical experts can set different transfer functions for every volume in a single view [68]. Jung et al. employed a novel visualization approach by integrating a visibility-driven transfer function specifically for PET/CT data [125]. Furthermore, they provided an intuitive region of interest selection tool for further exploration. Lawonn et al. recently developed an illustrative technique for focus-and-context rendering of PET/CT data [28]. In their approach, the functional information from the PET data is used as the focus, while the CT provides anatomical context using dynamic cutaway views (see Figure 6.3). Viola et al. presented illustrative techniques to visualize liver ultrasound combined with CT information serving as anatomical context [126]. They aim to shorten the long learning curve for ultrasound practitioners in training, as well as assist the interpretation of liver examinations.

CARDIOLOGY

Due to the severity and frequency of cardiac diseases this application area is of utmost importance. Essential diagnostic tasks include plaque assessment in the coronary arteries, detailed diagnosis of a heart infarction, e.g. the extent of infarct core, assessment of the heart valves, and assessment of abnormalities, such as congenital heart failures. Morphological information, extracted from CT and MRI, as well as functional information, e.g., wall motion extracted from ultrasound, are important.

An example of a multimodal visualization problem that was considered in visualization research is the use of perfusion data combined with other MR imaging modes of for the diagnosis of the coronary heart disease. Perfusion data of the heart, which indicates the blood perfusion in the myocardium and can be acquired from MRI or SPECT scanners, needs to be combined with anatomical data to provide morphology of the heart muscle and the coronary arteries. Similar to combining PET and CT data, the resolu-

tion of the perfusion data is much coarser. Moreover, the acquisition of cardiac perfusion data exhibits gaps, i.e., there are regions in the morphological data, where no corresponding perfusion data is available. Simply interpolating the missing information is not a satisfactory solution. Thus, the overlay of both information must convey where slices of the morphological data correspond to morphologic slices and where no such correspondence exists. Termeer et al. proposed a visualization to aid diagnosis of patients with Coronary Artery Disease using cardiac perfusion MRI data [129]. They extend the traditional bull's eye plot to a continuous volumetric version that reveals transmuralities of scar tissue and link this to an anatomical view of the heart. Transmurality indicates whether the whole wall of the (left) ventricle is affected by an infarction—information that is essential for prognosis and treatment. Hennemuth et al. developed a comprehensive approach to the analysis of contrast-enhanced cardiac MR images [127]. They propose a full workflow to align the different datasets, extract surfaces and analyze them in a combined way. Oeltze et al. presented an integrated multi-modal visualization of morphologic and cardiac perfusion data for the analysis of coronary artery disease [128]. Myocardial perfusion is measured using an MRI scanner and combined with CT Coronary Angiography (CTCA) which depicts the anatomy. In their visualization, colored icons, heightfields and lenses are used to visualize and explore the different parameters measured. Kirişli et al. combine CT Angiography (CTA) and SPECT myocardial perfusion imaging (MPI) in a visualization aimed at assessing coronary artery disease [130]. They evaluated the diagnostic value of a software-based image fusion system over conventional side-by-side analysis and found improved diagnostic performance when using their application.

6.6.3. TREATMENT PLANNING

Multimodal medical data visualization applications have been developed mainly in the areas of neurosurgical planning and guidance so far. To the best of our knowledge, there is no prior multimodal surgical planning application focusing on pelvic oncologic surgery. Some recent works have focused on radiotherapy treatment planning.

NEUROSURGERY

A primary focus in neurosurgical planning applications is identifying various surgically relevant structures that are at risk to be damaged, such as the functional areas in the gray matter and the white matter fiber tracts, and understanding how they are related to the lesion that needs to be removed. The surgically relevant structures can be acquired using several modalities such as CT and MRI, but also special MRI modes such as fMRI, DTI as well as multiple different scanning parameter combinations are often employed. Additionally, neurosurgeons need to plan the optimal or safe access path to such a lesion, with the least amount of damage to vital functional areas, derived from fMRI, and fiber tracts, derived from DTI. Visualization applications aimed at neurosurgical planning and guidance often focus on brain tumor resection and this optimal path planning, but also applications aimed at Deep Brain Stimulation (DBS) and Arteriovenous Malformations (AVM) surgery exist.

Jannin et al. fused several modalities and imaging modes from MRI for neurosurgery navigation [134]. They combined segmented structures and vessels with functional areas

from magneto-encephalography (MEG) and fMRI data in a neuronavigation application. As one of the most notable early techniques where surfaces are used for combined information, Stokking et al. presented a visualization method that combines functional input data and a surface extracted from anatomical data [123]. They map fMRI values onto a brain surface using the surface normal and evaluate their technique on both registered (S)PE(C)T/MRI and fMRI/MRI. This technique was an extension from earlier work on combining software-registered SPECT with a surface extracted from an MRI scan, by mapping functional values of the SPECT to the surface of the brain along the normal [124]. While these techniques were applied to the brain, the idea of using a surface can be transferred to other organs, such as the heart muscle. Rößler et al. described a GPU-based multi-volume rendering scheme that allows users to visualize an arbitrary number of volumes interactively (recall Section 5.1.1), and their technique was specifically focused on functional brain images [91]. Their tool was aimed at supporting cognitive neuroscientists in experimental studies and to communicate results to non-experts and employs a template brain with patient-specific fMRI data. Blaas et al. proposed an approach that fuses fMRI and DTI data for planning brain tumor resections [131]. In this fused volume, users can extract fiber bundles that pass through a region around the tumor. These bundles can then be explored by filtering on distance to the tumor, or by selecting a specific functional area using arbitrary convex geometries as selection criteria. Rieder et al. employed a combination of fMRI and DTI data for neurosurgical therapy planning [139]. Their application visualizes pathologies using a distance-based transfer function, as used in the work by Tappenbeck et al. [144], and only shows functional data in close proximity to the lesion. Furthermore, they increase depth perception by including a distance-ring, which visualizes how deep the lesion is situated in the brain from the current viewpoint. For the surgical planning itself, they provide access path visualization and rely on identification of superficial landmarks which can be translated to the per-operative context. A visualization approach for combined MRI and fMRI brain data was presented by Jainek et al. [69]. Various rendering styles were used, e.g. ambient occlusion, and illustrative techniques were employed to enhance the visual output. Following up on this work, Born et al. extended it to include DTI tracts, which reveal reconstructed nerve fibers that connect functional areas [133]. They also enhanced depth and shape perception by applying silhouettes and dithered half-toning. Janoos et al. proposed a method to visually analyze brain activity from fMRI data, with a special focus on temporal dependencies [121]. They propose a methodology to analyze the time dimension through volumes-of-interest, of which the selection is guided by a hierarchical clustering algorithm in the wavelet domain. They visualize these volumes-of-interest overlaid onto MRI data of the brain and show the cluster time-series of selected clusters in a separate view.

Diepenbrock et al. were the winners of the IEEE VIS Visualization contest on multi-modal visualization for neurosurgical planning [2]. They provided a 3D view inside the brain featuring an interactive probe containing structures acquired from fMRI, DTI and various additional MR imaging modes. Uncertainty in the fMRI and DTI data is revealed using uncertainty borders and a different rendering style. Furthermore, they provide a cylindrical access path projection representing the distance to the structures at risk along the planned path. A close-up view of the tumor allows a more detailed view on the

area that the surgeons are planning to resect.

Serra et al. developed a multimodal pre-operative neurosurgical planning system designed for use with a mirror-based Virtual Reality system workbench [141]. They render CT, MRI and MRA scans using 3D textures and allow the user to interact with the volumes using a tracked 3D pointer. Neubauer et al. proposed a surgical simulation application for endonasal transsphenoidal pituitary surgery, which is a minimally invasive endoscopic procedure [138]. They employ CT and MRI to simulate this endoscopic procedure for training and pre-operative planning purposes. For this, they used CT for skull anatomy, MRI for the tumor, optical nerve and pituitary gland, and contrast-enhanced CT and MRI to highlight the internal carotid artery. They provide interactive threshold adjustment to adjust the surface visualizations.

For planning neurosurgical procedures, Beyer et al. introduced a framework that can visualize multimodal volumes such as CT, MRI, fMRI, PET and Digital Subtraction Angiography (DSA) [25]. They propose a skull peeling algorithm that can automatically remove an occluding bony area to reveal the underlying brain. Furthermore, they developed a masked-based approach to show multiple modalities concurrently and a rendering technique to render binary segmented objects with a smoothed appearance when available. Within masked parts of the volume, different datasets are rendered, each with their own transfer function settings. Joshi et al. presented interaction techniques in a neurosurgical planning application that includes stereotactic navigation [113]. They visualize MRI, fMRI, DTI and SPECT in a visualization that allows the user to crop volumes using an interactive line widget. Kin et al. developed a neurosurgical planning tool specifically focussing on brainstem malformations in which they fuse MRI, CT, and 3D rotational angiography together [135]. They mainly rely on combining multiple surface reconstructions of the individual modalities, which demands a significant amount of preprocessing time in manually segmenting important structures.

Bock et al. recently presented a tool for planning and guiding deep brain stimulation (DBS) interventions by fusing multimodal data that includes uncertainty regions from the acquisition process [132]. In order to guide these procedures, CT and MRI are fused and visualized integrated with results from Microelectrode Recordings (MER), which measure the electric field in the brain intra-operatively. Their tool features a planning, recording and placement phase in which the corresponding steps of the intervention can be performed.

In the area of neurosurgical planning, Rieder et al. proposed a multimodal visualization of intracerebral pathological tissue [140]. They use multiple MRI sequences (T1, T1ce (contrast enhanced), T2, FLAIR) as the input for their visualization application and perform clustering to determine pathologic regions. Next, they blend these pathologic regions with the anatomical context information using an automatically calculated transfer function. Furthermore, they propose an automatic cutting tool and brain peeling to reveal hidden structures of interest.

Weiler et al. presented work on neurosurgical planning for treatment of Arteriovenous Malformations (AVMs) in the brain [142]. In these procedures, a precise identification of the arteries and veins is crucial to understand the complex inflow and outflow in these vascular pathologies. They combine several differently weighted MRI scans, such as T1-weighted images with and without contrast agent, arterial time-of-flight (TOF),

and MR-Venographies (MRV), into a single multi-volume visualization to facilitate understanding of the lesions angio-architecture. To prevent clutter, they propose a focus of attention technique, based on the distance to a point of interest, that allows the user to attenuate importance using transparency and saturation manipulation for structures outside the region of interest. Navkar et al. developed visualization tools for planning neurosurgical interventions with straight access, such as biopsies, deep brain stimulation and ablation of brain lesions [137]. For this, they generate various access maps based on vascular structures on the surface of the skin of the patients head for guidance in selecting a safe entrance point.

Much research has been done on the topic of visualizing fiber tracts from DTI scans, often using anatomical information from an MRI scan as context. Enders et al. proposed to visualize white matter tracts reconstructed from DTI with wrapped streamlines [119]. They generate surfaces that wrap around the convex hull of fiber bundles for a more intuitive representation of tracts and combine this with anatomical information from a T1-weighted scan. Merhof et al. proposed a new visualization technique for white matter tracts using triangle strips and point sprites [136]. Their novel DTI hybrid rendering approach speeds up the rendering process and can then be combined with DVR of anatomical data to provide additional information for surgical planning. Since a full overview of these visualizations is outside the scope of this survey, we refer to a recent survey written by Tobias Isenberg on illustrative visualization techniques for diffusion-weighted MRI tractography [145].

RADIOTHERAPY

Radiotherapy planning is generally based on CT, MRI and PET/CT scans. These volumes are used to define target areas for radiation but, also organs at risk that should not receive high radiation. Especially for target areas such as chest and upper abdomen, that move strongly due to breathing, the recent possibility to acquire 4D PET/CT data opens the possibility to capture and integrate the movement of tumors into the radiation target volume definition, building the basis for the dose calculation. In contrast to many other diagnostic and treatment planning procedures that occur under severe time pressure, radiation treatment planning is advanced and time-consuming. Target structures and structures at risk are segmented regularly to assess how they would be affected by a certain dose distribution. Thus, advanced multimodal visualization techniques would be applicable.

Schlachter et al. [29] showed that 3D and 4D visualization of images, combined with delineated regions, and the calculated dose (also available as volume dataset), complements the common slice views. This provides a fast overview over the spatio-temporal configuration of all delineated areas and the related dose distribution resulting finally in a faster quality checks of the radiation plan.

Magnetic Resonance Spectroscopy Imaging (MRSI) data provides metabolic information that quantifies the concentrations of multiple brain metabolites, such as Choline and Creatine, per voxel. Currently MRSI is only performed in a clinical research context, but not yet employed in clinical practice. In the work by Nunes et al., MRSI data is fused with multimodal radiology imaging in an integrated visual analysis system aimed at radiotherapy treatment [143]. They linked the medical imaging framework MITK and

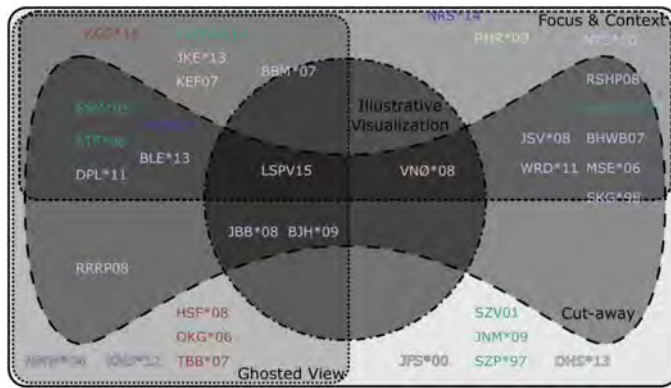


Figure 6.4: A Venn-Edwards diagram of the references categorized according to the applied visualization techniques. The color of the references corresponds to the application type colors: **vascular research**, **neuroscience**, **diagnostic oncology**, **diagnostic cardiology**, **neurosurgical treatment planning** and **radiotherapy planning**.

6

general purpose data exploration tool ComVis to analyze, relate and visualize MRSI data together with multimodal images.

6.6.4. DISCUSSION

An overview of all papers according to the application type and application area can be found in Table 6.1. Here, we stated if any segmentation is needed for processing the data in order to visualize it, which is important to estimate the amount of preprocessing an application requires. We mention which type of evaluation (if any) was performed to verify the utility of the applications. For this, we distinguish between quantitative (Qt) and qualitative (Ql) studies with domain experts (D), non-domain experts (N), or if the performance was tested (P). Furthermore, we mention the image modalities visualized and describe the visualization techniques that are applied for the visualization of the multimodal medical data. Additionally, we have composed a Venn-Edwards diagram of all application papers categorized according to the visualization techniques used and the application type in Figure 6.4.

From the table we can see that a fair amount of work has already been done on neurosurgical treatment planning, as well as work on guidance. Not much work has been done yet on radiotherapy planning, while this field especially uses multimodal data and is in need of suitable visualization methods. In radiotherapy, relevant structures are regularly segmented, making all described visualization techniques in Section 4 practically applicable. There has only been one visualization application paper combining ultrasound with another modality, while ultrasound has great potential and benefits that can be complimentary to other modalities, such as the ability to provide real-time information. The limited field of view and noise in ultrasound data could be alleviated by combining it with additional modalities, and could be of value also in for instance biopsy guidance. SPECT is involved only in cardiac diagnostic and neuro-science and -surgical

planning applications. A popular modality combination for both vascular research and oncologic diagnosis is PET/CT, which is used in all the application papers we listed. Due to the nature of MRI, fMRI and DTI data, these modalities are often involved in both neurosurgical treatment planning and research. For obvious reasons, fMRI is completely neuro-specific, but DTI is also clinically applied so far mainly in neurological contexts.

While the visualization techniques used per application area vary too much to make strong conclusions based on the limited number of samples, it is clear that ghosted views are often employed for diagnostic purposes (see Figure 6.4). Furthermore, cutaways are most successfully applied to neurosurgical planning, which makes sense due to the nested anatomical and pathological structures involved and the need for exact access path planning through the skull. There are five papers not employing smart visibility or illustrative techniques, mainly in a research context.

In most works qualitative evaluations with domain experts are presented, but quantitative evaluations as well as more elaborate clinical studies are often still needed. Furthermore, our overview revealed that only a few works applied illustrative visualization techniques [28, 69, 126, 133]. However, the application of illustrative visualization technique seems to be promising [65, 66] and we see a potential to use these techniques for further research on multimodal medical visualization. In multimodal medical visualizations, there are often many overlapping and nested structures, which causes occlusion problems that the abstraction in illustrative techniques can help alleviate. However, the segmentation required often limits the clinical uptake. It might also be helpful to employ novel focus-and-context techniques, by defining one modality as the focus and the other(s) as the context.

6.7. CONCLUSION AND FUTURE CHALLENGES

For many techniques described in this survey, it is not clear how they fit in clinical workflows, how much additional and relevant information they provide and if so, whether this justifies a potentially larger effort, e.g., due to the necessity to segment structures or adjusting complex 2D transfer functions.

Dealing with medical imaging data involves uncertainty in the form of imaging errors, for instance resulting from noise, field bias, patient motion, or imaging artifacts. Additionally, processing errors may occur due to errors in the segmentation or registration process. When visualizing multi-modal datasets resulting from combining acquisitions of multiple scanners, additional care has to be taken to visualize the uncertainty resulting from these processing errors in the registration process. This is a challenging and so far unsolved research challenge [146].

The existing multimodal visualization techniques need to be better evaluated and compared. On the one hand, evaluations that consider effectiveness in specific medical tasks are needed. On the other hand, perception-based evaluations are essential, e.g., to understand how effective certain emphasis techniques actually are, how well users could discriminate values with certain color scales (used as overlays) and how accurate they can locate functional abnormalities.

The existing techniques are spread over many research prototypes and commercial solutions. It is thus very difficult to compare them. A framework that integrates at least the most commonly used techniques and flexible parameterization options would be

very valuable. A comparison also benefits from benchmark tasks and data. A first step in this direction was the IEEE Vis. 2010 contest on multimodal brain data for neurosurgery planning. A few more such tasks, e.g., in cardiology, would be beneficial.

We have seen a multitude of visualization techniques. Interestingly, these are exclusively techniques that were developed first for displaying single datasets and then adapted and refined. An open question is whether truly multimodal techniques can be developed. Multimodal visualization obviously benefits from multiple views. More research is necessary to understand which views are essential, how flexible viewing configurations should be and which synchronization/coordination techniques are needed.

The visualization of single medical datasets may benefit from automatic viewpoint selection, i.e., choosing a good initial view on a dataset as starting point for further exploration [147–152]. Measures of visibility, size of objects and viewpoint stability are employed for the viewpoint selection. Automatic viewpoint selection is also promising for multimodal data visualization and requires careful adaptation of measures and application to selected case studies. Also the user interface must be considered to enable a physician to specify what is important for her as input for an algorithm.

Neurosurgery is among one of the most advanced user communities in this field, while other disciplines like radiotherapy are not yet exploiting the full potential of visualization techniques available today to support therapy planning and monitoring. We believe there are many applications in neurosurgery because of the relatively easy and reliable registration process, which can be automated, as well as specific application demands. First, CT and MRI contribute a significant amount of relevant and complementary information, while in lung surgery for example you can completely rely on CT. Second, neurosurgery has much higher accuracy demands compared to any type of abdominal surgery. A few millimeters in abdominal surgery usually do not matter, while in neurosurgery 2mm may be a lot.

Challenges also appear in case when the input does not consist of surface meshes. For volumetric data, some visualization techniques are more challenging than when applying them to surfaces. Feature lines are mostly defined on surfaces, but some methods can also be applied on volumetric data. Burns et al. [153] extended the suggestive contours to volumetric data. They also extract the lines as objects, and with this the lines could also be stylized. Kindlmann et al. [154] visualized ridges and valleys depending on curvature values. With an adapted transfer function, these lines could then be highlighted, which is similar to the approach by Lawonn et al. [28]. In contrast to feature lines, hatching approaches were also applied to volume datasets [155–159]. Mostly particles or points are placed in the dataset on a specific isovalue and then the points are traced along the principle curvature directions. Due to the challenging character of feature lines, these techniques were not commonly applied to multimodal data visualization so far. Although its potential was shown to illustrate surfaces [160] more work need to be done in this field.

Developing efficient and effective visualization methods with real impact on daily clinical or research routine of potential users requires a highly interdisciplinary effort fusing knowledge from Physicians, Visual Computing Specialists, HCI, Imaging and Image Analysis specialists. This becomes even more important with the increasing availability of new imaging modalities, that can be fused with imaging pipelines delivering

sets of highly heterogeneous data at different scales. A deep understanding of this data and its application context is necessary to generate useful results, on the other hand technical skills are getting more and more important, as the amount of data to be processed increases tremendously. As a consequence, multimodal visualization will be even more interdisciplinary than it has been ever before.

6.8. SUMMARY

We have presented a survey of the current state of the art in medical multimodal visualization. We introduced relevant medical imaging modalities and acquisition techniques, i.e., CT, MRI, fMRI, DTI, PET, and SPECT, and described their properties as well as the associated visualization challenges. Afterwards, we examined the current clinical workflow for the exploration and analysis of multimodal medical modalities. We summarized the requirements in designing a visualization technique to maximize the insights into relevant details in the depiction of multimodal medical data.

Subsequently, we highlighted the most common visualization techniques that support this visualization problem. These techniques need to incorporate heuristics for assessing the importance of information and emphasis techniques to adapt the importance. For this reason, smart visibility approaches, including focus-and-context techniques, ghosted views and cutaways are highly relevant for visualizing multiple volumes. While these techniques are frequently used in visualization research, they are not part of any commercial solution or available in radiology workstations. This is probably because these techniques require time-consuming preprocessing and are quite complex, as well as unfamiliar to the physicians.

Rendering multiple volumes is an associated research challenge for which many approaches have been developed. Furthermore, special interaction techniques have been designed for exploring multimodal datasets, which can be as important in providing insight in multimodal data as rendering refinements.

In the main part of our survey, we provided an overview of 35 visualization application papers. We defined the scope of the paper by focusing our survey on multimodal medical visualization applications associated to research, diagnosis and treatment planning or guidance. Within these categories, we found that the main application areas were related to oncology, cardiology, radiotherapy and neurology. We summarized our findings in a table featuring the three aforementioned categories, a further subdivision according to the medical application domain, whether preprocessing is required in terms of segmentation, image modalities used, and visualization techniques employed. We concluded from the table which application areas are active, upcoming or sufficiently researched. While a fair amount of work has been done on neurosurgical planning, radiotherapy planning seems to be a rising field with many opportunities for interesting research, and the same holds for combining ultrasound with other modalities. Additionally, we did not find any multimodal visualization applications combining atlas data with medical imaging acquisitions for surgical planning for pelvic oncological procedures.

Finally, we mentioned open problems in multimodal medical visualization that as a starting point for future research in this area. Evaluation of medical task performance and in the form of perceptual studies are needed to validate the effectiveness of visualization application designs. Furthermore, most developed techniques were extended

from existing single modality techniques, while truly multimodal techniques have yet to be developed. With the increase in the amount of data acquired and new modalities being brought into clinical practice, multimodal medical visualization remains a promising research area for future developments.

REFERENCES

- [1] K. Lawonn, N. N. Smit, K. Bühler, and B. Preim, *A survey on multimodal medical data visualization*, Computer Graphics Forum (2016), in Submission.
- [2] S. Diepenbrock, J.-S. Praßni, F. Lindemann, H.-W. Bothe, and T. Ropinski, *2010 IEEE visualization contest winner: Interactive planning for brain tumor resections*, IEEE Computer Graphics and Applications, 6 (2011).
- [3] B. Preim and C. P. Botha, *Visual Computing for Medicine: Theory, Algorithms, and Applications* (Morgan Kaufmann Publishers Inc., 2013).
- [4] N. Sauber, H. Theisel, and H.-P. Seidel, *Multifield-graphs: An approach to visualizing correlations in multifield scalar data*, Visualization and Computer Graphics, IEEE Transactions on **12**, 917 (2006).
- [5] R. Fuchs and H. Hauser, *Visualization of multi-variate scientific data*, Computer Graphics Forum **28**, 1670 (2009).
- [6] J. Blaas, C. P. Botha, and F. H. Post, *Interactive visualization of multi-field medical data using linked physical and feature-space views*, in *Proc. of EuroVis* (2007) pp. 123–130.
- [7] J. A. Maintz and M. A. Viergever, *A survey of medical image registration*, Medical image analysis **2**, 1 (1998).
- [8] J. P. Pluim, J. A. Maintz, and M. A. Viergever, *Mutual-information-based registration of medical images: a survey*, Medical Imaging, IEEE Transactions on **22**, 986 (2003).
- [9] O. Fluck, C. Vetter, W. Wein, A. Kamen, B. Preim, and R. Westermann, *A survey of medical image registration on graphics hardware*, Computer methods and programs in biomedicine **104**, e45 (2011).
- [10] B. Köhler, S. Born, R. E. van Pelt, U. Preim, and B. Preim, *A survey of cardiac 4D PC-MRI data processing*, (2015) pp. 139–148.
- [11] S. R. Cherry, *Multimodality imaging: Beyond PET/CT and SPECT/CT*, Seminars in nuclear medicine **39**, 348 (2009).
- [12] B. J. Pichler, A. Kolb, T. Nägele, and H.-P. Schlemmer, *PET/MRI: paving the way for the next generation of clinical multimodality imaging applications*, Journal of Nuclear Medicine **51**, 333 (2010).

- [13] O. A. Catalano, B. R. Rosen, D. V. Sahani, P. F. Hahn, A. R. Guimaraes, M. G. Vangel, E. Nicolai, A. Soricelli, and M. Salvatore, *Clinical impact of PET/MR imaging in patients with cancer undergoing same-day PET/CT: initial experience in 134 patients - a hypothesis-generating exploratory study*, *Radiology* **269**, 857 (2013).
- [14] A. Drzezga, M. Souvatzoglou, M. Eiber, A. J. Beer, S. Fürst, A. Martinez-Möller, S. G. Nekolla, S. Ziegler, C. Ganter, E. J. Rummeny, *et al.*, *First clinical experience with integrated whole-body PET/MR: comparison to PET/CT in patients with oncologic diagnoses*, *Journal of Nuclear Medicine* **53**, 845 (2012).
- [15] D. Townsend, J. Carney, J. Yap, and N. Hall, *PET/CT today and tomorrow*, *The Journal of Nuclear Medicine* **45**, 4S (2004).
- [16] M. S. Judenhofer, H. F. Wehrl, D. F. Newport, C. Catana, S. B. Siegel, M. Becker, A. Thielscher, M. Kneilling, M. P. Lichy, M. Eichner, *et al.*, *Simultaneous PET-MRI: a new approach for functional and morphological imaging*, *Nature medicine* **14**, 459 (2008).
- [17] T. Beyer, D. W. Townsend, T. Brun, P. E. Kinahan, M. Charron, R. Roddy, J. Jerin, J. Young, L. Byars, and R. Nutt, *A combined PET/CT scanner for clinical oncology*, *Journal of Nuclear Medicine: Official Publication, Society of Nuclear Medicine* **41**, 1369 (2000).
- [18] R. J. Hicks and M. S. Hofman, *Is there still a role for SPECT-CT in oncology in the PET-CT era?* *Nature reviews Clinical oncology* **9**, 712 (2012).
- [19] A. K. Buck, S. Nekolla, S. Ziegler, A. Beer, B. J. Krause, K. Herrmann, K. Scheidhauer, H.-J. Wester, E. J. Rummeny, M. Schwaiger, *et al.*, *SPECT/CT*, *Journal of Nuclear Medicine* **49**, 1305 (2008).
- [20] G. Mariani, L. Bruselli, T. Kuwert, E. E. Kim, A. Flotats, O. Israel, M. Dondi, and N. Watanabe, *A review on the clinical uses of SPECT/CT*, *European journal of nuclear medicine and molecular imaging* **37**, 1959 (2010).
- [21] M. Benamor, L. Ollivier, H. Brisse, G. Moulin-Romsee, V. Servois, and S. Neuenchwander, *PET/CT imaging: what radiologists need to know*, *Cancer Imaging* **7**, S95 (2007).
- [22] G. S. Lin, H. H. Hines, G. Grant, K. Taylor, and C. Ryals, *Automated quantification of myocardial ischemia and wall motion defects by use of cardiac SPECT polar mapping and 4-dimensional surface rendering*, *Journal of nuclear medicine technology* **34**, 3 (2006).
- [23] A. M. Treisman and G. Gelade, *A feature-integration theory of attention*, *Cognitive psychology* **12**, 97 (1980).
- [24] L. K. Griffeth, *Use of PET/CT scanning in cancer patients: technical and practical considerations*, *Proceedings (Baylor University. Medical Center)* **18**, 321 (2005).

- [25] J. Beyer, M. Hadwiger, S. Wolfsberger, and K. Bühler, *High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions*, *Visualization and Computer Graphics*, IEEE Transactions on **13**, 1696 (2007).
- [26] A. Rosset, L. Spadola, L. Pysher, and O. Ratib, *Navigating the fifth dimension: Innovative interface for multidimensional multimodality image navigation 1*, *RadioGraphics* **26**, 299 (2006).
- [27] J. W. Wallis, T. R. Miller, C. A. Lerner, and E. C. Kleerup, *Three-dimensional display in nuclear medicine*, *Medical Imaging*, IEEE Transactions on **8**, 297 (1989).
- [28] K. Lawonn, N. Smit, B. Preim, and A. Vilanova, *Illustrative multi-volume rendering for PET/CT scans*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.
- [29] M. Schlachter, T. Fechter, U. Nestle, and K. Bühler, *Visualization of 4D-PET/CT, target volumes and dose distribution: Applications in radiotherapy planning*, in *Proc. of MICCAI Workshop on Image-Guided Adaptive Radiation Therapy* (2014).
- [30] W. Wein, A. Khamene, D.-A. Clevert, O. Kutter, and N. Navab, *Simulation and fully automatic multimodal registration of medical ultrasound*, in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2007) pp. 136–143.
- [31] W. Wein *et al.*, *Multimodal Integration of Medical Ultrasound for Treatment Planning and Interventions*, Ph.D. thesis, Technische Universität München (2007).
- [32] W. E. Lorensen and H. E. Cline, *Marching cubes: A high resolution 3D surface construction algorithm*, in *ACM siggraph computer graphics*, Vol. 21 (ACM, 1987) pp. 163–169.
- [33] N. Otsu, *An automatic threshold selection method based on discriminate and least squares criteria*, *Denshi Tsushin Gakkai Ronbunshi* **63**, 349 (1979).
- [34] P.-S. Liao, T.-S. Chen, P.-C. Chung, *et al.*, *A fast algorithm for multilevel thresholding*, *J. Inf. Sci. Eng.* **17**, 713 (2001).
- [35] H. Seim, D. Kainmueller, H. Lamecker, M. Bindernagel, J. Malinowski, and S. Zachow, *Model-based auto-segmentation of knee bones and cartilage in MRI data*, in *Proc. of Medical Image Analysis for the Clinic: A Grand Challenge* (2010) pp. 215–223.
- [36] E. Mortensen, B. Morse, W. Barrett, and J. Udupa, *Adaptive boundary detection using 'live-wire' two-dimensional dynamic programming*, in *Proc. of Computers in Cardiology* (1992) pp. 635–638.
- [37] J. K. Udupa, S. Samarasekera, and W. A. Barrett, *Boundary detection via dynamic programming*, in *Proc. of Visualization in Biomedical Computing* (1992) pp. 33–39.
- [38] D. Selle, B. Preim, A. Schenk, and H.-o. Peitgen, *Analysis of vasculature for liver surgical planning*, *Medical Imaging*, IEEE Transactions on **21**, 1344 (2002).

- [39] T. Boskamp, D. Rinck, F. Link, B. Kümmerlen, G. Stamm, and P. Mildenerger, *New vessel analysis tool for morphometric quantification and visualization of vessels in CT and MR imaging data sets I*, *Radiographics* **24**, 287 (2004).
- [40] J. Serra, *Image analysis and mathematical morphology, v. 1* (Academic Press, 1982).
- [41] H. K. Hahn, M. Georg, and H.-O. Peitgen, *Fractal aspects of three-dimensional vascular constructive optimization*, in *Fractals in biology and medicine* (Springer, 2005) pp. 55–66.
- [42] M. Kass, A. Witkin, and D. Terzopoulos, *Snakes: Active contour models*, *International journal of computer vision* **1**, 321 (1988).
- [43] I. Bankman, *Handbook of medical image processing and analysis* (Academic Press, 2008).
- [44] S. K. Setarehdan and S. Singh, *Advanced algorithmic approaches to medical image segmentation: state-of-the-art applications in cardiology, neurology, mammography and pathology* (2012).
- [45] I. Viola and M. E. Gröller, *Smart visibility in visualization*, in *Proc. of Computational Aesthetics in Graphics, Visualization and Imaging* (2005) pp. 209–216.
- [46] E. M. Lidal, H. Hauser, and I. Viola, *Design principles for cutaway visualization of geological models*, in *Proc. of the Spring Conference on Computer Graphics* (2012) pp. 47–54.
- [47] M. Burns and A. Finkelstein, *Adaptive cutaways for comprehensible rendering of polygonal scenes*, *ACM Transactions on Graphics (TOG)*, **27**, 154 (2008).
- [48] C. Kubisch, C. Tietjen, and B. Preim, *GPU-based Smart Visibility Techniques for Tumor Surgery Planning*, *International Journal of Computer Assisted Radiology and Surgery*, 667 (2010).
- [49] K. Lawonn, S. Glasser, A. Vilanova, B. Preim, and T. Isenberg, *Occlusion-free Blood Flow Animation with Wall Thickness Visualization*, *IEEE Transactions on Visualization and Computer Graphics* **22 (1)**, 728 (2015).
- [50] I. Viola, A. Kanitsar, and M. E. Groller, *Importance-driven volume rendering*, in *Proc. of IEEE Visualization* (IEEE Computer Society, 2004) pp. 139–146.
- [51] G. Borgefors, *Distance transformations in digital images*, *Computer vision, graphics, and image processing* **34**, 344 (1986).
- [52] G. Rong and T.-S. Tan, *Jump flooding in gpu with applications to voronoi diagram and distance transform*, in *Proc. of the Symposium on Interactive 3D Graphics and Games* (2006) pp. 109–116.
- [53] K. Lawonn, R. Gasteiger, and B. Preim, *Adaptive surface visualization of vessels with animated blood flow*, *Computer Graphics Forum* **33**, 16 (2014).

- [54] R. Gasteiger, M. Neugebauer, C. Kubisch, and B. Preim, *Adapted surface visualization of cerebral aneurysms with embedded blood flow information*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.
- [55] S. Glaßer, K. Lawonn, T. Hoffmann, M. Skalej, and B. Preim, *Combined visualization of wall thickness and wall shear stress for the evaluation of aneurysms*, *Visualization and Computer Graphics*, IEEE Transactions on **20**, 2506 (2014).
- [56] A. Baer, R. Gasteiger, D. Cunningham, and B. Preim, *Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow*, *Computer Graphics Forum* **30**, 811 (2011).
- [57] W. Li, M. Agrawala, and D. Salesin, *Interactive image-based exploded view diagrams*, in *Proc. of Graphics Interface* (2004) pp. 203–212.
- [58] W. Li, M. Agrawala, B. Curless, and D. Salesin, *Automated generation of interactive 3d exploded view diagrams*, in *ACM Transactions on Graphics (TOG)*, Vol. 27 (ACM, 2008) p. 101.
- [59] S. Bruckner and M. Groller, *Exploded views for volume data*, *Visualization and Computer Graphics*, IEEE Transactions on **12**, 1077 (2006).
- [60] M. Ruiz, I. Viola, I. Boada, S. Bruckner, M. Feixas, and M. Sbert, *Similarity-based exploded views*, in *Proc. of Smart Graphics* (2008) pp. 154–165.
- [61] V. Interrante, H. Fuchs, and S. Pizer, *Illustrating transparent surfaces with curvature-directed strokes*, in *Proc. of IEEE Visualization* (1996) pp. 211–218.
- [62] S. M. F. Treavett and M. Chen, *Pen-and-ink rendering in volume visualisation*, in *Proc. of IEEE Visualization* (2000) pp. 203–210.
- [63] H. Hauser, L. Mroz, G. Italo Bischi, and E. Groller, *Two-level volume rendering*, *Visualization and Computer Graphics*, IEEE Transactions on **7**, 242 (2001).
- [64] E. B. Lum and K.-L. Ma, *Hardware-accelerated parallel non-photorealistic volume rendering*, in *Proc. of Non-photorealistic Animation and Rendering* (2002) pp. 67–74.
- [65] C. Tietjen, T. Isenberg, and B. Preim, *Combining silhouettes, surface, and volume rendering for surgery education and planning*, in *Proc. of EuroVis* (2005) pp. 303–310.
- [66] K. Lawonn, *Illustrative visualization of medical data sets*, Ph.D. thesis, University of Magdeburg (2015).
- [67] C. Tietjen, B. Meyer, S. Schlechtweg, B. Preim, I. Hertel, and G. Strauss, *Enhancing slice-based visualizations of medical volume data*, in *Proc. of EuroVis*.
- [68] J. Kim, S. Eberl, and D. D. Feng, *Visualizing dual-modality rendered volumes using a dual-lookup table transfer function*, *Computing in Science and Engineering* **9**, 20 (2007).

- [69] W. Jainek, S. Born, D. Bartz, W. Straßer, and J. Fischer, *Illustrative hybrid visualization and exploration of anatomical and functional brain data*, *Computer Graphics Forum* **27**, 855 (2008).
- [70] M. Cohen, *Focus and Context for Volume Visualization*, Phd thesis, University of Leeds (2006).
- [71] S. Bruckner, E. Gröller, K. Mueller, B. Preim, and D. Silver, *Illustrative focus+context approaches in interactive volume visualization*, in *Scientific Visualization: Advanced Concepts* (2010) Chap. 10.
- [72] Y. Wan and C. Hansen, *Fast volumetric data exploration with importance-based accumulated transparency modulation*, in *Proc. of IEEE/EG International Conference on Volume Graphics* (2010) pp. 61–68.
- [73] F. de Moura Pinto and C. Freitas, *Importance-aware composition for illustrative volume rendering*, in *Proc. of Conference on Graphics, Patterns and Images (SIB-GRAPI)* (2010) pp. 134–141.
- [74] R. Bramon, M. Ruiz, A. Bardera, I. Boada, M. Feixas, and M. Sbert, *Information Theory-Based Automatic Multimodal Transfer Function Design*, *Biomedical and Health Informatics, IEEE Journal of* **17**, 870 (2013).
- [75] P. Abellán, D. Tost, S. Grau, and A. Puig, *Regions-based illustrative visualization of multimodal datasets*, *Computerized Medical Imaging and Graphics* **37**, 263 (2013).
- [76] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose, *Toolglass and magic lenses: The see-through interface*, (1993) pp. 73–80.
- [77] R. Gasteiger, M. Neugebauer, O. Beuing, and B. Preim, *The FLOWLENS: A Focus-and-Context Visualization Approach for Exploration of Blood Flow in Cerebral Aneurysms*, *Visualization and Computer Graphics, IEEE Transactions on* **17**, 2183 (2011).
- [78] C. Schulte zu Berge, M. Baust, A. Kapoor, and N. Navab, *Predicate-based focus-and-context visualization for 3d ultrasound*, *Visualization and Computer Graphics, IEEE Transactions on* **20**, 2379 (2014).
- [79] K. Lawonn, M. Luz, B. Preim, and C. Hansen, *Illustrative visualization of vascular models for static 2D representations*, in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (2015) pp. 399–406.
- [80] J. Krüger, J. Schneider, and R. Westermann, *Clearview: An interactive context preserving hotspot visualization technique*, *Visualization and Computer Graphics, IEEE Transactions on* **12**, 941 (2006).
- [81] P. Hastreiter and T. Ertl, *Integrated registration and visualization of medical image data*, in *Proc. of Computer Graphics International* (1998) pp. 78–85.

- [82] M. Burns, M. Haidacher, W. Wein, I. Viola, and E. Gröller, *Feature emphasis and contextual cutaways for multimodal medical visualization*, in *Proc. of EuroVis*, Vol. 7 (2007) pp. 275–282.
- [83] R. A. Drebin, L. Carpenter, and P. Hanrahan, *Volume rendering*, in *ACM Siggraph Computer Graphics*, Vol. 22 (1988) pp. 65–74.
- [84] M. Levoy, *Display of Surfaces from Volume Data*, *IEEE Computer Graphics Applications* **8**, 29 (1988).
- [85] K. H. Höhne and R. Bernstein, *Shading 3D-images from CT using gray-level gradients*, *Medical Imaging*, *IEEE Transactions on* **5**, 45 (1986).
- [86] K.-H. Hohne, M. Bomans, U. Tiede, and M. Riemer, *Display of multiple 3D-objects using the generalized voxel-model*, in *Proc. of Medical Imaging II* (International Society for Optics and Photonics, 1988) pp. 850–854.
- [87] W. Cai and G. Sakas, *Data Intermixing and Multi-volume Rendering*, *Computer Graphics Forum* **18**, 359 (1999).
- [88] B. Wilson, E. B. Lum, and K.-L. Ma, *Interactive Multi-volume Visualization*, in *Computational Science*, Vol. 2330 (2002) pp. 102–110.
- [89] M. Ferre, A. Puig, and D. Tost, *A framework for fusion methods and rendering techniques of multimodal volume data*, *Computer Animation and Virtual Worlds* **15**, 63 (2004).
- [90] H. Hong, J. Bae, H. Kye, and Y. Shin, *Efficient Multimodality Volume Fusion Using Graphics Hardware*, in *Computational Science*, *Lecture Notes in Computer Science*, Vol. 3516 (2005) pp. 842–845.
- [91] F. Rößler, E. Tejada, T. Fangmeier, T. Ertl, and M. Knauff, *GPU-based multi-volume rendering for the visualization of functional brain images*, in *Proc. of SimVis* (2006) pp. 305–18.
- [92] F. Rößler, R. P. Botchen, and T. Ertl, *Dynamic shader generation for GPU-based multi-volume ray casting*, *Computer Graphics and Applications* **28**, 66 (2008).
- [93] R. Brecheisen, A. V. Bartroli, B. Platel, and B. M. ter Haar Romeny, *Flexible GPU-based multi-volume ray-casting*, in *Proc. of VMV* (2008) pp. 303–312.
- [94] B. Kainz, M. Grabner, A. Bornik, S. Hauswiesner, J. Muehl, and D. Schmalstieg, *Ray casting of multiple volumetric datasets with polyhedral boundaries on manycore gpus*, *ACM Transactions on Graphics (TOG)* **28**, 152 (2009).
- [95] E. Sunden, S. Kottraval, and T. Ropinski, *Multimodal volume illumination*, *Computers & Graphics* **50**, 47 (2015).
- [96] S. Lindholm, M. Falk, E. Sundén, A. Bock, A. Ynnerman, and T. Ropinski, *Hybrid data visualization based on depth complexity histogram analysis*, *Computer Graphics Forum* **34**, 74 (2015).

- [97] N. Schubert and I. Scholl, *Comparing GPU-based multi-volume ray casting techniques*, *Computer Science-Research and Development* **26**, 39 (2011).
- [98] J. Beyer, M. Hadwiger, and H. Pfister, *State-of-the-art in gpu-based large-scale volume visualization*, *Computer Graphics Forum* **34**, 13 (2015).
- [99] S. Gupta, K. Ramesh, and E. Blasch, *Mutual information metric evaluation for PET/MRI image fusion*, in *Proc. of Aerospace and Electronics Conference* (2008) pp. 305–311.
- [100] S. Lindholm, P. Ljung, M. Hadwiger, and A. Ynnerman, *Fused multi-volume DVR using binary space partitioning*, *Computer Graphics Forum*, **28**, 847 (2009).
- [101] R. Bramon, I. Boada, A. Bardera, J. Rodriguez, M. Feixas, J. Puig, and M. Sbert, *Multimodal data fusion based on mutual information*, *Visualization and Computer Graphics*, *IEEE Transactions on* **18**, 1574 (2012).
- [102] D.-J. Kim, B. Kim, J. Lee, J. Shin, K. W. Kim, and Y.-G. Shin, *High-quality slab-based intermixing method for fusion rendering of multiple medical objects*, *Computer Methods and Programs in Biomedicine* **123**, 27 (2016).
- [103] A. P. James and B. V. Dasarathy, *Medical image fusion: a survey of the state of the art*, *Information Fusion* **19**, 4 (2014).
- [104] A. Galande and R. Patil, *The art of medical image fusion: a survey*, in *Proc. of Advances in Computing, Communications and Informatics (ICACCI), International Conference on* (2013) pp. 400–405.
- [105] A. J. Fairfield, J. Plasencia, Y. Jang, N. Theodore, N. R. Crawford, D. H. Frakes, and R. Maciejewski, *Volume curtaining: a focus+ context effect for multimodal volume visualization*, in *Proc. of SPIE Medical Imaging* (2014) pp. 903527–903527.
- [106] I. H. Manssour, S. S. Furuie, S. D. Olabarriaga, *et al.*, *Visualizing inner structures in multimodal volume data*, in *Proc. of Computer Graphics and Image Processing* (2002) pp. 51–58.
- [107] D. Weiskopf, K. Engel, and T. Ertl, *Volume clipping via per-fragment operations in texture-based volume visualization*, in *Proc. of IEEE Visualization* (2002) pp. 93–100.
- [108] Å. Birkeland, S. Bruckner, A. Brambilla, and I. Viola, *Illustrative membrane clipping*, *Computer Graphics Forum* **31**, 905 (2012).
- [109] J. Kniss, C. Hansen, M. Grenier, and T. Robinson, *Volume rendering multivariate data to visualize meteorological simulations: A case study*, in *Proc. of the EG Visualization Symposium* (2002) pp. 189–194.
- [110] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller, *Information-based Transfer Functions for Multimodal Visualization*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.

- [111] R. Maciejewski, I. Woo, W. Chen, and D. S. Ebert, *Structuring feature space: A non-parametric method for volumetric transfer function generation*, Visualization and Computer Graphics, IEEE Transactions on **15**, 1473 (2009).
- [112] R. Bramon, M. Ruiz, A. Bardera, I. Boada, M. Feixas, and M. Sbert, *An information-theoretic observation channel for volume visualization*, Computer Graphics Forum **32**, 411 (2013).
- [113] A. Joshi, D. Scheinost, K. P. Vives, D. D. Spencer, L. H. Staib, and X. Papademetris, *Novel interaction techniques for neurosurgical planning and stereotactic navigation*, Visualization and Computer Graphics, IEEE Transactions on **14**, 1587 (2008).
- [114] T. Ropinski, I. Viola, M. Biermann, H. Hauser, and K. Hinrichs, *Multimodal visualization with interactive closeups*, in *Proc. of Theory and Practice of Computer Graphics Conference* (2009) pp. 17–24.
- [115] M. Haidacher, S. Bruckner, and M. E. Gröller, *Volume analysis using multimodal surface similarity*, Visualization and Computer Graphics, IEEE Transactions on **17**, 1969 (2011).
- [116] C. D. Correa and K.-L. Ma, *Visibility histograms and visibility-driven transfer functions*, Visualization and Computer Graphics, IEEE Transactions on **17**, 192 (2011).
- [117] S. Diepenbrock, S. Hermann, M. Schäfers, M. Kuhlmann, and K. Hinrichs, *Comparative visualization of tracer uptake in in vivo small animal PET/CT imaging of the carotid arteries*, Computer Graphics Forum **32**, 241 (2013).
- [118] T. Ropinski, S. Hermann, R. Reich, M. Schäfers, and K. Hinrichs, *Multimodal vessel visualization of mouse aorta pet/ct scans*, Visualization and Computer Graphics, IEEE Transactions on **15**, 1515 (2009).
- [119] F. Enders, N. Sauber, D. Merhof, P. Hastreiter, C. Nimsky, and M. Stamminger, *Visualization of white matter tracts with wrapped streamlines*, in *Proc. of IEEE Visualization* (2005) pp. 51–58.
- [120] A. F. van Dixhoorn, J. Milles, B. van Lew, and C. P. Botha, *Braincove: A tool for voxel-wise fMRI brain connectivity visualization*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.
- [121] F. Janoos, B. Nouanesengsy, R. Machiraju, H. W. Shen, S. Sammet, M. Knopp, and I. Á. Mórocz, *Visual analysis of brain activity from fMRI data*, Computer Graphics Forum **28**, 903 (2009).
- [122] T. K. Nguyen, H. Ohlsson, A. Eklund, F. Hernell, P. Ljung, C. Forsell, M. Andersson, H. Knutsson, and A. Ynnerman, *Concurrent volume visualization of real-time fMRI*, in *Proc. of IEEE/EG International Symposium on Volume Graphics* (Eurographics-European Association for Computer Graphics, 2010) pp. 53–60.

- [123] R. Stokking, K. Zuiderveld, and M. Viergever, *Integrated volume visualization of functional image data and anatomical surfaces using normal fusion*, *Human Brain Mapping* **12**, 203 (2001).
- [124] R. Stokking, K. J. Zuiderveld, H. E. H. Pol, P. P. van Rijk, and M. A. Viergever, *Normal fusion for three-dimensional integrated visualization of SPECT and magnetic resonance brain images*, *The Journal of Nuclear Medicine* **38**, 624 (1997).
- [125] Y. Jung, J. Kim, S. Eberl, M. Fulham, and D. D. Feng, *Visibility-driven PET-CT visualisation with region of interest (roi) segmentation*, *The Visual Computer* **29**, 805 (2013).
- [126] I. Viola, K. Nylund, O. K. Øye, D. M. Ulvang, O. H. Gilja, and H. Hauser, *Illustrated ultrasound for multimodal data interpretation of liver examinations*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.
- [127] A. Hennemuth, A. Seeger, O. Friman, S. Miller, B. Klumpp, S. Oeltze, and H.-O. Peitgen, *A comprehensive approach to the analysis of contrast enhanced cardiac mr images*, *Medical Imaging, IEEE Transactions on* **27**, 1592 (2008).
- [128] S. Oeltze, A. Kuß, F. Grothues, A. Hennemuth, and B. Preim, *Integrated visualization of morphologic and perfusion data for the analysis of coronary artery disease*, in *Proc. of EuroVis* (2006) pp. 131–138.
- [129] M. Termeer, J. O. Bescós, M. Breeuwer, A. Vilanova, F. Gerritsen, and M. E. Gröller, *Covicad: comprehensive visualization of coronary artery disease*, *Visualization and Computer Graphics, IEEE Transactions on* **13**, 1632 (2007).
- [130] H. A. Kirişli, V. Gupta, R. Shahzad, I. Al Younis, A. Dharampal, R.-J. van Geuns, A. J. Scholte, M. A. de Graaf, R. M. Joemai, K. Nieman, *et al.*, *Additional diagnostic value of integrated analysis of cardiac CTA and SPECT MPI using the SMARTVis system in patients with suspected coronary artery disease*, *Journal of Nuclear Medicine* **55**, 50 (2014).
- [131] J. Blaas, C. P. Botha, C. Majoie, A. Nederveen, F. M. Vos, and F. H. Post, *Interactive visualization of fused fMRI and DTI for planning brain tumor resections*, in *Proc. of SPIE Medical Imaging* (2007) pp. 65091P–65091P.
- [132] A. Bock, N. Lang, G. Evangelista, R. Lehrke, and T. Ropinski, *Guiding deep brain stimulation interventions by fusing multimodal uncertainty regions*, in *Proc. of IEEE Pacific Visualization Symposium (PacificVis)* (IEEE, 2013) pp. 97–104.
- [133] S. Born, W. Jaineck, M. Hlawitschka, G. Scheuermann, C. Trantakis, J. Meixensberger, and D. Bartz, *Multimodal visualization of DTI and fMRI data using illustrative methods*, in *Bildverarbeitung für die Medizin* (Springer, 2009) pp. 6–10.
- [134] P. Jannin, O. J. Fleig, E. Seigneuret, C. Grova, X. Morandi, and J.-M. Scarabin, *A data fusion environment for multimodal and multi-informational neuronavigation*, *Computer Aided Surgery* **5**, 1 (2000).

- [135] T. Kin, H. Nakatomi, M. Shojima, M. Tanaka, K. Ino, H. Mori, A. Kunimatsu, H. Oyama, and N. Saito, *A new strategic neurosurgical planning tool for brainstem cavernous malformations using interactive computer graphics with multimodal fusion images: Clinical article*, *Journal of Neurosurgery* **117**, 78 (2012).
- [136] D. Merhof, M. Sonntag, F. Enders, C. Nimsky, P. Hastreiter, and G. Greiner, *Hybrid visualization for white matter tracts using triangle strips and point sprites*, *Visualization and Computer Graphics*, *IEEE Transactions on* **12**, 1181 (2006).
- [137] N. V. Navkar, N. V. Tsekos, J. R. Stafford, J. S. Weinberg, and Z. Deng, *Visualization and planning of neurosurgical interventions with straight access*, in *Proc. of Information Processing in Computer-Assisted Interventions* (2010) pp. 1–11.
- [138] A. Neubauer, L. Mroz, S. Wolfsberger, R. Wegenkittl, M. T. Forster, and K. Bühler, *Steps-an application for simulation of transsphenoidal endonasal pituitary surgery*, in *Proc. of IEEE Visualization* (2004) pp. 513–520.
- [139] C. Rieder, F. Ritter, M. Raspe, and H.-O. Peitgen, *Interactive visualization of multimodal volume data for neurosurgical tumor treatment*, *Computer Graphics Forum* **27**, 1055 (2008).
- [140] C. Rieder, M. Schwier, H. K. Hahn, and H.-O. Peitgen, *High-quality multimodal volume visualization of intracerebral pathological tissue*, in *Proc. of Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*.
- [141] L. Serra, R. A. Kockro, C. G. Guan, N. Hern, E. C. Lee, Y. H. Lee, C. Chan, and W. L. Nowinski, *Multimodal volume-based tumor neurosurgery planning in the virtual workbench*, in *Proc. of Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (1998) pp. 1007–1015.
- [142] F. Weiler, C. Rieder, C. A. David, C. Wald, and H. K. Hahn, *AVM-Explorer: Multi-Volume Visualization of Vascular Structures for Planning of Cerebral AVM Surgery*, *Eurographics 2011 - Dirk Bartz Prize*, (2011).
- [143] M. Nunes, B. Rowland, M. Schlachter, S. Ken, K. Matkovic, A. Laprie, and K. Bühler, *An integrated visual analysis system for fusing mr spectroscopy and multi-modal radiology imaging*, in *Proc. of Visual Analytics Science and Technology (VAST), IEEE Conference on* (2014) pp. 53–62.
- [144] A. Tappenbeck, B. Preim, and V. Dicken, *Distance-based transfer function design: Specification methods and applications*, in *Proc. of SimVis* (2006) pp. 259–274.
- [145] T. Isenberg, *A survey of illustrative visualization techniques for diffusion-weighted MRI tractography*, in *Visualization and Processing of Higher Order Descriptors for Multi-Valued Data* (Springer, 2015) pp. 235–256.
- [146] H. K. Hahn, *Morphological Volumetry: Theory, Concepts, and Application to Quantitative Medical Imaging*, Ph.D. thesis, University of Bremen (2005).

- [147] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich, *Viewpoint selection using viewpoint entropy*, in *Proc. of VMV*, Vol. 1 (2001) pp. 273–280.
- [148] U. D. Bordoloi and H.-W. Shen, *View selection for volume rendering*, in *Proc. of IEEE Visualization* (2005) pp. 487–494.
- [149] K. Mühler, M. Neugebauer, C. Tietjen, and B. Preim, *Viewpoint selection for intervention planning*, in *Proc. of EuroVis* (2007) pp. 267–274.
- [150] P. Kohlmann, S. Brückner, A. Kanitsar, and A. Kanitsar, *Livesync: Deformed viewing spheres for knowledge-based navigation*, *Visualization and Computer Graphics*, IEEE Transactions on **13**, 1544 (2007).
- [151] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller, *Livesync++: Enhancements of an interaction metaphor*, in *Proc. of Graphics Interface* (Canadian Information Processing Society, 2008) pp. 81–88.
- [152] R. Ranon, M. Christie, and C. Lino, *Algorithms and Techniques for Virtual Camera Control*, (Eurographics Association, Lisbon, Portugal, 2016) p. t5.
- [153] M. Burns, J. Klawe, S. Rusinkiewicz, A. Finkelstein, and D. DeCarlo, *Line drawings from volume data*, *ACM Transactions on Graphics (Proc. SIGGRAPH)* **24**, 512 (2005).
- [154] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller, *Curvature-based transfer functions for direct volume rendering: Methods and applications*, in *Visualization and Computer Graphics*, *IEEE Transactions on* (2003) pp. 67–.
- [155] F. Dong, G. J. Clapworthy, H. Lin, and M. A. Krokos, *Nonphotorealistic rendering of medical volume data*, *IEEE Comput. Graph. Appl.* **23**, 44 (2003).
- [156] X. Yuan and B. Chen, *Illustrating Surfaces in Volume*, in *Proc. of EG/IEEE VisSym* (2004) pp. 9–16.
- [157] Y. Cai and F. Dong, *Surface hatching for medical volume data*, in *Proc. of International Conference on Computer Graphics, Imaging and Visualization* (2005) pp. 232–238.
- [158] R. v. Pelt, A. Vilanova, and H. v. d. Wetering, *GPU-based Particle Systems for Illustrative Volume Rendering*, in *Proc. of IEEE/ EG Symposium on Volume and Point-Based Graphics* (2008).
- [159] S. Busking, A. Vilanova, and J. J. van Wijk, *Particle-based non-photorealistic volume visualization*, *The Visual Computer* **24**, 335 (2008).
- [160] K. Lawonn, *Illustrative Visualization of Medical Data Sets*, Thesis, University of Magdeburg (2015).

7

REGISTRATIONSHOP: AN INTERACTIVE 3D MEDICAL VOLUME REGISTRATION SYSTEM

In medical imaging, registration is used to combine images containing information from different modalities or to track treatment effects over time in individual patients. Most registration software packages do not provide an easy-to-use interface that facilitates the use of registration. 2D visualization techniques are often used for visualizing 3D datasets.

RegistrationShop was developed to improve and ease the process of volume registration by adding 3D visualizations and intuitive interaction to an existing command-line interface registration framework [3], with the goal of bringing the registration required for our surgical planning application closer to being feasible in clinical practice. It supports several basic visualizations of 3D volumetric data. Interactive rigid and non-rigid transformation tools can be used to manipulate the volumes and immediate visual feedback for all rigid transformation tools allows the user to examine the current result in real-time. In this context, we introduce 3D comparative visualization techniques, as well as a way of placing landmarks in 3D volumes. Finally, we evaluated our approach with domain experts, who underlined the potential and usefulness of RegistrationShop.

This chapter is based on the following publication: **Noeska Smit**, Berend Klein Haneveld, Marius Staring, Elmar Eisemann, Charl Botha, and Anna Vilanova, "RegistrationShop: An Interactive 3D Medical Volume Registration System", *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, pp. 145-153. 2014 [1], which was based on work performed by Berend Klein Haneveld in his Master thesis [2].

7.1. INTRODUCTION

3D medical imaging techniques like Magnetic Resonance Imaging (MRI) and X-ray Computed Tomography (CT) are increasingly used to study and diagnose diseases of patients [4]. Often, follow-up scans are acquired to track the progression of a disease over time. However, comparing such datasets is difficult because the patient will never be in the exact same pose. In consequence, volumes are registered to establish image correspondences. Another application for registration is atlas-based segmentation [5], in which a segmented dataset is registered to an unseen dataset, to transfer the segmentation. In this thesis, we also require registration of atlas data with patient-specific scans, but aim at surgical planning purposes. Registration is also useful for combining multi-modal images like Positron Emission Tomography (PET)/CT acquisitions with MRI scans. Each acquisition technique provides specific information that can aid diagnosis when properly combined. Additionally, even different scanners and acquisition techniques might lead to important variations.

Despite the broad spectrum of use cases for registration, there are currently not many free-to-use tools available that would be suitable for novice users, which contributes to the fact that registration is not yet broadly used in a clinical setting. In particular, several available tools concentrate on the alignment of 2D slices, although registering volumes is a 3D operation. A transformation that appears perfectly aligned in 2D, might be misaligned in 3D. This restriction forces the user to construct and maintain a 3D mental model of the desired alignment transformations, which can be challenging.

Our goal is to provide a free and accessible 3D-3D registration solution. Our work, called *RegistrationShop*, facilitates the registration process by adding easy-to-use interaction techniques to existing state-of-the-art registration methods. Further, we provide interactive 3D visual feedback, which makes it easier for the user to recognize spatial relationships between datasets at a glance. Having a solution specifically designed for 3D-3D registration, leads to a less complicated interface compared to multipurpose tools and provides the user with a single solution to this tedious and difficult process. Our contributions are as follows:

- **The RegistrationShop system:** Our solution offers interactive transformation tools with instant visual feedback to iteratively improve the registration result. It introduces a 3D multi-volume rendering visualization to evaluate the registration in a single view as opposed to inspecting the result in linked individual views. RegistrationShop is distributed as open source under the MIT License and made available via GitHub.
- **The case-study-based evaluation:** We evaluated RegistrationShop via a user study with domain experts according to the guidelines for case study design by Yin [6].

In the following, we will discuss related work and compare their features (Sec. 7.2) before describing RegistrationShop (Sec. 7.3) and its implementation (Sec. 7.4). A user study is presented (Sec. 7.5) before concluding (Sec. 7.6).

	RS	AD	AL	FS	EZ	MT	MR	PM	SL	VV	VR
1. 3D volume visualization	✓	✓	✓	✗	✗	✓	✓ ^a	✗ ^b	✓	✓	✓
2. Multi-volume rendering	✓	✓	✓	✗	✗	✓	✗	✗ ^b	✗	✗	✓
3. Results inspection	2/3D	2/3D	3D	2D	2D	✗	2D	2/3D ^b	2D	✗	3D
4. 3D clipping of volumes	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✗
5. Manual transformation	✓	✗	✓ ^c	✓	✗	✗	✓	✓	✓	✗	✓ ^d
6. Landmark transformation	✓	✓ ^e	✗	✗	✗	✗	✗	✓ ^f	✓	✗	✗
7. Non-rigid registration	✓	✓	✓	✓	✓	✓ ^g	✓	✓ ^h	✓	✓	✓
8. Deformation visualization	✗	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗
9. GPU-accelerated	✓ ⁱ	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗
10. Transformation history	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
11. Free software	✓	✓	✓	✗	✓	✓	✗	✗	✓	✓	✓
12. Multiplatform	✓	✓	✓	✗	✓	✓	✗	✓	✓	✓	✓

Table 7.1: Features of RegistrationShop vs. related applications (alphabetical order): RS: RegistrationShop, AD: AMIDE, AL: AMILab, FS: FusionSync, EZ: Ezys, MT: MITK, MR: Mirada Medical, PM: PMOD, SL: Slicer, VV: VolView, VR: Voreen. ^a: MIP only. ^b: In external P3D tool. ^c: Functionality not working in current version 3.2.1. ^d: Interactive transformation matrix. ^e: Non-interactive landmark transformation wizard. ^f: Limited to only one set of landmarks. ^g: Not included by default in MITK workbench. ^h: Only available for brain normalization. ⁱ: Will become available with next Elastix release [7].

7.2. RELATED WORK

Table 7.1 lists specific strengths and shortcomings of existing applications featuring registration functionality. The listed features are deemed requirements for a medical volume registration system. An ideal system would support multiple visualization options for individual datasets and the registration result (features 1 to 4). Furthermore, user-support for transformation definitions, such as landmark placement, and non-rigid registration should be available (features 5 to 7). A related visualization to inspect deformations after a non-rigid registration process should also exist (feature 8). GPU acceleration is a desirable feature to improve the speed of the registration process (feature 9) and the software should be free and multi-platform to allow for bigger uptake (feature 11 and 12).

Interestingly, free software solutions are usually multi-platform. AMIDE (AD), a Medical Image Data Examiner, uses fiducial markers (landmarks) and the maximization of mutual information for the registration process. Only rigid body transformation is supported. The software can also be used for viewing medical imaging datasets [8]. AMILab (AL) lacks landmark-based registration support and is an image processing application with visualization capabilities [9]. It is a multipurpose and particularly supports multi-volume rendering. Slicer (SL) also allows to inspect, process and visualize 3D image data [10] and offers visualization (2D slices, 3D volume rendering) and registration (rigid and deformable) capabilities. It also supported multi-volume rendering in version 3.6, but this feature has been removed since version 4.0. Ezys (EZ) supports visualizations of multiple datasets and deformation fields. It is a non-linear 3D medical image registration program [11], featuring 2D visualizations of multiple datasets and the deformation field. Non-rigid registrations in Ezys are GPU accelerated. The comparison step in VolView (VV) [12], an open-source, advanced volume visualization tool, is a bit simpler; data sets are loaded next to each, but no multi-volume render possibilities exists. It




Figure 7.1: A high-level overview diagram of performing a registration with RegistrationShop.

also offers correlation based rigid registration and multi-modality registration methods.

General frameworks, such as Voreen (VR), facilitate open-source rapid application development for the interactive visualization and analysis of multi-modal volumetric datasets [13]. Multi-volume rendering is possible through the inclusion of a template project, but the only way of interactively registering multiple volumes is by manually editing transformation matrices. MITK (MT) is used for building medical imaging applications [14] and comes with the MITK workbench, which offers various functionalities in form of modules. Yet, the only available registration method is a non-rigid registration module — available only via custom builds.

7.3. REGISTRATIONSHOP

This section describes the design of RegistrationShop by explaining the visualization capabilities and available interaction tools for manual landmark-based and automatic transformations. A high-level overview diagram of the registration process using RegistrationShop is given in Figure 7.1. First, the volumes to be registered are imported. Next, a 3D visualizations of the structures of interest provided and the user applies transformation tools to align the moving volume. Finally, the current registration results are evaluated and either accepted, leading to a transformed volume, or the transformations are adjusted again. The following subsections describe the most important stages in more detail.

7

7.3.1. VISUALIZATIONS

The visualization stage plays a central role, as the user sees in real-time how their transformations affect the registration results. Even small changes can be investigated and by responding to them, users converge, ultimately, to a solution.

COLOR MAPPING

Throughout RegistrationShop the fixed volume is consistently shown as orange and the moving volume (to be registered) is shown in blue, as can be seen in Figure 7.2. This choice makes identifying the volumes easy in all stages. These specific colors were inspired by work on comparative 2D visualization [15]. Being complementary colors, adding them results in a shade of gray whenever both are of the same intensity. If one intensity is higher, a shade of blue or orange will appear, making differences, as well as the identification of the volume causing the color, easy to spot. A second reason for this color scheme is that blue and orange can be distinguished by most colorblind people.

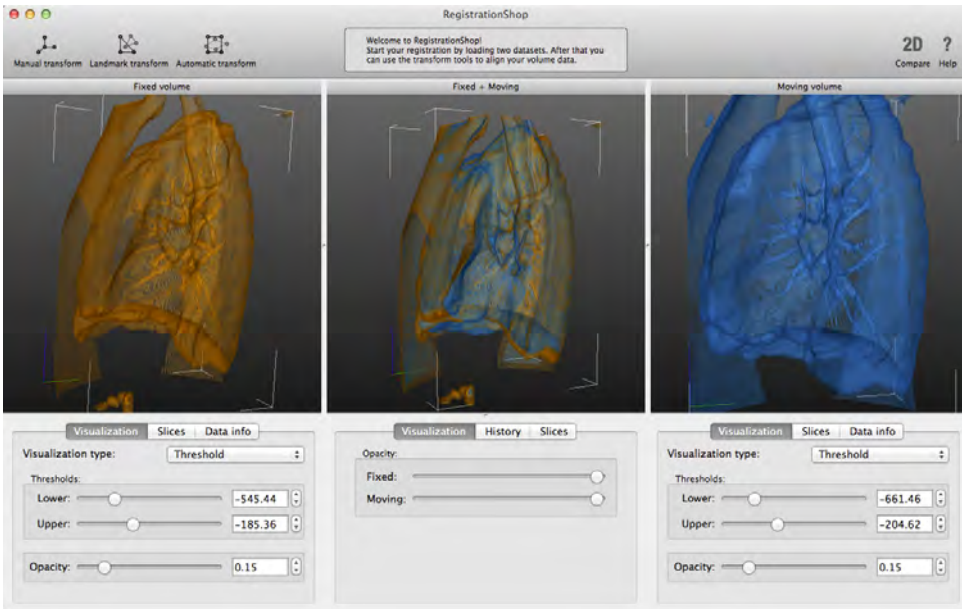


Figure 7.2: Screenshot of the main interface of RegistrationShop displaying two CT lung datasets. Left: baseline data loaded as the fixed volume with the associated visualization parameters below. Right: follow-up data loaded as the moving volume. Center: multi-volume rendering the two volumes in a single visualization to reveal the current registration result with blending parameters below. The datasets are available from the Elastix test suite.

SINGLE VOLUME VISUALIZATIONS

Two single-volume 3D-visualization options are available; the threshold visualization (THR) and Maximum Intensity Projection (MIP).

THR uses Direct Volume Rendering (DVR). To simplify controls for novice users, we offer only three sliders to define a transfer function. These modify a lower and upper threshold, as well as an opacity value. Figure 7.2 shows an example of a THR visualization; orange, resp. blue, indicates the fixed, resp. moving, volume.

The second method, MIP, displays only the voxels with highest intensity along a ray from the camera. As for THR, data can be thresholded to exclude some voxels. This option makes MIP more flexible if high values in the data are not the values of interest.

MULTI-VOLUME VISUALIZATION

The view in the center of RegistrationShop fuses the adjacent fixed and moving view (left, right, respectively). One parameter, which can be adjusted for this view, is the opacity for each of the volumes to modify the blending process. This opacity parameter setting is then multiplied by the opacity based on the transfer function. Consequently, three visualization mixes are possible: THR-THR, THR-MIP, and MIP-MIP.

THR-THR: The transfer functions of the adjacent views are applied to the volumes and, during ray tracing, voxels from both volumes are sampled alternately, starting with the fixed volume. The technique for compositing the sampled values is the same as the




Figure 7.3: Examples of visualizations in RegistrationShop. Left: a single volume rendered with MIP. Center: two volumes rendered combining MIP with THR. Right: two volumes rendered combining MIP with MIP. The two rightmost figures reveal that the two volumes are currently not aligned.

technique used for standard ray tracing. An example is shown in Figure 7.2.

MIP-THR: When MIP and THR visualizations are combined, during ray tracing, intermediate values from the two volumes are maintained: the current maximum intensity for the MIP volume and the accumulated color and opacity value for the THR volume. These values are composited until the opacity value for the THR volume becomes saturated. An example is shown in Figure 7.3.

MIP-MIP: The combination of two MIP visualizations uses a technique that extends the 2D image blending method introduced by Dzyubackyk et al [15]. Both of the volumes are separately traced with the MIP technique. The result of the fixed volume is transformed to orange while the result of the moving volume is transformed to blue. These values are then added in RGB space. When the voxels of the volumes match well, the resulting image will exist of mainly gray values. In regions with differences, blue or orange will indicate mismatches, drawing the attention of the viewer to these locations (compare Figure 7.3).

A known disadvantage of the MIP technique is that it removes depth cues and volumes might seem to overlap perfectly from a certain angle, while, from a different viewpoint, it becomes apparent that the volumes do not match well. Interaction with the visualization is therefore needed to verify the result. In the future, it would be interesting to evaluate the use of MIDA instead of MIP in this comparison setting [16]. This fusion technique is mainly applicable to datasets of the same modality, as it relies on data similarity. Further, as maximum values along the ray are output, differences in lower value ranges are harder to spot.

2D COMPARISON TOOL

Domain experts working with image volumes are used to exploring volumetric data via 2D slices through the volume. Providing only a 3D visualization would ignore their previous experience and expertise. Hence, we also include a 2D comparison tool to verify




Figure 7.4: 2D comparison tool with the lung datasets. Left: the fixed volume. Center: the current registration result. Right: the moving volume.

results acquired with the 3D registration.

A secondary reason for 2D visualization is to provide means for a detailed inspection, as it provides a direct view of the data. In a 3D visualization, the data is filtered through color and opacity assignment and details of the original datasets might no longer be visible in 3D. The layout of the views is the same as in the 3D viewer (Figure 7.4).

7.3.2. TRANSFORMATION TOOLS

To transform the moving volume, we included a manual, a landmark-based, and an automatic transformation tool. Whenever a tool is applied, the modification is kept in a history tab to allow the user to compare, or undo previous operations. This history tab stores all steps performed in the registration pipeline for the current project and can be used to visually inspect the results of previous steps.

MANUAL TRANSFORMATION TOOL

The manual transformation is used to coarsely align data sets and can serve as an initialization for subsequent registration steps. Here, the user interacts with a box widget, making interactive rigid transformations, rotations, and scaling of the moving volume possible.

LANDMARK-BASED TRANSFORMATIONS

Landmarks can be used to register mutual correspondence points in both volumes. A least-squares fit is then used to find a best match between corresponding landmark positions. The computation is interactive and builds upon the unit-quaternion algorithm by Horn [17], which offers the choice between rigid, similarity-based, or affine transformations.

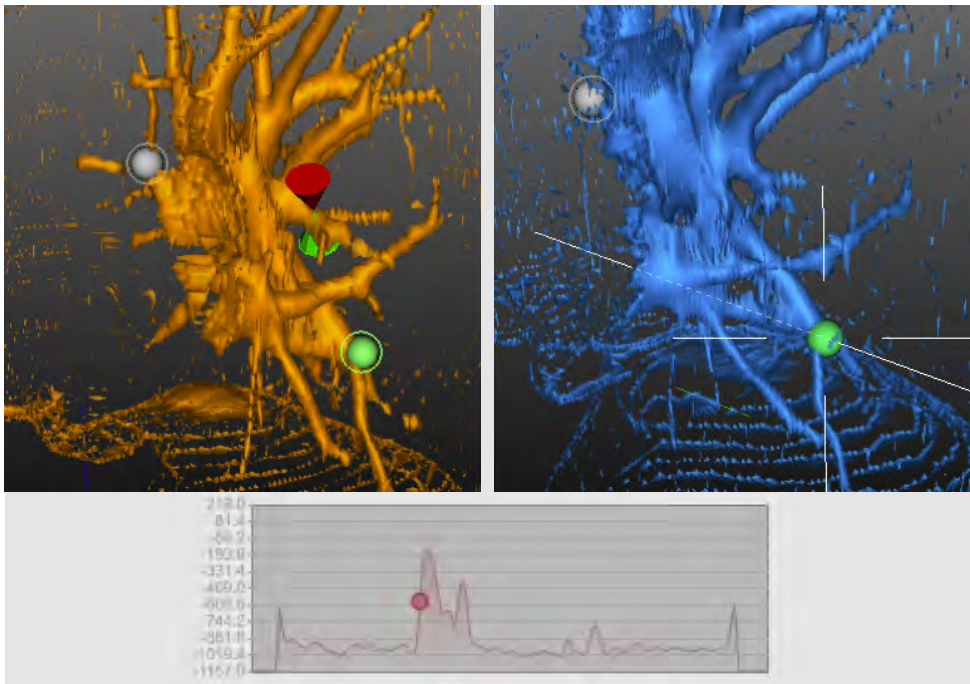


Figure 7.5: Landmark transformation tools. Left: surface picker (red and green cone) with an active (green) landmark and an inactive (gray) landmark. Right: two-step picker after shooting the picking ray. Bottom: associated value-profile along the ray for the two-step picker.

When specifying a landmark, the location is marked with a sphere and an additional circle around it. The sphere is drawn in the context of the volume, while the circle is drawn as an overlay, hence, keeping it always visible. Landmarks are specified in pairs and whenever a new landmark pair is indicated, the transformation is updated and its effect instantly shown in the multi-volume view.

Placing landmarks in a volume is not a trivial procedure: there are three dimensions in a volume, which are displayed on a two dimensional screen. Further, to ensure a wide use of our software, we restricted ourselves to mouse interaction, but, this choice, leads to 2D coordinates in screen space. Due to possible voxel transparency, it is impossible to know exactly, which voxel a user chooses based on a 2D DVR image. To remedy this problem, we offer two solutions: the surface picker and the two-step picker.

The **surface picker** traces a ray through the volume, originating from the mouse position, to determine the first voxel with an opacity value that exceeds a pre-defined threshold value. At this location, two cones, aligned with the normal (gradient) of this location, giving the impression that one cone appears to stick out of the surface, while the other seems to disappear behind it. The cones are rendered in the same context as the volume, giving the user a sense of the chosen voxel's depth (Figure 7.5).

The **two-step picker** addresses cases in which the surface picker potentially fails; e.g., if the opacity of the intended location is below the threshold, or a higher opacity re-

gion occludes it. The principle of the two-step picker relates to stereo reconstruction. In the first step, the user selects a location on screen to which the intended voxel projects, hereby, effectively restricting its location to the corresponding ray through the volume. In the second step, a different view is chosen and, again, the projected location of the intended voxel is clicked. The algorithm will then choose the point on the first line, closest to the second.

A sphere is used to indicate the selected location along the first line. Additionally, to support the user in the selection process, the depth profile along the first ray and the current selection is shown in the volume view (Figure 7.5). The intensity profile embeds the location in a context and can also be used directly to drive the selection.

AUTOMATIC TRANSFORMATION TOOL

The automatic transformation tool in RegistrationShop is a wrapper around the command-line version of Elastix, an open-source library with which rigid and non-rigid registration algorithms can be executed [3]. Elastix, while powerful, can be difficult to use for non-technical users due to the command-line interface and complexity of the parameter file choices. RegistrationShop provides default parameter files to start the registration process, which hides the technicalities from the user. Nonetheless, additional custom parameter files can be loaded or created. If an initial rigid transformation has been produced in RegistrationShop, it will be passed to Elastix to initiate the optimization. When Elastix finishes the registration, the created dataset is displayed automatically in RegistrationShop. The registration result can then be evaluated and corrected if needed.

7.4. IMPLEMENTATION

RegistrationShop uses several languages and frameworks. It is written in Python using the Visualization Toolkit (VTK) and PySide. VTK was chosen because of its volume-rendering and loading capabilities. PySide was used for building the user interface because it allows cross-platform development targeting all major platforms.

VTK does not support rendering of multiple volumes by default and we used a custom extension created by K. Krissian and C. Falcón-Torres [18]. The code has been adjusted and extended for all the possible combinations of DVR and MIP visualizations.

As indicated, Elastix is used for the non-rigid registration. It is based on The Insight Segmentation and Registration Toolkit (ITK), so it shares most of its functionality. Elastix is distributed under the simplified BSD-license. Elastix was the framework of choice as it has a permissive license.

Alternatively, we could have used ITK, which is an open-source, cross-platform segmentation solution [19], or DROP, which uses Markov Random Fields (MRF) to optimize for cost functions [20]. Compared to the gradient-descent method in Elastix, it performs well and is a promising alternative, but licensing allows only for research and non-commercial purposes.

7.5. EVALUATION

We relied on an evaluation by experts to evaluate if RegistrationShop reaches its goals of making the registration process easier and is able to produce the desired/intended

	MR	RA	I1	I2
Registration frequency	low	avg.	avg.	high
3D experience	low	avg.	high	high
Elastix experience	low	low	high	high
Technical knowledge	avg.	low	high	high

Table 7.2: Domain experts' prior experience (avg.: average).

result. We rely on two case studies with medical domain experts, who compared RegistrationShop to their current registration workflow.

7.5.1. CASE STUDY DESIGN

Our case-study design follows the guidelines stated by Yin for case study research [6]. The main evaluation question has been formulated as follows: *“Does RegistrationShop make performing registrations easier and help in producing the desired/intended result by providing 3D visualizations of the datasets in combination with simple interactive tools with instant user feedback?”*.

Participants were a medical radiation therapist (MR), a radiologist (RA) and two image-processing experts (I1, I2), who are all familiar with medical image registration. The first two experts are more familiar with the medical applications of registration, while the latter two are experts in registration itself. These experts have a reasonable coverage of users who might be interested in using RegistrationShop. Before starting, the domain experts were asked to answer several questions about their prior knowledge and experiences with registration, which allowed us to investigate the level of technical expertise in dealing with 3D visualization and interaction. Based on these questions, levels of prior experience were defined for all four domain experts (Table 7.2). MR does not perform registration in his current work or possess any prior knowledge on 3D visualization and interaction. While having some technical knowledge, MR is not familiar with Elastix. RA performs registrations automatically, using a radiologist workstation, and was also not familiar with Elastix. RA does use 3D visualization techniques occasionally, but is not familiar with technical terms. The two image processing experts I1 and I2 both are familiar with 3D visualizations, Elastix and technical terms. The only difference in prior experience between them is that I2 performs registrations more frequently.

Each evaluation session took 40 minutes and consisted of three phases. In the first phase, we showed the expert a live demonstration of the RegistrationShop system by registering test volumes containing misaligned 3D cubes. In the second phase, the expert was asked to perform a registration on a dataset from clinical practice. Two of the experts were asked to register full-body MRI scans of the same patient at different time-points, while the last two experts were asked to align baseline and follow-up CT-scans of a lung. These datasets were chosen, because they were closest to the type of studies currently performed by the users. We encouraged the experts to think out loud during this process and recorded the audio and screen. During the third phase, we asked the expert to express their level of agreement to the statements on a five-point Likert scale and to provide vocal commentary to clarify their responses.

Statement	MR	RA	I1	I2
1. Navigation is quick and easy.	4	2	2	3
2. Navigation controls are clear and easy to get used to.	4	4	2	3
3. Inspection of the data is simplified by the visualization capabilities.	4	4	2	4
4. The single-volume 3D visualizations help display internal structures of the datasets.	5	1	2	5
5. It is easy to make a meaningful 3D visualization with the available options.	5	1	3	2
6. The interactive box allows for quick and easy dataset alignment.	4	1	4	4
7. The interactive transformation matrix is a useful optional tool.	N/A	N/A	2	1
8. A manual transform made with RegistrationShop is a good initialization.	5	4	4	5
9. The surface picker allows for quick landmark placement in the intended locations.	3	3	2	3.5
10. The 3D visualizations enable placement of the landmarks in the intended locations.	4	4	4	3
11. Real-time updates in the combined view demonstrate landmark placement effects.	4	4	5	5
12. The two-step picker works where the surface picker would fail.	4	N/A	5	5
13. Making an initial transform for Elastix is easy with RegistrationShop.	N/A	N/A	N/A	2.5
14. Loading Elastix parameter files enables re-use of existing work.	N/A	N/A	N/A	5
15. The standard Elastix parameters are a good starting point for deformable registration.	N/A	N/A	4	5
16. Adjusting parameter values for Elastix is easier in RegistrationShop.	N/A	N/A	4	1
17. The 3D multi-volume visualization allows me to evaluate registration results.	5	4	2	5
18. The 2D comparison tool is indispensable to evaluate registration results.	5	4	4	5
19. The need for 2D visualizations is reduced by the 3D visualizations.	2	1	1	1
20. The history tab allows me to see if the registration improved with a registration step.	3	4	4	5
21. RegistrationShop provides me with all the tools I need to perform a registration.	2	2	2	5
22. I do not need additional technical knowledge to operate RegistrationShop	4	4	N/A	N/A
23. I see myself using RegistrationShop in the future.	3	3	4	5

Table 7.3: Domain expert level of agreement with statements based on a Likert scale. 1: Strongly disagree, 2: Disagree, 3: Neither agree nor disagree, 4: Agree, 5: Strongly agree, N/A: Not available. The x.5 values are given where the user did not feel the five-point Likert scale could reflect his opinion accurately.

7.5.2. CASE STUDY RESULTS

To examine our hypothesis, we decomposed the main evaluation question into several related propositions that the domain experts were asked to respond to after using RegistrationShop (Table 7.3).

3D VISUALIZATIONS

The 3D visualizations of the datasets were regarded as interesting and succeeded in providing a sense of the overall anatomical structure of the data and the spatial relations between the two volumes (question 17). Three users indicated that the inspection of the data is simplified by the visualization capabilities of RegistrationShop, although one image processing expert found 2D inspection easier (question 3) and also prefers to evaluate registration results in 2D (question 17). The combined viewer was successful in showing the spatial relation between the datasets, especially, because of the instant user feedback when working with the transformation tools. Because some domain experts were not familiar with interaction in 3D, they indicated that they would not consider navigation easy initially and expect a learning curve (questions 1, 2). Making a meaningful 3D visualization was also not considered easy by two experts due to a learning curve, but also lacking features, such as a reset functionality or optional parameter and camera linking (question 5). Two experts strongly agree that the single-volume 3D visualizations help to display internal structures of the data, but one expert pointed out that for the full-body MRI volumes, internal structures cannot be shown easily using volume rendering techniques. One expert indicated he could not visualize internal details of the lung dataset easily (question 4). While the 3D visualization is suitable for displaying anatomical structures, it is less suitable for detailed voxel-by-voxel inspection and comparison of the datasets. All experts agree that the 2D visualization included in Reg-

istrationShop is required for detailed inspection; consequently, the inclusion of the 2D comparison tool was deemed necessary (questions 18, 19). Incorporation of the 2D comparison tool directly in the 3D views was suggested as a potential improvement. Three experts suggested other improvements when asked if RegistrationShop provides them with all the tools they need for registration (question 21). Additional features included deformation-field visualization for non-rigid registrations, measurements and statistical reports of the datasets, and interactive segmentation.

TRANSFORMATION TOOLS

The transformation tools were mostly successful in helping the experts to create registrations, though one expert experienced problems when using the manual box widget for rotations (question 6). All experts agreed or strongly agreed that the manual transform tool is able to create a good initialization for subsequent registration steps (question 8). Landmark placement allowed the users to indicate the intended locations, though one expert would have liked to place landmarks in the 2D view as well (question 10). Especially landmark placement using the two-step picker was deemed useful; its value-profile visualization was considered helpful to determine the correct placement in 3D (question 12). However, the visualization of the surface picker and the landmarks should not obstruct the indicated location (question 9). The real-time transformation updates were appreciated by all and allowed the experts to understand landmark placement effects (question 11). The history tab allowed the experts to assess if a registration step had improved the registration result, and the combination with the real-time interaction updates was considered helpful by the experts (question 20).

7

ELASTIX

The image processing experts familiar with Elastix (I1 and I2) appreciated the interaction between RegistrationShop and Elastix and the choice of default and custom parameter files (questions 14, 15). The parameter editing was considered easier in RegistrationShop by one expert (question 16). The other expert has such extensive experience with Elastix, he found that it did not make working with Elastix easier for him personally. He did comment that RegistrationShop would be valuable to inexperienced users or users that infrequently do registrations. The benefit of using RegistrationShop over Elastix is that the results of an Elastix run can be visually inspected immediately after completion using RegistrationShop.

CASE STUDY CONCLUSION

Overall, RegistrationShop was well received by the users. The interface was praised by one of the non-technical experts for its simplicity and clear layout. The domain experts with limited technical experience felt that using RegistrationShop does not require additional technical knowledge, while the image processing experts could not answer as they already possess the corresponding technical knowledge (question 22). MR remarked that when previously using other registration software, it took four days to get acquainted with the tool, while RegistrationShop allows a user to get results quickly and intuitively. The image processing experts both could see themselves using RegistrationShop in the future, while the clinicians neither agreed nor disagreed with this state-

ment, because their current work does not require them to manually register datasets (question 23).

The user study indicates that in regard to our main evaluation question, several preliminary results can be reported. The 3D visualization is appreciated for a global overview in addition to the 2D views for detailed analysis. Furthermore, the simple interactive transformation tools got positive feedback from all four users. Finally, the instant visual user feedback was found helpful to evaluate the result of transformation interactions on the registration outcome. The non-technical experts found that they do not need additional technical knowledge to operate RegistrationShop. The image processing experts consider our work a promising volume registration tool with the potential to make registration more accessible to clinicians and researchers in the future.

7.6. CONCLUSION AND FUTURE WORK

We presented RegistrationShop, a 3D-3D volume registration system making use of 3D visualizations for the moving and fixed volumes alongside the current registration result. It provides a set of interactive tools coupled with instant visual feedback, which helps the user to iteratively register the datasets, as underlined by the appreciation of experts that evaluated our system. RegistrationShop has potential to bring complex registrations closer to clinical practice, for instance to register atlas data to patient-specific scans for pre-operative planning purposes, as we describe in Chapter 8.

Despite the overall positive outcome of our study, there is still room for future work, such as the integration of 2D views in the 3D volume, region-of-interest tools, such as special lenses or segmentation tools to visualize mutual information about the registration quality in certain regions. Similarly, other compositing operations [21] or transfer function definition could be used [22, 23]. We also envision using saliency and similarity measures to further improve the landmark placing process [24]. Furthermore, a larger follow-up user study is needed to further examine the clinical potential of RegistrationShop.

Visualizing the non-rigid deformation field might be helpful and lead to a more-likely acceptance in a clinical setting. Several users indicated this potential of our work, especially, as it makes the complicated task of registration less complex and accessible to inexperienced users.

RegistrationShop is freely available as an open source prototype and the source code of RegistrationShop can be viewed and downloaded on GitHub.

REFERENCES

- [1] N. N. Smit, B. K. Haneveld, M. Staring, E. Eisemann, C. P. Botha, and A. Vilanova, *RegistrationShop: An interactive 3d medical volume registration system*, *EG Workshop on Visual Computing in Biology and Medicine (VCBM)*, (2014).
- [2] B. Klein-Haneveld, *RegistrationShop: An interactive 3D medical volume registration tool*, Master's thesis, Delft University of Technology, the Netherlands (2014).
- [3] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, *Elastix: a toolbox*

- for intensity-based medical image registration, *Medical Imaging*, IEEE Transactions on **29**, 196 (2010).
- [4] R. Smith-Bindman, D. L. Miglioretti, E. Johnson, C. Lee, H. S. Feigelson, M. Flynn, R. T. Greenlee, R. L. Kruger, M. C. Hornbrook, D. Roblin, *et al.*, *Use of diagnostic imaging studies and associated radiation exposure for patients enrolled in large integrated health care systems, 1996-2010*, *Jama* **307**, 2400 (2012).
- [5] T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer Jr, *Quo vadis, atlas-based segmentation?* in *Handbook of Biomedical Image Analysis* (Springer, 2005) pp. 435–486.
- [6] R. K. Yin, *Case Study Research: Design and Methods* (Thousand Oaks, CA: Sage, 1994) pp. 1–53.
- [7] D. P. Shamonin, E. E. Bron, B. P. Lelieveldt, M. Smits, S. Klein, M. Staring, A. D. N. Initiative, *et al.*, *Fast parallel image registration on cpu and gpu for diagnostic classification of alzheimer's disease*, *Frontiers in neuroinformatics* **7** (2013).
- [8] A. M. Loening, S. S. Gambhir, *et al.*, *Amide: a free software tool for multimodality medical image analysis*, *Molecular imaging* **2**, 131 (2003).
- [9] K. Krissian, F. Santana, D. Santana-Cedr s, *et al.*, *Amilab software: medical image analysis, processing and visualization*, *Medicine Meets Virtual Reality–NextMed (MMVR)*. CA, USA: IOS Press **173**, 233 (2012).
- [10] S. Pieper, M. Halle, and R. Kikinis, *3d slicer*, in *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on* (IEEE, 2004) pp. 632–635.
- [11] A. Gruslys, S. Sawiak, and R. Ansorge, *3000 non-rigid medical image registrations overnight on a single pc*, in *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE* (IEEE, 2011) pp. 3073–3080.
- [12] *Kitware, inc. - volview*. URL: <http://www.kitware.com/opensource/volview.html> (1999), accessed: 2014-05-25.
- [13] J. Meyer-Spradow, T. Ropinski, J. Mensmann, and K. Hinrichs, *Voreen: A rapid-prototyping environment for ray-casting-based volume visualizations*, *Computer Graphics and Applications*, IEEE **29**, 6 (2009).
- [14] I. Wolf, M. Vetter, I. Wegner, M. Nolden, T. Bottger, M. Hastenteufel, M. Schobinger, T. Kunert, and H.-P. Meinzer, *The medical imaging interaction toolkit (mitk): a toolkit facilitating the creation of interactive software by extending vtk and itk*, in *Medical Imaging 2004* (International Society for Optics and Photonics, 2004) pp. 16–27.
- [15] O. Dzyubachyk, J. Blaas, C. P. Botha, M. Staring, M. Reijniere, J. L. Bloem, R. J. van der Geest, and B. P. Lelieveldt, *Comparative exploration of whole-body mr through locally rigid transforms*, *International journal of computer assisted radiology and surgery* **8**, 635 (2013).

- [16] S. Bruckner and M. E. Gröller, *Instant volume visualization using maximum intensity difference accumulation*, Computer Graphics Forum **28**, 775–782 (2009).
- [17] B. K. Horn, *Closed-form solution of absolute orientation using unit quaternions*, JOSAA **4**, 629 (1987).
- [18] K. Krissian and C. Falcón-Torres, *GPU volume ray casting of two volumes within VTK*, URL: <http://hdl.handle.net/10380/3350>, accessed: 2014-05-25.
- [19] T. S. Yoo, M. J. Ackerman, W. E. Lorensen, W. Schroeder, V. Chalana, S. Aylward, D. Metaxas, and R. Whitaker, *Engineering and algorithm design for an image processing api: a technical report on itk-the insight toolkit*, Studies in health technology and informatics , 586 (2002).
- [20] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, *Dense image registration through mrfs and efficient linear programming*, Medical Image Analysis **12**, 731 (2008).
- [21] N. Schubert and I. Scholl, *Comparing gpu-based multi-volume ray casting techniques*, Computer Science-Research and Development **26**, 39 (2011).
- [22] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller, *Information-based transfer functions for multimodal visualization*, in *Proceedings of the First Eurographics Conference on Visual Computing for Biomedicine*, EG VCBM'08 (Eurographics Association, 2008) pp. 101–108.
- [23] R. Bramon, I. Boada, A. Bardera, J. Rodriguez, M. Feixas, J. Puig, and M. Sbert, *Multimodal data fusion based on mutual information*, IEEE Transactions on Visualization and Computer Graphics **18**, 1574 (2012).
- [24] A. Azar, C. Xu, X. Pennec, and N. Ayache, *An interactive hybrid non-rigid registration framework for 3d medical images*, in *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on* (IEEE, 2006) pp. 824–827.

8

PELVIS: ATLAS-BASED SURGICAL PLANNING FOR ONCOLOGICAL PELVIC SURGERY

As stated in Chapter 1, due to the intricate relationship between the pelvic organs and vital structures, such as vessels and nerves, pelvic anatomy is often considered to be complex to comprehend. In oncological pelvic surgery, a trade-off has to be made between complete tumor resection and preserving function by preventing damage to the nerves. Since these autonomic nerves are not visible in pre-operative MRI scans or during surgery, avoiding nerve damage during such a surgical procedure becomes challenging.

In this chapter, we present visualization methods to represent context, target, and risk structures for surgical planning, relying on the VSP (see Chapter 3) and RegistrationShop (see Chapter 7). We employ distance-based and occlusion management techniques in an atlas-based surgical planning tool for oncological pelvic surgery. Patient-specific pre-operative MRI scans are registered to an atlas model that includes nerve information. Through several interactive linked views, the spatial relationships and distances between the organs, tumor and risk zones are visualized to improve understanding, while avoiding occlusion. In this way, the surgeon can examine surgically relevant structures and plan the procedure before going into the operating theater, thus raising awareness of the autonomic nerve zone regions and potentially reducing post-operative complications. Furthermore, we present the results of a domain expert evaluation with surgical oncologists that demonstrates the advantages of our approach.

This chapter is based on the following publication: **Noeska Smit**, Kai Lawonn, Annelot Kraima, Marco DeRuiter, Hessam Sokooti, Stefan Bruckner, Elmar Eisemann, and Anna Vilanova, "PelVis: Atlas-based Surgical Planning for Oncological Pelvic Surgery", accepted to appear, *IEEE Transactions on Visualization and Computer Graphics (Proceedings of Scientific Visualization 2016)*, 2017 [1].

8.1. INTRODUCTION

Oncological surgical procedures in the pelvic area are often performed to treat rectal, cervical and prostatic cancer. While removal of the tumor and a tumor-free circumferential margin is the main priority, the pelvic autonomic nerves should not be damaged in order to preserve function post-operatively, as described in Chapter 1.

The Virtual Surgical Pelvis (VSP), described in Chapter 3, is an atlas model containing surgical risk zones in which the autonomic nerves reside are indicated. By registering this atlas to patient-specific MRI scans, we are able to visualize the autonomic nerve zones in the context of an individual patient. After the registration process, we can construct a patient-specific 3D model of the surgically relevant structures that can be used in surgical planning and education for oncological surgeons in order to increase awareness of the regions at risk for nerve damage. Due to the anatomical complexity of the pelvis, many structures occlude a clear view on the structure that is to be resected, while they are still necessary as anatomical context. Furthermore, the proximity of the autonomic nerves to the target organ, as well as the distance of the tumor to the organ border are important aspects in planning the procedure. In our approach, we present several visualization methods to visualize target, risk and context structures for surgical planning. For this, we employ occlusion-management and distance-based visualizations to satisfy the surgical planning requirements. We demonstrate the utility of these methods in an interactive application designed for surgical planning of the total mesorectal excision procedure. In several linked 2D and 3D views, the surgeon is able to explore the spatial relations and distances between surgically relevant structures based on MRI and registered atlas information.

With this, our contributions are the following:

- We present visualization methods tailored to evaluate distances and to avoid occlusion in surgical-planning visualizations for oncologic procedures.
- We provide an interactive application based on these methods for surgical planning of the total mesorectal excision procedure.
- We present an evaluation with five domain experts and ten non-experts in which we demonstrate the utility of our approach.

The rest of the paper is structured as follows. We discuss related work (Section 8.2), followed by related medical background information (Section 8.3). We then describe a requirements analysis describing the results of our requirements elicitation from domain experts (Section 8.4). We present the proposed visualization methods to meet these requirements (Section 8.5), their implementation details (Section 8.6) and the obtained results (Section 8.7). The setup and results of our expert review with five surgeons and ten non-experts are described (Section 8.8), followed by a discussion (Section 8.9), before we conclude the paper with an outlook on potential future work (Section 8.10).

8.2. RELATED WORK

In this section, we describe work related to our method, such as visualization methods for surgical planning and general visualization techniques applicable in our context.

Much work has been done on applying visualization techniques to improve surgical planning, mainly in orthopedic, hepatic, facial and neurosurgery. An overview of multi-modal surgical planning applications is given in Chapter 6. A full overview of all surgical planning methods is out of the scope of this paper, but additional information can be found in the book by Preim and Botha [2].

Estimating distances between organs and relevant structures is an important task during surgical planning. Some previous works in surgical planning have used color mapping and isolines to encode distances. Marai et al. proposed a method for estimating joint contact areas in which they combine a colormap with isolines to visualize distances on the bone surfaces [3]. Dick et al. presented several methods for distance visualization in interactive 3D orthopedic implant planning [4]. They propose two methods to visualize the distance of the implant to the bone besides color mapping: slice-based distance visualization and glyph-based distance visualization. Both methods only visualize a single anatomical structure. Süßmuth et al. presented a color-encoded distance visualization of cranial nerve-vessel contacts for diagnosis and treatment of neurovascular compression (NVC) syndromes [5]. They employ a slice-based colored isoline visualization and also color-code the distances in 3D surface models. While the vessels were semi-automatically segmented, the nerves were added by manual segmentation. Rieder et al. presented a method where color and isolines were used to encode heat distribution successfully in radiofrequency ablation [6]. Using their GPU-based heat distribution calculation method, they visualize heating and cooling zones in 2D slices as well as 3D volume rendering. Since we have numerous structures to visualize and due to the complex pelvic anatomy, we employ isolines and color-coding on the resection target surface only in order to prevent further visual overload and clutter.

Regarding the visualization of distances to aid risk assessment, Hansen et al. employ discrete color coding to visualize robust safety margins for oncologic liver surgery [7]. Based on a distance calculation, they visualize which vascular structures would be affected by choosing a certain safety margin around the tumor in the planned resection. Besides the target and risk areas, no contextual anatomical structures are shown. Marshall et al. proposed a Proximity Map Projection (PMP) for interactive visualisation for surface proximity monitoring [8]. They apply their technique in the context of real-time MRI for the guidance of thermal therapy and show a mapping of the 3D distances to a 2D representation. Krüger et al. proposed an interactive visualization for oncological procedure planning in head and neck surgery [9]. They employ several visualization methods such as cutaway views, silhouettes and color-coded distances in an interactive surface-based application. Their work is the most closely related to ours, since it also involves complex anatomy and pathology with many structures in close spatial relations, though the excision goal is a tumor and not a complete organ.

To improve the reliability of our visualizations, we visualize an estimation of the uncertainty due to the registration. Djurcilov et al. presented methods for visualizing uncertainty in volume rendering [10]. They employed holes, depth-shaded holes, noise and texture to visualize uncertainty together with scalar volumetric data. Grigoryan and Rheingans employed point-based primitives to show surface uncertainty [11]. Their approach displaces points on surfaces where they are uncertain, resulting in a fuzzy appearance. Since we need to preserve the shape and view on the organ, while visu-

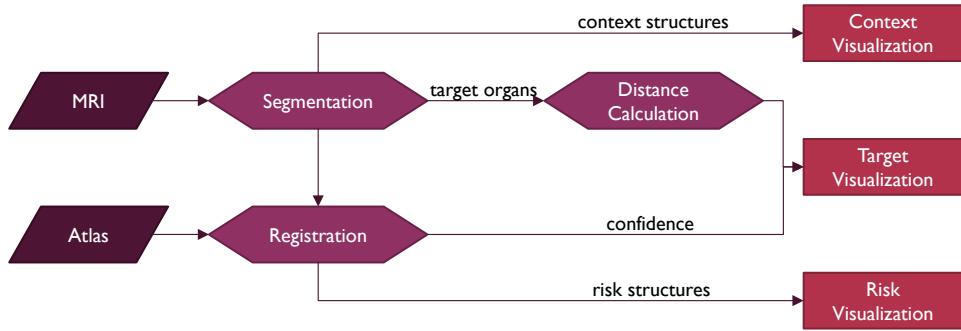


Figure 8.1: The PelVis pipeline.

alizing registration confidence, our approach employs screenspace-oriented methods. Whitaker et al. presented contour boxplots to visualize uncertainty in simulation ensembles [12]. Our contour confidence visualization was inspired by their visualization, though we do not base it on an ensemble of contours. While not aimed at surgical planning, Termeer et al. encoded myocardial perfusion in a comprehensive visualization of the coronary anatomy [13]. The blood supply was shown using a combination of colormapping and isolines. In under-perfused regions, a striped pattern was blended in with faded edges to indicate uncertainty. This pattern approach inspired our grid and halftone confidence visualizations, though ours are generated in screen space.

To the best of our knowledge, there is no related literature on surgical planning for pelvic oncologic procedures. Compared to other surgical planning applications, our approach adds information that is not visible in the pre-operative medical scans by involving atlas information. Furthermore, we propose several visual encodings that communicate distance and confidence information simultaneously and evaluate these qualitatively with domain experts and quantitatively with non-experts.

8.3. MEDICAL BACKGROUND

For treatment of rectal malignancies, the gold standard surgical procedure is the total mesorectal excision (TME) [14]. Though other procedures exist to treat rectal cancer, the TME has lower recurrence rates [15]. In the TME, the rectum is resected, including removal of the mesorectum, a layer of fatty tissue surrounding the rectum. The anal sphincter can be spared and connected to the remaining bowel through anastomosis in the low anterior variant of the procedure. In case the tumor is situated too close to the anal sphincter (within 5 cm), an abdominoperineal resection (APR) needs to be performed, resulting in the removal of the anal sphincter and a permanent stoma.

The quality of a TME procedure is evaluated using several factors influencing the recurrence and survival rates. After the procedure, the circumferential resection margin (CRM), the distance of the resection border to the tumor, is checked by a pathologist [16]. The influence of the circumferential resection margin on local recurrence, distant metastasis and survival rates was studied by Wibe et al. [17]. They concluded that a decreasing circumferential margin is associated with an exponential increase in the local recurrence

rates, metastasis and death. A margin of ≤ 1 mm is considered a negative prognostic factor for local recurrence [18]. Furthermore, the excised mesorectum specimen should be complete. If it is incomplete, there is an increased rate for local and distant recurrence [19].

While the recurrence and survival rates are crucial, preserving function is additionally an important factor in surgery. The TME procedure can result in urinal and fecal incontinence as well as sexual dysfunction if the autonomic nerves are damaged during the procedure.

The TME procedure can be performed in open surgery, in which the abdomen is cut open, or using a minimally invasive, or laparoscopic approach [20]. In laparoscopic surgery, also known as keyhole surgery, small incisions are made and the surgery is performed using a laparoscope, a flexible camera with a light source. Additionally, laparoscopic TME can also be performed aided by a robotic surgery [21]. While traditionally the TME procedure is executed in a cranio-caudal direction, there is a relatively new trend to approach the resection through the anus in the transanal minimally invasive surgery for total mesorectal excision (TAMIS-TME) [22].

Surgical planning for the TME procedure currently only consists of the analysis of a pre-operative MRI scan, which can be used to determine the type of surgery and in cancer staging [23, 24]. The complex anatomy of the pelvic area makes it challenging for the surgeon to determine the exact anatomical relations between structures in 3D during the operation, given the MRI scans. Often anatomy textbooks are additionally also used to complement the MRI information.

8.4. REQUIREMENTS ANALYSIS

In order to elicit the requirements for our surgical planning tool, we consulted with two anatomists and two surgeons. From these discussions, we formulated a list of surgically relevant structures and derived features. The following structures are surgically relevant during the planning phase of a TME procedure: (1) the pelvic bones, (2) the vagina/uterus or prostate, (3) the bladder, (4) the mesorectum, (5) the tumor, and (6) the autonomic nerves.

The mesorectum is important since it is the structure that is to be resected during the TME procedure. The organs closely arranged around the mesorectum, i.e., the internal genitalia and bladder serve as spatial context and are recognizable during surgery. The pelvic bones also serve as spatial context and additionally feature anatomical landmarks that can be identified even externally, such as the spina iliaca anterior superior. The tumor and its relation to the mesorectum is important to determine the required type of surgery. The autonomic nerves are structures that need to be avoided during surgery, since damaging them negatively affects functional outcome.

Among these structures, a distinction can be made between structures that serve as an anatomical context to aid spatial orientation (structures 1 to 3), structures that need to be resected (structures 4 and 5) and structures that pose a risk for poor functional surgical outcome (structure 6). Therefore, we classify these structures conceptually into the following categories: context, target and risk structures.

Based on these structures, there are several derived features that influence the decision of what surgical procedure to perform and approach: tumor diameter, distance of

the tumor to the anal sphincter, distance of the tumor to mesorectum border and distance of the mesorectum to the autonomic nerves. The tumor diameter and distance to the anal sphincter influence the staging and determine if an APR or low anterior resection have to be performed. The distance of the tumor to the mesorectal fascia influences if the circumferential resection margin is in danger of being positive, i.e., if tumor tissue is found in the resection border. The distance of the mesorectum to the autonomic nerves determines how careful surgeons need to be when cutting along that specific part of the mesorectum.

While the first three features are easily accessible from the MRI scan, the fourth feature is not, since the autonomic nerves are not visible in MRI scans. Our approach relies on registration to acquire these nerve zones and we must visualize the confidence in the registration success in order to not mislead the user. The MRI scan already provides much information, but for surgeons in training the translation between 2D MRI visualizations and 3D patient anatomy is more challenging.

It should be noted that, unlike for instance in biopsy planning or neurosurgery, in this procedure, there is no access path planning involved. The total mesorectum needs to be resected and there are no alternative approach paths among which the best must be selected. In our case we aim to increase awareness on the spatial extent of the patient-specific anatomy and pathology, as well as vital structures that are invisible both in MRI and during surgery and around which the surgeon has to be especially careful during the resection.

Based on this information, we formulated the following requirements in close collaboration with our domain experts:

- **Requirement 1:** The context and risk structures must be visualized in such a way that they do not occlude the target structures, while preserving shape perception.
- **Requirement 2:** The distance of the risk structures to the target structures and between the target structures must be perceivable from the visualization.
- **Requirement 3:** The relation between the MRI scan and 3D patient-specific models must be understandable from the visualization.
- **Requirement 4:** The user should be able to estimate the confidence in the result of the registration process.

By satisfying these four requirements our method aims to improve spatial understanding of anatomy and pathology, as well as risk assessment.

8.5. PELVIS

In this section, we describe our visualization design decisions with respect to the requirements and the components of our PelVis pipeline (Figure 8.1). Our method relies on MRI and atlas data as input and requires segmentation, registration and distance calculation as preprocessing steps (see Section 8.5.1). The resulting data and features are then used in three visualization methods aimed at visualizing context, target and risk structures (see Section 8.5.2). We discuss available interaction techniques in our application in Section 8.5.3.

8.5.1. PREPROCESSING

Since the risk structures, the regions where the autonomic nerves are located, are not visible in MRI scans, we make use of an atlas that contains these regions. To make the atlas patient-specific, we need to register the atlas labels to the MRI, a process which is described in Chapter 3. The inputs to our registration step are the label volume containing the segmented structures from the MRI and the label volume of the matching structures in the atlas. In this case, the atlas labels represent our moving image data, while the MRI labels are our target image data. As an initialization, we perform a rigid alignment. After this, we want to deform the atlas structures locally to fit the MRI as closely as possible using an automatic deformable B-spline registration [25]. The specific implementation we used employs an adaptive stochastic gradient-descent optimization [26]. Due to the use of gradient information, registering label volumes directly is unlikely to be successful. Therefore, we calculate distance fields on both label volumes and use these as the input for the deformable registration step. The resulting deformation field is then applied to the risk structures from the atlas to map them into the patient-specific MRI. More advanced techniques, which require less laborious manual work could be used for the segmentation. However, to test the concepts presented in this paper, we estimated that manual segmentation was the best option, since it does not introduce further inaccuracies that potentially arise from an automatic segmentation method. Due to the registration process, we might not be able to perfectly map all structures to each other, resulting in an uncertainty that we take into account in our visualization.

We rely on a simple indication of the registration success. Since we assume the proportional distance of the risk zones to the organs is preserved, we can estimate the registration success by checking how well the borders of the registered atlas organs align with the borders of the segmented MRI organs. To find this value, we check the distances for every point on the structure borders and scale them to a value between 0 and 1, where 0 indicates a perfect border alignment and 1 is a maximum difference. While more complex methods could be used to determine the confidence we have in the registration outcome, such as bootstrap resampling for uncertainty estimation [27], this is considered out of the scope of this paper.

Since visualization of the distances between the risk zones and target structures is a requirement (see Requirement 2), we calculate the distance between risk structures and target structures using an N-dimensional version of the Euclidian distance mapping proposed by Danielsson [28]. We calculate the distance field on the binary volumes containing the structures of interest. After the segmentation, registration and distance calculations are performed, we obtain the needed components that serve as the input to our visualization methods.

8.5.2. VISUALIZATION

Depending on the type of the structure, we have selected various visual encodings to support the user in relation to the requirements stated in the previous section. In this sub-section, we describe the design decisions made in visualizing the context, target and risk structures.




Figure 8.2: Four different representations of two context structures: opaque, semi-transparent, ghosted view and our approach.




Figure 8.3: Four different confidence visual encodings: a) blended gridlines, b) halftone, c) isoline confidence and d) isoline blur.

CONTEXT STRUCTURES

8

The context structures serve to provide anatomical context during the surgical planning and are arranged in multiple layers. According to Requirement 1, we need to visualize them in such a way that they do not obstruct the view on target structures, while maintaining their shape perception. This requirement rules out standard opaque rendering styles. Smart visibility techniques, which are illustrative and only convey the most important visual information via a high level of abstraction, can be used for occlusion management [29]. We combine two visualization techniques: a ghosted view to manage occlusion, and hatching to emphasize the shape and depth perception. The resulting visualization for the context organs is shown in Figure 8.2d.

While standard semi-transparent techniques could be employed, ghosted views have already shown to support accurate spatial analysis of vessels with aneurysms [30]. In our application domain, we have several nested structures with complex shapes, for which semi-transparent approaches fail to provide enough shape and depth information (see Figure 8.2a and 8.2b). Therefore, to prevent occlusion of the target structures, we render the context structures as a ghosted view. In our ghosted view, we use a view-dependent transparency rendering that is more opaque at regions facing away and more transparent at regions facing towards the viewer (see Figure 8.2c). We use an approach similar to the one proposed by Gasteiger et al. in which the opacity is determined based on the

Fresnel-reflection model [31].

Our ghosted view is able to prevent occlusion of the target structures, but makes shading perception more difficult than when using a standard semi-transparent representation. Both these rendering styles fail to convey the correct depth ordering. Therefore, we turn to illustrative techniques to emphasize the shape and depth ordering of the context structures. In the work by Interrante et al. [32], it was shown that the shape of an transparent surface and its relative depth distance from an underlying opaque object could be more accurately perceived when the surface is rendered with a sparse, opaque texture. Since we do not want to add to the pre-processing required for our method, we employ a simple hatching technique to convey shading and enhance the contour region perception. Hatching is able to provide shading cues without fully occluding the covered context structures or influencing the color perception of underlying structures. We chose to use image-based hatch strokes as developed by Lawonn et al. [33]. This hatching method is based on the dot product between the view direction and the surface normal and adds more strokes for 'darker' regions, where this dot product is small, when placing a light source at the camera position. To enhance the depth-ordering perception, we add a silhouette outline based on the same dot product (see Figure 8.2d). This approach for rendering silhouette outlines is based on the method by Gooch et al. [34].

TARGET STRUCTURES

The target structures are challenging to visualize in our application since they are nested; the tumor resides within the mesorectum. According to Requirement 2, we need to visualize both the distances of the risk structures to the targets, as well as the distances of the target structures among themselves, i.e., the tumor distance to the mesorectal wall. Furthermore, according to Requirement 4, we also need to visualize how confident we are in the calculated distances of the risk structures to the target structures, since they are the result of a registration process that introduces uncertainty.

After calculating the distance fields for both the risk structures and tumor in the pre-processing step, we visualize these by sampling the distances at the mesorectum border. For both distances, we employ a colormap from red to white to indicate proximity to the risk structures and the tumor respectively. We chose this colormap due to the intuitive meaning, i.e., red is dangerous and white is safe, and additionally because it does not interfere with the shading. To further enhance the distance perception, we add isolines at distances of interest. These isolines are formed by connection points of equal distance for several distances. Besides indicating these distances, these lines further emphasize the shape of the target structure, as contour lines such as these can enhance shaded surfaces to make a shape easier to perceive [35]. House and Ware state that contours can be thought of as a special type of texture that follows the shape of a surface [36]. In this way, we encode distances in a continuous representation using color mapping, as well as in a quantized manner using the isolines. While slice-based and glyph-based approaches worked well in the work by Dick et al. [4], the complex shape and spatial relationships of our target structures in relation to the risk zones, combined with the requirement to also visualize confidence, generates a situation in which clutter can become an issue. This makes a simple representation via color mapping and isolines a better visual encoding choice for our purposes.

Besides the distances, we also need to visualize the confidence that we have in the correctness of the risk-target distance calculation based on the registration success according to Requirement 4. For this, we calculate a local metric of registration confidence given by how well the registered structures of the atlas match up to the segmented MRI structures. We sample this metric at the mesorectal border and need to visualize the distance and the distance's confidence concurrently. While a separate view could be employed just to visualize confidence, our experts are only interested in the confidence and the distance simultaneously. A straightforward solution would be to map the confidence to color, but since we already encode distance using color, we looked into several alternatives that can be perceived simultaneously.

We developed two global visual encodings and two local visual encodings that are shown only near the isolines (see Figure 8.3). While textures could be employed to encode confidence [37], for our clinical context we do not want to perform texture-mapping beforehand and rely on screen-space solutions. For the first global option, we blend in a grid based on the amount of confidence. In regions where we are less confident, the grid is shown more clearly (see Figure 8.3a). The second option is to use a halftone pattern, where the size of the dots corresponds inversely to the confidence (see Figure 8.3b). In this case, large dots indicate that we are not confident about the distances in these areas. Both these options were selected in order not to interfere with the standard surface shading and hatching of the context structures.

Besides these two global representations, we also provide two localized visualization options. In the first, we show a 'confidence band' around the isolines (see Figure 8.3c). Here, a wider and lighter band indicates low confidence, while certain areas show a narrow dark band. In the second local option, we blur the isolines based on the amount of confidence (see Figure 8.3d). When confidence is low, the isolines are blended over a wider area than in areas where we are confident of the registration's success.

RISK STRUCTURES

The risk structures are similar to the context structures in that they should not occlude the target structures, according to Requirement 1. For this reason, we visualize them in a similar way by employing a ghosted view. The main difference in our application between the risk structures and the context structures, is that the risk zones are not an actual anatomical structure, but rather a zone in which the autonomic nerves are potentially located. We do not employ hatching to differentiate the look of the risk zones from the other context organs. However, to emphasize height information, which is important for the surgical approach in relation to the height of the tumor, we add contour bands at fixed intervals. These contour bands also serve to emphasize the different shapes of the risk zones for varying heights (see Figure 8.4).

8.5.3. INTERACTION

To satisfy Requirement 3, we visualize the MRI data and patient-specific 3D models in linked views. Furthermore, we enrich the MRI with contours of the segmented and registered structures by placing colored lines on the intersection of the MRI plane and the 3D models. The colors of these lines correspond to the structure color in 3D to further enhance the link between the views.




Figure 8.4: Visualization of a target structure and risk structure from a frontal and side view.

The visualization of the target structures shows the mesorectum in an opaque way and since the tumor is nested within, it is not directly visible. Simply making the mesorectum transparent to reveal the tumor does not work due to the number of overlapping organs. Instead, to show the spatial relation between the tumor and the mesorectal fascia, we employ a cutting and unfolding interaction technique. This 'exploded view' was inspired by the volumetric brain cleaving by van Dixhoorn et al. [38], which is based on the Hinge Spreader technique proposed by McGuffin et al. [39]. Since the human body is mainly symmetric along a sagittal plane, we place an interactive splitting plane that can be freely positioned in a region of interest. Then, the left and right halves of the structures can be interactively moved away from the center by rotation away from the center plane, similar to opening a book, combined with simultaneous translation (see Figure 8.6). The structures are rotated and translated subsequently in a staggered movement from outer to inner structures, which results in an interaction metaphor similar to peeling the model open layer by layer. In the center of the view, we place a copy of the uncut tumor for inspection along with the isolines that give an indication of the spatial extent of the mesorectum around the tumor. The unfolded visualization leaves space in the center for ghosted copies of individual structures to be shown when desired.

Based on the view during a surgical approach, we additionally define a default camera viewpoint from a similar perspective. In this 'surgical view', the pelvis is viewed from the top down, as is the case during both laparoscopic and open surgery. In this viewpoint, the unfolding movement is adjusted to take the current view into account.

8.6. IMPLEMENTATION

The required pre-processing MRI segmentation of the major organs was performed semi-automatically in AMIRA and MITK [40]. We used RegistrationShop [41] and Elastix [42] to perform the rigid and deformable B-spline registration respectively. After this, we reconstructed surfaces for the segmented organs and risk zones using DeVIDE [43]. Then, we calculated the distance fields from the nerve zones and tumor and stored this information in the mesorectum surface mesh. Finally, we calculated the distance confidence using a combination of DeVIDE and Matlab to compute the registration alignment mis-

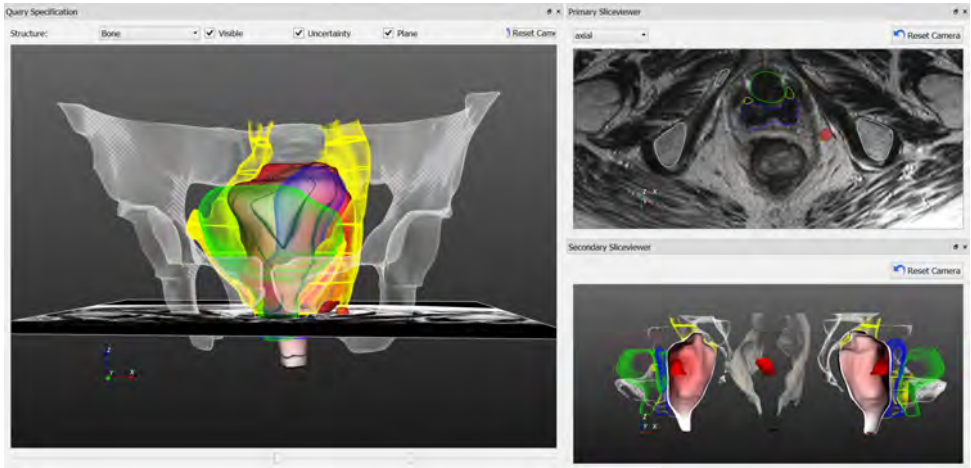


Figure 8.5: The PelVis prototype application featuring the 3D model on the left and a linked MRI and unfolded view visualization on the right.

match from the organs, and also stored these values in the mesorectum mesh. To get the distance confidence, we employ a simple method, where we calculate a distance field over all organ edges, and sample this at the mesorectal border and normalize it, leaving us with an indication of the registration organ alignment success.

We implemented the visualization and interaction methods in a prototype application in Python using the Visualization Toolkit (VTK) and Qt for the interface, extended with custom GLSL vertex and fragment shaders, with separate shaders for context, risk and target structures. The target structure shaders employed the values stored in the mesorectum mesh to generate the target structure distance and confidence visualization. For the MRI contours, we simply cut the structure surfaces at the location of the MRI plane. For the unfolding, we cut the surface data into a left and right half at the user-specific plane location, and rotate and translate the halves in opposite directions.

8.7. RESULTS

For our prototype application, we used the VSP atlas which is based on the Visible Korean female cryosectional dataset, as described in Chapter 3. We registered the atlas to five patient-specific pre-operative MRI scans of patients undergoing a TME procedure with various tumor locations and types and reconstructed the major surgically relevant structures and risk zones in 3D using our method.

A screenshot of the interface can be seen in Figure 8.5. In the 3D model the context structures are visualized in white (bone), green (bladder), and blue (vagina). The risk structures containing the autonomic nerves are shown in yellow, while the target structures are visualized in red/white (mesorectum) and red (tumor). The colors of the bone and risk zones are based on anatomical illustration standards, e.g., always coloring bone white and nerves yellow. Since the mesorectum and tumor are the target structures and this signifies importance, we use the color red to visualize these. Since there




Figure 8.6: Tumor distance and unfolded visualization for two datasets featuring a mid-rectal tumor and a low-rectal tumor respectively.

is no anatomical illustration standard for coloring the vagina and bladder, we chose colors that do not interfere with the perception of the underlying red color, and base our choices on the fact that red, green, yellow, and blue are hard-wired into the brain as primaries and should be considered first [35]. A linked 3D cursor (red sphere) relates the MRI data to the 3D model. When clicking on a structure in the 3D view, the MRI slice updates to the point corresponding to the clicked location. In the split view, the interior of the mesorectum can be explored while the cut structures can be displayed in their original form in the center on mouse hover. In the application, one of the four proposed distance confidence representations can be chosen to reveal the confidence in the outcome of the registration process, i.e., blended grid, halftone, contour confidence and contour blur. Individual structures can be made visible or invisible through selection. A surgical view is also available in which the structures are visualized from a familiar surgical orientation.

In Figure 8.6, two different datasets are shown based on scans of two patients. The first patient is a 71-year old female who has a mid-rectal tumor, while the second patient is a 92-year old female that has a low-rectal tumor. In comparison, the organ shapes between them vary a lot, as is typical for soft tissues such as these. Especially organs such as the bladder and rectum, which can have different levels of filling at the time of the scan, result in large size and shape variations. The tumor distance visualizations on the mesorectal wall reveal the mid- and low-rectal location of the tumor respectively.

Furthermore, the unfolded view shows that in the first dataset, the tumor is located on the ventral side of the mesorectal fascia, while the second dataset tumor is located more towards the back. The color-coding and unfolded visualization show the proximity of the tumor to the mesorectal wall and the chances of a clean circumferential margin.

8.8. EVALUATION

In this section, we describe the setup and results of our evaluation with five domain experts and ten non-domain experts.

8.8.1. EVALUATION SETUP

We evaluated the utility of our method with five domain experts: S1, S2, S3, S4 and S5. S1 is an oncologic surgeon who is specialized in the surgical treatment of rectal cancer. He has over 20 years experience as a surgeon. S2 is a medical resident who recently obtained her PhD in surgical anatomy. She plans to be a surgical oncologist and she is one of the domain experts involved in this project. S3 is a researcher and surgical oncologist in training who finishes her training next year. In her post-doctoral research, she focuses on imaging strategies for treatment response to therapy in rectal and breast cancer. S4 is a surgical oncologist with over 12 years surgical experience. S5 is a gynecologist in training with 7 years of experience in general medicine and surgery.

After a short demonstration of the features of our application, the participants were encouraged to interact with the tool. We allowed the participants to comment or pose questions during their session and afterwards asked several open questions in a semi-structured interview. Finally, we asked the participants indicate their (dis)agreement with 30 statements using a five-point Likert scale (see Table 8.1). The 30 statements were categorized into the following topics: General application (A), Context structure visualization (B), Target structure visualization (C), Risk structure visualization (D), MRI visualization (E), and Interaction (F).

In the form, the question order was randomized and symmetry was maintained by having equal numbers of positive and negative positions to prevent bias. However, for the ease of interpretation in the table, we rephrased the negative statements (indicated by a ★) to their positive form and inverted the scores, as described in the Sage handbook of methods in social psychology [44]. Additionally, participants were asked to rate the effectiveness of the four distance confidence visualization options in Figure 8.3 on a scale of 1 to 10.

To quantitatively evaluate the effectiveness of the four visual encodings of the combined distance and confidence, we performed an additional task-based evaluation with 10 non-expert participants. We presented 8 circular patches of parts of the mesorectum in the four confidence visualization styles. These patches were deemed safe or unsafe, either through their distance to the risk zones or due to high uncertainty. We asked the participants whether these regions are safe or not, and how sure they were about their answer. We measured the time that it took them to get to their answers.

Statement	S1	S2	S3	S4	S5
A1 The application can improve the planning of a TME procedure	3	5	4	5	4
A2 The application can help surgeons in training in preparing a TME procedure	5	5	5	5	4
A3 The application has value in the OR during the surgery★	3	3	2	2	2
A4 I would find this application useful in preparing a surgery★	4	5	5	4	5
A5 The application has added value over the current situation★	4	4	4	4	5
A6 With this, surgical complications during the procedure are easier to predict	2	4	4	4	4
A7 I would like to use this application to explain the procedure to my patients	5	4	2	1	5
B1 The context organs are clearly visualized	5	5	5	5	5
B2 The context organs do not hamper the view on the important structures★	5	4	4	5	4
B3 The context organs help me with spatial orientation	5	4	5	4	4
B4 The shape of the context organs is clearly visible ★	5	4	4	5	5
C1 The mesorectum is clearly visible at all times	5	5	4	5	4
C2 I can estimate the distance of the tumor to the mesorectal fascia well in 3D★	4	4	5	4	5
C3 I can estimate the distance of the risk zones to the mesorectum in 3D	4	5	4	4	5
C4 The uncertainty resulting from the registration can be clearly interpreted ★	4	4	4	4	4
C5 The contours help me estimate the distances to the mesorectum	4	3	5	5	4
C6 The color-coding helps me estimate the distances to the mesorectum	4	4	5	4	5
D1 The risk zones that contain the autonomic nerves are clearly visible in 3D★	4	4	4	4	5
D2 The lines in the risk zones help me estimate the height of structures	4	4	4	5	5
D3 The changes in risk zone shape per level are easier to interpret in 2D than 3D	4	1	1	2	2
D4 Showing the autonomic nerve zones is of added value for an experienced surgeon★	4	5	2	2	4
D5 Showing the autonomic nerve zones is of added value for a surgeon in training	5	5	5	5	5
E1 The contours on the MRI are of added value for an experienced surgeon★	3	4	4	2	4
E2 The contours on the MRI are of added value for a surgeon in training	4	5	5	5	5
E3 The MRI slice in the 3D visualization helps me with spatial orientation★	4	4	5	5	5
E4 The size of the tumor is the most clearly visible in the MRI	4	2	2	4	3
F1 Unfolding the structures helps me see the interior of the mesorectum★	4	4	4	3	4
F2 Unfolding helps me to estimate the distance of the tumor to the mesorectal fascia	4	4	5	3	4
F3 The surgical view is of added value in addition to the default view★	5	4	5	5	3
F4 The 3D cursor helps me connect the 3D and 2D visualizations	5	5	4	5	5

Table 8.1: User response to 30 statements on a 5-point Likert-scale: 1: Strongly disagree, 2: Disagree, 3: Neither agree nor disagree, 4: Agree, 5: Strongly agree. Negatively phrased statements in the original form are indicated by a ★ and their scores are inverted for ease of interpretation.

8.8.2. EVALUATION RESULTS

During the semi-structured interview, S1 mentioned that he sees strong potential in our work for surgeons in training. Especially if we process interesting or difficult cases in several pathology categories and represent them in our tool, it can help surgeons in training to better understand the anatomy and relations between the MRI and the 3D situation. Experienced surgeons can make the translation from MRI to 3D anatomy more readily. He stated that even for experienced surgeons, sometimes the 3D reconstructions can reveal details that are not immediately apparent from the MRI, due to cognitive blind spots and selective perception. In particular, he appreciated the, to him, novel distance contour representations and interactive unfolding. While he describes the registration uncertainty and influence on the distance to the risk zones as important, he states that surgeons should always be aware that the tool provides an indication of the areas, but no exact definitions. A remaining challenge to get the tool into clinical practice, according to him, is to further automate the registration process, so that MRI organ segmentation is no longer needed.

S2 sees potential in our work for both surgeons in training and experienced surgeons. She states that experienced surgeons can still benefit from our tool in difficult tumor cases. She especially appreciated the linked views between the MRI and the 3D representations. Her suggestions for improvement of the application are to include more

anatomical structures such as the levator ani muscle and the anal sphincter complex. Both these structures are visible from the MRI and available from the atlas. Furthermore, she would find it useful if surgeons were able to store viewing preferences.

S3 especially appreciates the contours indicating the structures in the MRI view and the linked 3D cursor. She thinks this really helps surgeons in training to relate the 2D MRI data to 3D anatomy. Furthermore, the unfolded view helps her understand the relation of the tumor to the mesorectal fascia in a clearer way. She also appreciates the surgical view and viewing the MRI and the 3D model from that perspective. She indicates that an extended version of the current application could also be useful during the operation itself additionally.

S4 states that he sees a clear benefit in our application for surgical education, as well as for surgical planning. He mentions that positive resection margins are often found on the ventral part of the mesorectum, near the vaginal wall/prostate, and that it is a surgically difficult area, which can be clearly visualized with our application. He emphasizes that the annotation of the tumor in the MRI must be done by a radiologist to prevent incorrect assumptions in planning the procedure. If that condition is met, he thinks it can raise awareness of surgeons in the pre-operative planning phase and that they can transfer this knowledge into the operating theater. In the future, he would like to see the application made available for surgical guidance, by tracking and registering the position and viewpoint of surgical tools.

S5 finds the application very useful and clear. She responded positively to the visualization of the context, target and risk structures. Furthermore, she sees a clear need for this application in surgical training. She finds the representations visually pleasing.

The level of agreement of the participants with the 30 statements that were presented can be seen in Table 8.1. In general, all participants agree that the application is useful in pre-operative planning and has added value over the current situation (A2, A4, A5, A6). S1 stated that the application can only improve TME planning in clinical practice if the registration and segmentation is further automated, therefore, he gives a relatively low score in A2. All users agree that the application is not yet ready to be brought into the OR at this point (A3), however, they see the potential for an extension of this work. Some surgeons would like to use the application to explain the procedure to the patients, while others think the patients already get overwhelmed by too much information and it would take too much time to explain what the application shows (A7).

The visualization of the context and target organs was positively to very positively evaluated by all participants (B1 to C6). Especially the visualization of context organs was found to be very clear by all. While the risk zone indication was found to be clear (D1, D2), not all participants agreed it would be useful to experienced surgeons, indicating they expect this information is already known to them (D4). Interestingly enough, the most experienced surgeons, S1 and S4, did evaluate this positively, even for experienced surgeons such as themselves. For surgeons in training, all participants strongly agreed that it would be useful to show these zones (D5). While S1 preferred the 2D MRI view to see the risk zone shape varying per level, all others preferred the 3D visualization (D3).

The MRI visualization was found useful in both 2D and 3D and useful especially to surgeons in training, though S4 pointed out that experienced surgeons will already be able to interpret this information (E1, E2, E3). While two of the participants preferred to

Table 8.2: Confidence visualization option ratings.

Representation	S1	S2	S3	S4	S5
Blended gridlines	7.00	7.00	7.00	6.00	8.00
Halftone	7.75	8.00	8.00	7.00	10.00
Isoline confidence	7.00	3.00	6.00	9.00	10.00
Isoline blur	8.25	7.00	5.00	8.00	8.00

Representation	Accuracy	Avg time	Confidence
Blended gridlines	0.80	10.17	3.01
Halftone	0.85	9.99	2.90
Isoline confidence	.95	10.73	2.79
Isoline blur	0.75	10.00	2.57

estimate the size of the tumor in the MRI, two preferred the 3D view for this purpose (E4). The unfolding interaction to reveal the relation between the tumor and the mesorectal fascia was rated positively by all surgeons (F1, F2). The surgical view was found beneficial by four out of five participants (F3). The 3D linked cursor between the MRI and 3D representation was found very helpful by all participants.

In Table 8.2, we present the responses of the expert participants regarding the visual encoding options to visualize the confidence of the risk zone distance on the mesorectum. In general, the response to all visual encodings was positive, but each participant had a personal preference towards one or two of the encodings. The option that was ranked the highest on average by the users is the halftone representation, while no participant had a real preference for the blended grid approach. All participants appreciate the clarity and intuitive understanding the halftone representation brings. The other representations were met with mixed responses. Especially the isoline confidence representation generated polarizing responses. While some participants thought it was very clear, others were confused by the view.

In Table 8.2, we present the conclusions of the task performance experiment with ten non-experts. The accuracy (number of correct answers divided by the total number of tasks) was highest for the isoline confidence, and worst for the isoline blur, but only by a small difference. Average times were also similar. The participants felt more confident in their answers with the global methods, but again by a small amount.

8.8.3. EVALUATION CONCLUSION

From Table 8.1, we conclude that the application was rated positively overall. Every participant saw the benefits of our application in improving surgical planning for surgeons in training. Furthermore, the surgeons indicated that for difficult cases, the application is also of value to experienced surgeons. The visualization of the context and target structures was found to be clear in both the enriched 2D MRI representation and the 3D models. While opinions varied on the utility of showing the risk zones to experienced surgeons, the visual representation was again found to be clear in both 2D and 3D. The MRI visualization and relation to the 3D anatomy was found to be a valuable addition. Furthermore, the interaction options helped the surgeons get insight into this relation, as well as the relation between the tumor and the mesorectal fascia. Among the confidence visualization options, the halftone representation was found to be the most effective by the experts on average. In the quantitative evaluation of these four

combined distance and confidence representations, there was no clearly better or worse performing representation, indicating that either a more thorough user study is needed with more data and participants, or that the choice of representation is mainly a personal preference.

8.9. DISCUSSION

In the analysis of this problem domain, we found several distinguishing features that set this particular application apart from other surgical-planning visualization applications. Especially the combination of complex anatomy and visualizing structures, which are invisible during surgery or in the scans, provides a challenge that distinguishes surgical planning for pelvic oncological procedures from surgical planning in other domains. While surgical planning visualization systems often focus on access planning, implant planning or reconstruction planning, in our case improving the spatial understanding of anatomy and pathology is the main focus. In the TME procedure specifically, due to the confined space and the fact that the complete mesorectum needs to be excised, the surgical planning phase is not about planning an access path, but more about choosing the type of surgery that is most appropriate for the pathology. Furthermore, the surgeon needs to be aware of the location of the autonomic risk zones surrounding the mesorectum in order to prevent damaging them. Since the zones in which the autonomic nerves reside in our atlas are defined quite coarsely, it is questionable if this information brings much new knowledge to experienced surgeons, who are already aware of these regions, but would like to see the exact spatial extent.

This particular problem domain poses several general visualization challenges. The complex pelvic anatomy, featuring many closely arranged structures that vary quite extensively in shape and size between patients, results in the need to illustrate the spatial anatomical context without distracting or occluding the structure that needs to be resected. For this, illustrative visualization methods with a high level of abstraction are especially suitable. Furthermore, the 'invisible' regions containing the risk structures and the distance of these regions to the target structure need to be emphasized without occluding the target structure. Finally, the confidence in the outcome of the registration process and the influence on these distances needs to be visualized in order to not mislead the user with a false sense of accuracy. Since this confidence is related to the distance, both information types need to be conveyed in visual channels that can be perceived simultaneously.

The limitations of our approach lie mainly in the atlas itself, the registration process and the amount of time required for the preprocessing steps. Our method relies on the risk zones defined by the atlas, which has limited accuracy, as mentioned in Chapter 3, e.g., it does not include multiple individuals and it exclusively contains female pelvic anatomy.

The registration confidence calculation could be improved to reveal more information about the uncertainty, by including for instance in which direction the mismatch occurred. This could provide more insight into the influence of the registration confidence on the risk zone distance calculation.

Currently there is pre-processing required in terms of segmentation, registration and distance calculation. While the distance calculation and registration processes them-

selves are fast, the manual segmentation of the MRI scan still takes up to 30 minutes. To bring our method into clinical practice for pre-operative planning purposes, this time needs to be reduced. We could achieve this by further automating the segmentation process. Once many MRI scans are registered to the atlas, we can perform MRI to MRI registration and employ atlas-based segmentation to automatically segment the major organs and bone structures from the MRI scan. Atlas-based segmentation is already successfully applied to brain and cardiac segmentation tasks [45].

Despite these limitations, as shown in the evaluation, we have provided a method that is already suitable for surgical planning, for surgeons in training and experienced surgeons alike, that forms a solid basis for further developments to bring it into clinical practice.

8.10. CONCLUSION

We presented PelVis, a method to visualize context, target and risk structures for preoperative planning of pelvic oncologic surgical procedures. Using occlusion-management and distance-based visualization techniques, the surgically relevant structures are visualized based on a patient-specific 3D model of the pelvis, acquired through atlas registration to pre-operative MRI scans. In an interactive application featuring several linked views, the spatial relations in the complex anatomy can be freely explored. From our evaluation we concluded that our prototype application has great potential in surgical planning and especially in surgical training for oncologic surgeons in training. However, before it is ready for clinical use in planning and surgical guidance, several steps still need to be taken. First, the user interface should be improved to fit the clinical requirements, e.g., minimum interaction required to get the desired information. We would also like to look into applying our methods to other procedures, since the methods are generally applicable and may be of value for instance in radiotherapy planning or even surgical simulation. Furthermore, this work could be extended to bring the application into the operating room and using instrument tracking to interactively update the views for surgical guidance.

REFERENCES

- [1] N. Smit, K. Lawonn, A. Kraima, M. DeRuiter, H. Sokooti, S. Bruckner, E. Eisemann, and A. Vilanova, *Pelvis: Atlas-based surgical planning for oncological pelvic surgery*, IEEE Transactions on Visualization and Computer Graphics (Proceedings of Scientific Visualization 2016) **23** (2017), accepted to appear.
- [2] B. Preim and C. P. Botha, *Visual Computing for Medicine: Theory, Algorithms, and Applications* (Newnes, 2013).
- [3] G. E. Marai, D. H. Laidlaw, Ç. Demiralp, S. Andrews, C. M. Grimm, and J. J. Crisco, *Estimating joint contact areas and ligament lengths from bone kinematics and surfaces*, Biomedical Engineering, IEEE Transactions on **51**, 790 (2004).
- [4] C. Dick, R. Burgkart, and R. Westermann, *Distance visualization for interactive 3D*

- implant planning*, Visualization and Computer Graphics, IEEE Transactions on **17**, 2173 (2011).
- [5] J. Süßmuth, W.-D. Protozerakis, A. Piazza, F. Enders, R. Naraghi, G. Greiner, and P. Hastreiter, *Color-encoded distance visualization of cranial nerve-vessel contacts*, International journal of computer assisted radiology and surgery **5**, 647 (2010).
- [6] C. Rieder, T. Kroeger, C. Schumann, and H. K. Hahn, *GPU-based real-time approximation of the ablation zone for radiofrequency ablation*, Visualization and Computer Graphics, IEEE Transactions on **17**, 1812 (2011).
- [7] C. Hansen, S. Zidowitz, M. Hindennach, A. Schenk, H. Hahn, and H.-O. Peitgen, *Interactive determination of robust safety margins for oncologic liver surgery*, International journal of computer assisted radiology and surgery **4**, 469 (2009).
- [8] D. F. Marshall, H. J. Gardner, and B. H. Thomas, *Interactive visualisation for surface proximity monitoring*, in *Proceedings of the 16th Australasian User Interface Conference (AUIC 2015)*, Vol. 27 (2015) p. 30.
- [9] A. Krüger, C. Tietjen, J. Hintze, B. Preim, I. Hertel, and G. Strauß, *Interactive visualization for neck-dissection planning*, in *EuroVis*, Vol. 5 (2005) pp. 295–302.
- [10] S. Djurcilov, K. Kim, P. Lermusiaux, and A. Pang, *Visualizing scalar volumetric data with uncertainty*, Computers & Graphics **26**, 239 (2002).
- [11] G. Grigoryan and P. Rheingans, *Probabilistic surfaces: Point based primitives to show surface uncertainty*, in *Proceedings of the conference on Visualization'02* (IEEE Computer Society, 2002) pp. 147–154.
- [12] R. T. Whitaker, M. Mirzargar, and R. M. Kirby, *Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles*, Visualization and Computer Graphics, IEEE Transactions on **19**, 2713 (2013).
- [13] M. Termeer, J. O. Bescós, M. Breeuwer, A. Vilanova, F. Gerritsen, M. E. Groller, and E. Nagel, *Visualization of myocardial perfusion derived from coronary anatomy*, Visualization and Computer Graphics, IEEE Transactions on **14**, 1595 (2008).
- [14] R. Heald, E. Husband, and R. Ryall, *The mesorectum in rectal cancer surgery—the clue to pelvic recurrence?* British Journal of Surgery **69**, 613 (1982).
- [15] R. Heald and R. Ryall, *Recurrence and survival after total mesorectal excision for rectal cancer*, The Lancet **327**, 1479 (1986).
- [16] I. D. Nagtegaal and P. Quirke, *What is the role for the circumferential margin in the modern treatment of rectal cancer?* Journal of Clinical Oncology **26**, 303 (2008).
- [17] A. Wibe, P. Rendedal, E. Svensson, J. Norstein, T. Eide, H. Myrvold, and O. Søreide, *Prognostic significance of the circumferential resection margin following total mesorectal excision for rectal cancer*, British Journal of Surgery **89**, 327 (2002).

- [18] J. S. Park, J. W. Huh, Y. A. Park, Y. B. Cho, S. H. Yun, H. C. Kim, W. Y. Lee, and H.-K. Chun, *A circumferential resection margin of 1 mm is a negative prognostic factor in rectal cancer patients with and without neoadjuvant chemoradiotherapy*, *Diseases of the Colon & Rectum* **57**, 933 (2014).
- [19] I. D. Nagtegaal, C. J. van de Velde, E. van der Worp, E. Kapiteijn, P. Quirke, J. H. J. van Krieken, *et al.*, *Macroscopic evaluation of rectal cancer resection specimen: clinical significance of the pathologist in quality control*, *Journal of Clinical Oncology* **20**, 1729 (2002).
- [20] J. Leroy, F. Jamali, L. Forbes, M. Smith, F. Rubino, D. Mutter, and J. Marescaux, *Laparoscopic total mesorectal excision (TME) for rectal cancer surgery: long-term outcomes*, *Surgical Endoscopy and Other Interventional Techniques* **18**, 281 (2004).
- [21] A. Pigazzi, J. Ellenhorn, G. Ballantyne, and I. Paz, *Robotic-assisted laparoscopic low anterior resection with total mesorectal excision for rectal cancer*, *Surgical Endoscopy and Other Interventional Techniques* **20**, 1521 (2006).
- [22] S. Atallah, B. Martin-Perez, M. Albert, G. Nassif, L. Hunter, S. Larach, *et al.*, *Transanal minimally invasive surgery for total mesorectal excision (TAMIS–TME): results and experience with the first 20 patients undergoing curative-intent rectal cancer surgery at a single institution*, *Techniques in coloproctology* **18**, 473 (2014).
- [23] G. Brown, A. Kirkham, G. T. Williams, M. Bourne, A. G. Radcliffe, J. Sayman, R. Newell, C. Sinnatamby, and R. J. Heald, *High-resolution MRI of the anatomy important in total mesorectal excision of the rectum*, *American Journal of Roentgenology* **182**, 431 (2004).
- [24] P. A. Georgiou, P. P. Tekkis, V. A. Constantinides, U. Patel, R. D. Goldin, A. W. Darzi, R. J. Nicholls, and G. Brown, *Diagnostic accuracy and value of magnetic resonance imaging (MRI) in planning exenterative pelvic surgery for advanced colorectal cancer*, *European Journal of Cancer* **49**, 72 (2013).
- [25] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes, *Non-rigid registration using free-form deformations: application to breast MR images*, *Medical Imaging, IEEE Transactions on* **18**, 712 (1999).
- [26] S. Klein, J. P. Pluim, M. Staring, and M. A. Viergever, *Adaptive stochastic gradient descent optimisation for image registration*, *International journal of computer vision* **81**, 227 (2009).
- [27] J. Kybic, *Bootstrap resampling for image registration uncertainty estimation without ground truth*, *Image Processing, IEEE Transactions on* **19**, 64 (2010).
- [28] P.-E. Danielsson, *Euclidean distance mapping*, *Computer Graphics and image processing* **14**, 227 (1980).
- [29] I. Viola and E. Gröller, *Smart visibility in visualization*, in *Computational Aesthetics* (2005) pp. 209–216.

- [30] A. Baer, R. Gasteiger, D. Cunningham, and B. Preim, *Perceptual evaluation of ghosted view techniques for the exploration of vascular structures and embedded flow*, in *Computer Graphics Forum*, Vol. 30 (Wiley Online Library, 2011) pp. 811–820.
- [31] R. Gasteiger, M. Neugebauer, C. Kubisch, and B. Preim, *Adapted surface visualization of cerebral aneurysms with embedded blood flow information*, in *VCBM* (Citeseer, 2010) pp. 25–32.
- [32] V. Interrante, H. Fuchs, and S. M. Pizer, *Conveying the 3d shape of smoothly curving transparent surfaces via texture*, *IEEE Transactions on visualization and computer graphics* **3**, 98 (1997).
- [33] K. Lawonn, S. Glaßer, A. Vilanova, B. Preim, and T. Isenberg, *Occlusion-free blood flow animation with wall thickness visualization*, *Visualization and Computer Graphics*, *IEEE Transactions on* **22**, 728 (2016).
- [34] B. Gooch, P.-P. J. Sloan, A. Gooch, P. Shirley, and R. Riesenfeld, *Interactive technical illustration*, in *Proceedings of the 1999 symposium on Interactive 3D graphics* (ACM, 1999) pp. 31–38.
- [35] C. Ware, *Information visualization: perception for design* (Elsevier, 2012).
- [36] D. House and C. Ware, *A method for the perceptual optimization of complex visualizations*, in *Proceedings of the Working Conference on Advanced Visual Interfaces* (ACM, 2002) pp. 148–155.
- [37] P. J. Rhodes, R. S. Laramée, R. D. Bergeron, T. M. Sparr, *et al.*, *Uncertainty visualization methods in isosurface rendering*, in *Eurographics*, Vol. 2003 (2003) pp. 83–88.
- [38] A. F. van Dixhoorn, J. Milles, B. van Lew, and C. P. Botha, *BrainCove: A tool for voxel-wise fMRI brain connectivity visualization*, in *VCBM* (2012) pp. 99–106.
- [39] M. J. McGuffin, L. Tancau, and R. Balakrishnan, *Using deformations for browsing volumetric data*, in *Visualization, 2003. VIS 2003. IEEE* (IEEE, 2003) pp. 401–408.
- [40] I. Wolf, M. Vetter, I. Wegner, T. Böttger, M. Nolden, M. Schöbinger, M. Hastenteufel, T. Kunert, and H.-P. Meinzer, *The medical imaging interaction toolkit*, *Medical image analysis* **9**, 594 (2005).
- [41] N. N. Smit, B. K. Haneveld, M. Staring, E. Eisemann, C. P. Botha, and A. Vilanova, *Registrationshop: An interactive 3d medical volume registration system*, *EG Workshop on Visual Computing in Biology and Medicine (VCBM)*, (2014).
- [42] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, *Elastix: a toolbox for intensity-based medical image registration*, *Medical Imaging*, *IEEE Transactions on* **29**, 196 (2010).
- [43] C. P. Botha and F. H. Post, *Hybrid scheduling in the DeVIDE dataflow visualisation environment*. in *SimVis* (2008) pp. 309–322.

- [44] C. Sansone, C. C. Morf, and A. T. Panter, *The Sage handbook of methods in social psychology* (Sage Publications, 2003).
- [45] M. B. Cuadra, V. Duay, and J.-P. Thiran, *Atlas-based segmentation*, in *Handbook of Biomedical Imaging* (Springer, 2015) pp. 221–244.

9

CONCLUSION

THE goals of this thesis are to unify existing surgically relevant anatomical knowledge, and to visualize this knowledge for educational and pre-operative planning purposes to improve functional outcome of the total mesorectal excision (TME). The previous chapters describe the steps we took towards reaching these goals. In the following, we discuss our main contributions and an outlook on new research challenges that arose from this thesis.

9.1. CONTRIBUTIONS

When we started looking at the requirements for the Virtual Surgical Pelvis (VSP), we found that an integration of multiple datasets is required in order to get a complete atlas. At the time, there was no way of storing, accessing and visualizing these heterogeneous datasets in an integrated fashion. Therefore, in Chapter 2, we proposed a representation that enables the integration of heterogeneous spatial and non-spatial data from different sources, and the relations between them, into a single model, all using standardized anatomical coordinates.

In Chapter 3, we describe the concrete development and medical application potential of the VSP: our virtual 3D anatomical atlas. While there were other 3D pelvic models at the time, which we described in a survey [1], none of them featured enough surgical detail to be useful for pre-operative planning and most of them were based on a generic artist impression of anatomy, rather than real anatomy. For this reason, we constructed a new model, based on real human cryosectional data, which provided a way to link the resulting 3D models directly to 2D high resolution image slices. These two chapters combined allowed us to reach our first goal: to unify existing surgically relevant anatomical knowledge.

After constructing the VSP, we wanted to make it available to medical students worldwide. For this, we developed the Online Anatomical Human (OAH), as described in Chapter 4. Via this web-interface, the relation between 2D and 3D anatomy can be explored interactively by the students. Furthermore, we added 3D annotation capabilities,

which allows us to quiz students or enrich the model based on more expert input. The alternative gesture-based input via the Leap Motion allows playful interaction and could be useful when using the OAH in an anatomical dissection class. The work described in this chapter of the thesis allowed us to reach the second goal: to share the knowledge gathered in the VSP for educational purposes.

While the VSP is based on multiple datasets, and can be deformed to take anatomical shape variations into account, there are also potential topological variations which may occur, for instance in arteries or veins. The current way to visualize these anatomical variations is via illustrations and text, and it is difficult to answer questions about similarity or to find intuitive groupings. For this reason, we developed an educational application that allows interactive visualization of anatomical variations in branching structures, as described in Chapter 5. This visualization was based on graph representations of the structures combined with a topology similarity metric. We provide a summary tree view, synthesized by a graph-matching algorithm, which allows users to explore local similarities and differences between multiple complex anatomical variations interactively.

We started the part of the thesis dealing with surgical planning with a survey of multimodal medical visualization in Chapter 6. While there are already surveys on many medical visualization topics and many application papers based on multimodal visualization, a review was still lacking. Even the comprehensive medical visualization textbook by Preim and Botha [2], was only briefly described as an add-on to a chapter on registration. In our survey, we provide an overview of multimodal acquisition techniques and current clinical workflow and relate this to requirements and challenges, as well as a taxonomy of multimodal visualization applications.

To use the VSP for surgical planning, registration of the atlas to patient-specific MRI was needed. We found that existing registration tools were mainly aimed at expert users, and it would be difficult to perform such a complex registration for clinicians. In Chapter 7, we developed an open source framework that combines simple 3D visualizations and interactive transformation tools to visually align two volumetric datasets, entitled RegistrationShop. Users can iteratively improve the registration result With instant visual feedback based on 3D multi-volume rendering. Domain experts evaluated our approach and underlined the potential and usefulness of RegistrationShop.

Once the VSP is registered to patient-specific MRI scans, it can be used for surgical planning. In Chapter 8 we developed an application, PelVis, for oncological pelvic surgery planning, which shows surgically relevant structures in linked 2D and 3D views. We present visualization methods tailored to represent combined distances to the nerve risk zones and confidence in the registration process, while avoiding occlusion and visual clutter. We evaluated PelVis with five surgeons, with varying levels of experience, and ten non-experts, which demonstrated the utility of our approach. This chapter made the first steps towards reaching the third goal: to improve pre-operative planning of the Total Mesorectal Excision surgical procedure. Before pre-operative planning can be truly improved in clinical practice, and the clinical goal of reducing surgical complications can be proven, further improvements and clinical trials are still needed.

9.2. OUTLOOK

In this thesis we have shown applications of the VSP in medical education and surgical planning. These applications would both benefit from further improvements made to the VSP. For instance, by adding a male dataset to the VSP, we would be able to visualize both genders for educational purposes. Additionally, surgical planning for male patients would be improved if a male atlas were available. With further cadaveric studies, the definition of the risk zones in the atlas could become more accurate. With each dataset we add to the VSP, we can cover more naturally occurring anatomical variations for both educational and surgical planning purposes.

In Chapter 6, we presented a survey on multimodal medical visualization, where we concluded that most of the current treatment planning techniques are aimed at neurosurgical planning, while there is not much work yet on other applications that might benefit from suitable visualization methods such as radiotherapy planning and oncologic surgery. Additionally, the research up until now is often not evaluated with respect to how the techniques can fit into the current clinical workflow and what additional benefits they provide for clinical practice. Evaluation in terms of relevant clinical user task performance and perceptual studies could help identify which techniques work well and what challenges still need to be solved before techniques are suitable for clinical use. Also, multimodal visualization offers unique opportunities to re-invent existing techniques such as automatic viewpoint detection taking the information the different modalities can offer into account, instead of just applying the same technique to both datasets.

As presented in Chapter 8, We have developed a pre-operative planning tool for the TME-procedure, but an open question remains how to bring the insights from a combination of the VSP atlas and patient-preoperative scan data into the operating room. While pre-operative planning is useful for the surgeons to explore the complex patient-specific anatomy and pathology, the surgeons expressed a need to bring this pre-operative knowledge into the operating theater for surgical training and guidance. I attended a laparoscopic pelvic oncologic surgical procedure and noticed that there is great potential for improving situational awareness using medical visualization techniques. Currently, only video information from the laparoscopic camera is presented to the surgeons and they rely on static MRI acquisitions or even anatomy textbooks to visualize anatomy and patient-specific pathology. Additionally, surgeons are expected to make a complex mental translation between the orientation of the camera with respect to the patient and their own orientation with respect to the camera, which is especially challenging for surgeons in training.

Improved surgical guidance using novel visualization techniques, such as automatic viewpoint selection, as well as the development of interaction techniques that are suitable for use in the OR, will make relevant information available to the surgeon, which fits the current situation in the theater. Such techniques have potential for improving functional outcome for the patients, by reducing chances of per-operative complications, as well as shortening the learning curve for surgeons in training.

REFERENCES

- [1] A. C. Kraima, N. N. Smit, D. Jansma, C. Wallner, R. Bleys, C. Van De Velde, C. P. Botha, and M. C. DeRuiter, *Toward a highly-detailed 3d pelvic model: Approaching an ultra-specific level for surgical simulation and anatomical education*, *Clinical Anatomy* **26**, 333 (2013).
- [2] B. Preim and C. P. Botha, *Visual Computing for Medicine: Theory, Algorithms, and Applications* (Newnes, 2013).

ACKNOWLEDGEMENTS

Doing a PhD felt like traveling a path for me. There are people who led me to this path, people without whom I would have never even begun to do a PhD. Then there were the people that guided and supported me along this path. And now, at the end of the PhD path, there are those who are helping me find a new path to continue my journey. Even though it is impossible for me to name all these people in these two pages, in this part of the thesis I would like to try to do so anyway. If I forgot you and you are reading this now, I am so sorry, thanks for reading though!

I would not have been able to finish this PhD without my excellent supervisory team, who guided my path, even when it was unclear for me where to go next. I started this PhD under the supervision of Charl Botha, and I can't thank him enough for being a scientific inspiration and endless source of enthusiasm. Dear Charl, unfortunately for me, you left academia, but I am grateful that I had the pleasure of working with you pre- and during the start of my PhD, also for the many sugarfixes and laughs we shared, 'we bloggen'! Shortly after Charl left, I was fortunate that Anna Vilanova took over as my daily supervisor and copromotor. I couldn't have wished for a better successor. Dear Anna, thank you so much, for your excellent supervision and guidance, but also for caring and understanding on a personal level. I will also never forget your great stories on Catalan Christmas traditions! Due to the interdisciplinary nature of my PhD project, I got to have the best promotors of both worlds, Elmar Eisemann in Delft and Marco de Ruiter in Leiden. Dear Marco, thank you for your supervision from the medical side. It was incredible to work with someone who is so passionate and excited about anatomical research, and also manages to spark this passion in others, such as myself. Dear Elmar, it was amazing to have you as my promotor, not only because of your excellence in research and sound advice, but also for all the nice discussions in the breaks on a wide range of topics, including the math-jokes! Thank you all!

I would like to thank my independent committee members, for accepting our invitation, and providing valuable feedback on my thesis. Special thanks to Erik Jansen, for being my first promotor, and providing extensive suggestions to strengthen this thesis. I would like to thank all collaborators and people that contributed to the papers that form the basis for this thesis, be it as co-authors or participants in evaluations.

My PhD project was funded by the STW, and I am thankful to them, our user committee, and especially to Shirley Baert, for their support. In this project, we worked closely together with amazing researchers at the Leiden University Medical Center, thanks to you all for this fruitful collaboration. Annelot Kraima, it was a great to have you as my fabulous PhD-project other half, and I can't recall a single shared presentation that was not absolutely rocked by team pelvis.

I would also like to extend my gratitude to the TU Delft Graduate School. Even though the Graduate School was only recently established when I started my PhD, I have benefited much from the services they provide.

The Computer Graphics and Visualization group in Delft, so many to thank, so little space! Thanks to you all for making this group such a fun place to work in! Some people I already thanked in other sections, but who is counting? Some of you were around when I started, thank you Erik, Frits, Peter van Nieuwenhuizen, Charl, Gerwin, Ruud, Bart, Thomas Kroes (my bf! Let's have a kiki! HW?), Peter Kok (.com), Ricardo, Ruben, Tim, Rafa, Francois, and Christian, for making it a great start for me in Delft. Others joined the group after I started, Elmar, Anna, Klaus, Changgong (my little brother!), Niels (Here come dat boi! HW?), Bert, Jean-Marc, Matthias, Jingtang (bedankt voor alles!), Timothy, Ben, Renata, Sergio, Pedro, Christopher, Pablo, Thomas Höllt, Nicola, Nestor, Leonardo, Stefanie, and Sandra, thank you all for all the good times we shared over the years! Berend and Cees-Willem, thank you for your awesome work and helping me, you have achieved so much in both your master theses, super proud of you guys! Outside the CVG group, but inside the TU Delft, I also had many a great coffeebreak with people from other groups. Raynor, thanks for all the good game discussions and laughs, Zeki, always great talking to and dancing with you, and Marco Loog, thanks for the great conversations we had, especially after one or more beers! So many people, so many good memories!

Besides local awesome people, due to the amazing visualization research community, I have had the pleasure of meeting many international awesome people at conferences or during research visits. Some of these meetings even lead to awesome collaborations! So a big thank you to my international co-authors: Bernhard Preim, Katja Bühler, Nils Lichtenberg, Christian Hansen, and Stefan Bruckner. Special thanks go out to Kai Lawonn (HW?), with whom I now have had awesome collaborations with on multiple occasions, be it on medvis papers, medvis.org, organizing medvis parties, or even medviz image contests; team medvis 4ever!

I would like to thank my current co-workers, the lovely visualization group at the University of Bergen, Norway, for welcoming me into their group and providing me with excellent new opportunities and a great start of my post-PhD life. In no particular order: Matthias, Jan, Chaoran, Sergej, Andreas, Ivo, Erlend, Åsmund, Július, Veronika, Helwig, and Stefan, thank you all for the very warm welcome to a very rainy, yet beautiful, city!

Finally, I would like to thank my friends and family, for always being there and encouraging me. My mother, besides doing all the things good mothers do, always managed to cheer me up when yet another paper, or even the same paper multiple times, got rejected by some mean reviewers. Thanks to my gaming father, I can now look totally cool (in some contexts at least) talking about childhood adventures on the Commodore 64. And then of course there were my first modding adventures with my brother Joeri in the original Doom, awesome times! You were all instrumental in piquing my interest in computer science at an early age already.

I'm saving the person I am most thankful to for last... I would like to thank my partner Gerrit. Without you and your encouragement and trust, I wouldn't have even dared to make the switch from being a radiographer to studying computer science, and definitely not doing this whole PhD thing. The PhD process includes some serious tests of endurance and emotional resilience, which would have been impossible for me to face alone, so thank you wholeheartedly for your endless support and patience.

CURRICULUM VITÆ

NOESKA Smit was born December 3rd, 1983, in Heenvliet, the Netherlands. After high-school, she worked and studied as a radiographer at the MCRZ in Rotterdam, the Netherlands. After obtaining her radiographer degree in 2005, she went on to study Computer Science at the Delft University of Technology in Delft, The Netherlands. During her studies she specialized in Computer Graphics and more specifically, Medical Visualization, and obtained her Master's degree in 2012.

After this, she continued as a PhD candidate in the Computer Graphics and Visualization group of the Electrical Engineering, Mathematics and Computer Science Faculty of the Delft University of Technology. During her research, which resulted in this thesis, she became active as an editor and contributor for medvis.org, the community blog on medical visualization research and its applications in practice. During her PhD, she visited several research groups in Bergen, Koblenz, and Magdeburg as a guest researcher. She supervised two master student theses and performed educational tasks for visualization courses. The Virtual Surgical Pelvis project won the ESSO Best Video 2nd Prize at the 34th Congress of the European Society of Surgical Oncology in 2014. Furthermore, PelVis was one of the winners of the MedViz image contest 2016, awarded at the joint MedViz conference and the Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM).

Noeska is currently working as a researcher in the Visualization group at the Department of Informatics in the University of Bergen, Norway.

LIST OF PUBLICATIONS


JOURNAL & REVIEWED CONFERENCE PUBLICATIONS

10. **Noeska Smit**, Kai Lawonn, Annelot Kraima, Marco DeRuiter, Hessam Sokooti, Stefan Brückner, Elmar Eisemann, and Anna Vilanova, "PelVis: Atlas-based Surgical Planning for Oncological Pelvic Surgery", Accepted, to appear, *IEEE Transactions on Visualization and Computer Graphics (Proceedings of the Scientific Visualization 2016)*, 2017.
9. Nils Lichtenberg, **Noeska Smit**, Christian Hansen, and Kai Lawonn, "Sline: Seamless Line Illustration for Interactive Biomedical Visualization", *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, pp. 133-142. 2016.
8. Annelot Kraima, Marloes Derks, **Noeska Smit**, Jacoba van Munsteren, Cornelis van de Velde, Gemma Kenter, and Marco DeRuiter, "Careful dissection of the distal ureter is highly important in nerve-sparing radical pelvic surgery: a 3D reconstruction and immunohistochemical characterization of the vesical plexus", *International Journal of Gynecological Cancer* 26 no. 5 (2016): 959–966.
7. Kai Lawonn, **Noeska Smit**, Bernhard Preim, and Anna Vilanova, "Illustrative Multi-volume Rendering for PET/CT Scans", *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, pp. 103-112. 2015.
6. Marjolein Kleppe, Annelot Kraima, Roy Kruitwagen, Toon Van Gorp, **Noeska Smit**, Jacoba van Munsteren, and Marco DeRuiter, "Understanding Lymphatic Drainage Pathways of the Ovaries to Predict Sites for Sentinel Nodes in Ovarian Cancer", *International Journal of Gynecological Cancer* 25, no. 8 (2015): 1405-1414.
5. **Noeska Smit**, Berend Klein Haneveld, Marius Staring, Elmar Eisemann, Charl Botha, and Anna Vilanova, "RegistrationShop: An Interactive 3D Medical Volume Registration System", *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, pp. 145-153. 2014.
4. Annelot Kraima, Marloes Derks, **Noeska Smit**, Conny Van Munsteren, Jacobus van der Velden, Gemma Kenter, and Marco DeRuiter, "Lymphatic drainage pathways from the cervix uteri: Implications for radical hysterectomy?", *Gynecologic oncology* 132, no. 1 (2014): 107-113.
3. Annelot Kraima, **Noeska Smit**, Daniel Jansma, Chris Wallner, Ronald Bleys, Cornelis van de Velde, Charl Botha, and Marco DeRuiter, "Toward a highly-detailed 3D pelvic model: Approaching an ultra-specific level for surgical simulation and anatomical education", *Clinical Anatomy* 26, no. 3 (2013): 333-338.
2. **Noeska Smit**, Annelot Kraima, Daniel Jansma, Marco DeRuiter, and Charl Botha, "The unified anatomical human (beta): Model-based representation of heterogeneous anatomical data", *Proceedings of the 3D Physiological Human (3DPH) Workshop, CASA*. 2012.

1. Anton Kerver, Gert-Jan Kleinrensink, **Noeska Smit**, Sverre Rabbelier, Ben Sedee, and Charl Botha, "Web-Based 'Computer Assisted Surgical Anatomy Mapping'", *Proceedings of the International Conference on Web Information Systems and Technologies (WEBIST)*, pp. 244-247. 2010.

SHORT PAPERS & POSTERS

5. **Noeska Smit**, Annelot Kraima, Daniel Jansma, Marco DeRuiter, Elmar Eisemann, and Anna Vilanova, "VarVis: Visualizing Anatomical Variation in Branching Structures", *Proceedings of EuroVis Short Papers*. 2016.
4. **Noeska Smit**, Cees-Willem Hofstede, Annelot Kraima, Daniel Jansma, Marco Deruiter, Elmar Eisemann, and Anna Vilanova, "The Online Anatomical Human: Web-based Anatomy Education", *Proceedings of EuroGraphics Education Papers*. 2016.
3. **Noeska Smit**, Cees-Willem Hofstede, Annelot Kraima, Daniel Jansma, Marco Deruiter, Elmar Eisemann, Charl Botha and Anna Vilanova, "The Online Anatomical Human: Anatomical Knowledge Exchange on the Web", *Poster Presentation IEEE Visualization*. 2014.
2. **Noeska Smit**, Thomas Kroes, Annelot Kraima, Daniel Jansma, Marco DeRuiter, Edward Valstar, and Charl Botha, "Towards Declarative Medical Visualization", *Poster Presentation EG VCBM (Eurographics Workshop on Visual Computing for Biology and Medicine)*. 2012.
1. **Noeska Smit**, Annelot Kraima, Marco DeRuiter, Daniel Jansma and Charl Botha, "A Unified Representation for the Model-based Visualization of Heterogeneous Anatomy Data", *Proceedings of EuroVis Short Papers*. 2012.



ISBN 978-94-6186-719-3



9 789461 867193 >