



Thèse

2016

Open Access

This version of the publication is provided by the author(s) and made available in accordance with the copyright holder(s).

---

## Efficient extraction of musculoskeletal structures from multi-channel MR images

---

Becker, Matthias

### How to cite

BECKER, Matthias. Efficient extraction of musculoskeletal structures from multi-channel MR images. Doctoral Thesis, 2016. doi: 10.13097/archive-ouverte/unige:89935

This publication URL: <https://archive-ouverte.unige.ch/unige:89935>

Publication DOI: [10.13097/archive-ouverte/unige:89935](https://doi.org/10.13097/archive-ouverte/unige:89935)

UNIVERSITÉ DE GENÈVE

Département d'informatique

FACULTÉ DES SCIENCES

Professeur José Rolim

FACULTÉ D'ÉCONOMIE  
ET DE MANAGEMENT

Institut de Science de Service Informationnel

Professeur Nadia Magnenat Thalmann

---

# **Efficient Extraction of Musculoskeletal Structures from Multi-Channel MR Images**

**THÈSE**

présentée à la Faculté des sciences de l'Université de Genève  
pour obtenir le grade de Docteur ès sciences, mention informatique

par

**Matthias Kai Holger BECKER**

de

Hanovre (Allemagne)

Thèse N° 4994

GENÈVE

Atelier d'impression ReproMail

2016





**UNIVERSITÉ  
DE GENÈVE**

FACULTÉ DES SCIENCES

**Doctorat ès sciences  
Mention informatique**

Thèse de *Monsieur Matthias BECKER*

intitulée :

**"Efficient Extraction of Musculoskeletal Structures from  
Multi-Channel MR Images"**

La Faculté des sciences, sur le préavis de Monsieur J. ROLIM, professeur ordinaire et directeur de thèse (Département d'informatique), Madame N. MAGNENAT THALMANN, professeure honoraire et codirectrice de thèse (Faculté d'économie et de Management), Madame E. ANGELINI, professeure assistante (Département traitement du signal des images, Institut Mines Telecom, Telecom Paris Tech, France et Institute for Translational Medicine and Therapeutics, Department of Surgery & Cancer, Faculty of Medicine, Imperial College London, United Kingdom) et Monsieur E. STINDEL, professeur (Laboratoire de traitement de l'information médicale, Centre hospitalier universitaire Morvan, Brest, France), autorise l'impression de la présente thèse, sans exprimer d'opinion sur les propositions qui y sont énoncées.

Genève, le 29 septembre 2016

**Thèse - 4994 -**

**Le Doyen**





---

# ACKNOWLEDGMENTS

---

This thesis was only possible with the help, input and support from many people. I would like to deeply thank Prof. Nadia Magnenat Thalmann for giving me the opportunity to work at MIRALab. Her insight and permanent support gave me the chance to pursue my research and to learn in an interdisciplinary environment. I would also like to express my thanks to the members of my jury for their time and effort in reviewing my manuscript: Prof. José Rolim (University of Geneva, Switzerland), Prof. Elsa Angelini (Telecom ParisTech), and Prof. Eric Stindel (LaTIM, Brest, France).

I would like to thank my current and former colleagues at MIRALab for the great time, collaborations and insightful discussions. In particular, I want to acknowledge Andra Chincisan and Dr. Hon Fai Choi, who went the way together with me and always have shown great support. I am grateful to have had the opportunity to work and publish together with Dr. Niels Nijdam and Dr. Lazhari Assassi. Many thanks to the other MIRALabians, current and former, to make my stay rewarding and pleasant: Yvain Tisserand, Simon Sénécal, Louis Cuel, Rui Antunes, Bart Kevelham and Maher Ben Moussa. Warm thanks to the designers, Marlène Arévalo-Poizat and Nedjma Cadi and to Lara Broi for the support in all administrative tasks. I would also like to thank Jérôme Schmid for his guidance at the start of my work.

The presented work was supported by the EU FP7 Marie Curie ITN MultiScaleHuman (grant ITN-2011-289897), which gave me the possibility to work with great teams from Germany, Italy, Portugal and Switzerland. I am thankful for the unique opportunity to get to know the ESRs and ERs, we have made a great team and will hopefully continue to work together. I am also grateful to the scientist in charge of the project partners and everyone else who made the meetings and secondments great stays. Many thanks to the University Hospital of Geneva and the group of Prof. Osman Ratib for the close collaboration in the image acquisition. In particular, I would like to acknowledge Dr. Bénédicte Delattre whose invaluable effort lead to the image data used in this work. Also, I would like to thank the volunteers in our study.

Finally, last but not least, I would like to thank my friends, for their unconditional support, encouragement and patience, and especially my family, my mother, my father and my brother.

**Matthias Becker**



---

# ABSTRACT

---

In a society with a growing share of elderly, musculoskeletal diseases (MSDs) have become a major burden not only for those affected but also for the health care system and have been found to be the main reason for work time loss in the European Union. MSDs are a group of disorders, including osteoarthritis, which affects body parts like bones, muscles, tendons and joints. They cause pain and reduce mobility and flexibility. A better understanding of these diseases can be achieved by acquiring knowledge of the patient-specific morphology of the involved anatomical structures. Magnetic resonance imaging (MRI) is a non-invasive technique for acquiring volumetric tomographic images. It allows the determination of anatomical parameters like volume, cross-sectional area, or pennation angle.

MRI, like other medical imaging techniques, is subject to constant improvement. This concerns, among others, image quality, resolution, acquisition duration and costs. It has led to an ever-growing number of acquisition with higher resolutions, resulting in big numbers of large data sets. The interpretation of MRI data is a lengthy process that requires special training and knowledge about anatomy and imaging properties. Depending on the acquisition parameters, analysing the MR images, e.g. to identify organs or pathological tissue, is a complex and time-consuming task. With higher resolutions, an increase in the total number of scans, and a reduction of medical personnel, there is less and less time left for image interpretation. This requires automated processing of the data to assist the medical personnel. This can be, for example, in the form of the detection of potentially pathological tissue or the extraction organs and other anatomical structures. Some MRI protocols can generate multi-channel images, which can highlight different anatomical structures in different channels.

We focus on the extraction of musculoskeletal structures (bones, muscles) from the lower limb. In our work, we have focused on three main objectives: development of an MRI protocol, processing and labelling of the image data, and the exploitation of multi-channel data during the segmentation of individual muscles. All these tasks have been addressed with a focus on efficiency.

We propose an MR acquisition protocol that generates seamless, high-resolution images of thigh and calf. Our setup goes beyond clinical routine and uses a special coil for signal amplification. We have developed a coil support structure that allows us to move the coil along the immobilised patient without touching. This prevents artefacts from movements between acquisitions.

The acquired image data is then processed, following our pipeline. This includes stitching, noise

---

and bias field reduction and the identification of different classes of image content. We present a method to identify air in the image. As the MR scanner has a big opening, large parts of the image is empty, showing only air. Using our air labels, we can crop the image and reduce the file sizes. This is particularly important for multi-channel images. An additional method allows the fast and reliable identification of muscle tissue. We propose a massively parallel version of these techniques.

Once the image data is acquired and processed, we identify the individual muscle. Using the labels and prior knowledge, we align a muscle template. This template is modified with our deformable model framework under the influence of image forces to match the actual anatomy. This is an iterative process and the image forces use the multi-channel image data to find relevant image features. We propose coupled multi-resolution deformable models that allow working on different resolutions in parallel.

This work is aimed at helping the understanding of diseases by providing personalised anatomical models. There is a large number of applications; we show a contribution towards the digital patient and the determination of personalised pennation angles.

---

# RÉSUMÉ

---

Dans une société où le nombre de seniors est en constante augmentation, les troubles musculo-squelettiques (TMS) sont devenus une charge considérable, non seulement pour ceux qui sont affectés mais aussi pour le système de santé. Il a été découvert que c'était la raison principale de perte de temps de travail dans l'Union Européenne. Les TMS sont un groupe de troubles, comme l'arthrose, qui affecte des parties du corps humain, comme les os, les muscles, les tendons ou les articulations. Ils causent des douleurs ainsi qu'une réduction de mobilité et de flexibilité. Une meilleure compréhension de ces maladies peut être accomplie en acquérant des connaissances morphologiques spécifiques de la structure anatomique du patient. L'imagerie par résonance magnétique (IRM) est une technique non-invasive qui permet d'acquérir des images volumétriques tomographiques. Cela permet de déterminer les paramètres anatomiques tels que le volume, les sections transversales ou l'angle de pennation.

L'IRM, tout comme d'autres techniques d'imagerie médicale, est sujet à une amélioration constante. Cela concerne, entre autre, la qualité d'image et sa résolution, le temps d'acquisition ainsi que son coût. Cela a conduit à un nombre croissant d'acquisitions ayant une grande résolution, ayant pour résultat l'augmentation des volumes de données. L'interprétation des données de l'IRM est un long procédé qui nécessite une formation particulière ainsi que des connaissances en anatomie et en imagerie. En fonction des paramètres d'acquisition, analyser les images IRM, par exemple identifier les organes ou les tissus pathologique, est une tâche complexe qui nécessite beaucoup de temps. Avec une résolution plus élevée, l'augmentation du nombre total de scans, ainsi que la réduction du personnel médical, il y a de moins en moins de temps disponible pour l'interprétation des images. Il est nécessaire d'automatiser le traitement des données afin d'assister le personnel médical. Cela peut être, par exemple, en détectant des potentielles pathologies des tissus ou en segmentant les organes ou d'autres structures anatomiques. Certains protocoles IRM peuvent générer des images à plusieurs canaux, ce qui permet de mettre en évidence différentes structures anatomiques à partir de plusieurs canaux.

Nous nous focalisons sur l'extraction de la structure musculo-squelettique (os, muscles) des membres inférieurs. Dans notre travail, nous nous sommes concentrés sur trois objectifs principaux : développer un protocole IRM, traiter et annoter les images, et d'exploiter les données multicanaux lors de la segmentation des muscles. Toutes ces tâches ont été traitées en se concentrant sur l'efficacité.

Nous proposons un protocole d'acquisition par résonance magnétique qui permet de générer

---

des images haute résolution, sans discontinuités, pour la cuisse et le mollet. L'installation va au-delà des habitudes cliniques et utilise une bobine spéciale afin d'amplifier le signal. Nous avons développé un support de bobine qui nous permet de bouger la bobine le long du patient immobilisé sans le toucher. Cela permet de prévenir certains artefacts dus aux mouvements entre les acquisitions.

Les images acquises sont ensuite traitées, en suivant notre pipeline. Cela comprend la fusion d'images, la réduction du bruit, la correction du biais et l'identification des différentes catégories de contenu de l'image. Nous présentons une méthode afin d'identifier l'air dans une image. Comme les scanners à résonnance magnétique ont une grande ouverture, une large partie de l'image est vide, montrant uniquement de l'air. En utilisant notre méthode de marquage de l'air, nous pouvons rogner l'image et réduire la taille des fichiers. C'est particulièrement important pour les images multicanaux. Une méthode supplémentaire permet une identification rapide et précise des tissus musculaires. Nous proposons une version extrêmement parallélisée de cette technique.

Une fois que l'image a été acquise et traité, nous identifions les différents muscles. En utilisant les annotations et les précédentes connaissances, nous alignons un modèle générique de muscles. Ce modèle a été modifié avec notre système de déformation sous l'influence d'images alignées par contrainte sur l'anatomie réelle. C'est un procédé itératif dans lequel la contrainte imposée aux images est basé sur la recherche de caractéristiques pertinentes sur des images multicanaux. Nous proposons d'utiliser des modèles déformables multi-résolution couplés permettant de travailler en parallèle sur plusieurs résolutions.

Ce travail a pour but d'aider la compréhension des pathologies en fournissant des modèles anatomiques personnalisés. Il y a un grand nombre d'applications, nous montrons notre contribution dans le domaine des patients digital et dans la détermination d'un angle de pennation personnalisé.

---

# CONTENTS

---

<b>Acknowledgments</b>	<b>i</b>
<b>Abstract</b>	<b>iii</b>
<b>Résumé</b>	<b>v</b>
<b>Table of Contents</b>	<b>vii</b>
<b>List of Figures</b>	<b>xi</b>
<b>List of Tables</b>	<b>xiv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	2
1.2 Context . . . . .	2
1.3 Objectives . . . . .	3
1.4 Contributions . . . . .	4
1.5 Thesis Organisation . . . . .	5
<b>2 Related Work</b>	<b>9</b>
2.1 Motivation . . . . .	10
2.2 Medical Imaging . . . . .	10
2.2.1 X-Ray Imaging . . . . .	10
2.2.1.1 Radiography . . . . .	11
2.2.1.2 Computed Tomography . . . . .	11
2.2.2 Magnetic Resonance Imaging . . . . .	13
2.2.2.1 History . . . . .	13
2.2.2.2 Working Principle . . . . .	13
2.2.2.3 Artefacts . . . . .	14
2.2.2.4 Protocols . . . . .	16
2.2.3 Other Methods . . . . .	17
2.2.4 Comparison . . . . .	19
2.3 General Segmentation Approaches . . . . .	20
2.3.1 Pixel- and Histogram-based Methods . . . . .	21
2.3.2 Region-based Methods . . . . .	22
2.3.3 Edge-based Methods . . . . .	22
2.3.4 Texture-based Methods . . . . .	23
2.3.5 Comparison . . . . .	23
2.4 Morphology-based Methods . . . . .	24



2.5	Classification . . . . .	25
2.6	Registration . . . . .	26
2.7	Graph Cuts . . . . .	27
2.8	Comparison . . . . .	28
2.9	Model-based Segmentation . . . . .	29
2.9.1	Snakes . . . . .	30
2.9.2	Level Sets . . . . .	32
2.9.2.1	Extensions . . . . .	33
2.9.3	Discrete Deformable Models . . . . .	34
2.9.3.1	Numerical Implementation . . . . .	35
2.9.3.2	Mesh Types . . . . .	36
2.9.3.3	Particle Systems . . . . .	36
2.9.4	Knowledge-based Deformable Models . . . . .	37
2.9.4.1	Feature Selection . . . . .	37
2.9.4.2	Construction . . . . .	37
2.9.4.3	Active Shape Models and Active Appearance Models . . . . .	38
2.9.4.4	Conclusion . . . . .	39
2.9.5	External Forces . . . . .	39
2.9.5.1	Basic Forces . . . . .	39
2.9.5.2	Advanced Forces . . . . .	40
2.9.5.3	Interactive Forces . . . . .	42
2.9.6	Conclusion . . . . .	43
2.10	Multi-Channel Image Processing . . . . .	43
2.11	Muscles of Thigh and Calf . . . . .	45
2.11.1	Muscles of Thigh . . . . .	46
2.11.2	Muscles of Calf . . . . .	47
2.12	Muscle Segmentation . . . . .	48
2.13	Conclusion . . . . .	50
<b>3</b>	<b>Image Acquisition</b> . . . . .	<b>53</b>
3.1	Motivation . . . . .	54
3.2	Image Acquisition . . . . .	54
3.2.1	MRI Setup . . . . .	55
3.2.2	Coils . . . . .	55
3.2.2.1	Posterior Coil . . . . .	56
3.2.2.2	Anterior Coil . . . . .	57
3.2.3	Sequences . . . . .	57
3.2.4	Challenges in the Image Setup . . . . .	58
3.3	Image Data . . . . .	59
3.4	Processing . . . . .	59
3.4.1	Stitching . . . . .	60
3.4.2	Air Filtering . . . . .	62
3.5	Evaluation of the Seamless Stitching and Processing . . . . .	63
3.6	Conclusion . . . . .	64
<b>4</b>	<b>Image Processing</b> . . . . .	<b>67</b>
4.1	Motivation . . . . .	68

4.2	Intensity Correction . . . . .	68
4.2.1	Inter-stack Correction . . . . .	68
4.2.2	Bias Field Correction . . . . .	69
4.3	Denoising . . . . .	72
4.4	Muscle Labelling . . . . .	74
4.4.1	Evaluation and Results . . . . .	75
4.5	Acceleration through Massively-Parallel Image Processing . . . . .	79
4.5.1	Platform Selection . . . . .	79
4.5.2	GPGPU Architecture . . . . .	80
4.5.3	Pixel-wise Operations . . . . .	81
4.5.4	Morphology Operations . . . . .	81
4.5.5	Connected Component Labelling . . . . .	83
4.5.6	Slice Histogram Generation . . . . .	85
4.5.7	Evaluation and Results . . . . .	85
4.6	Conclusion . . . . .	88
<b>5</b>	<b>Model Initialisation</b>	<b>89</b>
5.1	Motivation . . . . .	90
5.2	Bone Localisation . . . . .	90
5.2.1	Manual Initialisation . . . . .	90
5.2.2	Landmarks . . . . .	90
5.2.3	Edges and Image Features . . . . .	91
5.2.4	General Hough Transform . . . . .	93
5.3	Histogram Matching for Joint Localisation . . . . .	94
5.3.1	Evaluation and Results . . . . .	96
5.3.2	Bone Placement . . . . .	98
5.4	Registration of Muscle Atlas using Bone Location Information . . . . .	99
5.5	Label-based Muscle Mesh Correction . . . . .	100
5.5.1	Evaluation and Results . . . . .	102
5.6	Conclusion . . . . .	106
<b>6</b>	<b>Image Segmentation</b>	<b>107</b>
6.1	Motivation . . . . .	108
6.2	Deformable Model Framework Implementation . . . . .	108
6.2.1	Main Components of Deformable Models . . . . .	108
6.2.2	Structure and Components of the DMFramework . . . . .	109
6.2.2.1	Meshes . . . . .	109
6.2.2.2	Forces . . . . .	110
6.2.2.3	Image handling . . . . .	112
6.2.2.4	Physical simulation . . . . .	114
6.2.2.5	Output . . . . .	115
6.2.2.6	Implementation . . . . .	116
6.2.2.7	Parallelisation . . . . .	117
6.3	Evaluation . . . . .	118
6.3.1	Manual Segmentation . . . . .	118
6.3.1.1	Conclusions on Manual Segmentation Tools . . . . .	123
6.3.1.2	Post-processing of the Segmented Data . . . . .	123

6.3.1.3	Required Time . . . . .	124
6.3.2	Label-based Evaluation Metrics . . . . .	125
6.3.3	Results . . . . .	125
6.4	Webviewer for Remote and Multiplatform Viewing and Interaction . . . . .	127
6.4.1	Webviewer . . . . .	127
6.4.2	Model Streaming . . . . .	128
6.5	Multi-Resolution Mesh Coupling . . . . .	130
6.5.1	Concept of Coupled Meshes . . . . .	131
6.5.2	Coupling . . . . .	132
6.5.3	Synchronisation . . . . .	133
6.5.4	Interactive Segmentation . . . . .	134
6.5.5	Implementation . . . . .	135
6.5.5.1	Per Step Update Strategy (PSU) . . . . .	136
6.5.5.2	Scaled Updates Strategy (SU) . . . . .	136
6.5.5.3	Partial Update Strategy (PU) . . . . .	137
6.5.6	Evaluation and Results . . . . .	138
6.5.6.1	Deformable Model Configuration . . . . .	138
6.5.6.2	Synchronisation Evaluation . . . . .	138
6.5.6.3	Discussion . . . . .	139
6.6	Conclusion . . . . .	142
<b>7</b>	<b>Conclusions</b>	<b>143</b>
7.1	Motivation . . . . .	144
7.2	Applications . . . . .	144
7.2.1	Digital Patients . . . . .	144
7.2.2	Pennation Angle Measurements . . . . .	144
7.3	Discussion . . . . .	146
7.4	Contributions and Benefits . . . . .	146
7.5	Limitations and Future Work . . . . .	148
<b>A</b>	<b>List of Publications</b>	<b>151</b>
<b>B</b>	<b>Acronyms, Glossary and Technical Details</b>	<b>153</b>
B.1	Acronyms . . . . .	153
<b>C</b>	<b>Medical Appendix</b>	<b>156</b>
C.1	Medical Nomenclature . . . . .	156
C.2	Anatomy . . . . .	158
C.3	MRI information sheets . . . . .	164
C.4	MRI consent form . . . . .	167
<b>D</b>	<b>Technical Appendix</b>	<b>168</b>
D.1	Other Deformable Models . . . . .	168
D.2	Functionality exposed to Lua scripting . . . . .	169
	<b>Bibliography</b>	<b>173</b>

---

# LIST OF FIGURES

---

1.1	MRI usage in the USA according to OECD data . . . . .	3
1.2	Main objectives . . . . .	4
1.3	Main steps of the contribution . . . . .	5
1.4	Example of the multi-channel images: mDixon and T1 TSE . . . . .	5
1.5	Segmentation pipeline: From acquisition to segmentation . . . . .	6
2.1	Working principle of X-ray imaging and example image . . . . .	11
2.2	CT scanner and scan . . . . .	12
2.3	Motion artefact in MRI data . . . . .	15
2.4	Flow artefact in MRI data . . . . .	15
2.5	Inhomogeneities artefact in MRI data . . . . .	16
2.6	Ghosting artefact in MRI data . . . . .	16
2.7	Cross-section of the thigh using the Dixon protocol . . . . .	18
2.8	Thresholding segmentation principle . . . . .	21
2.9	Watershed segmentation of the Visible Human data set . . . . .	22
2.10	Region growing principle in a sample 2D image . . . . .	23
2.11	Segmentation of edges in a sample 2D image . . . . .	23
2.12	Morphological operations in a 2D image . . . . .	24
2.13	Example of classification . . . . .	26
2.14	Principle of registration . . . . .	26
2.15	Working principle of graph cuts . . . . .	27
2.16	Classification of Deformable Models . . . . .	30
2.17	Example of a snake . . . . .	31
2.18	Level set segmentation of the liver . . . . .	33
2.19	Example of an edge stopping level set . . . . .	34
2.20	Process of the iterative deformation of a surface using forces . . . . .	35
2.21	Triangle and simplex meshes . . . . .	36
2.22	Combination of pressure force and smoothing . . . . .	40
2.23	Example of intensity profiles sampled in an MR image of the femur . . . . .	41
2.24	Structure of Muscles . . . . .	46
2.25	Muscles of the thigh . . . . .	47
2.26	Muscles of the calf . . . . .	48
3.1	Dimensions and layout of the scanner board . . . . .	56
3.2	Illustration of the placement of the coil above and below the subject . . . . .	56
3.3	Board with blocks and coils installed on the scanner table . . . . .	57
3.4	Different modalities: mDixon and T1 TSE . . . . .	58

3.5	Sub-parts of the image acquisition and corresponding duration . . . . .	58
3.6	Preprocessing pipeline . . . . .	61
3.7	Stitching of the MRI stacks . . . . .	62
3.8	Different interpolation approaches between stacks . . . . .	62
3.9	Air masking process . . . . .	63
3.10	Artefact caused by subject motion between stacks . . . . .	64
3.11	Image data set before and after bias field correction and air labelling . . . . .	65
4.1	Slice-wise inter-stack correction of intensities . . . . .	69
4.2	Bias field correction using the N3 algorithm . . . . .	70
4.3	Bias field correction evaluation . . . . .	72
4.4	Noise removal in MR images . . . . .	73
4.5	Noise reduction evaluation . . . . .	74
4.6	Muscle detection process . . . . .	75
4.7	Evaluation of the muscle labelling . . . . .	77
4.8	Applications of the muscle labelling . . . . .	78
4.9	GPU structure with grid, blocks and threads . . . . .	81
4.10	Iterative process of merging labelled blocks . . . . .	84
4.11	Air processing times on CPU and GPGPUs . . . . .	87
4.12	Muscle processing times on CPU and GPGPUs . . . . .	88
5.1	SURF features and edges before filtering . . . . .	91
5.2	Influence of the radius in image smoothing filters . . . . .	92
5.3	Effect of multiple smoothing filter runs . . . . .	92
5.4	SURF features and edges in the smoothed image . . . . .	93
5.5	Example of the Hough Transform for lines . . . . .	93
5.6	Histogram of tissue and muscle volume per slice for five subjects . . . . .	94
5.7	Histogram registration error for different parameters . . . . .	97
5.8	Prototype user interface of the histogram matching . . . . .	98
5.9	Alignment of the femur in image and mesh using centroids . . . . .	99
5.10	Evaluation of the femur alignment . . . . .	100
5.11	Registration of volumetric mesh models to the labelled muscle images . . . . .	101
5.12	Shape of femur and corresponding convex hull . . . . .	102
5.13	Muscle mesh correction approach . . . . .	103
5.14	Muscle mesh correction application . . . . .	104
5.15	Visual comparison of the muscle mesh adaption . . . . .	105
5.16	Evaluation of muscle mesh correction for several muscles . . . . .	105
5.17	Impact of the number of rays during muscle mesh correction . . . . .	106
6.1	Rigid air and bones are used to limit muscle model movement . . . . .	111
6.2	Muscle label force working principle . . . . .	111
6.3	Pre-calculation process for the muscle force . . . . .	112
6.4	Final pre-calculation process for the muscle force. . . . .	113
6.5	Visualisation using VTK . . . . .	115
6.6	DMFramework architecture . . . . .	116
6.7	Architecture of the Deformable Model framework . . . . .	118
6.8	User interface of TurtleSeg . . . . .	120

6.9	User interface of Seg3D with a leg data set and several segments . . . . .	121
6.10	User interface of YaDiV . . . . .	122
6.11	Cleanup of overlapping segments . . . . .	124
6.12	Interpolation of missing slices . . . . .	124
6.13	Union, Intersection and Relative Complement for two segments . . . . .	125
6.14	Webviewer . . . . .	128
6.15	Client-server communication for mesh streaming . . . . .	130
6.16	Concept of the coupled models . . . . .	132
6.17	Comparison of the different synchronisation approaches . . . . .	133
6.18	The effect of the decimation on the model of the femur . . . . .	138
6.19	Timing measurements for the synchronisation approaches . . . . .	140
7.1	Data sources, modelling approaches and simulations for the digital patient . . . . .	145
7.2	Calculation of the pennation angle from MRI-based muscle models . . . . .	145
C.1	Anatomical directional terms . . . . .	157
C.2	Planes through the human body . . . . .	158
C.3	Femur . . . . .	159
C.4	Tibia and Fibula . . . . .	160
C.5	Knee joint with tendons, ligaments, and cartilage . . . . .	161

---

# LIST OF TABLES

---

2.1	Hounsfield unit values for selected substances . . . . .	12
2.2	Comparison of popular imaging techniques by their properties . . . . .	20
2.3	Comparison of general segmentation approaches . . . . .	29
2.4	Classification of multi-channel segmentation approaches . . . . .	45
3.1	Statistical data of the study participants . . . . .	60
3.2	Air filtering results of file size reduction . . . . .	64
4.1	Comparison of the automatic and manual muscle labelling . . . . .	76
4.2	Results for muscle labelling without and with fat reduction technique . . . . .	77
4.3	Muscle labelling comparison with SoA . . . . .	78
4.4	GPGPU specification . . . . .	86
4.5	Air processing times on CPU and GPGPUs . . . . .	86
4.6	Muscle processing times on CPU and GPGPUs . . . . .	87
5.1	Histogram matching results for the knee joint . . . . .	96
5.2	Histogram matching results for the hip joint . . . . .	97
5.3	Histogram matching results for the ankle joint . . . . .	97
5.4	Muscle mesh correction among subjects . . . . .	106
6.1	Comparison of different mesh libraries by features, license, and last update. . . . .	110
6.2	Comparison of medical image segmentation software tools . . . . .	119
6.3	Single channel results for RF, VM, VI, and VL in both subjects. . . . .	126
6.4	Multi-channel combination results for RF, VM, VI, and VL in both subjects. . . . .	126
6.5	Data types and required bandwidth for geometry streaming . . . . .	129
6.6	Lock timing measurements for multi-resolution mesh coupling . . . . .	141

---

## CHAPTER 1

## INTRODUCTION

---





## 1.1 Motivation

Musculoskeletal diseases (MSDs) are a major burden on society. This group of diseases often affects joints and is prevalent in the elderly. Knowledge of the individual morphology of the involved anatomical structures can not only lead to a better diagnosis and treatment but also to a better general understanding of these diseases. A promising approach is an examination that not only relies on one source of information but exploits multi-scale knowledge.


Musculoskeletal diseases are a group of disorders that affect different body parts such as bones, muscles, tendons, and joints. MSDs are often chronic and mostly occur in the elderly population. They cause pain, lead to a reduction of muscle performance, and eventually total incapacity of performing certain tasks. This reduction of mobility and flexibility causes the need for assistance and care.

The World Health Organisation (WHO) therefore has declared the last decade as the decade of bones and joints [CB12]. The health care costs resulting from MSDs are an economic burden, in particular with an ageing population. MSDs have been found to be the main reason of work-time loss in the EU [PTVvH<sup>+</sup>12]. The morphology of anatomical structures that play an major role at the joint level, captured e.g. through medical imaging, can help to explain some origins of MSD. MRI allows the determination of parameters like cross-sectional area, volume, pennation angle, from musculoskeletal structures [Seg07].

Medical imaging technologies have been vastly improving over the last decades. The number of scans increases as well as the resolution of these scans, leading to massive amounts of data generated. The invention of Magnetic Resonance Imaging (MRI) in 1972 gave a way to image inner organs without the hazard of ionising radiation as used by Computed Tomography (CT). The resulting increase of scans can be underlined by numbers from the OECD and is shown in Figure 1.1. The total number of MRI acquisitions has more than tripled from 9.1 million to 35 million over the last eighteen years and there are no signs of an end of this trend. Interpretation of the generated data is a lengthy process that has to be performed by trained professionals or radiologists, which are required to have a lot of experience. The main challenges lie in the image quality that depends on several factors like acquisition duration, acquisition parameters, imaged organs, patient condition, artefacts, noise, and implants. There is a high variety in the resulting images, therefore making it a complex task to identify structures in the images.

## 1.2 Context

MIRALab has been founded in 1989 by Prof. Magnenat-Thalmann with a core research focus on computer graphics and animation (e.g., virtual humans ([MTSC04, LMT08, KMT11]) and cloth simulation [VT00, VMT05, VMTF09]). This strong background has then allowed to pioneer research in anatomical modelling. Long experience in medical modelling and simulation has been gathered ([KBG<sup>+</sup>95, GKBMT96, MTWT98, YCGMMT03, MTYCS04, GMMT06, MTCS08, SKMT11]). The combination of a comprehensive examination of an articulation with detailed



**Figure 1.1** – MRI usage in the United States of America according to OECD data.

anatomical models and the simulation of soft tissues in motion ([Cha08, ACS<sup>+</sup>09, CAVMT09]) were carried out in several successful projects like CHARM [dMMoL<sup>+</sup>], Co-Me [VHM], 3D Anatomical Human (3DAH) [ACV<sup>+</sup>], and MultiScaleHuman (MSH) [HMLM<sup>+</sup>]. Recent research expands to calculation acceleration approaches like GPUs [SIG<sup>+</sup>11] or tools and techniques for collaborative working [HNS<sup>+</sup>10].

MIRALab is coordinating and leading the MultiScaleHuman project. This project gathers seven partners from four European countries to create a novel multi-scale approach to understanding musculoskeletal diseases. Five core biological scales (molecular, cellular, tissue, organ, and behaviour) are being exploited and the resulting knowledge is gathered in a visualisation and knowledge management system that helps physicians to improve their understanding, diagnosis, and treatment of MSDs in physiological human articulations. The project partners yield their experience in fields including molecular imaging, cellular biology, anatomical modelling, motion analysis, biomechanical and soft-tissue simulation, interactive visualisation, haptic interaction, and ontology-based knowledge management. With two hospitals and an industry partner, the MSH project covers the complete process from interacting with patients to providing a proof of concept towards further commercial applications.

### 1.3 Objectives

The main research objective of this thesis is to devise advanced image segmentation techniques to extract and create 3D models of musculoskeletal structures from images. This process includes the development of an image acquisition protocol, the image labelling and the actual segmentation with a focus on the use of efficient methods.

The resulting three main objectives, as shown in Figure 1.2, are based on each other:

**Development of an imaging protocol**


Development of an MR image acquisition protocol to seamlessly cover the full leg. This includes the selection and parametrisation of MR sequences as well as the physical setup of the coils needed. In the last step, the acquired individual and partially overlapping images have to be combined into a continuous data set.

**Advanced tissue labelling**

For the labelling, the goal is the identification of different tissue types (bones, muscles, fat) and surrounding air. This can be used for fast visualisation during the acquisition and the initialisation of segmentation approaches.

**Exploitation of multi-channel image data during the segmentation**

The segmentation process will be realised with deformable models, which can use prior knowledge about the structures. The focus lies on the segmentation of muscles in the human thigh. By using multiple inputs from the acquisition generated by different sequences, the multi-channel image data will be used for robust segmentation. These multi-channel inputs consist of mDixon images that can be decomposed into water, fat, in-phase and opposed-phase channels, and T1 weighted magnitude images.



**Figure 1.2** – The main objectives propose a pipeline from image acquisition and tissue labelling to the exploitation of multi-channel data for segmentation.


## 1.4 Contributions

Following the objectives defined above, we present contributions in three objectives in this thesis. We demonstrate an MR acquisition protocol that generates seamless, high-resolution images of the musculoskeletal system of thigh and calf. We then show how we can rapidly label different tissue types in these images. We then present the use of deformable models for efficient muscle segmentation. An overview of these contributions is shown in Figure 1.3.

In clinical routine, imaging is performed only on areas with pathologies. When full-body acquisitions are performed, they are low-resolution and often skip the lower limbs. All involved organs have to be taken into consideration to study musculoskeletal diseases of the knee. Therefore, the acquisition has to be seamless from ankle to hip. We have developed a setup that allows us to move the signal-amplifying coil while the subject rests immobilised, without contact or skin deformation through compression.


We identify different tissue types (bones, muscles, fat) and surrounding air. High-resolution, multi-channel image data drastically increases the memory consumption. We contribute a me-

thod that uses the air mask to crop the image and reduces the file size by 62 to 72%. Prior identification of tissue types is used to increase segmentation quality and is also used as an input for the initialisation. Our method for identifying muscle tissue is fast and reliable while maintaining correctness of the results.



**Figure 1.3** – Main steps of the contribution: Development of an imaging protocol (a) and processing of the resulting images (b,c), advanced tissue labelling (d) and identification of individual muscles (e).

Once the muscle tissue has been labelled, it has to be separated into the individual muscles. This process is realised with deformable models. The focus lies on the segmentation of muscles in the human leg. By using multiple inputs from the acquisition, coming from different sequences, the multi-channel image data is used for robust segmentation. These multi-channel inputs consist of mDixon images that can be decomposed into water, fat, in-phase and opposed-phase channels, and T1 weighted TSE magnitude images. An example of this multi-channel information can be seen in Figure 1.4. We have developed a framework to perform these tasks. The resulting methods and algorithms are evaluated on their efficiency. We propose a coupled multi-resolution segmentation process for speed-up.



**Figure 1.4** – Example of the multi-channel images: Dixon (in-phase (a), water (b), fat (c), and opposed-phase (d)) and T1 TSE (magnitude) (e).

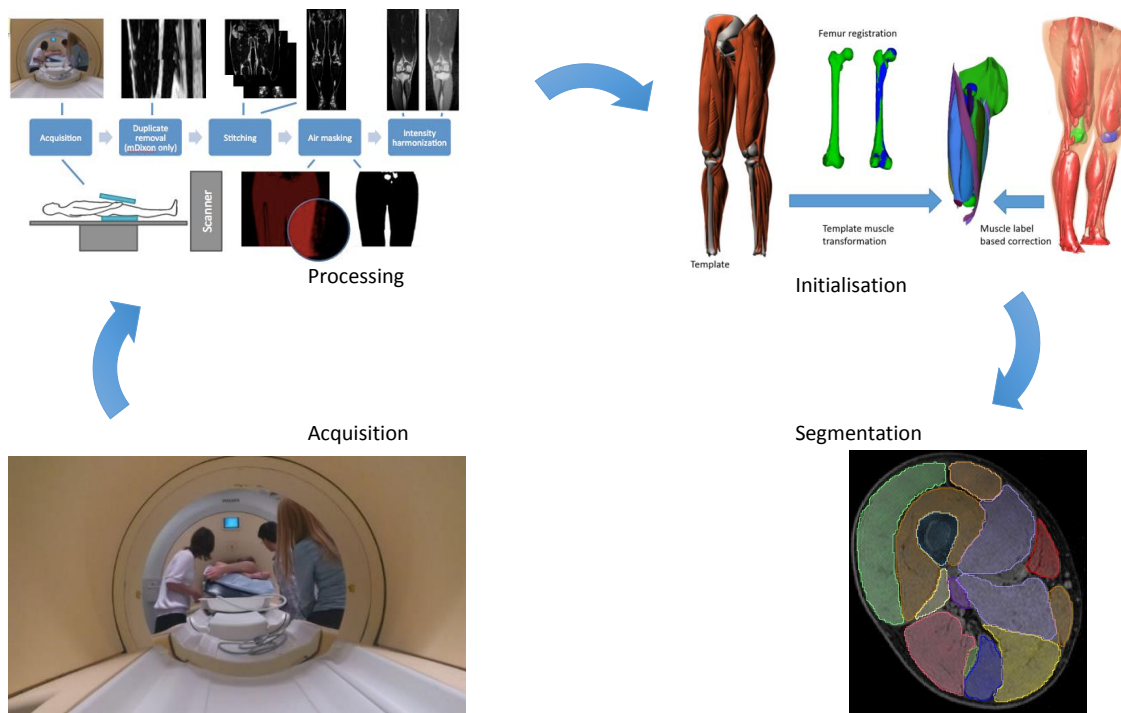
We have compared our results with manual segmentations to assess the accuracy. Medically experienced people have performed the manual segmentation, to enable the comparison of our contributions.

## 1.5 Thesis Organisation

This thesis is organised as follows: This first chapter contains an introduction that describes the motivation for this work. It explains the need for segmentation methods for medical images and the possible effects on the diagnosis of musculoskeletal diseases. We present the objectives and contributions and conclude with an overview of the thesis structure.

The second chapter discusses related work. We give an overview of medical imaging techniques and motivate our choices. Then, we present and compare general segmentation approaches. We show more advanced techniques, including registration and classification. We motivate the need to incorporate model information and present different methods (Snakes, Level Sets, and Deformable Models). We then explore how multi-channel image data can provide additional input for more accurate segmentation. As this thesis is focused on the human lower limb, we provide a short introduction to the affected anatomy and discuss segmentation approaches for muscle tissue. Finally, we conclude by summarising and categorising our findings.

In the following chapters, we present the individual contributions that form the pipeline from acquisition, processing and initialisation to segmentation. This pipeline is shown in Figure 1.5. In each chapter, we first introduce the methodology that we contribute and then present the results.



**Figure 1.5** – Overview of the complete pipeline with the steps from the acquisition, the image processing, the labelling of muscle tissue, registration for initialisation to the deformable model-based segmentation.

The third chapter focusses on the acquisition. We present our MRI setup, the coils and sequences used and discussed the initial processing steps, including stitching and air filtering. We introduce the data sets that will be used throughout this work and evaluate and discuss the acquisition results.

Chapter four discusses our proposed image processing. We show the denoising and intensity correction. We present our muscle tissue labelling approach and its evaluation. Finally, we show how it can be moved to the massively parallel architectures for more efficient processing.

The fifth chapter presents the model initialisation. Placing the models is an important step for the segmentation. We discuss several techniques and then present our approach for localising

joints. Following the localisation, the models need to be adjusted for individual subjects, we demonstrate a technique for that.

Chapter six is dedicated to the segmentation. We introduce our deformable model framework and present its structure and main components. After evaluating the framework, we demonstrate a webviewer for a more flexible visualisation of medical data sets. Lastly, we introduce coupled, multi-resolution meshes that form a parallel version of the pyramid approach to increase efficiency.

The last chapter discusses the results of this thesis. We summarise the presented methods and the achieved results. Two applications of our work are given. We show our contributions and give an outlook on possible extensions and additions.

The annexe provides additional material that had to be omitted in the main text. We provide an overview of medical nomenclature, a more comprehensive anatomy introduction and the information and consent sheets used in our study. We also present additional types of deformable models.




---

## CHAPTER 2

## RELATED WORK

---





## 2.1 Motivation

The extraction of musculoskeletal structures from medical images is a complex task. It can be decomposed into image acquisition and image segmentation.

Medical imaging aims to capture the body optically but also in ways that are imperceptible to the human. It ranges from the acquisition of spectrographic information of the skin to methods to capture the tissues inside of the body. Many applications require information about inner organs (e.g. liver) for diagnostic and therapeutic reasons. The visibility of fractures in X-ray images allows adequate surgery planning and preparation, more advanced MRI acquisition allow the detection of ligament ruptures. There exist different imaging approaches that carry different advantages and disadvantages. We will present and discuss several techniques and motivate our choice.

Medical images gain their value through the interpretation. Depending on the used techniques, the contained information can be very abstract, e.g. in an ultrasound image. While the analysis by radiologists still is the gold standard, the growing number of acquisitions has fostered the development of techniques for (semi-)automatic image analysis. We will give an overview of basic approaches and discuss more advanced approaches afterwards. We will continue by presenting deformable models, which benefit from the model information that they carry.

Some image acquisition techniques allow the generation of multi-channel information. We will discuss how to use this information for more reliable segmentation techniques that are more resistant to image artefacts.

In this thesis, we are looking into the segmentation of musculoskeletal structures, in particular, muscles. We, therefore, give an overview of the muscles of the lower limb in Section 2.11, before we discuss the current state of the art in muscle segmentation in Section 2.12.

## 2.2 Medical Imaging

Medical imaging provides a way to get an inside perspective of the human body without (or with very little) harm. Rather than using invasive methods like exploratory surgery, modern medical imaging provides the ability to detect bone fractures, tumours or metabolic diseases without touching the patient. In this section, we will explain the development from x-ray imaging to MRI (and nuclear medicine). After shortly covering further imaging methods, this section ends with a comparison of the presented techniques.

### 2.2.1 X-Ray Imaging

In this section, we will describe medical imaging methods that use X-rays. These ionising rays have been discovered by Wilhelm Conrad Röntgen<sup>1</sup> in 1895. They are capable of traversing dense

---

1. In 1901, Röntgen received the first Nobel Prize in Physics for this discovery.


structures and allow the imaging of structures beneath the surface. In the following sections, we describe radiography and its extension, Computed Tomography.

### 2.2.1.1 Radiography

In traditional radiography, the body is placed between an X-ray source and a detector. The X-rays are generated at the source, usually through the acceleration of electrons and their sudden deceleration at a target metal. During this deceleration, bremsstrahlung is emitted. If the acceleration is tuned to the correct value, X-rays are submitted as bremsstrahlung.

The X-rays are focused on the area of interest and traverse the body. Inside the body, they are partially absorbed. Therefore, X-ray imaging falls into the category of the ionising approaches. The absorption rates depend on the density of the traversed tissue and hence allow to infer the internal structure. After leaving the body, the rays are caught by a detector. Older systems still use specialised photographic films that need to be developed while newer systems use large arrays of digital photo-detectors. Digital acquisition is on the rise since it allows almost instant access to the results and can be directly integrated into the hospital IT [YR97].

X-ray imaging is commonly used in emergency care because it is very fast. It is good at highlighting bony structures [Lew04] and is often used to detect fractures [BDB<sup>+</sup>01]. An example of an X-ray image of the elbow joint is shown in Figure 2.1.



**Figure 2.1** – Working principle of X-ray imaging (left). The X-rays are created at the source, traverse the body and are absorbed in the detector. The resulting image (right) shows the accumulated density of the traversed tissue (Source: Eric Schmittenmaer, Flickr)

### 2.2.1.2 Computed Tomography

X-ray imaging only can show accumulated densities along a single direction; a 3D reconstruction is not possible. The 3D information is often needed for further diagnosis [HKB<sup>+</sup>06] and better analysis of partially covered organs. Creating several scans repeatedly from different directions has proven not to be practical.

To overcome this limitation, the **Computed Tomography** (CT) [Kal06] uses the Radon transform [Bey87]. A CT scanner and an example scan of the abdomen can be seen in Figure 2.2. During the acquisition, the scanner takes intensity images while rotating around the patient [Doi06]. The

acquisition results in a high number of scans from different angles. These scans are all taken along the transverse plane. The scans can be transformed into a sectional image using the radon transform. If this process is repeated along the superior-inferior axis, it is possible to create a full volumetric 3D reconstruction. These machines have been improved by acquiring 16 or up to 64 parallel planes at the same time. Further improvement allows a spiral shaped capture so that acquisition can be performed in one continuous motion, resulting in less stress for the patient. The **Cone Beam CT (CBCT)** [SF08] pushes the concept even further by using a cone-shaped set of x-rays and a 2D detector.



**Figure 2.2** – CT scanner (left) and resulting scan (right). Source: Images are public domain.

The resulting intensity values can be transformed into Hounsfield units (HUs) [Bro77]. HUs are well-defined and based on the intensity of water. An intensity value  $\mu_x$  can be converted into a HU using the following equation:

$$\text{HU} = 1000 \times \frac{\mu_x - \mu_{\text{water}}}{\mu_{\text{water}}} \quad (2.1)$$

**Table 2.1** – Hounsfield unit values for selected substances

Substance	HU
air	-1000
fat	-84
water	0
blood	30 - 45
muscle	40
soft tissue	100 - 300
bone	700 - 3000

A consistent mapping between HU is available for many substances, an extract of those is shown in Table 2.1. This mapping makes image interpretation and analysis less complex and more robust. Nonetheless, it must be noted that this mapping can still be invalidated through strong artefacts in the image, e.g. larger implants.

CT and CBCT are robust and fast technologies. They have their strengths in skeletal imaging but lack the details in soft tissue. The use of ionising radiation means an additional burden on the patients: In Switzerland, 173 cases of cancer caused by X-ray imaging have been reported. This

results in an 1% attributable risk [BD04].

## 2.2.2 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is an imaging technology that allows distinguishing between pathological and healthy tissue in clinical practice. It is non-invasive and very tolerable. Just in the United States of America, the number of MR scans has increased in the last seventeen years from 9 million a year to almost 34 million. MR scans allow a good distinction between different types of soft tissue [MZMM12, VPJ02, JGSL97]. Several kind of sequences [OAC<sup>+</sup>09, TTF<sup>+</sup>97, LHG<sup>+</sup>96] allow broad applications and the image quality can even be increased with contrast agents [RBM<sup>+</sup>05, SS08, TKR<sup>+</sup>08, Car06]. MR scans can also capture time-variant processes [FJT94]. It is possible to scan slices within a second to acquire dynamic image series, e.g. of a beating heart [CML<sup>+</sup>12]. To capture moving organs, e.g. the heart or the abdomen while breathing [FVC<sup>+</sup>11], the acquisition can be synchronised to external input like heart rate or breath rate.

Nonetheless, MRI also carries some disadvantages. MR scans usually take more time than other methods [VSG<sup>+</sup>15]. Due to the strong magnetic fields that are used, common contraindications are metal implants, pacemakers, or tattoos. Due to the small diameter of the machines (60cm, newer systems up to 75cm) claustrophobia can be a problem. Affected patients have to be monitored carefully or drugged. Although open MR scanners exist, they achieve a lesser image quality due to a lower magnetic field that is less homogeneous. MR scanners consume much energy (40-100kW) and are very noisy (exact amount depends on sequence), ear protection is required for patients during the scans. The combination of long acquisition times and high energy consumption makes MRI expensive.

### 2.2.2.1 History

The history of MR imaging goes back to 1952 when Herman Carr created a one-dimensional image in his Ph.D. thesis [Car52]. Raymond Damadian built the first MRI machine in 1972. It was targeted to find cancerous structures, based on an idea published in 1971 [Dam71]. Later it was found that this is impossible with the proposed method. In 1973, Paul Lauterbur created first 2D and 3D images using gradient coils and 1974 he produced a cross-sectional MR scan of a mouse. Finally, in 1977, Damadian, Minkoff and Goldsmith performed the first scan of a human. A comprehensive write-up of the history of MRI and the controversies around it can be found in [Kau14].

### 2.2.2.2 Working Principle

Magnetic resonance imaging uses nuclear magnetic resonance to create cross-sectional images. A strong, static magnetic field  $B_0$  aligns the nuclei of the atoms. The nuclei spin around themselves. A strong magnetic pulse deflects the nuclei so that they go into a precession. Due to the strong static field, they return to their initial position. The duration of the return depends on the

type of atom. During the return energy is emitted as a radio-frequency (RF) signal that can be detected. The return is controlled by two independent processes: T1 and T2 relaxation. They can be influenced by the repetition time  $T_R$  and the echo time  $T_E$ .

This is a very fragile process that needs to be carefully timed since only ten out of  $10^6$  protons give a signal. The signal-to-noise ratio is approximately proportional to field strength  $B_0$ , therefore stronger fields produce better images. The current generation of scanners has a field strength of 3T. The main engineering challenge is to keep the field homogeneous, which is crucial for image quality. Since there are always small inhomogeneities in the field, advanced bias field correction techniques [TG09b] have been proposed and are now common in MR image processing.

The spatial localisation of the received RF signals is done using additional field gradients that encode the position [LWS<sup>+</sup>12]. Therefore, a free selection of the cross-section planes is possible, although usually the common coronal, axial, and sagittal planes are used, as they are easier to address with the gradient coils.

The captured signal value can not be directly linked to tissue types; the interpretation is done using contrast and gradients. A strong signal is called *hyperintense* and shows very bright, weak signals are called *hypointense* and are shown dark. T1-weighted images make fat hyperintense, e.g. fatty structures like bone marrow, T2-weighted images highlight stationary liquids.


### 2.2.2.3 Artefacts

**Motion Artefacts** During the imaging process, the patient is expected to lay perfectly still. This can cause inconsistencies between acquisition steps, leading to seams at transitions. If motion artefacts occur within an acquisition step, the phase information is modified and causes a wrong representation [AHS<sup>+</sup>99]. A possible solution is the averaging of multiple acquisitions although this leads to longer acquisition times. Sometimes motion artefacts are caused by frequency coupling from the machine, e.g. through a mounted coil. Figure 2.3 shows such an artefact.


**Flow Artefacts** Even if the patient lies perfectly still, the organism is still running. Blood is being pumped through veins and arteries. This flow causes artefacts [EF90] like shown in Figure 2.4. Flow artefacts are a subclass of motion artefacts. These artefacts are hard to avoid. Their orientation can be controlled by adjusting the field gradient. Another approach is to adjust the capture to the pumping sequence. This comes with the cost of highly increased acquisition times.

**Field Inhomogeneity** By using stronger magnets, the volume usable for scanning has been increased. At the margins of the volume, field inhomogeneities can occur, as the field weakens further away from its centre. It will results in stripes [VPL07] on the resulting image, see figure 2.5. This artefact can also be caused by external magnetic field, which supports the need for a good shielding of the MRI machine.

**Metal Implants** Hip implants, pacemakers, and other implants can, although being certified for MR compatibility, cause artefacts [HWP<sup>+</sup>11]. The magnetic flow differs drastically in the metal part from the surrounding tissue and, therefore, causes field inhomogeneities with the consequences previously described.



**Figure 2.3** – Extreme case of motion artefact. The subjects movements destroy the spatial coherence of the measured signals.




**Figure 2.4** – The blood circulation through the aorta inside the knee joint causes the flow artefact shown in this figure.

**Partial Volume Artefacts** The Partial Volume Artefact [VVBV01] is caused by imaging voxels containing two different tissues and, therefore, possessing a signal average of both tissues. It blurs boundaries that usually would be clearly visible.

**Ghosting Artefacts** Ghosting artefacts can occur when the captured object stretches out of the selected field of view of the scanner [RABM97]. An example is shown in Figure 2.6.

**Chemical Shift Effect** This artefact results from the shift of tissues and causes a clearly visible, black delineation at the edge [SVS84].

**Folding artefact** Imaging a large structure with a too small field of view will result in parts of the images folding over to the other side.



**Figure 2.5** – Inhomogeneities artefact, a gradient along the muscles of both thighs is clearly visibly. It results from an inhomogeneous magnetic field.



**Figure 2.6** – Ghosting artefact in the lower half of the image: Instead of the femur a ghost of the other knee is seen.

#### 2.2.2.4 Protocols

Many parameters influence the resulting data set of an MRI scan. The two most prominent parameters are the  $T_1$  and  $T_2$  times. They can be tuned to highlight or suppress different tissue types. Their flexibility can be expanded by using tracers: Chemical compounds that are injected into the patient and accumulate in the region of interest. These tracers are chosen to have beneficial MR properties so that, for example, inflammations can be detected.

The combination of acquisition parameters, body part and, if necessary, tracer is called a protocol. There exist many standard protocols in clinical routine, allowing reproducible acquisitions of consistent quality. Protocols take a one-time effort for creation but then offer long-term benefit. In the following part, we will introduce three protocols that have proven to be useful for musculoskeletal imaging.

#### Modified Dixon (mDixon)

In 1984, Dixon [Dix84] has developed a modification of a spin echo pulse sequence. His approach allows generating images that show only water and only fat. The method has no significant

increase in acquisition time as well as comparable speed and signal-to-noise ratio. It increases the contrast between organs and allows better tissue characterisation.

The modified Dixon technique (mDixon) is a multipoint method for 3D water-fat magnetic resonance imaging. It uses an unrestricted echo time (TE) for more efficient scanning over traditional Dixon. The mDixon sequence consists of two parts. First, water and fat are attenuated in phase, resulting in the in-phase image. In a second step, the pulse is tuned to acquire fat and water out of phase. By voxel-wise adding the in-phase and opposed-phase images, a water showing image can be generated. By subtracting the opposed-phase image from the in-phase image, a fat highlighting data set can be created. Figure 2.7 shows an example of the four images that are acquired. In comparison with current 3D volumetric fat suppression methods, mDixon provides greater speed, improved signal-to-noise-ratio (SNR), multiple contrasts and more robust fat suppression. Today, the mDixon protocol is used for many applications like the detection of cancer metastasis [CKM<sup>+</sup>12], progress monitoring limb-girdle muscular dystrophy [WHC<sup>+</sup>13], and liver fat quantification [FRM<sup>+</sup>12]. Other applications are the attenuation correction for hybrid PET/MRI [BFS<sup>+</sup>12], and the assessment of muscle fat content [FPE<sup>+</sup>14].

### **T1 TSE (Turbo Spin Echo)**

Spin echo sequences are the most commonly used pulse sequence. The timing can be adjusted to weight T1, T2, proton, or spin density. A 90-degree pulse is followed by a series of 180-degree refocusing pulses that improve the SNR. An accelerated version is the turbo spin echo, which reduces the acquisition duration. In the T1 weighted case, the images show the different T1 relaxation times of the tissues. In this sequence, fat appears very bright (gives a high signal), and water very dark (gives a low signal). This is helpful for the identification of muscles, since the inter-muscular fat, which separates the individual muscle groups, is highlighted.

### **Knee Acquisition Protocol**


The anatomical structures at the joint level are very small and, therefore, a high resolution is needed to capture them. Higher resolutions always compromise between acquisition duration and SNR. For knee acquisitions, a specialised coil is being used. It wraps closely around the joint and provides improved SNR. The Philips VISTA SPAIR sequence is a 3D TSE sequence with Variable Flip Angle. It is an extension of the above describe TSE sequence that temporarily varies the flip angle between consecutive RF pulse to increase the SNR and to compensate phase encoding artefacts. With this sequence, it is possible to highlight cartilage and ligaments.

## **2.2.3 Other Methods**

For the sake of completeness, we will give a short overview of other common medical imaging techniques, covering positron emission tomography, ultrasonography, endoscopy, thermography and electrical impedance tomography.

**Positron Emission Tomography (PET)** is an imaging technique of the nuclear medicine. It allows





**Figure 2.7** – Cross-section of the thigh, showing femur and muscles, using the Dixon protocol: (a): in-phase, (b): opposed-phase, (c): water-emphasizing and (d): fat-emphasizing.

to capture quantitative information about the metabolic system. First, the patient is injected with a tracer, made of molecules labelled with positron emitting isotopes. Through metabolic processes, the tracer accumulates, for example, in tumours. The PET scanner then captures the gamma radiation from the decay of the isotopes. The natural decay rate makes the process time critical as the half-life time of common tracers ranges from 2 minutes ( $O_{15}$ ) through 20 minutes ( $C_{11}$ ) up to 110 minutes ( $F_{18}$ ). PET is usually done as a hybrid acquisition in combination with CT or MRI [JWN<sup>+</sup>08]. The low-resolution PET benefits from the spatial information from CT/MRI to localise the relative position of tracer accumulations in the body. An overview of the early developments in nuclear medicine has been provided by Brownell *et al.* Brownell1990.

**Ultrasonography** is a technology used to image structures using (ultra-)sound. It used for cardiac imaging and as well as prenatal diagnosis. Further organs include the thyroid, kidney and bladder. The spreading of the sound waves is limited by bones and gas, so areas of uses are limited. A sound wave is sent through the skin using a quartz crystal to create the images.

The sound wave is reflected at the transition boundary of different tissues. The reflection is received using linear or two-dimensional detectors; the reconstruction of images uses the signal run-time and exploits the Doppler effect. Ultrasonography can be performed in real-time to get a live video feed. First experimentation goes back to the 1950s; the first application was cancer diagnosis [WN51, Wat09]. However, US only works near the body surface and has a low SNR.

**Endoscopy** is a way to look inside the body. Through a set of fibres, light is sent into the body and the absorption is returned through a second set of fibres. Like this, organs can be diagnosed when they are accessible through natural cavities or minimal cuts. This technique has been extended to include small tools to perform minimal incision surgeries.

**Thermography**<sup>2</sup> collects the infrared radiation ( $9 - 14 \mu m$ ) emitted by objects, to get a temperature distribution image. The method is limited to the body surface and a thin layer below it. It allows detecting effects like raised temperatures and assists in mammographies. Thermography is usually used to extend the information gathered using other methods and also for patient monitoring.

**Electrical Impedance Tomography (EIT)**<sup>3</sup> was invented by Henderson *et al.* [HW78] in 1978, the first working device was built by Barber and Brown [BB84] in 1984. The principle idea is the distinction between tissues by their impedance. EIT is used for the detection of cancer in skin and breast and lung function monitoring. Using a belt of electrodes that is placed around the tissue to be imaged, small currents (a few  $\mu A$ ) with frequencies around 10-100 Hz are sent through the body and received by the other electrodes. The acquired data can be reconstructed using complex algorithms. In 2011, first commercial systems were offered to monitor the lung capacity of patients that receive intensive care and for breast cancer detection.

### 2.2.4 Comparison

The methods presented above all carry advantages and disadvantages. The proper choice of imaging technique depends on several factors: acquisition duration, region to be imaged, structures of interest, but also economic factors like availability and price. X-rays have been proven useful for minor fractures, CT is popular for more complex injuries and in emergency care. MR has strong capabilities in imaging soft tissues, especially inner organs like liver, kidney and muscular system. PET (Positron Emission Tomography) and PET/MR are used in nuclear medicine, usually to detect tumours. To monitor lung volume during surgery, EIT is a good candidate. US is commonly used for pre-natal and coronal examination. A comparison chart of the different properties of the presented techniques can be found in Table 2.2.

For musculoskeletal modelling, visibility of soft tissues is crucial. Furthermore, spatial information is required to locate structures and to measure their dimensions. The modelling of the lower limb requires imaging of a large area. Therefore, exposure should be kept low, and acquisition

2. Also known as infrared thermography (IRT), thermal imaging, digital infrared thermal imaging (DITI), and thermal video.

3. Also known as applied potential tomography, electrical impedance imaging, conductivity imaging or impedance computed tomography.

time (which is linked to price) should be kept low. X-ray, CT, and CBCT show low costs and good bone visibility. However, their invasiveness and low visibility of soft tissues like muscles and tendons rule out their use in musculoskeletal modelling. PET only shows low resolutions and is very invasive; EIT cannot image soft tissues, so these techniques cannot be used either. Ultrasound imaging shows muscle structures at a low cost, but is only applicable close to the body surface - the complete diameter of the lower limb cannot be covered.

**Table 2.2** – Comparison of popular imaging techniques by their properties to evaluate the applicability for musculoskeletal imaging (-: negative, o: neutral, +: positive).

	X-Ray	CT	CBCT	MRI	PET	PET/MR	US	EIT
Resolution	-	+	+	+	o	+	o	-
Exposure to radiation	-	-	-	+	-	-	+	+
Popularity	+	+	o	+	-	-	+	+
Price	+	o	o	-	-	-	+	o/-
Acquisition time	+	+	+	o	-	-	+	+
Soft tissue visibility	-	-	-	+	+	+	+	-
Bone visibility	+	+	+	o	-	o	+	-
Non-invasive	-	-	-	+	-	-	+	+

After considering the other options, MRI is the best fitting imaging technique for musculoskeletal modelling. Although it is more expensive and requires longer acquisition times, MRI provides higher resolutions, good soft tissue visibility and acceptable visibility of bones. MRI is non-invasive as it uses no ionising radiation and is wide-spread among hospitals. Furthermore, MRI allows tweaking the imaging protocols for specific applications. The tweaks can highlight, for example, the boundaries between soft tissue, which we will discuss later in this thesis. A good understanding of the imaging technique and the protocols used is a crucial factor for selection, design and configuration of segmentation approaches.

## 2.3 General Segmentation Approaches

Medical imaging nowadays is part of the routine in hospitals. The acquired images get better in quality and improvements are made to reduce the exposure of patients to radiation. Modern techniques allow for capturing small details and for differentiating between kinds of soft tissue. However, with the increased number of images and their higher resolution, interpretation has become a more complex task. Nonetheless, medical image segmentation is required for applications like radiotherapy, preoperative planning and postoperative evaluation.

Medical image segmentation, as described by Elnakib *et al.* [EGS<sup>+</sup>11] in a survey, is the process of identifying regions of interest in images. Approaches range from simple ones that only exploit intensity values or region information to model-based ones that include a priori knowledge. The images often suffer from noise, aliasing and anomalies or may contain gaps in boundaries, providing challenges that are hard to handle with non-model-based approaches.

Common segmentation approaches can be put into four categories, based on the features they use. Pixel- and Histogram-based methods directly work on the intensity value of a pixel or its

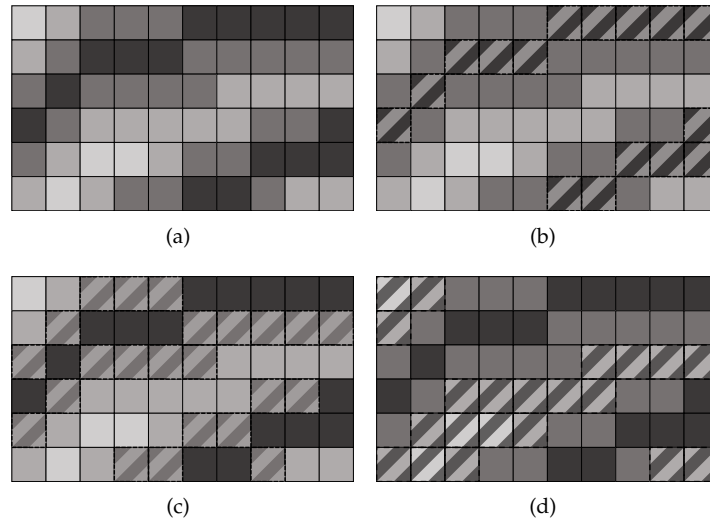
position in the histogram. The region-based methods exploit neighbourhood properties and are often used for tubular structures like arteries, veins or the colon. Edge-based methods perform edge detection on the image and then try to stitch the detected edges together. Texture-based methods analyse the textures in an image and therefore can detect structures with non-uniform intensities. A comprehensive survey of these and further methods has been made by Lorenz [ES03]. We describe and discuss these methods below.

### 2.3.1 Pixel- and Histogram-based Methods

Pixel-based methods directly operate on individual pixels (or voxels) and can be simple thresholds [SLJ13], ranges or more complex transfer functions between an intensity and the membership of a segment. The membership of a pixel/voxel with the intensity  $i$  can be tested using the function  $f(i)$ :

$$f(i) = \begin{cases} \text{true} & \text{if } i \geq \text{threshold} \\ \text{false} & \text{if } i < \text{threshold} \end{cases}$$

An example of threshold-based segmentation is shown in Figure 2.8. In a 2D example, different thresholds for low, medium and high intensities are visible.




**Figure 2.8** – Thresholding segmentation principle. From an original 2D image (a), pixels with low (b), medium (c) and high (d) intensities are selected and shown with stripes.

Thresholding is used on brain MR images [SLJ13] and to determine the volume of the left ventricular of the heart [JSR<sup>+</sup>08, NdGH09]. Due to the low spatial resolution and the absence of steep gradients, thresholding can also be used for segmenting PET images [JFE<sup>+</sup>07]. Newer approaches consider the input images to consist of two fuzzy sets [YD13], to reduce the need for a concrete threshold. Shahbazpanahi *et al.* [SRM14] propose to use a genetically learned transfer function for thresholding in 2D X-ray images for the segmentation of the lung volume. Thresholding is a simple approach that performs well on datasets with low complexity where the detection of very high (PET) or low values (X-ray) is needed. Since no neighbourhood information is used,

spatial model concepts can not be expressed. To overcome this, the watershed approach has been proposed.

The watershed approach tries to locate sinks, surrounded by high intensities in the image. This 2D driven approach is shown in Figure 2.9. It has been used for lung cancer segmentation [KTL15] and for segmenting the Visible Human data set [CWJ05]. More sophisticated approaches combine clustering, watershed and edge-based methods for brain segmentation [NOF<sup>+</sup>06]. While more robust than thresholding, the watershed approach can suffer from producing many small segments and requires further processing the re-established spatial coherence of segments.




**Figure 2.9** – Watershed segmentation of the Visible Human data set (a) and working principle (b). Source: Cates *et al.*, *Case study: an evaluation of user-assisted hierarchical watershed segmentation.*, [CWJ05], used with permission.

### 2.3.2 Region-based Methods


Region-based methods expand the segmentation around a neighbourhood based on seed points as shown in Figure 2.10. They are commonly used for slice-wise growing [AM14], vessel reconstruction [AM14] and the reconstruction of massively branched, tubular structures like arteries, veins, or  $\mu$ CT bone scans [RMPC02]. Other uses are the local application of image enhancements, e.g. in ultrasound images [KGG11]. While region-based methods include local neighbourhoods, they lack a more comprehensive understanding of the segmented structures.

### 2.3.3 Edge-based Methods

Segments in images are often limited by edges, areas with high gradients. Examples of such edges can be seen in Figure 2.11. These gradients can be detected using image kernels like the Canny-[Can86] or Sobel-[SF73] image operators. Edge-based methods are often combined with thresholding for the segmentation of the brain in MR images [TH11, TWM<sup>+</sup>00, JS12]. In CT, they have been used for kidney segmentation [MM15]. 2D images are widely segmented using edge-based methods, e.g. malaria cell detection in microscopy [SE15], the partition of photographs [FH04], or the segmentation of man-made structures in satellite imagery [MSK04]. However, the detected edges have to be processed, stitched and combined to create connected segments. This complex task becomes more expensive in 3D with planes instead of edges.



**Figure 2.10** – Region growing principle on a sample 2D image. First a seed point is selected (a), and then over a series of steps (b), (c), (d) the neighbourhood is explored and added (shown through stripes) if appropriate.



**Figure 2.11** – Fine (b) and more rough (c) segmentation of edges in a sample 2D image (a).

### 2.3.4 Texture-based Methods

Organs and other anatomical structures often carry a distinct texture on their surfaces. Texture analysis [DFRVM14] can be used for distinguishing between different organs. They combine different kernels and features for image classification [ZMLS06]. They can be used for images with repeated features, e.g. muscle fibres [KBFW07]. Common use cases are endoscopy images and skin analysis. MRI texture analysis is also used in brain imaging [KKGv01], abdominal imaging [WPNK08] and analysis of lesions in breast [NCY<sup>+</sup>08] and brain [KKGv01, KPG08]. However, skeletal muscles do not carry discernible texture features [BACP12b] which could be used for segmentation. Visible features in skeletal muscles are limited to diseases like lymphoma [CJP<sup>+</sup>10], inflammation [SKE<sup>+</sup>09] and others [TTK12].

### 2.3.5 Comparison

The presented general segmentation methods can be applied to different domains. Many are successful in highly specialised use cases. However, they require customization. Since they are very general, they need to be adapted for more complex use cases. We have found that for

musculoskeletal modelling in MR images, these methods are not sufficient. They either specialise on features not present (texture and edges), non-applicable model properties (tubular, region-based) or intensity properties that are not given in MRI (threshold, watershed).


## 2.4 Morphology-based Methods

Morphological operations are important techniques in processing image data. The theory behind these was introduced by Serra [Ser82] in 1982. They were first employed to analyse images by Haralick *et al.* [HSZ87]. Dilation and erosion are two main morphological operations. They apply a structuring element to modify the boundaries of objects in a binary image. Erosion is used to eliminate protrusions, to shrink or thin the objects. In opposite, dilation lengthens or thickens objects. The dilation of a binary image  $A$  by the structuring element  $B$  can be written as  $(A \oplus B)$  and the erosion  $(A \ominus B)$  can be defined as:

$$A \oplus B = \{z | (\widehat{B})_z \cap A \neq \emptyset\} \quad (2.2)$$

$$A \ominus B = \{z | (B)_z \subset A^c\} \quad (2.3)$$

with  $\widehat{B}$  as reflection of  $B$  and  $A^c$  as supplement of  $A$ . A more detailed description of their function and purpose in the medical domain can be found in [ZZM12]. The combination of erosion followed by dilation is called opening, a dilation followed by erosion is called closing. Dilation, erosion, opening and closing are demonstrated in Figure 2.12.



**Figure 2.12** – This figure depicts the effects of morphological operations on an image (a) using the structuring element (d). The shown operations are erosion (b), dilation (c), opening (e), and closing (f).

The application of morphological operations in (semi-) automatic medical segmentation approaches has shown great potential. Resmi *et al.* [RT12] demonstrate brain tumour segmentation in MR images, Jin *et al.* [JIL<sup>+</sup>03] show an approach for adipose tissue segmentation for whole body T1 MR images. A semi-manual software for fat and muscle segmentation that uses thresholds has been developed by Mattei *et al.* [MFC<sup>+</sup>06]. Appel *et al.* [AALB13] give an overview of

the challenges connected with biomaterials and discuss the different imaging techniques and applications.

Another application of morphological operations is edge detection. Yu-Qian *et al.* [YQWHZC<sup>+</sup>05] have shown this in medical images. The extra-ocular muscle boundary in CT images can be detected with the approach of Souza *et al.* [SR00] and Elliott *et al.* [EWG<sup>+</sup>97] have shown a semi-automatic method for human calf muscle volume measurement that combines morphological operations with boundary following. These methods are computationally expensive.

Other approaches use a combination of morphological operations and clustering, for example, the work of Barra *et al.* [BB02], or the registration of pre-segmented atlases as demonstrated by Leinhard *et al.* [LJR<sup>+</sup>08] for abdominal fat segmentation. Applications of morphological operations to other regions of the body include the work of Kullberg *et al.* [KAJF07] which specialises on the abdominal region and Kober *et al.* [KGZS11] for shoulder muscle segmentation.

Monziols *et al.* [MCB<sup>+</sup>06] have presented an approach for the quantification of different types of tissue in pig carcasses. They use morphological operations in MR images to detect subcutaneous fat and have developed a prediction model for belly fat. Kimori *et al.* [Kim11] have shown a method for contrast enhancement in mammographic, chest and retinal images. They use morphological operations, followed by contrast modification techniques. The work of Makrogianis *et al.* [MSF<sup>+</sup>12] is a method for measuring subcutaneous fat and muscle volume. They use morphological operations followed by the application of active contours. Their work is focused on single slices of the mid-thigh and the calculation of the active contours is complex. Csapo *et al.* [CMS<sup>+</sup>14] have investigated age-associated differences in triceps composition by testing a series of MR sequences that provide fat and water maps. Elliot *et al.* [EWRP13] use the variable projection (VARPRO) to find an iterative solution on the segmentation of cervical spine muscle fat. Thomas *et al.* [TNL<sup>+</sup>14] demonstrate an automated whole-body muscle segmentation in an adapted Dixon sequence. They use a non-rigid registration process followed by thresholding and fat removal and can distinguish between muscle compartments. The methods require time-consuming calculations. There are many groups working on fat labelling in MR images [VKA<sup>+</sup>13, HK13, PSR<sup>+</sup>13, TFM<sup>+</sup>13, AKP<sup>+</sup>14].


Overall, morphological operations are very capable when used for tissue type labelling. However, the identification of individual structures, e.g. single muscles, is not possible.

## 2.5 Classification

Image segmentation is the process of assigning labels to pixels or voxels. For every pixel/voxel additional information may be generated, such as gradients, additional channels, and derivatives. These features, after normalisation, can be considered a vector describing the location of the sample in a feature space. During the training of a classifier, training samples with manually assigned labels (=belonging to a specific segment) are used. The classifier tries to find a (higher dimensional) curve that separates the groups of samples. Once the curve is found, it can be used to test efficiently on which side new samples are located and to assign them corresponding



labels.




**Figure 2.13** – Example of classification. Samples (blue/x and orange/+) are plotted along two features (x- and y-axis) (a). To divide the two groups, a parable (black, dashed) is introduced (b) which allows a fast classification of samples as members of either group.

There exist many different approaches to classification. They include fuzzy c-means (FCM) clustering [HWWL14], decision forests [GCM<sup>+</sup>11], random ferns [PGC<sup>+</sup>11], random walkers [Gra05]. Hidden Markov models (HMM) can also be used for classification [IJKY06, AIA11]. MR phase information can be another input for classification [BFS<sup>+</sup>07].

## 2.6 Registration

In short, image registration is the process of finding a transformation ( $T$ ) that maps a reference image to a target image [GSKP11]. A fixed image, usually carrying an annotation, and a moving image, which should receive the annotations, are given. The registration is the process of finding a transformation that increases the correct overlap/overlay of the moving image on the fixed image. These transformations could be a translation, scaling, rotation or more complex ones. Figure 2.14 shows the principle of registration with a small example.



**Figure 2.14** – Principle of registration. A moving object (grey, a) with the bounding box should be overlaid on the fixed image (grey, c). The bounding boxes more clearly demonstrate the needed transformations. In this example, the transformation  $T$  (b) consists of a height scaling with the factor 1.5 and a rotation of  $20^\circ$ .

While commonly a regular grid is used to express the deformation of the transformation ( $T$ ), more complex meshes, tuned to specific use case, can be used [CGZM15]. This so called non-rigid registration can also be done using B-splines [KSP07]. Registration has a high computational

cost. Therefore, the application of multi-core CPUs or GPUs for massive parallelisation is an active area of research [SSKH10].


Registration may also be used to register a label image [ESS12]. During this process, an initial atlas is transformed to the target image and gives a segmentation. This segmentation from atlas registration can also be the input to further segmentation approaches [FWCO06, FBCO07].

Depending on reference and target image, the accuracy of the registration can vary within the image [GFW13]. Additionally, the registration process is very sensitive to changes in the subjects pose and size. The differences can only be expressed to a certain extent before the accuracy decreases drastically.

In the rest of this thesis, we will use the term registration to express the registration of label image and therefore, we consider it a segmentation technique.

## 2.7 Graph Cuts

Graph cuts are used to solve computer vision problems efficiently through global optimisation. A formulation as energy minimising approach (e.g. in level sets) can be reduced to the maximum-flow problem. The image is represented as a graph with nodes for the pixels/voxels and edges connecting adjacent nodes. The edges are weighted based on the intensity gradient between the pixels/voxels they connect. By selecting pixels/voxels inside and outside the segment, a flow can be assumed between these nodes. The max-flow-min-cut theorem states that the maximum flow is limited by a minimal cut. By solving this optimization problem, this cut, which equals the segment boundary between the inside and outside nodes, can be determined. An example illustration of a graph cut in a 2D image can be found in Figure 2.15.



**Figure 2.15** – Working principle of graph cuts. Pixels inside (S) and outside (T) the segment are defined and a graph, weighted by gradients, is built. The cut (black) equals the segment boundary.

Ababneh *et al.* [APG11] have presented an application of this approach for the segmentation of the knee bones in MR images. Zhou *et al.* [ZYZ10] have presented a graph-cut-based method for industrial image segmentation. Zhu-Jacquot *et al.* [ZJZ07] have proposed a method to combine graph-cuts with statistical shape priors. A similar approach has been presented by El-Zehiry *et al.* [EZE08]. Chen *et al.* [CUB<sup>+</sup>12] have demonstrated a combination of graph-cuts and active

appearance models to get more accurate segmentation results.

Vineet and Narayanan [VN08] demonstrated how to calculate graph cuts on the GPU. Their approach achieves better performance as it exploits the highly parallel structure of GPUs for concurrent calculation. Another optimisation has been proposed by Delong and Boykov in 2008 [DB08]. Their approach reduces the memory requirements for large data sets and achieves a nearly linear increase in computational speed with respect to the number of processors in multi-core systems. Lee *et al.* [LSP<sup>+</sup>10] have proposed Branch-and-MinCut in 2010. This approach combines graph cuts and other techniques like branch-and-bound to allow a global optimisation for a wide class of energies. This can be used to include shape- or colour-distribution priors.

## 2.8 Comparison

Until here, we have presented several segmentation approaches, general ones, registration, classification and more. To select a technique, we must evaluate if it can produce robust, reproducible, and correct results. Therefore, we have defined criteria tailored to our use case to compare the approaches.

The segmentation of the lower limb covers a large area and many different structures. Segmentation approaches have to be efficient and limit their computational costs to limit the total processing time. An approach can only be used if it is possible to integrate into practical scenarios. While this might be tackled in some cases by parallelisation and use of many-core CPUs or GPUs, foundational complexity in the algorithms will be a growing problem with higher image resolutions. The scope of the approaches can be either global or local. Local implies the use of only information which is located closely, e.g. neighbouring image elements, whereas global approaches also include a more comprehensive understanding of the structure to be segmented. An example would be spatial constraints, e.g. minimal diameter or maximal length. In clinical routine, the conditions of acquisitions can not be fully controlled. Subjects may have injuries or disabilities that prevent them from taking the optimal posture during the scans. Therefore, a segmentation approach has to be insensitive to (at least small) posture changes. Finally, the segmentation approach has to fulfil some sanity checks. The resulting model has to be a single, connected component, just like the structure to be segmented. A subdivided output with missing parts is not a desirable output. Finally, soft tissues usually have a smooth, gapless surface. A segmentation approach has to preserve these properties, stairs that go beyond the quantification of the image, should be avoided. A tabular comparison of the approaches with the presented criteria can be found in Table 2.3.

It is important to have smooth, connected results. Therefore, we can exclude the following approaches which cannot guarantee this: histogram-, region-, and texture-based as well as classification approaches. Edge-based and morphological approaches work locally and discard global connections. Registration is the only global approach but suffers from high computational costs and sensitivity to posture changes. This means that none of the presented approaches fulfils the selected criteria. To overcome this, a new group of approaches has been proposed: model-

**Table 2.3** – Comparison of presented segmentation methods using the criteria defined in this section. Desired values have been printed in bold letters.

	Insensitive to posture changes	Local / global	Computational costs	Connected results	Smooth results
Histogram	<b>Yes</b>	Local	<b>Low</b>	No	No
Region	<b>Yes</b>	Local	<b>Low</b>	<b>Yes</b>	No
Edge	<b>Yes</b>	Local	Med	Poss.	<b>Yes</b>
Texture	<b>Yes</b>	Local	Med	No	Poss.
Morphology	<b>Yes</b>	Mixed	<b>Low</b>	<b>Yes</b>	Poss.
Classification	<b>Yes</b>	Local	<b>Low</b>	No	No
Registration	No	<b>Global</b>	High	<b>Yes</b>	<b>Yes</b>

based segmentation. This class of approaches integrates specific model information, like shape, appearance, and rigidity and allows a more robust and flexible segmentation in changing imaging conditions. Deformable models are a popular technique of the model-based segmentation approaches.

## 2.9 Model-based Segmentation

Deformable models have been first proposed by Terzopoulos *et al.* in 1987 [TPBF87]. They provide a robust segmentation approach that uses bottom-up image-based constraints and top-down constraints from prior knowledge. Deformable models can be curves or surfaces (or of higher dimension, e.g. for temporal segmentation). They evolve under the influence of internal and external energies. The internal energy controls the curve's smoothness and the external energy aims to attract the model towards boundaries in the image domain. Deformable models are an interesting approach as they combine geometry (to describe the shape), physics (to simulate the behaviour) and approximation theory (for model fitting).

Deformable models have a broad range of possible applications. They have been used in computer graphics [NMK<sup>+</sup>06], to calculate the deformation of clothes [BW98] and in image composition [MB95]. Other uses include optical flow analysis for facial animation [DM00] and the determination of vehicle types [ZTHW12].

Due to their power and robustness, deformable models are often used in medical applications. For example, Deserno *et al.* [DLM<sup>+</sup>09] presented an application for segmenting the bony orbit while Heimann *et al.* [HMS<sup>+</sup>10] proposed a grand challenge for knee segmentation. Schmid [SKMT11] presented the segmentation of the hip bone, Snel *et al.* [SVG02] used a deformable model for wrist segmentation and Rafai *et al.* [RBH<sup>+</sup>00] proposed a method for skull segmentation. Another medical use is surgery simulation which has been reviewed by Meyer *et al.* [MLM<sup>+</sup>05]. Medical simulation approaches that use deformable models to describe the mechani-

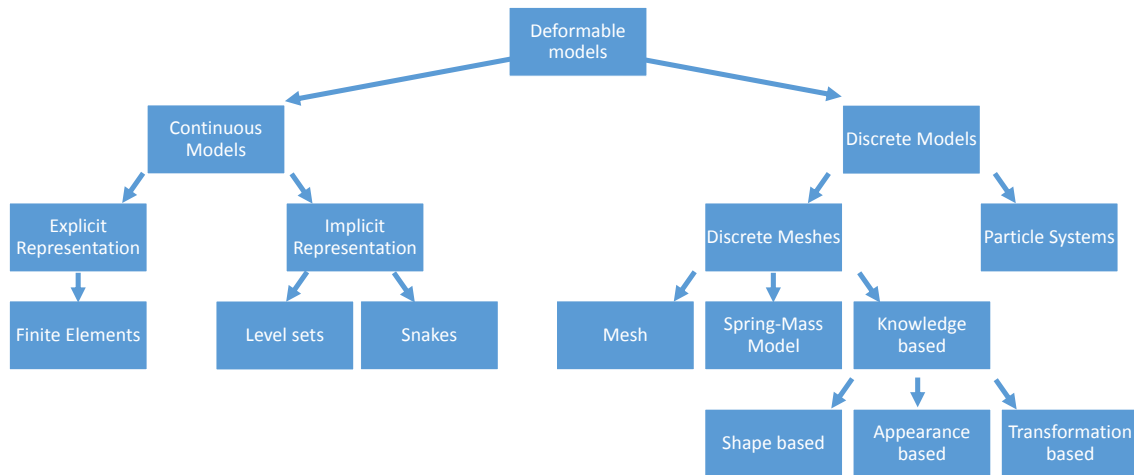
---

*This section was previously published as:*

M. Becker, and N. Magnenat-Thalmann, "Deformable Models in Medical Image Segmentation," in *3D Multiscale Physiological Human*, 1st ed. Springer-Verlag London, 2014, pp 81-106.

cal behaviour need proper evaluation. Marchal *et al.* [MADC08] proposed a new and open framework to combine several metrics and models to compare different algorithms.

Several surveys on deformable models have been published, e.g. by McNerney *et al.* [MT96] in 1996, Montagnat *et al.* [MDA01] in 2001 and Hegadi *et al.* [HKH10] in 2010. Moore *et al.* [MM07] provided a general survey of deformable models while Jain *et al.* [JZD98] have reviewed deformable template models, which are used for segmentation, image retrieval and video tracking.



**Figure 2.16** – Classification of Deformable Models, based on [HD11]

Heimann and Delingette [HD11] have proposed a scheme for organising the different kinds of deformable models (see Figure 2.16). They can be divided into discrete and continuous representations [MDA01, HD11]. The discrete models can be split into particle systems and meshes. Continuous representations can have either an explicit (snakes) or implicit (level sets) description of their surface. These different classes of deformable models will be surveyed in this overview. First, we will introduce active contour models, also known as snakes (Section 2.9.1). Next, several approaches to the level sets method will be discussed (Section 2.9.2), followed by a description of discrete deformable models (Section 2.9.3). We take a look at knowledge-based deformable models (Section 2.9.4) that are a specialisation. Other types of deformable models are discussed in the Appendix. An overview of external forces (Section 2.9.5) follows. In the conclusion, the most important aspects are summarised and outlooks are briefly discussed.

### 2.9.1 Snakes

Active contour models, usually called snakes, are a class of deformable models with a continuous representation based on an explicit surface description. They have been originally proposed by Kass *et al.* [KWT88] in 1988 and have been applied as a method for edge detection [LL93], motion tracking [Pet99], stereo matching [BD93] and interactive image interpretation [KWT88]. Other applications are shape modelling [TF88, MT95] and segmentation [LL93, DKY95].

The main concept is to use an energy-minimising parametric curve or a spline, reducing the problem to a minimisation problem. Snakes need to be initialised closely to the final shape to

ensure that they are attracted to the correct features.

The energy of a snake  $\alpha(s)$  can be written as

$$E_{\text{snake}}(\alpha(s)) = E_{\text{ext}}(\alpha(s)) + E_{\text{int}}(\alpha(s)) . \quad (2.4)$$

The external energy  $E_{\text{ext}}$  represents external constraints and image influence to get the contour pulled towards desired image features. The factor  $\omega_1$  is used to weigh the ratio between internal and external energy.


$$E_{\text{ext}}(\alpha(s)) = \int -\omega_1 |\nabla \alpha(s)| ds \quad (2.5)$$

The internal energy expresses smoothness and tension constraints:

$$E_{\text{int}}(\alpha(s)) = \frac{1}{2} \int \omega_2 \left| \frac{\partial \alpha(s)}{\partial s} \right|^2 + \int \omega_3 \left| \frac{\partial^2 \alpha(s)}{\partial s^2} \right|^2 ds \quad (2.6)$$

where  $\omega_2$  is the weight of the influence of stretching on the contour and  $\omega_3$  weighs the bending. An example of a snake for segmenting the femoral shaft is shown in Figure 2.17.

Xu *et al.* [XP98] have proposed the gradient vector flow (GVF, see section 2.9.5.2) as a new external image force to improve segmentation results. This concept has been extended by Cheng *et al.* [CFK04] to a directional GVF. In 1993 Ivbins and Porril [IP94] have shown a region growing segmentation that exploits a snake with pressure force and statistical characteristics of the image. Mitrea *et al.* [MBM<sup>+</sup>12] have reviewed snakes and proposed an iterative method for snake calculation. Wang *et al.* [WGY<sup>+</sup>07] showed a method for muscle extraction from the leg using snakes. Kauffmann *et al.* [KGG<sup>+</sup>03] proposed a snake-based method to quantify cartilage thickness and volume in MR images.



**Figure 2.17** – Segmentation of the femoral shaft in an MR image using snakes. The figure shows the initial shape and the progress after 1, 2, 3, 5, 7, 10, and 20 steps.

Ip and Shen [IS98] proposed affine invariant active contour models (AI snakes) in 1998. AI snakes are an efficient method of establishing a correspondence between model and data. The basis for their energy function is formed by local and global affine-invariant features.

In 1999, Vemuri and Guo [VG99] presented hybrid geometric active models. They introduced a hybrid geometric snake to allow for topology changes. Their model allows for the representation of global shapes with local details.

Another approach to topology adaptive snakes has been proposed by McInerney and Terzopoulos in 2000 [MT00]. Their T-Snakes offer topological flexibility and significantly extend conventional snakes. This approach has been extended further by others [DF01, OREG05].

## 2.9.2 Level Sets

Level sets have been introduced in 1988 by Osher and Sethian [OS88] to overcome the difficulties that snakes have with changes in topologies. Their work has proven to be a powerful tool in segmentation. The ability of level sets to perform automatic topology changes can be useful in many cases, especially for the segmentation of objects with high shape variations like shown by Israel-Jost *et al.* [IJDAB14]. A common segmentation application is the ventricles, as demonstrated by Angelini *et al.* [AHP<sup>+</sup>05]. Implementations can be found in tools like ITK-SNAP [YPH<sup>+</sup>06] and YaDiV [FBW11] or libraries like LSMLib<sup>4</sup>.

Like snakes, level sets can be regarded as continuous deformable models. To overcome the drawbacks of snakes and to allow topology changes, level sets make use of an implicit representation. A deformable surface  $S$  is implicitly represented as an iso-surface of a time-varying scalar function, embedded in 3D. In Figure 2.18, the principle is illustrated for the segmentation of the liver. The right side shows the 3D time-varying scalar function, on the left, the corresponding iso-surface is shown. The iso-surface is generated from the 3D surface. Therefore, topology changes are easily possible

Let  $\Phi : \Omega \times \mathbb{R} \rightarrow \mathbb{R}, \Omega \in \mathbb{R}^3$ . A level set  $S$  can be defined as an iso-hypersurface on  $\Phi$  as follows:

$$S = \{x \in \Omega | \Phi(x) = 0\} \quad (2.7)$$

We assume  $\Phi$  to be a distance map that is negative inside the level set and positive outside. To integrate a temporal evolution, we iteratively modify  $\Phi$  and get

$$\frac{\partial \Phi}{\partial t} + F |\nabla \Phi| = 0 \quad (2.8)$$

with  $F$  as speed function. The speed function controls the image influence on the evolution of the level set surface. This formulation leads to the need to solve partial differential equations (PDE). To define the speed function, we use the surface's normal  $N$  and curvature  $k$ , defined as:

4. <http://ktchu.serendipityresearch.org/software/lsmlib/index.html>

$$N = \frac{\nabla\Phi}{|\nabla\Phi|} \quad k = \operatorname{div} \left( \frac{\nabla\Phi}{|\nabla\Phi|} \right). \quad (2.9)$$

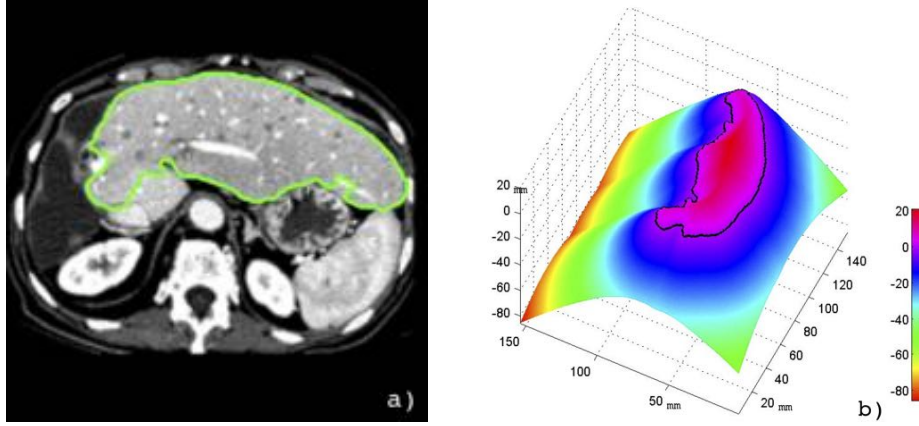


Figure 2.18 – Level set (b) and corresponding iso-surface (left) during liver segmentation.

Osher and Sethian [OS88] proposed a level set approach that exploits edges. Image gradients are used to slow down the expansion of the model. The speed function is defined as the sum of a curvature force and an expansion force. This force is scaled by a stopping function that makes use of the image gradient at this point.


Using these speed values, the new distance map can be updated. The resulting contour is the input for the next iteration step. This approach does not stop the contour at high gradients completely but only slows it down. Therefore, it is possible that after a large number of steps, the contour could overcome high gradients. This is a behaviour that should always be taken into consideration. Figure 2.19 shows an example of a patella segmentation using an edge stopping level set approach.

Instead of a gradient stopping function, an alternative approach defines an energy potential that is consequently minimised. The foundational work on this approach has been done by Mumford and Shah in 1989 [MS89]. The area and volume covered by the level set  $S$  are important in this approach and can be determined using the Heaviside function, which needs to be smoothed as described by Zhao *et al.* in [ZCMO96]. The calculation is performed in three steps. First, the average image intensity for voxels inside and outside the segment is calculated. The next step is to compute the energy of each voxel using the weighted area, volume and image intensities inside and outside the region of interest. Finally, the distances are being updated using these energies.

### 2.9.2.1 Extensions

Over the years, several extensions to level sets have been proposed. The normal iteration of level sets is limited to single voxel distances per step. Sethian [CIKW99] developed an extension for fast marching, allowing the contour to move more than one voxel per iteration step. This was previously prevented by numerical instabilities. In 2001, Chan and Vese [CV01] have presented





**Figure 2.19** – Example of an edge stopping level set (green) for segmentation of the patella after (a) 10, (b) 20, (c) 30, (d) 40, (e) 60 and (f) 80 steps. It moves towards the high gradient at the boundary of the bone.

an approach to level sets that is not based on an edge-function and can operate on very noisy, unsmoothed images. Zhang *et al.* [ZMSM08] showed a hybrid approach that uses both boundary and region information to get accurate results by reducing leaking problems. Yeo *et al.* [YXSN09] have proposed a new external force in 2009. Their geometric potential force is based on hypothesised interaction between relative geometry and image gradients. Caselles *et al.* [CKS95] have proposed geodesic active contours as a connection between snakes and geometric curves.

Other extensions focus on performance optimizations. In 1999, Adalsteinsson and Sethian [AS95] have proposed only to calculate the speed function for a narrow band around the current contour. Whitaker [Whi98] showed a computational method (sparse-field algorithm) that combines the level sets approach with the efficiency and accuracy of parametric representation. A solution to use level sets without the need for solving PDEs has been proposed by Shi and Karl in 2005 [SK05]. Finally, Unger *et al.* [UPB08] have extended geodesic active contours with GPU usage to improve performance.

### 2.9.3 Discrete Deformable Models

Unlike snakes and level sets, discrete deformable models do not use a continuous representation of their surface. Instead, they are modelled using particles or meshes (see section 2.9.3.2). In a discrete deformable model, forces are calculated in an iterative process and applied to the particles or vertices in the mesh. Several physical and numerical schemes can be used for this purpose while an abort criterion terminates the segmentation. When dealing with meshes, special attention has to be paid to local resampling, topology adaptation and the avoidance of (self-)

intersections and collisions. Kainmueller *et al.* [KLZ<sup>+</sup>08] proposed coupled deformable models. They exploited redundant structures that appear when segmenting multiple connected objects in parallel. They demonstrated a model for femur and ilium that reduces the need for expensive self-intersection testing. In their approach, the logical relationship can be retained using combined intensity profiles.

Bredno *et al.* [BLS03] gave a description of the components needed for a discrete deformable model. It consists of a surface, forces, a surface resolution adjustment approach and an abort criterion. In an iterative process, forces are calculated and applied. Afterwards, the surface is adjusted. The abort criterion defines the iteration end. This process is illustrated in Figure 2.20. The calculation of forces, e.g. the ones described in Section 2.9.5, is the first step. If a force is calculated for an edge or a triangle, it has to be distributed to the corresponding vertices. The next step is the movement of the vertices. The translation of each vertex is determined according to the selected numerical implementation. If the volume or shape of the models changes significantly, resampling processes are often needed. They adjust the mesh to ensure the geometrical quality of the triangles and modify the mesh in a way that triangles are small enough to fit into the desired structures. For example, the segmentation of a small vein needs a constant refinement as the model grows along the tubular structure. In a final step, a check is performed to see if the iteration process should be terminated. This can be triggered for example after a certain number of steps or once the vertices have reached equilibrium.




Figure 2.20 – Process of the iterative deformation of a surface (a-e) using forces.

### 2.9.3.1 Numerical Implementation

There are several possibilities to implement the physical properties of a discrete deformable model. The forces may be modelled directly or as Hookean spring forces, as presented by Gilles and Magnenat-Thalmann [GMT10] and Schmid [SKMT11]. In their formulation, a translation from the point  $x_i$  to  $x_j$  is weighted by a factor  $\alpha_h$ :

$$F_h = \alpha_h(x_j - x_i) \quad (2.10)$$

A common approach to describing the forces is by applying Newtonian physics, but a popular alternative is to use Lagrangian dynamics [SVG02]. Other approaches include the Finite Difference Methods (FDM) [TPBF87] and Finite Element Methods (FEM) [MT95, CCA92, CC93, TM91].

### 2.9.3.2 Mesh Types

Simple segmentation algorithms often create point clouds as a result. To initialise a deformable model, an initial shape has to be provided, which can be derived from point clouds. The conversion from a point cloud can be done using the marching cube algorithm as proposed by Lorensen and Cline [LC87] in 1987. The resulting mesh will be closed but also will contain a high number of faces, higher than needed in most cases. In this case, reduction techniques can be applied, but they will have an impact on the mesh quality. Miller *et al.* [MBL<sup>+</sup>91] have proposed another approach, the Geometrically Deformable Model (GDM) in which the mesh of a deformable model is grown into the point cloud. This results in a lower number of triangles while preserving the topology and details. Most discrete deformable models use traditional triangle meshes. Delingette [Del94, Del99] proposed to use a simplex representation to store the model geometry. A  $n$ -simplex mesh contains vertices that all have  $n+1$  distinct neighbours. Arbitrary topologies can be represented using 2-simplex meshes (each vertex has exactly three neighbours). Faces can have an arbitrary number of vertices and can be non-planar. A comparison between simplex and triangle meshes is shown in figure 2.21. Simplex meshes have been used for example in the segmentation of cardiac structures by Montagnat *et al.* [MD05] and for musculoskeletal structures in the lower limb by Gilles [GMMT06].




Figure 2.21 – Triangle (a) and simplex (b) meshes

To improve iteration performance, Lachaud and Montanvert [LM99] proposed a coarse to fine model while Snel *et al.* [SVG02] have demonstrated a multi-resolution scheme. Pons *et al.* [PB07] demonstrated an implicit triangular mesh that performs a Delaunay triangulation “on-the-fly” when needed. A permanent connection through continuous tessellation has been presented by Gilles *et al.* [Gil07].

### 2.9.3.3 Particle Systems

A different approach to discrete deformable models are particle systems which operate meshless. A collection of independent particles gets assigned physical properties, such as mass, position, speed and acceleration and evolve according to classical Newtonian mechanics.

Particle systems can represent arbitrary topologies but pose a great challenge when it comes to visualisation and the definition of surfaces. Szeliski and Tonnensen [ST92] have proposed to use oriented particles to represent surfaces. They use several techniques to derive a surface from the particles. This approach has been extended by Lombardo [Lom04] to define curvature and to get

an implicit surface description.

## 2.9.4 Knowledge-based Deformable Models

Deformable models require initial shape information to achieve better results. Knowledge-based deformable models are a group of approaches that aim to invoke more prior knowledge through additional features. This will increase the flexibility in poor images and the robustness against inaccurate initialisations. A comprehensive overview of knowledge-based deformable models has been given by Schmid [SKMT11] in 2011. A simple approach would be assigning labels with an associated behaviour to the model. Other models use the statistical or probabilistic variation of selected features. In this section, possibilities for feature selection are discussed, followed by the alignment process and construction of the statistical model. We conclude with the description of the two main approaches: active shape models and active appearance models.

### 2.9.4.1 Feature Selection

Features select the property of the models that will be exploited. They can be summed up in three categories: shape-, appearance- and transformation-based features.

**Shape-based Features.** Cootes *et al.* [CHTH94, CTCG95] have proposed shape-based features in 1994. They use prior shape information to improve results and limit shape variability to the shape variations from training data. A set of points, called landmarks, represents image features. They have to be (usually manually) selected from a large number of training data. Another approach for these point distribution models (PDM) can be to replace landmarks with parameters of a medial axis [YZG05].

#### **Appearance-based Features.**

To exploit image properties, appearance-based features can be such as intensity, gradient, texture, and momentum. Intensity profiles (IP, see section 2.9.5.2) are very important and commonly used. In 3D, appearance-based features suffer from an increased complexity (highly increased memory and computational effort).

#### **Transformation-based Features.**

A third group of approaches are transformation-based features; they are not used commonly. They exploit transformation parameters that can be used for the statistical deformation model. They have also been extended to support non-rigid transformations.

### 2.9.4.2 Construction

The creation process of statistical or probabilistic models of feature variation can be divided into two parts. Initially, the training models have to be aligned before the construction can be performed.

The alignment process aims at eliminating pose changes that are not feature related. To do this, commonly the iterative generalised Procrustes approach [Gow75] is applied.

$$T^* = \underset{T}{\operatorname{argmin}} \sum_i \|y_i - T(x)\|^2 \quad (2.11)$$

The transformation  $T^*$  is obtained by finding a transformation that minimises the distance between a model  $x$  and all other models  $y_i$ . The resulting transformation  $T^*$  is a combination of translation, rotation, and scaling. Once this transformation is applied, all models are in a common coordinate frame: the model space of the point distribution model. Although aligned, these models still have a substantial amount of shape variability.

The actual construction process uses the corresponding and aligned features and performs a principal component analysis (PCA). PCA has been introduced by Pearson [Pea01] in 1901. This technique has been used for the point distribution model (PDM) by Cootes and Taylor [CHTH94]. Given  $n$  shapes  $Y_1, Y_2, \dots, Y_n$ , the mean shape can be defined as

$$\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i . \quad (2.12)$$

The covariance matrix is defined as

$$S = \frac{1}{n-1} \sum_{i=1}^n (Y_i - \bar{Y})(Y_i - \bar{Y})^T \quad (2.13)$$

and its eigenvectors correspond to the largest eigenvalues of  $S$ . The eigenvectors describe the most significant variation modes from which a subset (much smaller than the number of points) is used.

This approach needs a sufficiently high number of training models. Otherwise, a poor generality is given, leading to poor adaptation to new data and enforcement of an over-constrained behaviour. Approaches to solving this problem include Bayesian inference [ER00], synthetic shape variations [CT95] or the derivation of new shapes from existing training shapes [TKLL06].

#### 2.9.4.3 Active Shape Models and Active Appearance Models

Active Shape Models (ASMs) use a statistical model of shape feature variations. To implement them, a model fitting process constrains movement to the modes defined by the trained models. In each iteration step of this process, the closest parametric model to the model  $X$  with the deformation  $dX$  is searched for and forces are applied to deform the model towards this shape. Given the model shape  $Y$ , the model  $X$  can be represented as follows:

$$X = MY + X_c \quad (2.14)$$

with  $M$  being the matrix for scaling and rotation and  $X_c$  the translation vector of the centre of  $X$ . The estimation of the pose adjustment can be done efficiently using a standard least-squares approach as described in [CTCG95]. An example for the segmentation of lung and cerebellum using

ASMs is given by Ginneken *et al.* [GFS<sup>+</sup>02]. Lamecker *et al.* [LKW<sup>+</sup>07] presented an approach to use ASMs for the segmentation of the bony orbit.

Active Appearance Models (AAM) have been proposed by Cootes *et al.* [CET98] in 1998. They combine a statistical model of the shape and the grey-level appearance. This extends the models for shape  $X_s$  with another one for the image texture  $X_t$ , assuming that all textures have been warped to the mean shape. Olabarriaga *et al.* [OBN04] have presented the use of AAMs for thrombus segmentation; Shen *et al.* [SLH11] have extended the approach to active volume models.

#### 2.9.4.4 Conclusion

Knowledge-based deformable models are a powerful tool for segmentation. They can express a wide range of valid shape variations and can enforce proper model evolution in areas of low information or in the vicinity of artefacts. However, they rely on large collection of required training data. While some training data can be derived from fewer models, the required amount of models is very high. Since there is very little imaging of the whole leg, there is not sufficient training data available for muscle segmentation. Furthermore, muscles are prone to high shape variations, which also depend on the pose during acquisition.

### 2.9.5 External Forces

Forces determine the deformation of the model. They bridge the gap between the model with its internal forces and the image domain. Over the years, these external forces have evolved from intensity or gradient based image forces to complex approaches using statistics or vector flow techniques. In this section, we will present some basic forces that are commonly used in deformable models. Advanced forces follow as an extension. We conclude with the description of interactive forces. All following forces will be calculated for a vertex  $x$  with the normal  $n_x$  and will be scaled by a factor  $\alpha$ .

#### 2.9.5.1 Basic Forces

##### Pressure Force

In areas with low image information, the model tends to stay at its current position as there are no image features to attract it. To overcome this issue and to reduce the need for very close initialisation, Cohen *et al.* [Coh91] have proposed a balloon or pressure force.

$$F_{\text{pressure}}(x) = \alpha_p n_x \quad (2.15)$$


It evolves the model along its surface normals and depending on the sign of  $\alpha_p$  the model can either shrink or expand.

### Laplacian Smoothing Force

Image noise or artefacts may stop small parts of the model. To prevent this from happening, model smoothing can be applied. An example for a common approach is the Laplacian smoothing:

$$F_{\text{smooth}}(x) = \alpha_s \frac{1}{|U_x(\Delta)|} \sum_{y \in U_x(\Delta)} y - x \quad (2.16)$$

It calculates the centroid of the neighbouring vertices and defines a force towards that point. Improvements have been proposed by Shen *et al.* [SHL<sup>+</sup>11]. Figure 2.22 shows an example for smoothing (transformation from cube to sphere) and pressure force (growing).



**Figure 2.22** – The combination of pressure force and smoothing transforms a cube into a sphere.

### Simple Image Force

A simple approach to include image features is the use of normalised intensity gradients that, for example, can be calculated using central differences. Along the vertex normal, several gradients are being calculated and the force will be directed towards the higher gradients:

$$F_{\text{image}}(x) = -\alpha_i n_x \sum_{\delta_{\min}}^{\delta_{\max}} k(\delta) * \Phi(x + \delta \cdot n_x) \quad (2.17)$$

where  $k(\delta)$  is a weighting factor on the distance and  $*$  is the convolution operator. An image interpretation factor  $\Phi$  can be calculated, in this case, the gradient.

#### 2.9.5.2 Advanced Forces

The basic forces can be extended by more advanced approaches to improve results and runtime.

### Image Intensity Profiles

Instead of simply using the image gradient, intensity profiles can be used to compare the intensity curve along the normals. By doing so, the desired structure can be described more powerful and it allows a distinction between noise-induced gradients and targeted edges. This approach is commonly used [BMVS02, CHTH94, CT01]. Clustering has been proposed to reduce the complexity of the high number of intensity profiles. For instance, Chung and Delingette [CD09] have

proposed an Expectation-Maximization-based algorithm for clustering and classification. An example of image intensity profiles in a magnetic resonance (MR) image of the femur is shown in Figure 2.23.




Figure 2.23 – Example of intensity profiles sampled in an MR image of the femur.

### Multiscale Gaussian Potential Force

The Gaussian Potential Force is designed to attract a model towards image features and is defined as

$$F_g(x) = \alpha_g |\nabla [G_\sigma(x) * I(x)]|^2 \quad (2.18)$$

with  $I(x)$  the intensity image,  $G_\sigma(x)$  a Gaussian function,  $\alpha_g$  a weight and  $*$  the convolution operator. This force has been extended by Terzopoulos *et al.* [TWK87] to a multi-scale scheme. To overcome the need of an initialisation close to the final contour, they propose to use a large initial value of  $\sigma$  to broaden the search space. Once equilibrium has been reached,  $\sigma$  could be decreased to maintain the accuracy of the original approach. Until now, no criterion has been established to determine when to reduce  $\sigma$ , limiting the utility of the multi-scale Gaussian potential force.

### Distance Potential Force

To extend the attraction range, Cohen and Cohen [CC93] have proposed to use a distance map in 1993. The values in this map are obtained by using either the Euclidian distance [Dan80] or Chamfer distance [Bor84] to calculate the distance between a voxel and the closest boundary point.

### Dynamic Distance Force

This force extends the distance potential force to include a larger spatial area around the surface [Sap95, AS99]. The dynamic distance has improved handling of boundary concavities. It is calculated by examining the image for features or gradients along the surfaces normal. The maximal search distance is limited by a threshold  $D_{\max}$ .



$$F_{\text{dynamic}}(x) = \alpha_d \frac{D(x)}{D_{\text{max}}} n(x) \quad (2.19)$$

The resulting force can pull the model towards distant image features. Nonetheless, the search is very time-consuming and has to be repeated during every step. Lowering the threshold can shorten the runtime but also reduces the attraction range for image features.

### Gradient Vector Flow

In 1998, Xu and Prince [XP98] have proposed a new external force model. Their Gradient Vector Flow (GVF) field is calculated as the diffusion of an intensity image. It allows a more flexible initialization and supports a more efficient convergence to concavities.

Ng *et al.* [NFO<sup>+</sup>07] present a medical image segmentation that uses a feature-based GVF snake. The iteration is stopped once the accuracy is sufficient by exploiting image features. The work of Zhao *et al.* [ZCZ11] has improved the dynamic GVF force field and introduced a strategy of deformable contour knots for a B-spline based model.

### Omnidirectional Displacements

When working with deformable surfaces, forces are commonly directed along a line, usually the surface's local normal. Kainmueller *et al.* [KLH<sup>+</sup>12] have proposed omnidirectional displacements for deformable surfaces (ODDS) that consider a sphere around each vertex. By doing so, a global optimisation can be performed. This technique has been proven to be useful in regions of high curvature, e.g. tips. They have also proposed a hybrid approach, *fastODDS*, to overcome the high memory and runtime requirements.

#### 2.9.5.3 Interactive Forces

Image artefacts, different protocols and implants can cause problems in the automated segmentation of medical images. In a clinical environment, operators can guide the deformable model to overcome these problems. The so-called interactive forces provide a link between real-time user input and model iteration.

Kass *et al.* [KWT88] have proposed two interactive forces. Spring forces are designed to pull the model in the direction of a point  $p$ . Their strength is proportional to the distance from  $p$ :

$$F_{\text{spring}}(x) = \alpha_s(p - x) \quad (2.20)$$

The opposite effect can be achieved using volcano forces. They push a model away from a point  $p$ :

$$F_{\text{volcano}}(x) = \alpha_v \frac{r}{|r|^3} \quad (2.21)$$

with  $r = x - p$ . Interaction forces often are only computed for a small neighbourhood to reduce computational costs. When using all points, the complexity is of order  $O(n^2)$ , while a small neighbourhood  $m$  with  $m \ll n$  can reduce the complexity to order  $O(n \cdot m)$ . Other approaches include smart paint [MS12], and live wire [PHA08]; an overview can be found in [ZX13].

### 2.9.6 Conclusion

We have presented an overview of the current state of deformable models in medical image segmentation. They are a powerful tool for the given image conditions and, therefore, play an important role. We have surveyed current snakes and level sets methods that are examples for continuous shape representation. These usually two-dimensionally used approaches achieve good results with high computational costs. However, they are highly dependent on a good initialization. Contrary to snakes, level sets can adapt to different topologies automatically. Discrete deformable models are three-dimensional approaches, which can be represented as particle systems or through different mesh types. They can also be based on several physical foundations. They provide great flexibility and have proven useful in many cases. Depending on the structure of interest, deformable models have to be carefully configured and parametrised. An extension to deformable models is the incorporation of prior knowledge to make them more robust in complex image conditions. These knowledge-based deformable models exploit prior knowledge on shape and appearance. They need a large training base and complex training but are very robust. We have presented several forces that cover the tools needed for creating deformable models.

## 2.10 Multi-Channel Image Processing

In the previous section, we have found that deformable models are a good segmentation tool that can provide comprehensive and robust results. However, like all methods, they rely on the amount of information carried by the input images. For optical acquisition, it is possible to increase this information by using colour- instead of grey-scale-images. With special acquisition sequences, it is possible to generate extra channels of information in MR scanners. This additional information benefits segmentation approaches, including the ones we have discussed in the previous sections. In this section, we will discuss how these methods benefit from multi-channel data and how they use it. We have grouped promising approaches based on the grouping in the previous sections: statistical, texture-based, model-based, registration, and other.

Statistical segmentation approaches work on image features. The feature space can be largely increased by using multi-channel input. Pauly *et al.* [PGC<sup>+</sup>11] use random ferns to produce multi-dimensional regression output that can be used to detect and locate organs. Their approach works on multi-channel dixon images. Geremia *et al.* [GMC<sup>+</sup>10] use spatial decision forests in multi-channel MRI data to detect multiple sclerosis lesions in the human brain. A probabilistic labelling approach using multi-shapes of pre-trained priors is presented by Changizi *et al.* [CH10]. Grady [Gra05] use a random walker and priors for image segmentation in different modalities. Pham [Pha03] present an unsupervised classification method that uses edge-clustering. Ange-

line *et al.* [AIJL02] propose an improved statistical model for multi-channel images. They use a fuzzy segmentation process and test it on the Visible Human data set and on brain MRI. Zhang *et al.* [ZBS01] demonstrate a segmentation of multi-channel brain images. They use a Hidden Markov Random Field model and the Expectation-Maximization Algorithm to generate their results. An irregular Markov Random Field has been proposed by Smits and Dellepiane [SD97]. They apply the MRF to handle properly low signal areas and use optical remote sensing data to validate their approach. Bourgeat *et al.* [BFS<sup>+</sup>07] propose to use the MR phase information as an additional input channel.

Derived from statistical segmentation are texture-based approaches. The additional information gives more means to express the texture and makes it easier to define and detect it. A study of local features and kernels has been done by Zhang *et al.* [ZMLS06]. They are used for texture classification and object categorization. Their work uses colour, hence multi-channel, images. Lakshmanan proposes [Lak01] a texture analysis method. They use a hierarchical, multi-scale scheme and apply it to the segmentation of real-world scenes. Randen and Husoy [RH94] propose a multi-channel filtering approach improved image texture segmentation. Dubuisson and Gupta [DJG00] propose to use texture fusion in aerial images for GIS applications.

Some work has been done on using model-based approaches in multi-channel images. Bogovic *et al.* [BPB13] present a geometric deformable model that can use multiple objects in parallel and works on colour images (multi-channel). Lehmann *et al.* [LBS03] give a review of active contours in medical image segmentation in single- and multi-channel images. They classify the methods by dimensionality, object representation and parameter selection. Xu *et al.* [XHY02] use a Markov random field based active contour to segment multi-channel images. They use colour images and multi-contrast MRI. A level set approach is followed by Vese and Chan [VC02] that uses multiple-phases. They test their framework on different MRI and colour images. Multi-channel brain data has been segmented by Vacheta *et al.* [VYBS95] using a deformable active contour model. Zhu *et al.* [ZLY96] propose to use multi-band information to improve image segmentation using snakes.

Multi-channel information may also be used to improve registration results. Heinrich *et al.* [HJB<sup>+</sup>12] introduce their modality independent neighbourhood descriptor that can be used for multi-modal deformable registration of an atlas. Multi-modal here means that either multiple data sets from different devices like CT and MRI, or different sequences of the same modality can be used. Their work shows better results than other techniques like entropy images or conditional mutual information. Guo [Guo06] shows the use of multiple modalities in GVF for image registration. Chang *et al.* [CSV00] present an active contour without edges which they apply on vector-valued (multi-channel) images. They present an example on the identification of nebula in colour space images. Skull segmentation is shown by Rifai *et al.* [RBH<sup>+</sup>00] using deformable models that compensate for partial volume effects in their MR images.

There are several other segmentation approaches that have been used with multi-channel images. Plaza *et al.* [PPVM06] present their work of using mathematical morphology for the segmentation

of multi-channel images. A hierarchical approach to colour image segmentation has been shown by Cheng and Sun [CS00]. They use homogeneity to identify the sub-parts of colour images. Heigl *et al.* [HPN99] work on the tracking of objects in multi-channel images. Malladi *et al.* [MSB96] discuss the use of the multiple channels of colour images in their min/max curvature flow filter. Pien and Gauch [PG94] also work on multi-channel brain data, using variational segmentation. The work of Fechter *et al.* [Fec11] is about a manual segmentation approach that has time as an additional channel.

**Table 2.4** – Classification of multi-channel segmentation approaches.

Approach	Publications
Statistical	[PGC <sup>+</sup> 11], [GMC <sup>+</sup> 10], [CH10], [Gra05], [AIJL02], [ZBS01], [SD97], [Pha03]
Texture-based	[ZMLS06], [Lak01], [RH94]
Model-based	[BPB13], [LBS03], [XHY02], [VC02], [VYBS95]
Registration	[HJB <sup>+</sup> 12], [Guo06], [CSV00], [RBH <sup>+</sup> 00]
Other	[PPVM06], [MSB96], [CS00], [HPN99], [PG94], [Fec11]


In conclusion, we have found that the use of multi-channel over single-channel input provides the segmentation approach with more information. This leads to more robustness and better results. Furthermore, some structures become visible which have been invisible in other channels. We list an overview of the approaches and publications in Table 2.4. The improved results, stemming from multi-channel input and the ability of MRI to generate this, motivate our choice to work with multi-channel MR images. In this section, we have shown that several classes of segmentation techniques benefit. However, while benefiting from multi-channel information, the disadvantages as presented in the previous sections, still prevail. Based on this conclusion, a model-based approach, in particular, deformable models, seems to be the most appropriate for the segmentation of musculoskeletal structures of the lower limb in multi-channel images.

While many segmentation approaches have been tested or used with multi-channel images, deformable models have not yet been used with them for muscle segmentation. Muscles are hard to delineate and current segmentation approaches are computationally costly (Section 2.12). Therefore, we expect the combination of deformable models, which have lower computational costs and are robust, with multi-channel image data, which carries additional information to help finding the fine delineation between individual muscles, to show improved results in comparison to the current state of the art. To further investigate the performance of our chosen approach, we will briefly present the anatomy of the lower limb with a focus on muscles before we discuss related work in muscle segmentation.

## 2.11 Muscles of Thigh and Calf

Muscles can be divided into three types: skeletal, cardiac (heart muscles) and smooth muscles (involuntary/subconscious contraction in the walls of hollow organs, e.g. blood vessels). Skeletal muscles can be contracted consciously and are attached to bones (through tendons). They are formed by many individual compartments, bundles of muscle fibres or myocytes, that all share the same orientation. Between these compartments can be adipose tissue. Skeletal muscles have

an origin and an insertion point. A visualisation of the structure of muscles can be seen in Figure 2.24.



**Figure 2.24** – Muscles are attached to the bones through tendons. They are composed of a large number of muscle fibres which are grouped into fascicles. To nurture the muscle, it is crossed by fine blood vessels. Source: NIH National Cancer Institute, USA

Muscles are stimulated by electrical impulses, transmitted through nerves. The maximum force possible depends on the cross-sectional area but is often increased through the use of levers, e.g. quadriceps through patella. Skeletal muscles usually work as a pair, so that there always is an opposite movement possible, e.g. for extension and flexion. For many movements more than just a pair of muscles is involved; muscles show a complex interplay forming groups for certain movements. This flexibility results in a wide range of motions that are possible with a limited number of muscles.

### 2.11.1 Muscles of Thigh

The thigh can be subdivided into three muscle compartments; anterior, medial and posterior compartment. Figure 2.25 shows the muscles of the thigh, from an anterior and posterior perspective as well as a traversal, mid-thigh profile.

The **anterior compartment** of the thigh contains the muscles *sartorius*, and the *quadriceps*, consisting of *rectus femoris*, *vastus lateralis*, *vastus intermedius*, and *vastus medialis*. Their main task is knee extension and hip flexion. They originate from hip and femur and insert to the tibia, via the patella where the quadriceps tendon turns into the patella ligaments.




Figure 2.25 – Muscles of the thigh. Source: Henry Gray (1918) *Anatomy of the Human Body*

The **medial compartment** contains the adductor muscles of the hip and is formed by the *gracilis* muscle and the *adductors*. Depending on the source, the *pectineus* and *obturator externus* muscles can also be considered to be part of the medial compartment. The *adductors* are formed by these muscles: *adductor longus*, *adductor brevis*, *adductor magnus*, and *adductor minimus*. Together they are used to pull the thigh medially, to support the trunk during sideways movements and they contribute to thigh flexion.

The **posterior compartment** consists of three muscles located in the posterior part of the thigh. These muscles are *semitendinosus*, *semimembranosus*, and *biceps femoris* with its long and short head. Together with their tendons, they are known as hamstrings. They connect knee and hip joint and play a crucial role in common movements like walking and running. They act as antagonist to the quadriceps in the deceleration of the knee extension.

The different compartments are separated by the *medial intermuscular septum* and the *lateral intermuscular septum*. The medial intermuscular septum of thigh separates the anterior compartment of the thigh from the medial compartment; it is located between the vastus medialis, and the adductors. The lateral intermuscular septum separates the anterior compartment of the thigh from the posterior compartment; it is located between the vastus lateralis and biceps femoris.

### 2.11.2 Muscles of Calf

The muscles of the calf can be divided into four compartments; the **anterior compartment**, the **lateral compartment**, and the **superficial and deep posterior compartment**. These compartments, along with a transversal profile are shown in Figure 2.26.

The **lateral compartment** contains two muscles, *fibularis longus* and *fibularis brevis*. They perform




Figure 2.26 – Muscles of the calf. Source: Henry Gray (1918) *Anatomy of the Human Body*

the extension and flexion of the foot.

The **anterior compartment** is formed by *tibialis anterior*, *extensor hallucis longus*, *extensor digitorum longus*, and *fibularis tertius*. They are used to dorsiflex and extend the digits and the great toe. The *fibularis tertius* dorsiflexes the ankle and aids in the eversion of the foot.

The **superficial posterior compartment** is formed by the *plantaris* muscle and the *triceps surae* muscle pair, consisting of the *gastrocnemius* and the *soleus*. The *accessory soleus* and the *Achilles tendon* are considered to be part of it. They stabilise the ankle and are an important part of forward locomotion. Jumping is another movement that relies on the contribution of the *triceps surae*.

The **deep posterior compartment** contains the *popliteus* muscle and the muscle in the *tarsal tunnel*, which are *flexor hallucis longus*, *flexor digitorum longus*, and *tibialis posterior*. The *popliteus* unlocks the knee during walking; the other allow the bending of the foot.

These compartments are separated by the **anterior, posterior, and transverse intermuscular septum**. The **Anterior crural intermuscular septum** separates the lateral from the anterior compartment of the calf. The **Transverse intermuscular septum** is located between the deep and superficial posterior compartments. It is connected to Tibia and Fibula. The **Posterior crural intermuscular septum** separates the *Soleus* from the deep posterior and the lateral compartment.

## 2.12 Muscle Segmentation

The reconstruction of anatomical structures from CT or MRI is a complex task, as demonstrated by Gargiulo *et al.* [GHR<sup>+</sup>14]. They present three studies on evaluating muscle density changes,

impact on bone and muscle after hip replacement and the electrical behaviour of the brain. The approaches for muscle segmentation fall into three main categories: classification, model-based, and atlas-based.

Classification approaches can be used to detect muscle tissue and with additional processing to identify muscles. Baudin *et al.* [BACP12b, BACP12a] propose a mathematical formulation that uses prior knowledge in a random walks formulation. Their statistical shape atlas is then coupled to it. Their results show potential but still need improvement. They use a graphical model to determine automatically seed locations. This work is based on their previous approach [BACP12c]. Kober *et al.* [KGZS11] show a method to identify muscle tissue in 4D MRI data of the upper arm. They use it for a visualisation of deformations during flexion. Mirzaalian *et al.* [MHHR11] use a manual segmentation of knee flexors and extensor for shape analysis in diagnosis. Barra and Boire [BB02] present a possibilistic clustering algorithm that reliably separates fat from muscle tissue. A compartmental labelling approach is tested by Thomas *et al.* [TNL<sup>+</sup>14]. Andrews *et al.* [ACH14] have proposed a new log-ratio transform for mapping probabilistic labels to vectors for learning statistical models. They present their segmentation on the thigh. A statistical model for the segmentation of the rectus abdominis muscle in CT images has been proposed by Kamiya *et al.* [KMZ<sup>+</sup>13]. The classification of muscle tissue in the thigh has been shown by Imamoglu *et al.* [IGTH<sup>+</sup>15] in T1 and T2 images. They achieve good results but do not identify individual muscles.

Model-based approaches carry a low amount of information about the structure to be segmented. A model-based approach for skeletal muscle segmentation in MR images has been proposed by Jolivet *et al.* [JDR<sup>+</sup>14]. They use a segmentation in several slices which is then interpolated through the data set. Firkbank *et al.* [FHWC01] show a method that uses active contours to segment the extra-ocular muscle. Andrews *et al.* [AHY<sup>+</sup>11] present a probabilistic approach for extension of knee extensor and flexor muscles. They use shape-prior that can indicate regions of uncertainty. Gilles [Gil07] proposed a deformable model for muscle segmentation in low-resolution MR images. However, it has long calculation times.

A recent work of Avendi *et al.* [AKJ16] explores the combination of deep learning (DL) and deformable models to achieve a fully automatic segmentation of the left ventricle. The DL is applied to the ground truth and learns the shape of the structure. This knowledge is applied during the deformable model segmentation.

Using an atlas with annotated muscles has been proposed by several groups for muscle segmentation. Ahmad *et al.* [AYDM14a, AYDM14b] present an approach for thigh where they identify muscle tissue in general and muscle compartments. Individual muscles are not addressed. Other works have addressed the individual muscle segmentation in the rat leg. The work by Sdika *et al.* [STL<sup>+</sup>15] present a multi-atlas approach that benefits from imposed plane restrictions which lead to high accuracy and reliability.

All segmentation approaches require manual segmentation for training and evaluation. Barnouin *et al.* [BBBV<sup>+</sup>13] assess the reliability and quality of manual segmentation of the rectus femoris



muscles. The results show that manual segmentation is a feasible way to evaluate the quality of automatic approaches. An approach to simplify the manual segmentation task has been proposed by Ray *et al.* [RMN<sup>+</sup>14]. Their framework is built on previous manual segmentations and a matching of that database using a cylinder projection of the thigh.

## 2.13 Conclusion

In this review of related work, we have looked at different kinds of segmentation approaches, how multi-channel data can improve segmentation results and current approaches to musculoskeletal modelling.

We have found general segmentation techniques (threshold-, region-, edge-, and texture-based) to be not sufficient for musculoskeletal segmentation. Morphological operations have been shown to be powerful in the labelling of tissue types. However, identification of individual structures is not possible. We have then discussed classification, registration and graph-cuts and found that all these approaches lack a comprehensive model understanding needed for muscle segmentation. Model-based approaches resolve this. Deformable models, in particular, are a powerful tool for model-based segmentation. We have described several types.

Modern medical imaging devices allow creating images that have multi-channel information. Multi-channels can result from colour images (RGB) or special acquisition sequences in MR scans (T1-, T2-weighted brain scans, mDixon scans). After reviewing the literature, we have found the additional information from multi-channels can improve segmentation outcome and the quality of results.

Finally, we have reviewed current approaches for muscle segmentation and have identified the main three categories they fall into classification, atlas-based and model-based approaches.

Several imaging techniques are widely available nowadays. We have reviewed them and their working principles and have found MRI to be most suitable musculoskeletal anatomical modelling. It can provide high-resolution data in reasonable durations and the multitude of protocols allows the highlighting of soft tissues. Furthermore, it is possible to create multi-channel image data in a single scan. This provides additional input without increasing acquisition time. Therefore, we propose to use the mDixon sequence which produces four different channels.

Based on our survey of related work, we have found that there is only little use of multi-channel images. While multi-channel data provides additional input for all kind of applications, most approaches use them for classification or texture analysis. We have found only little application with model-based techniques and none in combination with deformable models.

As we have found, deformable models are a robust and flexible technique for segmentation. We have presented related that uses deformable models for muscle segmentation. however, muscle segmentation has been found to be complex task due to the anatomical complexity. The contrast in soft tissues is often low and the distinction between individual muscles is complicated with very little visibility. We propose to combine deformable models with multi-channel image

forces for musculoskeletal segmentation as our contribution. Multi-channel information requires a reconsideration of the full pipeline from acquisition, processing, and initialisation to final segmentation. With the growing amount of data, special attention has to be paid to the increased memory and calculation requirements to ensure timely results can be integrated into clinical processes. We will present our approaches in the following chapter.



---

## CHAPTER 3

# IMAGE ACQUISITION

---



### 3.1 Motivation

Image acquisition is a crucial step in the anatomical modelling process. Mistakes and false assumptions during protocol design and application can cause problems in all the later stages and may require inefficient workarounds or could impact segmentation quality. After careful consideration of several imaging techniques (see Section 2.2), we have chosen MRI and have developed an acquisition protocol that is tailored for our application. To achieve best results, we have used surface coils for signal amplification and have designed a special support for them to allow us to generate seamless data sets. We are using two different acquisition protocols that highlight muscle boundaries to generate the multi-channel data used during segmentation. The raw image data from the scanner has to be processed to be used in the segmentation. The image data has been acquired in several stacks, due to the size of the coil used. We present a method stitch these stacks together. The total scanned volume is larger than the subject and related to the field of view of the scanner. We discuss an approach to detect the surrounding air and crop the image data so that only the subject remains to reduce the storage requirements. This is crucial for our high-resolution data that is multiplied by the number of channels.

### 3.2 Image Acquisition

Image acquisition is a crucial part of the process; all following steps rely on the acquired image data. The acquisition relies on various factors; together they are called protocol. Based on our analysis in Section 2.2, we have chosen MRI as imaging modality. It gives good contrast for soft tissues while avoiding harmful ionising radiation. Our goal is to define a protocol for the musculoskeletal imaging of the lower limb. MRI offers us the right contrast. Furthermore, for the protocol, we have to define the sequences, the resolution, the number of sub-scans, the position of the subject and if we use additional equipment like surface coils for signal amplification.

The acquisition protocol is a compromise of several variables:

- **Resolution:** The image resolution determines the visibility of structures. It has an impact on the duration of the acquisition as it increases with the resolution. Higher resolutions increase the level of noise. However, coils can be used to minimise that impact.
- **Signal to Noise Ratio (SNR):** The signal to noise ratio is influenced by the resolution and the duration of the scans. Higher resolutions usually carry more noise. Shortening acquisitions also increases noise, whereas longer sampling during the scan can reduce it.
- **Structures:** The choice of visible structures can influence the selection of sequences for highlighting. To gain better visibility, other structures can be suppressed, which must be considered during protocol design. The amount of needed detail and the size of structures directly influence the choice of resolution and consequent acquisition duration.
- **Coils:** Additional surface coils can be used during acquisition to increase the signal and reduce the noise. These coils can also increase the resolution or reduce the acquisition duration. However, they have to be placed close to the subjects and might cause deforma-

tions. Coils have a limited field of view which can cause the need to move them during acquisition.

- **Duration:** The duration of the acquisition is not only an economical cost factor but also determines the resolutions that can be achieved. Some protocols can be time-consuming and, therefore, might be limited in resolution.

Based on these factors, we have developed a setup and have selected coils and sequences in collaboration with experienced radiologists. During the two year process, many optimisations have been made that lead to the current iteration which we will describe in the following sections.

### 3.2.1 MRI Setup

We use a Philips Ingenuity TF PET/MR scanner available in the facilities of the University Hospital of Geneva (HUGE) for the development of the scanning protocol. This scanner combines a 3T MRI with a PET device, connected by a turning scanner table. The combination of the anatomical imaging from MRI and the PET metabolic information in the same coordinate system, allows advanced diagnostic options. However, in our study, we have focussed our work on the MRI data because the PET data is of a low resolution which prohibits contributions to the segmentation. The volume of interest has been defined to be the lower limb, starting at the femoral head and ending at the ankle. During a later evaluation with gait specialists, we have expanded the volume to include the hip. The subjects are placed in a supine position on the scanner table. This causes some deformation on the subjects soft tissues due to the pressure of the body. This has been ruled favourable over a standing MRI, which would suppress this deformation but only gives lower resolution results. The subjects have been placed on cushions for comfort. Small supports have been placed under the knee joints, causing a small bending; this allows the subjects to carry out small movements which are required for a study performed in parallel (see [JTD<sup>+</sup>14]). The subject is automatically moved, feet first, by the table through the MR scanner as its field of view is smaller than our volume of interest.

Initial scans have raised the need for a higher resolution to distinguish fine anatomical details. Therefore, we have chosen to use an additional surface coil for signal amplification.

### 3.2.2 Coils

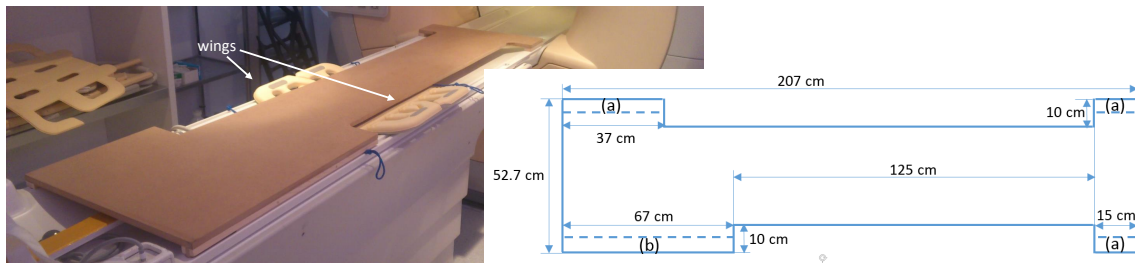
Acquisitions that cover large parts of the body pose a challenge to setups, especially when coils are used. Often these coils do not cover the complete volume of interest, and the acquisition has to be broken down into several sub-acquisitions, called stacks. The subject must not move in between to get a continuous image. Furthermore, special care has to be taken when using coils. Usually, they are put as close as possible to the subject, often directly on the skin. This results in deformations of the soft tissue not only on the surface but also on deeper layers like the muscles. When the coils are being moved, this deformation changes. This prevents the stitching into a single continuous image.

The coil used in this protocol consists of an anterior and posterior part. Our setup has two

separate parts to handle these coils. The posterior coil has been made movable by placing a board between the patient and the table, for the anterior coil we have developed another support structure. We will describe these structures in the following sections.

### 3.2.2.1 Posterior Coil


We have designed a board for the subject to lie on which is designed to allow moving the coil under the subject without body contact. It has the same dimensions as the scanner table and a key that fits the slot of the table. This way, it is fixed to the table and is stable while the table moves or the coil is repositioned.



**Figure 3.1** – Dimensions of the board have selected to match the table of the scanner. The indentation allows accommodating the wings of the coil, leaving additional space on one side for the connecting cable.

The board has indentations on both sides to accommodate the “wings” of the coil (see Figure 3.1). The indentation on the right side has been prolonged to fit the connecting cable of the coil. The length has been set to 125cm to cover subjects of any sizes from hip to toe.

Near the leg, the board is 32.7cm wide, which is sufficient for normal sized subjects. To increase the stability of the setup and to prevent the board from bending under the weight of a subject, flexible blocks are placed along a guide below the board (see Figure 3.2). They have rope handles that allow moving them so that they are always placed close to the coil. The complete board on a scanner table with a coil is shown in Figure 3.3.



**Figure 3.2** – Illustration of the placement of the coil (blue) above and below the subject with the help of the board (orange). The arrows indicate the movements of the coils that allow a complete acquisition of the complete leg.

### 3.2.2.2 Anterior Coil

The anterior coil has to be placed above the subject. As the dimensions of the hip are usually larger than the dimensions of the calf, the support structure for the anterior coil has to support different heights.

Our setup consists of two inverse u-shaped structures that hold the coil at its superior and inferior end. They are made of two wooden poles that are connected by a plastic screw. The poles have multiple holes (see Figure 3.3, left) for the screws at different heights to allow adjustments. The width can be controlled by the position of the screw nuts (see Figure 3.3, right). The poles fit into the comb structure of the “wings” of the coil. Velcro straps and the placing inside the combs give the structure enough stability.


### 3.2.3 Sequences

During the development of the protocol, several sequences have been evaluated based on the previously defined criteria (resolution, duration, visibility). Numerous tests have found that two sequences (T1 Turbo Spin Echo and mDixon) fit our criteria best. By using both, we can benefit from their advantages.

The two sequences which are used in the image acquisition - T1 Turbo Spin Echo (TSE) and mDixon - are using the following parameters:

- T1 TSE: TR = 633ms, TE = 20ms, resolution =  $0.67 \times 0.67 \times 4mm$ , field of view =  $430 \times 273 \times 400mm$ , and gap between slices = 1.33mm.

**Figure 3.3** – Scanning of a subject using the full set-up. The poles and their role in holding the anterior coil are clearly visible.






- 3D mDixon:  $TR = 4.6ms$ ,  $TE1 = 1.45ms$ ,  $TE2 = 2.8ms$ , resolution =  $0.78 \times 0.78 \times 1.50mm$ , and field of view =  $500 \times 336 \times 195mm$ .

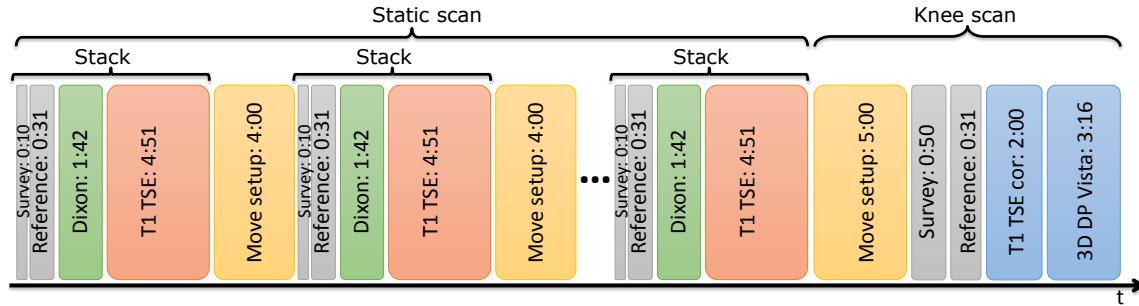
The T1 TSE scans have a slice thickness of  $4mm$  and between the individual slices is a gap of  $1.5mm$ . The mDixon slices carry information from a  $3mm$  slice. However, they overlap with  $1.5mm$  and the resulting slice thickness, therefore, is  $1.5mm$ .

The scans give different magnitude based images from mDixon and T1 TSE sequence. The mDixon magnitude images are the in-phase image  $I_I$ , the opposed-phase image  $I_O$ , the fat image  $I_F$  and the water image  $I_W$ . The T1 TSE scan produces a magnitude image  $I_M$ . Examples of these images can be seen in Figure 3.4.



**Figure 3.4** – Different modalities: mDixon (in-phase, water, fat, and opposed-phase) and T1 TSE (magnitude).

The total acquisition time depends on the leg length of the volunteer. The acquisition takes about seven minutes per stack, moving the setup can be done in four minutes. Subjects usually range from three to five stacks. This leads to total times from 29 minutes with three stacks to 52 minutes for five stacks. There are an additional five minutes for switching to the knee scan setup and seven minutes for the actual scan. The time line of typical scan can be seen in Figure 3.5.



**Figure 3.5** – The image acquisition consists of a static and a high-resolution knee scan. Survey and reference precede all scans. The static scan consists of stacks that include mDixon and T1 TSE scans. The knee scan combines a T1 TSE and a 3D DP Vista scan. Every item contains the duration in  $m:ss$  format.

### 3.2.4 Challenges in the Image Setup

The MRI setup has been improved iteratively during development, however as part of the process, some compromises had to be made. The posterior soft tissues are deformed due to supine position of the subjects. The image resolution has been limited by the duration of the acquisition. The used scanner has a tube diameter of  $60cm$ . This allows the majority of subjects to fit in there. In the presented setup, the support structure for the anterior coil takes up additional space

and prevents overweight subjects from participation. The range of motion of the posterior coil, currently limited by the board to 125cm, also limits the leg length of subjects.

### 3.3 Image Data

We have performed an MRI study with five volunteers. The study has been done at the University Hospital of Geneva, Switzerland (HUG) following the protocol described in Section 3.2. It was approved by the local Ethical Committee for Research On Humans (CEREH) of the HUG and the Swiss Agency for Therapeutic Products. Participants have been informed about the study and the goals using the sheets presented in Section C.3 and have given their consent on the form presented in Section C.4.

The subjects have been selected to fit the required criteria. They had to be between 18 and 35 years, male and needed to perform regularly physical activity: on average one hour per week. The activity was required so that the muscles are clearly shaped and the subjects were able to follow the required motions in the study design. The broader scope of the study was to research the impact of muscle injuries in athletes, so the age has been restricted to peak athleticism. Two of the subjects were active runners, two were soccer players and one followed light activity. This distribution allows us to see potential impacts of the different kinds of activities.

Due to more strict regulations on pregnancies, women had to be excluded. Furthermore, subjects were asked for other exclusion criteria stemming from the MRI (metal implants, metal ink tattoos, claustrophobia). Finally, all subjects were able to leave the study at any moment without any repercussions.


Following the acquisition, the data has been reviewed for quality and anonymised. Regarding accidental findings, the subjects were asked for their preference in the informed consent form.

Three stacks have been acquired for Subject 1 and 2. Subject 3 required four stacks. Before the acquisition of Subjects 4 and 5, it has been decided to extend the imaged volume to include the full hip. This has lead to a total number of five and respectively four stacks for Subjects 4 and 5.

During the acquisition, there was a movement in the left leg of Subject 1 between two stacks. This has caused to discontinuities in the dataset; we will discuss them in Section 3.5. The other subjects did not show similar problems. We have found that with taller subjects (e.g. Subject 3), the careful alignment on the centre of the scanner table is very important. Our coil support takes up additional space and must not touch the subject. Since the scanner opening is rather small to ensure a uniform magnetic field, there is only very little space (millimetres in some cases) left between the scanner and our setup.

### 3.4 Processing

The raw image data from the scanner has to be processed further to be used. The use of the surface coil leads to the acquisition of multiple stacks of images which first have to be stitched

**Table 3.1** – Statistical data (age, height, weight, activity) of the study participants.


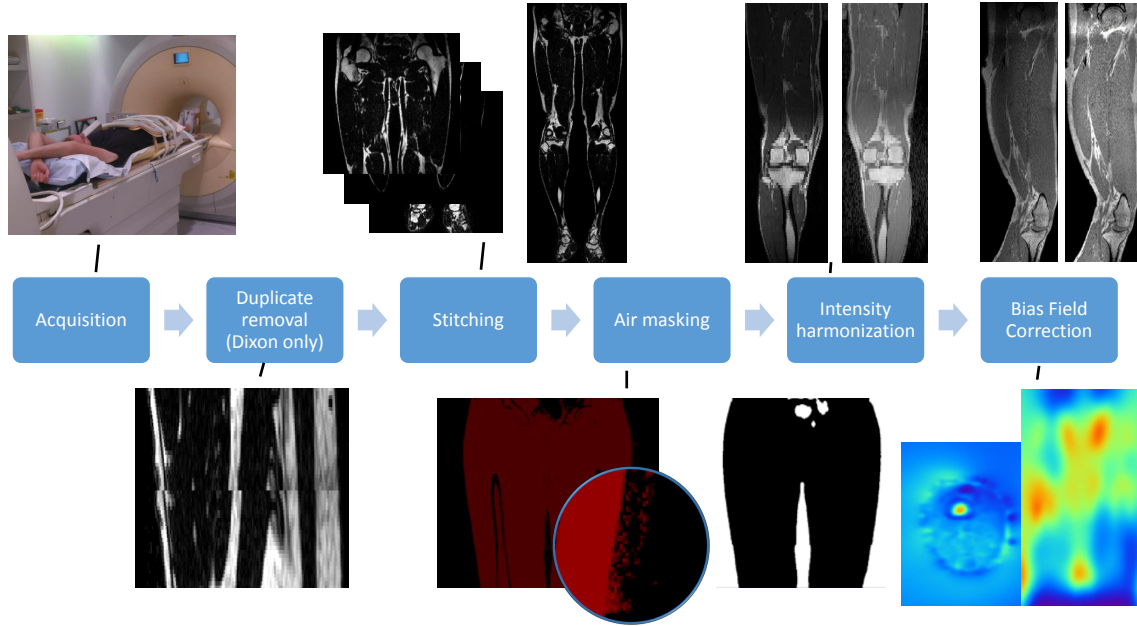
Subject	1	2	3	4	5
Age (years)	26	24	29	22	32
Height (cm)	175	172	197	176	173
Weight (kg)	85	65	84	68	58
Activity	None	Running	Running	Soccer	Soccer

into a single image. Then, a process to identify the air surrounding the subject is applied. It helps to reduce the file sizes and can mask in the following steps. This reduces calculation times and already allows to extract a surface model. This model can be used for visualisation and can be matched with body scans. Each stack carries an independent range of intensities, potentially leading to large jumps between stacks in the stitched image. We propose a method to reduce these variations by doing a slice-wise analysis of the intensity ranges. Due to small irregularities in the magnetic field of the scanner and the body moving through it, additional variations of the intensities can occur. We use established approaches to measure this bias field and remove its influence. The complete process is shown in Figure 3.6. In the following sections, we will describe the parts in detail.

### 3.4.1 Stitching

The use of the coil leads to splitting of the acquisition into several sub-acquisitions. The resulting images, stacks, have to be stitched back together to create the original image, which covers the full lower limb. Near the ends of the scanner, the magnetic field deteriorates, causes a bending of structures in the scans, which needs to be compensated for [WGZ<sup>+</sup>08, GWZ09]. In our scan setup, the length of the coil is significantly shorter than the scanner field. Therefore, no compensation is needed.

During the scans, when the coil is moved, special attention is paid to make sure that there is a small overlap between the stacks. However, the field of view of the scanner is not continuous between the stacks as the coil might have been moved sideways or the volume of interest is smaller, e.g. the ankles are smaller than the hip. During the scan, the subject is told not to move.




**Figure 3.6** – The preprocessing pipeline consists of several steps: Acquisition, stitching, air masking, intensity and bias field correction, noise reduction, and muscle tissue labelling.


Therefore, the subject is always at a well-defined position relative to the scanner table. Scanner table and scanner are connected to each other and the position of the scanner table relative to the scanner can be stored during the acquisition of the stacks. Usually, the first stack is used as the origin of the global coordinate system. Using the origin of each stack in the global coordinate system, their dimensions and spacing, we can calculate the bounding boxes of the stacks and the total bounding volume. Then, an image with properties of the bounding volume is created, and the stacks are copied into it. Areas that are not copied into mean that there is no image data available. This means, there are no structures, so that an intensity of 0, like no signal, can be inserted. Finally, the pair-wise overlap of the stacks is calculated and the datasets are interpolated between them using a linear approach. The working principle is illustrated in a 2D example in Figure 3.7.

The image data from the separate stacks is placed in a uniform global grid. This means the initial stack that defines the grid is perfectly aligned. For the following stacks, an error of half a voxel edge length may be introduced. This could be avoided by re-sampling the image data. However, this leads to interpolation errors and potential loss of details. Furthermore, our image data has a sufficiently high resolution, which keeps the offset in the sub-millimetre range:  $\leq 0.5/0.5/0.75mm$ . This offset is individual per stack and does not accumulate over the stacks of a data set.

Furthermore, we propose not to use linear interpolation, but to use the following weighting, as shown in Figure 3.8, which gives more weight to the stacks that the slices are closer to.



**Figure 3.7** – The bounding box of the stacks (a) is used for stitching. The stacks are copied into a new image (c) and linearly interpolated in overlapping volumes (b, shaded area).

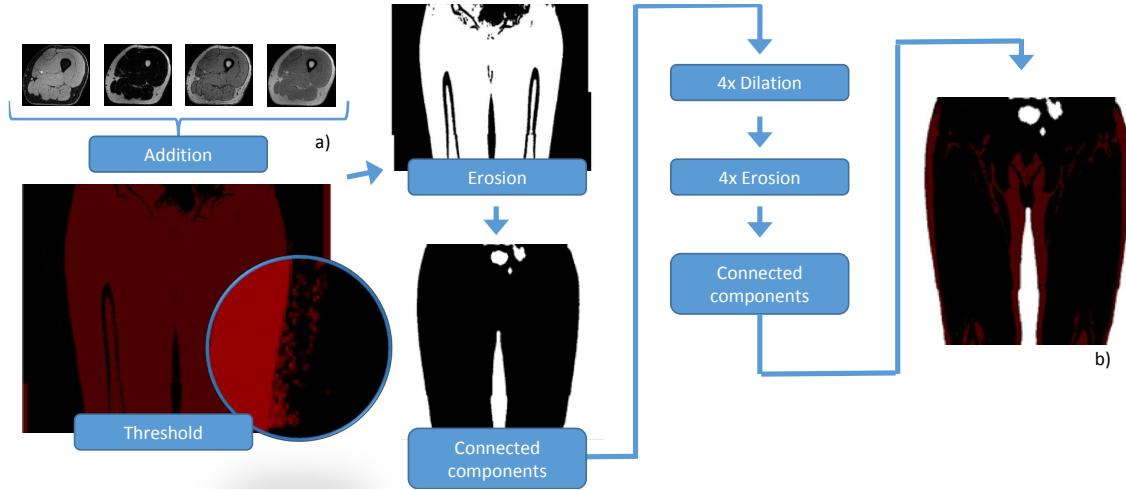


**Figure 3.8** – Different interpolation approaches between stacks: Linear (left) and non-linear (right).

### 3.4.2 Air Filtering

Large areas of the image data are covering air. These areas carry no relevant information and increase the image size unnecessarily. The field of the view, generating these areas, needs to be large during an acquisition to reduce image artefacts, like the partial volume effect. However, consequent image processing suffers from larger file sizes and the need to process more data, causing longer calculation times. Therefore, the reduction of air content in the image data is an important step.

We are using the pipeline shown in Figure 3.9 to create an air mask. We start by taking the four mDixon images and add them voxel-wise. In all four images, different tissue types have high-intensity values while the air always has low values. This way, we create an image with high intensities on tissue and low intensities on air. This process also amplifies the noise contained in the images, e.g. shown enlarged in the circle in Figure 3.9. Most of the noise image elements are small and size and not connected to the leg. A search for these connected components is performed and all components with a volume below a pre-defined threshold are being dropped. Our tests have shown that this task can be computationally expensive due to a high number of very small components. Therefore, we first run an erosion step that removes the majority those



**Figure 3.9** – The air masking pipeline takes all four mDixon channels (a) as input. It applies the depicted series of operations and creates the air mask. The result (b) shows the non-masked image parts (black) with overlaid fat (red).

components. As the resulting model can still have holes, we perform a series of dilation and erosion steps, also known as closing. The resulting mask is inverted from tissue representation to air and is subject to minimal size constraints.


The air mask is valuable for several processing tasks. It can be used to constrain filters to the actual image content, e.g. the noise reduction and bias field correction. As these methods use non-local techniques, the removed influence of the air parts leads to more accurate results. The field of view of the scanner is of a fixed size. Therefore, it does not adapt between thin and heavier subjects. The field of view is rarely fully filled. Therefore, much unused information is stored in the images. Using the air mask we were able to reduce the file sizes. In our datasets, we could reduce the file sizes to approximately one-third of the original size. When working with multi-channel data, this advantage applies to every single channel and therefore propagates. Smaller file sizes accelerate calculation and filtering times but also reduce memory pressure on older systems and make the use of graphics memory more feasible. The resulting mask can be used for visualisation; it generates an output similar to a body scanner [KHAK10].

### 3.5 Evaluation of the Seamless Stitching and Processing

Following the acquisition, the images have to be stitched and processed. In this section, we will discuss our findings and how the processing, in particular, the air filtering was applied.

During the image acquisition, as discussed in Section 3.3, Subject 1 has moved between the acquisition of two stacks. This has led to discontinuities, which became visible after the interpolation, as shown in Figure 3.10. The interpolation was initially configured to use the full overlap. Depending on the size of the subject and the distribution of stacks, the overlap could be up to 15 cm.

In cases of movements, like for Subject 1, the caused artefacts spread over a large region. There-



**Figure 3.10** – Subject moved during acquisition resulting in the two consecutive stacks not being aligned. During the interpolation, the above-seen artefact occurs.

fore, we have decided to reduce the slices that are used for interpolation to a maximum of ten. This limits the blurred region to only up to 15 mm. Using a non-linear interpolation method could help to reduce artefacts.

Using the air mask we were able to reduce the file sizes through cropping the volume of interest. In our datasets, we were able to reduce the file sizes to approximately one-third of the original size. This matter in particular for large, multi-channel datasets. Smaller file sizes accelerate calculation and filtering times but also reduce memory pressure on older systems and make the use of graphics memory more feasible. In our subjects, we were able to achieve reductions from 62% to 72%. The file sizes before and after cropping and the reduction can be seen in Table 3.2.

**Table 3.2** – Air filtering results of file size reduction.


	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5
Uncropped	1877 MiB	1973 MiB	2133 MiB	1998 MiB	2418 MiB
Cropped	701 MiB	571 MiB	728 MiB	684 MiB	658 MiB
Reduction	62%	71 %	66 %	66 %	72 %

We have applied the proposed bias field correction as well as the noise reduction. We discuss these steps in Sections 4.2 and 4.3. In Figure 3.11, we show the data set of Subject 1 before (a) and after the bias field correction (b). Furthermore, we show the air label in the dataset (c) and the 3D reconstruction of the surface.

### 3.6 Conclusion

The extraction of musculoskeletal structures from medical images is a complex task. This includes the imaging parameters and the consequent processing and filtering. First, we have presented our approach for the image acquisition. MR imaging shows satisfactory visibility with soft tissues. Our selected protocols, mDixon and T1 TSE, and the resolution are a good compromise between





**Figure 3.11** – Image data set of Subject 1 before (a) and after (b) bias field correction. The air label is shown in (c) and a 3D reconstruction in (d).

acquisition duration and image quality. To increase the SNR, we use a surface coil. We have developed a support structure to move it without touching the patient during the individual stacks.

In summary, we have contributed methods for acquiring multi-channel MR image data. We have demonstrated a new setup for continuous, gap-less multi-channel MR image acquisition using coils and subsequent processing of this multi-channel data. The acquisition has been used in a study with five subjects. Once the individual stacks have been acquired, we have stitched them. A large portion of the acquired volume is air, and we have proposed a fast technique to identify it. After the cropping of the volume of interest in the image data we were able to reduce the data set sizes by 66% on average. This reduces storage requirements and memory usage in multi-channel data sets.






---

## CHAPTER 4

# IMAGE PROCESSING

---



## 4.1 Motivation

Following the image acquisition, we need to process the image data. It carries artefacts from the scanner that we need to compensate for first. Another important step is the classification of tissue types, muscle tissue in our case. With the large multi-channel data sets, processing has to be efficient. Massively parallel computing can help to accelerate some of the tasks.

The intensity values of a voxel in an MR image do not directly related to the tissue but a radio response. This response varies on a large number of factors and can lead to different intensity ranges between different stacks. Therefore, we need to normalise the intensities between the stacks. Finally, small variations in the bias field of the scanner can skew the measured intensities. These variations are caused by imperfections in the main magnetic field of the scanner, which we have to compensate in software. Another artefact stemming from the acquisition is noise. We investigate techniques to suppress the noise in the image data as we cannot prevent it during the scans. In the next processing step, we use properties of the image data to label muscle tissue. This can be achieved by using morphological operators and the multiple channels which highlight different structures. These operations on large data sets benefit from many-core parallel processing. We propose a method to use GPGPUs for the processing.

## 4.2 Intensity Correction

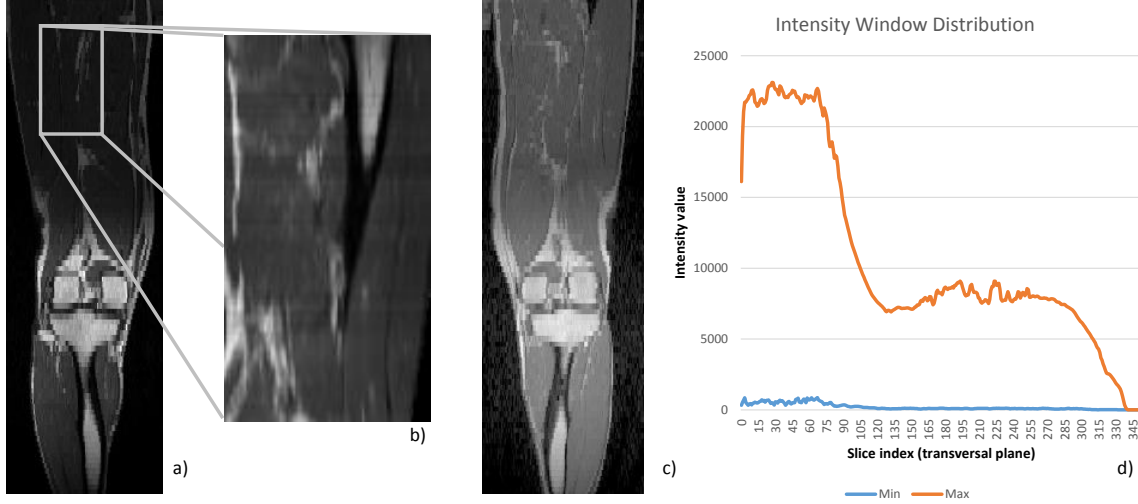
The splitting of the acquisition into several sub-acquisitions producing stacks can cause problems to keep consistency on the seams. Unlike CT with its Hounsfield units, MRI has no consistent intensity-tissue mapping. The range of intensities in an MR image depends on several factors, including the sequence used, the tissue imaged, and the field of view used. Besides inter-stack corrections, the intensities within a stack are also non-uniform. This is caused by the bias field of the scanner and requires proper correction. We will describe two approaches for the correction. First the inter-slice correction, which evaluates the range of intensities between neighbouring slices. The other approach is the bias field correction, which estimates the error in the magnetic field.

### 4.2.1 Inter-stack Correction

Different stacks can have vastly different ranges of maximal intensities per slice. We have found this to be in particular the case in the T1 TSE scans where the scans of the knee joint show higher intensities than the thigh. A plot of the minimal and maximal intensity value per slice can be found in Figure 4.1. Also, within stacks, the intensities are not uniform, our data shows jumps and spikes. This must be an artefact of the acquisition since the actual anatomical structures do not show these jumps.

We propose to correct these jumps by creating a uniform range of intensities per slice. First, we identify minimal and maximal intensity value per slice. We choose 0 to be the shared minimal intensity value and pick 10.000 (selected from the range of maximal values per slice) as the shared

maximal intensity. The result can be seen in Figure 4.1b. Due to the spikes in our image data, there are jumps between the slices. We propose to approximate the intensities by a continuous curve, leading to smoother results as shown in Figure 4.1c. This removes the spikes of the previous approach.



**Figure 4.1** – The minimum and maximum intensities can vary vastly between the individual stacks (a, d). Simple corrections can cause jumps between slices (b). Using a smoothed correction curve, these jumps can be prevented and the correction is more uniform (c).

This technique allows us to have a more uniform distribution of intensities among the slices of different stacks. It is very fast and can be easily parallelised. However, the resulting image data has lost some information about the actual range of intensities. Therefore, our technique reduces the information which can be derived from the image. Furthermore, although we uniform the intensities between slices, artefacts within slices are not addressed. Estimating the bias field is an approach that can correct these artefacts.

#### 4.2.2 Bias Field Correction


MR images carry noise that can be characterised by a non-uniform intensity distribution throughout the image. The main cause are inhomogeneities in the  $B_0$  field. These can have many reasons and are hard to prevent. One major source is the introduction of the subject into the  $B_0$  field as the body has a different magnetic flow than the surrounding air. Another source is a degradation of the  $B_0$  field near the boundaries of the field of view. These inhomogeneities result in intensity variations, dark spots, in the acquired images. These are complex to compensate since the inhomogeneities in the magnetic field cannot be predicted by a model. They depend on several parameters such as the object which is scanned, its position and its thickness.

The correction of this low-frequency bias field is a field of active research. In 1994, Wells III[WGKJ94] proposed to use the Expectation-Maximization (EM) algorithm to estimate tissue classes and the bias field in MR image data. Then, in 1997, Styner *et al.* [SG97] presented a polynomial model that uses (1+1)ES optimisation. They evolved their work into a new correction method called parametric bias field correction (PABIC), one of the dominant approaches used nowadays. The other

common approach is the non-parametric non-uniform intensity normalisation (N3) proposed by Sled *et al.* [SZE98]. Their method is independent from the used sequence and is an iterative approach to estimate bias field using tissue types. An ITK implementation has been proposed by Tustison *et al.* [TG09b]. A review of the different methods by Vovk *et al.* can be found in [VPL07].

The main concept of the non-parametric non-uniform intensity normalisation (N3) algorithm [SZE98] is to estimate the original signal (homogeneous field) iteratively by estimating the bias field without any prior assumptions. The non-uniformity of the intensities can be seen as a multiplicative source of error. Therefore, it is possible to find an estimation of the original image by estimating the bias field and division of the acquired MR image.

Sharpening the logarithmic image of the dataset gives a first approximation of the bias free image by only restoring the high frequencies. Removing them from the original image gives a good estimation of the bias field. This estimated bias field is removed from the original image. Iteratively, this process is repeated until the bias field estimation closely matches the real bias field. The final step is to divide the MR image by the bias field estimation in order to correct the non-uniform intensity in the dataset. An example for a correction in the opposed-phase images of our data sets can be found in Figure 4.2.



**Figure 4.2** – The unprocessed image data from the MR scanner shows large inhomogeneities (a). Using the N3 algorithm, we estimate the bias field (b) to correct the image. The resulting image (c) shows a more uniform distribution of intensities and clearer details at the boundaries of the field of view.

We have chosen to use the implementation by Tustison and Gee [TG09b]. It is a solid implementa-


tion that can be integrated easily with our DMFramework and gives access to many parameters, including the following:

- **Shrink factor:** Controls the sub-sampling of the input image to reduce the calculation duration.
- **Mask image:** Selects the areas that should be processed, this can be used to ignore air.
- **Number of iterations:** Maximum number of iterations to reach the final bias field estimation.
- **Number of fitting levels:** The algorithm obtains an estimation of the bias field by smoothing the subtraction of the uncorrected image by the estimated original image. For smoothing, B-splines are used. The number of hierarchical levels influence the details in the mesh resolution used for the estimation.

The mask and the number of iterations are crucial parameters for the applicability of the method. The mask can speed up the calculation and improve the quality of the output. The number of iterations can be used to stop the iteration when the process no longer converges. In order to estimate the noise, the algorithm uses a Gaussian model, which is characterised by the full width at half maximum (FWHM) of the Gaussian curve. Therefore, the estimation of this parameter is essential for output the algorithm. We have found that the quality of the filter increases with the value of FWHM. The uniformity is almost fully restored when FWHM is over 0.6, on closer inspection, the best results are obtained with FWHM values of 0.65. The fitting levels define the precision of the filtering since they define the hierarchy and hence the precision of the smoothed estimation of the bias field. We obtain the best results for a fitting level of 6 or 7.

With the parameters proposed so far, the calculation time is very high. The shrink factor is a parameter which can accelerate the calculation. However, using a too high value for the shrink factor (SF) reduces the quality of the filter output. In our tests, a SF of 8 results in the introduction of inhomogeneities. SFs over 5 do not have a significant impact on shortening the calculation duration. Values near 2-3 give better results at the cost of longer calculations. Therefore, we propose a shrinking factor of 4 or 5.

In our tests, we have found that the N3 algorithm is efficient and gives a satisfying results. The calculation takes time but can be performed autonomous. The user does not need to provide an estimation of the bias field. Also, the algorithm provides the estimated bias field which can then be used to correct quickly all channels in our multi-channel datasets. To evaluate the quality of the correction of the bias field, we measure the intensity distribution in the muscle tissue before and after correction. We use the data set of Subject 2, a shrink factor of 5, a fitting level of 6 and an FWHM of 0.65. The results can be found in Figure 4.3. It can be seen that the frequency with the bias field is distributed on a wide range of intensities. After the bias field correction, the frequency of only few intensities increases, resulting in a more homogeneous distribution within the muscle label.



**Figure 4.3** – Histogram based evaluation of the bias field reduction. The graph shows the intensities and their frequency with and without bias field. The histogram covers the volume labelled as muscle tissue, higher intensities have been omitted in the graph due to their low frequencies.

### 4.3 Denoising

MR signals are very weak and sensitive to noise. The use of surface coils makes the scanner more sensitive and allows higher resolutions. However, these higher resolutions either increase the acquisition time or the level of noise. Longer acquisitions are more stressful for the patients and cause higher costs.

The noise in images can impact processing and needs to be corrected to prevent it from misleading algorithms such as segmentation. The presence of noise has led to the need to filter the image. Many approaches have been proposed, and most of them are suited to specific niches. In 2005, Buades *et al.* [BCM05] proposed a new approach, the non-local means (NL-means, NLM) algorithm which is based on a non-local averaging of the voxels in a data set. The non-local aspect derives from calculating the denoised voxels from all voxels with a similar Gaussian neighbourhood. It has found adoption not only in MR image noise reduction but also in many other domains, including image and video processing. Tabelow *et al.* [TVP15] used a similar approach which exploits an image pyramid representation and a maximum likelihood estimation to determine the level of noise. Lie *et al.* [LHZZ13] propose to reduce of speckles in US images by using NLM filtering. They provide a GPGPU implementation and demonstrate significant accelerations.

Denoising filters are based on the re-evaluation of every voxel value in order to correct the error. Each voxel in the image is compared with the other voxels, so the algorithm can estimate a correct value. Neighbourhood filters, for instance, compare voxels with their neighbours. The non-local means algorithm expands the comparison to the full image. The filter assigns a new intensity to a voxel by analysing the similarity between that voxel and all others in the image. The new intensity is the weighted average of all voxels: the weight increases with the similarity between the voxels. More precisely, the NL-means method analyses the similarity between the neighbourhood of the processed voxel and the neighbourhood of all others. Taking into account all information of the entire image, helps the reconstruction of the texture after filtering.

Our MR images have a noise signal which overlays the original images and causes a granular texture. This is shown in Figure 4.4. We apply the NL-means algorithm to remove this noise

without changing the details in the texture which is needed for successful muscle segmentation.



**Figure 4.4** – Transverse section through the thigh (MRI) of the water-channel in the mDixon scan of Subject 2 showing the noise in MR images (a). The results of the noise reduction are shown in (b).

We use the NL means implementation of ITK with the following parameters:

- **Sigma**: The noise is considered as an additional independent and identically distributed (i.i.d.) Gaussian signal which needs to be estimated. Sigma is the standard deviation of the i.i.d. Gaussian model of the noise. The efficiency of the filter depends on the estimation of its value.
- **RSearch**: The algorithm looks for similar pixels which are located in a region described by this radius.
- **RComp**: In order to estimate the similarities between two pixels, the calculation are made on the pixel neighbourhood with a size of RComp.
- **H**: The degree of filtering allows to control the strength of the denoising.
- **PSTh**: The threshold used to discard insignificant pixels. It is used to accelerate the calculation.


Sigma is the most influential parameter in the algorithm, a wrong estimation will not remove the noise. We have found that the granular aspect of the muscles starts to be corrected and the best filtering is obtained with a sigma value of 40. For the filtering degree (H), an initial assumption would be that higher values would generate better results. However, [TVGPAFW12] propose to constrain H between 0.8 and 1.2. This confirmed that the denoising is more efficient with higher values of H. Some regions of the thigh still have the granular noise after the filtering with H values of 0.8 and 1. Based on this, we have chosen the value of H to be 1.2. The default value for the threshold (PSTh) gives good results, changing its value decrease the quality of the filtering. The muscle has been found to be more granular when PSTh is lower.

The two last parameters, RSearch and Rcomp, are crucial for the algorithm. They define the ex-



pansion of the comparison zone, which is the characteristic of the NL-means method. The quality of the details in the texture in the filtered image will depend on their values. As suggested by literature, a larger radius results in better results. However, if they are too large, the calculation time lengthens. By increasing these radius, the strength of the filtering increases. Values higher than 10/5 for RSearch/Rcomp cause very high the calculation times. Therefore, the best configuration is not the one with the best quality, but a compromise with the best quality for a reasonable calculation time. The configuration 10/5 gives the most promising results.

We have found the non-local means denoising method to be efficient on the MR images. The ability of the filter to remove the noise without altering the details in the images is an essential advantage in the processing phase before segmentation. To evaluate the quality of the denoising, we study the histograms of the muscle tissue before and after processing. We use the data set from Subject 2, a sigma value of 40, a filtering degree (H) of 1.2, a threshold (PSTh) of 2.3 and a RSearch/Rcomp of 10/5. The results can be found in Figure 4.5. It shows that the relevant intensities for the muscle labels are more compressed. Their frequency is higher among the main components and lower for less typical intensities. The SNR (as ratio between the average (muscle) values and their standard deviation [GC11]) reflects this with an improvement from 5.1 to 5.4.



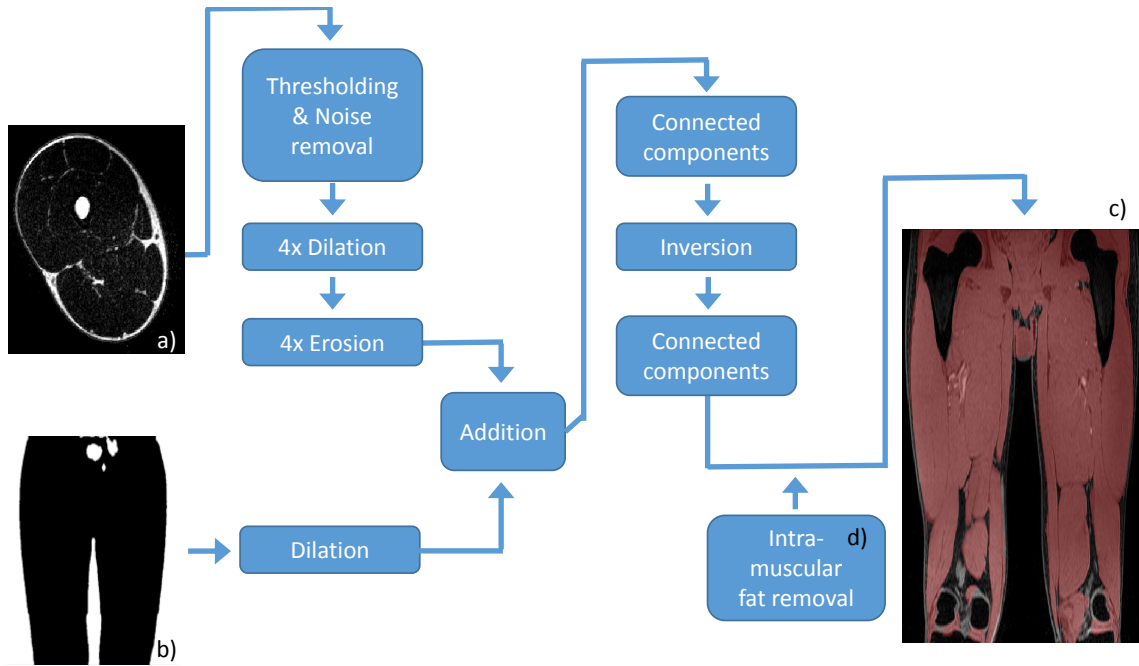
**Figure 4.5** – Histogram based evaluation of the noise reduction step. The graph shows the intensities and the frequency before and after denoising. The range of intensity has been optimised for the muscle tissue, the histograms have been constrained to the muscle labelled volume.

## 4.4 Muscle Labelling

The MRI mDixon sequence is fast and creates a multi-channel image, highlighting fat and water in separate channels. Selected combination of image operations in the multi-channel data can be used to label specific regions of interest (ROIs) or segments. These image operations include thresholding, morphological and connectivity operations. Based on that, we have developed a fully automated technique to label muscle tissue in the lower limb. It can be used in several applications to improve quantitative estimations in diagnosis.

Similar techniques have been used for fat labelling. Valentinitich *et al.* [VKA<sup>+</sup>13] demonstrate a method for detecting adipose tissue depots using morphological operations. They limit their work to the calf. Hu and Kan [HK13] describe the principle and applications of quantitative proton MR techniques to measure fat. Poonawalla *et al.* [PSR<sup>+</sup>13] have performed an evaluation study for their adipose tissue magnetic resonance imaging method. Thomas *et al.* [TFM<sup>+</sup>13] dis-

cuss whole-body fat distributions and imaging techniques. They show examples for different modalities and how the obtained data can be used for detecting fat distribution. Livingstone *et al.* [LBK<sup>+</sup>13] compare the results of hepatic and pancreatic fat measurements in mDixon and <sup>1</sup>H magnetic resonance spectroscopy. Ristic *et al.* [RBZ<sup>+</sup>12] have studied the connection between subcutaneous fat measurements and metabolic prediction. They use the waist circumference and ultrasound data. Addeman *et al.* [AKP<sup>+</sup>14] show an automatic adipose segmentation in the abdominal region. They use thresholding in mDixon images. Garteiser *et al.* [GDTG13] semi-automatically segment adipose tissue in mice. They apply thresholding after a Gaussian fitting of the pixel intensity histograms. Tang *et al.* [TSM11] have described an automatic technique for fat measurement in mice. It is based on morphological operations and fuzzy learning.



**Figure 4.6** – The muscle detection process uses the fat image (a) and the air mask (b) as input. The depicted series of operations creates the muscle label image (c). The intra-muscular fat removal step (d) is optional.

The muscle detection works similar to the air masking; the process can be seen in Fig 4.6. We start with the fat image and apply a minimal threshold to gain a fat mask. We perform an erosion to remove noise, followed by a series of dilation and erosion steps to close holes. We take the air image and dilate it several times so that it overlaps at the skin once we add the air and fat images. In the resulting image, we enforce a minimal size of connected components. After an inversion of the mask, we receive the muscle image, which we once more filter for component size to remove the risk of remaining artefacts. A last step can be used to reduce intra-muscular fat in the resulting image; this is optional as some applications benefit from hole-free results.

#### 4.4.1 Evaluation and Results

We have selected the subjects 1 - 4 from our study to evaluate our approach. We have performed a manual segmentation of multiple slices in the two subjects. The position of the slices is shown

Fig 4.7c. Table 4.1 shows that we have measured the volume of manual and automatic segmentation in these slices. Also, we have calculated the Sørensen-Dice coefficient, which has been adopted as standard in medical image segmentation evaluation when label images are used. An SDC of 0 expresses disjunct segments, a value of 1 expresses identity. For two segments  $S$  and  $T$ , the formula is the following:

$$SDC(S, T) = 2 \frac{|S \cap T|}{|S| + |T|}$$

**Table 4.1** – This table lists the results of the comparison of the automatic (auto) and manual (man) segmentation in several slices per subject. It provides the segment volume in ml and a comparison using the Sørensen-Dice coefficient.

Method	Subject 1						Subject 2					
	left leg			right leg			left leg			right leg		
	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC
Slice 1	0.79	0.84	0.96	0.88	0.96	0.92	0.97	0.99	0.98	0.96	0.98	0.97
Slice 2	11.60	12.23	0.97	11.37	12.30	0.97	11.88	12.09	0.97	11.85	12.24	0.97
Slice 3	0.49	0.59	0.90	0.44	0.53	0.91	0.75	0.80	0.95	0.72	0.76	0.94
Slice 4	0.44	0.54	0.88	0.51	0.57	0.92	0.59	0.70	0.90	0.61	0.73	0.89
Slice 5	14.99	15.93	0.95	15.90	16.66	0.95	16.15	16.75	0.95	16.15	16.50	0.96
Slice 6	21.86	23.07	0.97	22.37	23.60	0.97	21.96	22.67	0.97	21.17	21.92	0.97
Slice 7	23.84	25.35	0.96	23.68	24.72	0.97	23.30	24.65	0.96	22.31	23.67	0.96

Method	Subject 3						Subject 4					
	left leg			right leg			left leg			right leg		
	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC
Slice 1	0.96	0.94	0.98	0.91	0.95	0.97	0.77	0.81	0.97	0.69	0.80	0.93
Slice 2	14.29	14.99	0.97	13.80	14.06	0.98	14.18	14.50	0.98	13.18	13.62	0.98
Slice 3	10.64	11.03	0.97	10.07	10.40	0.97	11.39	11.92	0.97	11.11	11.58	0.97
Slice 4	0.56	0.65	0.89	0.52	0.65	0.86	0.74	0.90	0.91	0.70	0.74	0.91
Slice 5	15.86	16.28	0.97	14.36	14.41	0.95	20.01	20.99	0.97	19.78	20.01	0.97
Slice 6	21.31	22.32	0.97	20.72	21.55	0.98	27.83	28.61	0.98	26.66	27.83	0.98
Slice 7	25.37	26.58	0.97	25.29	27.19	0.96	28.93	30.32	0.97	27.16	28.93	0.97

Method	Subject 5					
	left leg			right leg		
	man (ml)	auto (ml)	SDC	man (ml)	auto (ml)	SDC
Slice 1	0.68	0.72	0.96	0.70	0.72	0.97
Slice 2	10.94	10.85	0.97	10.72	10.52	0.97
Slice 3	0.41	0.51	0.85	0.38	0.48	0.83
Slice 4	0.30	0.42	0.74	0.32	0.44	0.81
Slice 5	12.05	12.49	0.96	12.12	12.70	0.95
Slice 6	18.46	19.16	0.97	18.40	18.69	0.97
Slice 7	22.57	23.44	0.97	21.77	22.96	0.97

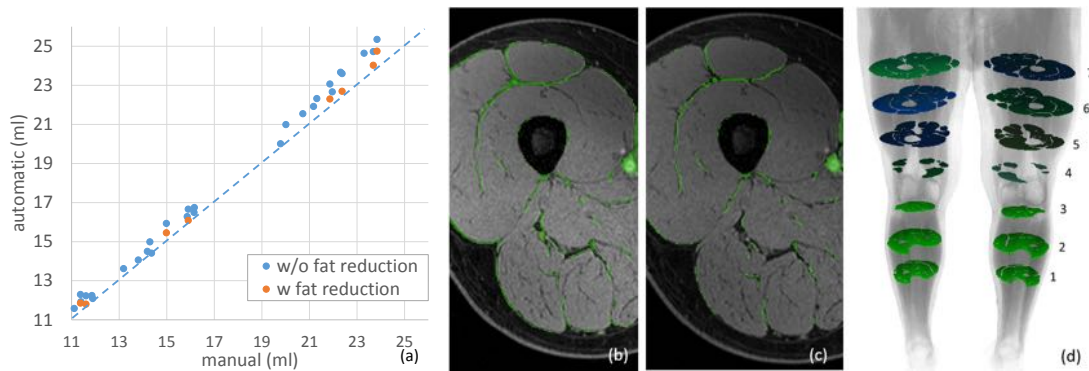
Our measurements can be split into two parts. The joint area, consisting of slices 3 & 4, is located close to the knee joint (as shown in Fig 4.7d) with its ligaments, tendons, menisci and other tissue types. Therefore, it is more complex to assign labels. The main area, with slices 1, 2, 5, 6, and 7, covers mid-thigh and calf. Table 4.1 shows the result without the fat reduction step; all areas combined have an average SDC of  $0.95 \pm 0.03$ . The joint area has an average SDC of 0.91 with a standard deviation of  $\pm 0.02$ ; the main area has an SDC of  $0.96 \pm 0.02$ .

**Table 4.2** – The results for the muscle labelling without and with fat reduction technique. The average SDC (with standard deviation) between manual and automatic technique is listed, as well as their difference.

Method	all		main area		joint area	
	average	std. dev.	average	std. dev.	average	std. dev.
without fat reduction	0.95	0.02	0.97	0.02	0.91	0.02
with fat reduction	0.96	0.02	0.97	0.01	0.93	0.01
difference	0.01	0.01	0.003	0.006	0.027	0.002

The use of the fat reduction results in an increased SDC as presented in Table 4.2. A more detailed analysis shows that this improvement mainly stems from the joint area, the main area only shows small improvements.

The comparison of manually and automatically determined volume is presented in Fig 4.7 (a). The measurements all show a bigger volume for automatic than manual labelling (equality indicated by the dotted line). A graphical overlay of the symmetrical difference between manual and automatic segments on the water image is shown in Fig 4.7 (b). It can be seen that the segment difference is mainly rooted in inter- and intra-muscular fat as well as veins and arteries.



**Figure 4.7** – Plot of manually and automatically segmented muscle volume (a), the results of the experiments with fat reduction are shown in orange, the ones without in blue. Symmetrical difference between manual and automatic segmentation overlaid on the water image without (b) and with (c) fat reduction. Position of the slices (1: bottom, 7: top) that were used for evaluation (d).

Several other groups have worked on muscle labelling of the thigh. Kang *et al.* [KPV<sup>+</sup>07] have proposed a fuzzy C-means approach (FCM), Pinti *et al.* [PHKTA06] have proposed a surface expansion-based segmentation method (SEB). The work of Urricelqui *et al.* [UMV09] present an unsupervised recursive algorithm that is based on histograms, thresholding, and connectivity (URA). Andrews *et al.* [AHY<sup>+</sup>11]. use a model-based technique with prior knowledge (MBPK). Imamoglu have presented saliency-based methods (SBM) [IGTG<sup>+</sup>14] and a fuzzy version (FSBM) of their approach [IGTH<sup>+</sup>15]. All works present their results in different formats. The F-measure is the harmonic mean of precision and recall; accuracy equals the recall. Another measure the coincidental pixel (CP) rate. We compare these approaches with our results in Table 4.3. We can see that our quantitative results are slightly higher while the computation time is lower.

During the evaluation, we have found that our method shows inferior results in the joint area. In these areas, many muscles end and are connected by tendons to the bones. Possible approaches


**Table 4.3** – Muscle labelling of the thigh in comparison with SoA. There have different evaluation chosen by these publications. However, we list these values to give an orientation.

Approach	Evaluation	Quantitative Results	Computation Time
FCM [KPV <sup>+</sup> 07]	Accuracy	0.91	NA
SEB [PHKTA06]	Accuracy	0.94	NA
URA [UMV09]	CP-Rate	over 90%	≈ 300s / slice
MBPK [AHY <sup>+</sup> 11]	DSC	0.95	NA
SBM [IGTG <sup>+</sup> 14]	F-measure	0.92	≈ 1s / slice
FSBM [IGTH <sup>+</sup> 15]	F-measure	0.92	≈ 10s / slice
Our method [BMT14]	DSC	0.97	≈ 0.1s / slice

would be an extension to tendons, e.g. visible with other sequences, or to rely on a small, local atlas. Our results currently identify if a tissue is a muscle or not.

The muscle labels created by our approach can be used for many purposes, including visualisation, clinical analysis, pose estimation and initialisation of further segmentation processes.

Medical image visualisation is a complex task. Volumetric data sets are not intuitive, and the navigation using several planes is time-consuming. Our presented technique allows the robust removal/masking of noisy air and the identification of fat and muscle structures. A real-time 3D rendering of muscle and fat tissue, enriched with two femur models, is shown in Figure 4.8. Indicated by arrows on the right image is visible how the muscle surface allows the visual differentiation of distinct muscles.



**Figure 4.8** – Applications of the muscle labelling. Simple gray-scale visualisation is commonly used in clinical environments (left). Our proposed technique could give additional instant feedback to radiologists. In the rendering, the distinction of individual surface muscles is possible (right).

In clinical environments, feedback for technicians and radiologists during image acquisition is very limited. Based on simple gray-scale visualisations, decisions on further acquisitions and/or different modalities have to be made. Our robust technique is fast and can be used for further highlighting of structures that could reduce the time per case for radiologists.

Another application is the registration of a pre-defined atlas. Using the air image and that atlas, the imaged body part can be identified and the pose of the subjects can be estimated. Furthermore, muscle labels can also be useful to initialise other segmentation methods to differentiate between distinct muscles, e.g. deformable models.

## 4.5 Acceleration through Massively-Parallel Image Processing

The processed image data can help radiologists during the acquisition. However, this is only possible if the data is available (almost) immediately. The time spent in the scanner not only is a burden for the patients but also increases costs for the hospital. Furthermore, the possible total number of patients is reduced, resulting in longer waiting lists.

Over the last decade graphical processing units (GPUs) have not only increased their performance but also have become more generic. These general-purpose GPUs (GPGPUs) allow access to their highly parallelized cores. In combination with high-speed memory, GPGPUs have proven to be able to reduce calculation times for image processing tasks. To meet our requirements for near-realtime feedback to radiologists, GPGPUs can be used for acceleration. Modern scanners already use GPGPUs for the image reconstruction, e.g. the Radon transform in CT or MR reconstruction [QHCJ15]. Therefore, GPGPUs are widely available in machines with direct access to the image data.

Graph processing is a popular task on GPGPUs. An overview, bottlenecks and fields for future research have been discussed by Xu *et al.* [XJA14]. Images can be considered a special case of graphs. Therefore, medical image processing can also be considered a good use case for GPGPU accelerated processing. Eklund *et al.* [EDFL13] and Smistad *et al.* [SFB<sup>+</sup>14] provide an overview currently used approaches. Al-Ayyoub *et al.* [AAADJ<sup>+</sup>15] present an overview GPGPU based fuzzy C-means approach, Schmid *et al.* [SIG<sup>+</sup>11] demonstrate an implementation of knowledge based deformable models. An additional way to further improve performance is to use multiple GPGPUs in parallel. Fang *et al.* [LMHF14] have demonstrated an approach for fast satellite image data segmentation by distributing images to multiple GPGPUs. Medical applications for the detection and segmentation of objects have been shown by Grace *et al.* [GFFK15].

In the following sections, we will discuss the platform selection for our approach, describe the general architecture of GPGPUs and present the implementation of the methods needed for air and muscle labelling.

### 4.5.1 Platform Selection

Besides multi-core platforms like OpenMP, Intel ArBB (Array Building Blocks) and TBB (Thread Building Blocks), two main platforms for many-core (GPGPU) programming exist: CUDA<sup>1</sup>, a proprietary platform from nVidia and OpenCL, that is standardised and maintained by the Khronos group. The most recent version of CUDA is 7.5, OpenCL has been recently published in version 2.1. While nVidia often updates CUDA to support new versions of their GPUs, OpenCL is platform agnostic and, therefore, needs fewer updates.

The selection between CUDA and OpenCL depends on a number of factors that go beyond pure performance. Many comparisons on typical applications have been proposed, for example by Fang *et al.* [FVS11] and Karimi *et al.* [KDH11]. While Karimi *et al.* find CUDA faster in data

1. Initially an acronym for Compute Unified Device Architecture

transfer and execution. Fang *et al.* initially also found CUDA to be faster, however, they were able to minimise the difference through optimisations. In comparison with multi-core systems, Sanchez *et al.* [SFSG12] have found that with small data sizes vector SIMD-CPU's can deliver performance results comparable to CUDA; this does not apply to large data sets. Finally, Shen *et al.* [SFVS13] have evaluated common strategies and challenges from an application-centric perspective in OpenCL systems. Many of their points apply to other many-core systems like CUDA.

Based on reviewing the literature, we have decided to use CUDA. It offers a slightly better performance and a broad range of compatible GPGPU's was available to us. While OpenCL is officially supported by nVidia, they only support version 1.1. However, starting with version 2.0, primitives for 3-dimensional data sets have been introduced that not only make development easier but also benefit from optimised memory access patterns. This is crucial in our large data sets. Finally, CUDA comes with extensive documentation, several openly available libraries and a large number of publications on different GPGPU applications.

## 4.5.2 GPGPU Architecture

GPGPU's have been designed for massively parallel data processing. While modern CPU's have few cores with long pipelines and complex instruction scheduling, GPGPU's have many cores with a much more simple architecture. These cores are designed to run short tasks with a limited instruction set. Since there is a large number of them, a high complexity would result in a prohibitive large number of transistors and large chip areas needed.

The code running on the GPGPU is called *kernel*. A GPGPU consists of multiple streaming multi-processors (SM). The total volume (called grid) is divided into blocks that are distributed across the SMs. These blocks consist of many threads and must be able to run independently; an ordered execution is not guaranteed. Threads within a block have shared memory and can be synchronised at certain points. Within a block, threads are divided into warps of 32. They share the same current instruction but have individual registers. The structure of the kernel code must be chosen carefully: too much different branching within warps hurts performance. The main architecture is shown in Figure 4.9.


Based on the block size and the current block and thread id, the coordinates of the voxel for processing can be calculated as follows:

$$x = blockIdx.x * blockDim.x + threadIdx.x$$

$$y = blockIdx.y * blockDim.y + threadIdx.y$$

$$z = blockIdx.z * blockDim.z + threadIdx.z$$

Kernels can only access GPGPU memory; all data has to be transferred to and from the host memory. This has to be included into performance measurements and can cancel the benefits from parallelization in the GPGPU for smaller data sets.



**Figure 4.9** – The grid represents the working task for the GPU. It can be broken down into blocks which can run on different streaming multiprocessors (SM). A block is composed of threads which run interleaved in parallel.

### 4.5.3 Pixel-wise Operations

This category contains the addition of images, multiplication and thresholding. These tasks operate pixel-wise independent from their neighbourhood. Therefore, their parallelization is straight forward. We have implemented these methods to either use a target image or to run in-place. The corresponding kernels are shown below. The high parallelism and efficient scheduling of warps allows running kernels on a single voxel only while still having a performance gain.

**Listing 4.1** – Kernels for adding two images with target image (top) and in-place (below).

```
__global__ void addImageKernel(short* out, short* in1, short* in2) {
    unsigned int idx = z*dim_x*dim_y + y*dim_x + x;
    out[idx] = in1[idx] + in2[idx];
}

__global__ void addImageInPlaceKernel(short* inOut, short* in2) {
    unsigned int idx = z*dim_x*dim_y + y*dim_x + x;
    inOut[idx] = inOut[idx] + in2[idx];
}
```

### 4.5.4 Morphology Operations

Initially, we have considered using the nVidia Performance Primitives<sup>2</sup>(NPP) as implementation for the morphological operations. While NPP shows 5-10 times speed improvements over CPUs, we cannot use it, because it is available only for 2D image data. Thurley and Danell [TD12] have shown another CUDA implementation of morphological operations. Although they also only work on 2D images, they confirm that speed improvements over CPUs are possible.

We propose a 3D implementation that is tailored to our use case. We need to apply erosion and dilation with a radius of 1 voxel. We have tested two different neighbourhoods (N6 and N26). As shown in the kernel code below, we operate pixel-wise. Through the structuring into blocks and warps, many of the neighbouring voxels are already in the caches. This reduces access time

2. <https://developer.nvidia.com/npp>



and accelerates the calculation. The in-place requires two runs. First, all elements that will be removed/added are identified. As the calculation is performed in parallel and not serial, later run threads could operate on image data was already modified. Therefore, we only mark the resulting changes (by placing the value 2) and check for it during execution. In a second run, the 2s are replaced with the desired value (erosion: 0, dilation: 1).

**Listing 4.2** – Kernels for N26-dilating a binary image with (top) target image and (bottom) in-place using the proposed two-step process.

```
__global__ void dilateImageKernelN26(short* outImg, short* inImg) {
    if (inImg[z*dim_x*dim_y + y*dim_x + x] == 1) {
        outImg[z*dim_x*dim_y + y*dim_x + x] = 1;
        return;
    }
    int counter = 0;
    for (test_x, test_y, test_z in N26(x,y,z)) {
        counter += inImg[test_z*dim_x*dim_y + test_y*dim_x + test_x];
    }
    if(counter > 0) {
        outImg[z*dim_x*dim_y + y*dim_x + x] = 1;
    } else {
        outImg[z*dim_x*dim_y + y*dim_x + x] = 0;
    }
}

__global__ void dilateImageKernelN26InPlace(short* inOutImg) {
    if (inOutImg[z*dim_x*dim_y + y*dim_x + x] == 1) {
        inOutImg[z*dim_x*dim_y + y*dim_x + x] = 1;
        return;
    }
    int counter = 0;
    for (test_x, test_y, test_z in N26(x,y,z)) {
        if (inOutImg[test_z*dim_x*dim_y + test_y*dim_x + test_x] != 2) {
            counter += inOutImg[test_z*dim_x*dim_y + test_y*dim_x + test_x];
        }
    }
    if(counter > 0) {
        inOutImg[z*dim_x*dim_y + y*dim_x + x] = 2;
    } else {
        inOutImg[z*dim_x*dim_y + y*dim_x + x] = 0;
    }
}
```

### 4.5.5 Connected Component Labelling

Connected component labelling, or blob detection, is a process that detects connected parts within an image. The connectivity depends on the use case, in 2D N4 and N8 are commonly used, in 3D N6 and N26. There exist several approaches to the problem, recent approaches operate serial and in two passes like the work of Wu *et al.* [WOS09] and Gupdate *et al.* [GPP<sup>+</sup>14].

Several approaches for CUDA implementations have been shown. They usually are implemented for 2D grids, e.g. the work by Kalentev *et al.* [KRKS11]. Other works, e.g. by Hawick *et al.* [HLP10] and Barnat *et al.* [BBBČ11] have evaluated the performance of CCL algorithms on CPU and GPGPU. They have found that optimised CPUs are often faster for low- and medium-resolution images and that the GPGPU is faster with growing data sets. These evaluations have been made with 2D data. We expect to benefit from a GPGPU implementation because our data sets are larger by magnitudes due to the third dimension.

The presented methods use the Union-Find method. The main part is a label equivalence tree (LET). Every label has an entry point to an equivalent label. Initially, they are the same. First, the image elements are labelled with continuous labels. Then, the neighbourhood for every image element is evaluated. If the intensities are the same, the LET is updated to the smaller label of both elements. Once this is completed, the LET contains indirect, tree-like reference to equivalent labels. Using the code shown in Listing 4.3 on each entry in the LET, these trees flattened. A final iteration over the image data allows looking up the minimal label for each element in the LET. This method produces connected labels that are not continuous.

**Listing 4.3** – Method used for flattening the label equivalence tree (LET).


```
function Find(LET, x)
    if LET[x] == x
        return x
    else
        return Find(LET[x])
function Union(LET, x0, x1)
    root0 = FindRoot(LET, x1);
    root1 = FindRoot(LET, x2);
    LET[root1] = min(LET[root1], root0);
    LET[root0] = min(LET[root0], root1);
```

The Union-Find method works well for small and medium sized 2D data sets, however, with our 3D data, we need a different approach. Stave *et al.* [ŠB11] propose a multi-step process: First classical union-find in small blocks, then merging to larger blocks using an LET, and finally, flattening the LET and applying it to the image data. We will describe the steps and our proposed implementation next.

#### Step 1: Initial labelling

We have selected a block size of  $16 \times 8 \times 8$ . This limits the merging steps to 7 in our data sets and

is still sufficiently fast. Within the block, every single voxel is processed by a single thread. The shared memory is allocated on a per block basis and large enough to hold an LET per block of our dimensions. Therefore, we can do fast updates of this local LET. At the end of the kernel, the local LET is merged with the global one. By doing this in a single copy-operations, we have reduced access to global memory, in particular, the fragmented read and write operations from updating the local LET. In Figure 4.10, we show the resulting block from labelling an inverted and thresholded IP image.



**Figure 4.10** – Visible label blocks from the initial labelling (a). During the merging of 4 blocks, twelve sides have to be processed (b).

### Step 2: Merging

During the merging phase, eight neighbouring cubes are combined into a single one by updating the labels along the neighbouring sides. This iterative process is repeated by halving the number of blocks and doubling the dimension of an individual block until it spans the full data set. A block is defined as four neighbouring cubes. These cubes share a total of twelve sides which can be seen in Figure 4.10 (right). The first dimension of the threads is assigned one of the sides (0..11). Once the blocks grow larger, additional threads can be created to sub-partition the sides of the cubes in the blocks. Since shared memory is not large enough to hold the LET, it is stored in global memory and entries require atomic updates (atomicMin).

### Step 3: Flattening and continuity of labels

After the threads have updated the global LET, it has to be flattened. This can be done in an iterative parallel process. The labels of flat LET are not continuous. This can be achieved through a relabelling. However, this is done best in a single, sequential process.

### Step 4: Label Transfer

Once the LET has been flattened, the labels can be directly transferred to the final image. In our

case, we know that the number of labels is low. Therefore, we can switch from a 4-byte number representation to a 2-byte representation. This reduces the memory consumption.

Another possibility is to check the size of the connected components. This could be done in a single kernel. Then LET could be updated for labels that are too small by setting the LET for these labels to 0 (background).

#### **Adjustment to large 3D data sets**

The presented method proposed by [SB11] is described to work with n-dimensional data sets, however, results have been presented only for 2D data sets. With increasingly large data sets, modifications have to be made. A data set with  $600 \times 600 \times 300$  requires 205 MB with a 16 bit representation. In a data set with potentially  $600 \times 600 \times 300$  labels, a 32 bit representation of the labels is needed. This increases the label image to 410MB. If we follow the approach with a global LET, it would require 810MB. Furthermore, we have to look at the uncropped image data, because the cropping is based on the air filtering, which requires the CCL. Since cropping based on the air reduces the data sets by approx 66%, we have to triple the actual data set sizes. This results in prohibitively large sizes of 1.2GB for the label image and 2.4GB for the LET.

Therefore, we propose a series of modifications to the original approach. We propose to perform the steps 3 and 4 after every merging step. Furthermore, we propose to limit the total number of labels per step and in general. This reduces the size of the LET drastically. It also reduces the costs for the Union-Find step. In the first stages, we propose to enforce already a minimum size per component. This keeps the total number low. The detection of component size is limited to blocks and only components that do not touch the boundaries of the block can be reliably removed if they are too small.

#### **4.5.6 Slice Histogram Generation**

For the initialisation process, we use a histogram of air and muscle tissue; a complete description can be found in Section 5.3. We can perform the generation of these histograms directly on the GPGPU as the data is already in memory. We take the air and the muscle labels as input. The number of blocks is identical to the number of slices. Since the histogram is slice-wise independent, the thread can use the shared memory which has faster access times. The threads process individual parts of the slice and store their results in the shared memory using an atomic operation. Once all threads have finished (detection through synchronisation), the result is transferred to global memory and can be used for the initialisation.

#### **4.5.7 Evaluation and Results**

In this section, we will evaluate the processing pipeline on CPU and GPGPUs. We aim to show the speed differences between multi-core (CPU) and many-core (GPGPU) systems and to evaluate how the pipeline can benefit from massively parallel processing.

We evaluate the pipeline for air and muscle labelling on one CPU type and five different GPGPUs.

**Table 4.4** – Specification of the GPGPUs used for evaluation.

	Gen.	Cores	Clock	Memory	Mem. clock	Bandwidth
GTX 470	Fermi	448	607 MHz	1280 MB	837 MHz	133 GB/s
GTX 660	Kepler	960	980 MHz	2048 MB	1502 MHz	144 GB/s
GTX 750Ti	Kepler	640	1085 MHz	2048 MB	1350 MHz	88 GB/s
M2050	Fermi	448	1150 MHz	3072 MB	773 MHz	144 GB/s
GRID K520	Kepler	1536	800 MHz	4096 MB	1250 MHz	160 GB/s


The CPU that we have used is an Intel Core i7 3770 of the Ivy Bridge generation with four cores and eight parallel threads. It is clocked at 3400 MHz. The caches are composed of 128 KB L1-, 1 MB L2-, and 8 MB L3-cache. The system was equipped with 16 GB RAM with a bandwidth of 12.8 GB/s. Based on our decision for CUDA, we are using only nVidia GPGPUs. Our five cards cover the Fermi and Kepler generation, with Kepler being the current generation at the time of writing. Except for the GTX 470, all cards have at least 2GB of memory. Within the Kepler generation, we have selected cards that show the range between large number of cores and higher clock rates. The GTX 750Ti represents cards with lower memory bandwidth. An overview of the specifications of the GPGPUs used can be found in Table 4.4.

With our setup, we can generate multiple data sets of different resolutions. To better understand the implications on the total duration of processing, we have tested the CPU and GPGPU pipeline on different resolutions of our five subjects. First, we have tested the original, uncropped data sets, which range from 469 to 605 MB. We have also tested the cropped data sets which range from 125 to 185 MB. Finally, we have tested a further reduced data set of subject 1, with only 86 MB.

**Table 4.5** – Total air processing times on our five subjects and reduced data sets. Missing results are caused by insufficient GPGPU memory to store the data sets and needed temporary information.

	Size	CPU	GTX 470	GTX 660	GTX 750i	M2050	Grid K520
Subject 1	469 MB	64 s	-	16 s	11 s	20 s	16 s
Subject 2	493 MB	68 s	-	16 s	12 s	20 s	16 s
Subject 3	533 MB	71 s	-	19 s	13 s	20 s	17 s
Subject 4	534 MB	67 s	-	20 s	13 s	20 s	17 s
Subject 5	605 MB	70 s	-	-	14 s	23 s	19 s
Subject 1 (crop)	175 MB	22 s	8 s	8 s	4 s	7 s	6 s
Subject 2 (crop)	125 MB	15 s	6 s	5 s	3 s	5 s	4 s
Subject 3 (crop)	182 MB	25 s	8 s	8 s	5 s	7 s	6 s
Subject 4 (crop)	185 MB	22 s	8 s	8 s	5 s	8 s	6 s
Subject 5 (crop)	165 MB	23 s	7 s	6 s	4 s	7 s	5 s
Subject 1 (half)	86 MB	11 s	4 s	3 s	2 s	4 s	3 s

While we have tried to evaluate all data sets on all GPGPUs, this was not possible due to memory constraints. For the proposed processing methods, we need to allocate at least two copies of a single channel of the data set in GPGPU memory. There is additional memory needed for the LET in the CCL step. Together with the memory used by the OS, the uncropped subject 5 could not be processed on the GTX 660 as well as all uncropped data sets on the GTX 470 which only has 1.25 GB of memory.



**Figure 4.11** – Total air processing times on our five subjects (S1 - S5) and reduced data sets (air cropped and halved). Missing results are caused by insufficient GPGPU memory to store the data sets and needed temporary information.


**Table 4.6** – Total muscle processing times on our five subjects and reduced data sets. Missing results are caused by insufficient GPGPU memory to store the data sets and needed temporary information.

	Size	CPU	GTX 470	GTX 660	GTX 750i	M2050	Grid K520
Subject 1	469 MB	88 s	-	30 s	22 s	43 s	30 s
Subject 2	493 MB	90 s	-	31 s	24 s	45 s	31 s
Subject 3	533 MB	98 s	-	35 s	25 s	45 s	34 s
Subject 4	534 MB	87 s	-	36 s	26 s	45 s	34 s
Subject 5	605 MB	108 s	-	-	29 s	51 s	38 s
Subject 1 (crop)	175 MB	33 s	15 s	14 s	8 s	16 s	11 s
Subject 2 (crop)	125 MB	24 s	11 s	9 s	6 s	11 s	8 s
Subject 3 (crop)	182 MB	31 s	15 s	13 s	9 s	17 s	12 s
Subject 4 (crop)	185 MB	31 s	16 s	13 s	9 s	17 s	12 s
Subject 5 (crop)	165 MB	32 s	14 s	12 s	8 s	15 s	11 s
Subject 1 (half)	86 MB	16 s	7 s	6 s	4 s	8 s	5 s

Our experiments show that the data processing time grows with the data set size. This applies to both, CPU and GPGPU. We have also found that GPGPU processing times are shorter than GPU times for air and muscle labelling. On the air labelling task, the GPGPU processing times are between  $\frac{1}{3}$  and  $\frac{1}{4}$  of the CPU times. These results are listed in Table 4.5 and shown in Figure 4.11. The improvements for the muscle labelling were smaller, but the GPGPU processing still ranges from  $\frac{1}{2}$  to  $\frac{1}{3}$  of the CPU times. These results are listed in Table 4.6 and shown in Figure 4.12.

During evaluation of these results, we have found that the CCL consumes a large share of the processing. The air labelling requires two CCL passes, however, for the muscle labelling, four passes are needed. Therefore, muscle labelling requires more time and has a smaller, but still large, gain over the CPU-only processing approach. While the CCL takes longer when performed on GPGPUs, we have found that operation like dilation, erosion, thresholding, and addition benefit. We assume that this stems from the fast, aligned memory access and the massively parallel processing of the GPGPU.

As part of the CCL implementation, we have included a step that reduces the number of labels by filtering for minimal size during the iteration. While this allows reducing the storage for the



**Figure 4.12** – Total muscle processing times on our five subjects (S1 - S5) and reduced data sets (air cropped and halved). Missing results are caused by insufficient GPGPU memory to store the data sets and needed temporary information.

LET, it requires processing time. We have tested two different approaches; the first one uses local memory and a fixed number of threads. This leaves too much work on the threads with growing block sizes. The second approach operates block independent and allows the distribution on many threads but it needs to use atomic operations. We have found that the Kepler GPGPUs have a fast implementation of atomic operations and showed speed improvements over the first approach. On Fermi GPGPUs instead, the use of atomic operations is too costly so that we have used the second approach.

In conclusion, we have presented a GPGPU implementation of our air and muscle labelling steps. We have achieved a reduction of the CPU processing to  $1/2$  to  $1/4$ . In clinical environments, GPGPUs are used for image reconstruction in MR and CT scanners. This provides a wide availability and an easier integration into clinical routine. We have identified CCL as area for future work. The serial implementation on the CPU shows superior performance. Therefore, we propose to evaluate the use of the CPU for CCL. The loss from data transfer from and to the GPGPU should be smaller than the gain from the CPU based CCL.

## 4.6 Conclusion

The acquired image data requires additional processing. In this chapter, we have presented the correction of non-uniform intensities on an inter- and intrastack basis, including bias field estimation and compensation. Another artefact in the image data is noise. We have discussed the non-local means approach for noise reduction. On the cleaned data, we propose a method to label muscle tissue efficiently using morphological operations. These labels can be used for initialisation and segmentation. Finally, we propose an acceleration of these steps by using GPGPU devices.

---

## CHAPTER 5

# MODEL INITIALISATION

---





## 5.1 Motivation

For a good segmentation, the initial models have to be placed close to their final position. Deformable models, in particular, benefit from this. Different techniques have been proposed, including manual approaches, edge-based methods and the Hough transform. However, those are either too simple or too computational heavy for the complex data sets that we have acquired. We propose a method that localises joints by using the histogram of the labelled slices of the image data. Once we have localised the joints, the bone models can be placed and segmented. We use the bones to register a template data set with the muscles of interest. To compensate for changes in fat/muscle ratio and total circumference of the limb, we correct the models using the muscle labels. We end the chapter with a conclusion.

## 5.2 Bone Localisation

There exist several approaches for the localisation of bones. In this section, we will present a selection of these approaches including manual and landmark-based initialisation as well as automatic methods using edge processing and the general Hough transform. Then, we present our proposed method.

### 5.2.1 Manual Initialisation

Manual initialisation relies on user input to create a proper initial state. The user is asked to manually place a model, to select seed points or to localise landmarks. There exist many techniques for the interaction, most common is the use of mouse and keyboard. However, more advanced techniques like haptic input devices can be used. Just like the other approaches, it is sensitive to the quality of the images. However, human intuition can help in heavily distorted data sets. The required manual interaction is time-consuming and needs the expertise of the user not only in anatomy but also in the software used. Therefore, manual initialisation is not feasible for repetitive clinical applications but is often used in research scenarios. An additional problem is the reduced reproducibility. There can be high inter- and intra-observer variation in repetitions. Many segmentation approaches are sensitive to initialisation. Therefore, results often depend on the experience of the user.

A hybrid approach relies on the user for an approximate initialisation and then refines algorithmically. This helps to reduce the search space and cuts down computational costs and required time. Furthermore, it improves reproducibility.

### 5.2.2 Landmarks


The use of landmarks is an extension of the manual approach. Instead of translating and scaling the whole model, the user selects a set of strong landmarks in the image, and the initial model is automatically adjusted to obtain the best fit. In this way, the user input and hence the margin of error can be significantly reduced. Landmarks should be positioned at significant areas that can

be easily detected in the image. Examples would be areas of high curvature, e.g. a fissure in a bone, or a measurable location, e.g. the middle of the femoral shaft. An initialization of a model using seeds has been proposed by Neuenschwander *et al.* [NFSK97]. The correct scaling and positioning of the model have to be determined from the seeds. This warping can be done using the Thin Plate Spine (TPS) transform. If the number of landmarks is very small and not sufficient for a confident pose estimation, different positions can be evaluated using a cost function that incorporates a penalty based on image properties. An example for that can be found in [Sch11].

Using landmarks at key anatomical positions allows the invocation of prior knowledge. Other approaches extend landmarks with the inclusion of spatial knowledge. For example, Fripp *et al.* [FCWO05] proposed a cartilage initialisation system that exploits prior knowledge of the bone location.

### 5.2.3 Edges and Image Features

Edges and other image features are needed for initialization techniques like the General Hough Transform (see next section). These approaches rely on features and edges that lie on the bone surface. Complex image features like SURF [BETV08] and SIFT [Low99] have shown poor results because they identify features that lie within the soft tissue as shown in Figure 5.1a. Edge detection algorithms like Canny or Sobel detect the bone boundaries but also create many results inside the soft tissue (see Figure 5.1b).




**Figure 5.1** – SURF features (a) and edges (b) before the application of smoothing filters.

We have evaluated several smoothing filters to reduce the number of unwanted elements in the image, for example, fine structures in fat tissue. The min-max curvature flow filter[MSB96] has shown the best results<sup>1</sup>. It can be parametrised by radius, iteration number and time step. Our


1. See <http://math.nist.gov/mcsd/savg/software/filters/smooth/index.html> for an extensive com-

evaluation indicates that a radius above 1 leads to an insufficient smoothing of soft tissue. A comparison of different radius values is shown in Figure 5.2.



**Figure 5.2** – Influence of the radius in image smoothing filters: 1 (a), 2 (b), 3 (c) and 4 (d) voxels.

The next parameter we have investigated is the time step size: We have found that a value near 0.0625 provides a good balance between satisfying results and noise. Larger numbers increase the artificial noise. The evaluation has further shown that an iteration count of 10-15 is sufficient. More iterations just increase the calculation times but have little or no effect on the resulting image. Instead of that, we have evaluated the effect of multiple runs (applying the filter to the filtered image). The results show an improvement in the second run and only minor changes during the third run, as we show in Figure 5.3.




**Figure 5.3** – Effect of multiple filter runs: 1 (a), 2 (b) and 3 (c) runs.

If we compare the results of edge filtering and SURF feature detection in Figure 5.4, we can see that the SURF features are still unusable for bone boundary operations. In opposite to that, edge detection now has reduced the number of edges that do not lie on bone boundaries.

The next step will be using the edges to perform a general Hough transformation (see next section). By performing the filtering, we have drastically reduced the number of feature points that has to be evaluated. Furthermore, we were able to reduce the number of unwanted edges that do not lie on bone boundaries. The results could be further improved by tuning the edge detection (different algorithms and parameters).

---


parison of smoothing filters.



**Figure 5.4** – SURF features (a) and edges (b) in the smoothed image. The total number of features is largely reduced.

#### 5.2.4 General Hough Transform

In 1962, Hough [Hou62] was granted a patent for *Method and Means for Recognising Complex Patterns*. His approach uses templates and a voting space. For each identified image feature all possible templates create votes for possible poses. The highest vote represents a set of parameters for the best corresponding template. An example of an original image and the corresponding voting space is shown in Figure 5.5. Illingworth and Kittler [IK88] have presented a review of the Hough transform and Khoshelham [Kho07] has demonstrated an extension to 3D object detection.



**Figure 5.5** – Example of the Hough Transform for lines. Original image (a) and voting space (b).

The general Hough transform can be used for the initialisation of deformable models. Van der Glas *et al.* [vdGVBV02] have demonstrated a method to detect ball joints. Seim *et al.* have

proposed approaches for the hip [SKH<sup>+</sup>08] and the knee [SKL<sup>+</sup>10]. In 2010, Ruppertshofen *et al.* [RLS<sup>+</sup>10] have proposed a discriminative approach for lower extremities, which has been extended to a multi-level scheme [RKL<sup>+</sup>11]. The main use of the Hough transform to detect bones in medical image data is on CT images, as they show clearer boundaries than MR images.


To use the Hough transform, a distinct shape is required. The work of Milner *et al.* [MB12] has shown that the average diameter of the femoral head is 42.1mm for women and 49mm for men. This concrete parametrisation gives the possibility to reduce the voting space for the Hough transform.

We have evaluated the possibility of using the Hough transform in our data sets. We have encountered two major problems with it. First, the visibility of the bone boundaries was too weak to generate satisfactory results. Second, while the femoral head has a ball shape of predefined size, other structures, e.g. the tibia lack such simple geometrical features. To describe the variations in the knee joint, a much larger voting space is needed, leading to weaker results with higher uncertainty, as the votes are more distributed than for a simple ball shape.

Therefore, we are exploring other approaches including the use of histograms as described in the next section.

### 5.3 Histogram Matching for Joint Localisation

The natural structure of muscles is constant among the population. They are very thin at their attachments and tendons and continue to grow in diameter towards the centre. As most muscles are attached near the joints, a recurrent pattern can be found among the population. The muscle volume per transversal slice increases towards the centres of calf and thigh. If we include fat, skin and other tissues, the same observation can be made. By sampling the slices along the inferior-superior axis, we can generate curves (histograms) as shown in Figure 5.6. As we compare them with the different subjects, they all carry the same characteristics. By aligning the curves to a template with prior known locations of the joints, we can approximate the location in our subjects.



**Figure 5.6** – Histogram of tissue (left) and muscle (right) volume per slice for five subjects.

Several approaches for aligning curves have been proposed, many of them intend to target

problems like optical character recognition [Mar98, SKK03] or handwritten signature matching [EFV07]. They use different techniques for aligning time series data like proposed by [WG97] and [Jam07]. The human leg contains three major joints: ankle, knee, and hip joint. These joints determine the position of femur and tibia, which we need for further processing steps. While exploring the different techniques presented before; we have found that we can generate satisfactory results by using a best neighbour registration approach. We have identified three actions to modify the histograms: superior-inferior axis translation (SIAT), superior-inferior axis scaling (SIAS) and slice volume scaling (SVS). The aim of the registration is to increase the overlap between the histograms. We define the template histogram  $H_t$ , which is fixed, and the sample histograms  $H_s$ , which is moving. The slice axis starts at the inferior end of the data set, at the feet. A common zero between the data sets can be assumed due to the air masking technique which we apply. We then identify the minimal and maximal shared slice index between both data sets:  $S_{min}$  and  $S_{max}$ . The overlap is now defined as:

$$\text{Overlap}(H_t, H_s) = \sum_{s=S_{min}}^{S_{max}} \min(H_t(s), H_s(s))$$

We then measure the overlap between the two histograms and test all possible operations (SIAT +/-, SIAS +/- and SVS +/-). If any of these provides a better overlap, the best one is chosen and applied to the sample histograms. Once no more improvement is found, we reduce the values for scaling and translation by 50% and repeat the process. This progressive refinement allows a fast first convergence followed by refinement for best results.

Further, we propose the following improvements. The length of legs can be easily estimated. This allows us to define an approximate position positions for the joints. We can then define a neighbourhood around them, e.g. with a width of  $1/4$  of the total leg length. These resulting sub-histograms prevent errors which could be introduced from unrelated areas of the histogram: e.g. optimisation for the hip joint could result in worse results for the ankle. Another approach to improve overlap amount is to adjust the volume per slice by introducing a shared minimum and maximum. Since we are interested only in the slice index but not the actual volume at that slice, we can freely modify these volumes (as long as they are propagated among all slices).

The initial parameters for translation and scaling have to be chosen carefully since they can have a large influence on the total error after histogram matching. We have performed an evaluation of a wide range of parameters, based on our subjects and a template data set.

Finally, we can calculate different histograms from the labelled image data. These options are air or total tissue, muscle tissue or total tissue minus muscle tissue. These different histograms show unique characteristics, e.g. there is little muscle tissue at the knee joint, this is a property which is constant for all subjects. Also, the shape of general muscles, e.g. hamstrings or quadriceps is common, although the total volume depends on the level of training/activity of the subject.

### 5.3.1 Evaluation and Results

In this section, we evaluate our proposed method for histogram based joint localization. We iteratively refine the histogram overlap using SIAT, SIAS and SVS. Once a refinement is no longer possible, we reduce the values for scaling and translation by 50%. The number of levels limits the repetition of this process. Once we have reached the finest level, we can apply the inverse operations to the slice positions of the ankle, knee and hip in the template to receive their position in the subject.

We have used our five subjects and an annotated template data sets to evaluate our joint localisation. During our evaluation, we have made the following observations. While it is possible to use muscle and tissue label histograms together, we have found that there exist pairs of joints and labels that produce better results than others. Our experiments have shown that muscle tissue gives best results for the knee. Hip joint and ankle can be located better using the total tissue label. While it is possible to use the full data set for the registration of the knee, we have found that we can achieve better results for ankle and hip joint by only using their close neighbourhood ( $\frac{1}{3}$  leg length around estimated joint location). During development, we assumed that an initial alignment of the slice volume through scaling (SVS) so that they share the same min-max values. However, our tests have found no improvements from that. Overall, the SVS only played a minor role in the quality of the results, disabling it resulted in only minor losses of accuracy.


The registration process runs independently for each subjects. To influence it, we can set the initial scaling values for SIAS and SVS, SIAT step size and the number of levels of refinement. We also apply a limit for the total translation and scaling. We have applied an exhaustive search to find the optimal parameters for all our subjects. The metric to evaluate the parameters was the sum of the absolute error per subject. We will present the individual results below, in Figure 5.7 we show the evaluation of the absolute error over a range of initial parameters for SIAT and SIAS. As the SVS has shown little or no effect on the output, we have omitted this parameter.

The knee histogram matching worked best with the muscle labels. We found an initial SIAT step size of 32 slices with a total maximum translation of 50 slices to work best. SIAS and SVS have both been initialised with the factor 0.01 ( $1 \pm 0.01$ ); the maximum accumulated scaling was limited to a factor of 0.2. For the knee, we have extracted a sub-histogram for matching. It was  $\frac{1}{4}$  of the data set length long and had  $\frac{1}{3}$  of the dataset length as offset, so it was ranging from  $\frac{4}{12}$  to  $\frac{4}{12}$  of the total data set). We show the results for every subject using these parameters in Table 5.1.

**Table 5.1** – Histogram matching results for the knee joint using the muscle labels.

Subject	1	2	3	4	5
Error (slices)	0	6	1	1	2
Error (mm)	0	9	1.5	1.5	3

The hip joint histogram matching to locate the femoral head worked best with the tissue labels. We found an initial SIAT step size of 7 slices with a total maximum translation of 50 slices to work best. SIAS has been initialised with the factor 0.01 ( $1 \pm 0.01$ ), the maximum accumulated



**Figure 5.7** – In this figure, we present an analysis of the total absolute error depending on the initial parameters of the registration process. The x-axis shows initial step size for the translation along the superior-inferior axis. The y-axis shows the initial scaling factor for the superior-inferior axis. Slice volume scaling has been omitted due to its only small influence on the results.

scaling was limited to a factor of 0.2. SVS was initialised with a factor of 0.07. Also for the hip joint, we have extracted a sub-histogram for matching. We have selected the superior 10% of the data set as our acquisition protocol cuts off above the femoral head. We show the results for every subject using these parameters in Table 5.2.

**Table 5.2** – Histogram matching results for the hip joint using the tissue labels.

Subject	1	2	3	4	5
Error (slices)	0	2	0	4	1
Error (mm)	0	3	0	6	1.5

The ankle histogram matching worked best with the tissue labels. We found an initial SIAT step size of 43 slices with a total maximum translation of 50 slices. SIAS has been disabled, SVS has been set to 0.15. For the ankle, we have extracted a sub-histogram for matching. It was  $\frac{1}{4}$  of the data set length long and started at the 30th slice of the dataset (from the inferior end). We had to omit these slices as the image quality was very low. As part of the acquisition protocol and as a compromise for the acquisition time, the ankle is not very visible in the data set and barely included. This has been a big challenge to get a close localization of the joint, in particular for Subject 3 and also for Subjects 4 and 5 where the ankle was within the first slices of the dataset. While the results for Subjects 4 and 5 still lie in a sufficient range for initialization, Subject 3 has a large error. We show the results for every subject using these parameters in Table 5.3.

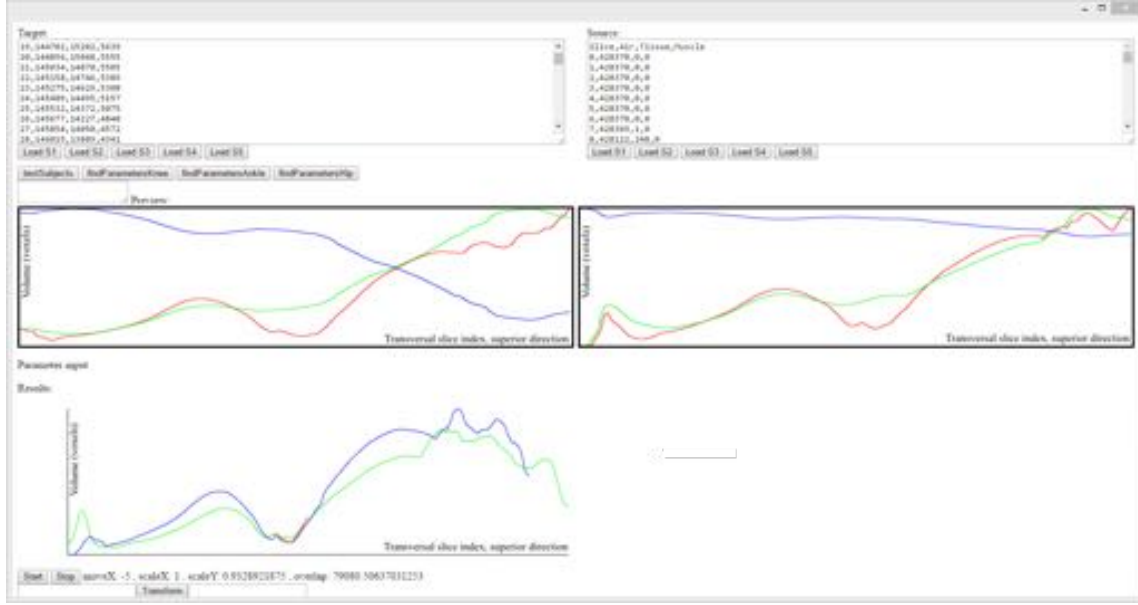
**Table 5.3** – Histogram matching results for the ankle joint using the tissue labels.

Subject	1	2	3	4	5
Error (slices)	5	3	25	10	11
Error (mm)	7.5	4.5	37.5	15	16.5

Initially, we have implemented the histogram matching in ECMA script, embedded in a web



page while using the HTML5 canvas element for visualisation. An example of that GUI can be seen in Figure 5.8. However, insufficient garbage collection during the exhaustive exploration of the parameter space has required the to a language with more control over memory allocations (C++). The implementation in C++ also allowed an integration in our DMFramework to make this feature available for the automated processing and segmentation pipeline.



**Figure 5.8** – Screenshot of the user interface of the initial histogram matcher interface. Three canvas (HTML5) elements present the histograms to the user. The top left histogram (a) is the fixed one with muscle (red), tissue (green) and air (blue). The top right (b) is the moving histograms. In the bottom histogram (c), we can see the result of the registration. The initial moving histogram (red) has been aligned (blue) to the fixed histogram (green).

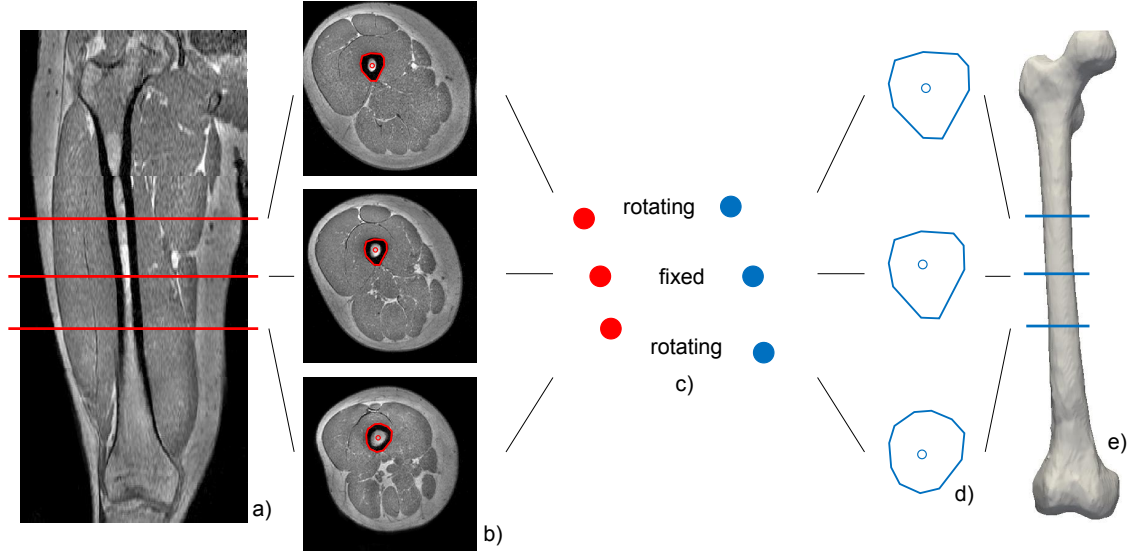
### 5.3.2 Bone Placement

Using the location of the joints, we can scale the models of the bones based on their length and place them along the superior-inferior axis. However, they are not properly located in the transverse plane yet. Tibia and femur show similar tubular structures, with both of them having a long shaft. Our approach is to locate this shaft in the MR images and to use it for the aligning of the meshes. In the following, we will describe the process using the femur as an example.

In the multi-channel image data from the mDixon scans, the in-phase scans show clear low intensities along the femoral shaft. To identify them in individual slices, we first remove the areas labelled as air. The raw muscle label also labels the cortical bone. Therefore, we can extract the low-intensity values inside the raw muscle label. We then filter it for the largest connected component to get the cut through the femoral shaft. This segment  $S_f$  is an approximation of the shaft, excluding the bone marrow. Finally, we calculate the centroid of  $C_f$  for the given slice.

The concrete placement of the subjects in the scanner leaves a small range of rotation for the bone, e.g. through different angles of the knee joint or rotation of the feet. To find a the correct placement of the mesh, we need at least three different points to align. The main idea is to take

the centroid of the femoral shaft in the image data  $C_f$  and the centroid of the femoral shaft in the mesh  $C_m$  (which can be easily calculated through a cut) and to align them. This principle is shown in Figure 5.9. We use the slice at the centre of the model on the superior-inferior axis and two slices  $1/5$  of model length inferior and superior to it. We align both sets of centroids using the centre point and optimise the rotation to match the remaining centroids best.



**Figure 5.9** – To align the femur model in the transverse plane, we extract three slices from the image data (a) and the mesh (e). From these, we extract the centroid of the femoral shaft (b,d). We overlay the middle centroids and optimise rotations to match the others (c). The resulting transformation is applied to the mesh.

We then use our deformable model framework DMFramework for the exact segmentation. In the mDixon in-phase image, the bone is easily visible. Therefore, it is sufficient to use this single-channel image and a robust gradient force to generate the bone segment.


Finally, we have used the centroid-based method for placing bones within the transverse plane. We present it on Subject 1 and compare it to a manual segmentation of the femur. The models have a mean distance of 4.27mm with a minimum of  $< 0.1$  mm and a maximum of 16.79 mm. The RMS is 5.48mm. This is a very good approximation; the remaining error will be removed during the actual segmentation task. The target and aligned femur can be seen in Figure 5.10. Further comparison has revealed an SDC of 0.72.

## 5.4 Registration of Muscle Atlas using Bone Location Information

For the initial placement of the muscle models, we use the femur. We can register the segmented femur with the template femur. The resulting transformation from template space to model space can then be applied to the template muscle models.

We are using an atlas based on the UltimateHuman<sup>2</sup> model and focus on the thigh. We apply

2. <http://www.cgcharacter.com/ultimatehuman.html>, based on MRI data



**Figure 5.10** – Evaluation of the femur alignment. The target (a) and aligned model (b) have been compared to show the distances from the surface of each other (c). Low distances are shown in red, medium in yellow and green and high distances in blue.


the automatic segmentation tools for the femur of Schmid *et al.* [SKMT11]. We then use the segmented femur for registration. First, the centroid of the femur in the atlas is translated to the centroid of the image femur. Then, we perform a best neighbour registration between both femur models. We restrict translation to 5cm per axis in total, rotations to  $20^\circ$ , and scaling to a range of 70-130%. The registration converges after 50 to 100 steps. The resulting transformation is combined with the initial translation and applied to the muscle meshes of the atlas. This process is illustrated in Figure 5.11.

## 5.5 Label-based Muscle Mesh Correction

The muscle labels can be used to register a muscle mesh atlas. For this, a given collection of muscle meshes (e.g. from prior segmentation, or other sources like the Visible Human Project [Ack98]) is transformed into image space so that it maximises overlap with the muscle volumes in the image. These atlases, also known as templates, are pre-registered in a first coarse step, followed by the fine corrections.

Often the actual circumference of the lower limbs differs from person to person, based on body weight and the amount of exercise. The muscle meshes, registered only based on the femur, are not adapted to that. They either stick out of the leg or fall short. Therefore, we propose the following method to correct the meshes by using the muscle labelling we previously created as a guide for correction. As an input, we use the air image, the muscle labels, the femur model and muscle meshes. The main idea is to slice-wise cast rays circularly from the centre of the femur and to correct the outermost muscle mesh contact to align with the outermost muscle label contact. This is done on every transversal slice along the femur.

In each slice, we first calculate the convex hull of the femur. Figure 5.12 shows that the shape of



**Figure 5.11** – Volumetric muscle models can be registered to the labelled muscle images. An initial correspondence is established using the previously segmented femur. The transformation gained from this step can be used to transform the muscle template model as well. After this step, the muscle labels can be used for corrections, e.g. for different circumferences of the thigh and different stages of physical training.


the femur can be very irregular, especially near the condyles and the femoral head. Our approach requires that the rays do not cross, which is only guaranteed on convex shapes. We are using 36 rays, every  $10^\circ$ , as can be seen in Figure 5.14. We always consider two consecutive rays at one time. The following explanations are shown in Figure 5.13, where they can be retraced more visually. The first ray hits the femur surface in  $A$ ; the second ray hits it in  $B$ . The outermost atlas mesh point for the first ray lies in  $C$  and in  $D$  for the second ray. The vectors  $\vec{a} = \vec{AC}$  and  $\vec{b} = \vec{BD}$  follow the rays. Let the mesh vertex inside the area between  $\vec{a}$  and  $\vec{b}$  be  $X$ . We can define the position of  $X$  with respect to  $\vec{a}$  and  $\vec{b}$ . Let  $\vec{g}$  be defined as  $\vec{BA}$ . We now define the  $\vec{h}$  that traverses  $X$  to be parallel to  $\vec{g}$ . Let  $S$  be the intersection between  $\vec{a}$  and  $\vec{h}$  and let  $T$  be the intersection between  $\vec{b}$  and  $\vec{h}$ . We can now define  $\vec{h}$  as  $\vec{TS}$ .

The points  $S$ ,  $T$ , and  $X$  can be expressed relatively to the vectors they lie on. Their scalars  $s$ ,  $t$ , and  $h_1$  and  $h_2$  are defined as

$$s = \frac{\|AS\|}{\|\vec{a}\|}, t = \frac{\|BT\|}{\|\vec{b}\|}, h_1 = \frac{\|XS\|}{\|\vec{h}\|}, h_2 = \frac{\|XT\|}{\|\vec{h}\|}.$$

When we follow both rays for the contact with the outermost labelled mesh, we find the points  $C'$  and  $D'$ . We then define the vectors  $\vec{a}' = \vec{AC'}$  and  $\vec{b}' = \vec{BD'}$ . Using the scaling factors  $s$  and  $t$ , we can define  $S'$  and  $T'$ :

$$S' = A + s \cdot \vec{a}', T' = B + t \cdot \vec{b}', \vec{h}' = \vec{T'S'}$$



**Figure 5.12** – The irregular shape of the femur (a) has many concave areas in the transversal cuts (b) along the superior-inferior axis. On these cuts, we can calculate the convex hull (c).

Finally, we can use the scaling factor  $h_1$  or  $h_2$  to find  $X'$ :


$$X' = T' + h_2 \cdot \vec{h}' = S' - h_1 \cdot \vec{h}'$$

We repeat this for all mesh vertices inside the circular segment, and then we proceed with the following segment. We apply this to all slices along the femur to retrieve the correction for all meshes in the muscle mesh atlas.

For rays facing the other leg, the outermost labelled muscle point may be falsely attributed to the other leg. To prevent this, we stop the search along the rays once air is encountered for the first time. This approach has shown problems when the fat of both thighs touches. Also, there is no air region in the pelvis, resulting in false corrections. Therefore, we slice-wise sample the transverse centroids along thigh and calf. If we take their average and connect them longitudinally, we get a line along which we can position a sagittal plane to stop the rays. This effectively separates both legs.

### 5.5.1 Evaluation and Results

To evaluate the outcome of correction, we run a visual comparison and we analyse the overlap of muscle label and mesh before and after correction. The results of the visual comparison can be



**Figure 5.13** – To correct the muscle mesh atlas we divide the femur ( $A, B$ ) into circular segments limited by  $\vec{a}$  and  $\vec{b}$ . The outermost points along them of the muscle mesh are  $C, D$  and in the muscle labels, it is  $C', D'$ . We use these points to calculate the correction  $X'$  for a point  $X$ .


found in Figure 5.15 (a, b, with hidden skin layer and c, d, with skin layer). In comparison with only the muscle mesh, the uncorrected version shows large displacements of the template outside of the muscle tissue. This is mainly visible on surface muscles as they are closer to the boundaries and more sensitive to anatomical differences between the subjects. However, this displacement also affects deep muscles, although it is not visible here. Depending on the variations between template and subject, the differences can be large. This is clearly visible in the comparison with the skin surface included.

In Figure 5.16, we have evaluated several muscles of the thigh, before (labelled as O, original) and after (labelled as C, corrected) the muscle label based mesh correction. We also list the difference. The evaluation has been carried out on Subject 3. From our results, we can see that deep muscles like the adductor brevis, the adductor magnus or the pectineus are already almost fully covered by the muscle label. After the correction, these values remain high and there little to no change.

For the surface muscles, large changes can be observed. As seen in Figure 5.16, muscles like biceps femoris, rectus femoris and gluteus maximus (40% are lying outside the imaged volume) benefit from the correction. They show vast improvements in overlap, with up to additional 30% of their volume included. In particular, the sartorius muscles with its thin size which wraps around the leg on the surface benefits, with an overlap increase from 50% before correction to 90% after successful correction. The results presented so far were acquired in Subject 3. We evaluated the other subjects to ensure similar correction results between the different anatomies in the subjects. As the main differences were shown in the surface muscles, we present a selection of them in Table 5.4 for each subject. These results support the findings in Subject 3 and show that the approach can be applied to subjects of different sizes, weight and physical fitness.

In our tests, we have used a default number of 36 rays per slices, or every  $10^\circ$ . The output depends on this, as the actual muscle contour is assumed to be linear between to rays. In Figure 5.17a, we illustrate two example in which this assumption is violated. The presence of concave or convex







**Figure 5.14** – The muscle mesh atlas correction is performed on circular sectors of the image. First, the convex hull of the femur is calculated (white). Then rays are cast from the centroid of the hull. The blue rays go to the outermost muscle tissue; the orange rays stop at the outermost muscle mesh vertex.

shapes can explain the difference from the 100% overlap that could be initially expected. We have evaluated different configurations of rays, 18, 36, 72 and finally 144 rays (every  $20^\circ$ ,  $10^\circ$ ,  $5^\circ$ , and  $2.5^\circ$ ). The results are plotted in Figure 5.17. Our results show that only 18 rays already give an improvement in overlap, whereas more than 36 rays only provide little additional improvement. The difference between 72 and 144 rays is minimal. However, the calculation is very efficient and additional rays only produce a little overhead. Therefore, a total of 72 rays seems a good selection.

In our proposed method, we only perform correction on individual slices. If mesh elements span among multiple slices, the slices in between are not taken into consideration. While it would be possible to perform a correction along the axes, the resulting output is sufficient for the



**Figure 5.15** – Visual comparison of the muscle mesh adaption. A) and b) show the muscle tissue; c) and d) add a (semi-transparent) tissue layer. A) and c) show the state before correction - the blue muscle meshes clearly are not fully inside the muscle tissue. B) and d) show the results after correction; the muscles meshes are now within the labels.



**Figure 5.16** – Evaluation of muscle mesh correction for several muscles. The graph shows the overlap of the original mesh (O) with the muscle label, the overlap after the correction (C) and the difference (diff). A perfect overlap is 1.0; this evaluation has been performed on Subject 3.


initialization step.

In conclusion, we have found that our proposed muscle label based mesh correction shows good results. We have shown a visual comparison that underlines the correction of misplaced muscles. A comparison of overlap between muscle labels and corrected meshes has found that deep muscles already show a good overlap, whereas surface muscles initially have a bad overlap. The overlap improves vastly using our correction technique. Nonetheless does our methods not produce a full overlap in all cases. This can be attributed to spikes and a low sampling rate during correction. However, this step is followed by the actual segmentation that will correct these problems. In pursuit of a fully automated segmentation pipeline, we can control for errors in the initial placement during this step. If bad values for overlap of deep muscles are found, this could trigger a manual inspection as this is unexpected.



**Table 5.4** – Muscle mesh correction of surface muscles among the subjects.

Muscle	Subject 1		Subject 2		Subject 3		Subject 4		Subject 5	
	before	after	before	after	before	after	before	after	before	after
biceps femoris	69%	90%	90%	97%	73%	93%	94%	97%	51%	98%
gluteus maximus	50%	81%	53%	88%	34%	64%	58%	83%	22%	84%
rectus femoris	81%	94%	86%	99%	79%	99%	96%	99%	42%	99%
sartorius	57%	78%	81%	93%	45%	88%	82%	93%	55%	95%

**Figure 5.17** – The number of rays used for muscle mesh correction can improve the overlap with the muscle labels. The graph (left) shows the overlap for different totals of rays used (18, 36, 72, 144). The figures on the right illustrate the effect of too few rays on the transversal (top) and sagittal (bottom) plane.

## 5.6 Conclusion


Deformable model segmentation requires proper initialisation, meaning that initial models have to be placed close to their approximate location. We investigate a slice-based histogram to identify the position of joints in the subjects. Using these positions, we can scale the models for femur and tibia. We identify the shaft of the bones in the transverse plane to get a proper placement. Then, we register a template of the leg, including muscles, using the femur as a reference. Throughout the population, the diameter of the leg varies, due to the amount of fat and training, resulting in changing amounts of muscle volume. By exploiting the previously created muscle labels, we correct the template to fit within the subject's leg.

---

## CHAPTER 6

# IMAGE SEGMENTATION

---



## 6.1 Motivation

Based on our discussion of segmentation techniques in Chapter 2, we have decided to use a deformable model approach for the segmentation. We present a framework that implements the core components and discuss the use of existing libraries to extend functionality. Our proposed framework relies on libraries like VTK, ITK, Bullet Physics and OpenMesh. Lua scripting allows controlling these components from external scripts. A specialised image manager handles the multi-channel images and custom forces access them during the segmentation process. Section 6.2 describes these core components of this framework. In Section 6.3, we describe our approach to evaluation. We present the tools and metrics for manual segmentation and then compare our DMFramework results to this ground-truth. To overcome system dependencies, we propose a web-based viewer for 3D multi-channel image data and segments and discuss a mesh streaming approach during segmentation. This work is presented in Section 6.4.

Segmentation of complex structures is often a multi-tiered process. First, a coarse model is used, which has low computational costs and allows fast adaptation. Through multiple stages, the model is refined and adapts to fine structures in the image. This refinement comes with an increase in computational cost and reduces the number of possible steps. This pyramid approach is one-way, and it is not possible to return to a more coarse model, e.g. when problems are only found at a fine level. We propose a method to overcome this by running all resolutions in parallel. The models are coupled with forces and allow the switching between resolutions at any time. We discuss our method and the used forces and synchronisation techniques in Section 6.5.

## 6.2 Deformable Model Framework Implementation

In this section, we introduce the design and implementation of our Deformable Model framework *DMFramework* that we propose for the segmentation of the individual muscles. First, we discuss the main parts of deformable models in general. Then, we present our design and describe implementation details. This includes meshes, image data handling, physics, forces and approaches to parallelisation.

### 6.2.1 Main Components of Deformable Models

We have selected Deformable Models (DM) for segmentation in this work. DMs provide reliable performance even in areas with little image information. Furthermore, they can carry an initial shape that can limit the possible changes of the model. A full discussion of the properties of DMs, related and detailed justification of our choice can be found in Sections 2.9.3 and 2.13.

Our thorough review of related work, e.g. from Bredno *et al.* [BLS03], has revealed the DMs are commonly divided into components. We have identified the following main components of DMs:

- Meshes, that model the anatomical structures
- Forces, that will be applied to the meshes

- Images, that provide input to the forces
- Physical simulation, that applies the forces to the mesh
- Output, to provide user feedback

In the next section, we will describe how we have addressed these main components. However, there are additional things to consider which we briefly discuss next.

Time and costs limit the amount of image data that can be acquired. While the core parts of the leg are fully acquired, a cut at the hip has to be made. This results in some of the muscle heads lying outside the volume. Proper boundary conditions are needed to handle these muscles. Abort criteria can be defined to limit the processing time. The deformable model process is iterative. Towards the end, the model converges and undergoes only small changes. These changes, at some point, fall below the voxel size and the calculation can be considered complete.

Clinical systems need to be very focused and should minimise user interaction. However, in our research environment, flexibility to test and develop new approaches is crucial. Therefore, we propose to include support for scripting. This way, different program flows can be evaluated and platform extension is easier. These scripts can also be narrowed down based on user interaction, giving more flexibility in targeted audiences. Many components of DMFramework also can be applied more generic and, therefore, can help during image processing and model initialization.

Finally, a thorough evaluation of the results is needed. Depending on the data set, user and application, the ground-truth to compare the segmentation results with may be provided in different formats. While our segmentation generates meshes, many clinical experts annotate the medical image data directly so that labelled voxel data is generated. In our DMFramework, we provide methods for mesh-to-mesh and label-to-label comparison. We also provide access to techniques to convert between both representations. In Section 6.3.2, we present evaluation metrics.

## 6.2.2 Structure and Components of the DMFramework

Our DMFramework has a layered structure to connect its components which are group by functionality and need to communicate. In this section, we will this structure in details and we will present and motivate our decisions for the individual components.

### 6.2.2.1 Meshes

We use meshes to represent the anatomical structures that we segment. These 3D representations not only represent the shape and structure, e.g. topology, but can also be used to carry additional information like labels, textures, and weights. Meshes can be represented in different ways, e.g. as simplex meshes or as triangle or quad meshes. In our applications, we have found triangle meshes to be useful. They show great flexibility and a wide support in existing libraries. Also, visualisation pipelines are usually optimised for triangle rendering.

Other than expressing a geometry, a mesh library should support loading from and saving to dif-

ferent formats. Models can stem from different sources, e.g. medical professionals or designers. They use vastly different tools and create different file formats. A wide support allows a better integration in a wide range of scenarios.

The internal representation of meshes in a library can be important for the performance. Optimisations often benefit one group of applications and hinder another. For example, the half-edge representation allows fast polygon traversal but increases memory consumption. The Laplacian smoothing as another example requires the processing of the neighbourhood of vertices in the mesh. To work efficiently, several algorithms for mesh processing are often needed, like decimation to create lower-resolution version, smoothing or spatial access structures like trees.

**Table 6.1** – Comparison of different mesh libraries by features, license, and last update.

Library	License	Changed	Algorithms
Trimesh2 <sup>1</sup>	GPL	recent	subdivision, smoothing, curvature estimation, triangle stripping, various other simple manipulations
OpenMesh <sup>2</sup>	LGPLv3	recent	smoothing, decimation, user-defined-attributes
Halfedge Mesh Lib <sup>3</sup>	unknown	1999	-
VCGLib <sup>4</sup>	GPL	recent	many
GTS <sup>5</sup>	LGPL	2006	Delaunay Triangulations, geometric refinement, coarsening, bounding boxes, kd-trees
CGAL <sup>6</sup>	(L)GPL / commercial	recent	several

Based on these requirements, we have evaluated a number of libraries. An overview of our findings can be found in Table 6.1. As a result, we have selected OpenMesh [BSBK02]. It offers support for a range of file formats and good control of the underlying data structures. Several algorithms for decimation are provided. The library is actively maintained and receives constant updates. The documentation is extensive.


### 6.2.2.2 Forces

We apply multiple forces to our models to transform them into their desired shape. These forces are divided into internal and external forces. Internal forces are applied through material properties by the physical simulation. External forces are based on other inputs, for example the image data. Some of these external forces have been presented in Section 2.9.5.

Several forces are required for satisfactory results. For internal forces, we use the shape preserving properties of the Bullet soft body implementation. We also apply Laplacian smoothing. To ensure a constant movement, we have added a small balloon force.


1. <http://gfx.cs.princeton.edu/proj/trimesh2/>
2. <http://www.openmesh.org/introduction/>
3. <http://www.cs.sunysb.edu/~gu/software/MeshLib/index.html>
4. [http://vcg.sourceforge.net/index.php/Main\\_Page](http://vcg.sourceforge.net/index.php/Main_Page)
5. <http://gts.sourceforge.net/index.html>
6. <http://www.cgal.org/>

In the previous sections, we have presented strategies for the robust segmentation of air and bones. Both have shown good results. Therefore, we can consider these valid and can assume that muscle must not enter air nor bone volume. We model them as rigid bodies with infinite mass as shown in Figure 6.1 for our simulation. They cannot be penetrated by the models. This also means that they cannot be deformed by other models and can be neither translated nor rotated through, for example, high impacts. Together, bone and air are used to exclude the models from prohibited, anatomically impossible regions.



**Figure 6.1** – Air and bones are modelled as rigid and fixed objects. They cannot be penetrated by the muscle models and prevent them from reaching unrealistic or impossible positions.

The muscle labels give an additional indication for the models. However, unlike bone and air, they cannot be considered rigid. While the labelling has shown good results, the proper labelling near the joints is challenging due to the many different types of tissues (see also Section 4.4) and fine structures. Therefore, we only use the muscle labels as an indication and pull the models to stay inside the muscle labels. For this, we pre-calculate a vector image where the vector values point towards the nearest muscle labelled voxels. The principle is shown in Figure 6.2.




**Figure 6.2** – Working principle of the muscle label force. Out the muscle labels (non-red), pre-calculated vector point towards the muscle label. Additionally, the distance to the next muscle label is stored for weighing.

Generating the pre-calculated image for muscle label force is a process that requires multiple

steps. First, we extract the boundary of the muscle label image. For the non-muscle voxels on the boundary, we calculate a vector pointing towards the muscle labels. For this, we traverse the N26 neighbourhood of the voxel. For each neighbour, that is part of the muscle label, we add a vector pointing towards it. The resulting vector then points towards the average muscle label location, as shown in Figure 6.3a. We normalise the vector and store it together with the distance to the muscle label (in this case, the distance is always one). In the special case of a non-muscle voxel surrounded by muscle voxels, the resulting vector will be zero (see Figure 6.3b). This can be neglected since this should not happen using our muscle labelling technique. Furthermore, a single voxel within the muscle label will not have a significant impact on the results.

After this first step, the voxels on the boundary of the muscle tissue contain vectors that points towards the muscle tissues as shown in Figure 6.3c. Next, we identify the boundary of the boundary and calculate the new vectors using the previously calculated vectors (see Figure 6.3d). We repeat this process until all voxels outside the muscle label have been processed. All voxels within the muscle label get a zero vector assigned.



**Figure 6.3** – Pre-calculation process for the muscle force. The initial labels are created based on the direction to the surrounding muscle labels (a). A special case is an isolated voxel inside the muscle label which carries a zero vector (b). After the first step, the boundary of the muscle label contains vectors pointing towards the muscle label (c). This process is iteratively repeated until all voxels are processed.

However, we have found that this iterative process is very time-consuming. It requires many traversals of the full data set. This means high memory access costs, resulting in a calculation that can take several minutes. Therefore, we propose to use the signed distance proposed by Maurer and Raghavan [MR03]. Their calculation can be performed in seconds and better suits our requirements. Next, we traverse the signed distance image and introduce zero-distances for muscle tissue and calculate local gradient vector. Figure 6.4 shows the vector field and the distance map for cut through the thigh.

Finally, we apply gradient-based image forces based on individual and combined channels in our image data. Forces and channels can be weighted. The balance is important to target the correct image features; inter-muscular fat should be targeted for a good segmentation, whereas intra-muscular fat is misleading and should be avoided.

### 6.2.2.3 Image handling

The medical image data is the single most important input for the image segmentation. Our choice for a library needs to fulfil the requirements not only from a segmentation point of view

but also for the processing steps described in the previous chapters.


Image data from many sources and processing tools can come in a large number of formats, like Dicom, MetaImage, NII, or even a stack of single 2D images. We need our image library to be able to load and save in a large number of formats to reduce the need for additional tools and intermediate file formats. This makes processing easier, faster and also eliminates a potential source of errors.

The image library should also support a large number of image filters for tasks like stitching (realised through pasting, resizing and interpolation) and more complex ones like noise reduction and even bias field correction. Other parts of the processing pipeline need support for filters like thresholding, erosion, dilation, multiplication and addition.

For the evaluation, we need to translate from the mesh to a label image. During the initialisation the inverse process can be required. Therefore, support for rasterisation/voxelisation and mesh reconstruction (e.g. marching cubes [LC87]) is needed. We want to calculate similarity measures between labelled voxel images. The library should have support for this or at least offer the means for a custom implementation.

We have shown in Section 5.3 that the histograms of air and muscle label images can be used for the initialisation. Therefore, the image processing library should have support for calculating these histograms. For the visualisation, the export of individual slices from the data sets is required.

In our review of available solutions, we have found that the Insight Segmentation and Registration Toolkit (ITK) [YAL<sup>+</sup>02] offers many filters for image manipulation and supports loading images of many different formats, e.g. DICOM, MetaImage. It is well-established and used worldwide by a large number of researchers and healthcare product developers. It is actively maintained and well-tested. While the heavy use of templates in the documentation examples are a high entry barrier, the gained flexibility outweighs this.



**Figure 6.4** – Final pre-calculation process for the muscle force. The resulting vectors point towards muscle tissue (a) and the distance map (b) are shown for a cut through the thigh.



#### 6.2.2.4 Physical simulation

The meshes and forces are combined during the physical simulation. During this step, the forces are applied to vertices of the meshes. While it would be possible to apply forces also to edges and faces, we have found that the use of only vertices keeps the code clean and is sufficiently powerful. The forces are used to deform the mesh towards its desired, anatomical shape. The deformation is subject to not only the forces but also the properties of the mesh. In the physical simulation, meshes can be either soft or rigid. Rigid meshes may be translated and rotated but cannot change their form. Fixed meshes are rigid and free from rotation or translation. For deformable models, soft meshes are used. These meshes can be deformed but are bound to certain specific properties that regulate the deformation and the influence of the force. Their properties include, but are not limited to:

- Mass of the object represented by the mesh  $[0, +\infty]$
- Linear stiffness coefficient  $[0, 1]$
- Area/Angular stiffness coefficient  $[0, 1]$
- Volume stiffness coefficient  $[0, 1]$
- Damping coefficient  $[0, 1]$
- Drag coefficient  $[0, +\infty]$
- Lift coefficient  $[0, +\infty]$
- Pressure coefficient  $[-\infty, +\infty]$
- Volume conservation coefficient  $[0, +\infty]$
- Dynamic friction coefficient  $[0, 1]$
- Pose matching coefficient  $[0, 1]$
- Rigid contacts hardness  $[0, 1]$
- Kinetic contacts hardness  $[0, 1]$
- Soft contacts hardness  $[0, 1]$

Neighbouring anatomical structures are in contact with each other. There is no empty space between them. They must also not inter-penetrate or self-penetrate. These are strict requirements for the validity of the results. Therefore, an important part of the physical simulation is the avoidance of these penetrations. This is realised through the process of collision detection and response. The above-listed properties control the reaction on two colliding meshes. This way, it is possible to have a larger muscle pushing a smaller one away. We can also define a strong response upon contact with rigid objects.

In our implementation, we consider the bones and the air rigid objects with an infinite mass. This ensures that they cannot be deformed (since we already know their exact shape, we do not need them to deform) or translated or rotated (the exact position is also already known). These rigid structures limit the range of motion for muscles and reduce the space for possible solutions.


The simulated deformation and the application and integration of forces is a continuous process. However, to be able to simulate it on computers, it has to be discrete. Therefore, the process is split into small time segments that are considered to be constant within. Our chosen simulation

engine, the Bullet physics library, assumes a simulation step to be 1/60th of a second. If we integrate larger times, they will be broken down into multiples of the base steps to ensure realistic results for force integration and collision response.

#### 6.2.2.5 Output

Direct feedback is needed for the development phase and the application phase in clinical environments. The users need to be able to assess the current progress of the segmentation and must be given the option to intervene, e.g. in a stalling process, to correct the initialisation in complex cases or to influence the segmentation through interactive forces.

We have designed a visualisation for Windows using the Visual ToolKit (VTK) [ABB<sup>+</sup>01]. It allows the 3D displaying of slices of the volumetric image data and the meshes. An additional 2D view allows seeing the slices with overlaid cuts of the meshes.



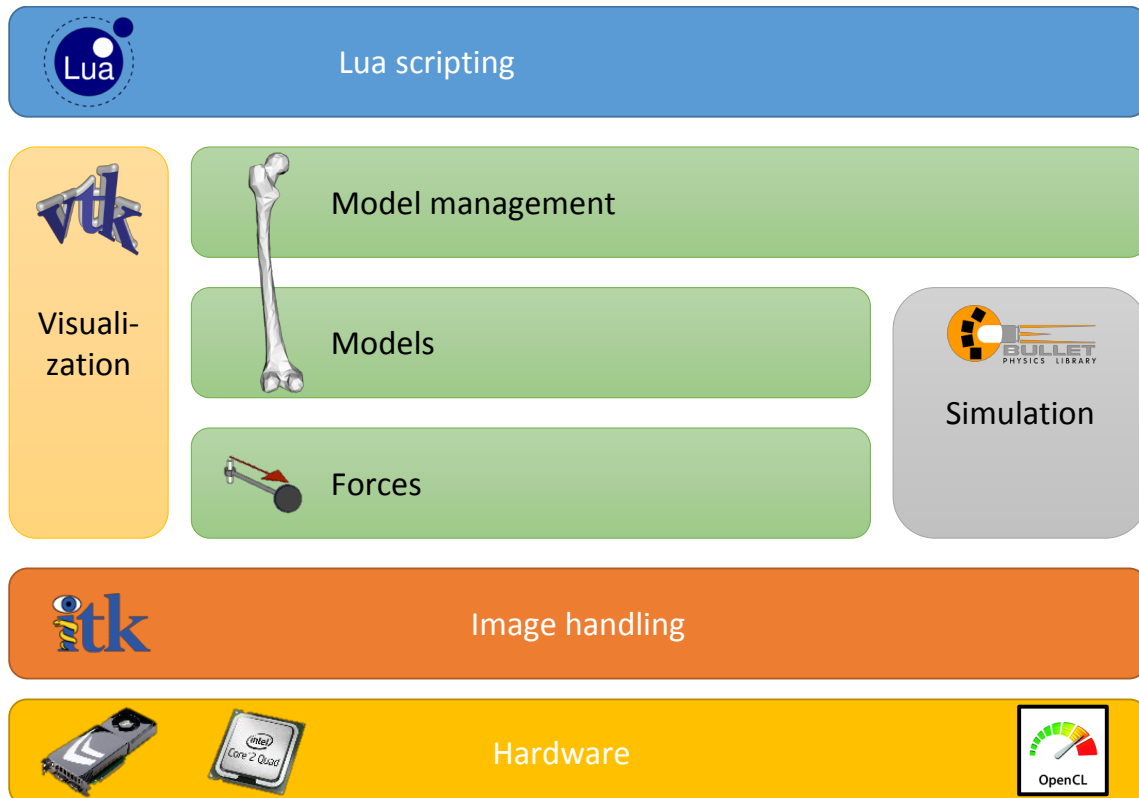
**Figure 6.5** – Visualisation of a femur model and a slice of the data set in 3D (left) and 2D (right) using VTK. In 2D, an outline of the femur model at the selected slice is shown.

While the visualisation using VTK offers a good first feedback, the options for more visual appealing output are limited (see Figure 6.5). The programming interface for interaction is not well-documented and requires a lot of supporting code. VTK is a large library that comes with many side-effects that require much work for integration into healthcare products, for example, all operations have to be run on the main thread. Therefore, we propose a different additional interface. We have integrated an interface that allows the streaming of meshes to arbitrary receivers. This could be web-based applications (see Section 6.4.2) or engines like Unity<sup>7</sup>, for example. This allows to leverage the power of other systems and avoids the implementation of niche features in our DMFramework. Another advantage is the possibility to create output that is specifically tailored to the target environment. A developer system requires different information than a radiologist's station.

7. <https://unity3d.com/>

### 6.2.2.6 Implementation

Based on the components and features defined in the previous section, we have designed our DMFramework to rely on well-tested components. Our segmentation framework is based on existing frameworks and structured into layers. We use VTK [ABB<sup>+</sup>01], ITK [YAL<sup>+</sup>02], OpenMesh [BSBK02], the Bullet physics library<sup>8</sup> and scripting using Lua [IdFFdF96]. We have paid special attention to be as platform independent as possible. Our DMFramework runs on MS Windows, Apple OS X and Linux.



**Figure 6.6** – The DMFramework architecture has been designed multi-layered. The basic components can rely different hardware and OS types. Custom components are used for model and force management. All parts can be scripted using the Lua language.

Besides existing libraries, we have contributed additional parts. The self-written image manager allows loading multiple data sets, for example from an mDixon sequence. It acts as a central storage for image data, data is loaded just once but can be accessed everywhere. This reduces memory consumption. Forces are loaded and managed using a self-written component.

The visualisation routines are based on the Visualisation Toolkit (VTK). The physical calculation of the forces is achieved using the Bullet physics engine. It offers well-tested code and highly configurable deformable soft-bodies. It also already contains methods for collision detection and response. The model manager, also self-written, allows to load multiple models and to connect them to different forces. It allows inside-outside testing of the model and many transforms like translation and scaling. It also supports mesh decimation to create reduced versions. For some

8. <http://bulletphysics.org/>

of its functionality, we rely on the OpenMesh framework, which has been selected after a comprehensive search. All the described components are scriptable using the language Lua. Figure 6.6 shows the components of the framework.

While the core of the framework is developed in C and C++, we provide Lua bindings for fast iterations and easy testing of new concepts. The main loop of the segmentation process is shown below. Prior to that code snippet, the images and models have been loaded, and the simulation has been initialised. In the loop, we first calculate the forces that are to be applied to the model, and then we integrate them using the simulation environment. Finally, the visualisation is updated.

**Listing 6.1** – Main loop for deformable model calculation

```
for i=1,1000 do
  if i%60*2 == 0 then
    print((i/60)*2, "seconds_simulated")
  end
  modelMan:CalcForces()
  sim:performSimulationStep(0.016666)
  visu:addModel(model,1.0,0.0,0.0)
end
```

A list of exposed Lua functionality can be found in the Appendix D.2. Our contributions to the framework have been developed in C++. It covers over 110 files and more than 15000 lines of codes. Additional contributions are the Lua scripts that combine the exposed functionality into applications or task in image processing and segmentation.

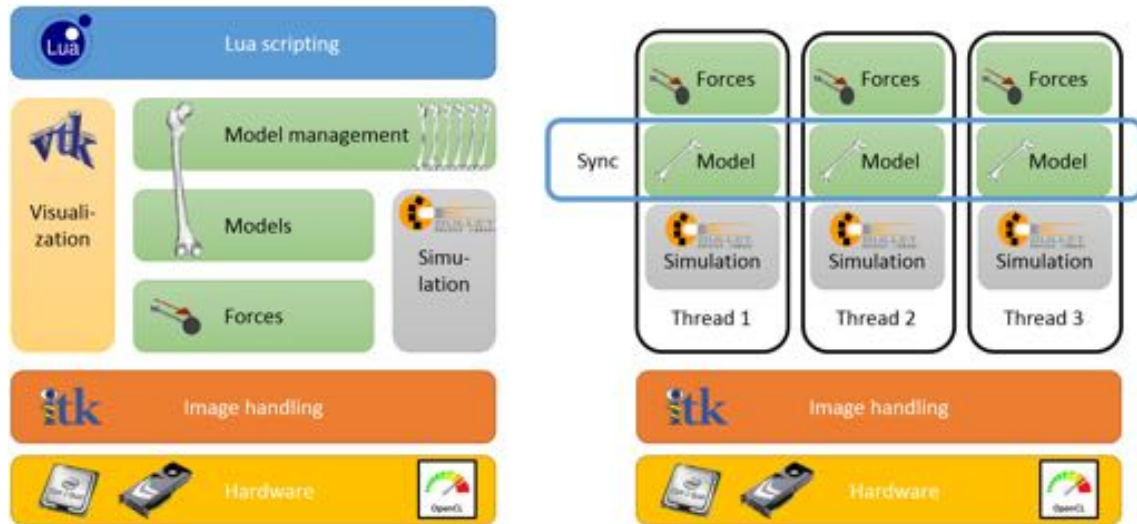
### 6.2.2.7 Parallelisation

Modern CPUs have large number of cores that allow parallel processing of data. Increasing the performance of individual cores has proven to be not economical since higher frequencies lead to increased power draw and heat generation. This problem can be addressed by splitting problems into smaller tasks that can be processed in parallel by multiple cores. This approach is taken even further by GPGPUs that combine a large number of reduced functionality cores with high bandwidth memory. To use CPUs and GPGPUs, different programming concepts are needed. There exist frameworks to bridge both platforms, like OpenCL<sup>9</sup>.

We have selected the components to be capable of using the underlying hardware at its best. The physics simulation library Bullet physics has been chosen since its pipeline also is available as OpenCL version which can exploit not only multi-core CPUs but also many-core GPGPUs. Similarly, ITK has been chosen not only because of the large number of supported file formats and filters but also because of its OpenCL support.

Image data sets can vary largely in size and can become very big, especially in multi-channel images. Therefore, we have designed the multi-threaded version of our framework to share the

9. <https://www.khronos.org/opencl/>



**Figure 6.7** – The Deformable Model framework used in this work has a layered architecture. Its single-process version (left) has several components (scripting, visualisation, model manager) which have been omitted in the presentation of the multi-threaded version (right side).

image data between the threads. Since all image operations are read-only, we do not need additional locking or access control. The image manager forms the shared foundation for the threads as seen in Figure 6.7. It is possible to run multiple independent simulations in parallel. In this case, each thread receives a different model and, therefore, needs its own model manager. To avoid interference, every model also uses its own simulation environment. Each thread that is created and started is assigned a unique ID accessible from the per-thread Lua scripting environment.

## 6.3 Evaluation

For the evaluation of our results, we use a ground-truth which represents the segmentation that is achievable through humans. While manual segmentation introduces some error sources of their own (such as dependency on individuals, limited visibility in image data, and experience), it is accepted as the standard for evaluation of medical image segmentation approaches. We will describe our experiences with this manual segmentation in the next section. Afterwards, we will discuss how to compare the manual and automatic segmentation results. We will discuss several metrics and evaluation approaches.

### 6.3.1 Manual Segmentation

In this section, we will describe how we have created the manual segmentation that will be used as a ground-truth for the evaluation. We begin by giving an overview different software tools for manual medical image segmentation and will discuss the three most promising candidates in more detail. We then discuss how the generated data has to be post-processed and present the required time for this process.

There exists a large number of software tools for medical image segmentation. They differ not only in their features but also in their licences and the availability of their sources (open or closed source). We have tested several tools, including Seg3D, TurtleSeg, YaDiV, 3D Slicer and MITK (Medical Image ToolKit). These candidates are compared in Table 6.2. We have discarded MITK and 3D slicer as they had complicated user interfaces and/or problems in loading our image format.

**Table 6.2** – Comparison of medical image segmentation software tools.

	References	Open source	License	Platforms
YaDiV	[FBW11]	open <sup>10</sup>	t.b.d.	Windows, OS X, Linux
Seg3D 2	[CIB13]	open	MIT	Windows, OS X, Linux
Turtleseg	[HYML05, THA11]	closed	commercial	Windows
Slicer	[FBKC <sup>+</sup> 12]	open	BSD style	Windows, OS X, Linux
Mitk	[WVW <sup>+</sup> 04]	open	BSD style	Windows, OS X, Linux

The three remaining tools, TurtleSeg, Seg3D, and YaDiV, showed some interesting features, that lead us to further investigate them. Detailed results can be found in the following sections.

### TurtleSeg

TurtleSeg<sup>11</sup> [HYML05, THA11] is an interactive tool that generates a 3D segmentation from a small number of contoured slices generated with minimal user guidance (or so intended). It was originally developed by the Medical Image Analysis Lab (MIAL) at the Simon Fraser University, Canada and the Biomedical Signal and Image Computing Laboratory (BiSICL) at the University of British Columbia, Canada, now it is a proprietary software owned by Oxipita Inc. A free test licence can be requested. The system requires the user to outline the desired structure on a few slices and tries to extrapolate from them. The system learns from the user input and proposes new slices that provide a larger increase in segmentation accuracy. The outlines, called contours, are generated using the Livewire algorithm, which follows gradients between support points.

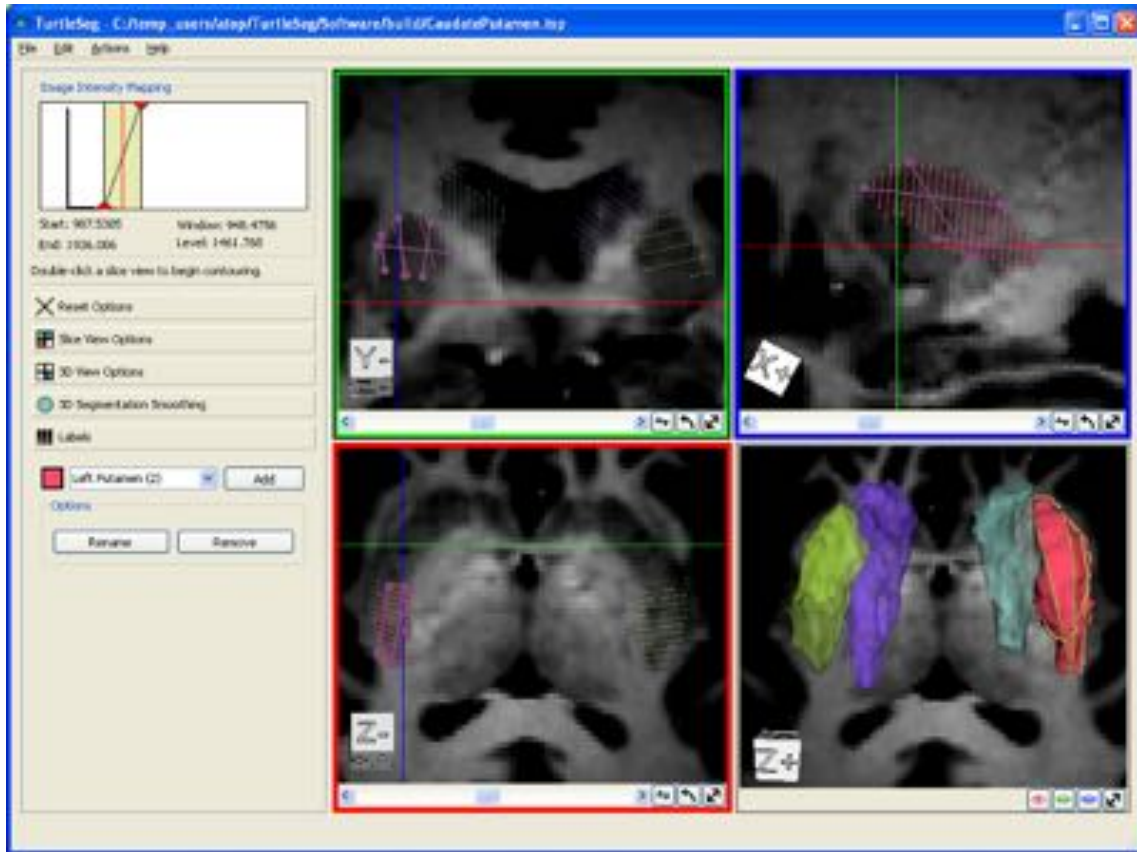
The user interface, as shown in Figure 6.8, shows 2D slices for outlining, a 3D rendering of the segment and only a few other elements, as the tool is fully focussed on the Livewire approach, which is the main strength of the tool.

The tool can automatically connect contours from different planes and slices and notifies the user in case it is not possible and allows the user to ignore this or to fix the contours manually. The tool generates additional fixed points on the contour that causes subsequent changes to contours to be more complex. The live wiring tool calculates differences of intensities and thus produces a smooth curve of the selected area. In areas where there is no sufficient gradient due to a lack of border contours, such as segmenting separate muscles of a muscle group, the Livewire tool had additional polygon functions (direct line connection between two points) and a free hand option (wired in direct relation to the motion of the mouse).

However, some major drawbacks were observed during evaluation. With a growing number of

10. Currently, still closed source, open source soon, the source code was available to the author.

11. <http://www.turtleseg.org/>



**Figure 6.8** – TurtleSeg presents a very focused user interface. Next to the visualisation of the image data are only a limited number of controls that are required for the Livewire control.

contours, the time for processing increased too fast, resulting in up to 30 minutes for a single muscle. Especially edits in later stages become very time-consuming and often a restart from scratch takes less time. During the loading of complex data sets, the software froze and no cancelling was possible. A large number of contours also increased the probability of crashes. Finally, the spotlight functionality did not propose helpful slices for the segmentation of long tubular structures, like muscles, and might be more useful for internal organs.

### Seg3D

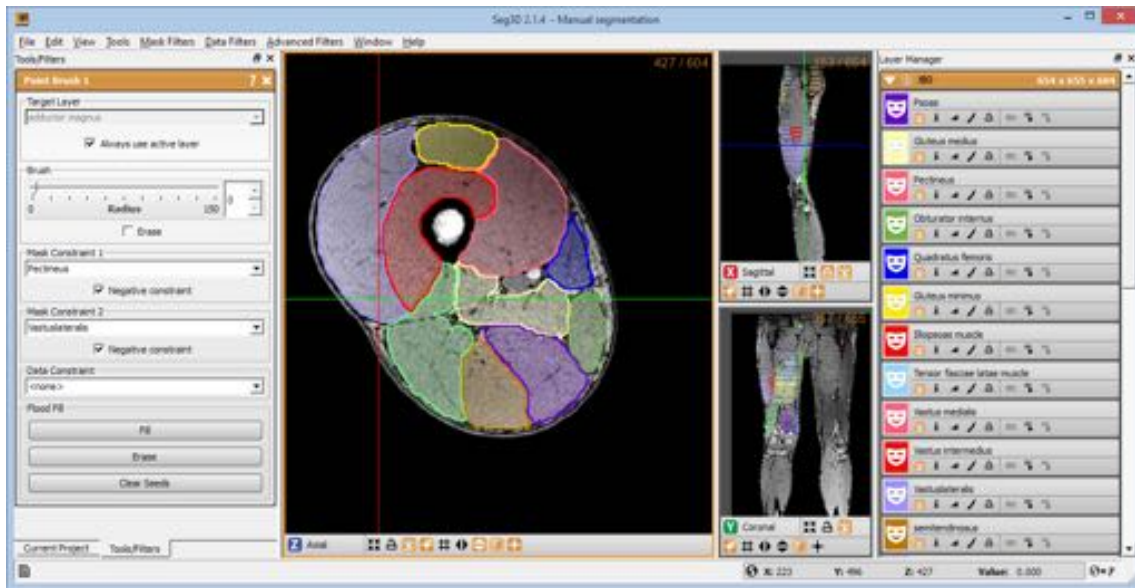
Seg3D<sup>12</sup> is a free volume segmentation and processing tool developed by the NIH<sup>13</sup> Center for Integrative Biomedical Computing at the University of Utah Scientific Computing and Imaging (SCI) Institute, USA [CIB13]. It provides a flexible interface for medical image data processing, including filtering and segmentation. Seg3D supports multiple biomedical image formats, runs in 64 bit and has a clean interface as shown in Figure 6.9. It is possible to show several projections of the image data in parallel. The interface always shows the list of all segments, allowing fast access and switching. Segments can be hidden or shown in a semi-transparent rendering. Also, the colour can be set. Finally, it is possible to load multiple images at the same time, e.g. different channels (fat and opposed-phase). This enables the user to switch between the channels to find

12. <http://www.sci.utah.edu/cibc-software/seg3d.html>

13. National Institute of Health, USA



the optimal contours for accurate processing.



**Figure 6.9** – User interface of Seg3D with a leg data set and several segments.

Polylines are the Seg3D equivalent to the Live Wiring found in TurtleSeg. However, they are restricted to straight lines between two selected points, resulting in rather rigid appearances in the 3D reconstructions. This collides with the continuous and round nature of most anatomical structures and results in small segmentation errors.

A beneficial function of the polyline tool is that it remains in the data set at the selected area when scrolling through the image slices which allows a faster processing of slice by slice. Only minor adjustments are needed in neighbouring slices; this reduces the time per slice in comparison to TurtleSeg, where the polyline was not transferred between slices. Sometimes the polyline function showed errors and misplaced fixed points on other polylines.

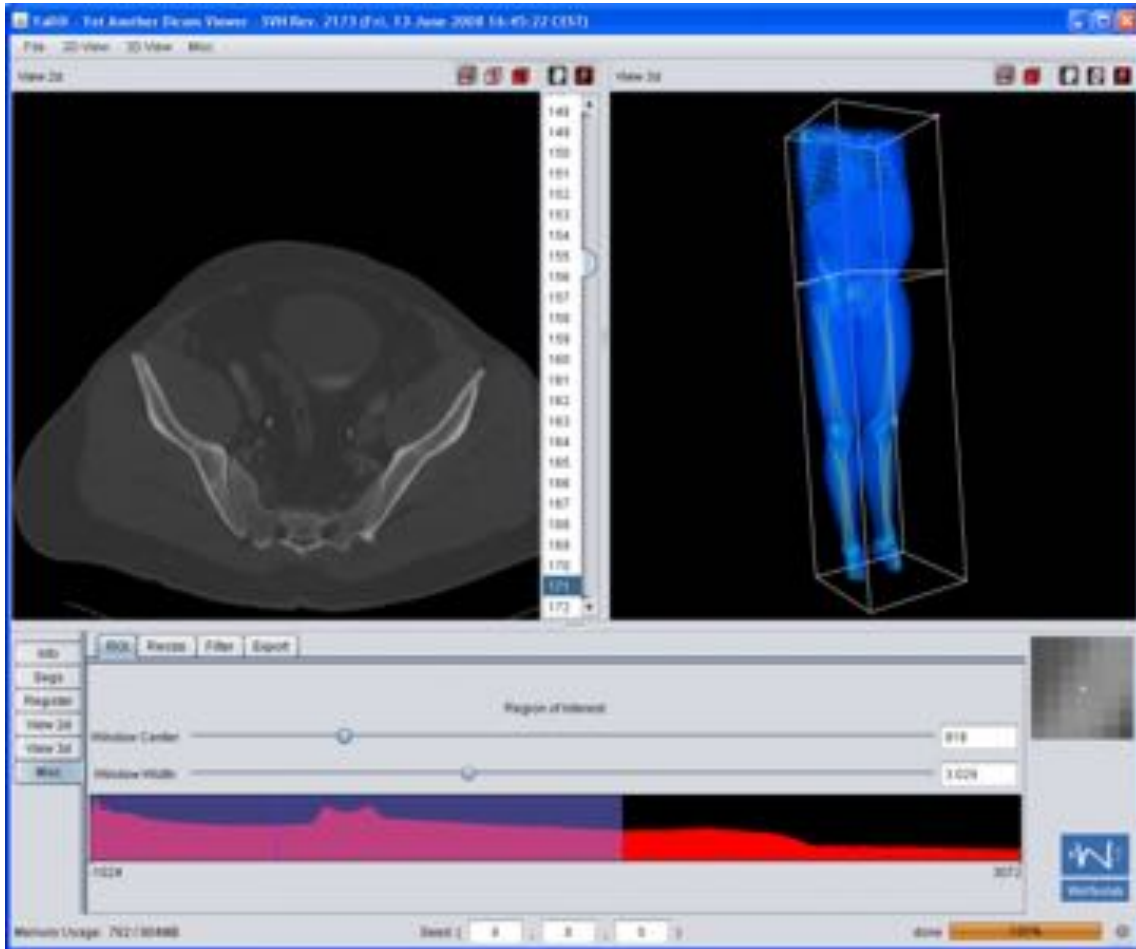
Unlike Seg3D, TurtleSeg supports interpolation between slices and does not require the processing for every single slice. Nonetheless, we have used external extrapolation techniques which we discuss in Section 6.3.1.2.

Seg3D provides a paint brush tool for manual segmentation. It allows making corrections to the results from the polylines and to segment finer structures with more details. Up to two segments can be selected as masks to prevent overlap between segments. However, with a large number of structures in the leg, a higher number would be more appropriate. The paint brush function works fairly simple and easy. To add an area to the segment, a user can use the left mouse button, the right button to remove it and the scroll wheel can be used to adjust the size of the paint brush. Mouse-based translation and zooming can be used through the shift button in paint brush mode.

## YaDiV



YaDiV<sup>14</sup>, short for *Yet Another Dicom Viewer*, is an open platform for 3D visualisation and segmentation of medical image data by Welfenlab, Leibniz University Hanover, Germany [FBW11]. It supports a small number of different medical image formats, multiple segmentation approaches (such as paint brush, range, region, level sets, and atlas), 2D and 3D visualisation, and segment visualisation. The user interface is clean, see Figure 6.10, but many important functions are hidden in sub-menus. YaDiV is considered a platform for academical research and development. It will be released later as open source. However, we had access to the source and had the possibility to extend it using its APIs.



**Figure 6.10** – The main user interface of YaDiV. The main functions can be accessed in the panel below the two main visualisations.

Dealing with a large number of segments, we have found the paint brush to be the most promising approach. Unlike the other tools, YaDiV offers two different brush shapes (square and circle), both in several sizes. Another interesting feature is the possibility to limit the paint brush to a certain intensity difference from the centre pixel of the brush. This makes it more robust and easier to use when segmenting along a steep gradient or edge and reducing the need for time-consuming, detailed corrections. During the segmentation, the list of segments is always accessible. It is also possible to perform operations on the segments, like adding or subtracting

14. <http://welfenlab.de/yadiv.html>

segments.

However, the lack of keyboard shortcuts made the process very time-consuming. Switching the view or slice using the keyboard could improve this. The biggest problems are the number of steps needed switch the brush size and the lack of a shortcut for switching the brush for adding to deleting mode. It would be possible to make these changes in the source, but overall, other tools seemed more optimised for the task.

#### 6.3.1.1 Conclusions on Manual Segmentation Tools


In conclusion, the presented tools all have their advantages and disadvantages. TurtleSeg presents a very helpful Livewire approach, but is not adapted to muscles and consumes an increasing calculation time. Also, the lack of stability is an issue. We are confident that this will be addressed in a future version. However, at the time of use, we had to decide against TurtleSeg. YaDiV is a flexible platform with many features. It allows a wide range of extensions and use histogram based limitations on the paint brush. The more complex than necessary user interaction and the resulting additionally needed time, forced us to exclude YaDiV. However, an inspection of the source code showed that the extensibility would allow an integration of features like keyboard shortcuts, but this is beyond the scope of this work. Finally, we have chosen Seg3D as the most efficient software tool during the evaluation. However, it would benefit from Livewire functionality. The lack of extrapolation and the long time required per slices, lead to the decision to skip several slices in areas with only little changes (middle of the thigh) and more slices near the joints where the anatomy is more complex.

#### 6.3.1.2 Post-processing of the Segmented Data


The evaluation data has been segmented manually. However, this is not free from error and requires some post-processing steps. First, we need to remove areas of accidental overlap in the data sets. Then, we need to fill the slices that were skipped to speed up the segmentation process. Finally, we have to convert the data into the required formats.

Sometimes, segments accidentally overlap after segmentation. This results from the manual work and cannot be fully avoided. However, during post-processing, we clean up this overlap by using the following approach: First, we identify the overlapping areas and remove them in both segments. We then, for each pixel in the boundary, determine the majority segment in the direct neighbourhood and assign the pixel correspondingly. We repeat this iteratively until all pixels are assigned to a segment. The whole process is shown in Figure 6.11.

We use the algorithm proposed Raya and Udapa [RU90] to fill the skipped slices. They use distance maps from cross-sectional boundaries and use them for interpolation. The algorithm allows reconstructing missing slices with high accuracy which can be seen in Figure 6.12. Boydev *et al.* [BPD<sup>+</sup>12] have integrated this algorithm in ITK. We use their implementation.



**Figure 6.11** – Our proposed method to clean up small overlaps in segments (a). First, the overlapping pixels are removed (b). Then, the boundary pixels are assigned to the majority of their neighbours (c). This is repeated iteratively until the hole is closed (d).



**Figure 6.12** – Interpolation of missing slices: Initial object (a), removed slices (b) and interpolated reconstruction (c). Source: Boydev *et al.* [BPD<sup>+</sup>12], CC-BY

### 6.3.1.3 Required Time

The manual segmentation process is very time-consuming. It depends on several factors:

- **Experience of the user** More experience allows a faster identification of anatomical structures.
- **Number of slices to be segmented** Higher resolution data sets with more slices require more time.
- **Number muscles to be segmented** If only few muscles need to be segmented, the overall time can be reduced.
- **Complexity of the anatomy** The anatomy varies in complexity, especially the higher complexity near the joints leads to a longer segmentation time.
- **Resolution of the data set** With higher resolutions, more details become visible that can help with identification of structures. However, it also results in more image elements that need to be labelled.
- **One or both legs needed** Based on the application, it might be sufficient to segment only one leg which can reduce the overall needed time greatly.
- **Development of fatigue in long segmentation session** The segmentation process requires a large amount of attention causing fatigue in large data sets. This either leads to longer segmentation time per slice or a higher rate of mistakes.

Based on these factors, we have experienced a segmentation time per slice and leg of three to five minutes. Given an average sized subject, we have around 550 slices. If we focus on the thigh, this results in 275 slices. We can segment only one out five slices. This gives a total of 55 slices. For 3, 4, and 5 minutes per slice the total segmentation time would be 165, 220, and 275 minutes. These approximations are higher in the initial data sets. They double when both thighs are needed. The


total number of slices for manual segmentation is actually higher due to the complex anatomy near the joints (hip, knee).

Barnoiun *et al.* [BBBV<sup>+</sup>13] have evaluated the outcome of manual segmentation study on the quadriceps in MR images. They have found an excellent reliability and therefore propose their methods to be applicable for future work. Other approaches try to evaluate (smaller) segmentation and annotation results through crowd sourcing, like Maier-Hein *et al.* [MHKR<sup>+</sup>15] or propose evaluation methods that work without ground truth, like Kohlberger *et al.* [KSA<sup>+</sup>12].

### 6.3.2 Label-based Evaluation Metrics

The evaluation of segmentation results is done by comparing it to a ground-truth, e.g. acquired by manual segmentation or through other established, verified segmentation tools. There exist many different approaches to compare different segments. Simply comparing the volume is not sufficient, all metrics should include the actual overlap. A common approach, the Jaccard index, uses the union overlap. In 1945, Sørensen and Dice [Dic45] had proposed their mean overlap coefficient: **Sørensen-Dice coefficient (SDC)**, which has been adopted as standard in medical image segmentation evaluation when label images are used. An SDC of 0 expresses disjunct segments, a value of 1 expresses identity. For two segments  $S$  and  $T$ , the formula is the following:

$$SDC(S, T) = 2 \frac{|S \cap T|}{|S| + |T|}$$



**Figure 6.13** – Union ( $\cup$ ), Intersection ( $\cap$ ) and Relative Complement ( $\setminus$ ) for two segments  $S$  and  $T$ .

The relation between the union ( $\cup$ ), intersection ( $\cap$ ) and relative complement ( $\setminus$ ) for the two segments  $S$  and  $T$  is shown in Figure 6.13. Other measures to quantify the error include the False Negative Error (FNE) and the False Positive Error (FPE):

$$FNE(S, T) = \frac{|T \setminus S|}{|T|} \quad FPE(S, T) = \frac{|S \setminus T|}{|S|}$$

We use the ITK implementation of these and additional metrics from Tustison *et al.* [TG09a].

### 6.3.3 Results

We use our DMFramework to evaluate the actual segmentation part by comparing the outcome of our segmentation approach to single-channel forces and a manually segmented ground-truth. The comparison is made using the Sørensen-Dice coefficient as overlap measure. We use the previously processed data sets and have configured our DMFramework with a small growing force, Laplacian smoothing, a gradient based image force and our muscle label force.

We selected one of the main groups of muscles in the thigh, the quadriceps, for this experiment. The quadriceps are composed of rectus femoris (RF), vastus lateralis (VL), vastus intermedius (VI), and vastus medialis (VM). First, we have segmented RF, VM, VI and VL using only a single channel (either  $I_I$ ,  $I_W$ ,  $I_F$ ,  $I_O$ , or  $I_T$ ). We present the results in Table 6.3. In both data sets, we can observe the highest results in the water image and we found the VL to be the hardest muscle to segment.

**Table 6.3** – Single channel results for RF, VM, VI, and VL in both subjects.

		In-phase	Water	Fat	Opposed-phase	T1/TSE
Subject 1	RF	0.78	0.85	0.72	0.79	-
	VM	0.78	0.81	0.76	0.79	-
	VI	0.76	0.78	0.72	0.77	-
	VL	0.68	0.74	0.71	0.71	-
Subject 2	RF	0.76	0.81	0.71	0.76	0.74
	VM	0.79	0.83	0.78	0.81	0.79
	VI	0.75	0.79	0.72	0.75	0.72
	VL	0.70	0.71	0.71	0.72	0.72

To test the opportunities of multi-channel images, we have repeated the segmentation of RF, VM, VI, and VL in our two subjects. This time, we have configured multiple image forces per muscle that operate on different channels. Together, they should combine information from the different channels. We have tested several combinations; the results can be found in Table 6.4. They show that  $I_F$  and  $I_W$  are important contributors to the results. We can also see that more channels lead to better results. However, the single water channel has shown better results for Subject 1 than a combination of channels.

**Table 6.4** – Multi-channel combination results for RF, VM, VI, and VL in both subjects.

		RF	VM	VI	VL
Subject 1	$I_F + I_W$	0.74	0.80	0.77	0.74
	$I_F + I_W + I_O$	0.83	0.80	0.76	0.74
	$I_F + I_I + I_O$	0.81	0.78	0.77	0.71
	$I_F + I_W + I_I + I_O$	0.82	0.79	0.78	0.72
Subject 2	$I_F + I_W$	0.71	0.77	0.76	0.72
	$I_F + I_W + I_O$	0.84	0.82	0.76	0.72
	$I_F + I_I + I_O$	0.81	0.81	0.75	0.72
	$I_F + I_W + I_I + I_O$	0.84	0.81	0.75	0.75

This experiment compares the use of single- and different-multi-channel configurations. We have found that the use of multi-channel image data can lead to better results. The selection of channels and the combination and weighting plays an important role. Individual channels can give misleading input. The fat channels, for example, show the inter-muscular fat which can be used to detect the muscle boundaries. However, depending on the amount of physical activity and age in the subject, larger amounts of intra-muscular fat might be present. This can indicate false boundaries and can be compensated by the additional use of other channels.

In general, we have found that the results improve with an increase in channel number. However, additional channels increase the force calculation times (mostly through memory access costs). The fat and opposed-phase channels contribute to good results whereas the use of the in-phase channels shows no or only small improvements. This is contrary to the observation of the single channel segmentation where the in-phase channel has created better results than the fat channel. In comparison to the related works, we found that models without ASMs like ours do not reach the accuracy of models that employ prior knowledge like statistical shape information or intensity profiles. With a large data base of subjects, future work should study the impact of multi-channel image data on these statistical approaches.

## 6.4 Webviewer for Remote and Multiplatform Viewing and Interaction

The direct visualisation of the deformable model segmentation is important during development and adoption of the segmentation process as discussed in Section 6.2.2.5. The proposed interface for output allows the easy integration of new visualisation approaches and extends our DMFramework with flexibility for visualisation and output.


### 6.4.1 Webviewer

We propose a web-based viewer for image data and streamed models to demonstrate the flexibility of our DMFramework. It can be seen in Figure 6.14. The viewer is fully implemented in ECMA Script and HTML5. It leverages from the Canvas element [Wha11] for the 2D rendering of image slices and uses WebGL (Khronos group, [Mar11]) and the THREE.js<sup>15</sup> library for 3D rendering. We use WebSockets (RFC 6455, [SA11]) for the communication with DMFramework to retrieve information about the environment (loaded data sets, models being processed). The main functionality lies in the transfer and rendering of large image data sets, windowing of image intensities, and the overlaying of additional information (e.g. meshes and cuts). The interface provides access to the data sets of our five subjects and allows the easy switching between them as well as the selection of different channels. Haehn *et al.* [HRA<sup>+</sup>14] have proposed the X ToolKit which provides image loading and rendering functionality. However, their proposed implementation seems targeted at smaller data sets and becomes unusably slow with our high-resolution multi-channel data sets. Therefore, we have optimised our implementation for large data sets. Other works target browsing large databases, like Shen *et al.* [SMZ<sup>+</sup>14]. They propose a method transfer image data in tiles. However, this requires a complex server setup.

Using ECMA Script and HTML allows us to target different platforms (such different OSes, and mobile platforms) and the easy adaptation of the user interface depending on the use case. Many new output devices (Hololens, Oculus) also support the displaying of websites and offer libraries for easy integration into the rendering process. The Webviewer also allows the simple

---

15. <http://threejs.org/>



**Figure 6.14** – Our proposed Webviewer. On the left, three different subjects and their data sets can be selected (a). Next to it, a 2D rendering of slices and models (b) is provided to control the 3D rendering on the right side (c).

integration of new interaction concepts, for example LeapMotion (hand tracking) and Myo (EMG arm tracking) provide ECMA Script access to the device data.

We use the canvas element for rendering (2D) slices of our image data. To achieve better performance, we use an off-screen canvas to capture the original slice with applied windowing information. Extraction and windowing from the image stack is a fast background task. The visible area of the slice can be controlled by the mouse and can be moved and zoomed. We store the transformation of the slice and apply it when we copy from the off-screen canvas to the visible canvas. Afterwards, we render additional annotations like markers, model cuts or text. WebGL and THREE.js provide the foundation for the 3D rendering of the scene. We render the visible slices using the off-screen (2D) canvases as textures. We use white lighting, show the bounding box of the image data and apply a trackball controller for navigation. We use triangle meshes to display the models from the DMFramework. They are supplied using the streaming technique described in the next section.

## 6.4.2 Model Streaming

While the image data is static and can be windowed locally without server contribution, the models are constantly changing throughout the simulation process. Prior work on this field has been proposed by Tang *et al.* [TRGP10] who have demonstrated a collaborative streaming for 3D shape deformations. Han *et al.* [HNS<sup>+</sup>10] and Nijdam *et al.* [NKHMT12] have demonstrated a client-server 3D data streaming system and have demonstrated it on a medical use case. Both approaches require the integration of their specialised software libraries. We propose a transfer built on standards like JSON and WebSockets. The feasibility of WebSockets for streaming has been demonstrated by Wessels *et al.* [WPJR11] although their work is more focussed on the streaming of server-side rendered image data.

The process of geometry streaming falls into two parts. In case of our DMFramework, the geometry of models currently does not change. Therefore, we split the streaming into a one-time transmission of the geometry (faces and initial vertex information) and a continuous stream of updated vertices. This reduces the required bandwidth, since the face information has to be transmitted only once.

Depending on the number of models and their resolution, the required bandwidth can be large and may saturate the network connection. Within our DMFramework the vertices are represented by double precision floating point number for a maximum accuracy. This is not required for other applications; many visualisation frameworks only use single precision floating point numbers without any visible reductions. In our application, the dimension of the working space is defined by the image data set. While we could place models with sub-voxel accuracy, this is not required for their visualisation. Based on the dimensions of our image data, a vertex-component could be reduced to 10 bits. However, for convenience, we choose the representation as short integer (2 bytes). The number of bytes to represent the vertex has a large influence on the required bandwidth. In Table 6.5, we present the data for a single model made of 1024 vertices with five updates per second. In real applications, the number of models is higher and the number of vertices also is often higher, whereas our tests have shown that fewer than five updates per second still can satisfy the user requirements, two to three updates seem sufficient.


**Table 6.5** – Data types and required bandwidth for different scenarios for geometry streaming of mesh with 1024 vertices and five updates per second.

Use case	Data type	Data type size	Total size	Data/s	Bandwidth/s
simulation	double	8 byte	24kB	120kB	960 kBit
visualisation	float	4 byte	12kB	60kB	480 kBit
image grid	short	2 byte	8kB	40kB	240 kBit
economic	bits	10 bit	3.75kB	18.75kB	150 kBit

We have designed the communication protocol between client and server to focus on the task. Once the clients connect, they can request information on the existing meshes (name, size, assigned colour where available). Next, they can switch to the geometry channel (a WebSocket can offer multiple channels) to retrieve the geometry. Finally, they open the channel named after the models to receive the stream of vertices. A simplified version of this process is presented in Figure 6.15.

The current implementation shows opportunities for future work. An important question is if the streaming of vertices would benefit from compression or if the additional overhead and delay would reduce the user experience. Based on the user focus, models could be streamed with different levels of detail. For now, the image is loaded fully before being displayed. A more advanced implementation could use streaming techniques like IIP protocol proposed by Husz *et al.* [HBH<sup>+</sup>12]. For the interaction with the models, interaction techniques like the ones surveyed by Olabarriaga *et al.* [OS01] could be implemented.





**Figure 6.15** – Communication between the server and a client for the streaming of mesh data. First, the client requests information about the available meshes (info message). The client then request the geometry for the meshes of interest. After receiving the geometry, the client connects to the mesh channels to receive mesh updates once they occur on the server.

## 6.5 Multi-Resolution Mesh Coupling

Magnetic resonance imaging (MRI) is widely used and creates higher and higher resolutions that result in larger data sets which require specialised processing. A common technique is the pyramid approach [AA84]. Several levels of detail are created by sub-sampling the input. Using a factor of two, a 2D image would quadruple at each level. With these levels, processing can be performed from coarse to fine, getting the main parts of the structure correctly outlined before working on details. This technique is commonly used in image registration [AA84]. In this case,

*This section was previously published as:*

*M. Becker, N. Nijdam, and N. Magnenat-Thalmann, "Coupling strategies for multi-resolution deformable meshes: expanding the pyramid approach beyond its one-way nature," in Int. J. Comput. Assist. Radiol. Surg., vol. 11, no. 5, pp. 695-705, May 2016.*

the parameters are adjusted for each level.

We propose a deformable model segmentation approach that works on multiple levels of detail in parallel. The meshes for each layer are synchronised with each other, breaking the serial structure of the pyramid approach. This changes the interaction pattern. The serial approach requires a scale based correction - influence points and corrections have to be set (and the mesh evolved) before switching to the next level of detail. Using our approach, the user can work location based, e.g. when segmenting the femur, focus on the femoral head first and work on the condyles later. This way, the user can switch between multiple levels of detail and can go to a coarser model without losing fine details previously created. The different meshes are deformed independently of each other (except for some form of synchronisation) and can be distributed across multiple cores or machines. We are therefore able to exploit multi-core machines even with algorithms that are not optimised for this.

Bogovic *et al.* [BPB13] describe a framework that allows the parallel segmentation of multiple, coupled objects. They demonstrate the feasibility of their level set based approach, with brain structures. In contrast, Changizi *et al.* [CH10] show a probabilistic multi-shape presentation. Their approach handles multiple related meshes and can be used to improve statistical models.

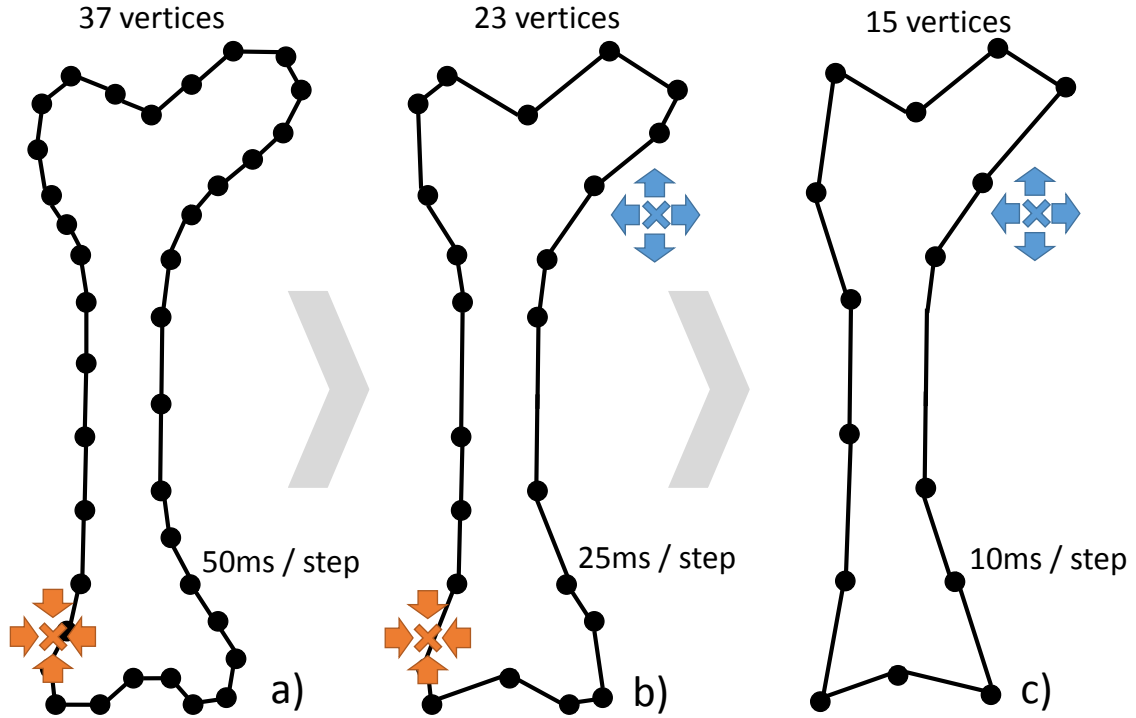
Edwards *et al.* [EWS<sup>+</sup>10] discuss the mesh consistency during parallel modifications. Lachat *et al.* [LDP14] present a method for the parallel re-meshing of distributed meshes. The work by Sumengen *et al.* [SEYB08] demonstrates a multi-resolution representation of an FEM model that can be deformed in a distributed environments. They subdivide the mesh and every client is simulating a part, the results are then shared with the other clients. On the client side it is then possible to generate representation at multiple resolutions. Also Tang *et al.* [TRGP10] demonstrate the distribution of mesh distributions through collaborative environments.

We have chosen deformable models for their flexibility and their well-established use in medical image segmentation. Furthermore, interaction methods have been defined for them. Some work has been done on parallel mesh access and distribution, and on the generation of multiple resolution representations. Nonetheless, an approach for coupled multi-resolution meshes that allows interaction is still missing.

### 6.5.1 Concept of Coupled Meshes

The coupling of multiple meshes is a complex process. We will discuss several approaches and their advantages and disadvantages. Since the meshes have different resolutions and are calculated independently from each other, some form of synchronisation has to be implemented. Lastly, we look at approaches for interactive segmentation and how they can be implemented with our multi-resolution approach.

For a given mesh  $M$ , its sub-versions for  $n$  levels are  $M_1, M_2, \dots, M_n$ . It is given that always  $M_1 \supseteq M_2 \supseteq \dots \supseteq M_n$ . A vertex  $V$  may be called  $V_i$  within the mesh  $M_i$ . This relation can be seen in Figure 6.16.



**Figure 6.16** – Concept of the coupled models. An initial model  $M_1$  (a) is decimated (b,c) to models  $M_2, M_3$  with fewer vertices. This reduces the calculation time per simulation step. Interaction is done through magnets and anti-magnets. Magnets (orange) are passed up to finer levels; anti-magnets are only propagated to more coarse levels.

### 6.5.2 Coupling

The meshes  $M_1, \dots, M_n$  have to be coupled so that changes in one mesh propagate to the others. In the following section, we discuss three different techniques: naive direct vertex copying, position averaging and finally spring forces between meshes.

#### — Direct Vertex Position Copying

A simple approach to couple the meshes is direct vertex copying. It can be done from coarse to fine or from fine to coarse. In either case, either  $M_1$  or  $M_n$  are favoured, since  $M_1 \supseteq \dots \supseteq M_n$ . Therefore any changes in the meshes  $M_2, \dots, M_{n-1}$  are discarded. This renders them useless and reduces the actual pyramid to two levels:  $M_1$  and  $M_n$ .

#### — Vertex Position Averaging

A more practical approach is to average the vertices across the models. The new average vertex  $V_{new}$  is calculated using weights  $\omega_i$ :

$$V_{new} = \frac{1}{n} \sum_{i=1}^n \omega_i V_i$$

with weights  $\omega_i$  so that

$$\sum_{i=1}^n \omega_i = 1$$

Using this approach, the position of the new vertex is the same for all meshes. The weighting can be used to increase the impact of certain meshes, e.g. higher resolutions. It is possible to adjust the weight dynamically, starting with high weight for the coarse model at the beginning and transfer the weight continuously to the finer models. Another possibility is to give the user control over the weights, either directly or through a generalised interface. In some applications, it can be useful to set the weights depending on the region. The femoral shaft, for example, has few details of importance so the coarse mesh could be favoured. The femoral head, on the other hand, has many detailed structures of high importance so that the finer meshes would be given more weight.

#### — Spring Force Coupling


The last approach we propose, uses the (weighted) average of the position of the vertices on the different meshes as described above and places a force between the actual vertices and the average position. This Hookean spring force is defined as follow:

$$F_{spring}(V_j) = \frac{1}{n-1} \sum_{i=1, i \neq j}^n \omega_i (V_j - V_i)$$

Also, with this approach weights ( $\omega_i$ ) can be used.

### 6.5.3 Synchronisation

The simulation of the deformation runs at different speeds, depending on the mesh resolution. Therefore, there must be a synchronisation between different models. It must be avoided that meshes are modified while forces are being calculated on their basis. We have developed different strategies, as shown in Figure 6.17, which we will present here. All are based on a central force (the sync force), that handles the synchronisation. After each simulation step, the models provide a new copy of their changed vertices. The main segmentation only has to be extended by that call. During the calculation of the forces, the sync force is queried. This call will be delayed internally until all models are in sync. The actual segmentation thread, therefore, is not aware of the other threads.



**Figure 6.17** – Comparison of the different synchronisation approaches with three models over time. The PSU strategy (a) causes a high amount of waiting time (grey) for the lower resolution models for synchronisation. The SU strategy (b), when carefully tuned, reduces these waiting times. The PU strategy (c) performs updates (arrows) in an unsynchronised manner, fully removing waiting times.

— **Per Step Updates Strategy**

The per step update strategy (PSU) ensures that all threads are always at the same iteration step. This is achieved by blocking the threads at the beginning of the force calculation until all of them have updated their meshes.

— **Scaled Updates Strategy**

The scaled update strategy (SU) is based on a per-thread cost estimation. While in the PSU strategy all models were synchronised with each other after every step, this strategy defines the number of steps per thread depending on the complexity of the model. Using a good estimation, waiting times can be reduced in comparison to the PSU strategy. A possible configuration for three threads would be the following: The first thread with a full resolution model is allowed to do one step, the second thread with a half-resolution model two steps, and the third with one third-resolution model can do three steps. In cases where the calculation scales linearly with the resolution, that configuration reduces waiting times.

— **Partial Updates Strategy**

The partial updates strategy (PU) largely reduces the synchronisation overhead. Model updates are stored immediately, and forces are calculated directly upon request. Locking is only used to ensure that no forces are being calculated while the meshes are updated. This approach allows the best performance but, depending on the thread scheduling, models can run apart massively.

#### 6.5.4 Interactive Segmentation

There exist many styles of interaction with deforming meshes. After careful evaluation, we have selected to implement springs and volcanos, which have been presented in Section 2.9.5.3. They are also known as magnets and anti-magnets. Magnets are designed to attract the mesh; anti-magnets are supposed to push the mesh away.

A magnet at the position  $P$  only interacts with the closest vertex  $V_i$  of the mesh. Using the scaling factor  $\alpha_{magnet}$ , they are defined as:

$$F_{magnet}(V_i) = \alpha_{magnet}(P - V_i)$$

Anti-magnets are designed to affect only near proximity, given by a radius. A vertex  $V_i$  inside this radius, having the distance  $r$  to the anti-magnet, receives the following force, scaled by  $\alpha_{antimagnet}$ :

$$F_{antimagnet}(V_i) = \alpha_{antimagnet} \frac{r}{\|r\|^3}$$

Magnets and anti-magnets are set/defined by the user on a specific level. Nonetheless, they should also influence the meshes at other resolutions. To reduce side-effects and unwanted behaviour, we have evaluated the semantics of the operations and defined the following rules:

1. Magnets are only passed down to more detailed levels. Magnets show the closest vertex

where to move to. While this is desirable at all levels, on more coarse levels, they will attach to a different vertex and can cause major, unwanted corrections.

2. Opposite to this, anti-magnets are only passed up to more coarse levels. When using anti-magnets, the user usually wants to correct a segmentation mistake and to push the model away from certain structures. This is a larger operation that works more efficiently on a coarse mesh. Nevertheless, the more detailed models also benefit from this, since they are coupled to the coarse models.

### 6.5.5 Implementation

The task of running multiple meshes in parallel and coupling their vertices requires careful synchronisation. It can be divided into two parts, the actual synchronisation using primitives like mutexes and the integration into the scripts.

**Listing 6.2** – Modified main loop for multi-threaded deformable model calculation.

```
modelSyncForce:registerThread ( threadID , model , indices )
for i=1,1000 do
  if i%60*2 == 0 then
    print ( threadID , " : " , ( i / 60 ) * 2 ,
              " seconds simulated " )
  end
  modelMan:CalcForces ( )
  modelSyncForce:calcForces
    ( threadID , femur )
  sim:performSimulationStep ( 0.016666 )
  modelSyncForce:updateThread
    ( threadID , femur )
  visu:addModel ( femur , threadID )
end
```

We have modified the main loop of the deformable model simulation to include calls to our force that performs the coupling (modelSyncForce). This force is implemented as a singleton and is accessible from all threads. Every thread initially has to register to this force, as seen on this first line of the source below. Threads register with their ID, their model and the vertex index mapping they use. The different resolutions of the meshes are created by decimation of the initial high-resolution mesh. All meshes share some vertices. To identify them, all vertices are mapped to the indices of the initial model. Since the modelSyncForce is calculated across threads, it is called separately from the model manager. Its parameters are the threadID and the model. After a simulation step is completed, the new vertices model are updated in the modelSyncForce. Depending on the synchronisation strategy, the execution can be interrupted in either the calculation or the update phase.

We have carefully selected the components of our single-threaded version framework to support different platforms, including Microsoft Windows, Apple OS X and Linux. For the multi-threaded version, we have selected boost's threads and synchronisation primitives because it allowed us

to keep the cross-platform compatibility. In the next section, we will describe how we have used mutexes to implement the three synchronisation strategies described in Section 6.5.3.

#### 6.5.5.1 Per Step Update Strategy (PSU)

The per step strategy requires synchronisation at two points: The forces can only be calculated once all models have been updated, and the updates can only be done once all forces have been calculated. The corresponding sample code is shown in Listing 6.3.

**Listing 6.3** – Usage of mutexes for locking in the per step update strategy.

```
mutex canUpdate;
mutex canCalculate;
map<Key, bool> threadUpdated;
map<Key, bool> threadCalculated;

void update(threadID, model) {
    while (!canUpdate.try_lock()) {
        yield();
    }
    canUpdate.unlock();
    threadCalculated[threadID] = false;
    //copy vertices here
    threadUpdated[threadID] = true;
    if (all_updated) {
        canUpdate.lock();
        canCalculate.unlock();
    }
}

void calculate(threadID, model) {
    while (!canCalculate.try_lock()) {
        yield();
    }
    canCalculate.unlock();
    threadUpdated[threadID] = false;
    //calculate force here
    threadCalculated[threadID] = true;
    if (all_calculated) {
        canCalculate.lock();
        canUpdate.unlock();
    }
}
```

#### 6.5.5.2 Scaled Updates Strategy (SU)

The scaled updates strategy resembles the per step update strategy, but some threads are entitled to more steps before sync (between sync it follows the partial update strategy). It is important that the last update is not called before all threads have reached it. In our implementation, we use a mutex per thread to block the (last) update until all threads have performed the required number of iterations. The sample code is presented in Listing 6.4.

**Listing 6.4** – Usage of mutexes for locking in the scaled update strategy.

```

map<Key,int> scalingFactor;
map<Key,int> threadCalculated;
map<Key,mutex> threadUpdated;

void updateThread(threadID, model) {
    if ((threadCalculated[threadID] >= (scaling_factor[threadID]))) {
        while (!updateLocks[threadID]->try_lock()) {
            yield();
        }
    }
    //copy vertices here
}

void calcForces(Key threadID) {
    //calculate force here
    threadCalculated[threadID]++;
    if (all_calculated) {
        //reset update counts
        for (Key k: threadCalculated) {
            threadCalculated[k] = 0;
            updateLocks[k]->unlock();
        }
    }
}

```

### 6.5.5.3 Partial Update Strategy (PU)

The Partial Update Strategy provides locking for concurrent access to the vertices, but it does not perform any synchronisation. If threads are stalled, it is possible that the other threads continue to run and scatter. The basic locking of access to the vertices is shown in Listing 6.5.

**Listing 6.5** – Usage of mutexes for locking in the partial update strategy.

```

mutex vertexLock;

void update(threadID, model) {
    vertexLock.lock();
    //copy vertices here
    vertexLock.unlock();
}

void calculate(threadID, model) {
    vertexLock.lock();
    //calculate force here
    vertexLock.unlock();
}

```



### 6.5.6 Evaluation and Results

We have tested our approach in different scenarios. We describe the data sets and medical field used for evaluation and evaluate the effect of the synchronisation overhead. The goal is to segment the femur in our data sets. We have evaluated three femur models, from low to high resolution. The low-resolution model had 344, the medium model had 1850 and the high-resolution model had 37010 vertices. Having these large variations in size, allows us to estimate the share of the time used for locking as well as the overall robustness.

#### 6.5.6.1 Deformable Model Configuration

We have used our deformable framework as described in Section 6.5.5. It has been configured with multiple forces: Laplacian smoothing and a constant growing force. We exploit the image data by using a gradient-based image force. For the coupling force, we have chosen the spring force, since it fits our requirements best.

Starting with the selected (high, medium, and low-resolution) model, we have created reduced versions of them. With  $n$  models, model  $M_i, i = 1..n$  only has  $1/i$  times the vertices of the initial model. We have used the built-in decimators of OpenMesh, in this case with collapse priority based on error quadrics. The effect of the decimation at different levels on our femur model are shown in Figure 6.18.



**Figure 6.18** – The effect of the decimation on the model of the femur. The original model (left) is reduced to fewer vertices while preserving shape features like the femoral head and the condyles.

#### 6.5.6.2 Synchronisation Evaluation

The synchronisation is the bottleneck of the whole calculation. The independent segmentation processes have to be coordinated following a specific strategy like the ones described in Section 6.5.5. In this section, we evaluate the three approaches on the impact of locking. We have used our femur model with three different resolutions and have run our framework with six threads per model. During the segmentation, we have measured the average times for updating the mesh

and calculating the forces. Furthermore, we have measured the time that was spent waiting for locks in these parts. The results are listed in Table 6.6 for the low-, medium- and high-resolution model.

The first result we found is that the partial update strategy always is the fastest. In the calculation time, it can be seen that the time linearly grows with model size (as indicated by the different threads). For the per step update strategy, the first thread has the shortest lock time. This can be explained by the fact that it requires the longest calculation and, therefore, the other threads already have reached the synchronisation point. With growing model size, the differences between the strategies become more visible. While, in the low-resolution model, the difference between PU, SU and PSU is in the area of ms, it grows at medium resolution. At high resolution, PU takes 7-11 ms, PSU between 30 and 115ms, and SU between 29 and 70 seconds. The locking time per thread increases with the resolution for PSU strategy. It can be seen that PU causes the least time waiting for locks.

#### 6.5.6.3 Discussion

The total calculation duration is sensitive to the model resolution. This can be seen from the timing data independent of the synchronisation strategy. The overall computation time (ignoring the locks) goes down with the resolution as expected.

Looking at the strategies, we found the partial update strategy to be the fastest. This is achieved by reducing the dependencies between the threads. The only locking is around updating and force calculation for memory safety, but there are no global synchronisation approaches. This is supported by the data that shows that the PU strategy has the shortest lock times. For the lower resolution threads, we can see a higher share of locking time. Since these threads are simulated faster, they reach the locks earlier. Furthermore, the higher resolution threads take longer for updating and force calculation and, therefore, block the locks longer.

The PSU strategy shows that the thread with the highest resolution spends the shortest time waiting for locks. Since this thread, by pure resolution, always takes the longest time to simulate, all other threads are waiting already when the locking point is reached.

The SU strategy is very sensitive to the proper selection of the scaling. We can see the expected timing results, all in the same region, for low- and medium-resolution model, but not for the high-resolution model. The large variation in values (19 s to 70 s) indicates that the scaling factors are not chosen well.

Overall, we found the PU strategy the most beneficial. It enabled a fast coupling with only very little synchronisation overhead. The behaviour follows user expectation and does not block interaction with the model through long locking times.

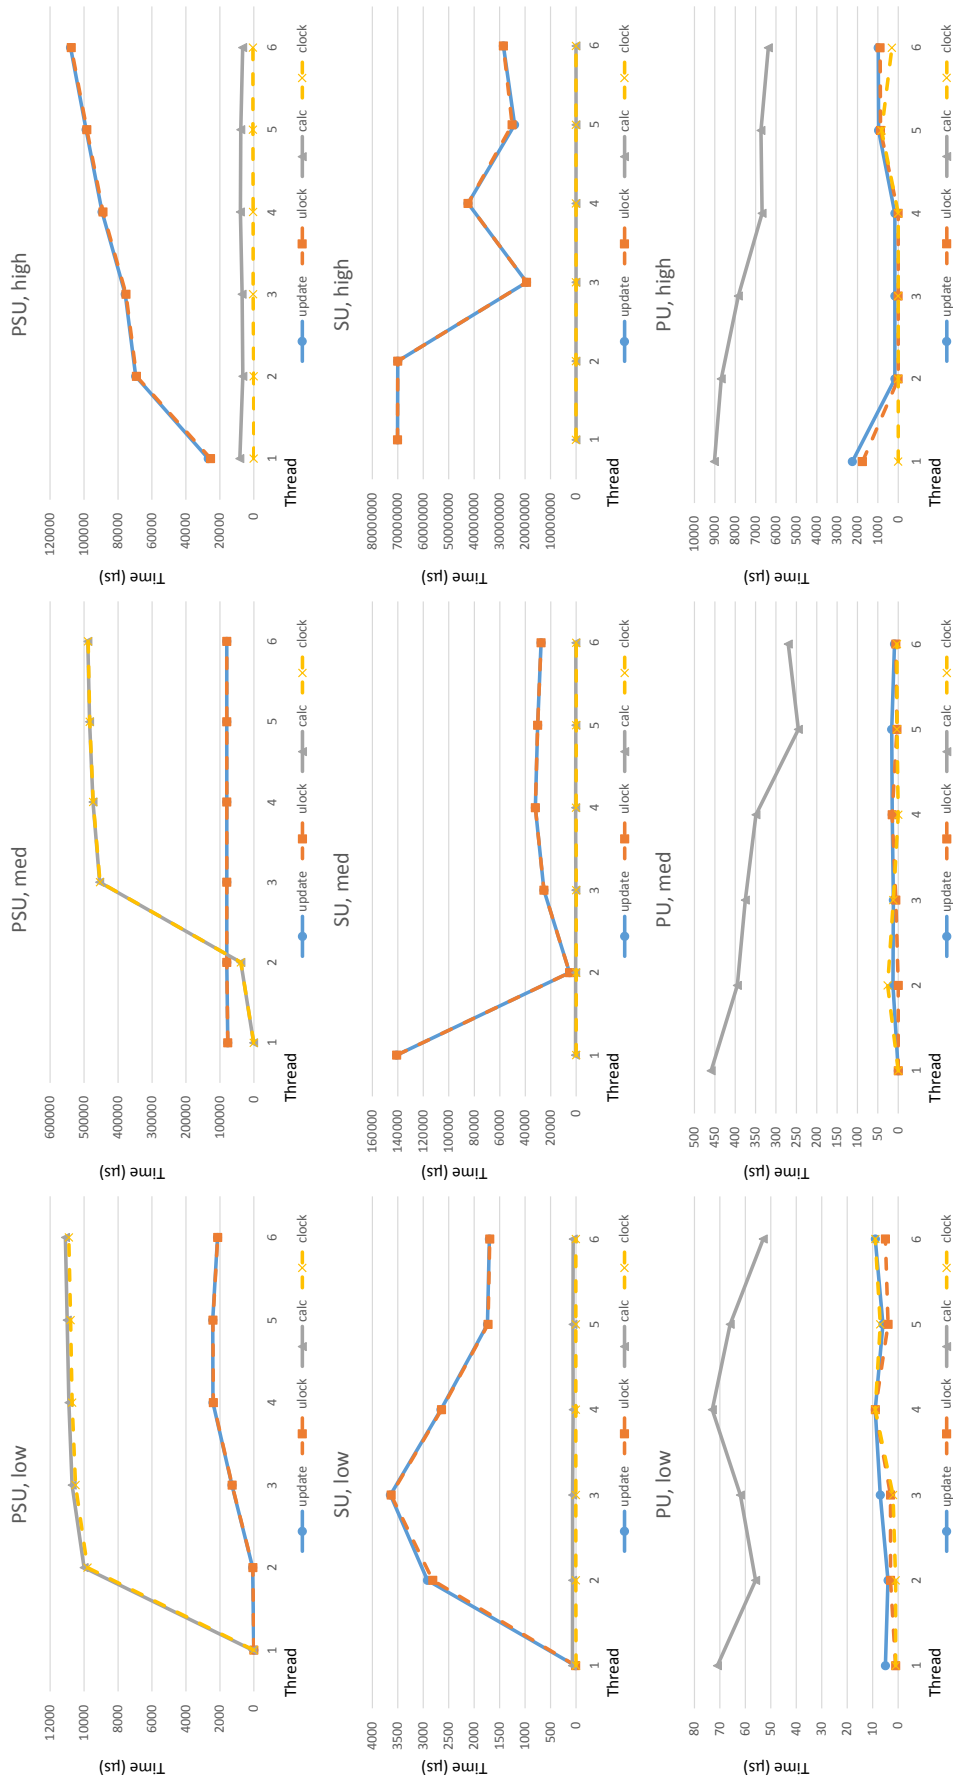


Figure 6.19 – Evaluation of the time spent on updating and calculating and the respective locks. Results are shown for low-, medium- and high-resolution models with the PSU, SU and PU strategy for a use case with six parallel threads.

Sync	Thread	low			med			high		
		update	lock	calc	update	lock	calc	update	lock	calc
PSU	1	5 $\mu$ s	0 $\mu$ s	72 $\mu$ s	76.9 ms	0 $\mu$ s	906 $\mu$ s	26.8 ms	25.4 ms	8216 $\mu$ s
PSU	2	65 $\mu$ s	62 $\mu$ s	10 ms	79.9 ms	9.8 ms	387 ms	69.5 ms	69.1 ms	6470 $\mu$ s
PSU	3	1287 $\mu$ s	1279 $\mu$ s	10.7 ms	79.7 ms	10.5 ms	454 ms	75.7 ms	75.2 ms	6849 $\mu$ s
PSU	4	2405 $\mu$ s	2391 $\mu$ s	10.9 ms	79.8 ms	10.7 ms	474 ms	89.3 ms	88.8 ms	7922 $\mu$ s
PSU	5	2421 $\mu$ s	2410 $\mu$ s	11 ms	79.9 ms	10.8 ms	484 ms	98.8 ms	98.3 ms	7714 $\mu$ s
PSU	6	2140 $\mu$ s	2129 $\mu$ s	11.1 ms	79.9 ms	10.9 ms	489 ms	107.9 ms	107.5 ms	6494 $\mu$ s
PU	1	5 $\mu$ s	1 $\mu$ s	71 $\mu$ s	0 $\mu$ s	1 $\mu$ s	459 $\mu$ s	2255 $\mu$ s	1754 $\mu$ s	9020 $\mu$ s
PU	2	4 $\mu$ s	3 $\mu$ s	56 $\mu$ s	13 $\mu$ s	1 $\mu$ s	395 $\mu$ s	8686 $\mu$ s	1503 $\mu$ s	167 $\mu$ s
PU	3	7 $\mu$ s	3 $\mu$ s	62 $\mu$ s	12 $\mu$ s	2 $\mu$ s	375 $\mu$ s	7851 $\mu$ s	1837 $\mu$ s	167 $\mu$ s
PU	4	9 $\mu$ s	9 $\mu$ s	73 $\mu$ s	15 $\mu$ s	9 $\mu$ s	349 $\mu$ s	6681 $\mu$ s	0 $\mu$ s	167 $\mu$ s
PU	5	6 $\mu$ s	4 $\mu$ s	66 $\mu$ s	16 $\mu$ s	7 $\mu$ s	245 $\mu$ s	6734 $\mu$ s	818 $\mu$ s	963 $\mu$ s
PU	6	9 $\mu$ s	5 $\mu$ s	53 $\mu$ s	9 $\mu$ s	9 $\mu$ s	270 $\mu$ s	6376 $\mu$ s	310 $\mu$ s	987 $\mu$ s
SU	1	15 $\mu$ s	7 $\mu$ s	70 $\mu$ s	141 $\mu$ s	0 $\mu$ s	500 $\mu$ s	70.1 s	70.1 s	10.1 ms
SU	2	2918 $\mu$ s	2813 $\mu$ s	66 $\mu$ s	5.2 ms	3 $\mu$ s	334 $\mu$ s	70.0 s	70.0 s	8.2 ms
SU	3	3646 $\mu$ s	3633 $\mu$ s	69 $\mu$ s	25.3 ms	4 $\mu$ s	324 $\mu$ s	19.4 s	19.4 s	6.1 ms
SU	4	2655 $\mu$ s	2644 $\mu$ s	58 $\mu$ s	32.1 ms	3 $\mu$ s	351 $\mu$ s	42.4 s	42.4 s	7.2 ms
SU	5	1739 $\mu$ s	1730 $\mu$ s	59 $\mu$ s	30.2 ms	4 $\mu$ s	268 $\mu$ s	24.1 s	25.1 s	6.1 ms
SU	6	1701 $\mu$ s	1694 $\mu$ s	56 $\mu$ s	27.6 ms	3 $\mu$ s	280 $\mu$ s	28.5 s	28.5 s	6.8 ms

**Table 6.6** – Timing measurement of the low-, medium- and high-resolution model for the total duration of the update, the time waiting for the locks during the update, the duration of the calculation of the force, and the time waiting for locks during the force calculation.

## 6.6 Conclusion

The extraction of musculoskeletal structures from medical images is a complex task. Deformable models are a strong tool for medical image segmentation and parallelisation shows opportunities to increase the efficiency. In this chapter, we have described our proposed techniques and algorithms for these challenges.

We have presented our DMFramework, an implementation of a deformable model framework. It combines the main components of DMs: meshes, forces, images, simulation and output in a scriptable, layered architecture. The components have been selected from widely adopted, established open libraries and attention has been paid to the possibilities for parallelisation.

Visualisation is an important way to give the user feedback. In our DMFramework we have included a VTK-based rendering of the output. To gain more flexibility and for easier updating, we propose a web-based viewer for medical image data and our deformable models. The viewer has been designed to deal with our large data sets. We also propose to stream updates of the meshes and have presented our approach.

Finally, we have proposed a way to increase the efficiency during interactive segmentation by running multiple meshes of different resolutions. This resembles the pyramid approach. However, it is possible to switch between the different levels of the pyramid at any time. This allows more flexibility and to change to the coarse model without having to restart.

In summary, we have contributed methods for working with multi-channel MR image data. We have presented a deformable model framework that is optimised for the efficient handling of multi-channel data and provides special forces. We have proposed a web-based viewer for our data sets and the streaming of meshes that allows more flexibility. Finally, we have explored how the traditional pyramid approach can be expanded to coupled multi-resolution models.

---

## CHAPTER 7

## CONCLUSIONS

---



## 7.1 Motivation

In this work, we have presented techniques for the efficient extraction of musculoskeletal structures from multi-channel MR images. We have discussed the image acquisition and the following processing. We proposed a segmentation using deformable models and have introduced several approaches for the initialisation. Finally, we have evaluated the possibility to increase the efficiency by using coupled multi-resolution models.

In this section, we will first demonstrate two applications for the proposed work. We continue by presenting our contributions and their benefits. Finally, we will discuss the limitations of our contributions and give an overview of possible future work.

## 7.2 Applications

There are many different possible applications for the reconstructed anatomical models that result from this work. In this section, we will present two examples that demonstrate the benefit that they can provide. The first example shows how digital patients can be recreated; the second example presents a concrete application on how understanding the muscle shape and the corresponding action line can be used to calculate a personalised pennation angle.

### 7.2.1 Digital Patients

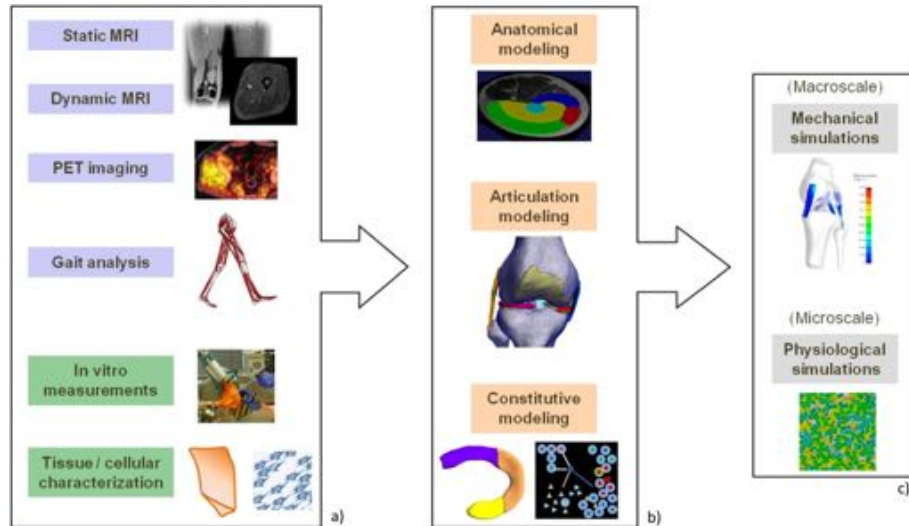
The number of image acquisitions is constantly growing and the resolutions are increasing. The resulting enormous amount of data requires automatic processing to be applicable to clinical decision making. Beyond a large variety of imaging modalities, many other data sources like gait data, and EMGs are acquired. For a more comprehensive integration of this data, the digital patient is a concept for the fusion of this data. This data can be acquired from several scales, from cellular to organ scale. Following the integration, the digital patient can then be used to simulate mechanical and physiological conditions under which articulations perform movements without putting the real patient in hazard. An overview of possible components from multiple sources and their integration can be found in Figure 7.1.

The full study can be found in:

H. F. Choi, A. Chincisan, M. Becker, and N. Magnenat-Thalmann, “Multimodal composition of the digital patient: a strategy for the knee articulation,” in *Vis. Comput.*, vol. 30, no. 6, pp. 1–11, May 2014.

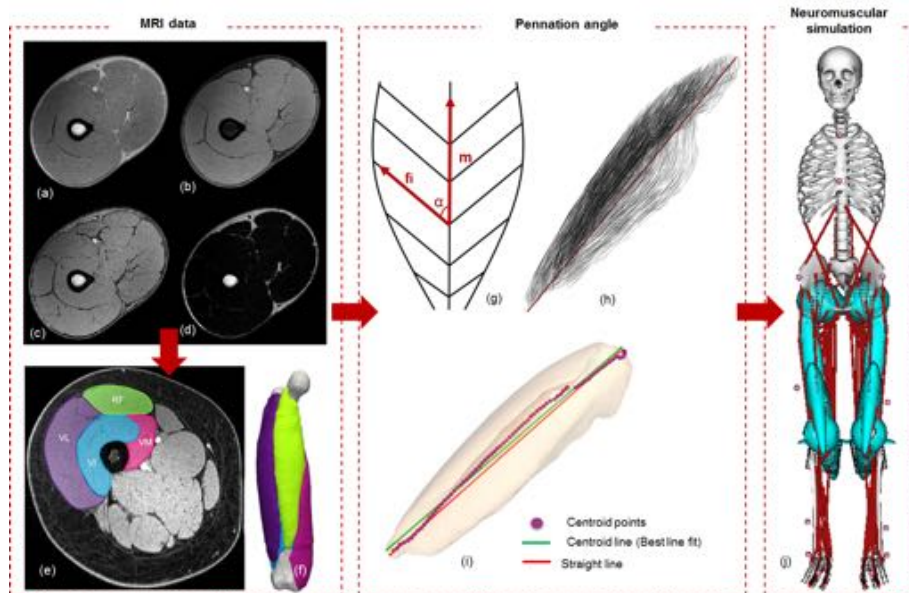
### 7.2.2 Pennation Angle Measurements

Gait analysis methods rely on the simulation of muscle forces. These simulations commonly use scaled generic models. The forces are influenced by the muscle line of action, the fibre length and the pennation angle (PA). The PA describes the angle between the fibre orientation and muscle line of action. It determines the maximum forces and the range of motion. Using our 3D muscle



**Figure 7.1** – Data sources (a), modelling approaches (b) and simulations (c) for the digital patient.

reconstruction, the fibre orientation and muscle line of action were calculated and the subject-specific PA was derived from them. It was then used for the neuro-muscular simulation, the different stages are shown in Figure 7.2. In an evaluation with three subjects, a variation of the PA between subject-specific and generic model was shown. This resulted in a significant difference in moments and muscle forces of the quadriceps, underlining the importance of subject-specific anatomy and more personalised approaches.



**Figure 7.2** – MRI data and our segmentation are used to create subject-specific models (a-f). These are used to calculate the subject-specific pennation angle (g-i) which improves the neuro-muscular simulation (j).

The full study can be found in:

A. Chincisan, K. Tecante, M. Becker, N. Magnenat-Thalmann, C. Hurschler, and H.F. Choi, "A Computational Approach to Calculate Personalized Pennation Angle based



on MRI: Effect on Gait Analysis,” in *Int. J. Comput. Assist. Radiol. Surg.*, vol. 11, no. 5, pp. 683-693, May 2016.

## 7.3 Discussion

In this thesis, we have addressed the extraction of musculoskeletal structures from MR images. MSDs are a common problem of the ageing population. To better understand these diseases, knowing the morphology of the involved anatomical structures can help to explain some of the origins. Medical imaging is used to acquire the data sets; we use multi-channel MRI. However, the required higher resolution and the total number of scans nowadays requires algorithmic support and processing to limit the time a radiologist needs for image analysis and interpretation. Therefore, we propose a system that combines image acquisition, processing and the segmentation of musculoskeletal structures. All these steps should be implemented efficiently. The scope of our work is the human leg. We can divide the work into three major objectives:

### 1. Development of an imaging protocol

Development of an MR image acquisition protocol to seamlessly cover the full leg. This includes the selection and parametrisation of MR sequences as well as the physical setup of the needed coils. In a last step, the acquired individual and partially overlapping images have to be combined into a continuous data set.

### 2. Processing and advanced tissue labelling

For the labelling, the main objective is the identification of different tissue types (bones, muscles, fat) and surrounding air.

### 3. Multi-channel image segmentation

The segmentation process has been realised with deformable models, which can carry knowledge about the structures. The focus lies on the segmentation of muscles in the human thigh. By using multiple inputs from the acquisition coming from different sequences, the multi-channel image data will be used for robust segmentation. These multi-channel inputs consist of mDixon and T1 weighted images.

All of these steps have been addressed while taking efficiency into account. This not only gives a faster process in general but is necessary with multi-channel data that is multiple times larger than single-channel data. Finally, an efficient implementation increases the applicability for clinical environments.

In the next sections, we will first discuss our contributions and then their limitations as well as future work.

## 7.4 Contributions and Benefits

In this section, we will discuss our contributions and their benefits. We will follow the main structure of our methods as outlined in this work.

### Multi-channel Image Acquisition

The imaging parameters are important to generate high-quality data sets that can be used for all the following steps. We have tested several sequences and propose to use mDixon and T1/TSE. These generate multi-channel data during a single acquisition and, therefore, offer more information. The total acquisition duration and the SNR are criteria for the acceptance of the protocol. We propose to use a body coil for signal amplification. It reduces the time needed and increases SNR. Since the coil does not cover the full limb, we have to move it. To avoid deformations and movements, we have developed a coil support that allows moving the coil independently from the subject.

### Image Processing Pipeline with Air and Muscle Labelling

The acquired image data requires processing to be usable for further segmentation steps. We have proposed a pipeline for processing. First, we have to stitch the individual stacks from the scanner. Next, we have proposed a technique based on morphological operations to identify air in the image data. This allows use to mask it during processing, resulting in shorter calculation times. It also allows cropping the data sets. This is even more important in multi-channel data like ours. Due to imaging constraints, our data sets are affected by the scanner bias field and carry noise. We use state of the art methods, to reduce these artefacts. Next, we have demonstrated the use of the air mask and the fat image to identify muscle tissue and have discussed the effects of intramuscular fat. These processing steps can be used for clinical applications, like direct feedback for radiologists during acquisition or as initialization for advanced segmentation techniques. Since powerful GPGPUs are usually installed for CT and MR reconstruction, we propose to use them for fast air and muscle labelling.

### Robust and Efficient Initialisation

For successful deformable model based segmentation, the initial models, e.g. from a template, have to be placed close to their final position. After evaluating several approaches, we have decided to contribute a novel three-level initialisation process. First, we use slice-wise histograms of air and muscle tissue. These histograms are aligned with a template. The characteristics of the muscle tissue distribution and total diameter of the leg from the air image, allow the identification of the slices of hip joint, knee joint, and ankle. Next, we use the muscle label and the in-phase image to locate the femoral shaft. With the location of the shaft and the hip and knee joint, the femur model can be placed. We then use our DMFramework for the segmentation. The process to register the muscle models is based on a generic template. First, we register the template femur with the segmented femur. We then apply the resulting transformation to the muscle models to place them around the segmented femur. Finally, we use ray based technique to ensure that models lie within the muscle labels. This is needed because of different muscle volumes between subjects, e.g. from different levels of activity.

### Multi-channel Image Segmentation

Based on our study of the related work, we have found deformable models to be a robust approach for image segmentation. We propose a deformable model framework called DMFramework. For its design, we have surveyed different approaches and have identified core compo-

nents. We rely on well-established frameworks, like ITK and VTK, for the components. This allows us to focus our work and to rely on the experience of the community. All components have been connected in a layered architecture. Special attention has been paid to ensure that all computationally intensive tasks can be parallelised and run on GPGPU systems through OpenCL support. The functionality of our DMFramework is exposed to Lua scripting. This allows us to quickly reconfigure the setup. We have added forces specialised on our use case. In particular, we exploit the muscle labels and the multiple channels of the mDixon images. Finally, we propose an web-based viewer for medical image data with geometry streaming for the visualisation of the deformable models.

### **Efficient Coupling of Multi-resolution Models**

The pyramid approach is commonly in large data sets. It allows to perform the calculation on reduced data sets that progressively increase in resolution up to the full resolution. We propose deformable models that operate on multiple levels in parallel. We have discussed approaches for the coupling through vertex propagation, averaging and springs. We have investigated the synchronisation of the parallel models and introduced three locking approaches. On the implementation side, we have shown how our DMFramework can be transformed into a multi-model framework.

## **7.5 Limitations and Future Work**

As discussed in the previous section, we have proposed several new contributions to efficient medical image processing. Like all work in progress, there are some limitations and opportunities for future work.

A challenge for future research is the development of more general approaches. Although very sophisticated and robust, deformable model applications are tuned to specific organs and well-defined image conditions. Knowledge-based models, for example, have to be trained with a large database of examples and image forces like the intensity profiles have to be recreated if imaging sequences are modified. Depending on the chosen approach, the initialisation can be a critical task. Speed and robustness still can be improved. Besides these general directions, we have identified the following limitations and possibilities for future work of our contributions.

### **Image Acquisition**

Our image acquisition protocol still requires a long time with respect to typical acquisition times in clinical environments. This is an obstacle towards integration into healthcare processes. This could be addressed by evaluating the need for the T1 image data. In our applications, it was mainly required for supporting the manual comparison; further work could be made to eliminate its need for the automated segmentation. This could significantly reduce acquisition times as the T1/TSE sequence requires five minutes per stack.

Our coil support system is currently only a prototype that is not ready for wide-spread use. Further work is needed to make it more durable and steady. It should also use faster locking systems for the traverse support. A series of preformed supports could be used to better accommodate

patient of different sizes.

### Image Data Processing

Currently, our stitching method can generate an error of half the slice thickness. With our high resolution data sets, this is acceptable. However, for lower resolutions it can be a problem. Careful consideration is needed to balance the loss of accuracy with the loss of quality from interpolating a full stack.

The air labelling has shown strong results. In subjects with thicker thighs, they sometimes are connected in the air labels. Future work could focus on removing this source of error. While the air allows cropping the data sets to reduce the overall file size, the bounding volume is determined by the dimension of the hip. However, with high memory pressure and large resolution, it could be beneficial to use a hierarchical storage tree, like an Octree, that allows the removal of air on a finer grid. This method would be beneficial below the hip, especially towards the ankles.

Our work on muscle labelling has been focused on the lower limb. It would be interesting to see the performance of our method on full-body mDixon scans. In particular, the area of the abdomen with its many different types of tissues could be a challenge. Furthermore, a larger study with a larger variation in subjects that allows statistical analysis, should be performed. For now, our method is specialised on mDixon images, since it exploits their multiple channels. A future research question would be the evaluation on how it can be adapted to other sequences that produce fat saturated images or other multi-channel image sequences. A way to improve the results would be to work on the segmentation and removal of veins and arteries. Further investigations of the joint area are possible to improve the local results. Our results currently identify if a tissue is a muscle or not. Future work could identify muscle groups.

Our GPGPU implementation has shown the possible gains in speed. We have found that CCL shows room for improvement and propose a joint solution which uses the CPU for CLL. Currently, the computationally expensive denoising has not been included. Future work should focus on the integration of the NLM algorithm; it has been shown to be possible by Cuomo *et al.* [CDP14] in 2D applications. Another topic would be porting the bias field estimation to a parallel approach.

### Model Initialisation

While the histogram based joint localisation has shown good results for hip and knee joint, the ankle has a lower accuracy. Future work should focus on optimisations for the ankle. Another question is the applicability of the proposed method to other limbs like the arms. They show similar anatomical compositions for muscles and a similar placement of bones and joints. Finally, the extension to the abdomen could be considered using the spine as reference, potentially with additional input from the ribs.

As discussed before, the muscle labelling is not fully accurate near the joints. Therefore, the muscle label based muscle mesh correction also shows some problems. Possible approaches would be an extension to tendons, e.g. visible with other sequences, or to rely on a small, local atlas.

**Multi-channel Image Segmentation**

The multi-channel image segmentation has to deal with the complex anatomy and little information from the image data. The accuracy of the results is limited, but could be improved through the use of Active Shape Models to support the muscle models in low information areas, e.g. boundary between vastus medialis and vastus intermedius. Another option would be to learn intensity profiles to exploit the image data beyond gradients. For these statistical approaches, more data sets would be needed. Future work should study the impact different resolutions and how different sequences could provide additional input.

The influence of the level of activities and development of the muscles as well the kind of sports performed could lead to different accuracies during the segmentation process. The amount of fat seems play an important role; little fat makes muscle boundaries less visible, too much fat accumulates inside the muscles, leading to false boundary signals. This requires further studying for confirmation and better understanding.

**Multi-Resolution Mesh Coupling**

For the future, calculation across multiple machines and collaborative interactions seem promising. To better understand the impact on users, a comprehensive user study should follow. For now, we have only tested our proposed method on a single machine. An interesting outlook would be to extend it to a multi-machine approach. Possible scenarios would be off-site vs on-site: A coarse simulation running locally allows smooth interaction whereas a more powerful remote server runs the fine simulation on the same data. Similarly, the coarse simulation could be run on a phone or tablet computer while the fine model is simulated in the hospital data centre or the cloud.

With model synchronisation already in place, another outlook would be collaborative interaction with the model. A distributed group of users could work in parallel on the model, if desired at different scales, and place interaction points. That way, additional knowledge could be included for critical parts of the segmentation process, e.g. in an image region with artefacts or unclear model boundaries.

---

## APPENDIX A

### LIST OF PUBLICATIONS

---

Most of the material presented in this thesis previously appeared in the following publications.

#### Peer-reviewed Journals

- **M. Becker**, and N. Magnenat-Thalmann, "Muscle Tissue Labeling of Human Lower Limb in Multi-Channel mDixon MR Imaging: Concepts and Applications," in *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, 2015, online. Impact Factor: 1.438
- **M. Becker**, N. Nijdam, and N. Magnenat-Thalmann, "Coupling strategies for multi-resolution deformable meshes: expanding the pyramid approach beyond its one-way nature," in *Int. J. Comput. Assist. Radiol. Surg.*, vol. 11, no. 5, pp. 695-705, May 2016. Impact Factor: 1.827
- A. Chincisan, K. Tecante, **M. Becker**, N. Magnenat-Thalmann, C. Hurschler, and H.F. Choi, "A Computational Approach to Calculate Personalized Pennation Angle based on MRI: Effect on Gait Analysis," in *Int. J. Comput. Assist. Radiol. Surg.*, vol. 11, no. 5, pp. 683-693, May 2016. Impact Factor: 1.827
- M. E. H. Wagner, N.-C. Gellrich, K.-I. Friese, **M. Becker**, F.-E. Wolter, J. T. Lichtenstein, M. Stoetzer, M. Rana, and H. Essig, "Model-based segmentation in orbital volume measurement with cone beam computed tomography and evaluation against current concepts," in *Int. J. Comput. Assist. Radiol. Surg.*, vol. 11, no. 1, pp 1-9, January 2016. Impact Factor: 1.827
- H. F. Choi, A. Chincisan, **M. Becker**, and N. Magnenat-Thalmann, "Multimodal composition of the digital patient: a strategy for the knee articulation," in *Vis. Comput.*, vol. 30, no. 6, pp. 1-11, May 2014. Impact Factor: 0.957

#### Peer-reviewed Conferences

- **M. Becker**, and N. Magnenat-Thalmann, "Muscle tissue labeling of human lower extremities in multi-channel mDixon MR imaging: Concepts and applications," in *Bioinformatics and Biomedicine (BIBM), 2014 IEEE International Conference on*, pp.279,284, 2-5 Nov. 2014, **Best Overall Paper Award**.
- D. García, B. M. A. Delattre, S. Trombella, S. Lynch, **M. Becker**, H. F. Choi, and O. Ratib, "Open framework for management and processing of multi-modality and multidimensional imaging data for analysis and modeling muscular function," *Proc. SPIE 9036, Medical Imaging 2014: Image-Guided Procedures, Robotic Interventions, and Modeling*, 90361W (March 12, 2014).

- L. Assassi, **M. Becker**, and N. Magnenat-Thalmann, “Dynamic Skin Deformation based on Biomechanical Modeling,” in *25th Annual Conference on Computer Animation and Social Agents (CASA 2012)* May 9-11, 2012 Nanyang Technological University, Singapore.

## International symposiums and workshops

- **M. Becker**, K.-I. Friese, F.-E. Wolter, N.-C. Gellrich, and H. Essig, “Development of a Reliable Method for Orbit Segmentation & Measuring,” in *2015 IEEE International Symposium on Medical Measurements and Applications (MeMeA 2015)*, 2015, pp. 285–290.

## Book Chapters

- **M. Becker**, and N. Magnenat-Thalmann, “Deformable Models in Medical Image Segmentation,” in *3D Multiscale Physiological Human*, 1st ed. N. Magnenat-Thalmann, O. Ratib, and H. F. Choi, Eds. Springer-Verlag London, 2014.

---

## APPENDIX B

# ACRONYMS, GLOSSARY AND TECHNICAL DETAILS

---

### B.1 Acronyms

Acronyms used throughout the thesis and their definition:

**AAM:** Active Appearance Model  
**ACL:** Anterior Cruciate Ligament  
**AI:** Affine Invariant  
**ArBB:** Array Building Blocks  
**ASSD:** Average Symmetric Surface Distance  
**ASM:** Active Shape Model  
**AWS:** Amazon Web Services  
**CAD:** Computer-Aided Diagnosis  
**CBCT:** Cone Beam CT  
**CT:** Computed Tomography  
**DM:** Deformable Model  
**EIT:** Electrical Impedance Tomography  
**EM:** Expectation Maximisation  
**EU:** European Union  
**FCM:** Fuzzy C-Means  
**FDM:** Finite Difference Method  
**FEM:** Finite Element Method  
**FNE:** False Negative Error  
**FPE:** False Positive Error  
**GDM:** Geometrically Deformable Model  
**GIS:** Geographical Information Systems  
**GPL:** GNU Public Licence  
**GPU:** Graphical Processing Unit  
**GPGPU:** General Purpose Graphical Processing Unit  
**GUI:** Graphical User Interface  
**GVF:** Gradient Vector Field



**HMM:** Hidden Markov Model  
**HU:** Hounsfield units  
**HUG:** University Hospitals of Geneva  
**ID:** Identifier  
**IP:** Intensity Profile  
**IRT:** Infrared Thermography  
**IT:** Information Technology  
**ITK:** Insight ToolKit  
**ITN:** Initial Training Network  
**LCL:** Lateral Collateral Ligament  
**LUA:** Lua scripting language  
**LET:** Label Equivalence Tree  
**MCL:** Medial Collateral Ligament  
**mDixon:** Modified Dixon  
**MR:** Magnetic Resonance  
**MRF:** Markov Random Field  
**MRI:** Magnetic Resonance Imaging  
**MSD:** Musculoskeletal Disease  
**MSH:** MultiScaleHuman  
**MSSD:** Maximum Symmetric Surface Distance  
**NIH:** National Institutes of Health  
**NLM:** Non-Linear Means  
**NPP:** nVidia Performance Primitives  
**OS:** Operating System  
**PA:** Pennation Angle  
**PCA:** Principal Component Analysis  
**PCL:** Posterior Cruciate Ligament  
**PDE:** Partial Differential Equation  
**PDM:** Point Distribution Model  
**PET:** Positron Emission Tomography  
**PSU:** Per Step Update Strategy  
**PU:** Partial Update Strategy  
**RAM:** Random Access Memory  
**RF:** Radio-Frequency  
**RGB:** Red, Green, Blue  
**RMS:** Root Mean Square error  
**SDC:** Sørensen Dice Coefficient  
**SIAS:** Superior-Inferior Axis Scaling  
**SIAT:** Superior-Inferior Axis Translation  
**SM:** Streaming Multiprocessors  
**SNR:** Signal-to-Noise-Ratio  
**SSM:** Statistical Shape Model  
**SU:** Scaled Update Strategy  
**SVS:** Slice Volume Scaling  
**TBB:** Thread Building Blocks  
**TE:** Echo Time

**TPS:** Thin Plate Spine

**TR:** Repetition Time

**TSE:** Turbo Spin Echo

**UML:** Unified Markup Language

**US:** Ultrasound

**VTK:** Visual ToolKit

**WHO:** World Health Organisation

---

## APPENDIX C

# MEDICAL APPENDIX

---

### C.1 Medical Nomenclature

Over the last centuries, major advances have been made in medicine. Cooperation is one of the important steps leading to it. Knowledge transfer has been done using books and now is extended via networks. To be able to apply results achieved by others correctly, a common language is required. Physicians need exact descriptions and have therefore created their own terminology.

#### C.1.1 Directional Terms

The terms for anatomical directions are applicable to human and animals as well. We will describe the interpretation for the human body. The following list includes the most common directional terms that are used in this work. The definitions are taken from Stedman's Medical Dictionary [Ste90] and are illustrated in Figure C.1.

##### **Superior**

Situated above or directed up. Cranial; in human anatomy, situated nearer the vertex of the head in relation to a specific reference point; opposite of inferior.

##### **Inferior**

Situated below or directed downward. In human anatomy, situated nearer the soles of the feet in relation to a specific reference point; opposite of superior.

##### **Posterior**

Dorsal; dorsalis; posticus; in human anatomy, denoting the back surface of the body. Often used to indicate the position of one structure relative to another, *i.e.*, nearer the back of the body.

##### **Anterior**

Ventral; ventralis; in human anatomy, denoting the front surface of the body; often used to indicate the position of one structure relative to another, *i.e.*, situated nearer the front part of the body.

##### **Medial**

Relating to the middle or centre; nearer to the median or midsagittal plane.

##### **Lateral**

On the side. Farther from the median or midsagittal plane.

##### **Left (Lateral)**

Left from the perspective of the body.

**Right (Lateral)**

Right from the perspective of the body.

**Proximal**

Proximalis; nearest the trunk or the point of origin, said of part of a limb, or and artery or a nerve, etc., so situated.

**Distal**


Distalis. Situated away from the centre of the body, or from the point of origin; specifically applied to the extremity or distant part of a limb or organ.

**Dorsal**

Tergal; pertaining to the back or any dorsum. In human anatomy, posterior.

**Ventral**

Pertaining to the belly or to any venter. Anterior.



**Figure C.1** – Anatomical directional terms. Source: CC-BY 3.0, *Blaussen.com staff*. "Blaussen gallery 2014". *Wikiversity Journal of Medicine*. DOI:10.15347/wjm/2014.010.

### C.1.2 Anatomical Axes

Besides the anatomical directions, the human body can be described along three axes. The **superior-inferior axis** goes from the superior end (the head) to the inferior end (the feet) of the human body. The **antero-posterior axis** connects the front of the body, anterior, and the back, posterior. The third axis can be described either by the uncommon left-right axis or by the **medio-lateral axis** that has its reference in the middle of the body and points laterally to left and to the right. The axes are also shown in Figure C.1.

### C.1.3 Planes

In medical images and anatomical descriptions, three reference planes are used. The normal of the *sagittal plane* points along the medio-lateral axis. The *coronal (or frontal) plane* divides the body into an anterior and posterior part. The *transverse plane* is perpendicular to the other planes and divides the body into a superior and an inferior part. Figure C.2 shows the different planes.




Figure C.2 – Planes through the human body: frontal (a), transversal (b) and sagittal (c) plane.

## C.2 Anatomy

The knee is the most complex joint in the human body. It is made up of four bones, thirteen muscles, and several different kinds of soft tissue. To define the nomenclature and anatomical context for this work, we will a short overview in this section. A more detailed explanation can be found in [Sta05]. We start with the bones of thigh and calf, and continue with their muscles. We will discuss the function and structure of the knee joint and describe the patella and several kinds of soft tissue: tendons, ligaments, cartilage and menisci.


### C.2.1 Bones of Thigh and Calf

Thigh and calf consist of three main bones: Femur, tibia, and fibula. Their main articulation is the knee joint. At its superior end, the femur is connected to hip through the hip joint. The inferior end of tibia and fibula are connected to the ankle for articulation of the foot. The three bones carry the body weight during gait and other movements. They are all long bones<sup>1</sup> with a shaft and two extremities at the ends. They have an outer shell of compact bone and cancellous bone inside the extremities as well as bone marrow inside the shaft.

#### C.2.1.1 Femur

The **femur**, also known as thigh bone, is the only bone in the thigh. It is the longest, strongest, and most heavy bone in the human. Anthropological studies have found the length of the femur usually is 26% of a person's total height. It consists of the body, shaped like a shaft, and the upper and lower extremity. Its strength is crucial since it carries the body weight and has attachments for the muscles used in gait. The femur and its lower and upper extremity can be seen in Figure C.3.

1. Other types of bones are: short, flat, irregular, and sesamoid.



**Figure C.3** – Femur. Source: Henry Gray (1918) *Anatomy of the Human Body*

The **upper extremity** of the femur consists of the femoral head and neck, and the greater and lesser trochanter. The head is located proximal, round, and projecting medially on its short neck. It articulates with the acetabulum, forming the hip joint. The neck connects the head to femoral body. The hip joint capsule attaches between head and neck. The vastus lateralis attaches to the *greater trochanter* of the upper extremity.


The **body** of the femur, the shaft, is long, slender, and almost cylindrical. It is bowed forward. The linea aspera, a ridge on the posterior side supports the strength of the bone. The vastus intermedius is attached to most of the anterior surface.

The **lower extremity** at the distal end of the femur is more massive than the head. It has two condyles (cylindrical shapes). The patella is located anterior of the condyles. The lower extremity has the cruciate ligaments attached and it articulates together with tibia and patella, therefore forming the knee joint.

### C.2.1.2 Tibia and Fibula

The **tibia**, also known as shin bone, is the larger of the two bones in the calf and second largest in the body (after femur). Is made up by its body and two extremities. Tibia and fibula and their extremities are shown in Figure C.4.

The upper extremity is defined by the *tibial plateau*, which is formed by the flattened condyles. It provides a bed for the femoral condyles to move on. In its middle is the *eminentia intercondylaris*, an elevation that divides the plateau into lateral and medial part.



**Figure C.4 – Tibia and Fibula.** Source: Henry Gray (1918) *Anatomy of the Human Body*

The plateau is delineated anterior and posterior by the *Area intercondylaris anterior* and *Area intercondylaris posterior*. Furthermore, it provides attachment points for the cruciate ligaments and the ligaments of the menisci. Located superior of the plateau is the infra-patellar adipoid.

The ligamentum patellae is connected at the *Tuberositas tibiae*, this is supported by the raw and bumpy surface of the tibia at that location. The body of the tibia is a shaft that is defined by an anterior rim *margo anterior* with direct bone-skin-contact. The shaft is shaped triangular, and its surfaces are divided into medial, lateral and posterior surface. The lower tibial extremity forms the ankle joint together with fibula and talus.

The **fibula**, or calf bone, is a bone located lateral of the tibia. It is smaller than the tibia. Its upper extremity is placed towards the back of head of tibia, whereas the lower extremity inclines forward. Although closely located, the fibula is not part of the knee joint. It does not carry a significant amount of weight, but it provides stability to the ankle and offers attachment points for several muscles.

Tibia and fibula are connected through two joints. The *superior tibiofibular joint* connects the fibular head and the lateral condyle of fibula and the *inferior tibiofibular joint* connects tibia and fibula superior of the ankle.

## C.2.2 Knee Joint

The knee joint is a hinge type synovial joint between thigh and calf. It is the largest and most complex joint in the human body. Located between femur, tibia, and patella, it carries the full body weight. Its main task is flexion and extension of the leg, it also allows small internal and external rotation. Due to its heavy usage in everyday tasks, it is prone to abrasive wear. Osteoarthritis developed over time affects the cartilage. Acute injuries also pose a risk.

Femur, tibia, and patella form the joint, the fibula is not a part of it. The bones are connected through ligaments to increase the stability but also to limit the range of possible motions. Especially the anterior and posterior cruciate ligaments are important for stability and coordination. The articulation areas are covered with cartilage to prevent painful, abrasive bone-to-bone contact. The half-moon-shaped menisci provide additional protection and damping. The joint has a capsule that is filled with synovial fluid.

In the following sections, we will describe important components of the joint: patella, tendons, ligaments, cartilage and menisci. An overview of the joint anatomy can be found in Figure C.5.




Figure C.5 – Knee joint with tendons, ligaments, and cartilage (*public domain*).

## C.2.3 Patella

The *patella*, also known as knee cap, is a bony structure, located at the anterior part of the knee joint. It is designed to protect the knee joint from frontal forces. It also allows the quadriceps femoris to improve their efficiency since it proves to be a strong lever.



The patella can be divided into two faces, *facies anterior* and *facies posterior*. The *facies anterior* is covered by the quadriceps femoris-tendon. It has a convex surface and numerous rough, longitudinal striae. Anterior of it is the *bursa subcutanea prepatellaris*. The *facies posterior* is covered by very thick joint cartilage, designed to manage multiple body weight loads. It is divided by a rim into two parts and has a smooth oval surface. Above it lies the suprapatellar adipoid.

The apex of the patella is pointed and gives attachment to ligamentum patellae. During the articulation of the limb, the patella follows the femur.

## C.2.4 Tendons

Tendons are fibrous tissue that connect muscles to bones. They are made of collagen fibres. Tendons usually have the shape of bands and have their name because their capability to tolerate tension. Many muscles have their tendons attached near the knee joint to enable its articulation, mainly flexion and extension. The most prominent tendon is the quadriceps tendons which connects the quadriceps muscle and the patella. The patellar tendon, connecting patella and tibia, is actually not a tendon but a ligament.

## C.2.5 Ligaments

Ligaments are fibrous tissue that connect bones to each other and are made of collagen fibers. They give stability, limit the range of motion and guide the force of muscles. Under tension, they can be stretched, but later return to their original shape. Ligament injuries and ruptures can be left untouched or surgically repaired, depending on their importance for joint stability.

Several ligaments play an important role in the stabilisation of the knee joint. They can be divided into tibiofemoral and patellofemoral ligaments.

The tibiofemoral ligaments consist of the cruciate and the collateral ligaments. The cruciate ligaments are formed by the *Ligamentum cruciatum anterius (ACL)* and the *Ligamentum cruciatum posterius (PCL)*. They connect femur and tibia, crossing each other like an X. The ACL is lateral, the PCL is medial. Together they stabilise the tibia relative to the femur and prevent it from slipping forward or backward. The *Ligamentum collaterale tibiale or medial collateral ligament (MCL)* is located on the medial side of the knee joint and connects femur and tibia. It is not very flexible and more prone to injury. The principal duty is to prevent medial directed forces that cause the valgus deformity. The *Ligamentum collaterale fibulare or lateral collateral ligament (LCL)* is the lateral counterpart of the MCL. It connects femur and fibula.

The patellofemoral ligaments include the *Ligamentum patellae*, which connects the patella to the tibial tuberosity. It is a continuation of the quadriceps femoris and is also known as anterior or patellar ligament. Below this ligament lies a large intrapatellar pad of fat.

These ligaments and their location inside the knee joint can be seen in Figure C.5.

## C.2.6 Cartilage

Cartilage is a connective tissue, it can be of three different types: elastic, hyaline, and fibro-cartilage. It is less hard than bone, but stiffer and less flexible than muscle tissue. It is formed by an extracellular matrix of collagen fibres. Cartilage shows viscoelastic properties. Articular cartilage (at an articulation, e.g. joint) prevents bone-to-bone contact that could damage the bone. The reduced friction makes it easy for bones to slide in joints.

In case of the knee joint, the cartilage can be divided into *articular cartilage* (hyaline cartilage) and *menisci* (fibrous cartilage, see next section). The articular cartilage covers the condyles of femur and tibia. The femoral cartilage is a single entity that covers the articulate surfaces of both condyles and the surface in front of the patella. The tibial cartilage covers the *superior articular surface* of the tibia which is divided into the separate *lateral* and *medial facet*.

### C.2.7 Menisci

The menisci are part of the fibrous cartilage. They are located between femur and tibia. They can be divided into *meniscus medialis* and *lateralis* which are half-moon-shaped. They are attached through ligaments to tibia and femur.

The menisci provide structural integrity during tension and torsion of the knee joint. Although they only partly divide femur and tibia, menisci play an important role in the distribution of the body weight in the joint. They also bridge the gap between the uneven, curved contact surfaces of the articulation.

## C.3 MRI information sheets

1/3

DEPARTEMENT D'IMAGERIE ET DES SCIENCES D'INFORMATION MEDICALE

Service de Médecine Nucléaire

Rue Gabrielle Perret-Gentil 4  
1211 Genève 14  
Tél. : 022 37 27 174  
Fax : 022 37 27 169



Genève, le 24.07.2013

### Formulaire d'information au volontaire (1.IRM)

**Concerne:** *Evaluation of leg muscular activity using Magnetic Resonance Imaging and  $^{11}\text{C}$ -acetate Positron Emission Tomography (PET/MRI). Pilot study.* [Evaluation de l'activité de la musculature de la jambe par imagerie par résonance magnétique et tomographie à émission de positons (PET/IRM) avec  $^{11}\text{C}$ -acétate. Etude pilote]. Détail de l'examen proposé avec ce formulaire, dans le cadre de l'étude décrit: examen IRM.

Monsieur,

Nous vous proposons de participer à un travail de recherche scientifique sur l'évaluation de l'activité musculaire de votre jambe, par imagerie par résonance magnétique (IRM).

**Contexte:** L'imagerie par résonance magnétique est utilisée dans la pratique clinique quotidienne pour la visualisation et la caractérisation de la morphologie des structures corporelles humaines.

Dans le cadre de cette étude, nous vous invitons à vous soumettre à un examen IRM qui sera utile au fin de tester et calibrer notre machine d'examen. Cette étude est réalisée conformément aux lois suisses en vigueur et dans le respect de principes reconnus au plan international, et elle a été approuvée par l'autorité compétente locale (Commission Centrale d'Ethique de la Recherche sur l'Etre Humain, CEREH).

Un nombre de cinq volontaires au total sera prévu pour se soumettre à un examen IRM. Vous gardez l'entière liberté d'accepter ou refuser cette proposition de participer à cette étude.

**Objectif de l'étude:** Cette étude a été conçue dans le cadre du projet de recherche MultiScaleHuman, qui est un projet Européen dans la catégorie "Marie Curie Research Training Networks" financé par l'Union Européenne [Seventh Framework Programme]. Ces "Networks" (réseaux) fournissent les moyens financiers pour permettre à des équipes de recherche européennes de niveau international de collaborer dans le cadre d'un programme d'étude dans un certain domaine. L'objectif scientifique du projet MultiScaleHuman est de réaliser un modèle du fonctionnement du corps humain adapté à chaque individu et, dans le cas particulier de notre étude, d'obtenir une meilleure compréhension de l'anatomie et du fonctionnement de la musculature et des articulations dans le but de pouvoir mieux traiter les pathologies et les affections ostéoarticulaires. Cette étude pourra aussi avoir des applications cliniques dans le suivi de chirurgie ou de traitements orthopédiques pour l'étude du fonctionnement du système musculaire et ostéo-articulaire dans la rééducation et réadaptation post chirurgicale.

**Hypothèse:** Cette étude pilote vise à:

- Evaluer la validité de la technique d'imagerie IRM et de la technique hybride PET/IRM avec de l'acétate marqué dans l'évaluation fonctionnelle de la musculature de la jambe et de son activité.

Evaluation of leg muscular activity using Magnetic Resonance Imaging and  $^{11}\text{C}$ -acetate Positron Emission Tomography (PET/MRI). Pilot study. [Evaluation de l'activité de la musculature de la jambe par imagerie par résonance magnétique et tomographie à émission de positons (PETIRM) avec  $^{11}\text{C}$ -acétate] (Version 1.4 du 24.07.2013).

2/3

- Accomplir l'objectif fixé par le projet Marie-Curie concernant le développement d'outils de visualisation et d'analyse visant à obtenir une meilleure compréhension de l'anatomie et du fonctionnement de la musculature et des articulations humaines.
- Mesurer le fonctionnement *in vivo* de la musculature de la jambe et vérifier que l'incorporation de l'acétate marqué reflète la force contractile musculaire effective et le travail produit du muscle, ayant pour objectif de mieux comprendre et traiter les pathologies correspondantes.

**Lieu de l'étude:** Service de Médecine Nucléaire (Prof.Osman Ratib, Chef de Service), Rue Gabrielle Perret-Gentil 4, 1211 Genève 14.

**Méthode:** Déroulement d'un examen IRM: l'enregistrement d'un examen IRM se déroule en position couchée et à l'intérieur d'un appareil d'IRM en forme de tunnel. La durée d'examen est entre 30 et 45 min au total.

Précautions:

- Rester à jeun au moins 4 heures avant l'examen.

**Critères d'inclusion:**

- Vous êtes un homme âgé plus de 18 ans et moins de 35 ans.
- Vous avez donné votre consentement à participer à cette étude en signant le "Formulaire de consentement éclairé".
- Vous pouvez supporter un temps d'immobilisation entre 30 et 45 min.

**Critères d'exclusion:**

- Vous avez subi des lésions neurologiques ou musculaires aux jambes.
- Vous avez des affections ou des morbidités pouvant affecter le protocole d'examen.
- Vous avez une maladie qui ne vous permet pas de supporter un temps d'immobilisation entre 30 et 45 min.
- Vous êtes porteurs d'un pacemaker ou d'autres implants métalliques, sauf si ceux-ci sont identifiés comme IRM compatibles.
- Vous êtes claustrophobe.

**Risques et bénéfices:** La résonance magnétique (IRM) vous expose à un champ magnétique important qui n'est pas nuisible en soi. Toutefois, pour pouvoir faire l'examen IRM il est important que vous ne soyez pas porteur d'un pacemaker ni avoir des implants métalliques sauf si ceux-ci sont identifiés comme IRM compatibles.

Cet examen évaluant l'anatomie de votre musculature de jambe ne pourra pas vous apporter un bénéfice direct. Mais cette étude pourra avoir des retombées importantes dans les domaines scientifique et clinique. Nous vous proposons cette étude car nous estimons le bénéfice potentiel pour une application future de cet examen, qui pourrait permettre une meilleure prise en charge des pathologies ou des traitements des affections musculaires et ostéo-articulaires et du suivi des effets des interventions chirurgicales et des phases de réadaptation post-chirurgicales.

**Traitement confidentiel des données recueillies dans cette étude:** Toutes les données et résultats observés dans le cadre de cette étude restent strictement confidentiels. Les dossiers de cette étude sont gardés sous clé pendant 10 ans et ensuite détruits. Toutefois, nous vous informons du fait que les membres de la commission d'éthique compétente ainsi que les spécialistes responsables travaillant pour les représentants des autorités Swissmedic et OFSP aient un droit de regard sur les données originales vous concernant pour procéder à des vérifications, ces informations restant cependant strictement confidentielles. D'autre part, les résultats de cette étude peuvent être utilisés dans des publications scientifiques à condition de respecter l'anonymat.

**Frais liés à l'étude:** Votre participation à cette étude n'entraînera aucun frais pour vous. L'examen IRM est pris en charge par un subside que nous avons obtenu pour cette étude.

Evaluation of leg muscular activity using Magnetic Resonance Imaging and <sup>11</sup>C-acetate Positron Emission Tomography (PET/MRI). Pilot study. [Evaluation de l'activité de la musculature de la jambe par imagerie par résonance magnétique et tomographie à émission de positons (PETIRM) avec <sup>11</sup>C-acétate] (Version 1.4 du 24.07.2013).

**Dédommagement:** Les frais de déplacement en transports publics en Suisse peuvent vous être remboursés sur simple demande. La participation à cet essai clinique ne donne droit à aucun dédommagement.

**Assurance:** En cas de dommages subis dans le cadre de l'étude, vous bénéficierez d'une assurance contractée par les HUG, selon les dispositions légales. Votre médecin investigateur entreprendra, le cas échéant, toutes les démarches nécessaires. Une assurance a été contractée pour couvrir cette responsabilité.

**Conditions de participation et obligation de rapporter d'éventuels effets secondaires:** Aucun risque n'est connu ou attendu. En cas d'observation d'effets indésirables dans un délai de 30 jours après l'examen et que vous pensez qu'ils pourraient avoir été liés à cette étude, nous vous demandons de nous les communiquer. L'investigateur vous informera aussi de toute nouvelle découverte qui pourrait influencer votre sécurité ainsi que votre déclaration de consentement. Il pourra également décider de vous exclure de l'étude si des nouveaux éléments étaient mis en évidence concernant votre état de santé ou votre éligibilité à participer à l'étude.

#### **Invitation de participer à l'étude sur base volontaire et liberté de retrait:**

Nous vous invitons donc :

- A participer à cette étude et à nous donner votre consentement en signant le "Formulaire de consentement éclairé" avant l'examen, et au plus tard le jour même de l'examen.
- A nous autoriser, le cas échéant, à utiliser les résultats anonymisés de cette étude pour la poursuite de notre recherche, sans que votre nom apparaisse, et à les publier dans des journaux scientifiques.
- A nous donner votre avis, en cas de découvertes fortuites qui pourraient contribuer à la prévention, au diagnostic et au traitement de maladies existantes ou peut-être futures, en choisissant entre les options suivantes: a) être informé(e) directement de ces résultats b) ne pas être informé(e), ou c) laisser à votre médecin traitant le soin de décider de ce qu'il convient de faire.

Important:

- Vous êtes libre de ne pas participer à cette étude.
- Vous avez le droit de vous retirer en tout temps de l'étude, sans avoir à vous justifier.
- Vous serez informé en cas de nouvelles connaissances pouvant remettre en cause votre participation à l'étude.
- L'investigateur de l'étude a la possibilité de vous exclure à tout moment, en vous précisant le motif.
- Vous pouvez vous adresser à tout moment au responsable de l'étude pour obtenir des informations complémentaires.
- L'investigateur respecte la confidentialité de vos données.

**Informations:** Pour toutes informations complémentaires, vous pouvez vous adresser au promoteur de l'étude, Prof. Osman Ratib, Chef du Service, Service de Médecine Nucléaire, HUG: Tél.: 022 372 71 45 Email: [osman.ratib@hcuge.ge](mailto:osman.ratib@hcuge.ge)

## C.4 MRI consent form

1/1

DEPARTEMENT D'IMAGERIE ET DES SCIENCES D'INFORMATION  
MEDICALE

Service de Médecine Nucléaire



**Evaluation of leg muscular activity using Magnetic Resonance Imaging and  $^{11}\text{C}$ -acetate Positron Emission Tomography (PET/MRI). Pilot study. [Evaluation de l'activité de la musculature de la jambe par imagerie par résonance magnétique et tomographie à émission de positons (PET/IRM) avec  $^{11}\text{C}$ -acétate. Etude pilote]**  
Détail de l'examen: examen IRM.

### Formulaire de consentement éclairé

Version 1.3 du 24.07.2013

- Le médecin signataire m'a informé(e) oralement et par écrit des buts de l'étude précitée, de son déroulement, des effets attendus, des risques et inconvénients possibles.
- J'ai lu et compris le formulaire d'information au patient version 1.4 du 24.07.2013 pour l'étude susnommée. J'ai reçu des réponses satisfaisantes aux questions concernant ma participation à cette étude. Je peux garder le dossier d'information du patient et je reçois une copie de ma déclaration écrite de consentement.
- J'ai eu suffisamment de temps pour réfléchir avant de prendre ma décision.
- Je sais qu'une assurance couvre les dommages qui pourraient survenir dans cette étude.
- J'accepte le fait que les spécialistes responsables travaillant pour les représentants des autorités Swissmedic et OFSP et de la commission d'éthique aient un droit de regard sur les données originales me concernant pour procéder à des vérifications, ces informations restant toutefois strictement confidentielles. Je suis d'accord que les données récoltées durant cette étude puissent être utilisées pour des publications scientifiques à condition de respecter l'anonymat.
- En cas de découvertes fortuites qui pourraient contribuer à la prévention, au diagnostic et au traitement de maladies existantes ou peut-être futures, je donne mon consentement à: ☐ être informé(e) directement de ces résultats ☐ ne pas être informé(e), ou ☐ laisser à votre médecin traitant le soin de décider de ce qu'il convient de faire.
- Je participe volontairement à cette étude. Je peux à tout moment retirer mon accord de participation à cette étude sans avoir à donner de raisons. Aucun inconvénient pour mon suivi médical ultérieur ne doit découler de cette décision.
- Je sais que les exigences et restrictions mentionnées dans le formulaire d'information du patient doivent être respectées durant l'étude. Dans l'intérêt de ma santé, l'investigateur peut à tout moment décider de m'exclure de l'étude.
- Je consens à participer,

Date : .....

Nom, Prénom du **patient** :                      Signature du **patient** :

.....

Date : .....

Nom du **médecin** investigateur :                      Signature du **médecin** investigateur

.....

Evaluation of leg muscular activity using Magnetic Resonance Imaging and  $^{11}\text{C}$ -acetate Positron Emission Tomography (PET/MRI). Pilot study. Version 1.3 (24.07.2013)

---

## APPENDIX D

# TECHNICAL APPENDIX

---

### D.1 Other Deformable Models

In this section, we will give a short overview of other types of deformable models that fall outside the classification described so far. We will discuss deformable Fourier models, the modal analysis and deformable superquadrics.

#### D.1.1 Deformable Fourier Models

In 1992, Staib and Duncan [SD92] have proposed to use a Fourier representation for deformable contours and surfaces. In the following, we will take a look at a closed Fourier contour, described as:

$$X(s) = \begin{bmatrix} X(s) \\ Y(s) \end{bmatrix} = \begin{bmatrix} a_0 \\ c_0 \end{bmatrix} + \sum_{k=1}^{\infty} \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} \begin{bmatrix} \cos 2\pi ks \\ \sin 2\pi ks \end{bmatrix} \quad (\text{D.1})$$

with the Fourier coefficients  $a_0, c_0, a_i, b_i, c_i, d_i$ . They are defined as

$$\begin{aligned} a_0 &= \frac{1}{2\pi} \int_0^1 X(s) ds, & a_k &= \frac{1}{\pi} \int_0^1 X(s) \cos 2\pi ks ds, \\ b_k &= \frac{1}{\pi} \int_0^1 X(s) \sin 2\pi ks ds \end{aligned} \quad (\text{D.2})$$

A smooth representation can be obtained by truncating the series. The shape translation is defined by the coefficients  $a_0$  and  $c_0$ . The subsequent terms follow the parametric form of an ellipse and can be mapped to the standard properties of an ellipse [SKBG96]. The parameters follow scale ordering, with low indices for global properties and high indices describing local deformations.

To incorporate prior knowledge, Staib and Duncan [SD92] use a Bayesian approach. A prior probability approach is defined by manual delineation and the consequent parametrisation of Fourier coefficients using a converted ellipse parameter set. This way, a mean and variance can be calculated for each parameter.

### D.1.2 Deformable Models Using Modal Analysis

Modal analysis has been introduced by Pentland and Horowitz [HP91]. Deformable models using it, are similar to deformable Fourier models but their base functions and nominal values are derived from templates. The objects consist of finite elements, stacked in the vector  $X$ . Displacements are stored in  $U$ , so that the new state after a deformation step is given by  $X + U$ . They are constrained using the following differential equation

$$M \frac{d^2 U}{dt^2} + C \frac{dU}{dt} + KU = f \quad (D.3)$$

with  $M$  as mass,  $C$  as damping and  $K$  as stiffness matrix. The vector  $f$  contains the external forces and  $f$  and  $U$  are defined as functions of time.

### D.1.3 Deformable Superquadrics

To incorporate local and global shape features, Terzopoulos and Metaxas [TM91] have proposed deformable superquadrics. These deformable models use a superquadric surface. While global deformations have a large influence on the global shape characteristics, local deformations capture the details. Local and global deformations are being performed in parallel. The closed surface  $x(u)$  in parametric coordinates  $u = [v, w]$  is expressed as:

$$x(u) = c + Rp(u) \quad (D.4)$$

with  $c$  a translation vector and  $R$  a rotation matrix. The model shape is expressed by  $p(u)$  and is the sum of the reference shape  $s(u)$  and the local deformations  $d(u)$ . Superquadrics are a popular extension to quadrics and can model many shapes with few parameters. Examples for superquadrics are superellipsoids, which are described in [BCA98].

## D.2 Functionality exposed to Lua scripting

Large parts of the functionality of our DMFramework are exposed to Lua scripting. We have identified five main categories:

1. Image Processing
2. Simulation
3. Models
4. Forces
5. Comparison

These categories are grouped in namespaces and described in the following sections. We have focussed on the main features, omitting minor functionality.

### D.2.1 Image processing

Image processing is a crucial step for medical image segmentation. We rely on the ITK library and extend it to our needs. We expose our needed functionality for flexible processing. For segmentation tasks, the main class is the Image Manager. It supports centralised loading of images. Since the images



are static and not modified during the actual segmentation process, multiple models, forces and even simulation environments can share the images and reduce the required memory. The image manager provides access to individual slices and supports the calculation of gradients.

We also provide the `ImageLoader` and `ImageWriter` to support a large range of image formats like `MetaImage`, `NII`, etc for loading and saving. We have extended the loader to also allow the loading of `Dicom` files from directories and have a rudimentary support for saving `Dicom` files. The functionality is provided for integer as well as floating point images. Finally, methods for creating new (empty) images of arbitrary sizes, e.g. for processing or stitching results, are available.

We have made the following custom filters available:

- **`RemoveDuplicateSlicesFilter::removeDuplicateSlices`** (*Image*, *NumberOfDuplicateSlices*): Remove duplicate slices that are a by-product of the `mDixon` scans.
- **`ImageStitchingFilter::stitchImages`** (*Image1*, *Image2*): Stitches two images using their offsets and interpolates overlapping volumes.
- **`DivideLegsInAirImage::divide`** (*AirImage*, *left*): Creates a mask for either the left or the right leg from an `AirImage`.

Other filters are wrappers and shown below:

- **`ImageAddFilter::addImages`** (*Image1*, *Image2*): Adds two images.
- **`BinaryThresholdFilter::thresholdImage`** (*Image*, *Threshold*): Performs a binary thresholding in `Image`.
- **`ConnectedComponentsFilter::findConnectedComponentsWithMinSize`** (*Image*, *MinSize*): Finds connected components of a minimal size (*MinSize*) in `Image`.
- **`BinaryErosionFilter::erodeImage`** (*Image*): Performs an erosion on the binary `Image`.
- **`BinaryDilationFilter::dilateImage`** (*Image*): Performs a dilation on the binary `Image`.
- **`PasteImageFilter::pasteImage`** (*Source*, *Target*, *Offset*, *Size*): Pastes the source image or parts into the target the target image.
- **`InvertIntensityImageFilter::invert`** (*Image*, *MaxIntensity*): Invert the intensity in `Image`. Uses the *maxIntensity* value to remove the need for an additional scan of `Image`.
- **`CropImageFilter::cropImage`** (*Image*, *Size*, *Offset*): Crops an image to the desired sub-image.
- **`ImageToModelFilter::createModelFromImage`** (*Image*, *Label*): Creates a 3D model from a `Label` in `Image` using the marching cubes approach.
- **`MultiplyImageFilter::multiplyImage`** (*Image*, *Factor*): Multiplies `Image` with a factor.
- **`BiasFieldImageFilter::getBiasFieldForImage`** (*Image*, *Mask*): Estimates the MR Bias field for `Image`. The mask can be supplied to reduce the calculation in areas like air which reduces the calculation time.

## D.2.2 Simulation

The simulation class capsules and configures the simulation environment. There can be several instances of the simulation class in parallel for different simulation environments, e.g. for multi-resolution mesh coupling.

- **`setupSimulation`** (): Sets up the simulation environment and configures it.
- **`tearDownSimulation`** (): Tears down the simulation environment and frees resources.
- **`performSimulationStep`** (*stepSize*): Performs a simulation step of the indicated length. Default values is 1/60s. Longer steps might be broken down into smaller steps to ensure continuity.

- **addModel** (*Model, Key*): Adds a model to the simulation environment. The model can be identified by the key.
- **addStaticPlane** (*center, orientation*): Adds a static plane to the simulation environment.
- **createModelFromOpenMesh** (*Mesh*): Creates a simulation model from an OpenMesh mesh.
- **kDP**: The damping coefficient used for simulation.
- **kDF**: The dynamic friction coefficient used for simulation.
- **kDG**: The drag coefficient used for simulation.

### D.2.3 Models

The Model Manager keeps track of the models in a simulation environment. It is the place for calculating forces and supports decimation. Model Managers cannot be shared between simulation environments.

- **AddFrom(DefmFile | OpenMesh | VTK | VTP)** (*Filename, Key*): Loads a model and stores it under the Key.
- **CalcForces** (): Calculates the forces for all models.
- **UpdateNormals** (): Updates the normals of all models, e.g. after a transform call.
- **RemoveAllModels** (): Removes all models from the Model Manager.
- **DecimateModel** (*Key*): Decimates the model Key to the specified number of vertices.

The Model class encapsulates the geometry and the forces of a model.

- **AddForce** (*Force*): Adds a force to the model.
- **RemoveAllForces** (): Removes all forces from the model.
- **ResetOutsidePointForces** (): Set the force to zero for all vertices outside the image volume.
- **Save(OpenMesh | VTK | VTP)** (*Filename*): Saves the model to a file.
- **Transform** (*Matrix4x4*): Transforms the model using the supplied matrix that can express translation, rotation and scaling.

### D.2.4 Forces

All forces share a common base class that defines a method used by the simulation for force calculation. Aside from that, there are no limitations. Some of the forces are the following:

- **GrowingForce**: This force creates a constant pull along the normals.
- **LaplacianSmoothingForce**: This force applies Laplacian smoothing.
- **SimpleImageForce**: The image gradient of a single image is used in this force.
- **FourImageImageForce**: This forces uses the weighted gradients of up to four images.
- **InfluencePointForce**: Influence points can be used to pull the model to or away from a point. This force implements this behaviour.
- **MuscleLabelForce**: This force uses a vector field to pull the models to stay within muscle tissue provided from our labelling technique.

### D.2.5 Comparison

For the comparison, we use a special image type: LabelImages. All label image metrics are encapsulated in the LabelImageMetrics class which is realised as a singleton. It offers several metrics for image to image comparison, below we have list some of the provided methods. The LabelImageMetrics class is a wrapper for the corresponding ITK functionality.

- **LoadLabelImage** (*Filename*): Loads a label image from a file.

- **SaveLabelImage** (*Filename, Image*): Saves a label image to a file.
- **GetFalseNegativeError** (*Image1, Image2, Label*): Calculate the false negative error for Label in two images (Image1, Image2).
- **GetFalsePositiveError** (*Image1, Image2, Label*): Calculate the false positive error for Label in two images (Image1, Image2).
- **GetTargetOverlap** (*Image1, Image2, Label*): Calculate the overlap of Label in two images (Image1, Image2).
- **GetJaccardCoefficient** (*Image1, Image2, Label*): Determines the Jaccard-Coefficient for Label in two images (Image1, Image2).
- **GetDiceCoefficient** (*Image1, Image2, Label*): Determines the Dice-Coefficient for Label in two images (Image1, Image2).

---

# BIBLIOGRAPHY

---

- [AA84] E. Adelson and C. Anderson, "Pyramid methods in image processing," *RCA Engineer*, vol. 29, no. 6, pp. 33–41, 1984. [Online]. Available: <https://alliance.seas.upenn.edu/{~}cis581/wiki/Lectures/Pyramid.pdf>
- [AAAD]<sup>+</sup>15] M. Al-Ayyoub, A. M. Abu-Dalo, Y. Jararweh, M. Jarrah, and M. A. Sa'd, "A GPU-based implementations of the fuzzy C-means algorithms for medical image segmentation," *The Journal of Supercomputing*, vol. 71, no. 8, pp. 3149–3162, 2015. [Online]. Available: <http://link.springer.com/10.1007/s11227-015-1431-y>
- [AALB13] A. a. Appel, M. a. Anastasio, J. C. Larson, and E. M. Brey, "Imaging challenges in biomaterials and tissue engineering." *Biomaterials*, vol. 34, no. 28, pp. 6615–30, sep 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23768903>
- [ABB<sup>+</sup>01] L. S. Avila, S. Barré, R. Blue, D. Cole, B. Geveci, W. A. Hoffman, B. King, C. C. Law, K. M. Martin, W. J. Schroeder, and A. H. Squillacote, *The VTK User's Guide*. Kitware Inc., 2001. [Online]. Available: [www.vtk.org](http://www.vtk.org)
- [ACH14] S. Andrews, N. Changizi, and G. Hamarneh, "The isometric log-ratio transform for probabilistic multi-label anatomical shape representation," *IEEE Transactions on Medical Imaging*, vol. 33, no. 9, pp. 1890–1899, 2014.
- [Ack98] M. J. Ackerman, "The visible human project," *Proceedings of the IEEE*, vol. 86, no. 3, pp. 504–511, 1998.
- [ACS<sup>+</sup>09] L. Assassi, C. Charbonnier, J. Schmid, P. Volino, and N. Magnenat-Thalmann, "From MRI to anatomical simulation of the hip joint," *Computer Animation and Virtual Worlds*, vol. 20, no. 1, pp. 53–66, jan 2009. [Online]. Available: <http://doi.wiley.com/10.1002/cav.266>
- [ACV<sup>+</sup>] AAU, CRS4, VRLab, INRIA, IOR, MIRALab, UCL, and VUB, "3D Anatomical Human project." [Online]. Available: <http://3dah.miralab.ch/>
- [AHP<sup>+</sup>05] E. D. Angelini, S. Homma, G. Pearson, J. W. Holmes, and A. F. Laine, "Segmentation of real-time three-dimensional ultrasound for quantification of ventricular function: a clinical study on right and left ventricles." *Ultrasound in medicine & biology*, vol. 31, no. 9, pp. 1143–1158, 2005.
- [AHS<sup>+</sup>99] D. Atkinson, D. L. G. Hill, P. N. R. Stoye, P. E. Summers, S. Clare, R. Bowtell, and S. F. Keevil, "Automatic compensation of motion artifacts in MRI," *Magnetic Resonance in Medicine*, vol. 41, no. 1, pp. 163–170, 1999.
- [AHY<sup>+</sup>11] S. Andrews, G. Hamarneh, A. Yazdanpanah, B. HajGhanbari, and W. D. Reid, "Probabilistic multi-shape segmentation of knee extensor and flexor muscles." *Medical image computing and computer-assisted intervention : MICCAI International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 14, no. Pt 3, pp. 651–8, jan 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22003755>
- [AIA11] S. AlZubi, N. Islam, and M. Abbod, "Enhanced Hidden Markov Models for accelerating medical volumes segmentation," in *2011 IEEE GCC Conference and Exhibition GCC*, Department of Electronic & Computer Engineering, School of Engineering and Design, Brunel University - West London. IEEE, 2011, pp. 287–290.
- [AIJL02] E. D. Angelini, C. Imielinska, Y. Jin, and A. F. Laine, "Improving statistics for hybrid segmentation of high-resolution multichannel images," *Medical Imaging 2002: Image Processing*, vol. 4684, pp. 401–411, may 2002. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.467182>

- [AKJ16] M. R. Avendi, A. Kheradvar, and H. Jafarkhani, "A Combined Deep-Learning and Deformable-Model Approach to Fully Automatic Segmentation of the Left Ventricle in Cardiac MRI," *Medical Image Analysis*, vol. 30, pp. 1–34, 2016. [Online]. Available: <http://arxiv.org/abs/1512.07951>
- [AKP<sup>+</sup>14] B. T. Addeman, S. Kutty, T. G. Perkins, A. S. Soliman, C. N. Wiens, C. M. McCurdy, M. D. Beaton, R. a. Hegele, and C. a. McKenzie, "Validation of volumetric and single-slice MRI adipose analysis using a novel fully automated segmentation method." *Journal of magnetic resonance imaging : JMRI*, vol. 00, jan 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24431195>
- [AM14] Z. Alom and M. Mostakim, "Automatic Slice Growing Method Based 3D Reconstruction of Liver with its Vessels," in *16th Int'l Conf. Computer and Information Technology*, no. March, 2014, pp. 8–10.
- [APG11] S. Y. Ababneh, J. W. Prescott, and M. N. Gurcan, "Automatic graph-cut based segmentation of bones from knee magnetic resonance images for osteoarthritis research," *Medical Image Analysis*, vol. 15, no. 4, pp. 438–448, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841511000223>
- [AS95] D. Adalsteinsson and J. A. Sethian, "A Fast Level Set Method for Propagating Interfaces," *Journal of Computational Physics*, vol. 118, no. 2, pp. 269–277, 1995. [Online]. Available: <http://math.berkeley.edu/~sethian/Papers/sethian.adalsteinsson.narrowband.ps.gz>
- [AS99] D. Adalsteinsson and J. A. Sethian, "The Fast Construction of Extension Velocities in Level Set Methods," *Journal of Computational Physics*, vol. 148, no. 1, pp. 2–22, 1999. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0021999198960909>
- [AYDM14a] E. Ahmad, M. H. Yap, H. Degens, and J. McPhee, "Enhancement of MRI human thigh muscle segmentation by template-based framework," in *2014 IEEE International Conference on Control System, Computing and Engineering (ICCSCE 2014)*. IEEE, nov 2014, pp. 405–410. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7072753>
- [AYDM14b] E. Ahmad, M. H. Yap, H. Degens, and J. S. McPhee, "Atlas-registration based image segmentation of MRI human thigh muscles in 3D space," *Progress in Biomedical Optics and Imaging - Proceedings of SPIE*, vol. 9037, 2014. [Online]. Available: <http://www.scopus.com/inward/record.url?eid=2-s2.0-84901812643&partnerID=tZOTx3y1>
- [BACP12a] P.-Y. Baudin, N. Azzabou, P. G. Carlier, and N. Paragios, "Automatic skeletal muscle segmentation through random walks and graph-based seed placement," in *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, may 2012, pp. 1036–1039. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6235735>
- [BACP12b] P. Y. Baudin, N. Azzabou, P. G. Carlier, and N. Paragios, "Prior knowledge, random walks and human skeletal muscle segmentation." *Medical image computing and computer-assisted intervention : MICCAI International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 15, no. Pt 1, pp. 569–76, jan 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23285597>
- [BACP12c] P. Baudin, N. Azzabou, P. Carlier, and N. Paragios, "Manifold-enhanced Segmentation through Random Walks on Linear Subspace Priors," in *Proceedings of the British Machine Vision Conference*, K. Bowden, Richard and Collomosse, John and Mikolajczyk, Ed. BMVA Press, 2012, pp. 52.1–52.10. [Online]. Available: <http://www.bmva.org/bmvc/2012/BMVC/paper052/abstract052.pdf>
- [BB84] D. Barber and B. Brown, "Applied potential tomography," *Journal of physics. E. Scientific instruments*, vol. 723, 1984.
- [BB02] V. Barra and J.-Y. Boire, "Segmentation of fat and muscle from MR images of the thigh by a possibilistic clustering algorithm." *Computer methods and programs in biomedicine*, vol. 68, no. 3, pp. 185–93, jun 2002. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/12074845>
- [BBBČ11] J. Barnat, P. Bauch, L. Brim, and M. Češka, "Computing strongly connected components in parallel on CUDA," *Proceedings - 25th IEEE International Parallel and Distributed Processing Symposium, IPDPS 2011*, pp. 544–555, 2011.
- [BBBV<sup>+</sup>13] Y. Barnouin, G. Butler-Browne, T. Voit, D. Reversat, N. Azzabou, G. Leroux, A. Behin, J. S. McPhee, P. G. Carlier, and J.-Y. Hogrel, "Manual segmentation of individual muscles of the quadriceps femoris using MRI: A reappraisal." *Journal of magnetic resonance imaging : JMRI*, vol. 00, pp. 1–9, nov 2013. [Online].

- Available: <http://www.ncbi.nlm.nih.gov/pubmed/24615897>
- [BCA98] E. Bardinet, L. D. Cohen, and N. Ayache, "A Parametric Deformable Model to Fit Unstructured 3D Data," *Computer Vision and Image Understanding*, vol. 71, no. 1, pp. 39–54, 1998. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1077314297905956>
- [BCM05] A. Buades, B. Coll, and J. Morel, "A non-local algorithm for image denoising," *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 2, pp. 60–65, 2005. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=1467423](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=1467423)
- [BD93] B. Bascle and R. Deriche, "Stereo matching, reconstruction and refinement of 3D curves using deformable contours," in *Computer Vision, 1993. Proceedings., Fourth International Conference on*, Berlin, 1993, pp. 421–430. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=378185](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=378185)
- [BD04] A. Berrington de González and S. Darby, "Risk of cancer from diagnostic X-rays: Estimates for the UK and 14 other countries," *Lancet*, vol. 363, no. 9406, pp. 345–351, 2004.
- [BDB<sup>+</sup>01] T. J. Blokhuis, J. H. D. De Bruine, J. a. M. Bramer, F. C. Den Boer, F. C. Bakker, P. Patka, H. J. T. M. Haarman, and R. a. Manoliu, "The reliability of plain radiography in experimental fracture healing," *Skeletal Radiology*, vol. 30, no. 3, pp. 151–156, 2001.
- [BETV08] H. Bay, A. Ess, T. Tuytelaars, and L. Vangool, "Speeded-Up Robust Features (SURF)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1077314207001555>
- [Bey87] G. Beylkin, "Discrete radon transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 35, no. 2, pp. 162–172, 1987.
- [BFS<sup>+</sup>07] P. Bourgeat, J. Frapp, P. Stanwell, S. Ramadan, and S. Ourselin, "MR image segmentation of the knee bone using phase information," *Medical Image Analysis*, vol. 11, no. 4, pp. 325–335, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841507000291>
- [BFS<sup>+</sup>12] Y. Berker, J. Franke, a. Salomon, M. Palmowski, H. C. W. Donker, Y. Temur, F. M. Mottaghy, C. Kuhl, D. Izquierdo-Garcia, Z. a. Fayad, F. Kiessling, and V. Schulz, "MRI-Based Attenuation Correction for Hybrid PET/MRI Systems: A 4-Class Tissue Segmentation Technique Using a Combined Ultrashort-Echo-Time/Dixon MRI Sequence," *Journal of Nuclear Medicine*, vol. 53, no. 5, pp. 796–804, 2012.
- [BLS03] J. Bredno, T. M. T. Lehmann, and K. Spitzer, "A General Discrete Contour Model in Two, Three, and Four Dimensions for Topology-Adaptive Multichannel Segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 5, pp. 550–563, 2003. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=1195990](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=1195990)
- [BMT14] M. Becker and N. Magnenat-Thalmann, "Muscle Tissue Labeling of Human Lower Extremities in Multi-Channel mDixon MR Imaging: Concepts and Applications," in *Bioinformatics and Biomedicine (BIBM), 2014 IEEE International Conference on*, 2014, pp. 279–284. [Online]. Available: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6999168&isnumber=6999111>
- [BMVS02] G. Behiels, F. Maes, D. Vandermeulen, and P. Suetens, "Evaluation of image features and search strategies for segmentation of bone structures in radiographs using Active Shape Models," *Medical Image Analysis*, vol. 6, no. 1, pp. 47–62, 2002. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/11836134>
- [Bor84] G. Borgefors, "Distance transformations in arbitrary dimensions," *Computer Vision Graphics and Image Processing*, vol. 27, no. 2, pp. 321–345, 1984. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/0734189X84901944>
- [BPB13] J. a. Bogovic, J. L. Prince, and P.-L. Bazin, "A Multiple Object Geometric Deformable Model for Image Segmentation," *Computer vision and image understanding : CVIU*, vol. 117, no. 2, pp. 145–157, feb 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23316110>
- [BPD<sup>+</sup>12] C. Boydev, D. Pasquier, F. Derraz, L. Peyrodie, A. Taleb-ahmed, and J.-p. Thiran, "Shape-based Interpolation of a Set of 2D Slices," *Insight Journal*, pp. 1–7, 2012.
- [Bro77] R. A. Brooks, "A Quantitative Theory of the Hounsfield Unit and Its Application to Dual Energy Scanning," *Journal of Computer Assisted Tomography*, vol. 1, no. 4, 1977. [Online]. Available: <http://journals.lww.com/jcat/Fulltext/1977/10000/>

- A[ ]Quantitative[ ]Theory[ ]of[ ]the[ ]Hounsfield[ ]Unit[ ]and.16.aspx
- [BSBK02] M. Botsch, S. Steinberg, S. Bischoff, and L. Kobbelt, "OpenMesh – a generic and efficient polygon mesh data structure," in *OpenSG Symposium*. Citeseer, 2002. [Online]. Available: <http://www.graphics.rwth-aachen.de/media/papers/openmesh1.pdf>
- [BW98] D. Baraff and A. Witkin, "Large steps in cloth simulation," in *Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98*. New York, New York, USA: ACM Press, 1998, pp. 43–54. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=280814.280821>
- [Can86] J. Canny, "A computational approach to edge detection." *IEEE transactions on pattern analysis and machine intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [Car52] H. Y. Carr, "Free precession techniques in nuclear magnetic resonance," PhD thesis, Harvard University, 1952.
- [Car06] P. Caravan, "Strategies for increasing the sensitivity of gadolinium based MRI contrast agents." *Chemical Society reviews*, vol. 35, no. 6, pp. 512–523, 2006.
- [CAVMT09] C. Charbonnier, L. Assassi, P. Volino, and N. Magnenat-Thalmann, "Motion study of the hip joint in extreme postures," *The Visual Computer*, vol. 25, no. 9, pp. 873–882, mar 2009. [Online]. Available: <http://link.springer.com/10.1007/s00371-009-0317-5>
- [CB12] P. Choong and P. Brooks, "Achievements during the bone and joint decade 2000-2010," *Best Practice and Research: Clinical Rheumatology*, vol. 26, no. 2, pp. 173–181, 2012.
- [CC93] L. Cohen and I. Cohen, "Finite-element methods for active contour models and balloons for 2-D and 3-D images," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 15, no. 11, pp. 1131–1147, 1993. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs\[ \]\\_all.jsp?arnumber=244675](http://ieeexplore.ieee.org/xpls/abs[ ]_all.jsp?arnumber=244675)
- [CCA92] I. Cohen, L. D. Cohen, and N. Ayache, "Using deformable surfaces to segment 3-D images and infer differential structures," *Comp Vision Graphics and Image Understanding*, vol. 56, no. 2, pp. 242–263, 1992. [Online]. Available: [http://dx.doi.org/10.1016/1049-9660\(92\)90041-Z](http://dx.doi.org/10.1016/1049-9660(92)90041-Z)
- [CD09] F. Chung and H. Delingette, "Multimodal prior appearance models based on regional clustering of intensity profiles." *Medical Image Computing and Computer-Assisted Intervention*, vol. 12, no. Pt 2, pp. 1051–1058, 2009. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20426215>
- [CDP14] S. Cuomo, P. De Michele, and F. Piccialli, "3D Data Denoising via Nonlocal Means Filter by Using Parallel GPU Strategies," *Computational and Mathematical Methods in Medicine*, vol. 2014, pp. 1–14, 2014. [Online]. Available: <http://www.hindawi.com/journals/cmmm/2014/523862/>
- [CET98] T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," in *Computer Vision — ECCV'98*, ser. Lecture Notes in Computer Science, H. Burkhardt and B. Neumann, Eds. Springer Berlin / Heidelberg, 1998, vol. 1407, pp. 484–498. [Online]. Available: <http://dx.doi.org/10.1007/BFb0054760>
- [CFK04] J. Cheng, S. W. Foo, and S. M. Krishnan, "Directional gradient vector flow for snakes," in *Proceedings of the Fourth IEEE International Symposium on Signal Processing and Information Technology 2004*. Ieee, 2004, pp. 318–321. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1433748>
- [CGZM15] Y. Cai, X. Guo, Z. Zhong, and W. Mao, "Dynamic meshing for deformable image registration," *Computer-Aided Design*, vol. 58, pp. 141–150, jan 2015. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S001044851400178X>
- [CH10] N. Changizi and G. Hamarneh, "Probabilistic multi-shape representation using an isometric log-ratio mapping." *Medical image computing and computer-assisted intervention : MICCAI International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 13, no. Pt 3, pp. 563–70, jan 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20879445>
- [Cha08] C. Charbonnier, "Analysis of extreme hip motion in professional ballet dancers," *Proceedings of the 10th International Symposium on 3D Analysis of Human Movement, Amsterdam, Netherlands*, pp. 3–4, 2008. [Online]. Available: <http://www.miralab.ch/repository/papers/496.pdf>
- [CHTH94] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The Use of Active Shape Models for Locating Structures in Medical Images," *Image and Vision Computing*, vol. 12, no. 6, pp. 355–366, 1994. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0262885694900604>

- [CIB13] CIBC, "Seg3D: Volumetric Image Segmentation and Visualization. Scientific Computing and Imaging Institute (SCI)," 2013. [Online]. Available: <http://www.seg3d.org>
- [CIKW99] P. G. Ciarlet, A. Iserles, R. V. Kohn, and M. H. Wright, Eds., *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, ser. Cambridge Monograph on Applied and Computational Mathematics. Cambridge University Press, 1999. [Online]. Available: <http://www.amazon.com/Level-Methods-Fast-Marching-Computational/dp/0521645573>
- [CJP<sup>+</sup>10] C. W. Chun, W.-H. Jee, H. J. Park, Y. J. Kim, J.-M. Park, S.-H. Lee, and S.-H. Park, "MRI Features of Skeletal Muscle Lymphoma," *American Journal of Roentgenology*, vol. 195, no. 6, pp. 1355–1360, 2010. [Online]. Available: <http://www.ajronline.org/doi/abs/10.2214/AJR.09.3904>
- [CKM<sup>+</sup>12] C. M. Costelloe, V. Kundra, J. Ma, B. a. Chasen, E. M. Rohren, R. L. Bassett, and J. E. Madewell, "Fast dixon whole-body MRI for detecting distant cancer metastasis: a preliminary clinical study," *Journal of Magnetic Resonance Imaging*, vol. 35, no. 2, pp. 399–408, 2012. [Online]. Available: <http://doi.wiley.com/10.1002/jmri.22815>
- [CKS95] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," in *Proceedings of IEEE International Conference on Computer Vision*, vol. 22. IEEE Comput. Soc. Press, 1995, pp. 61–79. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=466871>
- [CML<sup>+</sup>12] R. Chabiniok, P. Moireau, P. F. Lesault, a. Rahmouni, J. F. Deux, and D. Chapelle, "Estimation of tissue contractility from cardiac cine-MRI using a biomechanical heart model," *Biomechanics and Modeling in Mechanobiology*, vol. 11, no. 5, pp. 609–630, 2012.
- [CMS<sup>+</sup>14] R. Csapo, V. Malis, U. Sinha, J. Du, and S. Sinha, "Age-associated differences in triceps surae muscle composition and strength - an MRI-based cross-sectional comparison of contractile, adipose and connective tissue." *BMC musculoskeletal disorders*, vol. 15, no. 1, p. 209, jan 2014. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4072482&tool=pmcentrez&rendertype=abstract>
- [Coh91] L. D. L. Cohen, "On active contour models and balloons," *CVGIP: Image understanding*, vol. 53, no. 2, pp. 211–218, 1991. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/104996609190028N>
- [CS00] H.-D. D. C. H.-D. Cheng and Y. S. Y. Sun, "A hierarchical approach to color image segmentation using homogeneity," *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 9, no. 12, pp. 2071–2082, 2000. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=887975](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=887975)
- [CSV00] T. F. Chan, B. Sandberg, and L. a. Vese, "Active Contours without Edges for Vector-Valued Images," *Journal of Visual Communication and Image Representation*, vol. 11, no. 2, pp. 130–141, jun 2000. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S104732039990442X>
- [CT95] T. F. Cootes and C. J. Taylor, "Combining point distribution models with shape models based on finite element analysis," *Image and Vision Computing*, vol. 13, no. 5, pp. 403–409, jun 1995. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/026288569599727I>
- [CT01] T. F. Cootes and C. J. Taylor, "Statistical Models of Appearance for Computer Vision," *World Wide Web Publication February*, 2001. [Online]. Available: [http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.58.1455\(&rep=rep1&type=pdf](http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.58.1455(&rep=rep1&type=pdf)
- [CTCG95] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models - their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995. [Online]. Available: <http://www.isbe.man.ac.uk/{~}bim/Papers/cootes{~}cviu95.pdf>
- [CUB<sup>+</sup>12] X. Chen, J. K. Udupa, U. Bagci, Y. Zhuge, and J. Yao, "Medical Image Segmentation by Combining Graph Cuts and Oriented Active Appearance Models," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 2035–2046, apr 2012.
- [CV01] T. F. Chan and L. A. Vese, "Active contours without edges." *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–77, 2001. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18249617>
- [CWJ05] J. E. Cates, R. T. Whitaker, and G. M. Jones, "Case study: an evaluation of user-assisted hierarchical watershed segmentation." *Medical Image Analysis*, vol. 9, no. 6, pp. 566–578, 2005. [Online]. Available:



- <http://www.ncbi.nlm.nih.gov/pubmed/15919233>
- [Dam71] R. Damadian, "Tumor detection by nuclear magnetic resonance." *Science*, vol. 171, no. 976, pp. 1151–1153, 1971. [Online]. Available: <http://www.sciencemag.org/content/171/3976/1151.abstract>
- [Dan80] P. Danielsson, "Euclidean distance mapping," *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 227–248, 1980. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/0146664X80900544>
- [DB08] A. Delong and Y. Boykov, "A Scalable graph-cut algorithm for N-D grids," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, jun 2008, pp. 1–8.
- [Del94] H. Delingette, "Simplex meshes: a general representation for 3D shape reconstruction," in *Computer Vision and Pattern Recognition, 1994*, vol. 2214. IEEE Comput. Soc. Press, 1994, pp. 856–857. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=323913>
- [Del99] H. Delingette, "General Object Reconstruction based on Simplex Meshes," *International Journal of Computer Vision*, vol. 32, no. 2, pp. 111–146, 1999. [Online]. Available: <http://www.springerlink.com/index/h44l92r7n78u04t8.pdf>
- [DF01] A. V. Dias and S. S. Furuie, "Refinement of T-snakes using adaptive parameters," in *Proceedings XIV Brazilian Symposium on Computer Graphics and Image Processing*. IEEE Comput. Soc, 2001, p. 387. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=963093>
- [DFRVM14] A. Depeursinge, A. Foncubierta-Rodriguez, D. Van De Ville, and H. Müller, "Three-dimensional solid texture analysis in biomedical imaging: Review and opportunities." *Medical Image Analysis*, vol. 18, no. 1, pp. 176–96, jan 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24231667>
- [Dic45] L. R. Dice, "Measures of the amount of ecologic association between species," *Ecology*, vol. 26, no. 3, pp. 297–302, 1945. [Online]. Available: <http://www.jstor.org/stable/1932409>
- [Dix84] W. T. Dixon, "Simple proton spectroscopic imaging." *Radiology*, vol. 153, no. 1, pp. 189–194, 1984. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/6089263>
- [DJG00] M.-P. Dubuisson-Jolly and a. Gupta, "Color and texture fusion: application to aerial image segmentation and GIS updating," *Image and Vision Computing*, vol. 18, no. 10, pp. 823–832, jul 2000. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0262885699000505>
- [DKY95] R. Durikovic, K. Kaneda, and H. Yamashita, "Dynamic contour : a texture approach and contour operations," *The Visual Computer*, vol. 11, no. 6, pp. 277–289, 1995.
- [DLM<sup>+</sup>09] T. M. Deserno, W. Liao, M. C. Metzger, M. E. H. Kleiner, and D. Schulze, "Ein routine-integrierbares Planungswerkzeug zur operativen Rekonstruktion der Orbita," *GMS Medizinische Informatik, Biometrie und Epidemiologie*, vol. 5, no. 3, pp. 1–9, 2009.
- [DM00] D. DeCarlo and D. Metaxas, "Optical Flow Constraints on Deformable Models with Applications to Face Tracking," *International Journal of Computer Vision*, vol. 38, no. 2, pp. 99–127, 2000. [Online]. Available: <http://www.springerlink.com/index/T627164Q542G0857.pdf>
- [dMMoL<sup>+</sup>] E. des Mines de Nantes, MIRALab, T. U. of Lisbon, U. de las Islas Baleares, and U. of Karlsruhe, "CHARM: A Comprehensive Human Animation Resource Model, EC ESPRIT project 9036."
- [Doi06] K. Doi, "Diagnostic imaging over the last 50 years: research and development in medical imaging science and technology." *Physics in medicine and biology*, vol. 51, no. 13, pp. R5–R27, 2006.
- [EDFL13] A. Eklund, P. Dufort, D. Forsberg, and S. M. Laconte, "Medical image processing on the GPU - Past, present and future." *Medical image analysis*, vol. 17, no. 8, pp. 1073–94, dec 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23906631>
- [EF90] R. L. Ehman and J. P. Felmlee, "Flow artifact reduction in MRI: A review of the roles of gradient moment nulling and spatial presaturation," *Magnetic Resonance in Medicine*, vol. 14, no. 2, pp. 293–307, may 1990. [Online]. Available: <http://dx.doi.org/10.1002/mrm.1910140214>
- [EFV07] A. Efrat, Q. Fan, and S. Venkatasubramanian, "Curve matching, time warping, and light fields: New algorithms for computing similarity between curves," *Journal of Mathematical Imaging and Vision*, vol. 27, no. 3, pp. 203–216, 2007.
- [EGS<sup>+</sup>11] A. Elnakib, G. Gimel, J. J. Suri, A. El-baz, and G. Gimel'farb, "Medical Image Segmentation: A Brief Survey," in *Medical Image Segmentation*, A. S. El-Baz, R. Acharya U, A. F. Laine, and J. S. Suri, Eds. New York, NY: Springer New York, 2011, pp. 1–39. [Online]. Available:

- <http://www.springerlink.com/index/10.1007/978-1-4419-8204-9>
- [ER00] R. Everson and S. Roberts, "Inferring the eigenvalues of covariance matrices from limited, noisy data," *IEEE Transactions on Signal Processing*, vol. 48, no. 7, pp. 2083–2091, 2000. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=847792>
- [ES03] P. Elayaraja and M. Suganthi, "Survey on Medical Image Segmentation Methods," *International Journal of Advanced Research in Computer and Communication Engineering*, vol. 3, no. 11, pp. 3–8, 2003.
- [ESS12] M. Erdt, S. Steger, and G. Sakas, "Regmentation: A New View of Image Segmentation and Registration," *Journal of Radiation Oncology Informatics*, vol. 4, no. 1, 2012. [Online]. Available: <http://ojs.jroi.org/index.php/jroi/article/view/19>
- [EWG<sup>+</sup>97] M. A. Elliott, G. A. Walter, H. Gulish, A. S. Sadi, and D. D. Lawson, "Volumetric Measurement of Human Calf Muscle from MRI," *MAGMA*, vol. 5, no. 2, 1997.
- [EWRP13] J. M. Elliott, D. M. Walton, A. Rademaker, and T. B. Parrish, "Quantification of cervical spine muscle fat: a comparison between T1-weighted and multi-echo gradient echo imaging using a variable projection algorithm (VARPRO)." *BMC medical imaging*, vol. 13, no. 1, p. 30, jan 2013. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3847132&tool=pmcentrez&rendertype=abstract>
- [EWS<sup>+</sup>10] H. C. Edwards, A. B. Williams, G. D. Sjaardema, D. G. Baur, and W. K. Cochran, "SIERRA Toolkit Computational Mesh Conceptual Model," Sandia National Laboratories, Tech. Rep. March, 2010.
- [EZE08] N. El-Zehiry and A. Elmaghraby, "Graph cut based deformable model with statistical shape priors," in *2008 19th International Conference on Pattern Recognition*. Ieee, 2008, pp. 1–4. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4761726>
- [FBCO07] J. Fripp, P. Bourgeat, S. Crozier, and S. Ourselin, "Segmentation of the Bones in MRIs of the Knee Using Phase, Magnitude, and Shape Information," *Academic Radiology*, vol. 14, no. 10, pp. 1201–1208, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1076633207003443>
- [FBKC<sup>+</sup>12] A. Fedorov, R. Beichel, J. Kalphaty-Cramer, J. Finet, J.-C. Fillion-Robbin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, J. Buatti, S. Aylward, J. V. Miller, S. Pieper, and R. Kikinis, "3D Slicer as an Image Computing Platform for the Quantitative Imaging Network," *Magnetic resonance imaging*, vol. 30, no. 9, pp. 1323–1341, 2012.
- [FBW11] K.-I. Friese, P. Blanke, and F.-E. Wolter, "YaDiV—an open platform for 3D visualization and 3D segmentation of medical data," *The Visual Computer*, vol. 27, no. 2, pp. 129–139, 2011. [Online]. Available: <http://dx.doi.org/10.1007/s00371-010-0539-6>
- [FCWO05] J. Fripp, S. Crozier, S. Warfield, and S. Ourselin, "Automatic Initialization of 3D Deformable Models for Cartilage Segmentation," in *Digital Image Computing Techniques and Applications DICTA05*. IEEE Computer Society, 2005, pp. 513–518. [Online]. Available: <http://discovery.ucl.ac.uk/70778/>
- [Fec11] T. Fechter, "Deformation Based Manual Segmentation in Three and Four Dimensions," Diplomarbeit, Technische Universität Wien, 2011.
- [FH04] P. Felzenszwalb and D. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, pp. 167–181, 2004.
- [FHWCO1] M. J. Firbank, R. M. Harrison, E. D. Williams, and a. Coulthard, "Measuring extraocular muscle volume using dynamic contours." *Magnetic resonance imaging*, vol. 19, no. 2, pp. 257–65, mar 2001. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/11358663>
- [FJT94] K. J. Friston, P. Jezzard, and R. Turner, "Analysis of functional MRI time-series." *Human Brain Mapping*, vol. 1, no. 2, pp. 153–171, 1994. [Online]. Available: <http://doi.wiley.com/10.1002/hbm.460010207>
- [FPE<sup>+</sup>14] M. a. Fischer, C. W. a. Pfirrmann, N. Espinosa, D. a. Raptis, and F. M. Buck, "Dixon-based MRI for assessment of muscle-fat content in phantoms, healthy volunteers and patients with achillodynia: Comparison to visual assessment of calf muscle quality," *European Radiology*, vol. 24, no. 6, pp. 1366–1375, 2014.
- [FRM<sup>+</sup>12] M. A. Fischer, D. A. Raptis, M. Montani, R. Graf, P.-A. Clavien, D. Nanz, H. Alkadhi, and H. Scheffel, "Liver Fat Quantification by Dual-echo MR Imaging Outperforms Traditional Histopathological Analysis," *Academic Radiology*, vol. 19, no. 10, pp. 1208–1214, 2012. [Online].

- Available: <http://linkinghub.elsevier.com/retrieve/pii/S1076633212003066>
- [FVC<sup>+</sup>11] M. Filipovic, P. A. Vuissoz, A. Codreanu, M. Claudon, and J. Felblinger, "Motion compensated generalized reconstruction for free-breathing dynamic contrast-enhanced MRI," *Magnetic Resonance in Medicine*, vol. 65, no. 3, pp. 812–822, 2011.
- [FVS11] J. Fang, A. L. Varbanescu, and H. Sips, "A Comprehensive Performance Comparison of CUDA and OpenCL," *Parallel Processing (ICPP)*, pp. 216 – 225, 2011. [Online]. Available: <http://arxiv.org/abs/1005.2581>
- [FWCO06] J. Fripp, S. K. Warfield, S. Crozier, and S. Ourselin, "Automatic Segmentation of the Knee Bones using 3D Active Shape Models," in *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, 2006, pp. 167–170.
- [GC11] F. L. Goerner and G. D. Clarke, "Measuring signal-to-noise ratio in partially parallel imaging MRI." *Medical physics*, vol. 38, no. 9, pp. 5049–57, 2011. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3170395&tool=pmcentrez&rendertype=abstract>
- [GCM<sup>+</sup>11] E. Geremia, O. Clatz, B. H. Menze, E. Konukoglu, A. Criminisi, and N. Ayache, "Spatial decision forests for MS lesion segmentation in multi-channel magnetic resonance images." *NeuroImage*, vol. 57, no. 2, pp. 378–90, jul 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21497655>
- [GDTG13] P. Garteiser, S. Doblas, R. a. Towner, and T. M. Griffin, "Calibration of a semi-automated segmenting method for quantification of adipose tissue compartments from magnetic resonance images of mice." *Metabolism: clinical and experimental*, vol. 62, no. 11, pp. 1686–95, nov 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23890668>
- [GFFK15] C. Graca, G. Falcao, I. N. Figueiredo, and S. Kumar, "Hybrid multi-GPU computing: accelerated kernels for segmentation and object detection with medical image processing applications," *Journal of Real-Time Image Processing*, 2015. [Online]. Available: <http://link.springer.com/10.1007/s11554-015-0517-3>
- [GFS<sup>+</sup>02] B. V. Ginneken, A. F. Frangi, J. J. Staal, M. Bart, H. Romeny, M. A. Viergever, B. van Ginneken, and B. M. ter Haar Romeny, "Active Shape Model Segmentation with Optimal Features," *Medical Imaging, IEEE Transactions on*, vol. 21, no. 8, pp. 924–933, 2002.
- [GFW13] E. Gibson, A. Fenster, and A. D. Ward, "The impact of registration accuracy on imaging validation study design: A novel statistical power calculation." *Medical image analysis*, vol. 17, no. 7, pp. 805–815, apr 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23706752>
- [GHR<sup>+</sup>14] P. Gargiulo, T. Helgason, C. Ramon, H. Jónsson jr, and U. Carraro, "CT and MRI assessment and characterization using segmentation and 3D modeling techniques: applications to muscle, bone and brain," *European Journal of Translational Myology*, vol. 24, no. 1, pp. 55–62, mar 2014. [Online]. Available: <http://pagepressjournals.org/index.php/bam/article/view/bam.2014.1.55>
- [Gil07] B. Gilles, "Anatomical and Kinematical Modelling of the Musculoskeletal System from MRI," Ph.D. dissertation, University of Geneva, aug 2007.
- [GKBMT96] Gingins, P. Kalra, P. Beylot, and N. Magnenat-Thalmann, "Using VHD to build a comprehensive human model," in *The Visible Human Project Conf.*, 1996. [Online]. Available: <ftp://swanson.seas.upenn.edu/pub/dnm/siggr99/nadia/paper2.ps.gz>
- [GMC<sup>+</sup>10] E. Geremia, B. H. Menze, O. Clatz, E. Konukoglu, A. Criminisi, and N. Ayache, "Spatial decision forests for MS lesion segmentation in multi-channel MR images." *Medical image computing and computer-assisted intervention : MICCAI International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 13, no. Pt 1, pp. 111–8, jan 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21497655>
- [GMMT06] B. Gilles, L. Moccozet, and N. Magnenat-Thalmann, "Anatomical modelling of the musculoskeletal system from MRI," in *Proc. International conference on medical image computing and computer assisted intervention (MICCAI'06)*, vol. 4190, oct 2006, pp. 289–296.
- [GMT10] B. Gilles and N. Magnenat-Thalmann, "Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations," *Medical Image Analysis*, vol. 14, no. 3, pp. 291–302, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841510000150>
- [Gow75] J. C. Gower, "Generalized procrustes analysis," *Psychometrika*, vol. 40, no. 1, pp. 33–51, 1975. [Online].

- Available: <http://www.springerlink.com/index/10.1007/BF02291478>
- [GPP<sup>+</sup>14] S. Gupta, D. Palsetia, M. A. Patwary, A. Agrawal, and A. Choudhary, "A New Parallel Algorithm for Two-Pass Connected Component Labeling," in *Parallel & Distributed Processing Symposium Workshops (IPDPSW), 2014 IEEE International*. IEEE, 2014.
- [Gra05] L. Grady, "Multilabel Random Walker Image Segmentation Using Prior Models," *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1, pp. 763–770, 2005. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1467345>
- [GSKP11] B. Glocker, A. Sotiras, N. Komodakis, and N. Paragios, "Deformable medical image registration: setting the state of the art with discrete methods." *Annual review of biomedical engineering*, vol. 13, pp. 219–44, aug 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21568711>
- [Guo06] Y. Guo, "Medical Image Registration and Application to Atlas-Based Segmentation," Ph.D. dissertation, Kent State University, 2006.
- [GWZ09] B. Glocker, C. Wachinger, and J. Zeltner, "MRI Composing for Whole Body Imaging," *Bildverarbeitung für die Medizin*, 2009. [Online]. Available: [http://link.springer.com/chapter/10.1007/978-3-540-93860-6\\_85](http://link.springer.com/chapter/10.1007/978-3-540-93860-6_85)
- [HBH<sup>+</sup>12] Z. L. Husz, N. Burton, B. Hill, N. Milyaev, and R. a. Baldock, "Web tools for large-scale 3D biological images and atlases," *BMC Bioinformatics*, vol. 13, no. 1, p. 122, 2012. [Online]. Available: <http://www.biomedcentral.com/1471-2105/13/122?fmt{view=mobile>
- [HD11] T. Heimann and H. Delingette, "Model-Based Segmentation," in *Biomedical Image Processing*, ser. Biological and Medical Physics, Biomedical Engineering, T. M. Deserno, Ed. Springer Berlin Heidelberg, 2011, pp. 279–303. [Online]. Available: [http://dx.doi.org/10.1007/978-3-642-15816-2\\_11](http://dx.doi.org/10.1007/978-3-642-15816-2_11)
- [HJB<sup>+</sup>12] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Martin, F. V. Gleeson, S. M. Brady, and J. A. Schnabel, "MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration," *Medical Image Analysis*, vol. 16, no. 7, pp. 1423–1435, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361841512000643>
- [HK13] H. H. Hu and H. E. Kan, "Quantitative proton MR techniques for measuring fat." *NMR in biomedicine*, vol. 26, no. 12, pp. 1609–29, dec 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24123229>
- [HKB<sup>+</sup>06] S. S. Hong, A. Y. Kim, J. H. Byun, H. J. Won, P. N. Kim, M. G. Lee, and H. K. Ha, "MDCT of small-bowel disease: Value of 3D imaging," *American Journal of Roentgenology*, vol. 187, no. 5, pp. 1212–1221, 2006.
- [HKH10] R. Hegadi, A. Kop, and M. Hangarge, "A Survey on Deformable Model and its Applications to Medical Imaging," *IJCA, Special Issue on RTIPPR*, pp. 64–75, 2010.
- [HLP10] K. Hawick, A. Leist, and D. Playne, "Parallel graph component labelling with GPUs and CUDA," *Parallel Computing*, pp. 0–30, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167819110001055>
- [HMLM<sup>+</sup>] HUGE, MIRALab, LBB-MHH, CNR-IMATI, UMINHO, Softeco, and LUH, "MultiScaleHuman project." [Online]. Available: <http://multiscalehuman.miralab.ch/>
- [HMS<sup>+</sup>10] T. Heimann, B. Morrison, M. Styner, N. M., and W. S., "Segmentation of knee images: A grand challenge," in *Proc. MICCAI Workshop on Medical Image Analysis for the Clinic*, 2010, pp. 207–214.
- [HNS<sup>+</sup>10] S. Han, N. A. Nijdam, J. Schmid, J. Kim, and N. Magnenat-Thalmann, "Collaborative telemedicine for interactive multiuser segmentation of volumetric medical images," *The Visual Computer*, vol. 26, no. 6-8, pp. 639–648, apr 2010. [Online]. Available: <http://link.springer.com/10.1007/s00371-010-0445-y>
- [Hou62] P. V. C. Hough, "Method and means for recognizing complex patterns," pp. 225–231, 1962.
- [HP91] B. Horowitz and A. Pentland, "Recovery of non-rigid motion and structure," *Proceedings. 1991 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 325–330, 1991. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=139710>
- [HPN99] B. Heigl, D. Paulus, and H. Niemann, "Tracking points in sequences of color images," *Pattern Recognition and Image Understanding*, no. September, pp. 1–4, 1999. [Online]. Available: <http://scholar.google.com/scholar?hl=en{&}btnG=Search{&}q=intitle:Tracking+Points+in+Sequences+of+Color+Images{#}0>
- [HRA<sup>+</sup>14] D. Haehn, N. Rannou, B. Ahtam, E. Grant, and R. Pienaar, "Neuroimaging in the Browser

- using the X Toolkit," *Frontiers in Neuroinformatics*, no. 101, 2014. [Online]. Available: <http://www.frontiersin.org/neuroinformatics/10.3389/conf.fninf.2014.08.00101/full>
- [HSZ87] R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image analysis using mathematical morphology," *IEEE transactions on pattern analysis and machine intelligence*, vol. PAMI-9, no. 4, pp. 532–550, 1987. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=4767941](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=4767941)
- [HW78] R. P. Henderson and J. G. Webster, "An Impedance Camera for Spatially Specific Measurements of the Thorax," *IEEE Transactions on Biomedical Engineering*, vol. 25, no. 3, pp. 250–254, 1978.
- [HWP<sup>+</sup>11] B. A. Hargreaves, P. W. Worters, K. B. Pauly, J. M. Pauly, K. M. Koch, and G. E. Gold, "Metal-induced artifacts in MRI." *Ajr American Journal Of Roentgenology*, vol. 197, no. 3, pp. 547–555, 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21862795>
- [HWWL14] L. He, Y. Wen, M. Wan, and S. Liu, "Multi-channel features based automated segmentation of diffusion tensor imaging using an improved FCM with spatial constraints," *Neurocomputing*, feb 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0925231214002616>
- [HYML05] G. Hamarneh, J. Yang, C. McIntosh, and M. Langille, "3D live-wire-based semi-automatic segmentation of medical images," in *Medical Imaging*, J. M. Fitzpatrick and J. M. Reinhardt, Eds., vol. 5747, apr 2005, pp. 1597–1603. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.596148><http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1282186>
- [IdFFdF96] R. Ierusalimsky, L. H. de Figueiredo, W. C. Filho, and L. H. de Figueiredo, "Lua - An Extensible Extension Language," *Software: Practice and Experience*, vol. 26, no. 6, pp. 635–652, 1996.
- [IGTG<sup>+</sup>14] N. Imamoglu, J. Gomez-Tames, J. Gonzalez, D. Gu, and W. Yu, "Pulse-Coupled Neural Network Segmentation and Bottom-Up Saliency-On Feature Extraction for Thigh Magnetic Resonance Imaging Based 3D Model Construction," *Journal of Medical Imaging and Health Informatics*, vol. 4, no. 2, pp. 220–229, 2014. [Online]. Available: <http://www.ingentaconnect.com/content/asp/jmihi/2014/00000004/00000002/art00008>
- [IGTH<sup>+</sup>15] N. Imamoglu, J. Gomez-Tames, S. He, D. Y. Gu, K. Kita, and W. Yu, "Unsupervised muscle region extraction by fuzzy decision based saliency feature integration on thigh MRI for 3D modeling," *Proceedings of the 14th IAPR International Conference on Machine Vision Applications, MVA 2015*, pp. 150–153, 2015.
- [IJDAB14] V. Israel-Jost, J. Darbon, E. D. Angelini, and I. Bloch, "on the Implementation of the Multi-Phase Region Segmentation , Solving the Hidden Phase Problem," in *Image Processing (ICIP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 4338–4342.
- [IJKY06] M. Ibrahim, N. John, M. Kabuka, and A. Younis, "Hidden Markov models-based 3D MRI brain segmentation," *Image and Vision Computing*, vol. 24, no. 10, pp. 1065–1079, oct 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0262885606001107>
- [IK88] J. Illingworth and J. Kittler, "A Survey of the Hough Transform," *Comp Vision, Graphics, and Image Processing*, vol. 44, pp. 87–116, 1988.
- [IP94] J. Ivins and J. Porrill, "Statistical Snakes : Active Region Models," *Proceedings of the British Machine Vision Conference*, pp. 377–386, 1994.
- [IS98] H. H. S. Ip and D. Shen, "An affine-invariant active contour model (AI-snake) for model-based segmentation," *Image and Vision Computing*, vol. 16, pp. 135–146, 1998.
- [Jam07] G. M. James, "Curve alignment by moments," *The Annals of Applied Statistics*, vol. 1, no. 2, pp. 480–501, 2007. [Online]. Available: <http://projecteuclid.org/euclid.aoas/1196438028>
- [JDR<sup>+</sup>14] E. Jolivet, E. Dion, P. Rouch, G. Dubois, R. Charrier, C. Payan, and W. Skalli, "Skeletal muscle segmentation from MRI dataset using a model-based approach," in *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, apr 2014, pp. 1–8. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1080/21681163.2013.855146>
- [JFE<sup>+</sup>07] W. Jentzen, L. Freudenberg, E. G. Eising, M. Heinze, W. Brandau, and A. Bockisch, "Segmentation of PET volumes by iterative image thresholding." *Journal of nuclear medicine : official publication, Society of Nuclear Medicine*, vol. 48, no. 1, pp. 108–114, 2007.
- [JGSL97] E. F. Jackson, L. E. Ginsberg, D. F. Schomer, and N. E. Leeds, "A review of MRI pulse sequences and techniques in neuroimaging," *Surgical Neurology*, vol. 47, no. 2, pp. 185–199, 1997.

- [JIL<sup>+</sup>03] Y. Jin, C. Z. Imieli, A. F. Laine, J. Udupa, W. Shen, and S. B. Heymsfield, "Segmentation and Evaluation of Adipose Tissue from Whole Body MRI Scans," in *MICCAI*, 2003, pp. 635–642.
- [JS12] S. Jansi and P. Subashini, "Optimized Adaptive Thresholding based Edge Detection Method for MRI Brain Images," *International Journal of Computer Applications*, vol. 51, no. 20, pp. 1–8, 2012.
- [JSR<sup>+</sup>08] K. U. Juergens, H. Seifarth, F. Range, S. Wienbeck, M. Wenker, W. Heindel, and R. Fischbach, "Automated threshold-based 3D segmentation versus short-axis planimetry for assessment of global left ventricular function with dual-source MDCT," *American Journal of Roentgenology*, vol. 190, no. 2, pp. 308–314, 2008.
- [JTD<sup>+</sup>14] D. G. Juan, S. Trombella, B. M. a. Delattre, Y. Seimbille, and O. Ratib, "Study of skeletal muscle behavior by PET/MRI," *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 61–64, jun 2014. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6864304>
- [JWN<sup>+</sup>08] 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, K. Klingel, G. Reischl, S. Widmaier, M. Rocken, R. E. Nutt, H.-J. Machulla, K. Uludag, S. R. Cherry, C. D. Claussen, and B. J. Pichler, "Simultaneous PET-MRI: a new approach for functional and morphological imaging," *Nat Med*, vol. 14, no. 4, pp. 459–465, apr 2008. [Online]. Available: <http://dx.doi.org/10.1038/nm1700>
- [JZD]98] A. K. Jain, Y. Zhong, and M.-P. Dubuisson-Jolly, "Deformable template models: A review," *Signal Processing*, vol. 71, no. 2, pp. 109–129, dec 1998. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S016516849800139X>
- [KAJF07] J. Kullberg, H. Ahlström, L. Johansson, and H. Frimmel, "Automated and reproducible segmentation of visceral and subcutaneous adipose tissue from abdominal MRI," *International journal of obesity (2005)*, vol. 31, no. 12, pp. 1806–17, dec 2007. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17593903>
- [Kal06] W. A. Kalender, "X-ray computed tomography," *Physics in medicine and biology*, vol. 51, no. 13, pp. R29–R43, 2006.
- [Kau14] G. Kauffman, "Nobel Prize for MRI Imaging Denied to Raymond V. Damadian a Decade Ago," *The Chemical Educator*, vol. 4171, no. 14, pp. 73–90, 2014.
- [KBFW07] Y.-J. Kim, T. Brox, W. Feiden, and J. Weickert, "Fully automated segmentation and morphometrical analysis of muscle fiber images," *Cytometry Part A*, vol. 71A, no. 1, pp. 8–15, 2007. [Online]. Available: <http://dx.doi.org/10.1002/cyto.a.20334>
- [KBG<sup>+</sup>95] P. Kalra, P. Beylot, P. Gingsins, N. Magnenat-Thalmann, and P. Volino, "Topological modeling of human anatomy using medical data," *Computer Animation' ...*, 1995. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=393534](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=393534)
- [KDH11] K. Karimi, N. G. Dickson, and F. Hamze, "A Performance Comparison of CUDA and OpenCL," *arXiv:1005.2581v3*, no. 1, pp. 216 – 225, 2011. [Online]. Available: <http://arxiv.org/abs/1005.2581>
- [KGG<sup>+</sup>03] C. Kauffmann, P. Gravel, B. Godbout, A. Gravel, G. Beaudoin, J.-P. Raynauld, J. Martel-Pelletier, J.-P. Pelletier, and J. A. de Guise, "Computer-aided method for quantification of cartilage thickness and volume changes using MRI: validation study using a synthetic model," *Biomedical Engineering, IEEE Transactions on*, vol. 50, no. 8, pp. 978–988, 2003.
- [KGG11] N. Kanwal, A. Girdhar, and S. Gupta, "Region Based Adaptive Contrast Enhancement of Medical X-Ray Images," *2011 5th International Conference on Bioinformatics and Biomedical Engineering*, pp. 1–5, 2011.
- [KGZS11] C. Kober, L. Gallo, H.-F. Zeilhofer, and R. A. Sader, "Computer-assisted analysis of human upper arm flexion by 4D-visualization based on MRI," *International journal of computer assisted radiology and surgery*, vol. 6, no. 5, pp. 675–84, sep 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21350967>
- [KHAK10] M. Kasap, M. Holl, A. Aksay, and P. Kelly, "3D Realistic Animation of a Tennis Player," in *Proceedings of ENGAGE summer school*, 2010, pp. 1–13.
- [Kho07] K. Khoshelham, "Extending Generalized Hough Transform to Detect 3D Objects in Laser Range Data," *ISPRS Workshop on Laser Scanning and SilviLaser 2007, 12-14 September 2007, Espoo, Finland*, pp. 206–210, 2007.

- [Kim11] Y. Kimori, "Mathematical morphology-based approach to the enhancement of morphological features in medical images." *Journal of clinical bioinformatics*, vol. 1, no. 1, p. 33, jan 2011. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3275547&tool=pmcentrez&rendertype=abstract>
- [KKGv01] V. a. Kovalev, F. Kruggel, H. J. Gertz, and D. Y. Von Cramon, "Three-dimensional texture analysis of MRI brain datasets," *IEEE Transactions on Medical Imaging*, vol. 20, no. 5, pp. 424–433, 2001.
- [KLH<sup>+</sup>12] D. Kainmueller, H. Lamecker, M. Heller, B. Weber, H.-C. Hege, and S. Zachow, "Omnidirectional Displacements for Deformable Surfaces," *Medical Image Analysis*, vol. 17, no. 4, pp. 429–441, 2012.
- [KLZ<sup>+</sup>08] D. Kainmueller, H. Lamecker, S. Zachow, M. Heller, and H.-C. Hege, "Multi-object Segmentation with Coupled Deformable Models," in *Proceedings of the 12th Annual Conference on Medical Image Understanding and Analysis 2008*, 2008.
- [KMT11] M. Kasap and N. Magnenat-Thalmann, "Skeleton-aware size variations in digital mannequins," *The Visual Computer*, vol. 27, no. 4, pp. 263–274, feb 2011. [Online]. Available: <http://www.springerlink.com/index/10.1007/s00371-011-0547-1>
- [KMZ<sup>+</sup>13] N. Kamiya, C. Muramatsu, X. Zhou, H. Chen, R. Yokoyama, and T. Hara, "Model-based approach to recognize the rectus abdominis muscle in CT by use of a virtual image-unfolding technique," *IEICE TRANSACTIONS on Information and Systems*, no. 4, pp. 2–3, 2013.
- [KPG08] F. Kruggel, J. S. Paul, and H.-J. Gertz, "Texture-based segmentation of diffuse lesions of the brain's white matter," *NeuroImage*, vol. 39, no. 3, pp. 987–996, 2008. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1053811907008774>
- [KPV<sup>+</sup>07] H. Kang, A. Pinti, L. Vermeiren, A. Taleb-Ahmed, and X. Zeng, "An automatic FCM-based method for tissue classification: Application to MRI of thigh," in *2007 1st International Conference on Bioinformatics and Biomedical Engineering, ICBBE, 2007*, pp. 510–514.
- [KRKS11] O. Kalentev, A. Rai, S. Kemnitz, and R. Schneider, "Connected component labeling on a 2D grid using CUDA," *Journal of Parallel and Distributed Computing*, vol. 71, no. 4, pp. 615–620, 2011.
- [KSA<sup>+</sup>12] T. Kohlberger, V. Singh, C. Alvino, C. Bahlmann, and L. Grady, "Evaluating segmentation error without ground truth." *Medical image computing and computer-assisted intervention : MICCAI, International Conference on*, vol. 15, no. Pt 1, pp. 528–36, jan 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23285592>
- [KSP07] S. Klein, M. Staring, and J. P. W. Pluim, "Evaluation of Optimization Methods for Nonrigid Medical Image Registration Using Mutual Information and B-Splines," *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2879–2890, 2007. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18092588>
- [KTL15] S. S. Kanitkar, N. D. Thombare, and S. S. Lokhande, "Detection of Lung Cancer Using Marker-Controlled Watershed Transform," in *2015 International Conference on Pervasive Computing (ICPC)*, 2015.
- [KWT88] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes : Active Contour Models," *International Journal of Computer Vision*, vol. 331, no. 4, pp. 321–331, 1988. [Online]. Available: <http://dx.doi.org/10.1007/BF00133570>
- [Lak01] V. Lakshmanan, "A hierarchical, multiscale texture segmentation algorithm for real-world scenes," PhD, University of Oklahoma, 2001. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.203.763&rep=rep1&type=pdf>
- [LBK<sup>+</sup>13] R. S. Livingstone, P. Begovatz, S. Kahl, B. Nowotny, K. Straßburger, G. Giani, J. Bunke, M. Roden, and J.-H. Hwang, "Initial clinical application of modified Dixon with flexible echo times: hepatic and pancreatic fat assessments in comparison with (1)H MRS." *Magma (New York, N.Y.)*, dec 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24306514>
- [LBS03] T. M. Lehmann, J. Bredno, and K. Spitzer, "On the design of active contours for medical image segmentation. A Scheme for Classification and construction." *Methods of information in medicine*, vol. 42, no. 1, pp. 89–98, jan 2003. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/12695800>
- [LC87] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," *Computer*, vol. 21, no. 4, pp. 163–169, 1987. [Online]. Available: <http://portal.acm.org/>

- citation.cfm?id=37422
- [LDP14] C. Lachat, C. Dobrzynski, and F. Pellegrini, "Parallel Mesh Adaptation Using Parallel Graph Partitioning," in *Proc. of the 11th World Congress on Computational Mechanics*, 2014, pp. 1–12.
- [Lew04] R. a. Lewis, "Medical phase contrast x-ray imaging: current status and future prospects." *Physics in medicine and biology*, vol. 49, no. 16, pp. 3573–3583, 2004.
- [LHG<sup>+</sup>96] D. Levine, H. Hatabu, J. Gaa, M. Wendy Atkinson, and R. R. Edelman, "Fetal anatomy revealed with fast MR sequences," pp. 905–908, 1996.
- [LHZD13] L. Li, W. Hou, X. Zhang, and M. Ding, "GPU-based block-wise nonlocal means denoising for 3D ultrasound images," *Computational & Mathematical Methods in Medicine*, vol. 2013, p. 921303, 2013.
- [LJR<sup>+</sup>08] O. Leinhard, A. Johansson, J. Rydell, Ö. Smedby, F. Nyström, P. Lundberg, and M. Borga, "Quantitative abdominal fat estimation using MRI," in *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, Tampa, FL, 2008, pp. 1–4. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=4761764](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=4761764)
- [LKW<sup>+</sup>07] H. Lamecker, L. Kamer, A. Wittmers, S. Zachow, A. Schramm, H. Noser, and B. Hammer, "A method for the three-dimensional statistical shape analysis of the bony orbit," *Proc. Computer Aided Surgery around the Head (CAS-H)*, pp. 2–5, 2007. [Online]. Available: <http://www.zib.de/lamecker/publications/cash2007.pdf>
- [LL93] F. Leymarie and M. D. Levine, "Tracking deformable objects in the plane using an active contour model," in *Proceedings of IEEE International Conference on Multimedia Computing and Systems*, vol. 15, 1993, pp. 617–634. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=216733>
- [LM99] J. O. Lachaud and A. Montanvert, "Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction." *Medical Image Analysis*, vol. 3, no. 2, pp. 187–207, 1999. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/10711998>
- [LMHF14] Liuyang Fang, Mi Wang, Hexiang Ying, and Fen Hu, "Multi-GPU based near real-time preprocessing and releasing system of optical satellite images," in *2014 IEEE Geoscience and Remote Sensing Symposium*. IEEE, jul 2014, pp. 2467–2470. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6946972>
- [LMT08] E. Lyard and N. Magnenat-Thalmann, "Motion adaptation based on character shape," *Computer Animation and Virtual Worlds*, vol. 19, no. 3-4, pp. 189–198, 2008.
- [Lom04] J.-C. Lombardo, "Modélisation d'objets déformables avec un système de particules orientées," PhD, Université Joseph Fourier, Grenoble, 2004.
- [Low99] D. G. Lowe, "Object recognition from local scale-invariant features," *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, no. [8], pp. 1150–1157, 1999. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=790410>
- [LSP<sup>+</sup>10] S. Lee, H. Shim, S. H. Park, I. D. Yun, and S. U. Lee, "Learning Local Shape and Appearance for Segmentation of Knee Cartilage in 3D MRI," in *Proceedings of the 4th Medical Image Analysis for the Clinic - A Grand Challenge workshop, MICCAI*, 2010.
- [LWS<sup>+</sup>12] F. H. Lin, T. Witzel, G. Schultz, D. Gallichan, W. J. Kuo, F. N. Wang, J. Hennig, M. Zaitsev, and J. W. Belliveau, "Reconstruction of MRI data encoded by multiple nonbijective curvilinear magnetic fields," *Magnetic Resonance in Medicine*, vol. 68, no. 4, pp. 1145–1156, 2012.
- [MADC08] M. Marchal, J. Allard, C. Duriez, and S. Cotin, "Towards a Framework for Assessing Deformable Models in Medical Simulation," in *Biomedical Simulation*, ser. Lecture Notes in Computer Science, F. Bello and P. Edwards, Eds. Springer Berlin / Heidelberg, 2008, vol. 5104, pp. 176–184. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-70521-5\\_19](http://dx.doi.org/10.1007/978-3-540-70521-5_19)
- [Mar98] J. S. Marques, "A fuzzy algorithm for curve and surface alignment," *Pattern Recognition Letters*, vol. 19, no. 9, pp. 797–803, 1998. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0167865598000634>
- [Mar11] C. Marin, "WebGL Specification," Khronos, Tech. Rep., 2011. [Online]. Available: <https://www.khronos.org/registry/webgl/specs/1.0/>
- [MB95] E. N. Mortensen and W. A. Barrett, "Intelligent scissors for image composition," *Proceedings of the 22nd*



- annual conference on Computer graphics and interactive techniques SIGGRAPH 95*, vol. 84602, no. 801, pp. 191–198, 1995. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=218380.218442>
- [MB12] G. R. Milner and J. L. Boldsen, “Humeral and femoral head diameters in recent white American skeletons.” *Journal of forensic sciences*, vol. 57, no. 1, pp. 35–40, jan 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22074229>
- [MBL<sup>+</sup>91] J. V. Miller, D. E. Breen, W. E. Lorensen, R. M. O’Bara, and M. J. Wozny, “Geometrically deformed models: a method for extracting closed geometric models form volume data,” *SIGGRAPH Comput Graph*, vol. 25, no. 4, pp. 217–226, 1991. [Online]. Available: <http://portal.acm.org/citation.cfm?id=122742>
- [MBM<sup>+</sup>12] A. I. Mitrea, R. Badea, D. Mitrea, S. Nedeveschi, P. Mitrea, D. M. Ivan, and O. M. Gurzău, “Iterative methods for obtaining energy-minimizing parametric snakes with applications to medical imaging.” *Computational and mathematical methods in medicine*, vol. 2012, p. 918510, jan 2012. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3303766&tool=pmcentrez&rendertype=abstract>
- [MCB<sup>+</sup>06] M. Monziols, G. Collewet, M. Bonneau, F. Mariette, a. Davenel, and M. Kouba, “Quantification of muscle, subcutaneous fat and intermuscular fat in pig carcasses and cuts by magnetic resonance imaging.” *Meat science*, vol. 72, no. 1, pp. 146–54, jan 2006. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22061385>
- [MD05] J. Montagnat and H. Delingette, “4D Deformable Models with Temporal Constraints: Application to 4D Cardiac Image Segmentation,” *Medical Image Analysis*, vol. 9, no. 1, pp. 87–100, feb 2005. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/15581814><http://www.sciencedirect.com/science/article/pii/S1361841504000623>
- [MDA01] J. Montagnat, H. Delingette, and N. Ayache, “A review of deformable surfaces: topology, geometry and deformation,” *Image and Vision Computing*, vol. 19, no. 14, pp. 1023–1040, dec 2001. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0262885601000646>
- [MFC<sup>+</sup>06] J. P. Mattei, Y. L. Fur, N. Cuge, S. Guis, P. J. Cozzzone, and D. Bendahan, “Segmentation of fascias, fat and muscle from magnetic resonance images in humans: the DISPIMAG software.” *Magma (New York, N.Y.)*, vol. 19, no. 5, pp. 275–9, nov 2006. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17004065>
- [MHHR11] H. Mirzaalian, G. Hamarneh, B. HajGhanbari, and W. Reid, “3D Shape Analysis of the Knee Extensor and Flexor Muscles in Patients with COPD using Mesh Projection-based Features,” *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2011*, 2011. [Online]. Available: [http://www.cs.sfu.ca/~hamarneh/ecopy/miccai\[\\_\]meshmed2011.pdf](http://www.cs.sfu.ca/~hamarneh/ecopy/miccai[_]meshmed2011.pdf)
- [MHKR<sup>+</sup>15] L. Maier-Hein, D. Kondermann, T. Roß, S. Mersmann, E. Heim, S. Bodenstedt, H. G. Kenngott, A. Sanchez, M. Wagner, A. Preukschas, A.-L. Wekerle, S. Helfert, K. März, A. Mehrabi, S. Speidel, and C. Stock, “Crowdtruth validation: a new paradigm for validating algorithms that rely on image correspondences,” *International Journal of Computer Assisted Radiology and Surgery*, pp. 1201–1212, 2015. [Online]. Available: <http://link.springer.com/10.1007/s11548-015-1168-3>
- [MLM<sup>+</sup>05] U. Meier, O. López, C. Monserrat, M. C. Juan, and M. Alcañiz, “Real-time deformable models for surgery simulation: a survey.” *Computer Methods and Programs in Biomedicine*, vol. 77, no. 3, pp. 183–197, 2005. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/15721648>
- [MM07] P. Moore and D. Molloy, “A Survey of Computer-Based Deformable Models,” *International Machine Vision and Image Processing Conference IMVIP 2007*, pp. 55–66, 2007. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4318138>
- [MM15] M. M. Myint and T. Myint, “Effective Kidney Segmentation Using Gradient Based Approach in Abdominal CT Images,” in *Proceedings of 2015 International Conference on Future Computational Technologies (ICFCT’2015)*, 2015, pp. 130–135.
- [MR03] C. Maurer and V. Raghavan, “A linear time algorithm for computing exact Euclidean distance transforms of binary images in arbitrary dimensions,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 2, pp. 265–270, 2003. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1177156>

- [MS89] D. H. U. Mumford and J. N. U. Shah, "Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems," *Communications on Pure and Applied Mathematics*, vol. XLII, pp. 577–685, 1989.
- [MS12] F. Malmberg and R. Strand, "Smart Paint-A New Interactive Segmentation Method Applied to MR Prostate Segmentation," ... *MR Image Segmentation*, 2012. [Online]. Available: <http://www.cb.uu.se/~filip/papers/SmartPaint.pdf>
- [MSB96] R. Malladi, J. A. Sethian, and L. Berkeley, "Image Processing: Flows under Min/Max Curvature and Mean Curvature," *Graphical Models and Image Processing*, vol. 58, no. 2, pp. 127–141, 1996. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1077316996900118>
- [MSF<sup>+</sup>12] S. Makrogiannis, S. Serai, K. W. Fishbein, C. Schreiber, L. Ferrucci, and R. G. Spencer, "Automated quantification of muscle and fat in the thigh from water-, fat-, and nonsuppressed MR images." *Journal of magnetic resonance imaging : JMRI*, vol. 35, no. 5, pp. 1152–61, may 2012. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3319811&tool=pmcentrez&rendertype=abstract>
- [MSK04] M. Mueller, K. Segl, and H. Kaufmann, "Edge- and region-based segmentation technique for the extraction of large, man-made objects in high-resolution satellite imagery," *Pattern Recognition*, vol. 37, no. 8, pp. 1619–1628, 2004.
- [MT95] T. McInerney and D. Terzopoulos, "A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis." *Computerized Medical Imaging and Graphics*, vol. 19, no. 1, pp. 69–83, 1995. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/7736420>
- [MT96] T. McInerney and D. Terzopoulos, "Deformable Models in Medical Image Analysis: A Survey," in *Mathematical Methods in Biomedical Image Analysis, 1996., Proceedings of the Workshop on*, vol. 1. Elsevier, 1996, pp. 171–180. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/9873923>
- [MT00] T. McInerney and D. Terzopoulos, "T-snakes: topology adaptive snakes." *Medical Image Analysis*, vol. 4, no. 2, pp. 73–91, 2000. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/10972323>
- [MTCS08] N. Magnenat-Thalmann, C. Charbonnier, and J. Schmid, "Multimedia application to the simulation of human musculoskeletal system: A visual lower limb model from multimodal captured data," in *Proceedings of the 2008 IEEE 10th Workshop on Multimedia Signal Processing, MMSP 2008*, 2008, pp. 520–525.
- [MTSC04] N. Magnenat-Thalmann, H. Seo, and F. Cordier, "Automatic modeling of virtual humans and body clothing," *Journal of Computer Science and Technology*, vol. 19, no. 5, pp. 575–584, 2004.
- [MTWT98] W. Maurel, D. Thalmann, Y. Wu, and N. Thalmann, *Biomechanical models for soft tissue simulation*. Springer Berlin Heidelberg, 1998. [Online]. Available: <http://link.springer.com/content/pdf/10.1007/978-3-662-04603-6.pdf>
- [MTYCS04] N. Magnenat-Thalmann, L. Yahia-Cherif, and H. Seo, "Modeling anatomical-based humans," in *Proceedings - Third International Conference on Image and Graphics*, 2004, pp. 476–480.
- [MZMM12] F. E. Mullens, A. C. Zoga, W. B. Morrison, and W. C. Meyers, "Review of MRI Technique and imaging findings in athletic pubalgia and the "sports hernia"," *European Journal of Radiology*, vol. 81, no. 12, pp. 3780–3792, 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.ejrad.2011.03.100>
- [NCY<sup>+</sup>08] K. Nie, J.-H. Chen, H. J. Yu, Y. Chu, O. Nalcioğlu, and M.-Y. Su, "Quantitative Analysis of Lesion Morphology and Texture Features for Diagnostic Prediction in Breast MRI," *Academic Radiology*, vol. 15, no. 12, pp. 1513–1525, 2008. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1076633208003991>
- [NdGH09] K. Nassenstein, A. de Greiff, and P. Hunold, "MR Evaluation of Left Ventricular Volumes and Function Threshold-based 3D Segmentation Versus Short-Axis Planimetry," *October*, vol. 44, no. 10, pp. 635–640, 2009.
- [NFO<sup>+</sup>07] H. P. Ng, K. W. C. Foong, S. H. Ong, P. S. Goh, and W. L. Nowinski, "Medical Image Segmentation Using Feature-based GVF Snake," in *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, vol. 2007, jan 2007, pp. 800–803. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18002077>

- [NFSK97] W. Neuenschwander, P. Fua, G. Székely, and O. Kübler, "Velcro Surfaces: Fast Initialization of Deformable Models," *Computer Vision and Image Understanding*, vol. 65, no. 2, pp. 237–245, feb 1997. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1077314296905780>
- [NKHMT12] N. A. Nijdam, B. Kevelham, S. Han, and N. Magnenat-Thalmann, "An application framework for adaptive distributed simulation and 3D rendering services," *Proceedings of the 11th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry - VRCAI '12*, p. 103, 2012. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2407516.2407549>
- [NMK<sup>+</sup>06] A. Nealen, M. Müller, R. Keiser, E. Boxerman, and M. Carlson, "Physically Based Deformable Models in Computer Graphics," *Computer Graphics Forum*, vol. 25, no. 4, pp. 809–836, 2006. [Online]. Available: <http://doi.wiley.com/10.1111/j.1467-8659.2006.01000.x>
- [NOF<sup>+</sup>06] H. Ng, S. Ong, K. Foong, P. Goh, and W. Nowinski, "Medical Image Segmentation Using K-Means Clustering and Improved Watershed Algorithm," *2006 IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 61–65, 2006.
- [OAC<sup>+</sup>09] D. H. O'Donnell, S. Abbata, V. Chaithiraphan, K. Yared, R. P. Killeen, R. C. Cury, and J. D. Dodd, "Cardiac tumors: Optimal cardiac MR sequences and spectrum of imaging appearances," *American Journal of Roentgenology*, vol. 193, no. 2, pp. 377–387, 2009.
- [OBN04] S. D. Olabarriaga, M. Breeuwer, and W. J. Niessen, "Multi-scale Statistical Grey Value Modelling for Thrombus Segmentation from CTA," in *Proceedings of MICCAI*, vol. 3216. Springer, 2004, pp. 467–474.
- [OREG05] A. Oliveira, S. Ribeiro, C. Esperanca, and G. Giraldo, "Loop snakes: the generalized model," in *Ninth International Conference on Information Visualisation IV05*, 2005, pp. 975–980.
- [OS88] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, 1988. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/0021999188900022>
- [OS01] S. D. Olabarriaga and A. W. M. Smeulders, "Interaction in the segmentation of medical images: A survey," *Medical Image Analysis*, vol. 5, pp. 127–142, 2001.
- [PB07] J. P. Pons and J. D. Boissonnat, "Delaunay Deformable Models: Topology-Adaptive Meshes Based on the Restricted Delaunay Triangulation," in *2007 IEEE Conference on Computer Vision and Pattern Recognition*, vol. 13. Ieee, 2007, pp. 384–394. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4270044>
- [Pea01] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine*, vol. 2, no. 6, pp. 559–572, 1901. [Online]. Available: <http://stat.smmu.edu.cn/history/pearson1901.pdf>
- [Pet99] N. Peterfreund, "Robust tracking of position and velocity with Kalman snakes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 6, pp. 564–569, 1999.
- [PG94] H. Pien and J. Gauch, "Variational segmentation of multi-channel MRI images," *Image Processing, 1994. Proceedings of*, pp. 508–512, 1994. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs{\\_.}all.jsp?arnumber=413754](http://ieeexplore.ieee.org/xpls/abs{_.}all.jsp?arnumber=413754)
- [PGC<sup>+</sup>11] O. Pauly, B. Glocker, A. Criminisi, D. Mateus, A. M. Möller, S. Nekolla, and N. Navab, "Fast multiple organ detection and localization in whole-body MR dixon sequences." *Medical image computing and computer-assisted intervention : MICCAI, International Conference on*, vol. 14, no. Pt 3, pp. 239–47, jan 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22003705>
- [Pha03] D. Pham, "Unsupervised tissue classification in medical images using edge-adaptive clustering," *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)*, vol. 1, no. 2, pp. 3–6, 2003. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs{\\_.}all.jsp?arnumber=1279835](http://ieeexplore.ieee.org/xpls/abs{_.}all.jsp?arnumber=1279835)
- [PHA08] M. Poon, G. Hamarneh, and R. Abugarbieh, "Efficient interactive 3D Livewire segmentation of complex objects with arbitrary topology." *Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society*, vol. 32, no. 8, pp. 639–50, dec 2008. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18722750>
- [PHKTA06] A. Pinti, P. Hédoux, H. Kang, and A. Taleb-Ahmed, "An automated pixel classification method using surface expansion Application to MRI image sequence," *IMACS Multiconference on "Computational*

- Engineering in Systems Applications*", CESA, pp. 1514–1519, 2006.
- [PPVM06] A. Plaza, J. Plaza, D. Valencia, and P. Martinez, "Parallel segmentation of multi-channel images using multi-dimensional mathematical morphology," *Advances in Image and Video Segmentation*, pp. 270–291, 2006. [Online]. Available: <http://www.igi-global.com/chapter/parallel-segmentation-multi-channel-images/30637>
- [PSR<sup>+</sup>13] A. Poonawalla, B. Sjoberg, J. Rehm, D. Hernando, C. Hines, P. Irarrazaval, and S. Reeder, "Adipose tissue MRI for quantitative measurement of central obesity," *Journal of Magnetic Resonance Imaging*, vol. 37, no. 3, pp. 707–716, 2013. [Online]. Available: <http://onlinelibrary.wiley.com/doi/10.1002/jmri.23846/full>
- [PTVvH<sup>+</sup>12] A. Parent-Thirion, G. Vermeylen, G. van Houten, M. Lyly-Yrjänäinen, I. Biletta, and J. Cabrita, "Eurofund (2012) Fifth European Working Conditions Survey," Publications Office of the European Union, Luxembourg, Tech. Rep., 2012.
- [QHC]15] T. M. Quan, S. Han, H. Cho, and W.-K. Jeong, "Multi-GPU Reconstruction of Dynamic Compressed Sensing MRI," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, vol. 9350, 2015, pp. 238–245. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-24571-3>
- [RABM97] S. B. Reeder, E. Atalar, B. D. Bolster, and E. R. McVeigh, "Quantification and reduction of ghosting artifacts in interleaved echo-planar imaging," *Magnetic Resonance in Medicine*, vol. 38, no. 3, pp. 429–439, 1997.
- [RBH<sup>+</sup>00] H. Rifai, I. Bloch, S. Hutchinson, J. Wiart, and L. Garnerio, "Segmentation of the skull in MRI volumes using deformable model and taking the partial volume effect into account," *Medical Image Analysis*, vol. 4, no. 3, pp. 219–233, 2000. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/11145310>
- [RBM<sup>+</sup>05] M. Rohrer, H. Bauer, J. Mintonovitch, M. Requardt, and H.-J. Weinmann, "Comparison of magnetic properties of MRI contrast media solutions at different magnetic field strengths," *Investigative radiology*, vol. 40, no. 11, pp. 715–724, 2005.
- [RBZ<sup>+</sup>12] P. Ristic, D. Bokanjic, V. Zivkovic, V. Jakovljevic, M. Zdravkovic, J. Pejovic, D. Ristic, and J. Mladenovic, "Subcutaneous adipose tissue measurements and better metabolic prediction," *Central European Journal of Medicine*, vol. 8, no. 2, pp. 237–243, dec 2012. [Online]. Available: <http://www.springerlink.com/index/10.2478/s11536-012-0119-4>
- [RH94] T. Randen and J. Husoy, "Multichannel filtering for image texture segmentation," *Optical Engineering*, pp. 1–32, 1994. [Online]. Available: <http://opticalengineering.spiedigitallibrary.org/article.aspx?articleid=1073255>
- [RKL<sup>+</sup>11] H. Ruppertshofen, D. Künne, C. Lorenz, S. Schmidt, P. Beyerlein, G. Rose, and H. Schramm, "Multi-Level Approach for the Discriminative Generalized Hough Transform," in *CURAC 2011: 10. Jahrestagung der Deutschen Gesellschaft für Computer-und Roboterassistierte Chirurgie, Magdeburg, Germany, 15-16 September 2011*. CURAC, 2011.
- [RLS<sup>+</sup>10] H. Ruppertshofen, C. Lorenz, S. Schmidt, P. Beyerlein, Z. Salah, G. Rose, and H. Schramm, "Discriminative Generalized Hough transform for localization of joints in the lower extremities," *Computer Science - Research and Development*, vol. 26, no. 1, pp. 97–105, oct 2010. [Online]. Available: <http://www.springerlink.com/index/10.1007/s00450-010-0137-x>
- [RMN<sup>+</sup>14] N. Ray, S. Mukherjee, K. K. Nakka, S. T. Acton, and S. S. Blemker, "3D-To-2D Mapping For User Interactive Segmentation Of Human Leg Muscles From MRI," in *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*. IEEE, 2014, pp. 205–209.
- [RMPCO02] C. Revol-Muller, F. Peyrin, Y. Carrillon, and C. Odet, "Automated 3D region growing algorithm based on an assessment function," *Pattern Recognition Letters*, vol. 23, no. 1-3, pp. 137–150, 2002.
- [RT12] S. Resmi and T. Thomas, "A semi-automatic method for segmentation and 3D modeling of glioma tumors from brain MRI," *Journal of Biomedical Science and Engineering*, vol. 5, pp. 378–383, 2012.
- [RU90] S. P. Raya and J. K. Udupa, "Shape-based interpolation of multidimensional objects," *IEEE Transactions on Medical Imaging*, vol. 9, no. 1, pp. 32–42, 1990.
- [SA11] P. Saint-Andre, "RFC Standard 6455– The WebSocket Protocol," *RFC 6455 (Proposed Standard)*, pp. i –73, 2011. [Online]. Available: <http://datatracker.ietf.org/doc/rfc6455/>

- [Sap95] G. Sapiro, "Geometric partial differential equations in image analysis: past, present, and future," in *Proceedings International Conference on Image Processing*, vol. 3. IEEE Comput. Soc. Press, 1995, pp. 1–4. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=537565>
- [ŠB11] O. Štáva and B. Beneš, "Connected component labeling in CUDA," *GPU Computing Gems Emerald Edition*, pp. 569–581, 2011.
- [Sch11] J. Schmid, "Knowledge-based deformable models for medical image analysis," PhD thesis, University of Geneva, 2011.
- [SD92] L. H. Staib and J. S. Duncan, "Boundary finding with parametrically deformable models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 11, pp. 1061–1075, 1992. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=166621>
- [SD97] P. Smits and S. Dellepiane, "An irregular MRF region label model for multi-channel image segmentation," *Pattern recognition letters*, pp. 1133–1142, 1997. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167865597001013>
- [SE15] J. Somasekar and B. Eswara Reddy, "Segmentation of erythrocytes infected with malaria parasites for the diagnosis using microscopy imaging," *Computers & Electrical Engineering*, 2015. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0045790615001378>
- [Seg07] R. L. Segal, "Use of imaging to assess normal and adaptive muscle function." *Physical therapy*, vol. 87, no. 6, pp. 704–718, 2007.
- [Ser82] J. P. Serra, *Image analysis and mathematical morphology*, 1st ed. London, New York.: Academic Press, 1982.
- [SEYB08] S. Sumengen, M. T. Eren, S. Yesilyurt, and S. Balcisoy, "A multi-resolution mesh representation for deformable objects in collaborative virtual environments," in *Communications in Computer and Information Science*, vol. 21 CCIS, 2008, pp. 75–87.
- [SF73] I. Sobel and G. Feldman, "A 3x3 isotropic gradient operator for image processing." in *Hart, P. E. & Duda R. O. Pattern Classification and Scene Analysis*, pp. 271–272, 1973.
- [SF08] W. C. Scarfe and A. G. Farman, "What is cone-beam CT and how does it work?" *Dental clinics of North America*, vol. 52, no. 4, pp. 707–30, v, oct 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S001185320800044X>
- [SFB<sup>+</sup>14] E. Smistad, T. L. Falch, M. Bozorgi, A. C. Elster, and F. Lindseth, "Medical image segmentation on GPUs - A comprehensive review." *Medical Image Analysis*, vol. 20, no. 1, pp. 1–18, dec 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/25534282>
- [SFSG12] L. M. Sanchez, J. Fernandez, R. Sotomayor, and J. D. Garcia, "A Comparative Evaluation of Parallel Programming Models for Shared-Memory Architectures," *2012 IEEE 10th International Symposium on Parallel and Distributed Processing with Applications*, pp. 363–370, 2012. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6280314>
- [SFSV13] J. Shen, J. Fang, H. Sips, and A. L. Varbanescu, "An application-centric evaluation of OpenCL on multi-core CPUs," *Parallel Computing*, vol. 39, no. 12, pp. 834–850, 2013. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0167819113001014>
- [SG97] M. Styner and G. Gerig, "Evaluation of 2D/3D bias correction with 1+ 1ES-optimization," *Rapport de recherche*, no. October, 1997. [Online]. Available: [ftp://ftp.vision.ee.ethz.ch/publications/techreports/eth\\_{\\_}biwi\\_{\\_}00130.ps.gz](ftp://ftp.vision.ee.ethz.ch/publications/techreports/eth_{_}biwi_{_}00130.ps.gz)
- [SHL<sup>+</sup>11] T. Shen, X. Huang, H. Li, E. Kim, S. Zhang, and J. Huang, "A 3D Laplacian-driven parametric deformable model," in *Computer Vision (ICCV), 2011 IEEE International Conference on*, 2011, pp. 279–286.
- [SIG<sup>+</sup>11] J. Schmid, J. Iglesias Guitián, E. Gobbetti, N. Magnenat-Thalmann, and J. Guitian, "A GPU framework for parallel segmentation of volumetric images using discrete deformable models," *The Visual Computer*, vol. 27, no. 2, pp. 85–95, 2011. [Online]. Available: <http://dx.doi.org/10.1007/s00371-010-0532-0>
- [SK05] Y. Shi and W. C. Karl, "A Fast Level Set Method Without Solving PDEs," *Proceedings. (ICASSP '05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, vol. 2, pp. 97–100, 2005. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1415350>
- [SKBG96] G. Székely, A. Kelemen, C. Brechbühler, and G. Gerig, "Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and

- surface models." *Medical Image Analysis*, vol. 1, no. 1, pp. 19–34, 1996. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/9873919>
- [SKE<sup>+</sup>09] M. Schulze, I. Kötter, U. Ernemann, M. Fenchel, N. Tzaribatchev, C. D. Claussen, and M. Horger, "MRI Findings in Inflammatory Muscle Diseases and Their Noninflammatory Mimics," *American Journal of Roentgenology*, vol. 192, no. 6, pp. 1708–1716, 2009. [Online]. Available: <http://www.ajronline.org/doi/abs/10.2214/AJR.08.1764>
- [SKH<sup>+</sup>08] H. Seim, D. Kainmueller, M. Heller, H. Lamecker, S. Zachow, and H.-C. Hege, "Automatic Segmentation of the Pelvic Bones from CT Data Based on a Statistical Shape Model," in *Eurographics Workshop on Visual Computing for Biomedicine (VCBM)*, 2008, pp. 93–100. [Online]. Available: <http://www2.zib.de/visual/publications/sources/src-2008/vcbm2008.pdf>
- [SKK03] T. B. Sebastian, P. N. Klein, and B. B. Kimia, "On aligning curves," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 1, pp. 116–125, 2003.
- [SKL<sup>+</sup>10] 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 *Proceedings of the 4th Medical Image Analysis for the Clinic - A Grand Challenge workshop, MICCAI*, 2010.
- [SKMT11] J. Schmid, J. Kim, and N. Magnenat-Thalmann, "Robust statistical shape models for MRI bone segmentation in presence of small field of view." *Medical image analysis*, vol. 15, no. 1, pp. 155–68, feb 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/20951075>
- [SLH11] T. Shen, H. Li, and X. Huang, "Active Volume Models for Medical Image Segmentation," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 3, pp. 774–791, mar 2011.
- [SLJ13] G. E. Sujji, Y. V. S. Lakshmi, and G. W. Jiji, "MRI Brain Image Segmentation based on Thresholding," *International Journal of Advanced Computer Research*, vol. 3, no. 8, pp. 1–5, 2013. [Online]. Available: <http://wiselinjiji.com/journals/8.pdf>
- [SMZ<sup>+</sup>14] H. Shen, D. Ma, Y. Zhao, H. Sun, S. Sun, R. Ye, L. Huang, B. Lang, and Y. Sun, "MIAPS: a web-based system for remotely accessing and presenting medical images." *Computer methods and programs in biomedicine*, vol. 113, no. 1, pp. 266–83, jan 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24183386>
- [SR00] A. Souza and E. Ruiz, "Fast and accurate detection of extraocular muscle borders using mathematical morphology," in *Proceedings of the 22<sup>nd</sup> Annual EMBS International Conference*, Chicago, IL, 2000, pp. 1779–1782. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=900428](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=900428)
- [SRM14] S. Shahbazpanahi, S. Rahnamayan, and I. S. Member, "Finding Optimal Transformation Function for Image Thresholding Using Genetic Programming," in *Computational Intelligence for Multimedia, Signal and Vision Processing (CIMSIVP), 2014 IEEE Symposium on*, 2014, pp. 1–8.
- [SS08] F. G. Shellock and A. Spinazzi, "MRI safety update 2008: Part 1, MRI contrast agents and nephrogenic systemic fibrosis," *American Journal of Roentgenology*, vol. 191, no. 4, pp. 1129–1139, 2008.
- [SSKH10] R. Shams, P. Sadeghi, R. Kennedy, and R. Hartley, "A Survey of Medical Image Registration on Multicore and the GPU," *IEEE Signal Processing Magazine*, vol. 27, no. 2, pp. 50–60, 2010. [Online]. Available: [http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=5438962](http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=5438962)
- [ST92] R. Szeliski and D. Tonnesen, "Surface modeling with oriented particle systems," *ACM SIGGRAPH Computer Graphics*, vol. 26, no. 2, pp. 185–194, 1992. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=142920.134037>
- [Sta05] S. Standring, *Gray's Anatomy: The Anatomical Basis of Clinical Practice*, 39th ed. Churchill Livingstone, 2005, vol. 10. [Online]. Available: <http://www.amazon.com/dp/0443071683>
- [Ste90] T. L. Stedman, *Stedman's Medical Dictionary*, 25th ed. Baltimore: Williams & Wilkins, cop., 1990.
- [STL<sup>+</sup>15] M. Sdika, A. Tonson, Y. Le Fur, P. J. Cozzone, and D. Bendahan, "Multi-atlas-based fully automatic segmentation of individual muscles in rat leg," *Magma (New York, N.Y.)*, 2015. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/26646521>
- [SVG02] J. G. Snel, H. W. Venema, and C. A. Grimbergen, "Deformable triangular surfaces using fast 1-D radial Lagrangian dynamics—segmentation of 3-D MR and CT images of the wrist." *IEEE Transactions on Medical Imaging*, vol. 21, no. 8, pp. 888–903, 2002. [Online]. Available:

- <http://www.ncbi.nlm.nih.gov/pubmed/12472262>
- [SVS84] K. P. Soila, M. Viamonte, and P. M. Starewicz, "Chemical shift misregistration effect in magnetic resonance imaging." *Radiology*, vol. 153, no. 3, pp. 819–820, dec 1984. [Online]. Available: <http://radiology.rsna.org/content/153/3/819.abstract>
- [SZE98] J. G. Sled, a. P. Zijdenbos, and a. C. Evans, "A nonparametric method for automatic correction of intensity nonuniformity in MRI data." *IEEE transactions on medical imaging*, vol. 17, no. 1, pp. 87–97, feb 1998. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/9617910>
- [TD12] M. J. Thurley and V. Danell, "Fast morphological image processing open-source extensions for GPU processing with CUDA," *IEEE Journal on Selected Topics in Signal Processing*, vol. 6, no. 7, pp. 849–855, 2012.
- [TF88] D. Terzopoulos and K. Fleischer, "Deformable models," *The Visual Computer*, vol. 4, no. 6, pp. 306–331, 1988. [Online]. Available: <http://www.springerlink.com/content/r6m08783w2p326m5/>
- [TFM<sup>+</sup>13] E. L. Thomas, J. a. Fitzpatrick, S. J. Malik, S. D. Taylor-Robinson, and J. D. Bell, "Whole body fat: content and distribution." *Progress in nuclear magnetic resonance spectroscopy*, vol. 73, pp. 56–80, aug 2013. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/23962884>
- [TG09a] N. J. Tustison and J. C. Gee, "Introducing Dice, Jaccard, and Other Label Overlap Measures To ITK," *Insight Journal*, no. 2, 2009. [Online]. Available: <http://hdl.handle.net/10380/3141>
- [TG09b] N. Tustison and J. Gee, "N4ITK: Nick's N3 ITK implementation for MRI bias field correction," *Insight Journal*, pp. 1–8, 2009. [Online]. Available: [http://www.insight-journal.org/download/pdf/18880/IJ\[\\_\]640\[\\_\]1\[\\_\]n3.pdf](http://www.insight-journal.org/download/pdf/18880/IJ[_]640[_]1[_]n3.pdf)
- [TH11] M. Taghizadeh and M. Hajipoor, "A hybrid algorithm for segmentation of MRI images based on edge detection," *2011 International Conference of Soft Computing and Pattern Recognition (SoCPaR)*, pp. 107–111, 2011.
- [THA11] A. Top, G. Hamarneh, and R. Abugharbieh, "Active learning for interactive 3D image segmentation." *Medical image computing and computer-assisted intervention : MICCAI, International Conference on*, vol. 14, no. Pt 3, pp. 603–10, jan 2011. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22003749>
- [TKLL06] T. Tölli, J. Koikkalainen, K. Lauerma, and J. Lötjönen, "Artificially enlarged training set in image segmentation." *Medical Image Computing and Computer-Assisted Intervention*, vol. 9, no. Pt 1, pp. 75–82, 2006. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17354876>
- [TKR<sup>+</sup>08] K. M. Taylor, J. S. Kim, W. J. Rieter, H. An, and W. Lin, "Mesoporous silica nanospheres as highly efficient MRI contrast agents," *J Am Chem Soc*, vol. 130, no. 7, pp. 2154–2155, 2008.
- [TM91] D. Terzopoulos and D. Metaxas, "Dynamic 3D models with local and global deformations: deformable superquadrics," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 703–714, 1991. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=85659>
- [TNL<sup>+</sup>14] M. S. Thomas, D. Newman, O. D. Leinhard, B. Kasmai, R. Greenwood, P. N. Malcolm, A. Karlsson, J. Rosander, M. Borga, and A. P. Toms, "Test-retest reliability of automated whole body and compartmental muscle volume measurements on a wide bore 3T MR system." *European radiology*, may 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/24871333>
- [TPBF87] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically Deformable Models," *SIGGRAPH Comput. Graph.*, vol. 21, no. 4, pp. 205–214, aug 1987. [Online]. Available: <http://doi.acm.org/10.1145/37402.37427>
- [TRGP10] Z. Tang, G. Rong, X. Guo, and B. Prabhakaran, "Streaming 3D shape deformations in collaborative virtual environment," *2010 IEEE Virtual Reality Conference (VR)*, pp. 183–186, mar 2010. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5444793>
- [TSM11] Y. Tang, R. Simerly, and R. Moats, "An Automatic Technique for MRI Based Murine Abdominal Fat Measurement," *Radioengineering*, pp. 988–995, 2011. [Online]. Available: [http://radioeng.cz/fulltexts/2011/11\[\\_\]04\[\\_\]988\[\\_\]995.pdf](http://radioeng.cz/fulltexts/2011/11[_]04[_]988[_]995.pdf)
- [TTF<sup>+</sup>97] M. M. Thurnher, S. a. Thurnher, D. Fleischmann, A. Steuer, A. Rieger, T. Helbich, S. Trattinig, E. Schindler, and K. Hittmair, "Comparison of T2-weighted and fluid-attenuated inversion-recovery fast spin-echo MR sequences in intracerebral AIDS-associated disease," *American Journal of Neuroradiology*, vol. 18,

- no. 9, pp. 1601–1609, 1997.
- [TTK12] D. J. Theodorou, S. J. Theodorou, and Y. Kakitsubata, “Skeletal muscle disease: patterns of MRI appearances,” *The British Journal of Radiology*, vol. 85, no. 1020, pp. e1298–e1308, 2012. [Online]. Available: <http://www.birpublications.org/doi/abs/10.1259/bjr/14063641>
- [TVGPAFW12] A. Tristán-Vega, V. García-Pérez, S. Aja-Fernández, and C. F. Westin, “Efficient and robust nonlocal means denoising of MR data based on salient features matching,” *Computer Methods and Programs in Biomedicine*, vol. 105, no. 2, pp. 131–144, 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.cmpb.2011.07.014>
- [TVP15] K. Tabelow, H. U. Voss, and J. Polzehl, “Local estimation of the noise level in MRI using structural adaptation,” *Medical Image Analysis*, vol. 20, no. 1, pp. 76–86, feb 2015. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/25465845>
- [TWK87] D. Terzopoulos, A. Witkin, and M. Kass, “Symmetry-Seeking Models and 3D Object Reconstruction,” *International Journal of Computer Vision*, vol. 221, pp. 211–221, 1987.
- [TWM<sup>+</sup>00] H. Tang, E. X. Wu, Q. Y. Ma, D. Gallagher, G. M. Perera, and T. Zhuang, “MRI brain image segmentation by multi-resolution edge detection and region selection,” *Computerized Medical Imaging and Graphics*, vol. 24, no. 6, pp. 349–357, 2000.
- [UMV09] L. Urricelqui, A. Malanda, and A. Villanueva, “Automatic segmentation of thigh magnetic resonance images,” *World Acad Sci Eng Technol*, vol. 3, no. 10, pp. 953–959, 2009.
- [UPB08] M. Unger, T. Pock, and H. Bischof, “Continuous Globally Optimal Image Segmentation with Local Constraints,” in *Proceedings of the Computer Vision Winter Workshop 2008*, J. Pers, Ed., 2008.
- [VC02] L. Vese and T. Chan, “A multiphase level set framework for image segmentation using the Mumford and Shah model,” *International Journal of Computer Vision*, vol. 50, no. 3, pp. 271–293, 2002. [Online]. Available: <http://link.springer.com/article/10.1023/A:1020874308076>
- [vdGVBV02] M. van der Glas, F. M. Vos, C. P. Botha, and A. M. Vossepel, “Determination of position and radius of ball joints,” in *Proc. SPIE 4684, Medical Imaging 2002: Image Processing*, 157, 2002, pp. 1–7.
- [VG99] B. B. Vemuri and Y. Guo, “Hybrid Geometric Active Models for Shape Recovery in Medical Image,” in *Proc. Information Processing in Medical Imaging (IPMI ’99)*, 1999, pp. 112–125.
- [VHM] VRLab, HUGE, and MIRALab, “Co-Me, Swiss national project.” [Online]. Available: <http://co-me.ch/>
- [VKA<sup>+</sup>13] A. Valentinitich, D. C. Karampinos, H. Alizai, K. Subburaj, D. Kumar, T. M. Link, and S. Majumdar, “Automated unsupervised multi-parametric classification of adipose tissue depots in skeletal muscle,” *Journal of Magnetic Resonance Imaging : JMRI*, vol. 37, no. 4, pp. 917–27, apr 2013. [Online]. Available: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3573225&tool=pmcentrez&rendertype=abstract>
- [VMT05] P. Volino and N. Magnenat-Thalmann, “Implicit midpoint integration and adaptive damping for efficient cloth simulation,” in *Computer Animation and Virtual Worlds*, vol. 16, 2005, pp. 163–175.
- [VMTF09] P. Volino, N. Magnenat-Thalmann, and F. Faure, “A simple approach to nonlinear tensile stiffness for accurate cloth simulation,” *ACM Transactions on Graphics*, vol. 28, no. 4, pp. 1–16, aug 2009. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1559755.1559762>
- [VN08] V. Vineet and P. J. Narayanan, “CUDA cuts: Fast graph cuts on the GPU,” in *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW ’08. IEEE Computer Society Conference on*, jun 2008, pp. 1–8.
- [VPJ02] J. W. M. Van Goethem, P. M. Parizel, and J. R. Jenkins, “Review article: MRI of the postoperative lumbar spine,” *Neuroradiology*, vol. 44, no. 9, pp. 723–739, 2002.
- [VPL07] U. Vovk, F. Pernus, and B. Likar, “A review of methods for correction of intensity inhomogeneity in MRI,” *Medical Imaging, IEEE Transactions on*, vol. 26, no. 3, pp. 405–21, mar 2007. [Online]. Available: [http://www.ncbi.nlm.nih.gov/pubmed/17354645http://ieeexplore.ieee.org/xpls/abs/\\_all.jsp?arnumber=4114560](http://www.ncbi.nlm.nih.gov/pubmed/17354645http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=4114560)
- [VSG<sup>+</sup>15] P. Vemuri, M. L. Senjem, J. L. Gunter, E. S. Lundt, N. Tosakulwong, S. D. Weigand, B. J. Borowski, M. a. Bernstein, S. M. Zuk, V. J. Lowe, D. S. Knopman, R. C. Petersen, N. C. Fox, P. M. Thompson, M. W. Weiner, and C. R. Jack, “Accelerated vs. unaccelerated serial MRI based TBM-SyN measurements for clinical trials in Alzheimer’s disease,” *NeuroImage*, vol. 113, pp. 61–69, 2015. [Online]. Available:



- <http://linkinghub.elsevier.com/retrieve/pii/S1053811915002086>
- [VT00] P. Volino and N. M. Thalmann, "Implementing fast cloth simulation with collision response," *Proceedings Computer Graphics International 2000*, 2000.
- [VVBV01] M. J. P. Van Osch, E. J. P. a. Vonken, C. J. G. Bakker, and M. a. Viergever, "Correcting partial volume artifacts of the arterial input function in quantitative cerebral perfusion MRI," *Magnetic Resonance in Medicine*, vol. 45, no. 3, pp. 477–485, 2001.
- [VYBS95] C. Vacheta, B. Yvernault, K. Bhatta, and R. Smitha, "Automatic corpus callosum segmentation using a deformable active Fourier contour model," *Magnetic Resonance Imaging*, p. 1995, 1995. [Online]. Available: <http://www.nitrc.org/docman/view.php/466/1131/SPIE{ }2012{ }Vachet{ }CCSeg.pdf>
- [Wat09] G. Watts, "John Wild," *BMJ*, vol. 339, oct 2009. [Online]. Available: <http://www.bmj.com/content/339/bmj.b4428.abstract>
- [WG97] K. Wang and T. Gasser, "Alignment of Curves by Dynamic Time Warping," *The Annals of Statistics*, vol. 25, no. 3, pp. 1251–1276, 1997.
- [WGKJ94] W. M. Wells III, W. E. L. Grimson, R. Kikinis, and F. a. Jolesz, "Statistical Intensity Correction and Segmentation of MRI data," in *Visualization in Biomedical Computing*, R. A. Robb, Ed. International Society for Optics and Photonics, sep 1994, pp. 13–24. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=982936>
- [WGY<sup>+</sup>07] C. Wang, X. Guo, L. Yao, K. Li, and Z. Jin, "A Practical Method for Muscles Extraction and Automatic Segmentation of Leg Magnetic Resonance Images," *Complex Medical Engineering*, 2007. *CME 2007. IEEE/ICME International Conference on*, pp. 885–890, may 2007.
- [WGZ<sup>+</sup>08] C. Wachinger, B. Glocker, J. Zeltner, N. Paragios, N. Komodakis, M. S. Hansen, and N. Navab, "Deformable mosaicing for whole-body MRI." *Medical image computing and computer-assisted intervention : MICCAI, International Conference on*, vol. 11, no. Pt 2, pp. 113–21, jan 2008. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18982596>
- [Wha11] Whatwg.org, "The Canvas Element," 2011. [Online]. Available: <http://www.whatwg.org/specs/web-apps/current-work/multipage/the-canvas-element.html>
- [WHC<sup>+</sup>13] T. a. Willis, K. G. Hollingsworth, A. Coombs, M.-L. Sveen, S. Andersen, T. Stojkovic, M. Eagle, A. Mayhew, P. L. de Sousa, L. Dewar, J. M. Morrow, C. D. J. Sinclair, J. S. Thornton, K. Bushby, H. Lochmüller, M. G. Hanna, J.-Y. Hogrel, P. G. Carlier, J. Vissing, and V. Straub, "Quantitative muscle MRI as an assessment tool for monitoring disease progression in LGMD2I: a multicentre longitudinal study." *PloS one*, vol. 8, no. 8, p. e70993, 2013. [Online]. Available: [#pone.0070993.s001">http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0070993](http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0070993)
- [Whi98] R. T. Whitaker, "A Level-Set Approach to 3D Reconstruction from Range Data," *International Journal of Computer Vision*, vol. 29, no. 3, pp. 203–231, 1998. [Online]. Available: <http://portal.acm.org/citation.cfm?id=299660.299682>
- [WN51] J. J. Wild and D. Neal, "Use of High-Frequency Ultrasonic Waves for Detecting Changes of Texture in Living Tissues," *The Lancet*, vol. 257, no. 6656, pp. 655–657, mar 1951. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0140673651924038>
- [WOS09] K. Wu, E. Otoo, and K. Suzuki, "Optimizing two-pass connected-component labeling algorithms," *Pattern Analysis and Applications*, vol. 12, no. 2, pp. 117–135, jun 2009. [Online]. Available: <http://link.springer.com/10.1007/s10044-008-0109-y>
- [WPJR11] A. Wessels, M. Purvis, J. Jackson, and S. S. Rahman, "Remote Data Visualization through WebSockets," 2011 *Eighth International Conference on Information Technology: New Generations*, pp. 1050–1051, 2011. [Online]. Available: <http://ieeexplore.ieee.org/xpls/abs{ }all.jsp?arnumber=5945383>
- [WPNK08] J. Wu, S. Poehlman, M. D. Noseworthy, and M. V. Kamath, "Texture Feature based Automated Seeded Region Growing in Abdominal MRI Segmentation," 2008 *International Conference on BioMedical Engineering and Informatics*, pp. 263–267, 2008. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4549175>
- [WVW<sup>+</sup>04] 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," *Medical Imaging 2004*, pp. 16–27, 2004. [Online]. Available: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=840857>
- [XHY02] D. Xu, J. Hwang, and C. Yuan, "Segmentation of multi-channel image with markov random field based active contour model," *Journal of VLSI Signal Processing*, no. 31, pp. 45–55, 2002. [Online]. Available: <http://link.springer.com/article/10.1023/A:1014493104976>
- [XJA14] Q. Xu, H. Jeon, and M. Annavaram, "Graph processing on GPUs: Where are the bottlenecks?" in *2014 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, oct 2014, pp. 140–149. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6983053>
- [XP98] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow." *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 359–369, 1998. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/18276256>
- [YAL<sup>+</sup>02] 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," in *Studies in Health Technology and Informatics*, vol. 85, 2002, pp. 586–592.
- [YCGMMT03] L. Yahia-Cherif, B. Gilles, L. Moccozet, and N. Magnenat-Thalmann, "Individualized bone modelling from MRI: application to the human hip," *International Congress . . .*, 2003. [Online]. Available: [#0](http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Individualized+bone+modelling+from+MRI+application+to+the+human+hip)
- [YD13] J. Yang and T. Deng, "Optimal binary thresholding segmentation for medical images in rough fuzzy set framework," *Proceedings of the 2013 International Conference on Intelligent Control and Information Processing, ICICIP 2013*, pp. 638–643, 2013.
- [YPH<sup>+</sup>06] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, "User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability." *NeuroImage*, vol. 31, no. 3, pp. 1116–1128, 2006. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/16545965>
- [YQWHZC<sup>+</sup>05] Z. Yu-Qian, G. Wei-Hua, C. Zhen-Cheng, T. Jing-Tian, and L. Ling-Yun, "Medical images edge detection based on mathematical morphology." in *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, vol. 6, jan 2005, pp. 6492–6495. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17281756>
- [YR97] M. J. Yaffe and J. A. Rowlands, "X-ray detectors for digital radiography." *Physics in Medicine and Biology*, vol. 42, no. 1, pp. 1–39, 1997.
- [YXSN09] S. Y. Yeo, X. Xie, I. Sazonov, and P. Nithiarasu, "Geometric Potential Force for the Deformable Model," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2009, pp. 99.1–99.11.
- [YZG05] P. A. Yushkevich, H. Zhang, and J. C. Gee, "Statistical modeling of shape and appearance using the continuous medial representation." *Medical Image Computing and Computer-Assisted Intervention*, vol. 8, no. Pt 2, pp. 725–732, 2005. [Online]. Available: <http://discovery.ucl.ac.uk/192092/>
- [ZBS01] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm." *IEEE Transactions on Medical Imaging*, vol. 20, no. 1, pp. 45–57, 2001. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/11293691>
- [ZCMO96] H. K. Zhao, T. Chan, B. Merriman, and S. Osher, "A Variational Level Set Approach to Multiphase Motion," *Journal of Computational Physics*, vol. 127, no. 1, pp. 179–195, 1996. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0021999196901679>
- [ZCZ11] B. Zhao, S. Cheng, and X. Zhang, "A New Adaptive Deformable Model Using Gradient Vector Flow," *Integration The VLSI Journal*, pp. 472–482, 2011.
- [ZJZ07] J. Z.-J. J. Zhu-Jacquot and R. Zabih, "Graph Cuts Segmentation with Statistical Shape Priors for Medical Images," in *2007 Third International IEEE Conference on SignalImage Technologies and InternetBased System*. Ieee, 2007, pp. 631–635. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4618832>
- [ZLY96] S. C. Zhu, T. S. Lee, and A. L. Yuille, "Region competition: unifying snakes, region growing, energy/Bayes/MDL for multi-band image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence archive Volume 18 Issue 9, September 1996*, no. 6, pp. 416–423, 1996.

- [ZMLS06] J. Z. J. Zhang, M. Marszalek, S. Lazebnik, and C. Schmid, "Local Features and Kernels for Classification of Texture and Object Categories: A Comprehensive Study," pp. 213–238, 2006. [Online]. Available: <http://www.springerlink.com/index/10.1007/s11263-006-9794-4>
- [ZMSM08] Y. Zhang, B. J. Matuszewski, L. K. Shark, and C. J. Moore, "Medical Image Segmentation Using New Hybrid Level-Set Method," in *2008 Fifth International Conference BioMedical Visualization Information Visualization in Medical and Biomedical Informatics*, vol. 26. Ieee, 2008, pp. 71–76. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4618616>
- [ZTHW12] Z. Zhang, T. Tan, K. Huang, and Y. Wang, "Three-dimensional deformable-model-based localization and recognition of road vehicles." *IEEE Transactions on Image Processing*, vol. 21, no. 1, pp. 1–13, 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/21724513>
- [ZX13] F. Zhao and X. Xie, "An Overview of Interactive Medical Image Segmentation," *bmva.org*, vol. 2013, pp. 1–22, 2013. [Online]. Available: <http://www.bmva.org/annals/2013/2013-0007.pdf>
- [YZ10] J. Zhou, M. Ye, and X. Zhang, "Graph Cut segmentation with automatic editing for Industrial images," in *2010 International Conference on Intelligent Control and Information Processing*. Ieee, aug 2010, pp. 633–637. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5565294>
- [ZZM12] F. Zhao, J. Zhang, and Y. Ma, "Medical Image Processing Based on Mathematical Morphology," *Proceedings of the 2nd International Conference on Computer Application and System Modeling*, pp. 948–950, 2012. [Online]. Available: <http://www.atlantis-press.com/php/paper-details.php?id=2705>