

Variational and Statistical Shape Modeling for 3D Geometry Reconstruction

Hans Lamecker

Zuse-Institut Berlin (ZIB)

Dissertation

am Fachbereich Mathematik und Informatik
der Freien Universität Berlin

zur Erlangung des akademischen Grades
des Doktors der Naturwissenschaften (Dr. rer. nat.)

Berlin, April 2008

Betreuer

Prof. Dr. Dr. h.c. Peter Deuffhard

Freie Universität Berlin, Fachbereich Mathematik und Informatik
Präsident des Zuse-Instituts Berlin

Gutachter

Prof. Dr. Dr. h.c. Peter Deuffhard

Prof. Dr. Martin Rumpf (Institut für Numerische Simulation, Universität Bonn)

Datum der Disputation

17. Oktober 2008

Abstract

The reconstruction of geometric shapes plays an important role in many biomedical applications. One example is the patient-specific, computer-aided planning of complex interventions, which requires the generation of explicitly represented geometric models of anatomical structures from medical image data. Only solutions that require minimal interaction by medical personnel are likely to enter clinical routine. Another example is the planning of surgical corrections of deformities where the target shape is unknown. Surgeons are often forced to resort to subjective criteria. These applications still pose highly challenging reconstruction problems, which are addressed in this thesis.

The fundamental hypothesis, pursued in this thesis, is that the problems can be solved by incorporating a-priori knowledge about shape and other application-specific characteristics. Here, we focus mainly on the aspect of geometric shape analysis. The basic idea is to capture the most essential variations of a certain class of geometric objects via statistical shape models, which model typical features contained in a given population, and restrict the outcome of a reconstruction algorithm (more or less) to the space spanned by such models.

A fundamental prerequisite for performing statistical shape analysis on a set of different objects is the identification of corresponding points on their associated surfaces. This problem is particularly difficult to solve if the shapes stem from different individuals. The reason lies in the basic difficulty of defining suitable measures of similarity. In this thesis, we divide the correspondence problem into feature and non-feature matching. The feature part depends on the application, while the non-feature part can be characterized by a purely geometric description. We propose two different approaches. The first approach has proved useful in many applications. Yet, it suffers from some practical limitations and does not yield a measure of similarity. Our second, variational, approach is designed to overcome these limitations. In it, we propose to minimize an invariant stretching measure, constrained by previously computed features. An important property, which sets our method apart from previous work, is that it does not require the computation of a global surface parameterization.

Zusammenfassung

Die Rekonstruktion geometrischer Formen spielt eine wichtige Rolle in vielen biomedizinischen Anwendungen. Ein Beispiel ist die patientenspezifische, computergestützte Planung komplexer Interventionen, die die Erzeugung explizit repräsentierter geometrischer Modelle anatomischer Strukturen aus medizinischen Bilddaten erfordert. Nur Lösungen, welche geringster Interaktion durch das medizinische Personal bedürfen, können den Weg in die klinische Routine finden. Ein anderes Beispiel ist die Planung chirurgischer Korrekturen von Missbildungen, bei denen die Zielform unbekannt ist. Chirurgen sind häufig gezwungen, auf subjektive Kriterien zurückzugreifen. Anwendungsszenarien dieser Art werfen schwierige Rekonstruktionsprobleme auf, die in dieser Arbeit behandelt werden.

Die grundlegende Hypothese, auf welcher diese Arbeit aufbaut, ist, dass derartige Rekonstruktionsprobleme gelöst werden können, indem man a-priori Wissen über Form- und andere anwendungsspezifische Eigenschaften berücksichtigt. Wir konzentrieren uns besonders auf den Aspekt der geometrischen Formanalyse. Die Grundidee besteht darin, die wesentlichen Variationen einer bestimmten Klasse von geometrischen Objekten in einem statistischen Formmodell zu erfassen. Dieses enthält typische Merkmale einer gegebenen Population und schränkt das Resultat eines Rekonstruktionsalgorithmus (mehr oder weniger) auf den durch das Modell beschriebenen Lösungsraum ein.

Eine Grundlage für die statistische Analyse von unterschiedlichen Formen ist die Identifikation korrespondierender Punkte auf deren Oberflächen. Dieses Problem ist besonders schwierig, wenn die Formen von unterschiedlichen Individuen stammen. Die Ursache liegt in der grundlegenden Schwierigkeit, geeignete Ähnlichkeitsmaße zu definieren. In der vorliegenden Arbeit separieren wir das Korrespondenzproblem in die Identifikation von besonderen Merkmalen und nicht-merkmalsbehafteten Regionen. Während der Merkmalsanteil immer von einer spezifischen Anwendung abhängt, lassen sich nicht-merkmalsbehaftete Regionen nur geometrisch charakterisieren. Wir schlagen zwei verschiedene Verfahren vor. Der erste Ansatz hat sich in vielen Anwendungen als nützlich erwiesen, unterliegt jedoch praktischen Einschränkungen und liefert kein Ähnlichkeitsmaß. Unser zweites, variationelles Verfahren soll diese Eingrenzungen überwinden. Wir schlagen zur Minimierung ein invariantes, metrisches Verzerrungsmaß vor, unter der Berücksichtigung vorab berechneter Merkmale. Eine wichtige Eigenschaft unseres Verfahrens, die es von anderen wesentlich unterscheidet, ist, dass es keiner Berechnung einer globalen Oberflächenparametrisierung bedarf.

Acknowledgements

First of all, I wish to thank Hans-Christian Hege, head of the Department of Visualization and Data Analysis at the Zuse Institute Berlin (ZIB), for introducing me to the field of statistical shape modeling and 3D geometry reconstruction. Without his vision and constant support this work would not have been possible.

I am grateful to Prof. Dr. Dr. h.c. Peter Deuffhard, president of the ZIB, for suggesting the variational approach to shape matching. His guidance and continuous encouragement were indispensable in the realization of this work.

Originally, the work was inspired and also funded by DFG Collaborative Research Center 273 “Hyperthermia: Methodology and Clinics”. When the project ended in 2002, the work was placed in a more general setting within the DFG Research Center 86 MATHEON “Mathematics for Key Technologies: Modeling, Simulation and Optimization of Real-World Processes”.

I wish to thank Dr. Martin Seebass for acquainting me with medical image analysis and geometric anatomical modeling and for awakening my interest in the hyperthermia project. To Dr. Stefan Zachow, head of the Medical Planning Group at ZIB, I am indebted for introducing me to problems in computer-aided surgery and for his extensive support of this work.

One of the reasons for my joining the ZIB was the open-minded, motivating and supportive atmosphere created by all my colleagues in the Departments of Visualization and Data Analysis, Numerical Analysis and Modeling, and in the rest of the ZIB. My special thanks go to: Dr. Johannes Schmidt-Ehrenberg, Dr. Daniel Baum, Dr. Steffen Prohaska, Dr. Martin Weiser, Dr. Anton Schiela and many others for their helpful discussions; to Dr. Malte Westerhoff and Dr. Detlev Stalling for the development of the Amira software, which has served as the platform for this work; to Dagmar Kainmüller for her efforts in taking liver segmentation to new heights; to the students who have been involved in discussing and implementing ideas related to this work: Thomas Wenckebach, Philipp Beckmann, Britta Weber, Michael Zilske, Antonia Wittmers; and to all others at the ZIB who have made everyday work enjoyable and ensured its smooth running.

It has been an inspiration and a pleasure to work with Dr. Max Wardetzky, to learn about differential geometry, particularly the relationship between the

discrete and the continuous worlds, and to develop with him the method of variational surface matching, which is presented in this work.

I would also like to thank Prof. Dr. Konrad Polthier, chair of the Mathematical Geometry Processing Group at the Freie Universität Berlin, and his group for their cooperation, especially for sharing their expertise in the field of differential geometry. In addition, I extend my gratitude to Prof. Dr. Ulrich Pinkall, head of the Differential Geometry and Visualization Group at the Technische Universität Berlin, for his assistance with matters of differential geometry.

This work has also been motivated by problems with real-life applications. I wish to thank all partners and co-authors outside the ZIB for their collaboration and for granting an insight into their fields of expertise: Prof. Dr. Dr. h.c. Peter Schlag (Charité Berlin, Campus Buch) - computer-aided liver surgery planning; Dr. Johanna Gellermann and Prof. Dr. Peter Wust (Charité Berlin, Campus Virchow) - hyperthermia treatment planning; PD Dr. Hannes Haberl (Charité Berlin, Campus Virchow), Dr. Maja Zöckler and Dr. Malte Westerhoff - collaboration and support in the craniosynostosis project; PD Dr. Dr. Michael Stiller and Barbara Elsholtz (Charité Berlin, Campus Benjamin Franklin) - mandible-related projects; Dr. Markus Heller and Prof. Dr. Georg Duda (Charité Berlin, Center for Musculoskeletal Surgery) - the 2D-3D reconstruction project; and many others. Special thanks go to Thomas Lange (Charité Berlin, Campus Buch) for frank and fruitful collaboration over a period of many years.

This work has additionally benefited from numerous discussions with and helpful reviews by colleagues in the scientific community, to whom I owe my gratitude for appraising my work critically.

Last but not least, I thank my family for their support and enduring patience.

Contents

1	Introduction	1
1.1	Apples and Oranges	1
1.2	Quantitative Comparison of Shapes	2
1.3	Outline and Contributions	4
Part I	Shape Modeling	7
2	The Correspondence Problem for Surfaces	8
2.1	Introduction	8
2.2	Problem Statement	8
2.3	Geometric Characterization of Shape	9
2.4	Classification of Shape Matching Methods	13
2.5	Review of Geometric Shape Matching Algorithms	16
2.6	Consistent Surface Decomposition and Parameterization	18
2.7	Conclusions	24
3	Variational Surface Matching	25
3.1	Introduction	25
3.2	Measuring Geometric Deviation	25
3.3	Minimizing Stretching	28
3.4	Discrete Stretching Functional	29
3.5	First Variation of the Discrete Stretching Functional	33
3.6	Minimizing the Discrete Stretching Functional	34
3.7	Geodesics, Exponential Map and Surface Area for Meshes	37
3.8	Implementation and Numerical Results	41
3.9	Conclusions	46

4	Statistical Shape Modeling	48
4.1	Point Distribution Models	48
4.2	Shape Alignment	50
4.3	Statistical Shape Models	51
Part II	3D Geometry Reconstruction	56
5	Segmentation of Medical Images	57
5.1	Introduction	57
5.2	Problems and Challenges	58
5.3	Segmentation with Statistical Shape Models	59
5.4	Mandible from Low-Dose CT data	62
5.5	Liver Segmentation from CT data	66
6	Reconstruction from Pathological Shapes	76
6.1	Introduction	76
6.2	Model-based Surface Reconstruction	76
6.3	Reconstruction of Mandibular Dysplasia	79
6.4	Craniosynostosis	82
7	3D Shape Estimation from X-Ray Projections	86
7.1	Introduction	86
7.2	Reconstruction Process	87
7.3	Distance Measures	88
7.4	Reconstruction of Pelvic Bone	89
7.5	Conclusions	91
	Conclusions	93
	Appendix A Riemannian Geometry	94
	Appendix B Discrete Gauss Curvature	100
	Appendix C The Magnitude of the Area Gradient	101
	Bibliography	103

Chapter 1

Introduction

1.1 Apples and Oranges

How can different three dimensional objects be compared to each other in a quantitative fashion? What is the difference between them? The idiom *comparing apples to oranges* indicates that a quantitative comparison is sometimes not valid because two objects are not related. But can apples and oranges be related quantitatively? The answer is: yes - in several ways. Not only do apples and oranges have similar spectroscopic properties ¹, but we can compare their color, surface constitution, weight - or *geometric shape*.

Geometric shape is the common denominator by which all three dimensional objects are related. In many cases it can be measured much easier than other properties and it plays an important role in making out differences in our perception of things. D'Arcy Wentworth Thompson remarked in his famous book "On Growth and Form" (1917) that geometric shape, among the many different properties of three dimensional objects, plays a fundamental role:

"Cell and tissue, shell and bone, leaf and flower, are so many portions of matter, and it is in obedience to the laws of physics that their particles have been moved, moulded and conformed. They are no exceptions to the rule that God always geometrizes. Their problems of form are in the first instance mathematical problems, their problems of growth are essentially physical problems, and the morphologist is, ipso facto, a student of physical science."

The geometric description and quantitative relation of shapes is fundamental for the task of transforming shapes into one another, which is also known by its Greek word *metamorphosis* (μεταμόρφωση = change of shape) or its English variant *morphing*. Modern movie making often employs morphing techniques to render the seemingly impossible possible, e.g. transforming hu-

¹Apples and oranges show similar infrared transmission spectra. See Scott A. Sandford, NASA Ames Research Center: Apples and Oranges – A Comparison, Annals of Improbable Research (AIR), 1995.




Figure 1.1 Transformation of a human into a rhinoceros head.




Figure 1.2 When the tip of the nose $p \in M$ on the left head is matched with a point on the cheek $\tilde{p} \in \tilde{M}$ on the right head, shape interpolation may yield a head with two noses in an intermediate stage.

mans into animals as in Fig. 1.1 - by reducing humans and animals to geometric objects. The most crucial step in this process is matching corresponding points on different shapes correctly, as false identification may lead to unwanted results (Fig. 1.2).

Digital representations of shapes have opened up a variety of applications for quantitative shape comparison besides movie animation. This includes reconstruction, recovery, tracking, prediction or detection of shapes. Such techniques play a vital role in fields like computer graphics, computer vision, pattern recognition, robotics, image processing and many more.

The goal of this work is to contribute mathematical strategies to solve the problem of identifying corresponding points between different shapes on a *purely geometric basis*, and to show how to use these strategies in specific applications.

1.2 Quantitative Comparison of Shapes

In this work the word “shape” will refer to the geometric information represented by the boundary of a three-dimensional object. It will be used interchangeably with the word *surface*.

Throughout this work different surfaces are assumed to be topologically

equivalent. What else can we say about the kind of shapes we intend to compare? The answer comprises two important aspects:

Different subjects. We are interested in modeling variations of a three dimensional object with certain functional or semantic properties from different individual origins. These origins could be organs from different patients or different biological compounds. Such shapes may exhibit large geometric variations. Furthermore, physical deformation models of such objects are not known in general.

Discrete shapes. The shapes we will consider in this work are acquired from physical measurements (e.g. a tomographic device, laser scanner, microscope, etc.) and have potentially undergone some complex reconstruction process. They are therefore discrete in nature and subject to errors of often unknown size.

The goal is to devise methods that capture the spatial variability of different shapes with the properties mentioned before. This is based on three tasks:

1. The ability of identifying corresponding points on different shapes.
2. The alignment of different shapes into a common frame of reference.
3. The interpolation between different shapes (morphing).

The terms *matching* or *registration* are often used to denote either the first task alone or the first two tasks together. Unfortunately, this ambiguity sometimes causes confusion. The first problem however - also referred to as the *correspondence problem* - is the most crucial task in order to capture true geometric variations, as false identification of point pairs may lead to unreasonable results in the alignment or interpolation process (Fig. 1.2).

The matching problem can be approached from three different angles: one is to determine corresponding points by spatially aligning two shapes first and then computing a - in general - nonlinear transformation between them, e.g. by minimizing surface distance; the second is to compute pairs of corresponding points *invariantly*, i.e. independent of the location and size in three dimensional space, and to use this information to align the shapes; the third is a mixture of both.

As the shapes to be compared are acquired from different individuals at different points in space and time, and may spatially vary significantly, three-dimensional spatial overlap of the shapes, such as closest points on different surfaces, in general is not sufficient to correctly identify corresponding points. Therefore, the focus of this work is on computing correct correspondences from an *invariant* point of view.

How can corresponding points on different surfaces be characterized? Naturally, salient *features* that are evident on different surfaces should be put into correspondence. Such features may be of geometric nature, like e.g. regions with large curvature. But it may also be specified by other information from the perceptual or semantic context of the shape. For example, it may be some muscle attachment site on a bone, specified by the location where a vessel enters some organ, or additional image-based features in a photographic texture. Another problem is that, in general, for a feature that is present on one shape there may not exist an obviously associated feature or any feature at all on the other shape.

Solving the problem of surface matching as stated before is an ambitious goal and is still not completely solved. The paradigm followed in this work is to separate two difficult problems, that is

- (a) the identification and matching of common salient features, and
- (b) the matching in regions around these features

of different objects. In this work we will focus on part (b) of the problem. Although both issues are often treated simultaneously, part (a) really is the subject of its own field of research and generally requires separate treatment. It is also strongly dependent on the type of shape to be treated and the underlying application.

Identifiable salient features that are common to a set of individual representatives of a given object will generally be sparsely distributed on the shapes. We therefore consider them as constraints in part (b) of the matching process. Once the features have been specified the only invariant information left to consider is of geometric nature. The geometric content can roughly be split into two ingredients: metric information and extrinsic curvature. Extrinsic curvature often is what we perceive as the feature of a surface, and it will thus be treated that way.

One important application of shape matching is the statistical analysis of the variability of a certain object between many different individuals within some data base. It requires matching of a large number of shapes. This task may be performed simultaneously (*group-wise* matching) or sequentially by matching *pair-wise* only. Group-wise matching is computationally much more demanding than pair-wise matching. In this work, we will be concerned with pair-wise matching only.

1.3 Outline and Contributions

Part I - Shape Modeling (Chapters 2 through 4) of this thesis deals with surface matching methods and statistical shape modeling.

In Chapter 2 we formalize our view on the problem of surface matching. We introduce basic geometric notions and concepts which play a central role in this work. We classify and review related work on registration and surface matching, and finally identify the methods most closely related to our work. At the end of the chapter we describe a specific method, which we call *consistent surface decomposition and parameterization*, in more detail. This interactive approach is conceptually simple and can be applied to solve a large class of real-world matching problems, such as the generation of statistical shape models. It has originally been invented by Malte Zöckler et al. [84] to produce morph sequences. We realized its usefulness for the solution of the correspondence problem for anatomical shapes and the statistical analysis of shape variations. We replaced the original barycentric patch parameterization method with shape preserving and mean-value based convex-combination maps in order to approximately realize the idea of minimizing geometric distortion. Additionally, we allow patches to touch themselves, which potentially reduces the number of cuts that have to be performed. These extensions have been published by the author of this thesis and co-workers in [71, 72, 75].

Both from a theoretical as well as from a practical point of view the methods discussed in Chapter 2 have certain drawbacks, which are discussed at the end of Chapter 2. Therefore, we describe a new variational approach to shape matching in Chapter 3, which is based on minimizing isometric stretching between two surfaces that are represented as triangle meshes. The main contributions of our approach are

- a discrete multi-level representation of the bijective correspondence map, where the discretization is *de-coupled* from the resolution of the underlying surfaces.
- a new discrete functional that *robustly measures isometric distortion* between two triangle meshes based on the 1-skeleton of the discrete correspondence map.
- a *manifold-based optimization* method, which does not require the computation of a global surface parameterization.

We describe a prototype implementation of the optimization algorithm, show some first results and discuss the advantages and limitations of our approach.

Chapter 4 deals with the generation and analysis of statistical shape models as a basis for the applications discussed in Chapters 5 through 7. The basic concept of statistical shape models was proposed by Cootes et al. [28]. The setup of statistical shape models requires the solution of the correspondence problem. Based on such solutions shapes can be transformed into elements of a vector space and traditional statistical methods can be applied. The main goal of such an analysis is to capture the *essential degrees of freedom*

of variations within a well-defined training database, e.g. by using principal component analysis.

Part II - 3D Geometry Reconstruction (Chapters 5 through 7) of this thesis contain a variety of new contributions to problems of three-dimensional geometry reconstruction. Most of the contents of these chapters have been published by the author of this thesis and co-workers in different peer-reviewed fora (see citations below). The mathematical tool that is central to all these applications is the statistical shape model. Its use is motivated by the insight that the incorporation of a-priori knowledge about shape and other application-specific features allows to progress in the field of geometry reconstruction problems.

Chapter 5 presents a framework for automatic segmentation of three-dimensional medical images, which allows to incorporate additional knowledge about imaging characteristics. Our approach is capable of robustly segmenting medical images even in the presence of noise or artifacts, as is the case for low-dose CT data in dental imaging (see Sec. 5.4 and [79]). The method for automatic segmentation of liver tissue from contrast-enhanced CT data (see Sec. 5.5 and Kainmüller et al. [63]) is almost as accurate as human performance. With the proposed approach we were able to win the MICCAI 2007 contest “Segmentation in the Clinics - a Grand Challenge”. The method extends previous work by the author of this thesis and co-workers [71, 73, 74].

Chapter 6 discusses a method for generating surgical reconstruction proposals for complex interventions. Two applications are presented: the reconstruction by mandibular dysplasia (see Sec. 6.3 and Zachow et al. [144]), and craniosynostosis planning (Sec. 6.4). In the case of skull deformities due to craniosynostosis a first clinical trial could be performed based on our method [51, 76, 78, 83].

For diagnostic purposes X-ray imaging is a widely used modality. We show that it is possible to estimate the three-dimensional shape of anatomical structures from such two-dimensional images using statistical shape models [77].

Part I

Shape Modeling

Chapter 2

The Correspondence Problem for Surfaces

2.1 Introduction

There exists a huge variety of approaches for establishing correspondence between surfaces, for which it is almost impossible to give a full account. In this chapter we will first formalize what we understand as the basic requirements of surface matching in a principal way. As we are interested in purely geometric matching, we summarize the most basic definitions and facts about geometric characterization of surfaces. We will try and state what kind of approaches exist for solving the correspondence problem and discuss their properties and limitations. Finally, we focus on a particular approach that has originally been developed to produce morph sequences, but has turned out to be a powerful tool to produce surface matchings in real-world applications. It satisfies most of the stated requirements; yet, it has some drawbacks that will be addressed in the next chapter.

2.2 Problem Statement

The fundamental requirement in registration or matching of two different surfaces is the identification of corresponding points. In this work, we will consider the problem of pair-wise matching of surfaces, which are isometrically embedded into \mathbb{R}^3 .

Definition 2.1 (Surface). *A two dimensional compact manifold $M \subset \mathbb{R}^3$ with a boundary, which is connected and oriented, is called a surface. If the manifold is C^k -differentiable, then M is called a C^k surface.*

Definition 2.2 (Homeomorphism). *Let M and \widetilde{M} be two surfaces. A map $\varphi : M \rightarrow \widetilde{M}$ is called a homeomorphism if it is bijective, continuous and has a continuous inverse φ^{-1} . We say that M and \widetilde{M} are homeomorphic.*

Definition 2.3 (Diffeomorphism). *If a homeomorphism φ and its inverse φ^{-1} are differentiable, it is called a diffeomorphism. The set of diffeomorphisms from M to \widetilde{M} will be denoted by $\text{Diff}(M, \widetilde{M})$. We say that M and \widetilde{M} are diffeomorphic.*

Solving the *correspondence problem* is the task of computing a homeomorphism (or diffeomorphism)

$$\varphi : M \rightarrow \widetilde{M}$$

between two homeomorphic surfaces M and \widetilde{M} , meeting the following criteria:

Invariance. The correspondence map is independent of the location and size of the surfaces M and \widetilde{M} in \mathbb{R}^3 .

Symmetry. If φ is the result of matching M to \widetilde{M} and ψ the result of matching \widetilde{M} to M then $\varphi \circ \psi = \text{id}$.

It should be possible to integrate specific requirements of the underlying application into an algorithm for computing the correspondence map. Generally, it is desirable to achieve the following:

Plausibility. The correspondence map matches salient features on M and \widetilde{M} in the context of a specific application, e.g. the tip of the nose on two faces, and performs *reasonable* interpolation between the features.

Efficiency. The method for computing φ is conceptually simple, i.e. its parameters are easy to interpret and to adjust to given specific applications.

2.3 Geometric Characterization of Shape

As mentioned before, in this work we focus on the correspondence problem on a purely geometric basis. To this end, we will introduce here the basic ingredients needed for the discussion of related work as well as our proposed method in Chap. 3. For further details consult any textbook on differential geometry and Riemannian geometry such as [33, 34].

2.3.1 Fundamental Geometric Concepts and Notation

Definition 2.4 (Riemannian Metric). *A Riemannian metric g is a correspondence that associates to each point $p \in M$ a symmetric, bi-linear, positive definite form $g_p : T_p M \times T_p M \rightarrow \mathbb{R}$, that varies smoothly with respect to p .*

2. THE CORRESPONDENCE PROBLEM FOR SURFACES

The pair (M, g) is called a Riemannian surface. The metric g allows to measure lengths, areas or angles. All quantities that depend on the metric g only are *intrinsic*, i.e. they do not depend on the embedding of the surface. Any smooth surface can be equipped with a Riemannian metric, e.g. with the Euclidean inner product in \mathbb{R}^3 . In the remainder of this work, all manifolds will be considered Riemannian.

Definition 2.5 (Shape Operator). *Let $n : M \rightarrow \mathbb{R}^3$ denote the unit normal vector field on M . The shape operator (or Weingarten map) $S_p : T_p M \rightarrow T_p M$ is defined by the differential of the unit normal field $S_p = -dn_p$.*

S measures the degree to which the normal is changing in a direction tangential to the surface M . In general, it is an *extrinsic* quantity, i.e. it depends on the embedding of the surface in \mathbb{R}^3 .

Definition 2.6 (Second Fundamental Form). *Since S is a self-adjoint linear map [33], there exists a symmetric bi-linear form $\Pi : T_p M \times T_p M \rightarrow \mathbb{R}$ that is associated with S by $\Pi(X, Y) = g(S(X), Y)$. Π is called the second fundamental form.*

The quantity $\Pi(X, X)$ for a unit vector $X \in T_p M$ is the *normal curvature* along a curve passing through p tangential to X . The eigenvectors of the shape operator S are called principal curvature directions, their associated eigenvalues principal curvatures κ_1, κ_2 . Two important curvature scalars are the *Gauss curvature* κ and the *mean curvature* H

$$\kappa := \det S = \kappa_1 \kappa_2 \tag{2.1}$$

$$H := \operatorname{tr} S = \frac{1}{2}(\kappa_1 + \kappa_2) \tag{2.2}$$

Gauss's Theorema egregium states that the Gauss curvature κ depends only on the Riemannian metric and its derivatives and not on the second fundamental form. This implies that κ is an intrinsic property. The mean curvature H , on the other hand, is an extrinsic quantity.

2.3.2 Congruency

For the remainder of this section we will assume that all maps φ are diffeomorphisms, i.e. $\varphi \in \operatorname{Diff}(M, \widetilde{M})$. In this case it is possible to define the differential of φ at $p \in M$ as the linear map (see Fig. 2.1)

$$d\varphi_p : T_p M \rightarrow T_{\varphi(p)} \widetilde{M}. \tag{2.3}$$

The pullbacks of the metric \widetilde{g} and the second fundamental form $\widetilde{\Pi}$ of the




Figure 2.1 The action of $d\varphi_p$: tangent vectors X, Y in $T_p M$ are mapped to the tangent space $T_{\varphi(p)} \widetilde{M}$.

surface \widetilde{M} by the map φ are defined as

$$(\varphi^* \widetilde{g})_p(X, Y) := \widetilde{g}_{\varphi(p)}(d\varphi_p(X), d\varphi_p(Y)) \quad (2.4)$$

$$(\varphi^* \widetilde{\Pi})_p(X, Y) := \widetilde{\Pi}_{\varphi(p)}(d\varphi_p(X), d\varphi_p(Y)) \quad (2.5)$$

where $X, Y \in T_p M$.

Definition 2.7 (Congruency). *Two surfaces M and \widetilde{M} are said to be congruent, if they differ only by an Euclidean transformation (rigid motion or translation and rotation).*

The following theorem characterizes surfaces up to congruency (refer to do Carmo [33], Spivak [119] and Ciarlet [24] for a more concise exposition and proofs):

Theorem 2.8 (Fundamental Theorem of Surface Theory). *(1) Existence: Let g and Π be two differentiable forms, that satisfy the Gauss and Codazzi-Mainardi equations. Then there exists a surface M with metric g and second fundamental form Π . (2) Uniqueness: Let $g = \varphi^* \widetilde{g}$ and $\Pi = \varphi^* \widetilde{\Pi}$ for two surfaces M and \widetilde{M} and a diffeomorphism φ . Then M and \widetilde{M} are congruent.*

Theorem 2.8 implies that two surfaces M and \widetilde{M} are congruent if their Riemannian metric and second fundamental forms are equal. Yet, there exists some redundant information described by the compatibility conditions (Gauss and Codazzi-Mainardi equations, see [33]): the Gauss curvature, for instance, is an intrinsic property. This naturally leads to the question whether metric information in combination with some scalar curvature value suffices to guarantee congruency. Bonnet [15] noticed the mean curvature H (or, equivalently, any other extrinsic scalar curvature) does the job in most cases. However, counter-examples do exist that are called *Bonnet mates*:

Definition 2.9 (Bonnet mates). *The surfaces M and \widetilde{M} with $g = \varphi^*\widetilde{g}$ are called Bonnet mates if they are not congruent, but their mean curvatures H and \widetilde{H} are equal.*

Instead of classifying all surfaces that do not have Bonnet mates it is possible to construct all surfaces which do have Bonnet mates [65]. It could not be proved yet whether any of these surfaces is closed. On the other hand it was shown that for compact surfaces there are at most two surfaces with a given metric and mean curvature [50]. Uniqueness can be established under various global assumptions (see Kamberov and Kamberova [64]).

2.3.3 Isometry

Instead of considering congruency, a weaker characterization of surface geometry may be based on the concept of

Definition 2.10 (Isometry). *The map φ is called an isometry if the length of any rectifiable curve on M is the same as that of its image on \widetilde{M} . If such a map exists, M and \widetilde{M} are said to be isometric.*

Congruency implies isometry, but the converse is generally not true. Important notions related to isometry are based on analyzing properties of the length function $L(\gamma)$ of curves γ on M .

Definition 2.11 (Geodesics). *A geodesic $\gamma : \mathbb{R} \rightarrow M$ is a locally shortest curve on M , i.e. γ is a critical point of the length function $L(\gamma)$ with respect to variations tangential to M that leave the endpoints of γ fixed.*

Definition 2.12 (Distance). *The distance $d : M \times M \rightarrow \mathbb{R}$ between two points $p, q \in M$ is defined as $\inf_{\gamma} L(\gamma)$ for all rectifiable curves $\gamma : [0, 1] \rightarrow M$ with $\gamma(0) = p$ and $\gamma(1) = q$.*

There exist different equivalent characterizations of geodesics, e.g. that the geodesic curvature κ_g vanishes along γ [69, §50]. Another useful notion is the metric distortion tensor [58]:

Definition 2.13 (Metric Distortion Tensor). *There exists a symmetric positive definite 2×2 matrix field Q_p for $p \in M$, uniquely defined almost everywhere, such that*

$$g_p(Q_p(X), Y) = \widetilde{g}_{\varphi(p)}(d\varphi_p(X), d\varphi_p(Y)) \quad (2.6)$$

for all vector fields X, Y at point p on the surface M .

There are different equivalent characterizations of isometry:

Lemma 2.14 (Characterizations of Isometry). *The following statements are equivalent (refer to [69, §57] for a proof):*

1. M and \widetilde{M} are isometric.
2. $\exists \varphi \in \text{Diff}(M, \widetilde{M}) : d(p, q) = \widetilde{d}(\varphi(p), \varphi(q))$ for all $p, q \in M$.
3. $\exists \varphi \in \text{Diff}(M, \widetilde{M}) : g_p = (\varphi^* \widetilde{g})_p$ for all $p \in M$.
4. $\exists \varphi \in \text{Diff}(M, \widetilde{M}) : Q_p = \text{id}$ for all $p \in M$.

2.4 Classification of Shape Matching Methods

In many cases, correspondence as well as alignment and interpolation are treated simultaneously within a single approach. Such methods are often tailored to specific applications and thus are not easily adaptable to other types of problems. It is impossible to give a complete list of works concerned with or contributing to the problem of surface matching. Therefore we try to our best to give a classification of existing approaches most closely related to this work - and remain aware that within the scope of the present work we were forced to be selective to a certain degree.

There exist some surveys on shape matching and correspondence: Lazarus and Verroust [80] review 3D matching with a focus on shape transformations, Alt and Guibas [5] focus on geometric techniques based on the Hausdorff distance, Audette et al. [8] place their overview in the context of medical imaging, while Veltkamp and Hagedoorn [134] emphasize similarity measures. More recently, Planitz et al. [101] have presented an attempt to classify a set of correspondence methods, which are popular in the pattern recognition community. Most authors agree that matching algorithms can be classified via four different elements:

1. shape representation
2. type of transformation
3. similarity measure
4. algorithm

Other aspects are matching of partial, multiple or non-homeomorphic shapes, none of which fall under the problem described in Sec. 2.2. They are not considered in this work.

2.4.1 Shape Representations

A surface M may be represented either parametrically as the image of a map $S : \mathbb{R}^d \rightarrow \mathbb{R}^3$, or implicitly as the level-set $\{x \in \mathbb{R}^3 : F(x) = 0\}$ of a smooth function $F : \mathbb{R}^3 \rightarrow \mathbb{R}$. Such a function F is typically computed via the distance transform [29]. Implicit surfaces [14, 98, 114] are widely used in computer graphics and image processing, as they allow for handling topological changes in a simpler way than it is the case with parametric surfaces. However, in shape matching they increase the complexity of the problem by introducing an additional dimension.

Meshes, i.e. piecewise linear polygonal surfaces, are a special case of non-smooth parametric surfaces heavily used in computer modeling and graphics, due to conceptual simplicity that allows for flexible and efficient processing [19]. Smooth parametric surfaces, such as spline or nurbs surfaces [100], are commonly used in computer aided engineering [111]. Other parametric representations include: spherical harmonics [45], medial axis representations [17] or point-based surfaces [6]. In this work, we will be dealing with triangle meshes when it comes to computational issues.

Furthermore, shapes may be represented approximately by other mathematical objects, e.g. graph, skeleton or point representations. It is beyond the scope of this work to discuss methods based on representations other than surfaces, although such methods do have their eligibility.

2.4.2 Transformations

The correspondence map is not always computed directly as a map $\varphi : M \rightarrow \widetilde{M}$ between manifolds. It is also rather common to compute a map $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ between the embeddings of the shapes in their ambient space. The term *registration* is often used to refer to this setting. In general, ϕ is non-linear. One distinguishes between non-parametric (see e.g. Modersitzki [92] for a detailed exposition) and parametric methods, e.g. based on thin-plate [16] or Bezier splines [109, 67].

Methods that operate directly on the surfaces and not their ambient space can be divided into two classes: (a) A common intermediate parameterization of both M and \widetilde{M} is computed and the problem is solved on the parameter domain [84, 48, 82]. (b) A map between the manifolds is computed directly [116, 20].

2.4.3 Similarity Measures

Let us consider a parametric representation of two surfaces $S, T : \mathbb{R}^d \rightarrow \mathbb{R}^3$ for the moment. The only basic ingredients of geometric similarity measures are the values of S, T and their derivatives. Probably the best known and

most widely used measure is surface distance [12, 23], i.e. a measure that involves only evaluations of S and T itself and as such is not invariant. Measures based on first- and higher-order derivatives, however, are invariant. *First-order* measures essentially quantify metric deviations (stretching): e.g. geodesic distances [20, 130] or elastic deformations such as the Green-St Venant-Kirchoff energy [99]. *Second-order* derivatives measure curvature deviations (bending) of the surfaces. Terzopoulos et al. [125] introduced a measure that is often called thin-shell energy, assessing the discrepancy between the first and second fundamental forms at corresponding points. Similar energies are used by Gu and Vemuri [48] and Litke et al. [82].

In the literature those basic ingredients are combined into a variety of different similarity measures. Depending on the underlying application one part of such a measure is called the *regularization* while the rest is referred to as the *feature* term. It is the task of the feature term to match characteristic regions of the shapes, while the regularization shall ensure well-posedness and certain degrees of regularity of the solutions.

The measures can also be interpreted as scalar fields on the surfaces, which allow to adopt the similarity measures known from image registration [146] to surface matching, as was done by Wang et al. [136] for mutual information, a popular information theoretic measure in image registration. However, this approach requires a common global parameterization of both surfaces. Furthermore feature descriptors (also termed indicators or fingerprints) may be used to measure similarity [108, 91, 61]. A survey of such methods is presented by Mikolajczyk and Schmid [88].

There exist many approaches that are beyond pure geometric characterization of shapes. Although not in the focus of this work, we will mention two different approaches that combine geometric with statistical information:

Pennec et al. [99] introduce the statistical Riemannian elasticity framework which combines an elastic energy with statistical information from a set of learned diffeomorphisms. This approach requires that a number of prior matchings have been performed on a carefully selected training set, which reflects the properties of the current shapes to be matched. It has not been studied how the training data effects the result.

Twining et al. [128] combine statistical information based on the information theoretic concept of minimum description length (MDL) [31] with a fluid-based registration approach [22]. Here the idea is to describe the shapes as viscous fluids using the Navier-Stokes equation to drive a deformation. The MDL approach is a group-wise matching approach that requires the solution of a costly optimization problem. Its focus is on generating statistical shape models (see Chap. 4).

2.5 Review of Geometric Shape Matching Algorithms

We restrict our discussion here to algorithms that deal with invariant surface matching based on purely geometric considerations. We have identified the following methods as being most closely related to our work.

2.5.1 Matching based on Surface Parameterization

Both Gu and Vemuri [48] and Litke et al. [82] parameterize meshes M and \widetilde{M} onto a common base domain and compute the correspondence map $\phi : \Omega \rightarrow \Omega$ on the parameter domain $\Omega \subset \mathbb{R}^2$. This setup has the advantage that efficient numerical optimization schemes are far simpler to devise on analytically accessible parameter domains (e.g. plane or sphere) than over arbitrary curved domains. Particularly, the resampling of the parameter domain is generally much simpler, thus it makes the decoupling of the discretization of the original meshes from a discretization of ϕ much easier. Both approaches are symmetric, and require the specification of two parameters. Additional constraints may be incorporated.

Gu and Vemuri [48] use the method of conformal parameterization [49] for meshes of spherical topology, and apply a quasi-Newton method to minimize their energy. The energy is given by

$$E(\phi) = \alpha \int_{\Omega} \|\lambda(x) - \widetilde{\lambda}(\phi(x))\|^2 dx + \beta \int_{\Omega} \|H(x) - \widetilde{H}(\phi(x))\|^2 dx$$

where λ is the conformal factor, measuring stretch and computed as the result of a conformal surface parameterization method. H is the mean curvature of a surface.

Litke et al. [82] use the method of Clarenz et al. [25] to parameterize meshes with disc-like topology. They apply a finite element discretization and optimize the matching energy using a multi-scale approach. Their energy is similar to that of Gu and Vemuri [48] and is given by

$$E(\phi) = \alpha \int_{\Omega} W(\text{tr}Q(x), \det Q(x)) \sqrt{g(x)} dx + \beta \int_{\Omega} \|H(x) - \widetilde{H}(\phi(x))\|^2 \sqrt{g(x)} dx$$

where $Q(x)$ is the metric distortion tensor at a point $x \in \Omega$ (see Def. 2.13 on page 12) and the energy density W accounts for length and area distortion.

2.5.2 Intersurface Mapping

The restriction to disc- or sphere-like topology is overcome by the method of Schreiner et al. [116], which does not require a common parameterization.

They base their method on a so-called meta-mesh which may be automatically embedded on both surfaces M and \widetilde{M} and thus allows to directly evaluate their distortion energy on corresponding planar faces of the meta-mesh. They construct a first-order symmetric similarity measure based on the L^2 stretch metric [112] per triangle, which smoothly penalizes scale distortion. The match is optimized by local relaxation steps of the vertices of the meta-mesh on different resolutions of the input meshes. The relaxation of the meta-mesh vertices is a complicated process which requires careful treatment in order to maintain bijectivity of the map. Moreover, the discretization of the correspondence map ϕ is coupled directly to the resolution of the input meshes, and not to the error in the similarity measure.

The following two approaches were originally invented to detect isometric deformations of objects, but can generally be used for shape matching as well. Both methods can handle partial matchings and are intended to operate in an unsupervised fashion.

2.5.3 Multi-dimensional Scaling

The approach of Elad and Kimmel [38] is based on the concept of isometric embedding and solved by applying a multi-dimensional scaling technique (MDS). It was extended by Bronstein et al. [20] to non-Euclidean embeddings and thus termed generalized multi-dimensional scaling (GMDS). It works directly on the manifolds and does not require the computation of an intermediate parameterization. The idea of measuring similarity is based on the (non-local) characterization of isometry given in Lemma 2.14, item 2 (page 13):

$$E(\varphi) = \sup_{p,q \in M} |d(p,q) - \widetilde{d}(\varphi(p), \varphi(q))| \quad (2.7)$$

where d denotes geodesic distance on M . Since this problem is apparently intractable, a modified functional is minimized based on mapping a discrete number of point pairs $\widetilde{p}_i = \varphi(p_i)$ only. The problem is reduced to the finite dimensional problem

$$E(\widetilde{p}_1, \dots, \widetilde{p}_N) = \sum_{i>j} |d(p_i, p_j) - \widetilde{d}(\widetilde{p}_i, \widetilde{p}_j)|^2$$

which is minimized using a gradient descent algorithm combined with a multi-resolution strategy to avoid getting stuck in insignificant local minima. The geodesic distances d, \widetilde{d} are pre-computed for the mesh vertices and stored in a matrix. The values of d on the interior of the triangles are interpolated using a so-called three-point geodesic distance approximation described by Bronstein et al. [21]. The number of embedded points is typically of the order of 10-100. Meshes may not grow too large in order to be able to store the geodesic distance matrix. The correspondence map between the sample points

is interpolated by inverse square distance-weighting. A remarkable fact is that the geodesic distance can be estimated robustly even on noisy triangle meshes.

2.5.4 Probabilistic Matching

Anguelov et al. [7] define a probabilistic model over all correspondences $\tilde{p}_i = \varphi(p_i)$, which encodes geodesic distance constraints as well as other high-level surface features (spin images, see Johnson [62]). They derive a Markov network, which models a joint probability distribution over the correspondence pairs and apply the loopy belief propagation algorithm [142] to solve for the most likely configuration. No details on how to compute the geodesic distances are given in this paper. Yet an interpolation of the correspondence map could be performed in the same manner as in the above approach. The main limitation of this approach is the fact that it makes the assumption of (approximate) preservation of geodesic distance. A second limitation of the approach is that it assumes that the data mesh is a subset of the model mesh.

2.6 Consistent Surface Decomposition and Parameterization

In this section we will discuss one specific geometric shape matching algorithm, which has proved successful in many real-world applications, especially for generating statistical shape models (Chap. 4) of anatomical structures. It can be considered as an extension to the surface parameterization approaches (Sec. 2.5.1), that on the one hand can handle arbitrary topological meshes. On the other hand it is conceptually and computationally much simpler than the intersurface-mapping algorithm (Sec. 2.5.2).

2.6.1 Motivation

A general problem in surface matching is the combined consideration of features and regularizations. As mentioned in the introductory Chap. 1.2, we strive to devise a framework where features are incorporated as hard constraints in a regularization process. This way the definition of features is completely decoupled from the regularization criterion. Variational methods that use additive coupling of regularization and feature terms (soft constraints) suffer from the difficulty of specifying the coupling constant. It may be possible to add hard constraints to unsupervised methods (Sec. 2.5.3 and 2.5.4), like GMDS or probabilistic matching, but this has not yet been dealt with.

Typical features or constraints in matching applications are corresponding points, corresponding lines or whole corresponding regions where exists no detailed knowledge of how to map the interior of such regions or lines. Hence, our idea to decompose the surfaces into corresponding parts in a consistent

way and match these parts separately by considering suitable compatibility conditions along the boundary of those regions. This makes sense if matching the parts as well as the subsequent assembly is significantly simpler than solving the whole problem at once. Moreover, this implies that the features should be contained in the set of boundaries of these parts, not on their interior. This idea was inspired by the work of Malte Zöckler et al. [84]. Their goal was the generation of visually pleasing morph sequences. Our contribution is the realization that such an approach can be applied to solve the correspondence problem as posed in Sec. 2.2.

2.6.2 Consistent Decomposition

There are several ways to define such decompositions. A major goal, however, is that the matching of the parts and their assembly can be performed for arbitrary surfaces and allows to incorporate features in a simple manner. For instance, a simple scheme can be devised for parts that can efficiently be parameterized globally onto common domains. In this case, the matching $\varphi|_R$ of two corresponding regions $R \subset M$ and $\tilde{R} \subset \tilde{M}$ with given parameterizations $x : \Omega \rightarrow R$ and $\tilde{x} : \Omega \rightarrow \tilde{R}$ can be computed by concatenating their parameterizations:

$$\begin{array}{ccc}
 R & \xrightarrow{\varphi|_R = \tilde{x} \circ x^{-1}} & \tilde{R} \\
 & \swarrow x & \nearrow \tilde{x} \\
 & \Omega &
 \end{array} \tag{2.8}$$

The most efficient parameterization methods for given base domains exist for regions that are homeomorphic to a disc. For triangle meshes there exist a variety of different methods for computing parameterizations of disc-like surfaces, see e.g. Floater and Hormann [43] for a classification and overview over current approaches. Conformal methods are widely used as they minimize angular distortion, which is desirable when dealing with triangle meshes. Harmonic maps [37] using finite element discretization [36] have the drawback of not guaranteeing bijectivity. This property is explicitly enforced in the so-called convex-combination map approach [41]:

Let $P = \{p_1, \dots, p_n, p_{n+1}, \dots, p_N\}$ be the coordinates of the vertices of a triangulated surface patch R . Assume that p_1, \dots, p_n are internal nodes, while p_{n+1}, \dots, p_N lie on the boundary ∂R . Furthermore let $u_i \in \Omega \subset \mathbb{R}^2$ ($i = 1, \dots, N$) be the corresponding coordinates in the parameter domain Ω . The map

$$x^{-1} : P \rightarrow \Omega, p_i \mapsto u_i = x^{-1}(p_i)$$

2. THE CORRESPONDENCE PROBLEM FOR SURFACES

is called a convex combination map, if

$$u_i = \sum_{j=1}^N \lambda_{ij} u_j \quad (2.9)$$

for all internal nodes $i = 1, \dots, n$, and

$$\sum_{j=1}^N \lambda_{ij} = 1 \quad , \quad \lambda_{ij} \begin{cases} = 0 & (i, j) \notin E \\ > 0 & (i, j) \in E \end{cases} \quad (2.10)$$

where E is the set of all edges of the triangulation. It was shown by Floater [41] that convex combination maps do not produce fold-overs in the planar domain Ω , and every point u_i lies in the convex hull of $\partial\Omega$. The linear system (2.9) is sparse and can be solved efficiently even for large triangle meshes. The most simple choice for $\lambda_{ij} = 1/n_i$ are barycentric weights, where n_i are the number of neighbors of point p_i in the triangle mesh. However, the freedom to choose λ_{ij} can be exploited to minimize geometric distortion. We use the so-called mean value coordinates by [42], which satisfy the mean value theorem of harmonic functions.

Let us now formalize what we mean by consistent decomposition of two surfaces, and how the parts R are assembled to produce a continuous correspondence map φ from the set $\{\varphi|_R : R \subset M\}$. Our basic assumption is that each region R is homeomorphic to a disc. Let $K = (V, E, C)$ be an abstract 2-dimensional cell complex with

$$\begin{aligned} V &= \{1, \dots, n\} \quad \text{the set of vertices,} \\ E &= \{(i, j) \mid i, j \in V\} \quad \text{the set of edges and} \\ C &= \{(i_1, \dots, i_k) \mid i_1, \dots, i_k \in V, k > 2\} \quad \text{the set of 2-dimensional cells.} \end{aligned}$$

Each vertex i has a corresponding point p_i embedded in M and \tilde{p}_i embedded in \tilde{M} . The geometric realization of an edge $e \in E$ is a curve $\gamma_e : \mathbb{R} \rightarrow M$, that does not intersect any other curve $\gamma_{e'}$ except at their endpoints. Cutting the surface M along the union of embedded curves $\Gamma := \bigcup_{e \in E} \text{Im}(\gamma_e)$ produces a decomposition of M into regions R_c , equivalent to the geometric realizations of cells $c \in C$.

Definition 2.15 (Consistent Surface Decomposition). *We call a decomposition of M valid, if $M \setminus \Gamma = \bigcup_{c \in C} R_c$ degenerates into disconnected surface patches each homeomorphic to a disc. Two valid decompositions based on the same cell complex K are said to be consistent. Note that a cell of of such a decomposition may touch itself, see Fig. 2.2.*

For each pair of corresponding patches R_c and \tilde{R}_c a common base domain $\Omega_c \subset \mathbb{R}^2$ is defined. The shape of the domain influences the amount of distortion introduced by the parameterization method. Generically, the unit disc




Figure 2.2 A special decomposition of a double-torus. The whole surface is cut into a single region homeomorphic to a disc. The picture was taken from Erickson and Har-Peled [39].

may be chosen, but other constructions are conceivable as well. In order to achieve continuity of the correspondence map \widehat{m} across boundary curves γ_e , vertices on M must be mapped to vertices on \widetilde{M} , while boundary curves must be mapped to common intervals. This must be considered in the construction of the boundary $\partial\Omega_c$ of the common base domains *before* parameterization can be performed. A natural way of fixing the location of the vertices on $\partial\Omega_c$ is common arc-length parameterization, i.e. averaging the lengths of the boundary curves $\underline{\gamma}_e$ and $\widetilde{\gamma}_e$ relative to the total length of curves bounding a cell c in M and \widetilde{M} , respectively. The final correspondence map φ at point p is now given by φ_c for c such that $p \in R_c$.

2.6.3 Implementational Issues

In order to implement the above method triangle meshes are cut along user-defined curves on surfaces. This is achieved by constrained Delaunay triangulation of the original mesh (Fig. 2.3a) per triangle along the curve (Fig. 2.3b). Our implementation is based on the *Triangle* library by Shewchuk [117].

This procedure generally introduces triangles with bad aspect ratio, which have to be avoided in the further course of processing, e.g. to improve the condition of the linear system (2.9) in the mesh parameterization process. We solve this problem by re-meshing the surface (Fig. 2.3c), while preserving feature lines. It is based on local surface operations only and does not require previous parameterization of the surface. Please refer to Zilske et al. [145] for details.

In our algorithm, we allow surface regions to touch themselves, which po-

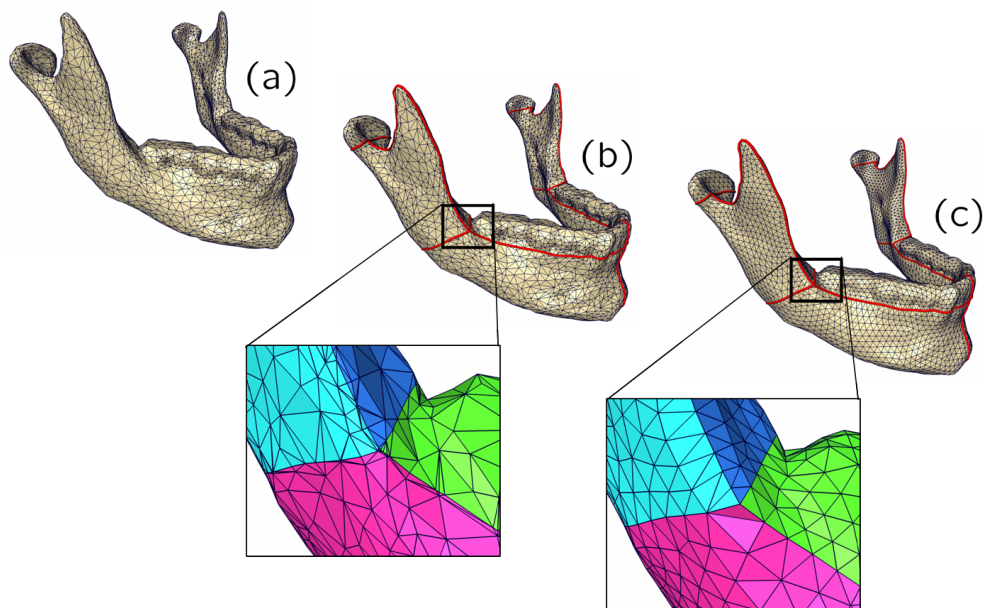


Figure 2.3 Decomposition of a surface mesh: (a) Original mesh, (b) Definition of cutting curves: the close-up shows the mesh after re-triangulation along the cuts. (c) Re-meshing of the mesh such that the geometry of the cutting curves is preserved while the triangle quality is improved.

tentially reduces the number of cuts that have to be performed. In general, every closed surface may even be cut into a single region homeomorphic to a disc (Fig. 2.2).

2.6.4 Discussion

Fig. 2.4 shows the decomposition of a pelvic bone and a liver surface. The pelvis is cut into 11 regions, the liver is composed of 4 regions. Generally not all cuts can anatomically be well-defined, as shown in Fig. 2.5. Some cuts are needed merely for achieving topological consistent decomposition. As cuts on different shapes will be matched and, generally, have different (uncorrelated) arc-length parameterizations, unwanted correspondences may be introduced along the cuts. More details about the decomposition of specific shapes are given in Chap. 5 through 7.

The fact that the parameterization of the cutting lines is fixed independently of the interior of the regions, leads to potentially large distortions in the resulting map, see Fig. 2.6. A possible remedy is a post-processing step that improves the mesh quality iteratively by moving vertices tangentially along the surfaces. This idea is part of the approach taken by Schreiner et al. [116] and is closely related to the concept of surface re-meshing.




Figure 2.4 Decomposition of a pelvic bone surface into 11 regions and a liver surface into 4 regions.




Figure 2.5 Some cuts must be defined for purely technical reasons. Such cuts are not well defined anatomically and thus possibly introduce unwanted correspondences.




Figure 2.6 Distortion of the correspondence map introduced by the method of consistent surface decomposition and parameterization: (a) mesh M , (b) mesh \tilde{M} with the same mesh topology as M , (c) large distortion of triangles occur mainly around the region boundaries.

Another limitation of this approach is that it is not obvious how to obtain a quantitative similarity measure within this construction. Currently, the surface decomposition process is interactive, which can be regarded as a drawback in certain applications. Efforts to further automatize this task have been proposed in the literature [104, 116]. Extensions have been and still are investigated by two diploma theses co-supervised by the author of this thesis [10, 141]. For a wide range of applications, however, such as statistical shape modeling (see Chap. 4), the possibility of interaction is beneficial, as it decouples feature identification from regularization.

2.7 Conclusions

Each discussed method for computing corresponding points on different shapes has its strengths and weaknesses, and has been shown to be successful on certain types of data. The goal must be to combine their strengths into one approach that meets the desired requirements stated in Sec. 2.2 and overcomes at least some of the issues of the discussed methods. These are - briefly summarized - the following:

- One class of methods relies on global parameterizations of the surface. This is only tractable for simple topologies (disc, sphere). For surfaces of higher genus the computation of atlases seems prohibitive.
- Non-local similarity measures, as in the MDL approach, limit the problem size which the optimization method is able to handle.
- Working directly on the surface mesh does not decouple mesh resolution from the discretization of φ .
- Splitting the matching problem into sub-problems, dealing with topologically simpler shapes, in order to handle arbitrary topology, makes it difficult to obtain a quantitative similarity measure, introduces unwanted correspondences and possibly large distortions around the interfaces between the different regions (Fig. 2.6).

In the next chapter we propose an attempt to solve some of these problems. The advantages of our approach will be that it can handle arbitrary topology without computing surface parameterizations explicitly. Furthermore it is based on a local characterization of isometric deviation. Our discrete distance functional approximates a smooth functional, and can be evaluated robustly on triangle meshes.

Chapter 3

Variational Surface Matching

3.1 Introduction

In this chapter we will introduce a variational approach to shape matching. The underlying idea is to measure only the deviation between the geometric content of surfaces. Basically there are two ingredients: metric and extrinsic curvature. Extrinsic curvature, however, is a property of a surface that is often understood to represent specific features of a shape. As stated before, such features may not exist on all shapes to be compared and it may not necessarily be desirable to match regions with similar extrinsic curvature. Hence we will measure metric deviations only.

First we derive a distance functional for smooth surfaces and examine the structure of the optimization problem. In real applications, however, we face two major problems: (1) Even for smooth surfaces of high genus it is virtually impossible to repeatedly compute surface parameterizations (atlases). (2) We are dealing with discrete surfaces. These aspects lead us to proposing a discrete distance measure which can be evaluated robustly on discrete meshes. We show how to compute its first variation and describe a minimization-on-a-manifold scheme, which does not require the computation of surface parameterizations. Finally, we present a prototype implementation and show some preliminary results of our method.

3.2 Measuring Geometric Deviation

As stated in Sec. 2.2, we are interested in characterizing surfaces not only modulo translation and rotation but also scale. Let A and \tilde{A} denote the total surface area of M and \tilde{M} , respectively. We can always assume scale invariance by simply re-scaling M and \tilde{M} such that $A = \tilde{A} = 1$.

3. VARIATIONAL SURFACE MATCHING

Let (M, g) and $(\widetilde{M}, \widetilde{g})$ be two Riemannian surfaces and

$$\varphi \in \mathcal{M} := \text{Diff}(M, \widetilde{M}) \quad (3.1)$$

a diffeomorphism. We want to derive a (family of) functional(s)

$$F : \mathcal{M} \rightarrow \mathbb{R}$$

such that its minimum

$$F(\varphi) \rightarrow \min \quad \text{subject to additional constraints} \quad (3.2)$$

is attained when the two surfaces M and \widetilde{M} are *geometrically most similar*. How can geometric similarity (or equivalently geometric deviation) be measured quantitatively?

Let Π and $\widetilde{\Pi}$ be the second fundamental forms of M and \widetilde{M} , respectively. An obvious way of measuring geometric deviation - neglecting symmetry considerations for the moment - could be

$$F(\varphi) = \alpha \int_M \|g_p - (\varphi^* \widetilde{g})_p\|^2 dA + \beta \int_M \|\Pi_p - (\varphi^* \widetilde{\Pi})_p\|^2 dA \quad (3.3)$$

where $\|\cdot\|$ denotes some suitable operator norm and dA the area element of M . The parameters $\alpha, \beta \in \mathbb{R}_{>0}$ weights metric and curvature deviation.

Functional (3.3) was introduced in the seminal work by Terzopoulos et al. [125]. It is closely related to the theory of elastic shells [24] and can be regarded as a simplified version of Koiter's energy for a nonlinear elastic shell. In our setting, the surface (M, g, Π) corresponds to the deformed shell, while $(\widetilde{M}, \widetilde{g}, \widetilde{\Pi})$ corresponds to the reference configuration. Although functional (3.3) is not derived from the theory of elastic materials, it does simulate

$$\text{stretching : } \alpha \|g_p - (\varphi^* \widetilde{g})_p\|^2 \quad \text{and} \quad (3.4)$$

$$\text{bending : } \beta \|\Pi_p - (\varphi^* \widetilde{\Pi})_p\|^2 \quad (3.5)$$

of an infinitely thin shell. Its parameters α and β , however, do not have a direct physical interpretation.

In computer graphics applications, functional (3.3) is used to model surface deformations. In this field, one is usually interested in generating physically plausible and aesthetically pleasing deformations at interactive rates. Therefore, functional (3.3) is often replaced by a linearized version and minimized on a given triangle mesh, which stays fixed during the deformation. See Botsch and Sorkine [18] for a review of such methods. For the task of computing correspondences such an approach does not seem appropriate, as the discretization of φ depends on the resolution of the surface M .

The bending term (3.5) is minimized when regions of similar curvature (and curvature directions) are matched. As pointed out in the introduction




Figure 3.1 Part of a cylinder M is matched with a copy \tilde{M} of itself that is rotated by 90° . The corner mappings are fixed. The bending term $\|S_p - (\varphi^* \tilde{S})_p\|$ in equation (3.5) drives the match φ of the dashed line on M towards the dotted line on \tilde{M} .

(see Chap. 1) such regions are often regarded as salient features of a shape. However, a feature that is present on one shape may not have one obvious associated feature or any feature at all on the other shape. Since the feature definition is *context-dependent*, we have decided to treat feature detection separately.

In order to illustrate the effect of the bending term, we consider a situation where no salient features shall be present on the two surfaces M and \tilde{M} depicted in Fig. 3.1. Both shapes are identical parts of a cylinder that are matched under the following constraints: boundaries of M shall be mapped to boundaries of \tilde{M} by identifying the corners of the shapes (A, B, C, D with $\tilde{A}, \tilde{B}, \tilde{C}, \tilde{D}$) and matching the thick black lines L and \tilde{L} . Without the bending term, the dashed lines on both shapes would be matched. Including the bending term, however, favors solutions φ that rotate the dashed line *towards* the dotted line. According to our paradigm, this may not be the desired solution when no additional features are given on the interior of the shapes.

Therefore, in this work, we will neglect extrinsic curvature and seek to minimize metric distance in a symmetric way, as defined in Sec. 2.2. We propose to minimize

$$F(\varphi, \varphi^{-1}) = \int_M \|g_p - (\varphi^* \tilde{g})_p\|^2 dA + \int_{\tilde{M}} \|\tilde{g}_q - ((\varphi^{-1})^* g)_q\|^2 d\tilde{A} \quad (3.6)$$

with additional constraints. Here, dA and $d\tilde{A}$ denote the area element of M and \tilde{M} respectively. In many cases, constraints are necessary to guarantee the existence of a unique solution. In this work we will assume that “sufficient” constraints are given. In the case of surfaces with boundaries, the mapping

of boundaries on M to boundaries on \widetilde{M} must be specified as constraints to ensure that φ is a diffeomorphism. Further constraints may be derived from an analysis of extrinsic curvature measures or other non-geometric information. Bihun and Chicone [13] discuss a similar functional to equation (3.6) based on the Jacobian of φ and prove the existence of minimizers.

3.3 Minimizing Stretching

In this section we will discuss the theoretical setting of the variational problem stated in equation (3.2):

$$F(\varphi) \rightarrow \min \quad , \quad \varphi \in \mathcal{M} := \text{Diff}(M, \widetilde{M})$$

such as the stretching functional defined in equation (3.6). Recall that M and \widetilde{M} are Riemannian surfaces as given in Def. 2.1. For the relevant background on Riemannian geometry refer to Appendix A.

The space \mathcal{M} is not a linear space, i.e. one cannot simply perform addition of two elements $\phi, \psi \in \mathcal{M}$ such that $\phi + \psi \in \mathcal{M}$. Yet, for the purpose of optimization, we must work with linear spaces.

The key to tackle this problem lies in realizing that \mathcal{M} is a manifold modeled on a Banach space [70]. Its tangent bundle $T\mathcal{M}$ can be identified with the tangent bundle $T\widetilde{M}$ of the target surface \widetilde{M} , see [13]. Furthermore, \mathcal{M} is a differentiable manifold [96] and can hence be equipped with a Riemannian metric. For more details about infinite dimensional manifolds and spaces of maps refer to the literature on global nonlinear analysis (e.g. Nitecki [96] and references therein).

For functionals F defined on a Riemannian manifold a Taylor expansion was derived by Smith [118], Chap. 4, §2. The variation $V_\varphi \in T_\varphi\mathcal{M}$ of F in a neighborhood of $\varphi \in \mathcal{M}$ is given by

$$F(\exp_\varphi(V_\varphi)) = F(\varphi) + \nabla F(V)(\varphi) + \frac{1}{2}\nabla^2 F(V, V)(\varphi) + O(\|V_\varphi\|^3)$$

where $V_\varphi \in T_\varphi\mathcal{M}$ is a tangent vector at $\varphi \in \mathcal{M}$ and ∇ is the covariant derivative of F . The vector field V on \mathcal{M} is *adapted* to V_φ , i.e. for each point ψ in the neighborhood of φ we have $V_\psi = \tau(V_p)$, where τ denotes the parallel translation of V_p along the unique geodesic connection between p and q . For the definition of geodesics and parallel translation refer to Sec. A.

The standard approach to performing numerical computations for such functionals would be to derive the Euler-Lagrange equations, and express all quantities in local coordinates. This requires the computations of an atlas structure for the manifold \mathcal{M} , the specification of a Riemannian metric on \mathcal{M} as well as the computation of atlases for both surfaces M and \widetilde{M} . In the case where

a common global parameterization of M and \widetilde{M} can be computed, i.e. for simple topologies such as disc or sphere [48, 82], the variational problem can be written as

$$F(r) = \int_M \|g(x) - \widetilde{g}(r(x))\|^2 \sqrt{g(x)} dx$$

where $r : \Omega \rightarrow \Omega$ is a map on the common parameter domain $\Omega \subset \mathbb{R}^2$ of M and \widetilde{M} , and $M = x^{-1}(\Omega)$ and $\widetilde{M} = \widetilde{x}^{-1}(\Omega)$ are the local coordinates. However, in the case of arbitrary topologies, common local coordinates need to be computed repeatedly depending on the current value of φ . It is not clear whether such an approach is tractable and can lead to efficient numerical schemes. Yet, at least from a theoretical point of view this approach should be pursued further. Some progress was already reported in [13].

In this work, however, we will take a different approach for two reasons: (1) The surfaces we are dealing with are a result of physical measurements (e.g. laser scanners) and/or complex reconstruction processes (e.g. manual segmentation and marching-cubes surface reconstruction from medical image data) and are therefore given as discrete data, more specifically as triangle meshes. There exists no correlation between the sampling of both M and \widetilde{M} . (2) M and \widetilde{M} can be of arbitrary topology. The explicit computation of atlases currently seems prohibitive for the reason mentioned above.

How can we represent the correspondence map φ on polygonal meshes? How can we measure or estimate stretching, i.e. the metric g and \widetilde{g} and their pullbacks (see equation 2.4), reliably from discrete data?

For triangle meshes there exists no canonical smooth approximation scheme. In fact, two different schemes may lead to completely different values of the metric, surface area, curvature, etc. On the other hand, if a given sequence of polygonal meshes converges to a smooth surface M in Hausdorff distance and the normal field converges, then (discrete) quantities like intrinsic distance, area, mean curvature, geodesics converge too [58, 137]. Hence, the discrete setting has a *smooth limit*.

Therefore, our goal is to develop a discrete representation of φ in conjunction with a method that is capable of reliably estimating geometric quantities on triangle meshes of arbitrary topology. Such an approach is presented in the next sections.

3.4 Discrete Stretching Functional

We propose to discretize φ by decomposing the shapes into cells c_i and approximating the integrand of F , i.e. the *stretching density*

$$f(\varphi) := \|g_p - (\varphi^* \widetilde{g})_p\|^2 \quad (3.7)$$

3. VARIATIONAL SURFACE MATCHING

by a constant f_i in each cell c_i , thus turning the infinite dimensional optimization problem into a finite one:

$$F(\varphi) = \int_M f(\varphi) dA = \sum_i \int_{c_i} f(\varphi) dA \stackrel{\varphi \approx \varphi_h}{\approx} \sum_i f_i \text{area}(c_i)$$

The discrete map φ_h is represented via the following construction. Let $K = (V, E, C)$ be an abstract 2-dimensional cell complex with

$$\begin{aligned} V &= \{1, \dots, n\} \quad \text{the set of vertices,} \\ E &= \{(i, j) \mid i, j \in V\} \quad \text{the set of edges and} \\ C &= \{(i_1, \dots, i_k) \mid i_1, \dots, i_k \in V, k > 2\} \quad \text{the set of 2-dimensional cells.} \end{aligned}$$

The surfaces M and \widetilde{M} are decomposed by embedding the cell complex K into M and \widetilde{M} , i.e. by specifying *valid* (definition see below) geometric realizations ψ and $\widetilde{\psi}$ of K in M and \widetilde{M} :

$$\begin{array}{ccc} M & \xrightarrow{\varphi_h} & \widetilde{M} \\ & \swarrow \psi & \nearrow \widetilde{\psi} \\ & K & \end{array} \quad (3.8)$$

The geometric realizations ψ and $\widetilde{\psi}$ are defined as follows: Let the image of the vertices V on M and \widetilde{M} be denoted by the sets of points

$$\begin{aligned} P &= \{p_1, \dots, p_n\} \quad \text{with } p_i \in M \\ \widetilde{P} &= \{\widetilde{p}_1, \dots, \widetilde{p}_n\} \quad \text{with } \widetilde{p}_i \in \widetilde{M} \end{aligned}$$

The geometric realizations of an edge $e = (i, j) \in E$ is defined as the shortest curve γ_{ij} connecting the points p_i and p_j . Its length is denoted $L_e = d(p_i, p_j)$. The geometric realization of each cell $c \in C$ is thus called a *geodesic cell* with an area denoted by $\text{area}(c) = A_c$. A cell $c = (i_1, \dots, i_k)$ is bounded by the set of curves $\gamma_{i_1 i_2}, \gamma_{i_2 i_3}, \dots, \gamma_{i_{k-1} i_k}$. Fig. 3.2 illustrates the embedding process. The resulting embeddings constitute a consistent surface decomposition of M and \widetilde{M} (cf. Def. 2.15).

In this setting, the discretization parameter h is given as

$$h = \max \left\{ \max_{e \in E} L_e, \max_{e \in E} \widetilde{L}_e \right\}. \quad (3.9)$$

Definition 3.1 (Valid Geometric Realization). *A geometric realization ψ of an abstract cell complex K on a surface M is called valid, if the image of the geometric realization $\psi(K)$ forms a geometric cell complex K_ψ which is topologically equivalent to K , i.e. consists of the same sets of vertices, edges and cells.*




Figure 3.2 Embedding a common cell complex K into two surface M and \tilde{M} : the vertices i, j, k are mapped to p_i, p_j, p_k and $\tilde{p}_i, \tilde{p}_j, \tilde{p}_k$ respectively. E.g. vertex i and k are connected by an edge (i, k) , whose geometric realization are the geodesics γ_{ik} and $\tilde{\gamma}_{ik}$. The cell $c = (i, j, k)$ is bounded by its three adjacent geodesics.

For the purpose of this work we will assume that valid geometric realizations exist for both M and \tilde{M} . As M and \tilde{M} are homeomorphic this is certainly the case for $h \rightarrow 0$. In order to complete the construction of the correspondence map $\varphi_h : M \rightarrow \tilde{M}$ we need to specify how to map the interior of geodesics and geodesic cells from M and \tilde{M} (interpolation scheme). In fact, we will not specify this interpolation scheme, as the functional we propose will be evaluated only on the 1-skeleton of the map φ_h given by the vertices and edges of the embedded cell complex. In order to extend the map from the 1-skeleton to the whole surfaces, continuity must be ensured across the embedded edges and inside cells. Otherwise, we leave the interpolation undefined. The method of consistent surface decomposition and parameterization described in Sec. 2.6.2 can be directly applied to obtain the full correspondence map φ , as it approximately minimizes stretching inside the cells.

In the case of two planar surfaces and a complex K , which is simplicial, the construction of a suitable distance measure based on the 1-skeleton of φ_h reduces to comparing two flat triangles meshes with Euclidean straight edges (Fig. 3.3). In this case the stretching density (3.7) is indeed constant on each cell, because the map from cell c to \tilde{c} is linear. It can be measured in different ways, e.g. by comparing edge lengths L_e with \tilde{L}_e or areas A_c with \tilde{A}_c (or eigenvalues of the linear map, see [112]).

We will use an analogy to the planar case to setup our discrete matching functional. The basic idea is to approximate stretching density by a constant value per cell by measuring only length deviations of the cell boundaries and area deviations between the interior of the cells c and \tilde{c} embedded on M and

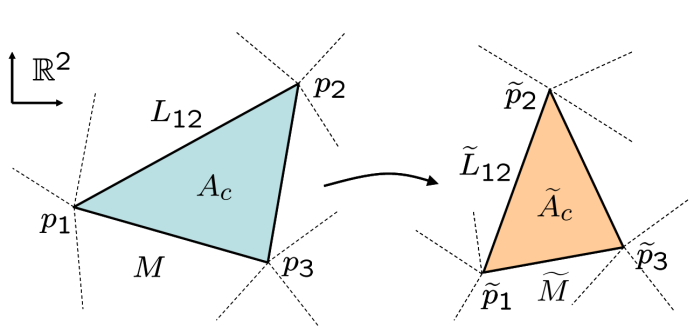


Figure 3.3 Stretching between two planar triangle meshes (simplicial cell complex embedded into the plane) is constant per cell, because the the map from cell c to \tilde{c} is linear. It can be measured by comparing edge lengths L or areas A .

\tilde{M} , thereby neglecting further geometric details on the interior of the cell embeddings. We propose:

$$F_h(P, \tilde{P}) = \lambda_L \cdot \sum_{e \in E} k \left(\frac{L_e}{L}, \frac{\tilde{L}_e}{\tilde{L}} \right) + \lambda_A \cdot \sum_{c \in C} k \left(\frac{A_c}{A}, \frac{\tilde{A}_c}{\tilde{A}} \right) \quad (3.10)$$

with the *discrete stretching density*

$$k(x, y) := \left(1 - \frac{x}{y} \right)^2 + \left(1 - \frac{y}{x} \right)^2 \quad (3.11)$$

λ_L and λ_A are parameters to weight edge length against area deviation, while $L = \sum_e L_e$ and $A = \sum_c A_c$ denote total edge length and total surface area, respectively.

Functional (3.10) measures length deviations of corresponding geodesics and area deviations of corresponding geodesic cells on M and \tilde{M} . Functional (3.10) has the following desired properties (see Sec. 2.2):

- (i) F_h is **invariant under Euclidean transformations**, because it depends on differential geometric - even intrinsic - properties only.
- (ii) F_h is **invariant under rescaling** of both M and \tilde{M} , since the arguments of k are dimension-less.
- (iii) F_h is **symmetric** with respect to exchanging M and \tilde{M} , i.e. $F_h(P, \tilde{P}) = F_h(\tilde{P}, P)$, since $k(x, y) = k(y, x)$.

3.5 First Variation of the Discrete Stretching Functional

We have transformed the shape matching problem into the finite dimensional optimization problem

$$F_h : \mathcal{M} \rightarrow \mathbb{R}$$

defined on the $2n$ -dimensional product manifold

$$\mathcal{M} := \left(\prod_{i=1}^n M \right) \times \left(\prod_{i=1}^n \widetilde{M} \right)$$

As described before in Sec. 3.3 variations are defined in the tangent space of \mathcal{M} , which is given as a Cartesian product [97, p. 24]:

$$T\mathcal{M} = \left(\prod_{i=1}^n T_{p_i} M \right) \times \left(\prod_{i=1}^n T_{\tilde{p}_i} \widetilde{M} \right)$$

In order to minimize F_h we need to compute its variation ∇F_h with respect to some vector field $V \in T\mathcal{M}$. Due to the Cartesian product structure of \mathcal{M} and isomorphism, the variation can be split into independent variations for some $v \in T_{p_j} M$ (and $\tilde{v} \in T_{\tilde{p}_j} \widetilde{M}$, respectively). This essentially requires computing the variation of the length of a geodesic $\nabla_v L_e$ and of the area of a geodesic cell $\nabla_v A_c$ with respect to some vector $v \in T_{p_j} M$.

We fix some notations needed to produce the desired results (see Fig. 3.4): let $w_{ij} \in T_{p_j} M$ denote the unit vector tangentially to the geodesic between

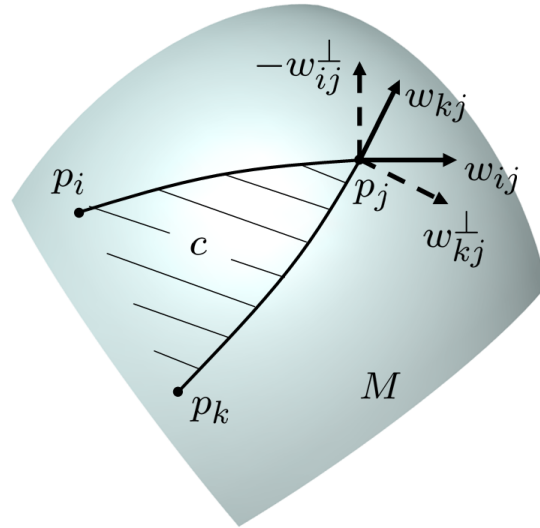


Figure 3.4 Gradient computation of F_h at point p . Explanation see text.

(and oriented from) p_i and p_j . Let $w_{ij}^\perp \in T_{p_j}M$ be the vector w_{ij} rotated by 90° clockwise (in the orientation of the surface M). For a cell c containing adjacent vertices i, j, k the gradients in $T_{p_j}M$ can be written as

$$\nabla L_e(p_j) = w_{ij} \tag{3.12}$$

$$\nabla A_c(p_j) = \sigma_{kj} w_{kj}^\perp - \sigma_{ij} w_{ij}^\perp \tag{3.13}$$

The gradient of the length L_e of a geodesic from p_i to p_j points tangentially along the geodesic and has unit length, denoted w_{ij} . The gradient of the area A_c at p_j has two terms, for both geodesics bounding cell c ending in p_j (see Fig. 3.4). Each term is the gradient of the area swept by varying the geodesic by displacing p_j tangentially while fixing p_i . Since the area swept by a geodesic $\gamma_{i,j}$ is zero when p_j is displaced tangentially to the geodesic, i.e. along w_{ij} , its gradient must be normal to the geodesic at p_j , i.e. denoted w_{ij}^\perp . Its magnitude σ_{ij} is given by the derivative of the integrated normal components $u : [0, L] \rightarrow \mathbb{R}$ of the variation vector field

$$\sigma = \frac{d}{d\epsilon} \left(\int_0^L u(s) ds \right) \tag{3.14}$$

constrained by

$$u(0) = 0 \quad \text{and} \quad u(L) = \epsilon. \tag{3.15}$$

The variation vector field for a family of geodesics is called *Jacobi field*. Its normal component u is characterized by the ordinary differential equation

$$u''(s) = -\kappa(s)u(s) \tag{3.16}$$

where κ is the Gauss curvature along the curve. Thus, the solution of the boundary value problem (3.16) with (3.15) substituted into (3.14) yields the magnitude of the area gradient.

As examples, we consider the case of $\kappa = 0$ (planar surfaces) and $\kappa = 1$ (sphere). For $\kappa = 0$ the boundary value problem is solved by $u(s) = \epsilon \frac{s}{L}$. Hence, we have $\sigma = L/2$, which is the well known Euclidean expression. For $\kappa = 1$ we have $u(s) = \epsilon \frac{\sin(s)}{\sin(L)}$ and $\sigma = \frac{1 - \cos(L)}{\sin(L)} = \frac{L}{2} + \frac{3}{24}L^3 + \dots$. For small values of L , σ exhibits the expected asymptotic behavior towards the Euclidean case.

For triangle meshes, the computation of the point-wise Gauss curvature κ is described in Appendix B. An algorithm for computing σ for arbitrary Gauss curvature κ is presented in Appendix C.

3.6 Minimizing the Discrete Stretching Functional

We are faced with the problem of minimizing a function $F_h : \mathcal{M} \rightarrow \mathbb{R}$ over a finite-dimensional manifold, that is not given analytically, i.e. no atlas is

given. As mentioned before it is desirable to avoid the explicit computation of such an atlas. Therefore, we propose a method where the minimization is directly performed on the manifold.

3.6.1 Optimization on a Manifold

So far, we have only the gradient of F_h and no higher derivatives at our disposal. Hence, we will adopt a gradient-based line-search method to minimize F_h . A nonlinear conjugate gradient descent for optimization on a manifold was described by Smith [118]:

Algorithm 3.1. (Nonlinear Conjugate Gradient Descent on a Manifold)

1. Let $x_0 = (p_1, \dots, p_n, \tilde{p}_1, \dots, \tilde{p}_n)_0 \in \mathcal{M}$.
Compute $d_0 = -\nabla F_h(x_0) \in T_{x_0}\mathcal{M}$ and set $k = 0$.
2. Compute t_k so that $F_h(\exp_{x_k}(t_k d_k)) \leq F_h(\exp_{x_k}(t d_k))$ for all $t \geq 0$.
3. Set $x_{k+1} = \exp_{x_k}(t_k d_k)$.
4. Set $d_{k+1} = -\nabla F_h(x_{k+1}) + \beta_k \tau(d_k)$, where τ is the *parallel transport* from x_k to x_{k+1} , and β_k may be given by the Polak-Ribière formula

$$\beta_k = \frac{\langle \nabla F_h(x_{k+1}), \nabla F_h(x_{k+1}) - \tau(\nabla F_h(x_k)) \rangle}{\langle \nabla F_h(x_k), \nabla F_h(x_k) \rangle}$$

5. Stop if convergence is reached or increment k and go back to step 2.

It was shown by Smith [118] that Alg. 3.1 has the same rate of super-linear convergence as its Euclidean analogue. Convergence in the case of functions defined on Riemannian manifolds is defined as follows: let $\{x_k\}$ be a Cauchy sequence in \mathcal{M} that converges to x . The sequence is said to be convergent of the order n if there exist an integer m and a constant $\delta \in [0, 1)$ such that $d(x_{k+1}, x) \leq \delta d(x_k, x)^n$ for all $k \geq m$, where d is the distance on \mathcal{M} . Likewise, convergence could be shown for other optimization schemes such as Newton's method [44, 118, 129] or trust region methods [2].

Several authors [85, 4, 59, 89] showed that the convergence properties are preserved when the exponential update is relaxed to the general notion of retraction: Let $R_x : T_x\mathcal{M} \rightarrow \mathcal{M}$ be a first-order approximation to the exponential map, i.e. for $v \in T_x\mathcal{M}$ let $\exp_x(v) = R_x(v) + O(\|v\|^2)$. This means that geodesics can be replaced by any curve tangential to the search direction in the update step (step 2 and 3 in Alg. 3.1). Similarly, parallel transport τ (step 4 in Alg. 3.1) might be replaced by the more general notion of vector transport, but this is ongoing research [1, 3].




Figure 3.5 A hierarchy of cell complexes (grids): K^1 covers the whole shape with one cell, it is refined by adding one vertex in the interior of the shape and connecting it to all boundary vertices (K^2), resulting in a simplicial complex. Two additional refinements (K^3 and K^4) are performed by 1:4 splitting of each triangular cell.

3.6.2 Multi-level Optimization

We adopt a multi-level approach with Gauss-Seidel smoothing. We start with a coarse cell complex K^1 (few vertices and cells) on level 1 and minimize F_h . The coarse level solution is then prolonged to the next finer level K^2 and so on, resulting in a hierarchy of cell complexes:

$$K^1 < K^2 < \dots < K^n$$

The prolongation operator performs 1:4 refinement of triangular shaped cells by connecting the midpoints of edges with a shortest path. For cells with more than 3 vertices a new vertex is added on the interior of the cell and connected to all cell vertices, resulting in a triangular shaped cells on the finer levels (Fig. 3.5).

As a multi-grid smoother, we apply a Gauss-Seidel relaxation scheme. This means, that F_h is minimized *successively* with respect to a single variable p_i or \tilde{p}_i , while fixing all other variables, using Alg. 3.1. Constraints can be incorporated by fixing certain vertices p_i or \tilde{p}_i in the current cell complex K . Furthermore edges which lie on the boundaries or on non-manifold regions on M must be mapped to boundaries or non-manifold regions on \tilde{M} respectively. The bijectivity of φ_h is guaranteed throughout the algorithm by using the following line-search strategy in step 2 of Alg. 3.1:

The exponential map in a point $p \in M$ is computed (see Sec. 3.7.1) along the negative gradient d_k until it touches any other embedded geodesic curve γ in the star of the vertex $i \in V$ (see Fig. 3.6). The functional is evaluated at a discrete number n_{LS} of points evenly spaced along the resulting geodesic. Around the minimum the process is repeated for the interval between neighboring samples recursively until a given accuracy h_{\min} is reached. The value t_k is given by the optimal position on the geodesic. This way bijectivity is preserved.




Figure 3.6 Line minimization at p along the conjugate gradient direction d_k in iteration k : The functional F_h is (recursively) evaluated at a discrete number of points along the geodesics γ emanating in p (see text) to find the minimum and a better value for p .

3.7 Geodesics, Exponential Map and Surface Area for Meshes

In order to implement Alg. 3.1, we need to specify the computation of geodesics on M and \widetilde{M} . It is only at this point that we really need to fix the representation of the surfaces M and \widetilde{M} in order to present concrete algorithms. As mentioned before, the surfaces are given as triangle meshes.

On smooth surfaces, geodesics are straightest and locally shortest curves. They generalize the concept of Euclidean straight lines. On smooth surfaces they satisfy two properties: (1) Geodesics solve the initial value problem which states that from any point on a surface there starts a unique geodesic in any direction. This is required for evaluating the exponential map in the algorithm of Sec. 3.6. (2) The length minimizing property provides a solution of the boundary value problem of connecting two given points on a surface with a locally shortest curve. This is needed to evaluate the discrete stretching functional (3.10) and its variations (3.12) and (3.13).

3.7.1 Exponential Map

The algorithm presented in this section is based upon the work of Polthier and Schmieß [102] and their definition of straightest geodesics on meshes:

Definition 3.2 (Straightest Geodesics). *A curve γ on a triangle mesh M is called a straightest geodesic if for each point $p \in \gamma$ the left and right curve angles θ_l and θ_r at p are equal.*

This definition uniquely solves [102] the initial value problem for triangle




Figure 3.7 Computation of the straightest geodesic on a triangle mesh based on the definition of equal left θ_l and right θ_r angles (see Def. 3.2 and Alg. 3.2). Three cases have to be distinguished: point inside a triangle (a), on an edge (b) and on a vertex (c).

meshes, and thus allows to evaluate the exponential map $\exp_p(v)$ at any given point $p \in M$ in tangential direction $v \in T_pM$. $\exp_p(v)$ is simply the endpoint of the straightest curve γ emanating in p in direction v with length $\|v\|$. The algorithm for computing γ is as follows:

Algorithm 3.2. (Straightest Geodesics)

1. Let $q_0 = p$ and set $k = 0$.
2. Unfold the local neighborhood $U \subset M$ of the endpoint q_k of the curve γ into tangent space $T_{q_k}M$ as follows: If q_k is located inside a triangle the local neighborhood is the triangle itself and unfolding is trivial. If q_k lies on an edge the neighboring triangles are mapped to the tangent plane isometrically. If q_k coincides with a mesh vertex all adjacent triangles are flattened to the tangent plane by preserving angles relative to the total angle at that vertex (see Fig. 3.7).
3. If $k = 0$ compute q_1 as the intersection of the ray emanating from p in direction v with the boundary of the unfolded neighborhood U^* . If $k > 0$ compute q_{k+1} as the intersection of the ray passing from q_{k-1} through q_k mapped to U^* with the boundary of U^* . With this construction subsequent curve endpoints q_{k-1} and q_k touch a common mesh triangle and therefore q_{k-1} is always contained in the local neighborhood of q_k .
4. The current straightest path γ is the piecewise linear curve connecting the points q_0, \dots, q_{k+1} . Stop if γ has length $\|v\|$ or increment k and go back to step 2.

In the unfolded local neighborhood U^* this algorithm constructs equal left and right angles of 180° at each point q_k . Angles are preserved exactly when reversing the unfolding process (step 2) in case of triangles and edges. At

vertices angles are scaled uniformly and hence Def. 3.2 is adhered to by the algorithm, i.e. θ_l and θ_r are equal everywhere along the curve.

3.7.2 Length Minimizing Geodesics

Geodesics between two given points on triangle meshes are often computed only on the edge-graph of the triangle mesh using Dijkstra’s algorithm [32]. This algorithm can be implemented with a complexity of $O(n \log n)$ using a Fibonacci heap, when n is the number of vertices in the mesh. One basic operation is essentially equivalent to one addition. Backtracing is trivial: simply indexing the heap. This efficiency and simplicity for implementation has made this algorithm so popular. However, the approximation based on such geodesics may be very bad depending on the triangulation of the surface. Hence, two variants have been proposed in the literature for the computation of geodesics that cut across triangles in the mesh hence producing much better approximations of geodesics. Kimmel and Sethian [68] employ a variant of the fast-marching method to compute approximate geodesics on meshes. Surazhsky et al. [122] present a practical implementation of the algorithm of Mitchell et al. [90]. All these algorithms possess the same asymptotic run-time as Dijkstra’s algorithm. Yet, the basic operations in setting up the distance heap and backtracing are more difficult to implement and costly. They involve computations of propagating intersecting lines and intervals, which rely on multi-precision floating point computations. Alternatively, Martinez et al. [86] propose an iterative local minimization scheme in combination with an initialization based on the fast marching method.

Conceptually, we propose a similar approach to the work of Martinez et al. [86]. The algorithm we propose was conceived by Polthier et al. [103]. It provides a highly efficient way to compute locally shortest paths on meshes that cut across mesh triangles. The algorithm iteratively minimizes the length of an initially given path connecting two points on the mesh:

Algorithm 3.3. (Shortest Geodesics)

1. Let γ be an arbitrary initial curve on M connecting the point p with q . Set $k = 0$.
2. Compute the triangle strip S of M that contains γ . Unfold that triangle strip isometrically into the plane. Note that the unfolded strip S^* may overlap.
3. Compute the shortest path from p to q within the unfolded triangle strip S^* . This can be done efficiently in $O(m)$ -time using the algorithm of Lee and Preparata [81], where m is the number of vertices in the strip. Replace γ with the projection of the shortest path in S^* back to M .

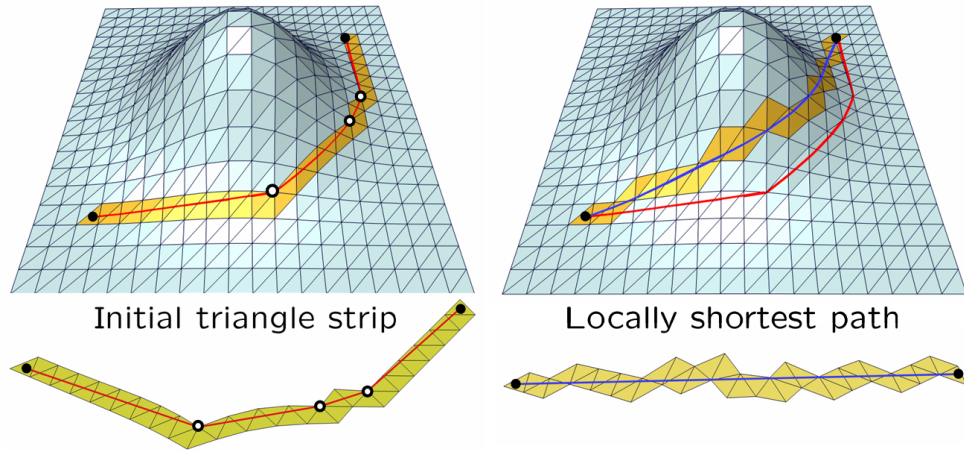


Figure 3.8 Iterative computation of the shortest geodesic on a triangle mesh. Starting from some initial triangle strip, the shortest path within this strip is computed. The triangle strip is updated, when the path touches its boundary. See Alg. 3.3 for more details.

4. Let Θ denote the total angle at a vertex of the mesh M . The current path γ may touch the boundary of the triangle strip S only at mesh vertices. If any of these vertices is (a) spherical ($\Theta < 2\pi$), or (b) hyperbolic ($\Theta > 2\pi$) and either the left or right curve angles θ_l, θ_r are not in the range $[\pi, \Theta - \pi]$ then the path will become shorter when the strip is redirected around that vertex star. Replace S with the triangle strip that is redirected at such a vertex, increment k and go back to step 3. If no such vertex exists stop.

See Fig. 3.8 for an illustration of the algorithm. The number of iterations required depend on the quality of the initial curve guess. We use Dijkstra's algorithm on the triangle graph (dual of the edge graph) of the mesh to generate the initial curve. Since in our algorithm one iteration can be performed in $O(m)$ -time, we have observed that Alg. 3.3 runs much faster than the method of Surazhsky et al. [122]. Furthermore, the algorithm is easy to implement. Note, that the solution depends on the initialization of the path: If, for instance, a shorter path around the other side of the peak in Fig. 3.8 exists, yet, the initial path lies on the opposite side, it cannot be detected with this method.

3.7.3 Computing Surface Area

After the geodesics have been updated, the surface area of the geodesic cells are computed. Generally geodesics run across faces of the triangle mesh. For faces, which are crossed by geodesics, the portion of the face area for each

adjacent cell is computed. For all other faces, their associated geodesic cell is determined and its area added to the area of total area of the cell. This requires virtually cutting the mesh along the geodesics and flood-filling the regions between the cuts to find all faces that belong to a given cell. In general, an arbitrary number of geodesics may run through a given face, which is cut by a geodesic. Hence, care must be taken that such faces are decomposed consistently. We decompose each face into a number of closed polygons, that are assigned uniquely to the geodesics cells.

3.8 Implementation and Numerical Results

We have implemented a prototype version of the proposed algorithm. Currently, some interaction is required and some parameters have to be adjusted manually. It is subject to future work to remedy such issues. The focus of this section is to demonstrate the feasibility of the proposed method.

3.8.1 Software Framework

The implementations have been done within the development framework of the Amira 3D visualization and modeling software [121]. Amira offers a triangle mesh data structure as well as a variety of read/write routines, visualization and computational modules for triangle meshes. For the task of surface matching the following extensions have been implemented:

- Consistent decomposition and parameterization of meshes (Sec. 2.6).
- Data structure for curves on meshes.
- Editor for interactive modification of curves on meshes.
- Shortest and straightest geodesics on meshes (Sec. 3.7)
- Computation of the area of a geodesics cell (Sec. 3.7.3).
- Gauss curvature on a mesh (Appendix B).
- Multi-level minimization of the discrete stretching functional (3.10) as described in Sec. 3.5 and Sec. 3.6.

3.8.2 Parameters and Interaction

Currently, the initial cell complex K^1 must be interactively embedded on both meshes M and \widetilde{M} . This can be done efficiently using the editor developed for defining and modifying curves on meshes. Practically, this is achieved by (a) specifying all embedded vertices, (b) connecting them with geodesic

or boundary/non-manifold curves according to the edge set E^1 in K^1 and (c) ensuring correspondence between geodesic cells. The parameters of the method are:

- **Functional** parameters allow to weight the influence of length λ_L versus area λ_A distortion.
- **Line-search** parameters adjust the accuracy of the one-dimensional minimum search. n_{LS} is the initial number of samples inspected along the gradient curve, while h_{\min} determines the desired accuracy.
- The number of levels in the **multi-level** hierarchy depends on the desired accuracy h_{\min} .
- The **convergence** criterion for stopping the conjugate gradient descent is given as a tolerance ϵ on the relative change in functional value.

At this stage, these parameters have to be fixed at the start of the algorithm. In future work, strategies for adaptively adjusting them to the geometric details of the shapes to be matched shall be investigated. Furthermore, the cell complex is currently refined uniformly in the multi-level hierarchy. Adaptive methods shall be devised in the future based on suitable error estimators.

3.8.3 Numerical Results

In this section we will illustrate different aspects and properties of the proposed method.

Multi-level Hierarchy

The multi-level scheme allows to adjust the level of geometric details that is to be resolved in the matching process. Fig. 3.9 illustrates this optimization on different levels for $\lambda_L = \lambda_A = 1$. Refining the grid K resolves more geometric details of the shapes, hence the optimal functional value will generally increase on a new level. Furthermore, the points along the boundaries are currently fixed on each level in our implementation. This introduces additional stretching. In the future this constraint shall be relaxed: generally points along the boundaries should have the freedom to move along the boundaries.

Robustness to Noise

The adjustment of the scale parameter h (or the coarseness of the cell complex) allows to control the robustness with respect to noise that is present in the data. If h is significantly larger than the noise present the matching can be computed robustly. Comparing Fig. 3.10 with Fig. 3.9 qualitatively shows

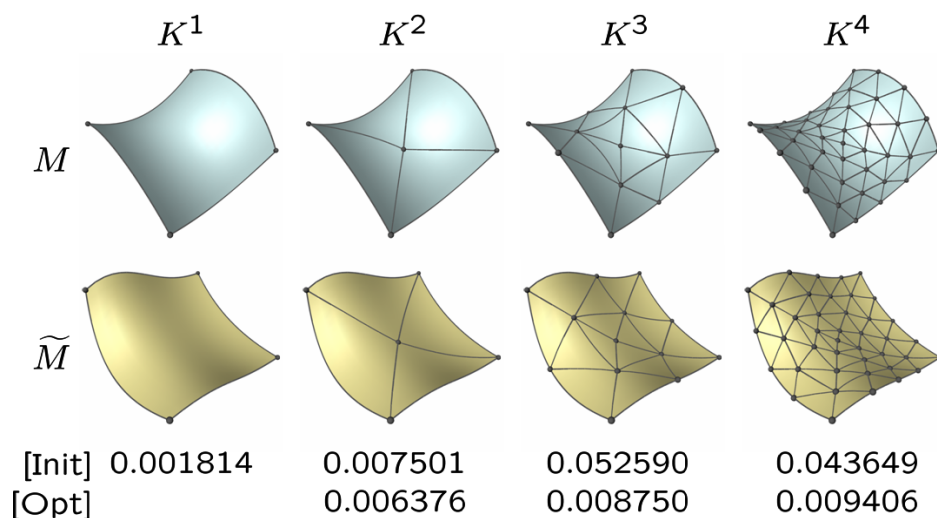


Figure 3.9 Multi-level optimization of F_h on four levels K^1, \dots, K^4 . The top row show the value of F_h after prolongation, the bottom row after optimization on that level. The images show the optimized result on each level.

that the matchings on level K^2 and even K^3 are not too different, starting from the same initial value.

Length versus Area Distortion

With $\lambda_A = 0$ the stretching functional represents a mass-spring system embedded in a curved domain with non-zero rest lengths. Fig. 3.11 illustrates the effect of neglecting the area term particularly for non-planar shapes.

Realistic Example

Fig. 3.12 shows the matching of two faces. Initially six points along the boundary are fixed (two points at each ear and two points at the mid-sagittal (or median) plane, and one at the tip of the nose). This initial embedding is performed manually, but could probably be automatized without too much difficulty. Since all cells are triangular they are refined by 1:4 splitting, while the nose tip is fixed on all levels. The optimization on level 2 (row 3) took about 10 minutes, while on level 3 (row 4) it took about one hour with the current implementation.

Qualitative Performance Analysis

Each evaluation of the stretching functional requires the computation of geodesics and cell areas for the edges and cells of the current cell complex on




Figure 3.10 Stretching optimization ($\lambda_L = \lambda_A = 1$) in the presence of noise: For large h (coarse grid) the discrete correspondence map can be computed quite robustly. The smaller h becomes the more sensitive to the noise in data the placement of the connecting geodesics gets (see close-ups).




Figure 3.11 Length versus area distortion: A deformed half-ellipsoid is matched with a plane. The interior point on the plane is fixed to illustrate the behavior of optimizing the corresponding vertex on the ellipsoid: With $\lambda_A = 0$ (a) the optimal position is such that the adjacent geodesics wind themselves around the ellipsoid. With $\lambda_A = 1$ the solution is much more intuitive.




Figure 3.12 Matching of two faces with $\lambda_L = \lambda_A = 1$: in each row the faces are shown from three different viewpoints (right/center/left). The second row shows the initially embedded grid K^1 . The vertex on the nose is fixed throughout the optimization. The third row shows the result after one refinement step K^2 , the third row after a second refinement K^3 .

two surfaces. Whenever a vertex is relocated on one shape the geodesics to its adjacent points as well as the surface area of its adjacent cells are recomputed (local update). For convex optimization the conjugate gradient descent requires two line-search step to reach the minimum. We observed, that locally only rarely more than two steps are necessary.

The computational effort for computing geodesics depends on the number of triangles N_M of the underlying mesh as well as on the value of the grid resolution h of the cell complex. Large h and high mesh resolution potentially require more iterations, as the triangle strip needs to be redirected more often, depending on the initial curve. The run time for the area computation depends on the number of triangles enclosed within a cell, and the number of triangles cut by the geodesics. This means, that for large h local functional updates are expensive, yet there are few (small number of vertices). For small h , local updates are less expensive (shorter geodesics, smaller cells), yet there are much more of them. In the future, the performance shall be analyzed in a quantitative fashion.

As usual in local relaxation schemes, the maximum change in vertex relocation is h . This means, that for a small values of h the convergence may be very slow, as the error must propagate across many layers of vertices. This motivates the combination of Gauss-Seidel relaxation with a multi-level hierarchy, where the optimization is started with large h .

On the one hand, the goal of future efforts will be directed on improving the numerical scheme with the aim to reducing the number of functional evaluations. On the other hand, there is much potential for improving the implementation.

3.9 Conclusions

We have developed a multi-level approach to shape matching which minimizes geometric stretching between two surfaces M and \tilde{M} . The continuous correspondence map φ is approximated by a discretization φ_h via embedding a common grid K into both surfaces. This de-couples the discretization of φ from the discretization of the surfaces. The matching functional is defined by the 1-skeleton of φ_h only, which allows to selectively adapt the method to the desired level of geometric content of the surfaces. This is useful particularly in the presence of noise in the data.

Our optimization-on-a-manifold approach of the discrete matching functional does not require the computation of surface parameterizations. This is an advantage when a global parameterization is not available. The proposed optimization strategy also assures that φ_h is always bijective.

These benefits are achieved at the expense of a rather costly functional evaluation: each update requires the computation of geodesics as well as of the

area of cells bounded by these geodesics on both shapes. We have presented algorithms that perform these operations with a reasonable computational effort. Future effort will be directed towards optimizing the efficiency of these methods.

From a practical point of view, there remain several open issues. Our approach requires an initial embedding of the grid K into the surfaces M and \widetilde{M} . How can such embeddings be achieved efficiently for arbitrarily given shapes? Another important aspect for making this approach applicable is an effective treatment of constraints. Currently, boundary points are fixed. This introduces unnatural stretching. Yet, it should be possible to prescribe corresponding input curves and allow grid vertices to move along such curves (e.g. boundary curves), or even other types of constraints. Furthermore, constraints should be definable by the user in an intuitive way. Another important aspect for future work is to improve the numerical approach by incorporating adaptivity, higher-order variations and improved line-search strategies. It shall be analyzed more thoroughly how the discrete relates to the continuous setting.

Chapter 4

Statistical Shape Modeling

4.1 Point Distribution Models

A major goal of statistical analysis is to transform measured data into more compact or visibly accessible representations. This often includes reducing the dimensionality of the original data and extracting only the essential degrees of freedom contained in the data.

Statistical analysis of shapes has attracted considerable attention in the fields of computer vision and image understanding through the work of Cootes et al. [28]. They have introduced the concept of *point-distribution models* and applied it for locating structures in medical images (active shape models).

The basic concept behind statistical shape models is to capture the most characteristic geometric variations contained in some given data base, often called *training set*. The main obstacle in constructing such models lies in the registration or matching of the training shapes, which is necessary in order to represent the shapes within a common vector space. This problem is referred to as the correspondence problem.

In the early works in the 1990s correspondence was often generated manually, and for closed two-dimensional contours only. Since then much work has been devoted to solving the correspondence problem both in 2D and in 3D, as well as on improving other aspects of statistical shape modeling [27, 26]. An overview of these developments is in preparation [53].

In principal, all non-linear matching or registration methods for three-dimensional shapes can be used to solve the correspondence problem, completely de-coupled from the application to statistical shape modeling. The methods we use are presented and discussed in Sec. 2.6 and Chap. 3. Another approach is based on the idea to compute the correspondences in such a way that certain properties of a statistical model are optimized [31]. Such an approach leads to a group-wise registration problem and it requires suitable regularization that takes into account the geometric nature of the shapes

[128]. Our variational method from Chap. 3 could be combined with such an approach.

Throughout this chapter we will assume that the training set consists of surfaces M_i with $i = 1, \dots, n$, and that continuous correspondence maps $\varphi_{ij} : M_i \rightarrow M_j$ have been computed between them.

Point distribution models (PDM) represent each shape M_i in terms of a description or shape vector

$$v_i \in \mathbb{R}^{3m}, \quad i = 1, \dots, n \quad (4.1)$$

which contains the coordinates of three-dimensional points

$$v_{ij} = (x_{ij}, y_{ij}, z_{ij})^T \in \mathbb{R}^3, \quad j = 1, \dots, m \quad (4.2)$$

sampling the shape M_i , so that

$$v_i = (v_{i1}^T, \dots, v_{im}^T)^T \quad (4.3)$$

Given a set of correspondence maps $\varphi_{ij} : M_i \rightarrow M_j$ the sampling could be fixed on any shape M_i to be the set of points

$$P_i = \{v_{ik} \in \mathbb{R}^3 : k = 1, \dots, m\} \quad (4.4)$$

Then the set of corresponding points on M_j is given by

$$P_j = \{v_{jk} := \varphi_{ij}(v_{ik}) \in \mathbb{R}^3 : k = 1, \dots, m\} \quad (4.5)$$

With this construction there is some bias in the choice of the points that represent the shapes in the PDM. The denser the shape is sampled, however, the better the PDM represents the real shape (Fig. 4.1), and the bias gradually becomes unimportant. Equation (4.5) clearly shows that the details of the sampling process - including the identification of corresponding points - influences the result of the statistical analysis. One can picture the points as vertices of different triangle meshes with the same mesh topology embedded on different surfaces.

Statistical analysis may be performed on alternative shape representations, cf. Sec. 2.4.1, which may be better suited to specific classes of problems. Examples are Fourier based shape descriptors [120], spherical harmonics [66] or angle based representation, which play an important role in molecular analysis [113]. In this thesis, we will work with a point-based representation, as this is the most commonly encountered representation in the field of geometry reconstruction.




Figure 4.1 Densely sampled point-based representation of two head models with corresponding points.

4.2 Shape Alignment

In order to perform meaningful statistical analysis, the shape vectors v_i must correspond in a reasonable (e.g. anatomical) way and must be aligned in a common coordinate system. In general, these two goals are accomplished independently of one another: methods that minimize spatial distance between shapes may simultaneously compute an alignment, while invariant methods (like ours) do not take the spatial location into account. In the latter case, the shape vectors must be aligned afterwards.

Then the task of alignment consists of finding a set of linear (affine, Euclidean or similarity) transformations T_i for each shape v_i such that some cost function D is minimized. A common procedure is generalized partial Procrustes analysis [47], where the distance is given by

$$D(\{T_i\}) = \sum_{i=1}^n |T_i(v_i) - \bar{v}|^2 \quad (4.6)$$

where $\bar{v} = \sum_{i=1}^n v_i/n$. $T(v_i)$ is an abbreviated notation for a linear transformation, that is applied to each point v_{ij} of v_i separately. Let

$$T(v_i) := (T(v_{i1}), \dots, T(v_{im})) \quad (4.7)$$

with

$$T(v_{ij}) := Av_{ij} + b$$

and $A \in \mathbb{R}^{3 \times 3}$, $b \in \mathbb{R}^3$. Furthermore the transformations T_i may be restricted to similarity (rotation, scale and translation) or Euclidean (rotation and translation) transformations. We will denote the space of three-dimensional linear transformations $\text{Lin}(3)$. The following iterative approach minimizes D :

Algorithm 4.1. (Iterative Alignment of Multiple Shapes)

1. Translate the center of gravity of each shape v_i to the origin.
2. Choose one shape v_j as an initial estimate for the mean shape \bar{v} .
3. Align all shapes v_i to the current mean shape \bar{v} by computing

$$T_i = \underset{T}{\operatorname{argmin}} |T(v_i) - \bar{v}|^2 \quad \text{for } i = 1, \dots, n \quad (4.8)$$

4. Update all $v_i \leftarrow T_i(v_i)$ and recompute the mean as $\bar{v} = \sum_i v_i/n$.
5. Repeat step (3) and (4) until convergence is achieved.

For affine transformations T , equation (4.8) amounts to solving a linear system. In the case of Euclidean transformations T , the problem is known as ordinary partial Procrustes analysis [124, 35] and is solved by singular value decomposition. Since D is decreased monotonically at each iteration, this iterative minimization scheme converges to a local minimum.

4.3 Statistical Shape Models

Once the shape vectors are aligned in a common coordinate system, statistical analysis can be performed. One of the most commonly used methods is based on the assumption of a multi-variate Gaussian distribution of the training shapes. It is known as principal component analysis (PCA):

Algorithm 4.2. (Principal Component Analysis)

1. Compute the mean shape (see Fig. 4.2)

$$\bar{v} = \frac{1}{n} \sum_{i=1}^n v_i \in \mathbb{R}^{3m}$$

2. Compute deviations from the mean shape

$$D = ((v_1 - \bar{v}), \dots, (v_d - \bar{v})) \in \mathbb{R}^{3m \times n}$$

3. Compute the covariance matrix

$$C = \frac{1}{n} DD^T \in \mathbb{R}^{3m \times 3m}$$

4. Compute eigenvalues and eigenvectors of C

$$Cp_k = \lambda_k p_k$$

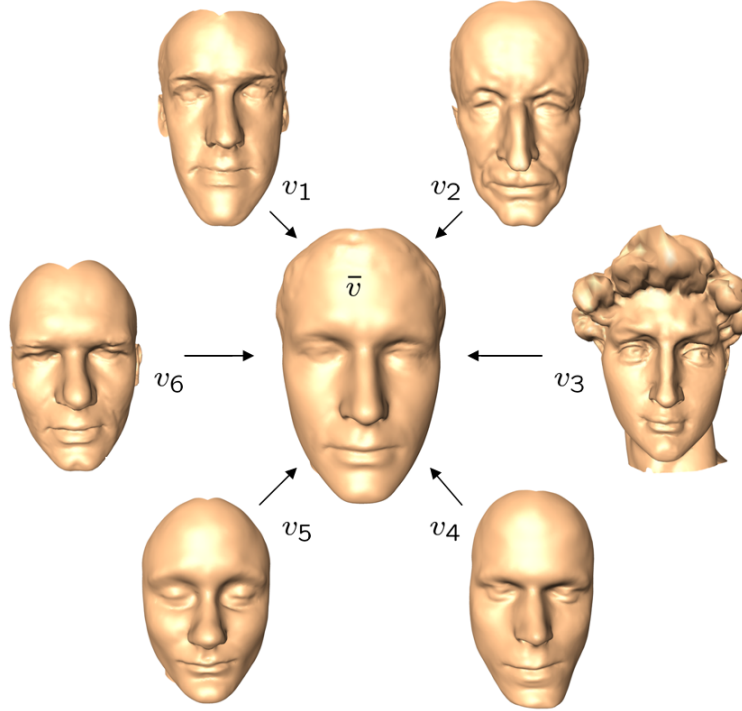


Figure 4.2 A training set $\{v_i\}$ of $n = 6$ faces and the corresponding average shape \bar{v} . Each model consists of $m \approx 16000$ points.

In physics, p_k are referred to as the principal axes of inertia, where each shape v_i is considered a point of mass 1. C is a positive semi-definite matrix, λ_k is the variance of the training data in the direction of p_k . When there are fewer samples in the training set than dimensions of the shape vectors, i.e. $3m > n$, then $3m - n$ eigenvalues will vanish. The remaining eigenvalues/-vectors of C can efficiently be computed as follows. Consider

$$\frac{1}{n} D^T D q_k = \mu_k q_k$$

and multiply both sides from the left by D to obtain

$$C(Dq_k) = \mu_k(Dq_k). \quad (4.9)$$

The computation of the eigenvectors q_k and -values μ_k of $\frac{1}{n} D^T D \in \mathbb{R}^{n \times n}$ is more efficient than that of $C \in \mathbb{R}^{3m \times 3m}$. From equation (4.9) we conclude that $p_k = Dq_k$ and $\lambda_k = \mu_k$.

PCA belongs to a family of methods known as *factor analysis* (FA). Such methods can be classified as linear or non-linear, reflecting whether the shapes

can be represented by a linear or non-linear function of the statistical components. Linear models, such as PCA, are often easier to interpret than others.

Another linear FA method is *independent component analysis* (ICA): its underlying assumption is that the data cannot be modeled by a Gaussian distribution. ICA looks for components that are statistically independent without assuming Gaussian distributions. There exist many different approaches for estimating ICA, e.g. nonlinear decorrelation by maximum likelihood methods or information theoretic approaches based on mutual information. A detailed discussion of these methods goes beyond the scope of this work. Please refer to Hyvärinen et al. [60] for more details. In contrast to PCA, ICA does not provide a compact representation of the variability contained in the training data, as there are as many independent components needed to reconstruct the training data as the number of data.

Yet another linear FA method worth noting here is called *principal factor analysis* (PFA) [9]. PFA models the covariance between the variables as opposed to the total variance in the data (PCA). While PCA determines the factors which account for the total (unique and common) variance in the set of variables, PFA determines the least number of factors which can account for the common variance (correlation).

Statistical shape models attempt to capture and model the variability contained in a given training set of shapes $V = \{v_i \in \mathbb{R}^{3m} : i = 1, \dots, n\}$. In this work we define a *statistical shape model* as the map

$$S : \mathbb{R}^d \times \text{Lin}(3) \rightarrow \mathbb{R}^{3m}$$

such that

$$(b, T) \mapsto S(b, T) = T \left(\bar{v} + \sum_{k=1}^d b_k p_k \right) = T(\bar{v} + Pb) \quad (4.10)$$

The parameters of the model are modes of variation $P = (p_1, \dots, p_d) \in \mathbb{R}^{3m \times d}$ with $p_k \in \mathbb{R}^{3m}$ derived from the statistical analysis on the training set and the mean shape $\bar{v} \in \mathbb{R}^{3m}$, as defined in Sec. 4.3. The eigenvectors are sorted in descending order of the magnitude of the corresponding eigenvalues

$$\lambda_k \geq \lambda_{k+1} \geq 0$$

The number of modes $d \leq n - 1$ included in the model may be less than all modes available from the statistical analysis. The variables of the models are the *shape weights* $b \in \mathbb{R}^d$ and a linear *transformation* $T \in \text{Lin}(3)$. This linear transformation may be a Euclidean or similarity transformation. If all eigenvectors are used, then the shape weights b for a training data set v_i can exactly be computed as

$$b = P^T(v_i - \bar{v}) \quad (4.11)$$




Figure 4.3 All $d = 5$ eigenvectors of the face model derived from the training set in Fig. 4.2. Column k shows the shape $S(b)$ associated to the minimal (top) and maximal (bottom) shape weight b_k present in the training data (computed from equation (4.11)), while $b_j = 0$ for $j \neq k$.

since $P^T P = \text{id}$. Fig. 4.3 shows the variability of a shape model of human faces.

4.3.1 Quality of Statistical Shape Models

The following properties of statistical shape models can be used to assess the quality of a model. They are defined by Davies [30]:

Generality

The generalization ability of a model measures its capability to represent unseen instances of the class of object modeled. This is a fundamental property as it allows a model to learn the characteristics of a class of objects from a limited training set. If a model is over-fitted to the training set, it will be unable to generalize to unseen examples. The generalization ability of each model is measured from the training set using *leave-one-out* reconstruction. A model is built using all but one member of the training set and then fitted to the excluded example. The accuracy to which the model can describe the unseen example is measured and the process is repeated excluding each example in turn. The approximation error is averaged over the complete set of trials.

Specificity

A specific model should only generate instances of the object class that are similar to those in the training set. It is useful to assess this qualitatively by generating a population of instances using the model and comparing them to the members of the training set.

Compactness

A compact model is one that has as little variance as possible and requires as few parameters as possible to define an instance. This suggests that the important information is captured in a plot of cumulative variance.

Statistical models can be used in a variety of applications. While the above criteria may yield good indications of the quality of a model, the decisive moment is the quantitative performance of the model in an application. Three selected applications are presented in the next chapter.

Part II

3D Geometry Reconstruction

Chapter 5

Segmentation of Medical Images

5.1 Introduction

In medicine, modern three dimensional imaging techniques such as computer tomography, magnetic resonance or even three dimensional ultrasound have made it possible to create detailed and complex geometric models of individual anatomical structures. Such models form the basis for computer-aided diagnosis and therapy planning. In biology, more and more three dimensional imaging methods such as confocal microscopy or electron tomography are established which allow the study of cellular or even sub-cellular structures. Therefore segmentation has become a major focus of many researchers worldwide who strive to develop efficient, robust and automatic methods, that meet the highly demanding requirements of the daily routine.

The creation of a geometric model of a certain structure from three dimensional image data requires its delineation from surrounding structures and the background. This process is referred to as image segmentation.

Due to the different characteristics of available imaging modalities and the variety of possible structures of interest a single method for automatic segmentation cannot be expected to succeed in all problems. Yet, even for one imaging modality automatic segmentation methods based on low-level vision methods do not produce satisfactory results, except in special cases. Therefore, interactive methods are routinely used in many cases. These are often time-consuming and their outcome is irreproducible.

The fundamental hypothesis is that automatic segmentation can be achieved by incorporating a-priori knowledge about geometric shape characteristics and image characteristics of the objects to be segmented.

5.2 Problems and Challenges

Image segmentation is the task of partitioning a given image $I : \mathbb{R}^3 \rightarrow \mathbb{R}$ into contiguous regions representing individual objects - in medical images e.g. anatomical entities. This step is usually followed by image classification which assigns these regions to classes, based on features that are extracted from these regions. The label associated to each region may or may not belong to some semantic space. Here we will consider the two steps merged into one: The segmentation process assigns a label to each pixel/voxel, thereby creating a labeled partition of the image domain, i.e. a map from the image domain to a set of labels. The semantic definition of the label is determined by the application. For instance in a medical application dealing with bones, it must be defined whether bone and bone marrow shall be assigned the same label or not.

Since distinct materials often cannot be discriminated based on intensity values only, much more complicated segmentation operators are needed, which depend not only on the intensity value of the voxel to classify, but also on its neighborhood, i.e. on textural and other contextual information. Standard image segmentation techniques base their decisions on homogeneity measures, like e.g. intensity variance and texture similarity (similarity of local statistical properties), and measures signaling the presence of edges, like e.g. intensity gradients.

However, in many practical cases detecting edges and clustering homogeneous regions is not sufficient for the following reasons:

- Images are noisy, and so are homogeneity measures and edge information. This means, there are regions where homogeneity and edge presence are either pretended or camouflaged.
- Due to object properties or imaging characteristics, individual objects may appear inhomogeneous and separated by artificial edges.
- Due to object properties or imaging characteristics, different neighbored objects may appear homogeneous and not separated by edges.

See Fig. 5.1 for an illustration of these problems. A consequence is a lack of robustness of automatic segmentation methods of the kind mentioned above.

Deformable models [87] of the object to be segmented have proven to be helpful in overcoming some of the problems. For instance they help to bridge regions where the signal is deficient. The basic principle behind all these approaches is some dynamic evolution of the deformable model governed by internal and external forces. Internal forces constrain the elasticity of the model (length, area, curvature) while external forces attract the model to features in the image. However, this approach allow shapes to evolve that definitely are




Figure 5.1 Left: CT of the liver/heart/lung region, Right: T1-weighted MR data from the abdomen.

not *legal* instances of the object considered (lack of specificity). Therefore, image segmentation procedures are sought that utilize *a-priori knowledge* about average shape and shape variability of the objects to be segmented. This may for instance be achieved by modeling the physical behavior of objects to be segmented based on a single example of a shape [95]. However, such properties are often not available, highly inaccurate or the methods are computationally prohibitive.

5.3 Segmentation with Statistical Shape Models

An alternative is the use of statistical shape models (SSM), see Chap. 4. Here the variability of an object is captured statistically from a suitable training set. The segmentation then consists of computing the set of variables of the shape model such that the model yields a true representation of the object in the image data I to be segmented. The segmentation strategy described here was essentially introduced by Cootes et al. [28].

Let $R \in \mathbb{R}^{3m}$ denote a shape vector of some surface, e.g. the vector of the coordinates of m vertices on some triangle mesh, to be segmented from the image data I . Segmentation using a statistical model $S(b, T)$ (see equation (4.10)) can be formulated as the optimization problem

$$(b^*, T^*) = \operatorname{argmin}_{b, T} |R - S(b, T)|^2 \quad (5.1)$$

The final segmentation is an approximation to R given by $R^* = S(b^*, T^*)$. However, the location and shape of R is only encoded implicitly in the image data I . Therefore the computation of R^* proceeds iteratively. Let $R^i = S(b^i, T^i)$ denote the segmentation in iteration i :




Figure 5.2 Left: A shape model $S(b, T)$ of the liver is initialized in the bounding box of a CT data set. Middle: The optimization of the transformation parameters T has converged. Right: The final segmentation after also optimizing the shape parameters b .

Algorithm 5.1. (Model-based Image Segmentation)

1. $R^0 := S(0, T^0)$.
2. Compute a *displacement* vector field $\Delta R \in \mathbb{R}^{3m}$ defined on the current segmentation R^i , i.e. a vector $\Delta r_k \in \mathbb{R}^3$ is assigned to each vertex $k \in \mathbb{N}$ of the surface R^i . It describes the desired deformation of the model towards the (unknown) surface R in the underlying image data I .

3. Project the displacements onto the SSM by solving the optimization problem

$$(b^{i+1}, T^{i+1}) = \underset{b, T}{\operatorname{argmin}} |(R^i + \Delta R) - S(b, T)|^2. \quad (5.2)$$

4. Update $i \leftarrow i + 1$ and return to step (2) if convergence has not been achieved, i.e. if $|S(b^i, T^i) - S(b^{i+1}, T^{i+1})| > 3m \cdot \epsilon$; else return $b^* = b^i$ and $T^* = T^i$.

The evolution of the algorithm is illustrated in Fig. 5.2.

The crucial ingredient in this algorithm is the computation of the displacement vector field $\Delta R \in \mathbb{R}^{3m}$ (*displacement strategy*). This computation is based on some *intensity or appearance model* of the underlying type of image data to be segmented. Such a model depends on the knowledge about imaging characteristics as well as image properties of the object to be segmented. For example, the segmentation of bone is performed differently for CT or MRT data. Apart from the shape knowledge incorporated in the statistical shape model, this is the second point where available a-priori knowledge from the application can enter and help to improve the segmentation.

A generic displacement strategy was proposed by Cootes et al. [28]. For each training data set, not only the shape but also the image data is sampled in a neighborhood of the current surface of the shape model. This may e.g. be done by extracting a pre-defined number of image data samples along a




Figure 5.3 One way to compute the displacement vector field ΔR is by analyzing image data profiles at each vertex of the shape model normal to its surface. Left: a model of a human mandible inside a CT image data with an exemplary profile. Middle: the CT sampled along the profile shown in the right image. Value 25 on the x-axis corresponds to the point on the surface, 50 is furthest outside. The peak indicates the bone boundary in the image data: a good candidate for a desired displacement. Right: The resulting displacement vector field ΔR , where the length of each displacement is color-coded on the surface (red=outwards, blue=inwards).

one-dimensional profile of a given length normal to the surface of the shape (see Fig. 5.3). At each point of the surface, a *statistical profile model* of the image data distribution is computed. In the segmentation process, several such profiles are extracted from the image data to be segmented at each point and their Mahalanobis distance to the statistical profile model is computed. The best match determines the normal displacement ΔR^i . However, this procedure is not always applicable. For instance if the shape model was not generated on the same data basis as the data to be segmented. Other displacement strategies, as for example heuristic strategies, have to be devised.

We add some comments about Alg. 5.1. The optimization (5.2) in step (3) is performed either with respect to

- (1) the position parameters T (position adjustment), or
- (2) the shape parameters b (shape adjustment)

The first case is equivalent to the shape alignment problem (4.8), hence amounts to solving a linear system of equations or a singular value decomposition, depending on the type of transformation allowed. In the second case, a linear system of equations of dimension d (= number of allowed shape modes in the model) has to be solved. The iterative segmentation process may generically be initialized by positioning the average shape ($b = 0$) in the center of




Figure 5.4 Three different slices from low-dose CBVT data.

the bounding box of the image data I . Better initialization strategies should be devised depending on the application.

The iterative segmentation process can be augmented with a hierarchical structure in order to improve the performance and increase the robustness of the method. The parameters that can be adapted are

- the sequence of position and shape adjustments,
- the number of modes d of the shape model S , and
- the resolution of the image data to be considered.

Transformation parameters should always be optimized, before shape parameters are optimized. It may be useful to start shape optimization with only few shape modes on a coarse version of the image data and successively increase d and refine the image data. This is a common approach in image registration in order to prevent getting stuck in local minima.

The following sections contain different examples of shape model based segmentation of different anatomical structures. All shape models were constructed using the method of consistent surface decomposition and parameterization (Sec. 2.6).

5.4 Mandible from Low-Dose CT data

Cone-beam volumetric tomography (CBVT) is a widely available technology and a suitable foundation for three-dimensional diagnoses and planning in cranio-maxillofacial surgery [115]. Such scanners can operate with a significantly reduced patient's exposure to radiation compared to conventional CT. As a side effect of the low dose, however, such images are often noisy, and metal artifacts are present (Fig. 5.4). One of the major applications for CBVT is dental imaging. The surgical procedure of placing dental implants requires careful




Figure 5.5 Training set of mandibles (without teeth region), lower right: average shape.

preoperative planning. The surgical plan is guided by prosthetic considerations and anatomical structures, which limit the volume into which implants can be inserted. The basis for computer aided planning is a segmentation of the mandibular bone in the data. Due to the problems mentioned, this task is difficult to automate with low-level techniques, such as thresholding, region growing or morphological filters.

The statistical shape model of the mandible used in this application consists of 13 individual mandible shapes (Fig. 5.5), reconstructed interactively from conventional CT data. Each surface is decomposed into 8 regions in a symmetric manner: lower corpus mandibulae, teeth region, ramus mandibulae and the caput mandibulae (Fig. 5.6). The teeth region (two patches) is not considered in the statistical analysis because, in general, teeth differ in their number and the topology of the reconstructed geometry (contact vs. non-contact) from patient to patient. Moreover, the displacement model for the segmentation is not designed to include this region.

The design of the displacement model for segmenting mandibular bone from low-dose CT data is based on the analysis of one dimensional profiles along the normal of the shape model's surface. Typically, such profiles show two distinct peaks that indicate the transition from the callus and the marrow of the bone (see Fig. 5.7). In the upper mandibular regions the bone is rather thin and often exhibits only a single strong peak. Thus, the deformation vector at a




Figure 5.6 Patch decomposition of mandibular bone for the construction of the statistical shape model.




Figure 5.7 Design of deformation model. Left: two profiles ($L = 4$ cm) normal to the surface. Middle: profile in upper mandibular region (single peak). Right: profile in lower mandibular region (twin-peak).

given vertex is determined as follows:

- sample a profile normal to the current surface of length (such that half of the profile is located on the inside and on the outside of the surface)
- apply smoothing on the profile, e.g. by Gauss or median filtering
- detect the two major peaks
- if the magnitude of the peaks differs significantly, discard smaller of the two
- move outwards from the right peak to the point of inflection

The teeth region is not considered in this analysis, because of metal artifacts inherent in this region.

For the evaluation of the segmentation process 15 data sets from a NewTom DVT scanner were available. These were segmented interactively by anatomical experts and serve as the gold standard for the evaluation of the proposed




Figure 5.8 Mean and maximum surface distances between gold standard and results of segmentation and surface reconstruction for each of the 15 NewTom data sets.

methods. For comparison, the statistical model is directly fitted to the surface reconstructed from the segmentation of each of 15 data sets by solving the minimization problem

$$(b^*, T^*) = \operatorname{argmin}_{b, T} d_2(R, S(b, T))$$

Here $d_2(S, S')$ denotes the squared surface distance between S and S' . For details on how to solve this optimization problem refer to Chap. 6.

Initially the shape model is placed in the center of the bounding box of the CT data. The number of sampled points along each profile is kept fixed at all times and the segmentation process is divided into two phases. In the first phase, only the lower corpus mandibulae is matched to the CT data, i.e. the deformation vectors of all other patches equaled 0, was set to 5 cm. This yields a good initialization for the second phase, where the ramus mandibulae is also segmented, and is reduced to 2 cm - resulting in a more accurate deformation of the model and at the same time preventing profiles in the upper mandibular region from (falsely) detecting the maxilla. The accuracy of the segmentation is measured by computing the surface distance between the automatic result with the gold standard. Results are shown in Fig. 5.8.

The results of the direct surface reconstruction indicate the best possible reconstruction to be achieved with the available statistical shape model. The deviations between the segmentation and the surface reconstruction are due to shortcomings in the segmentation process, i.e. inaccuracies in the deformation model and/or the segmentation strategy. Improvements are subject to future work.

However the results show that statistical models of 3D shapes offer a promising approach for automating the segmentation of volumetric data. Incorporat-

ing a-priori knowledge about the object and the data seems a feasible way to segment image data in the presence of high noise-to-signal ratios or artifacts - as is the case in low-dose CT data. The usage of a statistical 3D shape model of the mandible is a suitable approach due to the mandible's characteristic shape and well defined topology.

While the proposed method may require some manual interaction for post-processing, we expect that with increasing number of samples in the training set, interaction will be required only in a minority of cases. Since the model-based approach yields a good initialization for such interaction, this may even be avoided completely by means of locally elastic deformations, thus providing a fully automated segmentation method.

5.5 Liver Segmentation from CT data

5.5.1 Overview

The main issue in automatic liver segmentation from contrast-enhanced CT data is that the intensity values of the liver tissue are often similar to those of some surrounding anatomical structures like stomach, pancreas, kidney and muscles. Approaches which are only based on local intensity or intensity gradient features are usually not sufficient to differentiate between liver tissue and other anatomical structures in those regions.

In order to alleviate this problem a-priori knowledge about the typical shape of a liver is incorporated into the process to constrain the segmentation process where the image information is not reliable. The shape may be constrained by a single template [93], a statistical shape model (SSM) [28] or even more flexible deformable models. Combinations of these approaches have also been presented [138, 55].

Although the SSM based segmentation provides robustness, as was shown in the previous Sec. 5.4, it often lacks flexibility to accurately model the desired segmentation. To overcome this limitation we augment the statistical model based segmentation with a constrained free-form step with the aim of further reducing manual correction.

We present a simple heuristic model of the typical intensity distribution around the liver boundary based on a nonlinear isotropic diffusion filtering of the original CT data. The model also considers the potential presence of tumors inside the liver. An algorithm for computing displacement vectors is derived that drives the surface deformation. General intensity features [93, 138] or a statistical model of the intensity distribution [54] have previously been used in the literature.

In order to achieve fully automatic segmentation both the set of parameters required for the computation of the displacements as well as the initial posi-

tioning of the model in the image data will be estimated automatically from the image data.

5.5.2 Statistical Liver Model

The statistical liver model (see Fig. 5.9) was generated from 102 different data sets and has $m \approx 7000$ vertices. The patch decomposition of the liver surfaces is chosen to be along lines of high curvature (feature lines), since these lines represent anatomical features that can be found on each liver. We have divided the surfaces into four patches:

- the lower border of the left lobe (LL),
- the lower border of the right plus the caudate lobe (LR+CL),
- the lower border of the quadratic lobe (QL) and
- the whole upper part of the liver surface (U)

The feature lines are drawn on the surface by manually selecting points where two or more patch boundaries would meet (branch points). Some intermediate points along the feature lines are added manually. These points are then automatically connected by computing geodesic shortest paths between them. In almost all cases we use a metric that favors paths along lines of high curvature. Only in one region we use a pure distance measure. Following this procedure the user has to specify no more than 10 landmarks per surface, resulting in 4 patches, divided by 6 patch boundaries and 4 branch points. The first three modes of the model are visualized in Fig. 5.10.

5.5.3 Constrained Free-Form Segmentation

Although the SSM based segmentation provides robustness, it may lack flexibility to accurately model the desired segmentation R . In order to overcome this limitation we perform a so called free-form segmentation step, starting from the previous result R^* as computed by Alg. 5.1. However, we constrain the free-form segmentation to a narrow band around R^* to prevent too large deviations from the previously computed shape. The free-form segmentation is performed as follows:

Algorithm 5.2. (Free-Form Segmentation)

1. $R^0 := R^*$.
2. Compute two different vector fields on the current segmentation R^i : A *displacement* vector field ΔR and a *smoothing* vector field ΔS , pulling each vertex towards the barycenter of its 1-ring neighborhood.




Figure 5.9 Surface decomposition of the liver.




Figure 5.10 Variability of a statistical model of the liver shape made from 43 training data sets: in the left column the eigenmode with the largest variance λ_1 is varied between $\pm 3\sqrt{\lambda_1}$, in the second and third column the modes with the second and third largest variance are shown respectively.

3. Compute a *resulting* vector field ΔV . For an individual vertex of the surface it is given by $v = w_r \Delta r + w_s \Delta s$. The weights depend on the cosines $c_{rs} = \cos(\angle(\Delta r, \Delta s))$ as follows:

- (1) $|\Delta p| = 0 \Rightarrow w_s := 0.15 \cdot (1 - c_{pr}), w_r := 1 - w_s, w_p := 0$.
- (2) For $a, b \in \{s, r\}, a \neq b$: If $|\Delta a| = 0, c_{ab}$ is not defined $\Rightarrow w_a := 0$ and $c_{ab} := 0$. Set the remaining weight accordingly.

The resulting vector v is trimmed such that it does not leave the prescribed narrow band.

4. Set $R^{i+1} := R^i + \Delta V$. Perform this addition iteratively for all vertices, such that updated position $x^{i+1} := x^i + v$ of an individual vertex does not produce self-intersections in the surface. Then R^{i+1} also has no self-intersections.
5. Update $i \leftarrow i + 1$ and return to step (2) if stopping criterion has not been met (same ϵ as in Alg. 5.1). Otherwise, or if a maximum number of iterations has been reached, return.

5.5.4 Computation of the Displacement Field ΔR

The displacement vector field ΔR is computed on the basis of the image data I , after applying a nonlinear isotropic diffusion filter [139]: The filtered image is the solution of the partial differential equation

$$\partial_t u = \operatorname{div} (g(|\nabla u_\sigma|^2) \nabla u) \quad \text{with } u(x, 0) = I(x) \quad (5.3)$$

and diffusivity function

$$g(s) := \begin{cases} 1 & (s \leq 0) \\ 1 - \exp\left(\frac{-3.315}{(s/\lambda)^4}\right) & (s > 0) \end{cases} \quad (5.4)$$

A displacement vector Δr_k is computed for each vertex k of the surface by analyzing a 1D intensity profile at vertex position x along the (unit) surface normal u . This profile is sampled equidistantly over a length L at the set of points:

$$P = \{x_i := x + [(i - 1)/(N_p - 1) - 1/2] \cdot L \cdot u \text{ with } i = 1, \dots, N_p\}.$$

The result of the analysis will yield the displacement vector $\Delta r_k = w(x_n - x)$, with $x_n \in P, n \in \{1, \dots, N_p\}$ and a confidence weight w , at the vertex k of the surface. Fig. 5.11 shows some typical profiles in different anatomical regions.

The basic assumptions for the computation of the displacement are based on the following simple *model* for the intensity distribution around the liver boundary: The intensity inside the liver lies in the range $G_L = [g_L - t_L, g_L + t_L]$, where g_L denotes the average liver intensity and t_L the corresponding

tolerance. Analogously, the intensity of tumors (potentially) contained in the liver are in the range $G_T = [g_T - t_T, g_T + t_T]$. Furthermore, there exist bounds g_{\min} and $g_{\max} := g_L + 3t_L$ below resp. above which the intensity indicates definitely non-liver and non-tumor tissue. The liver boundary is characterized by a minimal absolute value of the intensity gradient $d_{\min} > 0$. A sample point along a profile is only considered inside the liver/tumor if at least $c_{\min} := N_p/5$ consecutive neighbors on the profile lie within G_L/G_T . Based on these assumptions we propose the following strategy to compute $x_n \in P$ from the intensity profile at a vertex:

Algorithm 5.3. (Displacement Computation)

1. Initialize $n := (N_p + 1)/2$ and $w := 1$.
2. If $I(x_n) \notin G_L \rightarrow$ Determine largest $i < n$ with $I(x_i) \in G_L$ and set $n := i$. If there is no such point i , set $w := 0$.
3. If $I(x_n) \in G_L \rightarrow$ Count number c of consecutive $i \leq n$ with $I(x_i) \in G_L$. If $c \geq c_{\min}$, define x_n to be inside liver.
4. If x_n outside liver but $I(x_n) \in G_L \rightarrow$ Count number c of consecutive $i \geq 1$ with $I(x_i) \in G_L$. If $c \geq c_{\min}$, set $n := c$ and define x_n to be inside liver.
5. If x_n outside liver \rightarrow Set $w := 0$.
6. If $I(x_n) \notin G_L \rightarrow$ If tumor present, search for an x_i that is inside tumor. Use the same strategy as for the liver, but with G_T instead of G_L . If such an x_i exists, set $n := i$ and $w := 0.75$.
7. If $I(x_n) \notin G_L \cup G_T \rightarrow$ Determine all consecutive $i \leq N_p$ with $I(x_i) < g_{\min}$. If smallest such $i < n$, set $n := i$, $w := 0.75$ and return.
8. If $w = 0 \rightarrow$ Return.
9. Now x_n is either inside liver or inside tumor. The remaining steps are equal in both cases. \rightarrow Find first point $i > n$ with either $|I'(x_i)| > d_{\min}$ or $I(x_i) \notin G_{L/T}$. If $I(x_i) \notin G_{L/T}$, set $n := i - 1$, else $n := i$.
10. If $|I'(x_n)| > d_{\min} \rightarrow$ Find point of inflection x_i with smallest $i > n$, then find largest $j < i$ such that $I(x_j) > g_{\min}$. Set $n := j$.
11. Find largest $i < N_p$ with $I(x_i) < g_{\max}$. If $i < n$, set $n := i$. This prevents from moving too far into regions such as kidney or heart.




Figure 5.11 Exemplary intensity profiles. (a) G_L , (b) g_{\min} , (c) g_{\max} , (d) current vertex position x , (e) suggested new position x_n , (f) profile plot.

5.5.5 Estimation of Intensity-based Parameters

The intensity parameters, which are required for computing the displacement field ΔR , are estimated based on an analysis of two different histograms H_1 and H_2 of the preprocessed image data. The basic idea is to discriminate the major liver peak(s) from minor yet distinct other peaks, which indicate the presence of tumor tissue. H_1 is the histogram on the volume inside some given liver surface, while H_2 considers a slightly *enlarged* liver volume by growing the current liver surface by 5 voxels. The histograms are evaluated only on voxels with an intensity in the range of $[0, 300]$ HU, which is assumed to cover both liver and tumor tissue. A weighted sum of 10 Gaussians P_i ($i = 1, \dots, 10$) is fitted to each of the two histograms using the Expectation Maximization (EM) algorithm on a Gaussian mixture model. As a result, we obtain weights w_i , means μ_i , standard deviations σ_i and the so-called *peak height* $h_i = w_i/\sigma_i$ for each Gaussian P_i . First, the tumor intensity range G_T is computed based on histogram H_1 :

Algorithm 5.4. (Estimation of Segmentation Parameters I)

1. Identify P_i in H_1 with the highest h_i . P_i is assumed to model the intensity distribution inside the liver (liver peak).
2. Identify $N := \{j \mid |\mu_i - \mu_j| < \max(15, \min(20, 3\sigma_i)), h_j > 0.02 \cdot h_i\}$ as *nearby liver peaks*. Set $l_1 := \min_{j \in N} \{\mu_j - \sigma_j \cdot \max(1, \min(3, 50 \cdot h_j^2/h_i^2))\}$ as the *lower liver boundary*. Set the *upper liver boundary* u_1 accordingly.
3. Compute the *approximate total peak height* as $r := \sum_{j \in N} h_j$. Note that r is not necessarily a peak height of our Gaussian mixture.
4. Now identify *potential tumor peaks* $T := \{j \mid \mu_j < l_1, h_j > 0.05 \cdot r\}$. If $|T| = 0$, assume that there is no big tumor. Otherwise identify $\sigma_{\min} := \min_{j \in T} \{\sigma_j\}$. Set $T^* := T \setminus \{j \in T \mid \sigma_j > 2\sigma_{\min}\}$ and $t := \arg \max_{j \in T^*} \{\mu_j\}$. P_t is assumed to model the intensity distribution inside tumors.
5. Fix the tumor intensity range G_T with $g_T := \mu_t$ and $t_T := \min(20, 3\sigma_t)$.

Since Histogram H_1 possibly excludes important regions of the liver, which are not covered by the current segmentation, we analyze the Gaussians fitted to H_2 in a second step to see whether the liver intensity range needs to be enlarged:

Algorithm 5.5. (Estimation of Segmentation Parameters II)

1. Identify liver intensity boundaries l_2 and u_2 in H_2 just like l_1 and u_1 in H_1 .
2. Identify a tumor Gaussian P_{t_2} in H_2 just as in H_1 . Accept P_{t_2} only if none was found in H_1 , and if $\mu_{t_2} > 0.5(l_2 + u_2) - 30$. Then set $g_T := \mu_t$ and $t_T := \min(20, 3\sigma_{t_2})$. This reduces the risk of including fat to the liver range.
3. Set $l := l_2$ and $u := \max(u_1, u_2)$. If $l_1 < l_2$ and not $g_T > l_1$, set $l := l_1$.
4. Fix $g_L := 0.5(l + u)$ and $t_L := 0.5(u - l)$, and $d_{\min} := 0.5t_L$.

5.5.6 Estimation of the Initial Position

It is assumed that the general position of the patient in the CT scanner is known and correct from the DICOM header (e.g. FFS = Feet First Supine). The idea of the position initialization is to robustly detect the lower rim of the right lobe of the lung and to position the liver model below it. First, all connected components with intensity values less than -600 HU are determined and the largest two components (left and right lobe), which are adjacent to the upper border of the image volume, are selected. Next, the lobe component on the right-hand side is projected in patient axis direction from feet to head and the center and orientation of the resulting lung area is determined. The liver model is now translated and oriented according to the back-projected center point and orientation of the lower rim of the right lobe.

5.5.7 Segmentation Algorithm

The segmentation algorithm consists of a series of steps combining the methods presented in the sections above, see Tab. 5.1. Each step is one of the following methods:

- position initialization (**Init**)
- estimation of parameters (**Estimate**), see Alg. 5.4 and Alg. 5.5
- optimization of the shape model (**Opt-SSM**), see Alg. 5.1
- constrained free-form deformation (**Opt-CFFD**), see Alg. 5.2
- or some other processing step explained in the table.

Step	Type	Details
1	Preproc.	two passes of nonlinear diffusion filtering of I with $t = 30$, $\Delta t = 5$, $\lambda = 10$. Pass 1 with $\sigma = 3$, and pass 2 with $\sigma = 1$.
2	Init	position T^0
3	Estimate	compute G_L, G_T, d_{\min}
4	Opt-SSM	only position (rigid + isotropic scaling $\in [0.5, 1.5]$), $L := 50$ mm, $N_p := 50$, $g_{\min} := 0$
5	Estimate	recompute G_L, G_T, d_{\min}
6	Opt-SSM	position and shape parameters (5 modes), $L := 60$ mm and $N_p := 60$ without tumor, $g_{\min} := 0$, $\epsilon := 0.1$ mm
7	Estimate	recompute G_L, d_{\min} only
8	Opt-SSM	position and shape parameters. The number of used modes m depends on the slice thickness z of the original ct data: Without tumor $m = \min(50, \max(20, -75z + 57.5))$, with tumor $m = \min(30, \max(20, -75z + 57.5))$. $g_{\min} := g_L - 3t_L$ without tumor and $g_T - t_T$ with tumor, $\epsilon := 0.05$ mm
9	Estimate	recompute G_L, d_{\min} only
10	Opt-SSM	same as step 7 but with $L := 40$ mm and $N_p := 40$ with tumor
11	Opt-SSM	Only performed if no tumor: $L := 20$ mm, $N_p := 40$.
12	Remeshing	The surface of the SSM is scan-converted, interior holes are filled, a new surface mesh is generated as input for the following steps.
13	Opt-CFFD	narrow band radius $r_B := 10$ mm, $L := 30$ mm, $N_p := 60$, $\epsilon := 0.04$ mm, maximum number of iterations $:= 30$.
14	Opt-CFFD	Only performed if no tumor was detected: $L := 10$ mm, $N_p := 50$, otherwise same as last OPT-CFFD step.
15	FillAllSlices	Fill 2D holes in xy-slices of scan converted result from step 14

Table 5.1 Segmentation algorithm. Parameters remain the same if not noted otherwise.

For each step the parameters for the specific method or other comments are listed in the third column labeled *details*.

5.5.8 Results and Conclusions

We presented an algorithm for automatic segmentation of the liver in CT data. Its main components are statistical shape and constrained free-form

5. SEGMENTATION OF MEDICAL IMAGES

Dataset	Overlap Error		Volume Diff.		Avg. Dist.		RMS Dist.		Max. Dist.		Total Score
	[%]	Score	[%]	Score	[mm]	Score	[mm]	Score	[mm]	Score	
1	6.3	76	-2.3	88	0.9	77	2.0	73	20.6	73	77
2	11.5	55	-10.8	43	1.6	60	2.4	66	17.0	78	60
3	4.7	82	-1.6	91	0.9	77	2.0	72	21.4	72	79
4	7.2	72	-0.5	98	1.2	70	2.4	67	18.4	76	76
5	8.8	66	-6.1	68	1.5	62	2.5	65	21.5	72	66
6	7.0	73	-5.6	70	1.3	68	3.3	54	36.5	52	63
7	6.3	76	-1.9	90	1.1	74	3.1	57	25.6	66	73
8	4.5	82	-0.2	99	0.7	82	1.5	79	13.5	82	85
9	4.1	84	-1.0	95	0.5	88	1.0	86	15.9	79	86
10	9.2	64	-5.7	70	1.3	67	2.3	68	19.1	75	69
Average	7.0	73	-3.6	81	1.1	72	2.3	69	20.9	72	73

Table 5.2 Results of the comparison metrics and scores for all ten test cases.




Figure 5.12 From left to right, a sagittal, coronal and transversal slice from a relatively easy case (1, top), an average case (4, middle), and a relatively difficult case (3, bottom). The outline of the reference standard segmentation is in red, the outline of the segmentation of the method described in this paper is in blue. Slices are displayed with a window of 400 and a level of 70.

deformations. The computation of the displacement vectors are based on a simple intensity model that may also apply to other segmentation problems.

Our method was evaluated within the MICCAI 2007 *Workshop on 3D Segmentation in the Clinic* [132], where it competed against nine other automatic segmentation algorithms from all over the world. The evaluation was performed on 10 CT data sets, where radiological experts manually outlined liver contours for all images in transversal slice-by-slice fashion. These segmentations were used as the ground truth. For quantitative comparison, a scoring system was developed that is based on five different discrepancy measures between automatically computed segmentations and the ground truth: (1) volumetric overlap error, (2) relative volume difference, (3) average symmetric surface distance, (4) root mean square symmetric surface distance, and (5) maximum symmetric surface distance. All measures were weighted equally to produce a score of 100 for a perfect match and 0 for complete failure. Additionally, the score was calibrated by a second rater, such that a score of 75 can be regarded as roughly equivalent to human performance.

Fig. 5.12 and Tab. 5.2 show qualitative and quantitative results of our method. The not yet optimized run-time of the algorithm was 15 minutes per liver on an Intel 3.2 GHz processor. The overall performance of our method has a score of 73. There has been no failure. Our method achieved the best overall score of all automatic methods, both in the on-site as well as in the pre-workshop competition. A comparison article is in preparation [56].

Segmentation errors sometimes occur in regions where anatomical structures with very similar intensity values are located close to the liver, in particular if parts of those structures can be captured within a typical liver shape: lower part of vena cava (Fig. 5.12 top, right), duodenum (Fig. 5.12 middle, right), heart, muscles, stomach and pancreas. The consideration of tumors inside the liver during the segmentation process is very important and works very well (Fig. 5.12 middle and bottom row). At high noise levels in the original image data the isotropic nonlinear diffusion process is stopped before reaching the liver boundary (e.g. case 2). Some deviations are caused by incorrect manual segmentations (e.g. Fig. 5.12 middle row, middle column, near the gallbladder).

The initiative of providing a common pool of test data as well as well-defined measures for evaluation is indispensable for further progress in the field of automatic image segmentation. Yet it remains difficult to establish the correct ground truth, especially for the liver, due to the difficulty to clearly define the exact boundary of the liver tissue (e.g. at the inner side, where the portal vein enters the liver or the vena cava). Although fully automatic liver segmentation algorithms are desirable in the clinical routine, methods that reduce and facilitate manual interaction will remain essential in case of failures.

Chapter 6

Reconstruction from Pathological Shapes

6.1 Introduction

In planning surgical interventions, a central issue is how to reconstruct pathological deformities or missing anatomical structures, when the original *normal* situation is unknown. In many cases, objective criteria to guide the re-modeling and re-shaping process are missing, and surgeons are forced to withdraw to subjective, e.g. aesthetic, assessment. This prevents impartial control of therapeutic success and aggravates guidance and instruction of the remodeling process for inexperienced surgeons.

Statistical models of healthy shapes offer the possibility for automated and reproducible reconstructions of pathological shapes R . The idea is to fit the statistical model $S(b, T)$ to match R in regions where there is no pathology and/or which will be unaffected by the surgical intervention. The resulting best fit R^* to R will provide a patient-specific yet objective proposal for the re-modeling process.

6.2 Model-based Surface Reconstruction

From a mathematical point of view, the difference to the segmentation process is mainly two-fold: (a) the shape to be reconstructed is explicitly known, yet (b) there exists no correspondence between the shape model $S(b, T)$ and R , i.e. we cannot easily compute $R - S$. Since R may be a subset or superset of the object represented by the shape model, the computation of correspondences is not easily accomplished.

Hence, we must minimize some surface distance between R and S :

$$(b^*, T^*) = \operatorname{argmin}_{b, T} d(S(b, T), R) \quad (6.1)$$

such that $R^* = S(b^*, T^*)$ is the best fit for R .




Figure 6.1 Point-to-surface distance is not symmetric.

Given two surfaces S and R the *point-to-surface* distance $d(p, R)$ between a point $p \in \mathbb{R}^3$ on surface S and the surface R is given by

$$d(p, R) = \min_{q \in R} |p - q|$$

Based on this definition surface distance measures can be constructed. In order not to lose important information, such measures should be symmetric, i.e. $d(S, R) = d(R, S)$ (see Fig. 6.1). We use the square symmetric surface distance defined as

$$d_2(S, R) = \frac{1}{|S|} \int_S d(p, R)^2 dp + \frac{1}{|R|} \int_R d(q, S)^2 dq \quad (6.2)$$

where $|S|$ denotes the surface area of S and dp its area element.

In the case of triangle meshes the surfaces are represented by a set of vertices $p_i \in \mathbb{R}^3$ with $i = 1, \dots, n_v$, and the integrals in equation (6.2) are replaced by sums over all vertices:

$$\frac{1}{|S|} \int_{p \in S} d(p, R)^2 dS \rightarrow \frac{1}{|S|} \sum_{i=1}^{n_v} d(p_i, R)^2 \cdot \Delta_i$$

where Δ_i is an area measure around sample point p_i , e.g. one third of the area of the neighboring triangles. For triangle meshes $d(p_i, R)$ can efficiently be evaluated using octree data structures [133]: First, all triangles of the surface R are stored in the octree. Next, only those triangles in octree cells nearest to the point p_i are considered for point-to-triangle distance computations. Finally, the smallest of these distances yields the desired solution.

We use an iterative procedure to solve the optimization problem (6.1), where we separate optimization with respect to b and T :




Figure 6.2 Computation of the gradient of the symmetric root-mean-square distance. Left: from S to R . Right: from R to S

Algorithm 6.1. (Surface Reconstruction)

1. Let (b^i, T^i) denote the solution at time i .
2. Update transformation parameters

$$T^{i+1} = \operatorname{argmin}_T d_2(S(b^i, T), R)$$

by applying the *iterative closest point* (ICP) algorithm [12].

3. Update shape weights

$$b^{i+1} = \operatorname{argmin}_b d_2(S(b, T^{i+1}), R)$$

by applying quasi-Newton optimization [105].

4. Update $i \leftarrow i + 1$ and go back to step (1) if convergence has not been achieved, i.e. $\|(b^i, T^i) - (b^{i+1}, T^{i+1})\| > \epsilon$.

In order to apply quasi-Newton optimization in step (3) we must compute the gradient of $d_2(S(b, T), R)$ with respect to the shape weights b :

Let p_i denote the vertex positions of S and q_j the vertex positions of R (see Fig. 6.2, left). The closest point of $p_i \in S$ is denoted $p'_i \in R$ and their distance is given by

$$d(p_j, R) = |p_i - p'_i|$$

Since p_i is given by the shape model as

$$p_i = \bar{v}_i + \sum_k b_k w_{k,i}$$

where $\bar{v}_i \in \mathbb{R}^3$ and $w_{k,i} \in \mathbb{R}^3$ denote the coordinates of the i -th sample point of the average shape and the eigenvector of the covariance matrix respectively, we obtain

$$\frac{\partial}{\partial b^k} d(p_i, R)^2 = 2 d(p_i, R) \cdot w_{k,i}$$

We approximate the derivative of the second summand in (6.2) (see Fig. 6.2, right) as follows: For each vertex $q_j \in R$ the closest point $q'_j \in S$ is computed. The closest vertex $p_l \in S$ to q'_j is used to approximate the distance $|q_j - q'_j|$ by $|q_j - p_l|$:

$$d(q_j, S(b, T))^2 = |q_j - q'_j|^2 \approx |q_j - p_l|^2$$

The finer the triangle mesh becomes the more accurate this approximation will become. As $p_l = \bar{v}_l + \sum_k b_k w_{k,l}$ we obtain

$$\frac{\partial}{\partial b^k} d(q_j, S)^2 \approx 2 |q_j - p_l| \cdot w_{k,l}$$

The distance d_2 is only one of many different measures that may be used. The basic idea is always to *project* the pathological shape R onto the statistical shape model S . The choice of a suitable distance depends on the application. We present two different medical applications of this technique.

6.3 Reconstruction of Mandibular Dysplasia

Patients with distinct craniofacial deformities or missing bony structures require a surgical reconstruction that in general is a very complex and difficult task. The main reasons for such malformations, as show in Fig. 6.3, are tumor related bone resections or craniofacial microsomia [106, 135]. In cases where the reconstruction cannot be guided by the symmetry of anatomical structures, it becomes particularly challenging. In such a case, the surgeon has to compare the individual pathological situation with a mental image of a regular anatomy in order to modify the affected structures accordingly. For such a surgical therapy osteotomies are typically performed with either subsequent osteodistraction [46] or osteosynthesis after relocation of bony segments [52], sometimes even in combination with selective bone and soft tissue augmentation. In more than 15 cases of mandibular dysplasia and hemifacial microsomia that we have planned so far [143, 140], any kind of guideline for the perception of a designated objective was highly desired. Hence, the aim of our work is to provide a statistical 3D shape model of a human mandible, that will serve as a template for individual treatment planning. Depending on gender and age, different models might be developed as well.

The statistical model used is the same as in Sec. 5.4, yet with the teeth region included. First experiments with this rather small amount of mandibular




Figure 6.3 Three cases of hemifacial microsomia with evidently malformed mandibles




Figure 6.4 top) selection of different shapes of the training set, bottom) three major modes of variation of the average mandible

shape samples already show a broad range of typical variations. With only 10 characteristic shape modes we are able to distinguish between the height of the rami mandibulae, the mandibular angle, the length of the vertical branches, the width of the entire mandible, the radius of the mandibular arch, the shape and the size of the condyles (Fig. 6.4).

After optimization of the rigid transformation T in combination with the shape weights b with regard to a minimal distance d_{rms} between the relevant part of the mandible that is to be reconstructed and the shape model, a mean distance d_{mean} between 1.2 and 1.5 mm with a median of 1.0 to 1.2 mm was achieved (cf. Tab. 6.1). For more than 70, up to 83% of the selected surfaces the deviation was below 2 mm, and only 2–6.5% of the surfaces were

Patient	d_{mean} [mm]	d_{median} [mm]	d_{max} [mm]	$d_r(2\text{ mm})$ [%]	$d_r(3\text{ mm})$ [%]	$d_r(4\text{ mm})$ [%]
P ₁	1.2	1.0	10.4	16.8	5.0	2.0
P ₂	1.4	1.0	9.2	22.7	11.4	6.4
P ₃	1.5	1.2	8.3	29.0	13.5	5.8

Table 6.1 Statistics on the deviation between the surfaces of the shape model and the malformed mandible within a region of interest




Figure 6.5 Template generation for patient P₁, see Tab. 6.1: a) hypoplastic mandible, b) average mandible shape, c) adaptation of the shape model to the right part of the mandible, d) 3D template for mandible reconstruction

deviating more than 4 mm. The maximum distance of 8.3–10.4 mm between the adjusted shape models and the individual mandibles originates from the fact, that a statistical shape model consisting of only 11 samples is far from being representative to describe all variations of a human mandible. However, the value for d_{max} is expected to diminish with a larger training set.

For each of the three pathological cases in Fig. 6.3 we were able to find a suitable candidate from our statistical shape model, using the optimization described above. An example for the first patient is shown in Fig. 6.5. The morphological difference of the optimally aligned two shapes can be measured in size and volume, thus either indicating the thickness of bone augmentation or being suited as a template for autologous bone grafts, the configuration of titanium plates or even the fabrication of individual prostheses.

The statistical 3D shape model of the human mandible seems to be a valuable planning aid for surgical reconstruction of bone defects. This is particularly useful for severe cases of hemifacial microsomnia, as shown in our three examples. With a best matching candidate of the shape model, regarding the size and the shape of available bone, a surgeon gets a good mental perception of the reconstruction that is to be performed.




Figure 6.6 MRI data of normally developed skulls.

6.4 Craniosynostosis

Premature ossified cranial sutures of infants (craniosynostosis) often lead to skull deformities in the growth process. This can lead to increased intracranial pressure, vision, hearing, and breathing problems. Since research on the correction of underlying disorders on the cellular level is still being carried on patients with craniosynostosis depend on surgical intervention for preventing or reducing functional impairment and improving their appearance. The most commonly used surgical procedure consists of bone fragmentation, deformation (reshaping) and repositioning based on standards developed by Paul Tessier [126] and refined by Daniel Marchac and Dominique Renier [107]. A major problem is the evaluation of the aesthetic results of reshaping the cranial vault in small children as the literature does not provide sufficient criteria for assessing skull shape during infancy. A definition of the correct target shape after surgery is missing. The most important and in many cases only indication of the best possible approximation of the skull shape to the unknown healthy shape is left to the subjective aesthetic assessment of the surgeon.

In order to establish objective criteria for the reshaping process, we propose to perform statistical analysis of normally developed cranial shapes. The idea is to compute an average shape and the most characteristic variations from a training set of skulls. Pathological shapes are then projected onto the space spanned by the healthy shapes. The resulting shape will provide a patient-specific proposal for the remodeling process.

A statistical shape model has been created from 21 MRI data sets (patient age: 3 to 10 months) of normally developed skulls (Fig. 6.6). These have been segmented interactively by an anatomical expert. From these segmentations polygonal surfaces have been reconstructed. A typical result is shown in Fig. 6.7 (a). Next, the relevant region for the surgical intervention has been determined. Therefore, four landmarks have been defined on the reconstructed skull surfaces, Fig. 6.7 (b):




Figure 6.7 Surface reconstruction of neurocraniums of infants for statistical shape modeling. (a) reconstruction of the bone surface from segmented data, (b) landmarks, (c) determining cutting planes, (d) final skull patches used to generate the statistical analysis.

- (1) meatus acusticus externus: the entries to the auditory canals.
- (2) nasion: foremost point of the sutura naso-frontalis in the mid-sagittal plane.
- (3) sella: center of the sella turcica (hypophysis)
- (4) occiput: palpable elevation of the os occipitale in the mid-sagittal plane.

These landmarks define a set of planes through which the surface mesh is cut (Fig. 6.7 (c)). The resulting surface of the affected neurocranial region serves as input for the statistical analysis and consists of four sub-regions (Fig. 6.7 (d)): right and left orbital regions, left and right upper skull region.

The completeness of the model was tested in a leave-one-out experiment (cf. Sec. 4.3.1) on all 21 data sets available: on average the model is capable of approximating any other arbitrary skull shape with an error of 0.7 ± 0.2 mm (mean symmetric surface distance).

The 3D cranial model serves as a template for the reshaping process, by finding an optimal fit of any of its variations to a given malformed skull. Usually, no pre-operative MRI scan is available for the infant patients (mostly under the age of one year) in order to avoid unnecessary anaesthesia. Hence the matching of the model towards the pathological skull of the patient is performed by non-invasively measuring anthropometric distances that are not affected by the surgical intervention:

- (1) width between both entries of the auditory canals
- (2) distance from nasion to occiput
- (3) height between vertex and the midpoint of the line between the auditory canals



Figure 6.8 Three different views of a patient with trigonocephaly (ossification of the suture running down the midline of the forehead) - before surgery.




Figure 6.9 Cutting lines indicated on the skull, removed frontal skull region before the reshaping, facsimile of shape model instance on which bone parts are reshaped.




Figure 6.10 Bone stripe before and after reshaping, result of reshaping process on model.

These distances are extrapolated to the skull surface by approximating the skin and skull thickness. The shape model instance that best fits these measurements is selected as a template for the reconstruction process. The resulting shape instance represents an individual interpolation of all shapes contained in the training set.

In a first clinical application, the statistical model was pre-operatively matched to a patient using the method described in section 2.3. From this computed shape model instance a life-size facsimile of the skull was built and taken to the operating room to guide the reshaping process. Fig. 6.8 through 6.12 describe and illustrate the surgical procedure and the role of the statistical skull model (photos taken by F. Hafner).

The application of a statistical shape model as a tool for guiding the skull




Figure 6.11 Microplates for fixating bone pieces on remaining skull are also shaped on the model, result after fixation of reshaped bone on skull.




Figure 6.12 Comparison between pre- and post-operative situation (from left to right): patient 2 months before surgery, immediately before surgery, fac-simile of the target shape derived from the statistical model, patient immediately after surgery, 3 weeks after surgery.

reshaping process in cases of craniosynostosis has proven successful, as shown in a first clinical evaluation. Statistical shape models are capable of providing objective, yet patient-specific criteria for the reshaping process. At the same time they accelerate the process of reshaping as they prevent mistakes or uncertainties followed by time consuming corrections.

We want to extend the number of samples in the training set to improve the completeness of the statistical shape model. It will have to be examined whether the model can be used for segmentation purposes as well, because this is the most time-consuming task in the model generation pipeline.

In this work, the matching of the model was carried out on the basis of a few landmark measurements. In the future we want to explore the possibilities of 3D surface scanners to acquire pre-operative patient data as well as post-operative data for long-term validation.

Chapter 7

3D Shape Estimation from X-Ray Projections

7.1 Introduction

X-ray projection images still play a crucial role in diagnosis and surgery or therapy planning. Accurate computer-aided pre-operative planning, however, requires the knowledge of the 3D-geometry of the anatomy. The problem addressed in this work is how to reconstruct the a-priori unknown 3D-geometry of objects from 2D-projection images.

One of the most prominent applications is the treatment of degenerative joint diseases by artificial hip joint replacement. A large number of hip prostheses are implanted per year. Due to the aging of the population a strong rise in this number is expected for the future. While it is known that the loads on the hip joint play a crucial role for the long-term function and successful performance of artificial joints, there exists no reliable data about expected joint loads for surgical planning. Computer-assisted planning shall help to further improve the treatment in order to assure an individual and optimal biomechanical reconstruction of the hip. To this end, arising forces and loads before and after the intervention shall be computed via validated biomechanical models. Most commonly, the basis for such simulations are X-ray projections. While substantial data for the adjustment of biomechanical model parameters can be extracted from coronal X-ray projections [57], a more accurate analysis of the loads requires the knowledge of the 3D-geometry of the bone and muscles.

A large portion of the work concerning 2D/3D-registration (cf. van de Kraats et al. [131] and references therein) is based on the assumption that there exist pre-operative 3D-data of the patient, which shall be registered to data acquired intra- or post-operatively. This work however addresses the problem of reconstructing 3D-objects from 2D-data, where no such reference is available. Among many differently parameterized deformable surface models (see. Montagnat et al. [94] for an overview), models that incorporate a-priori knowledge about typical shape variations occurring in the object to be re-




Figure 7.1 3D-reconstruction from X-ray projections.

constructed seem to be most suitable for this task due to their robustness. Miscellaneous works build upon this idea:

Fleute and Lavallée [40] use a 3D-statistical shape model of the distal femur for registration with X-ray projections from a C-arm. They minimize the distances between the contours of the model surface and the contours formed by a discrete number of projection rays within the X-ray acquisition setting using the ICP algorithm. Benameur et al. [11] build and use a 3D-statistical shape model of vertebrae for registration and segmentation of X-ray projections. The registration consists of a minimization of an image edge potential, which measures the distance of the projected contours to the contours in the X-ray projections. Tang and Ellis [123] generate a hybrid shape model for the reconstruction of femurs from X-ray data. The correlation between simulated thickness images of the shape model and the X-ray projections serves as the similarity measure for the optimization.

7.2 Reconstruction Process

The basic idea is to project the statistical shape model S into the X-ray plane and measure the distance of the projected shape to the shape in the projection image, i.e. the X-ray projection in our case (Fig. 7.1).

More formally: for a given camera calibration K and a shape model $S(b, T)$ a 2D projection image $P(b, T, K)$ is computed and compared to the X-ray projection X . The reconstruction can now be formulated as the optimization problem

$$(b, T, K)^* = \operatorname{argmin}_{b, T, K} D(P(b, T, K), X) \quad (7.1)$$

where the distance measure D quantifies the mismatch between the two-dimensional images P and X .

The distance function D is non-linear and generally exhibits many local minima. Thus, we use a gradient-descent evolution for the minimization of D : $\dot{x}(t) = -\nabla D(x)$ with some initial value $x = (b_0, T_0, K_0)$. We replace the gradient ∇D by a more suitable search direction, which is computed via a method that evaluates the distance measure at the bounds of a large interval, hence avoiding the solution from being stuck in local minima. In addition, a multi-resolution approach is adopted by performing the registration in a data pyramid: the silhouettes are considered at different resolutions during the course of the optimization. This also speeds up the computation time considerably [110].

For a given camera calibration K (location and orientation of the X-ray source w.r.t. the image acquisition planes) and a given instance of the shape model $S(b, T)$, a thickness image of the shape model in the image acquisition plane is computed by computing the propagation length of simulated rays through the volume enclosed by the shape model (Fig. 7.2). This can efficiently be accomplished using graphics hardware acceleration.

7.3 Distance Measures

The most obvious and fully automatic choice for D would be an intensity correlation between the thickness and the X-ray projection (cf. Tang and Ellis [123]), e.g. by sum of squared differences or mutual information. However, our experiments showed unsatisfactory results on the pelvic data. The approach produces mismatches on the inside of the pelvis since it does not take into account inhomogeneities there (Fig. 7.2 and 7.3). Another cause for problems is the additional structure that is present on the outside in the X-ray data but not in the thickness image.

The second most obvious choice consists in measuring the distance between the edge-maps (cf. Benameur et al. [11]). This approach is almost automatic, as it involves adjusting some parameters of the Canny-filter. Let

$$d(x, s') = \min_{x' \in s'} \|x - x'\| \quad (7.2)$$

be the distance between a point x in the set of contour s of the thickness image and the set of contours s' of the X-ray projection. Then

$$D = \int_{x \in s} d(x, s')^2 dx + \int_{x \in s'} d(x, s)^2 dx. \quad (7.3)$$

In order to efficiently evaluate this distance, the distance map of both the model and the X-ray contours are computed.




Figure 7.2 Canny edge detection in thickness images.




Figure 7.3 Left: Canny edge detection in X-ray projections. Right: Silhouette of the pelvic bone.

The consideration of full edge-maps, however, failed to produce acceptable results, due to the fact that many contours have no well-defined corresponding counterpart (see Fig. 7.2). Hence we propose to consider the silhouettes of the objects instead of the full edge-maps s and s' (Fig. 7.3), which alleviates the problems mentioned.

In contrast to the cited previous work we propose to measure the distance between the model and the data based on their silhouettes. From both the thickness and the X-ray projections accentuated contours are extracted using a Canny [123] edge detector (Fig. 7.2). Finally, the silhouette is extracted from the projection images. The silhouette can be automatically computed from the thickness images: it is its zero level-set, i.e. it separates pixels with thickness > 0 from pixels with thickness $= 0$. However, the silhouette must be determined interactively in the X-ray projections: wrong edges are discarded or missing edges are added to the silhouette (demonstrated in Fig. 7.3 on real X-ray projections). The final contours are rasterized as 2D-images.

7.4 Reconstruction of Pelvic Bone

We have generated a statistical shape model of the pelvic bone from 23 CT data sets of male patients [75] using the method of consistent surface decomposition and parameterization. Each data set has been segmented manually by labeling the pelvic bone regions. Before surface reconstruction a label based interpolation algorithm has been applied to reduce the effect of anisotropy of the voxels (slice thickness of 5 mm, resolution of 1.4 mm in axial slices). For reasons of efficiency all surfaces have been simplified by reducing the number of triangles, obtaining meshes with about 25000 triangles and 12500 nodes. Each pelvic bone surface has been divided into 11 patches (Fig. 7.4):

- the promontorium (P)
- the frontal/ventral sacrum (FS, VS)
- the frontal right/left upper ilium (FRUL, FLUL)

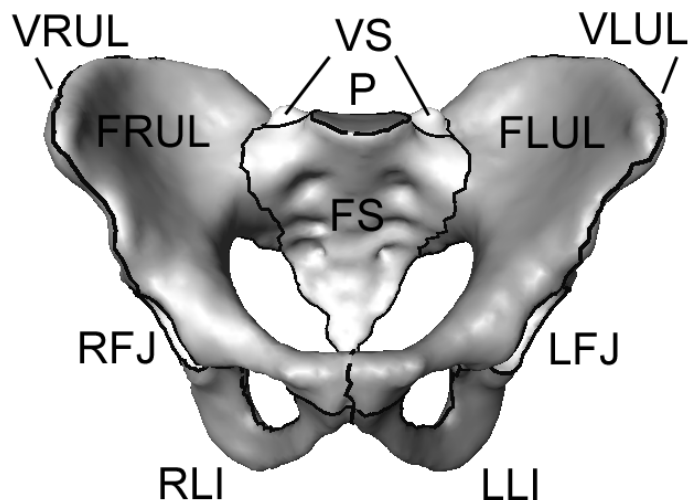


Figure 7.4 Surface decomposition of the pelvic bone.

- the ventral right/left upper ilium (VRUL, VLUL)
- the right/left femur joint (RFJ, LFJ)
- the right/left lower ilium (RLI, LLI)

The user has to specify about 40 landmarks per surface, resulting in 11 patches, bounded by 40 patch boundaries and 25 branching points. Manual interaction amounts to about 30 minutes per training data.

The main modes of variation of the pelvic bone (see Fig. 7.5) model can be well interpreted as: (1) scaling in z-direction, (2) widening and bending of the ilium and (3) change of length of the sacrum.

For validation of the proposed method 23 CT data sets of the abdomen without bone defects have been used (resolution $1 \times 1 \times 5$ mm). For all data sets manual segmentations of the pelvic bone have been available as a gold standard for quantitative evaluation. The goal of this evaluation is to examine the accuracy achievable with the proposed method. The main ingredients of this method to be tested are the statistical shape model and the similarity measure. Therefore we assume that the camera calibration K and the linear transformation T is known and need not be optimized. For evaluation purposes these parameters are known from the generation of the simulated X-ray projections. In fact these assumptions often are met in real world: in conventional surgical planning calibrated X-ray projections can be generated under standardized acquisition conditions [127].

All experiments were conducted for one X-ray projection (coronal: CO) and two X-ray projections (CO plus sagittal: CO-SA). The error was measured by




Figure 7.5 Variability of a statistical model of the pelvic bone shape made from 23 training data sets: in the left column the eigenmode with the largest variance λ_1 is varied between $\pm 3\sqrt{\lambda_1}$, in the second and third column the modes with the second and third largest variance are shown respectively.

computing the symmetric mean surface distances between optimization result and gold standard. Three different experiments were performed (see Tab. 7.1 for results):

- (a) *Leave-All-In Test (LI)*: As a test of the similarity measure and the optimization strategy to capture the true 3D-shape, the reconstruction was carried out with a statistical shape model that contained the shape to be reconstructed. The error should ideally reduce to zero.
- (b) *Leave-One-Out Test (LO)*: The shape to be reconstructed was removed from the shape model. This represents the “real-world” situation.
- (c) *Surface-Optimization (SO)*: As a reference value for the leave-one-out test the shape model was directly matched to the gold standard surfaces in leave-one-out test, where the surface distance between the two shapes was directly minimized [75]. This yields the optimal result to be achieved with a given statistical shape model in the leave-one-out test (b).

7.5 Conclusions

In this work a new method for the reconstruction of unknown 3D-shapes from X-ray projections was presented. The method was validated based on synthetic X-ray data (from pelvic CT) with known camera calibration. The

7. 3D SHAPE ESTIMATION FROM X-RAY PROJECTIONS

Test	Mean [mm]	Median [mm]	Max [mm]
LI (CO)	$1,5 \pm 0,5$	$1,2 \pm 0,3$	$8,8 \pm 3,3$
LI (CO-SA)	$1,3 \pm 0,5$	$1,1 \pm 0,5$	$7,9 \pm 2,6$
LO (CO)	$2,6 \pm 0,4$	$2,1 \pm 0,3$	$17,6 \pm 6,5$
LO (CO-SA)	$2,4 \pm 0,4$	$2,0 \pm 0,3$	$14,9 \pm 3,1$
SO	$2,0 \pm 0,2$	$1,6 \pm 0,2$	$13,3 \pm 2,6$

Table 7.1 2D-3D reconstruction of the pelvic bone. Experimental results (mean values and standard deviations across 23 data sets)

mean values of the measured errors even for a single projection image (coronal) were up to a few millimeters within a range where one can expect to obtain a sufficient estimate of the 3D-geometry with respect to the application of computing load conditions for biomechanical studies [57]. This shall be validated on clinical data in the future. For a complete analysis of the loads occurring in the context of artificial hip joint replacements the method shall be extended to include the femur in a next step. The leave-one-out test shows that the largest portion of the error stems from the incompleteness of the statistical model. Hence the model shall be enlarged by more training data sets in the future. For the clinical application of the method it must be examined to what extent the method can cope with occlusions, artifacts in the X-ray projection or pathological situations, e.g. such as degenerative changes of the pelvic bone (Fig. 7.6). To this end it must be investigated if the method yields sufficient results in cases of incomplete silhouette information. A potential increase in accuracy consists in incorporating further information from the X-ray data and the shape model into the similarity measure for the registration. Hence future work shall attempt to combine silhouette information with modified thickness images into the registration process.



Figure 7.6 X-ray projections with occlusions and pathologies.

Conclusions

This thesis has contributed to the solution of some open problems in the field of 3D geometry reconstruction based on the consideration of a-priori knowledge about shape variability and other application-specific properties.

The central mathematical tool, used in this work, is a statistical model of inter-individual shape variability. In order to establish correspondences between different shapes, two approaches have been developed. With our first approach, large numbers of training data can be processed efficiently to generate shape models. Secondly, initial steps towards a multi-level treatment of variational surface matching have been presented. Both approaches are restricted to homeomorphic surfaces, and hence cannot deal with missing parts or topological variations. However, they can handle surfaces of arbitrary topology and can be robustly implemented for triangle meshes.

Statistical shape models have successfully been applied to image segmentation. They increase the robustness at the expense of flexibility, and even accuracy. It remains an issue how to trade-off these antagonistic properties in general. Simultaneously, the deformation models that drive the reconstruction processes leave a lot of space for further improvement and for incorporation of additional application-specific knowledge. However, statistical shape models yield good initializations for additional post-processing. It has been shown for particular anatomical structures that unsupervised segmentation procedures of medical image data can be devised that are almost as accurate as human performance, even in the presence of noise and artifacts. In addition, it has been shown that it is even possible to estimate a three-dimensional shape from two-dimensional projection images, like X-ray.

Moreover, statistical shape models also provide a very useful means for *surgical* reconstruction of missing, malformed or other pathological anatomical structures. Patient-specific, yet objective surgical reconstruction proposals based on such models have been established as a planning criterion for complex surgical interventions, like individual implant or transplant design.

Appendix A

Riemannian Geometry

In this appendix we state elementary definitions and facts from Riemannian geometry. The exposition follows Smith [118], as well as O'Neill [97] and do Carmo [33].

Smooth Manifolds

Let M be a set and E some Banach space. A *smooth atlas* on M is a collection of *charts* (U_i, x_i) (with i ranging in some indexing set \mathcal{I}) satisfying the following properties:

M1 All $U_i \subset M$ and $M = \bigcup_{i \in \mathcal{I}} U_i$.

M2 $x_i : U_i \rightarrow E$ is a bijection onto an open subset $x(U_i)$ of E . For any i, j , $x(U_i \cap U_j)$ is open in E .

M3 The map $x_j \circ x_i^{-1} : x_i(U_i \cap U_j) \rightarrow x_j(U_i \cap U_j)$ is a smooth bijection for each pair of indices i, j .

Let $x : U \rightarrow E'$ be a topological isomorphism from an open set $U \subset M$ onto an open set $E' \subset E$. The chart (U, x) is called *compatible* with the atlas $\{(U_i, x_i)\}$ if each map $x_i \circ x$, defined on a suitable intersection as in M3, is a smooth isomorphism. Two atlases are said to be compatible if each chart of one is compatible with the other. An equivalence class of smooth atlases on M defines a *smooth manifold* on M . We say that M is *modeled* on E . If $E = \mathbb{R}^n$ for some fixed n , then we say that the manifold is *n-dimensional*. The inverse of a chart map $x^{-1} : E \rightarrow U$ is often called a *parameterization* or *system of coordinates* for $U \subset M$.

Maps and Functions

Let M, N be two smooth manifolds, and $\varphi : M \rightarrow N$ a map. φ is called smooth if, given $p \in M$, there exists a chart (U, x) at p and a chart (V, y) at $\varphi(p)$ such that $\varphi(U) \subset V$ and

$$y \circ \varphi \circ x^{-1} : x(U) \rightarrow y(V)$$

is smooth. It follows from M2, that this definition holds for any choice of charts (U, x) and (V, y) with $\varphi(U) \subset V$. Let $C^\infty(M, N)$ denote the set of smooth maps between two manifolds. The special case $C^\infty(M, \mathbb{R})$ will be denoted the set of smooth functions on M .

Tangent Bundle, Vector Fields

A smooth map $\gamma : I \rightarrow M$ on some open interval $I \subset \mathbb{R}$ is called a *curve* in M . Let $\gamma(0) = p \in M$. The tangent vector $V_p := \gamma'(0)$ to the curve γ is a function $V_p : C^\infty(M, \mathbb{R}) \rightarrow \mathbb{R}$ given by

$$V_p(f) = \frac{d(f \circ \gamma)}{dt}(0) \quad , \quad f \in C^\infty(M, \mathbb{R})$$

The *tangent vector* at p is the tangent vector to some curve γ with $\gamma(0) = p$. The set of all tangent vectors to M at p is called the *tangent space* $T_p M$. It is a vector space over \mathbb{R} and the Leibniz rule holds, i.e.

$$V_p(fg) = V_p(f)g(p) + f(p)V_p(g)$$

The *tangent bundle* is defined as

$$TM := \{(p, V_p) : p \in M, V_p \in T_p M\}$$

The tangent bundle can be provided with a smooth structure and hence be made a manifold itself (see [33], Ex. 4.1).

A *vector field* $V : M \rightarrow TM$ on a manifold M is a map that assigns to each point $p \in M$ a tangent vector V_p to M at p . If $f \in C^\infty(M, \mathbb{R})$, then Vf denotes the function on M given by

$$(Vf)(p) = V_p(f) \quad \text{for all } p \in M$$

The set of all smooth vector fields

$$F(M) = C^\infty(M, TM)$$

on M is a module over the ring $C^\infty(M, \mathbb{R})$ with the following operations:

$$\begin{aligned} (fV)_p &= f(p)V_p \\ (V + W)_p &= V_p + W_p \quad \text{for all } p \in M \end{aligned}$$

Covariant Derivative

An *affine connection* D on a smooth manifold M is a map

$$D : F(M) \times F(M) \rightarrow F(M)$$

characterized (but *not* uniquely) by the following axioms:

D1 $D_{fV+gV'}W = fD_VW + gD_{V'}W$

D2 $D_V(W + W') = D_VW + D_VW'$

D3 $D_V(fW) = (Vf)W + fD_VW$ for $f \in C^\infty(M, \mathbb{R})$.

A vector field D_VW is called a *covariant derivative of W with respect to V* for the affine connection D . As shown by do Carmo [33] (sec. 2) an affine connection “furnishes a [intrinsic] manner of differentiating vectors along curves” by satisfying D1 and D2. Axiom D3 allows to show that an affine connection is a local notion.

The definition of the covariant derivative can be extended to arbitrary tensor fields (see O’Neill [97], sec. 3.2). Particularly, for the $(0,0)$ -tensor $f \in C^\infty(M, \mathbb{R})$ we have

$$Df(V) := D_Vf = Vf \tag{A.1}$$

Moreover, this allows to compute higher-order covariant derivatives of f by defining: $D^n f := D \circ \dots \circ Df$. The second covariant derivative of f is a $(0,2)$ -tensor (see O’Neill [97], sec. 3.11) given by

$$D^2f(V, W) = D_W(Df)(V) = WVf - (D_WV)f \tag{A.2}$$

Geodesics and Parallelism

Let $\gamma : I \rightarrow M$ be a smooth curve with tangent vectors $V_{\gamma(t)} = \gamma'(t)$. The curve γ is called a *geodesic* if $D_VV = 0$ for all $t \in I$. Now let $W_{\gamma(t)} \in T_{\gamma(t)}M$ be a smooth family of tangent vectors defined along γ . The family $W_{\gamma(t)}$ is said to be *parallel* along γ if $D_VW = 0$ for all $t \in I$.

For every $p \in M$ and $V_p \neq 0$ in T_pM there exist a unique geodesic $t \rightarrow \gamma(t)$ such that $\gamma(0) = p$, $\gamma'(0) = V_p$ and 1 is in the domain of γ . We define the *exponential map*

$$\exp_p : T_pM \rightarrow M \quad \text{by} \quad \exp_p(V_p) := \gamma(1)$$

for all $V_p \in T_pM$. We write for $\gamma(t) = \exp_p(tV_p)$. The image of a star-shaped neighbourhood N_0 of the origin of T_pM is called the *normal neighbourhood*,




Figure A.1 Relation between parallelism and covariant differentiation.

denoted by N_p . It can be shown to always exist.

Given a curve $\gamma : I \rightarrow M$ with $\gamma(0) = p$ for each V_p in $T_p M$ there exists a unique family $V_{\gamma(t)} \in T_{\gamma(t)} M$ ($t \in I$) of tangent vectors parallel along γ such that $V_{\gamma(0)} = V_p$. This parallelism induces an isomorphism

$$\tau_h : T_p M \rightarrow T_q M \quad \text{by} \quad \tau_h(V_p) := V_{\gamma(h)}$$

for $\gamma(h) = q$ and $h \in I$. $\tau_h(V_p)$ is called *parallel transport* of V_p along γ . A vector field W on a normal neighbourhood N_p is called *adapted* to the tangent vector $V_p \in T_p M$ if $W_q = \tau_h(V_p)$ is a parallel translation along the unique *geodesic* segment joining p and q .

Parallel transport and covariant differentiation are related in the following way (see Fig. A.1) Let V be a vector field on M and γ be an integral curve of V . Then for vector field W on M :

$$(D_V W)_p = \lim_{h \rightarrow 0} \frac{1}{h} (\tau_h^{-1}(W_{\gamma(h)}) - W_p)$$

This relation holds if W is replaced by an arbitrary tensor on M .

Riemannian Manifolds

A *Riemannian structure* g_p is a $(0, 2)$ -tensor that assigns to each point $p \in M$ a positive definite symmetric bilinear form

$$g_p : T_p M \times T_p M \rightarrow \mathbb{R}$$

This makes the tangent space $T_p M$ a Hilbert space. For a smooth connected manifold M the pair (M, g) is called a *Riemannian manifold*. For vector fields we have

$$g : F(M) \times F(M) \rightarrow C^\infty(M, \mathbb{R})$$

A. RIEMANNIAN GEOMETRY

For every p in M the Riemannian structure provides an inner product on T_pM . For $V_p \in T_pM$ the norm is given by

$$\|V_p\| = g_p(V_p, V_p)^{1/2}.$$

Note that a bilinear symmetric form can be obtained from a given norm using the polarization identity

$$g_p(V_p, W_p) = \frac{1}{2} (\|V_p + W_p\|^2 - \|V_p\|^2 - \|W_p\|^2)$$

The length of a curve $t \rightarrow \gamma(t)$ in M with $t \in [a, b]$ is defined as

$$L(\gamma) = \int_a^b \|\gamma'(t)\| dt$$

For $\gamma(a) = p$ and $\gamma(b) = q$ the *Riemannian metric*

$$d(p, q) = \inf_{\gamma} L(\gamma)$$

yields a metric on M . It can be shown that the distance d between two points is the length of the shortest geodesic connecting these two points. Moreover we have

$$d(p, \exp_p(tV_p)) = t \cdot \|V_p\| \tag{A.3}$$

For a vectorfield V and some curve γ with $\gamma(0) = p$ and $\gamma'(t) = V_{\gamma(t)}$, we have

$$\|V_p\| = \left. \frac{d}{dt} \right|_{t=0} d(p, \gamma(t)) \tag{A.4}$$

Riemannian Connection

For every Riemannian manifold (M, g) there exists a unique connection ∇ , called the *Levi-Civita or Riemannian connection*, satisfying

D4 $[V, W] = \nabla_V W - \nabla_W V$

D5 $Ug(V, W) = g(\nabla_U V, W) + g(V, \nabla_U W)$

for $U, V, W \in F(M)$. D4 is the definition for a *symmetric* affine connection. D5 is equivalent to the statement that the connection ∇ is *compatible* with the metric g , i.e. for any smooth curve γ on M and any two parallel vector fields $V, W \in F(M)$ along γ we have $g(V, W) = \text{constant}$ (see do Carmo [33], sec. 2.3). In other words, parallel transport is an isometry.

For manifolds embedded in Euclidean space, the Levi-Civita connection is the directional derivative in the embedding space followed by an orthogonal projection onto T_pM (see O'Neill [97], sec. 4.3).

For the Riemannian connection we can define the gradient and Hessian operator as follows:

$$\nabla f(V) = g(\text{grad}f, V) \quad \text{with} \quad \text{grad}f \in F(M)$$

and

$$\nabla^2 f(V, W) = g(\text{Hess}f(V), W) \quad \text{with} \quad \text{Hess}f : F(M) \rightarrow F(M)$$

The following relationship holds (see O'Neill [97], sec. 3.11)

$$\text{Hess}f(V) = \nabla_V(\text{grad}f)$$

Due to the symmetry property D4, the second covariant derivative of f is symmetric, and hence the associated Hessian is a symmetric operator.

Appendix B

Discrete Gauss Curvature

Let T be a triangle mesh with vertices $v_1 \dots, v_n$. Let ϕ_i be the piecewise linear function with value 1 at vertex v_i and 0 at all other vertices at the mesh (Lagrange basis). The mass matrix is given by

$$M_{ij} = \int_M \phi_i(p) \phi_j(p) dA \quad (\text{B.1})$$

If v_i and v_j are neighbored vertices in the mesh then M_{ij} is $1/12$ the area of the triangles adjacent to edge (i, j) . Diagonal entries M_{ii} are $1/6$ the area of all triangles adjacent to vertex v_i . All other entries are zero.

The integrated Gauss curvature K_i at interior and boundary vertices i is defined by the angle excesses

$$K_i = 2\pi - \theta_i \text{ (interior)} \quad \text{and} \quad K_i = \pi - \theta_i \text{ (boundary)}$$

where θ_i is the sum of all angles adjacent to vertex v_i . Hence, the pointwise Gauss curvature κ_i at vertex v_i can be defined via the solution of the sparse linear system

$$\boxed{M\kappa = K} \quad (\text{B.2})$$

where $K = (K_1, \dots, K_n)^T$ and $\kappa = (\kappa_1, \dots, \kappa_n)^T$. This definition for the pointwise Gauss curvature satisfies the Gauss-Bonnet theorem

$$2\pi\chi = \sum_i K_i = \text{id}K = \text{id}M\kappa = \int_M \text{id}\kappa dA$$

Appendix C

The Magnitude of the Area Gradient

We describe here the computation of the integral (3.14)

$$\sigma = \frac{d}{d\epsilon} \left(\int_0^L u(s) ds \right) \quad (\text{C.1})$$

by numerically solving a boundary problem for the normal geodesic variations u along a given geodesic $\gamma : [0, L] \rightarrow M$ on a surface M . Discretizing u in (3.16) and the boundary constraints (3.15) by equally deviding the interval $[0, L]$ into n smaller subintervals yields the following system of equations:

$$\begin{aligned} \frac{1}{h^2}(u_{i-1} - 2u_i + u_{i+1}) + \kappa_i u_i &= 0 \\ u_0 &= 0 \\ u_n &= \epsilon \end{aligned} \quad (\text{C.2})$$

with $h = L/n$ and $\kappa_i = \kappa(ih)$ with $i = 0, \dots, n$ is the Gauss curvature at discrete sample points along the curve γ .

System (C.2) is a tridiagonal linear system $Au = b$ with

$$A = \begin{pmatrix} 1 & 0 & & & \\ 1 & (h^2\kappa_1 - 2) & 1 & & 0 \\ & 1 & (h^2\kappa_2 - 2) & 1 & \\ & & & \ddots & \\ & 0 & & 1 & (h^2\kappa_{n-1} - 2) & 1 \\ & & & & 0 & 1 \end{pmatrix}, \quad (\text{C.3})$$

and $u = (u_1, \dots, u_n)^T$ and $b = (0, \dots, \epsilon)^T$. We approximate the integral in equation (C.1) using the trapezoidal rule:

$$\int_0^L u(s) ds \approx h \left(\frac{u_0 + u_n}{2} + \sum_{i=1}^{n-1} u_i \right) \quad (\text{C.4})$$

C. THE MAGNITUDE OF THE AREA GRADIENT

Due to the linear dependence of u on ϵ the differentiation of (C.4) with respect to ϵ is equivalent to setting $\epsilon = 1$ in the linear system (C.2). Then σ is given by the formula in (C.4). Note, that the system (C.2) is non-symmetric and generally indefinite.

Bibliography

- [1] P.-A. Absil. Vector transport on manifolds. Talk at ICIAM 2007, July 2007. URL http://www.inma.ucl.ac.be/~absil/Talks/ICIAM070717_oom_05.pdf.
- [2] P.-A. Absil, C. G. Baker, and K. A. Gallivan. Trust-region methods on riemannian manifolds with applications in numerical linear algebra. *Foundations of Computational Mathematics*, 2006. accepted for publication.
- [3] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, January 2008.
- [4] R. L. Adler, J.-P. Dedieu, J. Y. Margulies, M. Martens, and M. Shub. Newton’s method on riemannian manifolds and a geometric model for the human spine. *IMA J. Numer. Anal.*, 22:359 – 390, 2002.
- [5] H. Alt and L. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In *Handbook of Computational Geometry*, pages 121–153. Elsevier Science Publishers B.V. North-Holland, 1999.
- [6] N. Amenta and Y. J. Kil. Defining point-set surfaces. *ACM Trans. Graph.*, 23(3):264–270, 2004.
- [7] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, H. Pang, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proc. of the Neural Information Processing Systems (NIPS) Conference*, 2004.
- [8] M. Audette, F. Ferrie, and T. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis*, 4(3):201–217, 2000.
- [9] M. G. Ballester, M. Linguraru, M. R. Aguirre, and N. Ayache. On the adequacy of principal factor analysis for the study of shape variability. In J. Fitzpatrick and J. Reinhardt, editors, *SPIE Medical Imaging ’05*, volume 5747, pages 1392–1399. SPIE Publishing, 2005.

- [10] P. Beckmann. Consistent cell decomposition of homeomorphic simplicial surfaces. Master's thesis, Freie Universität Berlin, 2005.
- [11] S. Benameur, M. Mignotte, S. Parent, H. Labelle, W. Skalli, and J. de Guise. 3D/2D registration and segmentation of scoliotic vertebrae using statistical models. *Computerized Medical Imaging and Graphics*, 27(5):321–337, 2003.
- [12] J. P. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. of Pattern Anal. and Mach. Intell.*, 14(2):239–256, 1992.
- [13] O. Bihun and C. Chicone. Distortion minimal morphing i: The theory for stretching. *ArXiv Mathematics e-prints*, 2006.
- [14] J. Bloomenthal. *An introduction to implicit surfaces*. Morgan-Kaufman, San Francisco, 1997.
- [15] O. Bonnet. Mémoire sur la théorie des surfaces applicables. *J. Éc. Polyt.*, 42, 1867.
- [16] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, 1989.
- [17] G. Borgefors, I. Nyström, and G. S. di Baja. Computing skeletons in three dimensions. *Pattern recognition*, 32(7):1225–1236, 1999.
- [18] M. Botsch and O. Sorkine. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):213–230, 2008.
- [19] M. Botsch, M. Pauly, C. Rössl, S. Bischoff, and L. Kobbelt. Geometric modeling based on triangle meshes. SIGGRAPH 2006 Course Notes, 2006.
- [20] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Calculus of non-rigid surfaces for geometry and texture manipulation. *IEEE Trans. Visualization and Computer Graphics*, 13(5):902–913, 2007.
- [21] M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Efficient computation of isometry-invariant distances between surfaces. Technical Report CIS-2006-02, Dept. of Computer Science, Technion, Israel, 2006.
- [22] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *Image Processing, IEEE Transactions on*, 5(10):1435–1447, 1996.

-
- [23] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89:114–141, 2003.
- [24] P. G. Ciarlet. *An Introduction to Differential Geometry, with Applications to Elasticity*. Springer Dordrecht, 2005.
- [25] U. Clarenz, N. Litke, and M. Rumpf. Axioms and variational problems in surface parameterization. *Computer Aided Geometric Design*, 21 (8): 727 – 749, 2004.
- [26] T. Cootes. Statistical models of appearance for computer vision. Ongoing report draft, 2004. URL http://www.isbe.man.ac.uk/~bim/Models/app_models.pdf.
- [27] T. Cootes and C. J. Taylor. Statistical models of appearance for medical image analysis and computer vision. In *Proc. SPIE Medical Imaging*, 2001.
- [28] 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*, 12:355–366, 1994.
- [29] O. Cuisenaire. *Distance Transformations: Fast Algorithms and Applications to Medical Image Processing*. PhD thesis, Université catholique de Louvain, 1999.
- [30] R. H. Davies. *Learning Shape: Optimal Models for Analysing Natural Variability*. PhD thesis, University of Manchester, 2002.
- [31] R. H. Davies, C. Twining, T. Cootes, J. C. Waterton, and C. J. Taylor. A minimum description length approach to statistical shape modelling. *IEEE Transactions on Medical Imaging*, 21(5):525–537, 2002.
- [32] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. URL <http://jmvidal.cse.sc.edu/library/dijkstra59a.pdf>.
- [33] M. P. do Carmo. *Differential geometry of curves and surfaces*. Prentice Hall, 1976.
- [34] M. P. do Carmo. *Riemannian Geometry*. Birkhäuser, 1992.
- [35] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. John Wiley & Sons, 1998.
- [36] T. Duchamp, A. Certain, T. DeRose, and W. Stuetzle. Hierarchical computation of PL harmonic embeddings. Technical report, University of Washington, July 1997.

- [37] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *ACM SIGGRAPH*, volume 29, pages 173–182, 1995.
- [38] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(10):1285–1295, 2003.
- [39] J. Erickson and S. Har-Peled. Optimally cutting a surface into a disk. In *Symposium on Computational Geometry*, pages 244–253. ACM, 2002.
- [40] M. Fleute and S. Lavallée. Nonrigid 3-d/2-d registration of images using statistical models. In *Proc. MICCAI*, volume 1679, pages 138–147, 1999.
- [41] M. S. Floater. Parameterization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14(3):231–250, 1997.
- [42] M. S. Floater. Mean value coordinates. *Computer Aided Geometric Design*, 20:19–27, 2003.
- [43] M. S. Floater and K. Hormann. Surface parameterization: a tutorial and survey. In N. A. Dodgson, M. S. Floater, and M. A. Sabin, editors, *Advances in Multiresolution for Geometric Modelling*, Mathematics and Visualization, pages 157–186. Springer, Berlin, Heidelberg, 2005.
- [44] D. Gabay. Minimizing a differentiable function over a differential manifold. *Journal of Optimization Theory and Applications*, 37(2):177–219, 1982.
- [45] G. Gerig, M. Styner, M. E. Shenton, and J. A. Lieberman. Shape versus size: Improved understanding of the morphology of brain structures. *LNCS*, 2208:24–32, 2001.
- [46] P. Glat and J. McCarthy. Distraction of the mandible, experimental studies. In J. McCarthy, editor, *Distraction in the Craniofacial Skeleton*, pages 67–203. Springer-Verlag, 1999.
- [47] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
- [48] X. Gu and B. C. Vemuri. Matching 3d shapes using 2d conformal representations. In C. Barillot, D. Haynor, J. F. e Cunha, and P. Hellier, editors, *Proc. MICCAI*, volume 3216 of *LNCS*, pages 771–780. Springer, 2004.
- [49] X. Gu, Y. Wang, T. Chan, P. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. In *Proc. IPMI*, volume 2732 of *LNCS*, pages 172–184. Springer, 2003.

-
- [50] J. H. Blaine Lawson and R. D. A. Tribuzy. On the mean curvature function for compact surfaces. *J. for differential geometry*, 16:179–183, 1981.
- [51] H. Haberl, B. Hell, Maja Zöckler, S. Zachow, H. Lamecker, A. Sarrafzadeh, B. Riecke, W. R. Langsch, P. Deuffhard, J. Bier, and M. Brock. Technical aspects and results of surgery for craniosynostosis. *Zentralblatt für Neurochirurgie*, 65(2):65–74, 2004.
- [52] F. Härle, M. Champy, and B. C. Terry. *Atlas of Craniomaxillofacial Osteosynthesis*. Thieme, 1999. Drawings by A. Reinhardt.
- [53] T. Heimann and H.-P. Meinzer. Statistical shape models for 3d medical image analysis: a review. *Medical Image Analysis*, 2008. submitted for publication.
- [54] T. Heimann, I. Wolf, and H.-P. Meinzer. Active shape models for a fully automated 3d segmentation of the liver - an evaluation on clinical data. In R. Larsen, M. Nielsen, and J. Sporring, editors, *Proc. MICCAI*, volume 4191 of *LNCS*, pages 41–48. Springer-Verlag, 2006.
- [55] T. Heimann, S. Münzing, H.-P. Meinzer, and I. Wolf. A shape-guided deformable model with evolutionary algorithm initialization for 3d soft tissue segmentation. In *Proc. IPMI*, volume 4584 of *LNCS*, pages 1–12. Springer-Verlag, 2007.
- [56] T. Heimann, B. van Ginneken, M. Styner, Y. Arzhaeva, V. Aurich, C. B. Christian Bauer, R. Beichel, G. Bekes, F. Bello, G. Binnig, A. B. Horst Bischof, P. M. M. Cashman, Y. Chi, A. Córdova, B. Dawant, M. Fidrich, D. F. Jacob Furst, L. Grenacher, J. Hornegger, D. Kainmüller, H. K. Richard I. Kitney, H. Lamecker, T. Lange, J. Lee, B. Lennon, R. Li, H.-P. M. Senhu Li, G. Németh, D. S. Raicu, A.-M. Rau, E. van Rikxoort, L. R. Mikaël Rousson, K. A. Saddi, G. Schmidt, D. Seghers, A. Shimizu, E. S. Pieter Slagmolen, G. Soza, R. Susomboon, A. Wimmer, and I. Wolf. Comparison and evaluation of methods for liver segmentation from CT datasets. *Medical Image Analysis*, 2008. in preparation.
- [57] M. Heller, G. Bergmann, G. Deuretzbacher, L. Durselen, M. Pohl, L. Claes, N. Haas, and G. Duda. Musculo-skeletal loading conditions at the hip during walking and stair climbing. *J Biomech*, 34(7):883–893, 2001.
- [58] K. Hildebrandt, K. Polthier, and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata*, 123(24):89–112, 2006.

- [59] K. Hüper and J. Trumppf. Newton-like methods for numerical optimization on manifolds. In *Proc. 38th Asilomar Conference on Signals, Systems and Computers*, pages 136–139, 2004.
- [60] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley & Sons, 2001.
- [61] A. Ion, G. Peyré, Y. Haxhimusa, S. Peltier, W. G. Kropatsch, and L. Cohen. Shape matching using the geodesic eccentricity transform - a study. In *Proc. of Austrian Association for Pattern Recognition*, 2007.
- [62] A. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, 1997.
- [63] D. Kainmüller, T. Lange, and H. Lamecker. Shape constrained automatic segmentation of the liver based on a heuristic intensity model. In B. v. G. T. Heimann, M. Styner, editor, *Proc. MICCAI Workshop 3D Segmentation in the Clinic: A Grand Challenge*, pages 109–116, 2007. URL <http://www.zib.de/lamecker/publications/miccai2007.pdf>.
- [64] G. Kamberov and G. Kamberova. Ill-posed problems in surface and surface shape recovery. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, volume 1, pages 552–557, 2000.
- [65] G. Kamberov, F. Pedit, and U. Pinkall. Bonnet pairs and isothermic surfaces. *Duke Math. J.*, 92(3):637–643, 1998.
- [66] A. Kelemen, G. Székely, and G. Gerig. Three-dimensional model-based segmentation of brain MRI. *IEEE Trans. on Medical Imaging*, 18(10):828–839, 1999.
- [67] A. Khamene, L. Schwarz, D. Zikic, F. Azar, E. Rietzel, and N. Navab. A unified and efficient approach for free-form deformable registration. *IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, 2007.
- [68] R. Kimmel and J. A. Sethian. Computing geodesic on manifolds. *Proc. US National Academy of Science*, 95:8431–8435, 1998.
- [69] E. Kreyszig. *Differential Geometry*. Courier Dover, 1991.
- [70] A. Kriegl and P. Michor. *The convenient setting of global analysis*. Number 53 in Mathematical Surveys and Monographs. American Mathematical Society, 1997.
- [71] H. Lamecker, T. Lange, and M. Seebaß. A statistical shape model for the liver. In T. Dohi and R. Kikinis, editors, *Proc. Medical Image Computing and Computer Assisted Intervention*, volume 2489 of *Lecture Notes in*

Computer Science, pages 422–427. Springer, 2002. URL <http://www.zib.de/lamecker/publications/miccai2002.pdf>.

- [72] H. Lamecker, T. Lange, and M. Seebaß. Erzeugung statistischer 3D-Formmodelle zur Segmentierung medizinischer Bilddaten. In T. W. et al., editor, *Proc. Bildverarbeitung für die Medizin*, Informatik aktuell, pages 398–403. Springer, 2003. URL <http://www.zib.de/lamecker/publications/bvm2003a.pdf>.
- [73] H. Lamecker, T. Lange, M. Seebaß, S. Eulenstein, M. Westerhoff, and H.-C. Hege. Automatic segmentation of the liver for the preoperative planning of resections. In J. W. et al., editor, *Proc. Medicine Meets Virtual Reality*, volume 94 of *Studies in Health Technologies and Informatics*, pages 171–174. IOS press, 2003. URL <http://www.zib.de/lamecker/publications/mmvr2003.pdf>.
- [74] H. Lamecker, T. Lange, and M. Seebaß. Segmentation of the liver using a 3d statistical shape model. Technical report, Zuse Institute Berlin, 2004. URL <http://www.zib.de/PaperWeb/abstracts/ZR-04-09/>.
- [75] H. Lamecker, M. Seebaß, H.-C. Hege, and P. Deuffhard. A 3d statistical shape model of the pelvic bone for segmentation. In J. Fitzpatrick and M. Sonka, editors, *Proc. SPIE - Medical Imaging: Image Processing*, volume 5370, pages 1341–1351, 2004. URL <http://www.zib.de/lamecker/publications/spie2004.pdf>.
- [76] H. Lamecker, Maja Zöckler, H. Haberl, S. Zachow, and H.-C. Hege. Statistical shape modeling for craniostylosis planning. In *Proc. Advanced Digital Technology in Head and Neck Reconstruction*, page 64, 2005. URL <http://www.zib.de/lamecker/publications/Banff2005.txt>.
- [77] H. Lamecker, T. H. Wenckebach, and H.-C. Hege. Atlas-based 3D-shape reconstruction from x-ray images. In *Proc. Int. Conf. of Pattern Recognition (ICPR2006)*, volume Volume I, pages 371–374. IEEE Computer Society, 2006. URL <http://www.zib.de/lamecker/publications/icpr2006.pdf>.
- [78] H. Lamecker, S. Zachow, H.-C. Hege, and Maja Zöckler. Surgical treatment of craniostylosis based on a statistical 3D-shape model: First clinical application. *Int. J. Computer Assisted Radiology and Surgery*, 1(1):253–254, 2006. URL <http://www.zib.de/lamecker/publications/cars2006b.pdf>.
- [79] H. Lamecker, S. Zachow, A. Wittmers, B. Weber, H.-C. Hege, B. Elsholtz, and M. Stiller. Automatic segmentation of mandibles in low-dose CT-data. *Int. J. Computer Assisted Radiology and Surgery*, 1(1):

- 393–395, 2006. URL <http://www.zib.de/lamecker/publications/cars2006a.pdf>.
- [80] F. Lazarus and A. Verroust. Three-dimensional metamorphosis: a survey. *The Visual Computer*, 14:373–389, 1998.
- [81] D. Lee and F. Preparata. Euclidean shortest path in the presence of rectilinear barriers. *Networks*, 13:393–410, 1984.
- [82] N. Litke, M. Droske, M. Rumpf, and P. Schröder. An image processing approach to surface matching. In M. Desbrun and H. Pottmann, editors, *Symposium on Geometry Processing*, volume 255 of *ACM International Conference Proceeding Series*, pages 207–216. Eurographics Association, 2005.
- [83] Maja Zöckler. *Modellgebundene Cranioplastie - Operationstechnik zur Umformung frühkindlicher Schädeldeformitäten unter Verwendung dreidimensionaler Standardformmodelle aus MRT-basierten Rekonstruktionen nicht deformierter Kinder*. PhD thesis, Humboldt Universität Berlin, 2006.
- [84] Malte Zöckler, D. Stalling, and H.-C. Hege. Fast and intuitive generation of geometric shape transitions. *The Visual Computer*, 16(5):241–253, 2000.
- [85] J. H. Manton. Optimisation algorithms exploiting unitary constraints. *IEEE Transactions on Signal Processing*, 50(3):635–650, March 2002.
- [86] D. Martinez, L. Velho, and P. C. Carvalho. Geodesic paths on triangular meshes. In *Proc. 17th Brazilian Symp. Computer Graphics and Image Processing*, pages 210–217, 2004.
- [87] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Medical Image Analysis*, 1(2):91–108, 1996.
- [88] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(10):1615–1630, 2005.
- [89] S. A. Miller and J. Malick. Newton methods for nonsmooth convex minimization: connections among u-lagrangian, riemannian newton and sqp methods. *Math. Program. Ser. B*, 104:609–633, 2005.
- [90] J. Mitchell, D. Mount, and C. Papadimitriou. The discrete geodesic problem. *SIAM J. of Computing*, 16(4):647–668, 1987.

-
- [91] N. J. Mitra, L. Guibas, J. Giesen, and M. Pauly. Probabilistic fingerprints for shapes. In A. Sheffer and K. Polthier, editors, *Symposium on Geometry Processing*, volume 256 of *ACM International Conference Proceeding Series*, pages 121–130. Eurographics Association, 2006.
- [92] J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, 2004.
- [93] J. Montagnat and H. Delingette. Volumetric medical images segmentation using shape constrained deformable models. In *CVRMed*, pages 13–22, 1997.
- [94] J. Montagnat, H. Delingette, and N. Ayache. A review of deformable surfaces: topology, geometry and deformation. *Image and Vision Computing*, 19(14):1023–1040, 2001.
- [95] C. Nastar and N. Ayache. Fast segmentation, tracking and analysis of deformable objects. In *Conf. on Comp. Vision ICCV*, pages 275–279. IEEE Computer Society Press, 1993.
- [96] Z. Nitecki. *Differentiable Dynamics: an Introduction to the Orbit Structure of Diffeomorphisms*. MIT Press: Cambridge, Mass., 1971.
- [97] B. O’Neill. *Semi-Riemannian Geometry*. Academic Press, 1983.
- [98] N. Paragios, M. Taron, X. Huang, M. Rousson, and D. Metaxas. *Statistics and Analysis of Shapes*, chapter On the Representation of Shapes Using Implicit Functions., pages 167–200. Springer Verlag, 2006.
- [99] X. Pennec, R. Stefanescu, V. Arsigny, P. Fillard, and N. Ayache. Riemannian elasticity: A statistical regularization framework for non-linear registration. In J. Duncan and G. Gerig, editors, *Proc. MICCAI*, volume 3750 of *LNCS*, pages 943–950. Springer, 2005.
- [100] L. Piegl. On nurbs: A survey. *IEEE Computer Graphics and Applications*, pages 55–71, 1991.
- [101] B. M. Planitz, A. J. Maeder, and J. A. Williams. The correspondence framework for 3d surface matching algorithms. *Comput. Vis. Image Underst.*, 97(3):347–383, 2005.
- [102] K. Polthier and M. Schmies. Straightest geodesics on polyhedral surfaces. In H. C. Hege and K. Polthier, editors, *Mathematical Visualization*, page 391. Springer, 1998.
- [103] K. Polthier, E. Preuss, and U. Pinkall. Locally shortest paths on triangle meshes. personal communication, 2006.

- [104] E. Praun, W. Sweldens, and P. Schröder. Consistent mesh parameterizations. In *ACM SIGGRAPH*, pages 179–184, 2001.
- [105] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C – The Art of Scientific Computing*. Cambridge University Press, 1986.
- [106] S. Pruzansky. Not all dwarfed mandibles are alike. In Williams and Wilkins, editors, *The Clinical Delineation of Birth Defects*, volume 5 of *Birth Defects*, pages 120–129, 1969.
- [107] D. Renier, E. Lajeunie, E. Arnaud, and D. Marchac. Management of craniosynostoses. *Childs Nerv Syst*, 16(10-11):645–658, 2000.
- [108] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-beltrami spectra as "shape-dna" of surfaces and solids. *Computer-Aided Design*, 38(4): 342–366, 2006.
- [109] D. Rückert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, , and D. J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Trans Med Imaging*, 18(8):712–721, 1999.
- [110] D. Rueckert. Non-rigid registration: Techniques and applications. In J. V. Hajnal, D. L. G. Hill, and D. J. Hawkes, editors, *Medical Image Registration*. CRC Press, 2001.
- [111] D. Salomon. *Curves and Surfaces for Computer Graphics*. Springer, 2006.
- [112] P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *ACM SIGGRAPH*, pages 409–416, 2001.
- [113] J. Schmidt-Ehrenberg. *Analysis and visualization of molecular conformations*. PhD thesis, Freie Universität Berlin, 2008.
- [114] B. Schölkopf, F. Steinke, and V. Blanz. Object correspondence as a machine learning problem. In *Proc. Int. Conf. on Machine Learning*, pages 776–783. ACM, 2005.
- [115] A. Schramm, M. Rücker, N. Sakkas, R. Schön, J. Dükerb, and N.-C. Gellrich. The use of cone beam ct in cranio-maxillofacial surgery. In H. Lemke, K. Inamura, K. Doi, M. Vannier, and A. Farman, editors, *CARS 2005: Computer Assisted Radiology and Surgery*, volume 1281 of *Int Congress Series*, pages 1200–1204. Elsevier, 2005.
- [116] J. Schreiner, A. Asirvatham, E. Praun, and H. Hoppe. Inter-surface mapping. *Trans. Graph.*, 23(3):870–877, 2004.

-
- [117] J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications*, 22(1-3):21–74, May 2002.
- [118] S. T. Smith. *Geometric optimization methods for adaptive filtering*. PhD thesis, Harvard University, 1993.
- [119] M. Spivak. *Comprehensive Introduction to Differential Geometry*. Publish Or Perish, 1979.
- [120] L. Staib and J. Duncan. Boundary finding with parametrically deformable models. *IEEE PAMI*, 14(11):1061–1075, 1992.
- [121] D. Stalling, M. Westerhoff, and H.-C. Hege. *Amira: A Highly Interactive System for Visual Data Analysis*, chapter 38, pages 749–767. Elsevier, 2005. URL <http://www.zib.de/visual/publications/sources/amira-overview.pdf>.
- [122] V. Surazhsky, T. Surazhsky, D. Kirsanov, S. J. Gortler, and H. Hoppe. Fast exact and approximate geodesics on meshes. *ACM Trans. Graph.*, 24(3):553–560, 2005.
- [123] T. S. Tang and R. E. Ellis. 2d/3d deformable registration using a hybrid atlas. In *Proc. MICCAI*, volume 3750, pages 223–230, 2005.
- [124] J. F. ten Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 42(2):267–276, 1977.
- [125] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. In *ACM SIGGRAPH*, pages 205–214, 1987.
- [126] P. Tessier. Total facial osteotomy. crouzon’s syndrome, apert’s syndrome: oxycephaly, scaphocephaly, turricephaly. *Ann Chir Plast*, 12(4): 273–286, 1967.
- [127] B. The, R. Diercks, R. Stewart, P. van Ooijen, and J. van Horn. Digital correction of magnification in pelvic x rays for preoperative planning of hip joint replacements: theoretical development and clinical results of a new protocol. *Med Phys*, 32(8):2580–2589, 2005.
- [128] C. Twining, R. Davies, and C. Taylor. Non-parametric surface-based regularisation for building statistical shape models. In *Information Processing in Medical Imaging*, pages 738–750, 2007.
- [129] C. Udriste. *Convex functions and optimization methods on Riemannian manifolds*, volume 297 of *Mathematics and Its Applications*. Kluwer Academic Publishers, 1994.

- [130] M. Vaillant and J. Glaunes. Surface matching via currents. In G. E. Christensen and M. Sonka, editors, *Proc. Information Processing in Medical Imaging*, volume 3565 of *LNCS*, pages 381–392. Springer, 2005.
- [131] E. van de Kraats, G. Penney, D. Tomazevic, T. van Walsum, and W. Niessen. Standardized Evaluation Methodology for 2D-3D Registration. *IEEE Trans Med Imaging*, 24(9):1177–1190, 2005.
- [132] B. van Ginneken, T. Heimann, and M. Styner. 3d segmentation in the clinic: A grand challenge. In T. Heimann, M. Styner, and B. van Ginneken, editors, *Proc. MICCAI Workshop 3D Segmentation in the Clinic: A Grand Challenge*, pages 7–15, 2007. URL <http://mbi.dkfz-heidelberg.de/grand-challenge2007/>.
- [133] J. Veenstra and N. Ahuja. Line drawings of octree-represented objects. *ACM Trans. Graph.*, 7(1):61–75, 1988. ISSN 0730-0301.
- [134] R. C. Veltkamp and M. Hagedoorn. *Principles of Visual Information Retrieval*, chapter State-of-the-art in shape matching, pages 87–119. Springer, 2001.
- [135] A. Vento, R. LaBrie, and J. Mulliken. The O.M.E.N.S. classification of hemifacial microsomia. *The Cleft Palate-Craniofacial Journal*, 28:68–76, 1991.
- [136] Y. Wang, M.-C. Chiang, and P. M. Thompson. Mutual information-based 3d surface matching with applications to face recognition and brain mapping. In *Proc. of the 10th IEEE International Conference on Computer Vision (ICCV'05)*, volume 1, pages 527–534. IEEE Computer Society, 2005.
- [137] M. Wardetzky. *Discrete Differential Operators on Polyhedral Surfaces - Convergence and Approximation*. PhD thesis, Freie Universität Berlin, 2006.
- [138] J. Weese, M. Kaus, C. Lorenz, S. Lobregt, R. Truyen, and V. Pekar. Shape constrained deformable models for 3d medical image segmentation. In M. Insana and R. Leahy, editors, *Proc. IPMI*, volume 2082 of *LNCS*, pages 380–387. Springer-Verlag, 2001.
- [139] J. Weickert, B. M. ter Haar Romeny, and M. A. Viergever. Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Trans. Image Proc.*, 7(3):398–410, March 1998.
- [140] A. Westermarck, S. Zachow, and B. Eppley. 3D osteotomy planning in maxillofacial surgery, including soft tissue prediction. *Journal of Craniofacial Surgery*, 16(1):100–104, 2005.

-
- [141] A. Wittmers. Übertragung von Pfadnetzen auf triangulierten Oberflächen unter Erhaltung der Topologie. Master's thesis, Freie Universität Berlin, 2008. in preparation.
- [142] J. Yedidia, W. Freeman, and Y. Weiss. *Exploring Artificial Intelligence in the New Millennium*, chapter Understanding belief propagation and its generalizations. Science & Technology Books, 2003.
- [143] S. Zachow, A. Trepczynski, E. Gladilin, R. Sader, and H.-F. Zeilhofer. 3D osteotomy planning in cranio-maxillofacial surgery: Experiences and results of surgery planning and volumetric finite-element soft tissue prediction in three clinical cases. In H. Lemke, editor, *CARS 2002: Computer Assisted Radiology and Surgery*, pages 983–987, 2002.
- [144] S. Zachow, H. Lamecker, B. Elsholtz, and M. Stiller. Reconstruction of mandibular dysplasia using a statistical 3d shape model. In H. L. et al., editor, *Proc. Computer Assisted Radiology and Surgery*. Elsevier, 2005. URL <http://www.zib.de/lamecker/publications/cars2005.pdf>.
- [145] M. Zilske, H. Lamecker, and S. Zachow. Adaptive remeshing of non-manifold surfaces. *Computer Graphics Forum*, 27(2), 2008. accepted for publication.
- [146] B. Zitová and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21:977–1000, 2003.