Medial Shape Description of Variable Biological Objects

Martin Andreas Styner

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

> Chapel Hill 2008

> > Approved by: Guido Gerig, Advisor Stephen Pizer, Reader J. S. Marron, Reader Jack Snoeyink, Reader Jeffrey Lieberman, Reader

© 2008 Martin Andreas Styner ALL RIGHTS RESERVED

Abstract Martin Andreas Styner: Medial Shape Description of Variable Biological Objects. (Under the direction of Guido Gerig.)

This dissertation describes a novel shape description scheme that incorporates variability of an object population into the generation of a characteristic 3D shape model. Knowledge about the biological variability of anatomical objects is essential for statistical shape analysis and discrimination between healthy and pathological structures. The proposed shape representation is based on a fine-scale spherical harmonics (SPHARM) description and a coarse-scale m-rep description. The SPHARM description describes the object boundary as a weighted series of spherical harmonics. The correspondence on the boundary is defined by a first-order ellipsoid normalized parameterization. The medial m-rep description is composed of a net of medial primitives with fixed graph properties. A m-rep model is computed automatically from the shape space of a training population of SPHARM objects. Pruned 3D Voronoi skeletons are used to determine a common medial branching topology in a stable way. An intrinsic coordinate system and an implicit correspondence between objects are defined on the medial manifold. My novel representation scheme describes shape and shape changes in a meaningful and intuitive manner. Several experimental studies of shape asymmetry and shape similarity in biological structures demonstrate the power of the new representation to describe global and local form. The clinical importance of shape measurements is shown in the presented applications.

The contributions made in this dissertation include the development of a novel automatic pruning scheme for 3D Voronoi skeletons. My experiments showed that only a small number of skeletal sheets are necessary to describe families of even quite complex objects. This work is also the first to *compute* a common medial branching topology of an object population, which deals with the sensitivity of the branching topology to small shape variations. The sensitivity of the medial descriptions to small boundary perturbations, a fundamental problem of any skeletonization technique, is approached with a new sampling technique.

ACKNOWLEDGMENTS

First of all I would like to thank those who directly contributed to the research in this dissertation. The most important source of ideas and guidance is Guido Gerig, my PhD-advisor. His enthusiasm and positive attitude towards my research was important for this dissertation as well as for my sanity. He was also the one who inspired me in the first place to start the adventure of producing a dissertation. I hope that we will continue to collaborate in the future.

I am also very thankful to the co-students at my former lab in Zürich. Especially important for the scientific aspects of this dissertation are Christian Brechbühler, Markus Näf, Dominique Attali and Andras Kelemen. Their work forms the principal foundation for this dissertation. I am further very thankful to Steve Pizer, who welcomed me very warmly in Chapel Hill. His m-rep description stands as the other portion of this dissertation's foundation. I would like to thank Steve for his insightful comments and discussions about the m-rep and other medial shape descriptions. Sarang Joshi was also helpful especially regarding discussions about the m-rep fitting procedure.

Dr J. Lieberman, Dr. M. Chakos and the UNC image analysis lab at Psychiatry, UNC Hospitals, kindly provided the original MR and segmentations of the hippocampi. Dr. R. Kikinis and Dr. M. Shenton, Brigham and Women's Hospital, Harvard Medical School in Boston, provided the original MR and segmentations of the hippocampusamygdala study. I further acknowledge Dr. D. Weinberger and Dr. D. Jones, NIMH Neuroscience in Bethesda, for providing the twin datasets that were segmented in the UNC image analysis lab. The dissertation research and the used facilities has been funded in part by the NCI grant CA 47982 and the UNC Intel computer grant.

In my non-scientific office life, many people both at my former lab as well as at my current lab helped me a great deal to enjoy my daily work. In the starting years of my PhD it was mainly Christian Stöcklin and Daniel Welti who were grateful victims, boosting my ego by getting 'x-blast'ed after a talk-rich coffee break. Here in Chapel Hill, it was Paul Yushkevich who helped me get through the day by discussing life, universe and everything else during our Franklin-Street-lunch meals.

A major change in environment is never easy while doing a dissertation, but I did it anyway and moved from Zürich to Chapel Hill. Guido, Maya, Neelesh and Shanti Gerig helped me a lot in doing so. Their strong support was very important in a time of change. Stephen Aylward provided his friendship and helped to get me established here. I was here only for 3 months, when I found the woman of my life. Maya is the friend, girlfriend, wife that every student should have in order to survive the sometimes despairing life of a graduate student.

Finally, I thank all my friends and family, especially my mother who was always supportive of the crazy ideas that poured out of my head.

All these people made it possible for me to produce this dissertation. They make me feel that if I had the chance to step back in time and reconsider my decisions, I would again take the path of this dissertation and enjoy it a second time.

Table of Contents

LI	ST C	OF TABLES
LI	ST C	OF FIGURES
LI	ST C	OF SYMBOLS xvi
1	Intr	$ m oduction \ . \ . \ . \ . \ . \ . \ . \ . \ 1$
	1.1	Dissertation contributions
		1.1.1 Claims of this dissertation
	1.2	Guide to the chapters
2	Shaj	pe descriptions: general considerations, correspondence and anal-
	ysis	
	2.1	Shape description via the 3D surface boundary
		2.1.1 Parametric surface description
		2.1.2 Non-parametric surface descriptions
	2.2	Shape description via the medial manifold

	2.4	Shape	analysis	14
	2.5	Selecte	ed properties of shape descriptions	17
		2.5.1	Localization: local versus global	17
		2.5.2	Scale: fine versus coarse	19
		2.5.3	Boundary versus medial shape description	19
		2.5.4	Constraints: unconstrained versus constrained	21
	2.6	Proper	ties of a description suitable for shape analysis $\ldots \ldots \ldots \ldots$	21
3	A c	ombine	ed medial-boundary shape description	23
	3.1	SPHA	RM	23
		3.1.1	Spherical harmonics descriptors	24
		3.1.2	Computation scheme for SPHARM	26
		3.1.3	Parameterization of SPHARM by optimization	27
		3.1.4	SPHARM correspondence	28
		3.1.5	Mean Squared Distance (MSD) between SPHARM objects $\ . \ . \ .$	29
		3.1.6	Point Distribution Model (PDM) from SPHARM	30
	3.2	M-rep		31
		3.2.1	Medial atom properties	32
		3.2.2	Correspondence via m-rep	35

	3.3	Combi	ining SPH	ARM and m-rep	35
		3.3.1	Computi	ng the combined medial-boundary shape description of	
			an indivi	dual object	35
4	Con	nputin	g a comr	non m-rep model of an object population	38
	4.1	Overvi	iew		38
	4.2	The sh	nape space	e for an anatomical object	41
	4.3	Brancl	ning topol	ogy for a common medial model	45
		4.3.1	Branchin	g topology for individual objects	46
			4.3.1.1	General view on Voronoi skeletons	46
			4.3.1.2	Voronoi skeleton generation from SPHARM	49
			4.3.1.3	Grouping Voronoi faces to sheets	52
			4.3.1.4	Pruning Voronoi sheets	55
			4.3.1.5	Pruning by a real contribution to the medial manifold	57
			4.3.1.6	Pruning by volumetric contribution to the whole object .	58
			4.3.1.7	Merging Voronoi sheets	61
		4.3.2	Common	branching topology via common spatial frame	66
			4.3.2.1	Common frame for spatial correspondence between me-	
				dial sheets	67
			4.3.2.2	Identifying corresponding sheets in a common spatial frame	69
			4.3.2.3	Computation of the common branching topology \ldots .	70
	4.4	Comp	utation of	the grid sampling for a common medial model	72
		4.4.1	Grid sam	pling of a single medial sheet	73

		4.4.2	M-rep computation from a sampled medial sheet
		4.4.3	Minimal grid sampling for a single object
		4.4.4	Minimal grid sampling for an object population
	4.5	Fittin	g an m-rep into a SPHARM object
5	Clir	nical aj	oplications
	5.1	Exam	ple 1: Hippocampus structures from a schizophrenia study 86
		5.1.1	First hippocampus case
		5.1.2	Second hippocampus case
	5.2	Exam	ple 2: Hippocampus-amygdala complex from a schizophrenia study 93
	5.3	Exam	ple 3: Lateral ventricles from a mono/dizygotic twin study 98
		5.3.1	Asymmetry and similarity analysis of a monozygotic twin pair 103
		5.3.2	Monozygotic (MZ) vs. dizygotic (DZ) twin pairs
			5.3.2.1 Volume analysis $\ldots \ldots 105$
			5.3.2.2 Shape analysis via SPHARM
			5.3.2.3 Shape analysis via m-rep
6	Dis	cussior	a and conclusions 117
	6.1	Comp	utation time
	6.2	Stabil	ty
	6.3	Homo	logy

	6.4	Conclusions	123
		6.4.1 Scientific contributions of this dissertation	124
	6.5	Future work	127
7	API	PENDIX 1: Properties of the Principal Component Analysis	
	(PC	(A)	129
B	BLI	OGRAPHY List of Tables	133

	List of Figures
	tion scheme
6.1	Estimated computation time for all stages in the m-rep model computa-
5.3	MZ/DZ class statistics of the m-rep properties, no scaling 106
5.2	MZ/DZ class statistics of the m-rep properties, volume scaling $\ . \ . \ . \ . \ 106$
5.1	Error ranges for E_i and E_{pop} for left and right m-rep model
3.2	Summary of properties of SPHARM and m-rep
3.1	Explicit expressions of the spherical harmonics up to degree 3 25

2.1	Four different shape descriptions of a human left hippocampus	•	•	•	•	•	•	18
2.2	Lateral ventricles of two monozygotic twins							20

3.1	The SPHARM shape description of a human left hippocampus/amygdala $% \mathcal{A} = \mathcal{A} = \mathcal{A} + $	
	complex shown at 4 different degrees	24
3.2	Visualization of SPHARM correspondence	28
3.3	Icosahedron subdivision for different levels of sampling	30
3.4	The m-rep shape representation without and with implied boundary in	
	an example of a human left hippocampus/amygdala complex	32
3.5	Two types of medial atoms: internal atom - end atom	33
3.6	Computation of the combined description for a study population of objects	37
4.1	Three steps to automatically compute a m-rep model in a stable way. $\ .$.	39
4.2	The proposed shape space is spanned by the first few eigenmodes of de-	
	formation	40
4.3	Subset of the object set from a PCA shape space for a hippocampus-	
	amygdala population	41
4.4	Two examples of the (v_1, v_2) PCA space from synthetic datasets with	
	non-Gaussian distributions.	43
4.5	Three examples of the (v_1, v_2) PCA space from real datasets	43
4.6	Extraction of a common branching topology	44
4.7	Icosahedron subdivision of a human hippocampus	48
4.8	Preprocessing of Voronoi skeletons	49
4.9	Failed preprocessing of Voronoi skeletons	51

4.10	Progression of the initial grouping algorithm	52
4.11	Creation of sheet-internal branches in initial grouping algorithm \ldots .	53
4.12	Final pruning step that prunes sheets consisting only of one Voronoi face	56
4.13	Application of the pruning scheme to a human left hippocampus-amygdala	58
4.14	Steps of my pruning scheme	59
4.15	Pruning scheme applied to the average hippocampus-amygdala objects of the left and right brain hemisphere	60
4.16	Pruning scheme applied to a lateral ventricle pair	61
4.17	Application of the pruning scheme to various brain structures	63
4.18	Pruning scheme applied to the object set of a left hippocampus-amygdala	
	population	64
4.19	Topology matching scheme applied to the object set	65
4.20	Schematic overview of matching procedure	67
4.21	Schematic algorithm for computing the common branching topology	70
4.22	Computation of common m-rep model with minimal grid dimensions	71
4.23	Computing the sampled medial sheet axis	73
4.24	Visualization of the sampling method	74
4.25	M-rep approximation errors in the initial position and in the final position	
	after deformation	78

4.26	Different stages of the m-rep fit for an individual object o_i	81
5.1	Visualization of the first hippocampus case	85
5.2	Visualization of thickness asymmetry along the medial axis for first case .	86
5.3	Visualization of thickness asymmetry in the first case	87
5.4	Visualization of thickness asymmetry in the first case	87
5.5	Visualization of second hippocampus case	88
5.6	Visualization of thickness asymmetry in the second case	89
5.7	M-rep description of left hippocampus-amygdala population $\ldots \ldots \ldots$	92
5.8	Application of the scheme to the populations of the left and right hippocampu amygdala	ıs- 93
5.9	Three-dimensional rendering of the skin surface (transparent) and the lateral ventricles.	95
5.10	Visualization of the lateral ventricles of all twin pairs	97
5.11	Approximation error for the object set members along the first two eigen- modes	98
5.12	Two examples of fitting the common m-rep model to the study population	98
5.13	Visualization of the lateral ventricles of two MZ twins and its correspond- ing m-reps	99
5.14	Asymmetry and similarity analysis of the thickness: difference of thickness at medial atoms	100

5.15	Volume difference between lateral ventricles of twin pairs
5.16	Statistics for \sqrt{MSD} shape similarity between lateral ventricles of twin
	pairs
5.17	Class statistics of the pure m-rep properties for objects individually scaled
	to their volume
5.18	Class statistics of the combined m-rep properties for objects individually
	scaled to their volume
5.19	Class statistics of the pure m-rep properties for unscaled objects 109
5.20	Class statistics of the combined m-rep properties for $\mathit{unscaled}$ objects $% \mathcal{O}(\mathcal{O})$. 110
5.21	Locations of significant difference between MZ and DZ twins
6.1	Common branching topology for two different orderings of the matching
	procedure
6.2	Schematic 2D visualization of the homology for the medial description
	and its implied boundary in two cases
7.1	Principal component analysis as basis rotation

LIST OF SYMBOLS

1D	One dimensional
2D	Two dimensional
3D	Three dimensional
\mathbb{R}	Set of all real numbers
\mathbb{R}^+	Set of all positive real numbers
SO(3)	Group of rotations in 3D
Y_l^m	Spherical harmonic of degree l and order m , introduced section 3.1.1
\mathbf{c}_i	SPHARM coefficient vector for the i -th object, introduced section 4.2
$\overline{\mathbf{c}}$	SPHARM coefficient vector for the average object, introduced section 4.2
$m_{i,j}$	M-rep atom at grid index i, j , introduced section 3.2
x	position in 3D, introduced section 3.2
r	local thickness, introduced section 3.2
λ_i	i-th PCA eigenvalue, introduced section 4.2
v_i	i-th PCA eigenvector, introduced section 4.2
MSD	Mean Squared Distance, introduced section 3.1.5
MAD	Mean Absolute Distance, introduced section 4.4
MAD_{norm}	Normalized Mean Absolute Distance, introduced section 4.4
r_{avg}	Average thickness of the i -th object, introduced section 4.4
$r_{avg,pop}$	Average thickness in the whole population, introduced section 4.4
E_i	Normalized approximation error using r_{avg} , introduced section 4.4
E_{pop}	Normalized approximation error using $r_{avg,pop}$, introduced section 4.4

Chapter 1

Introduction

In many aspects of modern medicine, medical imaging provides crucial information for diagnosis, treatment and disease related investigations. Many diseases are studied using images of Magnetic Resonance Imaging (MRI), Computed Tomography (CT) or ultrasound imaging. Quantitative studies of anatomical objects are of high importance to investigators of neurologic diseases. Many quantitative studies rely on the statistical analysis of volume measurements |48, 57, 31, 2|. The goal therein is to detect volumetric changes between patients and healthy controls. The idea to extend the analysis to incorporate shape features has been longstanding, but it was technically not feasible. To date only a small number of shape analyses have been performed due to the complexity of this task. A quantitative shape analysis would be of special interest in neurologic diseases presenting morphological changes in brain structures like schizophrenia, autism, epilepsy or Alzheimer's disease. The brain anatomy is thought to change with the progression of these diseases. The successful detection of brain shape changes could be used for diagnostic detection, for early intervention and for an improved understanding of the disease. It is thus important for a shape analysis method not only to demonstrate that there are shape differences, but also where and how these differences manifest themselves.

Why do we need more than volume measurements? It is clear that the shape information captures additional information not covered by volume measurements; yet it is not obvious that this additional information is relevant. In section 5.3, I present an example in which I was able to detect a ventricular shape difference between populations of monozygotic and dizygotic twins, whereas these populations fail to show significant differences in their volume measurements. The results of the shape analysis revealed that monozygotic twins have more similarly shaped lateral ventricles than do dizygotic twins or non-related subjects. This result might be of clinical importance, for example in the analysis of discordant monozygotic twin studies. In discordant twins the studied disease is manifested in only one of the twins while the other twin is healthy.

While humans have no apparent difficulty describing shape features and capturing shape differences of 3D objects, researchers in computer vision face the usual problem: humans *can* do it, but nobody knows exactly *how*. The representation and the analysis of objects has proven to be a challenging problem. No general shape description exists that solves all shape related tasks. Shape descriptions are thus tailored for specific tasks. In this dissertation, I will present a shape description scheme that is especially suitable for shape analysis. A description that is suitable for shape analysis might not be appropriate for doing object visualization and vice versa.

In my shape description scheme, a shape analysis should be able to capture coarse and fine scale shape differences. These shape differences should be intuitively captured. In medial shape descriptions, the object is represented by a skeletal graph that is topologically equivalent to the original object. Medial descriptions have the advantage over many other shape descriptions that the medial features capture shape properties in a more intuitive and meaningful manner. However, medial descriptions are sensitive to small changes on the boundary. This makes them suited to describe shape at a coarse scale and less well suited at fine scale. Many other researchers proposed to describe fine scale shape properties using boundary descriptions. In this dissertation, I develop a scheme that uses the boundary descriptions of a population of objects to compute a stable coarse-scale medial description by incorporating this population's shape variability. This medial description is an m-rep description expressing shape differences as changes of intuitive local shape features. To compute a stable m-rep description, a fine-scale medial description called Voronoi skeletons is used as an intermediate representation. Voronoi skeletons are sensitive to boundary perturbations like all fine-scale medial descriptions. Thus, a regularization step is needed to remove those parts of the skeleton that are irrelevant to the object's shape. This regularization is called pruning. A lot of research has been done on 3D Voronoi skeletons, but to my knowledge an automatic pruning scheme has not formerly been developed.

The difficult problem of correspondence has to be solved as a part of a shape analysis method. Correspondence defines homologous points between different objects, enabling the computation of shape comparisons. I propose to solve the correspondence problem on the boundary using a parametrized description, in which the same parameter values across different objects results in corresponding points. This parametrized boundary description is called SPHARM. It represents the object as a weighted series of spherical harmonic basis functions. The correspondence on the boundary can then be propagated to the medial manifold of the object.

1.1 Dissertation contributions

In this dissertation, I present a novel framework for building models to be used for shape analysis. The framework is based on a combination of the medial m-rep and the boundary SPHARM description. A novel method computes a common m-rep model automatically based on a population of similar objects in a stable fashion. The shape variability of this population is thereby incorporated into the m-rep model. The shape analysis of these models yields localized shape changes that can be intuitively understood.

1.1.1 Claims of this dissertation

- #1: A new shape description scheme based on a combined m-rep and SPHARM description is developed. It incorporates prior statistical knowledge about the object variability. The shape description scheme is suitable for shape analysis and thereby allows new insights and paths of exploration in various fields of morphological research.
- #2: A common 3D medial branching topology can be computed for a population of objects in a stable way. This common medial branching topology is based on Voronoi skeletons and is necessary to deal with the sensitivity of the branching topology to small shape variations.
- #3: A novel, general scheme is developed that automatically prunes Voronoi skeletons of 3D objects.
- #4: Only a small number of skeletal sheets are necessary to describe the branching topology of complex anatomical brain objects. The complexity of the common branching topology is of the same magnitude as the individual branching topologies.
- #5: The sensitivity of medial representations to small boundary perturbations can be dealt with using a novel medial sampling technique and refinement by m-rep deformation.
- #6: The shape of an anatomical object captures clinical information that is superior to volume measurements. My new shape description scheme can provide global and local measurements. Additionally, disease related effects can be measured by shape analysis that cannot otherwise be measured with volume measurements.
- #7: Shape analysis demonstrates that the lateral ventricles of monozygotic twins are significantly more similarly shaped than those of dizygotic twins. The shape anal-

ysis also yields the locations of significant shape difference between the two twin populations.

1.2 Guide to the chapters

This dissertation is organized in 6 Chapters, followed by an appendix and references. Chapter 2 is dedicated to a brief review of shape representations, their properties and other related work. Chapter 3 discusses the novel shape description scheme. The scheme to compute a common m-rep model for an object population is described in chapter 4. Chapter 5 presents three applications of the combined shape description. Chapter 6 discusses issues of stability and homology of the developed scheme. In this chapter, I also summarize this work and present some possible future extensions and applications. The appendix discusses the mathematical properties of the Principal Component Analysis.

Chapter 2

Shape descriptions: general considerations, correspondence and analysis

This chapter provides a summary of 3D shape descriptions using either the object boundary or the medial surface. Shape description necessitates facing the issue of correspondence, a subject discussed in section 2.3. Next, an overview of related work in shape analysis is presented. In order to derive properties that are needed for a shape description suited for shape analysis, selected properties of shape descriptions are investigated in the following section. This investigation leads to a list of properties that outlines the requirements for my proposed shape description scheme.

2.1 Shape description via the 3D surface boundary

2.1.1 Parametric surface description

In this dissertation, I focus on objects of spherical topology. For such objects, a parametric surface description is a mapping of a two-dimensional (u, v) parameter space of a 2D manifold to the 3D-Euclidean space. This section describes approaches using this type of representation.

Superquadrics: One popular approach to parametric surface description is su-

perquadrics, as described by Bajcsy and Solina[8]. Bajcsy and Solina describe implicit definitions of ellipsoids and toroids. Only the superquadric ellipsoids are of interest here as toroids are not suited for the representation of most biological objects. Superquadric ellipsoids are defined by

$$\mathbf{v}(a, a_1, a_2, a_3, \varepsilon_1, \varepsilon_2, u, v) = a \begin{pmatrix} a_1 C_u^{\varepsilon_1} C_v^{\varepsilon_2} \\ a_2 C_u^{\varepsilon_1} S_v^{\varepsilon_2} \\ a_3 S_u^{\varepsilon_1} \end{pmatrix} \quad -\pi \le v < \pi$$

$$(2.1)$$

where $S_w^{\varepsilon} = \operatorname{sgn}(\sin w) |\sin w|^{\varepsilon}$ and $C_w^{\varepsilon} = \operatorname{sgn}(\cos w) |\cos w|^{\varepsilon}$. Further, $a, a_1, a_2, a_3 \ge 0$ are scale parameters also defining the aspect ratios and $\varepsilon_1, \varepsilon_2 \ge 0$ are "squareness" parameters. Superquadrics are not able to describe complex objects and are thus rarely used for shape analysis of biological objects.

Fourier descriptors: Surfaces in 3D can be represented by a series expansion of parametric coordinate functions in 2D parameter space $\mathbf{X} = (x(u, v), y(u, v), z(u, v))$, where u and v vary over the surface. Along with spherical harmonics, discussed in section 3.1, Fourier basis functions are a possible choice for the series expansion. Staib and Duncan [65, 66] used Fourier representations for surface finding in image volumes. Each coordinate function $v_1(\theta, \phi)$, $v_2(\theta, \phi)$ and $v_3(\theta, \phi)$ is represented by

$$v(\theta,\phi) = \sum_{m=0}^{2K} \sum_{l=0}^{2K} \lambda_{m,l} (a_{m,l} \cdot \cos 2\pi m\theta \cdot \cos 2\pi l\phi + b_{m,l} \cdot \sin 2\pi m\theta \cdot \cos 2\pi l\phi + c_{m,l} \cdot \cos 2\pi m\theta \cdot \sin 2\pi l\phi + d_{m,l} \cdot \sin 2\pi m\theta \cdot \sin 2\pi l\phi),$$

$$(2.2)$$

where

$$\lambda_{m,l} = \begin{cases} 1 \text{ for } m = 0, \ l = 0 \\ 2 \text{ for } m > 0, \ l = 0 \text{ or } m = 0, \ l > 0 \\ 4 \text{ for } m > 0, \ l > 0 \end{cases}$$

and $\theta, \phi \in [0; 2\pi)$. Cutting off the decomposition at different values for K results in different levels of detail. The Fourier representation is not limited to closed surfaces. It can also represent open surfaces, tube surfaces and torus surfaces. The drawback of the representation of closed surfaces is that such surfaces are treated as tubes whose ends close up to a point. Thus, one of the axes is straight and the axes are treated unequally, leading to sensitivity to small correspondence errors. This makes this description less well suited for shape analysis. When dealing with surfaces of complex objects, the surface parameterization also poses other problems, as discussed in later in section 3.1.3.

Wavelets: Another possible set of basis functions are wavelets [47]. In contrast to Fourier basis functions, wavelets have compact support in the frequency and in the spatial domain. The wavelet transform of the input signal is computed using a filter bank that splits a signal into subsampled low pass and high pass bands. This procedure is iteratively repeated for the low pass band. Wavelet based descriptions have not been used for shape analysis since the correspondence problem remains difficult to solve in a wavelet based description.

Spherical harmonics SPHARM: Spherical harmonics are another basis function for a parametric surface description. This description, called SPHARM, is described in section 3.1, as it is part of my shape description scheme.

2.1.2 Non-parametric surface descriptions

A large class of surface descriptions are non-parametric and describe the surface via a series of primitives located on the boundary. The primitives most often used are points, triangles, quadrilateral meshes and simplex meshes.

Surface points, Point Distribution Models (PDM): Several researchers have used boundary points and their distribution (Point Distribution Model = PDM) for shape description and analysis, e.g. Bookstein in 2D [16, 17], Rangarajan [55] and Cootes [26] in 3D. These approaches rely on an appropriate sampling of the object boundary. The work of Cootes and Taylor has shown that even for a large number of points a statistical shape analysis can be done. PDM's have been used in shape analysis as is discussed in section 2.4.

Triangulation and meshes: The description of surfaces via triangulation or quadrilateral meshes is the main technique used in computer graphics. These descriptions are rarely used for shape analysis. Delingette [30] introduced simplex meshes, which are topologically dual to a triangulation, as a possible shape description. The mesh is adaptive and can change density and topology. The main advantage of the simplex mesh description is that missing data can be interpolated using smoothness, density and geometric constraints of the mesh. Approaches for volumetric and shape measurement of the simplex mesh have been developed. Since the correspondence problem remains unsolved due to the adaptive nature of the mesh, no shape comparative analysis has been done so far.

2.2 Shape description via the medial manifold

One of the most extensively studied shape description in 2D is the medial axis transform (MAT) originally proposed by Blum [11]. Blum claims that medial descriptions are based on the idea of a biological growth model and a 'natural geometry for biological shape.' The idea is to represent the object by a fully connected skeletal graph. The medial axis in 2D captures shape intuitively and can be related to human vision (see Burbeck [22] and Siddiqi [62]). The terms *prairie fire* transform, *medial axis* transform, *symmetry axis* transform, and *skeleton* transform have been used in the literature almost interchangeably and refer to the same basic shape description concept.

The purpose of the medial axis transform is to extract a skeletal figure from the object. The formation of this skeleton can be explained with the prairie fire analogy. Let the object be composed of flammable dry grass, and initiate a fire simultaneously over the

whole boundary of the object. This fire will propagate towards the center of the object. At some points, called quench points, the fire fronts will meet and extinguish themselves. The skeleton of the object is defined as the connected collection of these quench points. If the distance to the original boundary is recorded at every quench point, then the object can be fully reconstructed from the skeleton and thus no information is lost via the medial axis transform. Fire front propagations have been computed using distance transform methods [18, 58], by application of the shrinking/thinning operation [45, 75, 76] or by solving a curve evolution equation system (see below).

The disadvantage of the medial axis transform is its sensitivity to small noise on the object boundary. Boundary noise of small amplitude might produce a quite large skeletal change. August investigated these skeletal changes [6, 7]. A common practice to deal with the boundary noise sensitivity is to smooth the boundary prior to skeleton generation. August shows that even smoothing itself can introduce new skeletal branches. He also showed that the changes in the branching topology are located in regions of ligature, which is a term introduced by Blum to describe the locations on the skeleton influenced by concave boundary sections. In general, the branching topology is unstable. This also means that a graph representation based on skeletal *branching nodes* is unstable and less well suited for shape analysis.

Zucker and Kimia [43, 62] proposed smoothing via a diffusion equation. The skeleton is thereby generated via the solution of a curve evolution equation having two terms of which one is related to a geometric heat diffusion (parabolic α - term) and the other to a reaction (hyperbolic β -term):

$$\begin{cases} C_t = (\beta - \alpha \kappa)N \\ C(s, 0) = C_0(s) \end{cases}$$

Other curve evolution based skeletonization methods were proposed by Siddiqi [63, 61] using a hyperbolic evolution with a nonlinear Gaussian smoothing term. He designed a hierarchical graph description [59], called shock graphs, for shape analysis. These shock graphs are based on the skeletal branching nodes and local geometric measures at these nodes. A measure was defined to quantify the similarity between different objects. The skeletonization was only recently extended to 3D, but the high complexity and the high degree of ambiguity of the 3D skeletal graph has limited the extension of the graph formation stage approach.

The representation of the skeleton via its Voronoi graph has been pursued by several researchers. The process of smoothing the boundary in front-propagation-based skeletons is analogous to the pruning process in Voronoi skeletons. In section 4.3.1.1, I discuss the Voronoi skeleton generation and the pruning mechanisms in detail.

The idea of a fixed branching topology for the medial description of similar objects was evaluated by Golland [38, 37] and Pizer [53]. Imposing a *fixed* branching topology on the medial description solves the problem of boundary noise sensitivity. The question arises of how well a fixed topology represents individual objects. This subject will be discussed in more detail later in this document because it is essential to the applicability of the proposed shape description scheme. Golland fits a 2D skeleton with given, fixed topology into an object's distance transform in a snake-like fashion. This approach cannot be extended straightforwardly to 3D, neither has it been shown to handle branching skeletons. Pizer takes a multi-scale viewpoint. He proposes the m-rep description, which fits a medial model via its implied boundary to the object boundary with given aperture. This can be done with a set of apertures to create a more robust multi-scale medial description. The m-rep description is described in more detail in section 3.2.

2.3 Correspondence

Correspondence, which defines the homology of points between different objects, is extremely important in order to do comparisons and generate statistics. However, there is no agreement on what correspondence exactly should be and how it can be measured.

Correspondence can be established discretely or continuously. In the discrete case the objects are represented as sets of primitives and the correspondence is defined by uniquely assigning primitives in different sets to each other. In the continuous case parameterizations of the surfaces are defined such that same parameter values parameterize corresponding locations. The correspondence defined by the SPHARM description is continuous, while the correspondence on the m-rep is discrete.

The straightforward method to solve the 2D discrete correspondence is the manual selection of a number of landmarks on each object (see [15], [14]). This method clearly requires extensive user input. Also, the extension to 3D is difficult because a greater number of points has to be considered and identifying landmarks in 3D is difficult for a human. Non-manual methods have been proposed, such as the Softassign Procrustes matching algorithm [55], which tackles the problem of finding correspondences in two point sets and identifying outliers. A correspondence match matrix and the associated Procrustes distance are optimized iteratively until convergence is obtained. Another approach to solve the 3D discrete correspondence is the iterative closest point (ICP) algorithm [35]. The ICP algorithm assigns iteratively the closest points of 2 point sets to match. The definition of the distance between 2 points can be extended to include such terms as curvature, or local image properties, as shown by Caunce [23]. The method for establishing discrete correspondence in m-reps is straightforward and is explained in section 3.2.2.

The most common approach for the approximation of continuous correspondences in 2D is an arc-length parameterization. The sample curves are parameterized by the arc-length in a given interval. Points with the same parameter on different curves are then taken to be corresponding. A different approach is taken by Tagare [73, 74], who defines the correspondence between closed curves C_1 and C_2 as a subset of the product space $C_1 \times C_2$. Postulating that correspondence Φ must be a bi-morphism, he derives that Φ is a regular curve in $C_1 \times C_2$ such that its projections on C_1 and C_2 are non-decreasing with respect to arc-length. Taking this as a constraint he obtains the correspondence by minimizing a shape dissimilarity function based on the difference of the curve normal derivatives. Kotcheff [44] presents a different algorithm that finds a correspondence via maximizing the amount of the variability described by the first few PCA components of the shape vector covariance matrix. The high dimensional optimization is tackled via a genetic algorithm.

The problem of 3D continuous correspondence is hard to tackle since most of the continuous methods in 2D and discrete methods proved difficult to extend. One of the few existing solutions is the straightforward expansion of the arc-length parameterization correspondence to surfaces described by Brechbühler [21]. This is the correspondence used in this dissertation between objects described by SPHARM. Section 3.1.4 describes this method in more detail.

2.4 Shape analysis

This section gives a brief overview of shape analysis techniques. A more detailed overview was published by Loncaric [46].

As early as the first decade of the last century researchers were starting to investigate quantitative shape analysis. D'Arcy Thompson [77], an early 20th century morphologist, pioneered the method of transformation grids. His goal was not to describe shape qualitatively or quantitatively but rather to formulate a method that could measure shape change. That is, given two objects, he was primarily interested in comparing their differences. Thompson's work was novel for its analysis of complex biological processes from a mathematical and physical viewpoint. The last chapter in his groundbreaking book [77], called "On the Theory of Transformations, or the Comparison of Related Forms", is probably the best known and has direct bearing on all image warping methods.

It was Bookstein [13, 16] who developed methods to analyze shape changes quantitatively based on principal warps. This principal warp analysis is basically a Principal Component Analysis (PCA) of the thin-plate spline deformation matrix between two point sets. These point sets are first registered using the Procrustes method. The main problem with this method is the sensitivity to the choice of landmarks, which are assumed to be without error. Dryden [32] established ways to deal with this using S-estimators.

Davatzikos et al [29, 28] proposed morphometry via a spatially normalizing elastic transformation. Inter-subject comparisons were made by comparing the individual transformations. The method is applied in 2D to a population of corpora callosa. A similar approach in 3D has been chosen by Joshi [40] and Miller [49, 50] to compare hippocampi. Using the viscous fluid transformation proposed by Miller [24], inter-subject comparisons were made by analyzing transformation fields. The *analysis* of transformation fields in these two methods has to cope with the high dimensionality of the transformation and the sensitivity to the initial position. Although the number of subjects in the studied populations is low, both show a relatively stable extraction of shape changes. Csernansky [27] used these methods in his widely recognized article about hippocampal shape differences in schizophrenia. Toga [52] was using a similar method in analyzing brain growth processes and brain diseases. Toga also published a good overview [78] of shape analysis methods that study deformation fields.

Cootes et al [25] proposed the Active Shape Model to generate statistical models of 2D and 3D point sets. A statistical model is built after applying the Procrustes algorithm to the PDM. The PDM and its associated PCA is then used for segmentation and shape analysis. One of the main problems in dealing with 3D PDM's is the identification of corresponding surface points.

Inspired by the success of Cootes Active Shape Models, Kelemen [42] adapted this

scheme to be used with the SPHARM description. The statistical description of SPHARM is used in this dissertation and is described in detail in section 4.2.

The INRIA group developed a feature based concept using crest lines to represent shape variability and to build statistical atlases (Subsol [70]). The approach results in a point to point correspondence and uses a space spline transformation to warp 3D image data.

The use of medial descriptions for shape analysis has been proposed by Golland [38] and Pizer [53]. Both Pizer and Golland propose a sampled medial model that is fitted to individual objects. By holding the topology of the model fixed, an implicit correspondence between objects is given. Yushkevich [79] used 2D m-reps to find local changes in populations effectively. M-reps and its shape properties are discussed in more detail in section 3.2.

Kimia and Giblin [36] have proposed a medial hypergraph in 3D. They showed that the hypergraph completely characterizes the shape of an object. Similar to work in 2D by Siddiqi et al [60], this hypergraph could be used for object recognition and design. To our knowledge, no studies have been done towards using the medial graph/hypergraph directly for shape analysis.

In this dissertation I present a new approach to shape analysis using both boundary and medial description [68, 69]. The medial description is computed automatically from a shape space based on a given population.

2.5 Selected properties of shape descriptions

This section studies a selection of general shape description properties that can be used to broadly categorize most shape description methods. The objects of interest are biological objects that were segmented by selecting a 'region of interest' from volumetric medical images resulting in a binary segmentation. Fig. 2.1 displays the properties of the shape descriptions that are most important in this dissertation. In the following sections, the effect of the presence of biological variability is emphasized.

Definition: Efficiency of a shape description The efficiency of a shape description is defined as the ratio of non-redundant information to the size of the description. A shape description is called *efficient* if objects are described with a given accuracy by *concise* sets of parameters or features. A finely sampled description is thus less efficient than a coarser sampled description if both describe the same object with the same accuracy.

Definition: Local/global shape description The terms local and global shape description are used in this dissertation to express whether a representation captures the object as a set of primitives with locality or as a set of parameters without locality.

2.5.1 Localization: local versus global

Each global 3D shape descriptions in this chapter is based on a 2D-parameterization (u, v), like Staib and Duncan's sinusoids [66] and Brechbühler's spherical harmonics [21] (SPHARM). A specific object is described by a set of coefficients weighting a given set of basis functions. Shape properties like derivatives can be computed analytically from the functional parameterization. Deformations are not well localized in a global description but are rather distributed over the whole set of coefficients. Thus, changes of the coefficients cannot be interpreted intuitively. Moreover, local shape changes can

lead to changes of the whole set of coefficients.



Figure 2.1: Four different shape descriptions of a human left hippocampus: A. spherical harmonics (SPHARM): global, fine scale, boundary. B. Dense Point Distribution Model (PDM): Local, Fine Scale, Boundary. C. Dense Voronoi Skeleton: local, fine scale, medial (color coded thickness). D. m-rep (color/radius coded thickness, atoms = dots, links = purple, implied boundary = blue) : local, coarse scale, medial.

Local shape descriptions are composed of primitives, such as points, edges or faces. Changes to the primitives are captured locally and can be visualized and understood intuitively. Shape properties, such as derivatives or curvatures, are defined by the relationships between these primitives. A dense sampling is required to describe these relationship and thus also the shape properties accurately.

A finely sampled local shape description is not as efficient as a global description since it describes the same object by a larger number of parameters. Sparse sampling can be used to achieve a more efficient local description, if we are not interested in fine scale shape properties or if we can determine them by additional prior information.

As I deal with biological objects, I aim to pinpoint deformations intuitively as changes of anatomical landmarks, and this clearly favors a local description. Global descriptions are favored by the need for an efficient and accurate computation of geometric shape properties, which are used for registration and establishing correspondence between homologous points of different objects.

2.5.2 Scale: fine versus coarse

In medical image analysis studies, 3D objects are most often defined a binary segmentations of 'regions of interest' in volumetric images. In the present routine analysis, such anatomical objects are segmented based on human expert interaction. These segmentations are often processed as if free of error. Because fine scale descriptions reconstruct the object accurately, they are perceived to be anatomically correct. However, the presence of noise, partial volume effects, intensity inhomogeneities and other artifacts suggests that the view of an error-free object is misleading. A fine scale description is not efficient since it reconstructs the objects to an unnecessarily high degree of precision based on the non-accurate manual segmentation. Statistical shape analysis, detection and discrimination of shape changes, favors an efficient description. On the other hand, we would like to be able to precisely pinpoint significant shape changes, requiring a representation of anatomical details. Thus, the choice of scale can be interpreted as balancing the tradeoff between the efficiency of description and precise localization.

2.5.3 Boundary versus medial shape description

The 3D shape of an object can be described by its boundary or by the medial manifold. Main advantages of medial descriptions are the intuitive capturing of shape information (see Fig. 2.2) and the separate characterization of the local shape properties: location, orientation and thickness. The local orientation of a medial description captures first order properties of the medial manifold, which are directly related to those on the boundary. The main disadvantages of medial descriptions include their inability to capture non-symmetric information and the sensitivity to small changes on the boundary (see also section 2.2). Because the non-symmetric part of shape can be regarded as being less stable, some researchers view this property as an advantage rather than a disadvantage. Considering the presence of biological variability, the sensitivity to small changes cannot be left unsolved. In particular, a statistical analysis of a set of medial manifolds based on similar biological objects would be very challenging if the branching topology were not the same for all objects.

I propose that boundary descriptions are well suited for fine scale and medial descriptions are well suited for coarse scale descriptions (see also Pizer [53]). If we can solve the problem of the branching topology sensitivity, the medial description could also be well suited for statistical shape analysis.

Twin A Twin B

Figure 2.2: Lateral ventricles of two monozygotic twins. The objects are similar, but twin A has a larger right ventricle (Volume L/R = 0.75). Twin B shows a reversed volumetric symmetry (Volume L/R = 1.26). The medial description (bottom, Voronoi skeletons), with color coded thickness, captures more intuitively the local shape structure of the object than the boundary representations (top).

2.5.4 Constraints: unconstrained versus constrained

Most shape descriptions can be constrained by incorporating prior knowledge. For example model based segmentation techniques generate descriptions that are constrained, if the segmentation does not allow free-form deformations. The segmentation is guided by constraints on the object's geometric and intensity distribution. In general, an unconstrained description describes the object as precisely as possible, whereas a constrained description might deviate from this 'best-precision' description. If the constraints are of statistical nature that incorporate knowledge about the object's variability, then the generated description is usually more robust. Constraining shape descriptions can also help to overcome some of the description's inherent disadvantages, for example constraining the branching topology for medial descriptions. Pathological cases, when constraints disallow an appropriate shape description, are usually detectable.

2.6 Properties of a description suitable for shape analysis

As a consequence of the discussion in the previous sections of this chapter, the following properties are desired in a shape description suitable for shape analysis :

- An efficient description is needed to reduce the dimensionality of the feature space.
- The shape features and their changes should be precisely localized.
- The shape features and their changes should be meaningful and invoke an intuitive understanding of the local and global form.
- The description should be stable in the presence of shape variability since I deal with populations of similar objects.
- The description should offer means to establish an appropriate correspondence.

The different properties cannot be achieved simultaneously in a single description. A medial description for example offers meaningful features, but is quite unstable in the presence of shape variability. A sampled description is only efficient if the sampling is low, but the resulting coarse scale description does not guarantee anatomical correctness. A parametrized description is often efficient even at a fine scale, but the features and especially their changes are not intuitive. I thus propose to have two descriptions, each one incorporating some of the desired properties. By combining two descriptions, some of the inherent disadvantages can be overcome and all requirements can be met in the combined description. Such a combined description is presented in the next chapter.
Chapter 3

A combined medial-boundary shape description

I propose a novel shape description that combines the boundary-based spherical harmonic description (SPHARM) [20, 21] and the m-rep description of a net of medial primitives [53]. This combination is especially well suited for doing statistical shape analysis. It pools the advantages of the individual descriptions and overcomes some of their inherent disadvantages. Both coarse and fine scale features are captured efficiently. Correspondence is defined on both descriptions and thus a shape analysis can directly be applied. Shape changes can be intuitively described using the m-rep description. The m-rep model can be determined stably by incorporating variability using the SPHARM description.

In this chapter, I first describe the SPHARM and m-rep descriptions individually. Then, I describe how they are combined.

3.1 SPHARM

The SPHARM description is a hierarchical, global, multi-scale boundary description that can only represent objects of spherical topology. The basis functions of the parameterized surface are spherical harmonics. Kelemen [42] demonstrated that SPHARM can be used to express shape deformations. Truncating the spherical harmonic series at different degrees results in object representations at different levels of detail, as it shown in Fig. 3.1. SPHARM is a smooth, accurate fine-scale shape representation, given a sufficiently small approximation error.

In the next sections, I briefly describe the mathematical properties of spherical harmonic descriptors, the computation scheme for a SPHARM description, and the parameterization computation. Also, I discuss how to establish correspondence between different objects described by SPHARM. This correspondence is used to compute the Mean Squared Distance measures of shape difference and to generate a spherically uniform point sampling.



Figure 3.1: The SPHARM shape description of a human left hippocampus/amygdala complex shown at 4 different degrees (1, 3, 6, 10 harmonics).

3.1.1 Spherical harmonics descriptors

This section discusses the mathematical properties of the spherical harmonic basis functions. It gives a summary of spherical harmonic descriptors as they are presented in Brechbühler's dissertation [20].

Spherical harmonic basis functions Y_l^m , $-l \leq m \leq l$ of degree l and order m are

defined on $\theta \in [0; \pi] \times \phi \in [0; 2\pi)$ by the following definitions [54]:

$$Y_{l}^{m}(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_{l}^{m}(\cos\theta) e^{im\phi}$$
(3.1)

$$Y_l^{-m}(\theta,\phi) = (-1)^m Y_l^{m*}(\theta,\phi) , \qquad (3.2)$$

where Y_l^{m*} denotes the complex conjugate of Y_l^m and P_l^m the associated Legendre polynomials

$$P_l^m(w) = \frac{(-1)^m}{2^l l!} (1 - w^2)^{\frac{m}{2}} \frac{d^{m+l}}{dw^{m+l}} (w^2 - 1)^l.$$
(3.3)

Table 3.1 lists explicit expressions for the spherical harmonic functions up to degree 3. The Cartesian notion reveals that the spherical harmonics are polynomials in the 3D space (u_0, u_1, u_2) .

l	m = 0	m = 1	m = 2	m = 3
0	1			
1	$\cos heta$	$e^{i\phi}\sin heta$		polar
2	$-1+3\cos^2\theta$	$e^{i\phi}\cos heta\sin heta$	$e^{2i\phi}\sin^2\theta$	
3	$-3\cos\theta + 5\cos^3\theta$	$e^{i\phi}(1-5\cos^2\theta)\sin\theta$	$e^{2i\phi}\cos\theta\sin^2\!\theta$	$e^{3i\phi}\sin^3\theta$
0	1			
1	u_2	$u_0 + i u_1$		Cartesian
2	$-1+3 u_2^2$	$(u_0 + i u_1) u_2$	$(u_0 + i u_1)^2$	
3	$-3u_2+5u_2^3$	$(u_0 + i u_1)(1 - 5 {u_2}^2)$	$(u_0 + i u_1)^2 u_2$	$\left(u_0 + iu_1\right)^3$
0	$1/\sqrt{4\pi}$			
1	$\sqrt{3/4\pi}$	$\sqrt{3/8\pi}$		
2	$\sqrt{5/16\pi}$	$\sqrt{15/8\pi}$	$\sqrt{15/32\pi}$	
3	$\sqrt{7/16\pi}$	$\sqrt{21/64\pi}$	$\sqrt{105/32\pi}$	$\sqrt{35/64\pi}$

Table 3.1: Explicit expressions of the spherical harmonics up to degree 3, in both polar and Cartesian form due to Brechbühler. The last part of the table gives the common normalizing constants, e.g. $Y_1^0 = \sqrt{3/4\pi}u_2$.

To express a surface using spherical harmonics, the three coordinate functions are

decomposed and the surface $\mathbf{v}(\theta, \phi) = (x(\theta, \phi), y(\theta, \phi), z(\theta, \phi))^T$ takes the form

$$\mathbf{v}(\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathbf{c}_{l}^{m} Y_{l}^{m}(\theta,\phi) , \qquad (3.4)$$

where the coefficients \mathbf{c}_l^m are three-dimensional vectors due to the three coordinate functions. The coefficients \mathbf{c}_l^m are obtained by solving a least-squares problem. Therefore, the values of the basis functions are gathered in the matrix $\mathbf{z} = (z_{i,j(l,m)})$ with $z_{i,j(l,m)} = Y_l^m(\theta_i, \phi_i)$, where j(l,m) is a function assigning an index to every pair (l,m)and i denotes the indices of the n_{vert} points to be approximated. The coordinates of these points are arranged in $\mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{n_{vert}})^T$ and all coefficients are gathered in $\mathbf{c} = (\mathbf{c}_0^0, \mathbf{c}_1^{-1}, \mathbf{c}_1^0, \dots)^T$. The coefficients that best approximate the points in a leastsquares sense are obtained by

$$\boldsymbol{c} = (\boldsymbol{z}^T \boldsymbol{z})^{-1} \boldsymbol{z}^T \boldsymbol{v} . \tag{3.5}$$

Using spherical harmonic basis functions, we obtain a hierarchical surface description that includes further details as more coefficients are considered. This is illustrated in Fig. 3.1.

3.1.2 Computation scheme for SPHARM

The objects of interest are usually manually or semi-automatically segmented by a human expert, resulting in a voxel representation. The voxel representation has to be preprocessed to fulfill the precondition of sphere topology, e.g. via a closing operation or a smoothing filter. As I am only interested in its boundary, the voxel representation is converted to a polygonal surface mesh that serves as input for the optimization procedure. The optimization computes an appropriate (θ_i, ϕ_i) parameterization, which is used for the computation of the spherical harmonic coefficients. For shape normalization, the SPHARM coefficients are then adjusted to the first order ellipsoid regarding translation, rotation and parameter correspondence. Other normalizations are also possible, e.g. to an externally defined frame using anatomical landmarks.

3.1.3 Parameterization of SPHARM by optimization

The appropriate parameterization of the points of a surface description is a key problem. Every point *i* of the point cloud that will be approximated by the surface description is to be assigned a parameter vector (θ_i, ϕ_i) . For surfaces of spherical topology, the natural parameter space is the unit sphere with polar coordinates. A homogeneous distribution of the parameter space is essential for the decomposition of the surface. This is also necessary for an appropriate approximation of corresponding points, as described in the next section 3.1.4. We give here a brief summary of the surface parameterization procedure of Brechbühler.

A bijective mapping of the surface to the unit sphere is created, i.e., every point on the surface has to map to exactly one point on the sphere, and vice versa. The main idea of the procedure is to start with an initial parameterization. This initial parameterization is optimized so that every surface patch gets assigned an area in parameter space that is proportional to its area in object space.

First, an initial mapping from object to parameter space is constructed using discrete Laplace's equations to solve the corresponding Dirichlet problem. To obtain a homogeneous distribution of the parameter space over the surface, the initial parameterization is modified in a constrained optimization procedure considering two criteria:

- 1. Area preservation: Every object region must map to a region of proportional area in parameter space.
- 2. Minimal distortion: Every quadrilateral should map to a spherical quadrilateral in parameter space.

Brechbühler establishes constraints for area preservation, while the distortion of the mesh serves as the objective function during optimization. The optimization solves the resulting system of nonlinear equations by linearizing them and taking Newton steps.

3.1.4 SPHARM correspondence



Figure 3.2: Visualization of the SPHARM correspondence. A first order ellipsoid and six left lateral ventricles are displayed. The surface net shows the (θ_i, ϕ_i) parameterization (same parameters = same homologous points). The ridges on the first order ellipsoid are the equator and $\{0, \pi/2, \pi, 3 \cdot \pi/2\}$ meridian lines in all objects. The equator and meridian lines are emphasized in different colors. The poles are at the crossing of the meridian lines.

The scheme for establishing correspondence between objects described by SPHARM is a 3D extension of the 2D arc-length shape parameterization (see also Székely[72]). The first step is a homogeneous distribution of the parameter space over the surface, a step done in the parameterization optimization (see section 3.1.3). in the second step the parameterization is rotated in the parameter space for normalization. This rotation is based on the first order ellipsoid, which is computed from the first three SPHARM coefficients. The result of the rotation satisfies the following properties:

- The parameter locations of the poles of the first order ellipsoid match with the poles of the sphere.
- The parameter locations of the 3 main ridges of the first order ellipsoid are moving along the equator, and the 0 and π meridians of the sphere.

Correspondence is established by taking two points with the same parameter vector (θ_i, ϕ_i) on different entities of an object to be a corresponding pair.

3.1.5 Mean Squared Distance (MSD) between SPHARM objects

The correspondence between objects described by SPHARM allows the computation of distance measures between two objects. The orthogonality of the spherical harmonic basis functions allows Parseval's theorem to be used to compute the root Mean Squared Distance (\sqrt{MSD}) between two objects directly from their coefficients via a difference calculation. A correction is needed since the squared spherical harmonic basis functions do not integrate to 1 but to 4π .

$$MSD = \frac{1}{4\pi} \cdot \sum_{l=0}^{\inf} \sum_{m=-l}^{l} ||\mathbf{c}_{1,l}^m - \mathbf{c}_{2,l}^m||^2$$
(3.6)

If we want to determine error measures other than \sqrt{MSD} , we first need to appropriately sample the spherical parameterization (θ_i, ϕ_i) . Using the (θ_i, ϕ_i) point-to-point correspondence described in the previous section, error measures like the Mean Absolute Distance (MAD) or the Hausdorff-distance can be computed straightforwardly.

3.1.6 Point Distribution Model (PDM) from SPHARM

From the SPHARM description we can compute a Point Distribution Model (PDM) of the surface by sampling the parameterization. Equidistant sampling of each parameter leads to a dense sampling around the poles ($\theta = 0, \theta = \pi$) and a coarse sampling around the equator ($\theta = \pi/2$). This fact can be explained by the poles being mapped to all points having $\phi = 0...2\pi$ and $\theta = 0$, or $\theta = \pi$ (see also Fig. 3.2). The object is thus inhomogeneously sampled, as the parameterization was chosen such that areas in parameter space are proportional to areas in object space. Using a uniform icosahedron subdivision shown in Fig. 3.3, however, we gain a good approximation of a homogeneous sampling of the spherical parameter space and thus also of the object space.

Using the pre-computed parameter locations (θ_i, ϕ_i) from the icosahedron subdivi-



Figure 3.3: Icosahedron subdivision for different levels of sampling. From left to right: Base icosahedron, subdivision factors 2, 4 and 6.

sion, we can compute the PDM directly from the coefficients, as the parameter locations stay constant for all objects at a given subdivision level. The sampling points \mathbf{x}_i at the locations (θ_i, ϕ_i) are obtained using equation (3.4):

$$\mathbf{x}_i = \sum_{l=0}^{K} \sum_{m=-l}^{l} \mathbf{c}_l^m Y_l^m(\theta_i, \phi_i) \; .$$

Setting $\mathbf{x} = (\mathbf{x}_1^T, \mathbf{x}_2^T, \dots, \mathbf{x}_n^T)^T$ we write

$$\mathbf{x} = \mathbf{A}\mathbf{c}\,,\tag{3.7}$$

where the matrix \boldsymbol{A} consists of the spherical harmonic basis function values $Y_l^m(\theta_i, \phi_i)$, one for each dimension.

In this dissertation the sampling of the PDM is high (subdivision level 15-20), and it has a better localization of shape changes than SPHARM but is a less efficient shape description. The PDM is used in this dissertation mainly for the creation of the Voronoi skeleton representation.

3.2 M-rep

A m-rep is a linked set of medial primitives (see Pizer et al [53]) called medial atoms, $m = (x, r, \underline{F}, \theta)$. The atoms are formed from two equal length vectors and are composed of 1) a position x, 2) a width r, 3) a frame $\underline{F} = (\mathbf{n}, \mathbf{b}, \mathbf{b}^{\perp})$ implying the tangent plane to the medial manifold and 4) an object angle θ . The properties of medial atoms are described in further detail in the next section. The medial atoms are grouped into figures. A figure is defined as an unbranching medial sheet formed by a planar graph of medial atoms connected by intra-figural links. Figures are connected via inter-figural links. The connections of the medial atoms and the figures form a graph with edges representing either inter- or intra-figural links. In the remainder of this text, I will refer to that graph by the term 'medial graph'. In the generic case, the graph of the whole mrep is overlapping when displayed in a 2D diagram, i.e. the medial graph is non-planar. An example of an m-rep is shown in Fig. 3.4.

The m-rep description is a local and medial shape description. In my approach, I choose a small number of medial atoms, which leads to a coarse scale description.



Figure 3.4: The m-rep shape representation without (left) and with implied boundary (right) in an example of a human left hippocampus/amygdala complex. Three figures with differently colored intra-figural links are shown. The medial atoms are red dots and the implied boundary is represented by blue dots.

I compute the individual m-rep description by deforming an m-rep model into an SPHARM described boundary while constraining the medial graph to a fixed configuration (see section 4.5). The branching topology and the sampling of the medial atoms are thereby kept constant. Every object is thus expressed by the same medial graph varying only the parameters of the individual medial atoms. Therefore, I propose that a statistically derived common m-rep model is computed for each anatomical structure separately. These m-rep models incorporate the biological variability of the modeled object.

3.2.1 Medial atom properties

This section summarizes the properties of medial atoms following Fletcher and Pizer [33].

Extending Blum's point of view to 3D, each point on the medial axis of an object represents the center of the largest inscribing sphere to that object. Thus, the medial representation is implied by the boundary. The m-rep point of view is quite different since the medial representation implies the boundary. A medial atom thereby serves as a descriptor of the relation between the location of an implied 'maximal sphere' and the position and normal at the two boundary points $\mathbf{x}_{R}, \mathbf{x}_{L}$ on this sphere. At locations where the medial manifold ends, the medial atom models the boundary with 3 points and is called an end atom. The additional third point \mathbf{x}_{E} does not have to lie on the sphere. Fig. 3.5 illustrates the 2 different types of medial atoms.



Figure 3.5: Two types of medial atoms (visualized in 2D for simplicity). Left: internal atom. Right: end atom. All medial atoms encapsulate a position \mathbf{x} on a medial axis, a width r, a frame \underline{F} and an bisector angle θ . The implied sphere touches the boundary at least at 2 points $\mathbf{x}_L, \mathbf{x}_R$. An end atom is a medial atom with a component that defines an additional boundary point \mathbf{x}_E . Pictures by P. Yushkevich.

More formally, a medial atom m of the space $\mathbb{M} = (\mathbb{R}^3 \times SO(3) \times \mathbb{R}^+ \times \mathbb{R})$ can be defined as

$$m = (\mathbf{x}, r, \underline{F}, \theta) \tag{3.8}$$

where **x** is a position on the medial axis, r is the radius of the inscribing sphere, θ is

one half of the angle between the two vectors to the boundary, and \underline{F} is a 3D frame located at **x** and fitted to the geometry of the medial axis. The frame is a rotation of the standard Euclidean basis and thus an element of SO(3), which is a schematic of the group of rotations in three dimensions. The frame is composed of three orthonormal vectors,

$$\underline{F} = \{\mathbf{b}, \mathbf{b}^{\perp}, \mathbf{n}\} \tag{3.9}$$

These vectors define the local compass for the medial surface at the point \mathbf{x} , where $\{\mathbf{b}, \mathbf{b}^{\perp}\}$ span the tangent plane at \mathbf{x} , and \mathbf{n} is its normal. Moreover, the vector \mathbf{b} is chosen as the bisector of the two vectors pointing to the boundary. This is the direction of greatest narrowing of the implied sphere, given by the relation

$$\Delta r = -\mathbf{b} \cdot \cos\theta \tag{3.10}$$

Therefore, the vector \mathbf{b}^{\perp} is in the direction of constant r, i.e., no widening or narrowing. The two vectors pointing to the boundary can be defined as

$$(\mathbf{s}, \mathbf{p}) = (\mathbf{x}_L - \mathbf{x}, \mathbf{x}_R - \mathbf{x}) = (r(\mathbf{b} \cdot \cos\theta + \mathbf{n} \cdot \sin\theta), r(\mathbf{b} \cdot \cos\theta - \mathbf{n} \cdot \sin\theta))$$
(3.11)

From Blum's point of view, the maximal inscribing sphere on the edge of the medial manifold will no longer touch the boundary of the object at exactly two points but rather touches at one point of multiplicity 2. At that point of contact, a crest position on the implied boundary, the corresponding medial atom will have an angle θ equal to zero. Such an atom is ill suited to image analysis due to the instability of its dependence on a single image point. To avoid this case, an end atom is defined to imply a crest position at

$$\mathbf{x}_E = \mathbf{x} + \eta r \mathbf{b} \tag{3.12}$$

with **b** as surface normal. The parameter η defines the pointiness, i.e., elongation of the

crest region. This elongation factor is restricted to lie within $[1, 1/\cos\theta]$. When $\eta = 1$, the end atom implies a circular cross-section; $\eta = 1/\cos\theta$ yields a sharp corner.

3.2.2 Correspondence via m-rep

Since the medial graph is fixed when the m-rep model is deformed into individual cases, an inherent correspondence is given between objects whose m-rep is derived from the same medial model. We label each medial atom of the m-rep model with an unique index. Thus, when comparing different m-rep objects, medial atoms with the same index are considered homologous and corresponding. This homology using my proposed scheme is discussed in more detail and shown in section 6.3.

3.3 Combining SPHARM and m-rep

The proposed shape description is a combined description consisting of a fine-scale SPHARM description of the boundary and a coarse-scale m-rep description of the medial manifold. The properties of the two parts of the combined shape description are described in table 3.2. The two descriptions complement each other. The m-rep description is suitable for doing shape analysis since it has a low number of features. These features and their changes are intuitive. The SPHARM description is used to stabilize the unstable branching topology of the m-rep. It does so by defining a smooth shape space and a spatial correspondence for branching topology comparisons, as described in the next chapter.

3.3.1 Computing the combined medial-boundary shape description of an individual object

Starting from a voxel-based segmentation of the object, the SPHARM description is computed as described in section 3.1.2. A previously computed m-rep model is then

SPHARM	m-rep	
High number of features.	Low number of features	
E.g. hippocampus: 507 coefficients	8x3 m-rep: 24 thickness +	
	72 location + 72 orientation features	
Captures fine scale shape changes effi-	Captures coarse scale shape changes ef-	
ciently.	ficiently.	
Analytical computation of geometric	Discrete computation of geometric	
properties.	properties.	
No intuitive interpretation of coeffi-	Features are intuitive. Local changes in	
cients and its changes since global fea-	thickness, location and orientation.	
tures.		
Stable in presence of shape variation.	Branching topology unstable in pres-	
	ence of shape variation. An appropri-	
	ate fixed branching topology is needed	
3D correspondence by parameter align-	3D correspondence given if medial	
ment of first ellipsoid poles and ridges.	graph (figures $+$ atoms) is constant.	

Table 3.2: Summary of properties of SPHARM and m-rep

fit into the boundary described by SPHARM. A good initial estimate of the m-rep is obtained using the boundary correspondence of the SPHARM associated with the m-rep model. The fitting process yields the deformed individual m-rep description of the object. The proposed description is then the combination of the SPHARM and the deformed m-rep. Fig. 3.6 shows how the combined description of a population of objects is computed.

The major problem associated with this scheme is the determination of an appropriate m-rep model. This m-rep model should incorporate the biological variability of the object and should be computed automatically and in a stable way. The solution to this problem, which is described in the next chapter, is the main contribution in this dissertation.



Figure 3.6: Computation of the combined description for a study population of objects. First the SPHARM description is computed and then the m-rep is determined via fitting a previously computed m-rep model into the SPHARM boundary. The m-rep model was built from a training population. Shape analysis is applied to a study population to derive new findings.

Chapter 4

Computing a common m-rep model of an object population

4.1 Overview

In the proposed combined shape description, the main problem that remained was the computation of an m-rep model in the presence of biological shape variability. Given an object population for an anatomical object, how can we determine a common m-rep model automatically and in a stable way?

To date the m-rep model building process has been conducted manually by human interaction. A model was derived from one sample object, representative of a population of biological objects. Thus, it was assumed that the sample's topology and geometry represent a set of similar objects with respect to the population and that a human expert can reliably define the appropriate topology and geometry. My new scheme presented here moves towards a more stable statistical description by taking into account a whole population and by automatically deriving a 'common' m-rep model from statistical observations of the shape properties.

My scheme can be subdivided into 3 steps and is laid out in Fig. 4.1. The input to the scheme is a representative training population of objects for the anatomical object to be modeled. First, the scheme computes a smooth SPHARM shape representation for every individual object in the training population. A Principal Component Analysis (PCA) computes a shape space from the SPHARM objects, as described later in section 4.2. This shape space incorporates the major part of the shape variability in the training population. All subsequent computations are performed on this shape space. The second step of the scheme computes the medial branching topology of the common m-rep model using Voronoi skeletons. The third step computes the minimal sampling of the common branching topology with a predefined maximal approximation error in the shape space. The detailed algorithms of the topology computation are described in section 4.3, and those of the optimal sampling computation are described in section 4.4. Section 4.5 describes how the m-rep model estimated from the optimal sampling is fitted into a SPHARM object for an appropriate m-rep representation. Issues of the established homology and the scheme's stability are discussed later in chapter 6.



Figure 4.1: Three steps to automatically compute a m-rep model in a stable way. Given is a training population of objects for an anatomical structure. Step 1. Computation of a shape space by PCA, section 4.2. Step 2. Extraction of a common medial branching using pruned Voronoi skeletons, section 4.3. Step 3. Computation of the minimal sampling, section 4.4.

4.2 The shape space for an anatomical object

Figure 4.2: The proposed shape space is spanned by the first few eigenmodes of deformation. The first two dimensions of the 6-dimensional shape space from a hippocampusamygdala study are visualized.

As a crucial step in my scheme, I define a shape space for all computations. The shape space for an anatomical object is derived from a training population. The shape space is used for all subsequent computations of my scheme, instead of applying the computations directly to the training population. The shape space smoothes the shape variability in the training population, thus making the computations of my scheme more stable. I assume that the shape space is an appropriate representation of the object's biological variability. The Principal Component Analysis (PCA) is applied to SPHARM objects of the training population to compute the eigenmodes of deformation (see also appendix 7). The shape space is then defined by the average object and the major eigenmodes of deformation.

The shape space computation starts from the smooth SPHARM shape representation \mathbf{c}_i for each individual shape, as described in section 3.1. PCA is then applied to the population of coefficients resulting in a average coefficient vector $\mathbf{\bar{c}}$ and the eigenmodes of deformation $\{(\lambda_1, \mathbf{v}_1) \dots (\lambda_{n-1}, \mathbf{v}_{n-1})\}$. The first few eigenmodes $\{(\lambda_1, \mathbf{v}_1) \dots (\lambda_k, \mathbf{v}_k)\}$ are selected and the remaining eigenmodes are discarded. The first k eigenmodes are chosen to cover at least 95% of the population's variability.

Formally, the PCA eigenmode selection is defined as follows:

$$\Sigma = \frac{1}{n-1} \sum_{i} (\mathbf{c}_i - \bar{\mathbf{c}}) \cdot (\mathbf{c}_i - \bar{\mathbf{c}})^T$$
(4.1)

$$0 = (\Sigma - \lambda_i \cdot I_n) \cdot \mathbf{v}_i; i = 1 \dots n - 1$$

$$(4.2)$$

$$0.95 \leq (\sum_{i=0}^{\kappa} \lambda_i) / (\sum_{i=0}^{n-1} \lambda_i)$$
 (4.3)

$$Space_{shape} = \{ \bar{\mathbf{c}} \pm 2 \cdot \sqrt{\lambda_i} \cdot \mathbf{v}_i \}; i = 1 \dots k.$$
 (4.4)

A discrete description of the shape space is gained by sampling it either uniformly or probabilistically. These samples form an object set that is a representative sampling of the shape space. All subsequent computations are then applied to this object set. The number of objects in this set can be considerably higher than the original number of objects in the training population. An example of such an object set is presented in Fig. 4.3.



Figure 4.3: Subset of the object set from a PCA shape space for a hippocampusamygdala population: Average case (left) with deformations along all selected eigenmodes.

As with any sampling method, information is lost by the discretization of the shape space. Fig. 4.4 presents 2 examples of non-Gaussian distributions in the PCA shape space. Regular sampling on these PCA shape space would result in members of the object set that do not represent objects similar to the original population. For such shape spaces, a special sampling method would be needed in order to create an object set that appropriately represents the population. In all my studies with real datasets, the training populations appear to have a nearly Gaussian distributions in the PCA shape space (see Fig. 4.5). Thus, a straightforward sampling method was chosen computing the object set by uniformly sampling 5 objects along every eigenmode. The appropriateness of computed object set should be checked for every shape space.



Figure 4.4: Two examples of the $(\mathbf{v}_1, \mathbf{v}_2)$ PCA space from synthetic datasets with non-Gaussian distributions. Left: bimodal distribution, Right: banana shaped/multimodal distribution.



Figure 4.5: Three examples of the $(\mathbf{v}_1, \mathbf{v}_2)$ PCA space from real datasets. On the left, the population of 40 (left and mirrored right) lateral ventricles is displayed (Weinberger dataset). In the middle, the population of 46 hippocampus-amygdala (left and mirrored right) and on the right, the population of 172 hippocampi (left and mirrored right) are displayed.

4.3 Branching topology for a common medial model

In the first part of this section (the first step in Fig. 4.6), the computation of the branching topology as a set of medial sheets for an individual object is described. In the second part (the second step in Fig. 4.6), these medial sheets are compared in the object set to compute a common branching topology.



Figure 4.6: Extraction of a common branching topology. After computing the branching topologies as sets of medial sheets, all relevant sheets are combined in a common model topology. The medial sheets are computed using pruned Voronoi skeletons.

Two topology terms are used in this section. First, the *branching topology* of the medial model represents the figures and the intra-figural links of the medial graph. Second, the *topology of the skeleton* is related to its genus, which is determined by the number of distinct parts, holes, cavities and inclusions. As only skeletons from objects described by SPHARM are generated, the topology of the skeleton is spherical.

4.3.1 Branching topology for individual objects

The branching topology of an individual object is represented by a set of medial sheets from the pruned Voronoi skeleton. A general view on Voronoi skeletons is presented first in section 4.3.1.1. The inner 3D Voronoi skeleton is calculated from a shape described by SPHARM using the scheme described in section 4.3.1.2. My new pruning scheme for a Voronoi skeleton is the topic of the remainder of section 4.3. The pruning scheme starts with grouping the Voronoi skeleton into a set of non-branching, non-self-intersecting medial sheets (see section 4.3.1.3). A medial sheet is a 2D manifold that comprises a group of directly connected Voronoi faces. The initially large number of medial sheets is reduced to a small number via several pruning steps incorporating heuristics about areal and volumetric contributions (see section 4.3.1.4 - 4.3.1.6). My pruning scheme thereby only prunes whole medial sheets and no single Voronoi vertices. The pruning scheme is enhanced by a merging step to further reduce the number of medial sheets (see section 4.3.1.7). The computed set of medial sheets represent the branching topology.

4.3.1.1 General view on Voronoi skeletons

Voronoi skeletons exploit the relation between the Voronoi diagram of a point set and its skeleton. That Voronoi Diagram is dual to another fundamental structure in computational geometry, the Delaunay Triangulation.

The Voronoi Diagram of a discrete, n-dimensional point set is a partition of the space into cells such that each cell of the partition contains exactly one member of the point set and is the locus of all points which are nearer to this member than to any other. In this way the space is divided into convex cells that are not always finite. In 3D the Voronoi cell is called a Voronoi polyhedron and the member of the point set that is inside the Voronoi cell is called the generating point. A face that separates two Voronoi polyhedra is called a Voronoi face. A Voronoi face has equal distance to the generating points of its two neighboring Voronoi polyhedra. The edges of a Voronoi

face are called Voronoi edges. They have equal distance to all generating points of the Voronoi polyhedron they belong to. Finally the vertices of the Voronoi polyhedra are called Voronoi vertices. Again they are equidistant from all generating points of the Voronoi polyhedra they belong to.

The Voronoi skeleton concept can be easily understood by a discrete prairie-fire analogy, similar to the well-known continuous one. The only difference is that not the continuous outline of the object but a discrete number of boundary points are set on fire. If the fire is evolving isotropically, circular fire-fronts are generated at each of these boundary points. These fire-fronts will quench at the boundaries of the point set's Voronoi cells.

Boissonnat and Kofakis [12] showed that the skeleton of an object that is described by a set of boundary points can be approximated from a subgraph of the point set's Voronoi diagram. Brandt and Algazi [19] studied the quality of this approximation in terms of the sampling density. Using morphological methods in 2D, Schmitt [56] has shown that if the density of the boundary point set uniformly goes to infinity, then the corresponding Voronoi diagram converges to the exact skeleton after the bisector of adjacent generating points are removed. The *exact* skeleton is defined as the correct skeleton of the boundary described by the *connected* point set. Amenta [1] states that this is not true for the 3D case. The 3D Voronoi diagram based on a uniform infinite sampling can have Voronoi vertices that are not part of the exact skeleton. Thus, when generating 3D Voronoi skeletons, the vertices of the skeleton have to be checked as to whether they are part of the exact skeleton. Since the exact skeleton is not known, heuristic criteria are used to do the check. The topology of the exact skeleton is usually known since it is equal to the topology of the original object.

The generation of Voronoi skeletons can be described by the following four major steps:

1. Approximation of the object's boundary by a sufficiently dense set of generating

points. The extraction of the generating points is usually performed from a discrete image raster and only the subsequent steps are computed with real numbers in \mathbb{R}^3 . In contrast, in this dissertation, all steps of the Voronoi skeleton generation are performed with real numbers (see section 4.3.1.2) since the generating points are sampled from the continuously defined SPHARM.

- 2. Generation of the Voronoi diagram of these generating points. It is sufficient to compute the Voronoi diagram inside the object if we are only interested in the 'inner' skeleton. The 'inner' Voronoi diagram can be determined in several ways:
 - The Voronoi diagram is intersected with the object
 - Only Voronoi vertices are chosen that are inside the object, and only edges connecting two such vertices are included.
 - Only Voronoi vertices are chosen that correspond to those Delaunay triangles that are fully inside the object.

Like other researchers (Naef, Attali), I chose to use the second of these methods (see section 4.3.1.2).

- 3. Removal of Voronoi vertices that are not part of the exact skeleton. Since the exact skeleton is not known, this can only be done in an approximate fashion (see section 4.3.1.2).
- 4. Extraction of the skeleton from the inner Voronoi diagram as a subgraph. We will call this process pruning or regularization. This is a crucial and the most difficult part of the whole procedure (see section 4.3.1.4).

4.3.1.2 Voronoi skeleton generation from SPHARM

SPHARM is a global shape representation, from which we cannot directly calculate a Voronoi skeleton. Since Voronoi skeletons are computed from the Voronoi graph of a



Figure 4.7: Icosahedron subdivision of a human hippocampus. The PDM (A) of a human hippocampus-amygdala is displayed, which was determined via a spherical icosahedron subdivision (B) of the SPHARM parameterization.

point set, we need to determine a point set from the SPHARM object. Székely [71] claims that for a correct computation of the Voronoi skeleton the point set should be uniformly sampled. A fine sampling of the SPHARM object determines a PDM via an icosahedron subdivision of the SPHARM parameter space (see Fig. 4.7), as described in section 3.1.6. The PDM can be calculated to an arbitrary precision by changing the level of the icosahedron subdivision.

The Voronoi diagram is calculated by an implementation of the incremental 4Dconvex-hull approach (see Attali [4]). Then the Voronoi diagram is partitioned into an inner and an outer part. The inner Voronoi diagram is defined by all Voronoi vertices and Voronoi elements attached to theses vertices that are inside the object. The Voronoi diagram is quite stable against noisy boundary perturbations because the SPHARM is a smooth, continuous description and the PDM is a fine sampling of this smooth description.

In 2D the Voronoi vertices approximate the medial axis of the curve described by the given point set. Unfortunately this is not necessarily the case in 3D. The Voronoi diagram in 3D can have vertices that are not part of the exact skeleton. In my experiments, I encountered several Voronoi skeletons that have a different topology than the exact skeleton. For these cases, increasing the density of the sampling does not change the incorrect topology of the skeleton. My pruning scheme described in section 4.3.1.4 cannot remove such incorrect vertices since the scheme preserves the topology of the skeleton. Thus, I developed a preprocessing step for the pruning scheme that enables the removal of such vertices. Additionally, this preprocessing step removes all skeletal vertices that are not adjacent to Voronoi faces. This removes all single Voronoi vertices and edges from the skeleton.



Figure 4.8: Preprocessing of Voronoi skeletons: A: Single Voronoi vertices (arrow) and single Voronoi edges (white lines). B: Topology violation: Arrow points at skeleton topology violation, a non-prunable inclusion (bubble). C: After the removal of one Voronoi vertex (arrow) using the closest-to-surface-heuristic, the topology violation is no longer present.

The preprocessing step ensures that the skeleton has no cavities or inclusions. This leads to a correct skeleton topology, if there are no holes present. Since the skeleton topology of objects described by SPHARM is spherical, a topology preserving pruning algorithm should theoretically be able to remove *all* except one Voronoi vertex. If that is not possible, a topology violation is found so there are vertices that cannot be pruned. These vertices are not part of the exact skeleton and thus are incorrect. It is impossible to determine which Voronoi vertices of the skeleton are correct and which are incorrect. Therefore, I implemented a procedure for the successive removal of nonprunable Voronoi vertices based on heuristics. Two heuristics were tested as mentioned below. The heuristic selects a single non-prunable Voronoi vertex, which is removed. The preprocessing algorithm then checks whether the remaining skeleton can be pruned to a single vertex via the topology preserving pruning. If this is not the case, then another non-prunable Voronoi vertex is removed and the check is performed another time. This remove-and-check procedure is performed until the resulting skeleton can be reduced to a single vertex via the topology preserving pruning.

Other researchers proposed a heuristic based on a match between the orientation of the dual Delaunay-Tetrahedron for each Voronoi vertex with the correct surface normal. This heuristic produced visually bad results and motivated me to develop a different heuristic. This new heuristic removes successively the Voronoi vertex that is closest to the surface. The closest-to-surface-heuristic turned out to be effective. Only a small number of Voronoi vertices are removed using this heuristic. After using this preprocessing step, all remaining Voronoi vertices theoretically are prunable.





before correction

after correction

Figure 4.9: Failed preprocessing of Voronoi skeletons: Correcting a topology violation via removal of vertices can create a new topology violation. On the left, one can see that the topology violation was due to an incorrect Voronoi diagram (a non-generic case that could not be handled by the applied 4D-Complex-Hull-algorithm).

In all my tests I observed one case in which the removal heuristic generated a new topology violation (see Fig. 4.9). The nature of the newly created violation is different than the original one. Instead of an inclusion or cavity, there are holes in the medial manifold. In this specific case, both heuristics lead to holes. Also, the initial topology violation was probably due to an incorrect computation of the Voronoi diagram, which might be due to a non-generic case that could not be handled by the applied 4D-Convex-Hull-algorithm. However, such cases did not occur in the shape space that I propose.

4.3.1.3 Grouping Voronoi faces to sheets

This step computes the branching topology of the skeleton as a set of non-branching medial sheet. The algorithm presented in this section groups the initially unordered Voronoi skeleton into a set of medial sheets. In his dissertation, Naef [51] proposed such a grouping algorithm for Voronoi skeletons. His grouping algorithm consists of two steps: A) an initial grouping step, which results in a set of medial groups but with many of these groups still having branches; B) a refinement step that breaks up the branching groups into non-branching sub-groups.

My algorithm uses the same idea as Naef for the initial grouping step but with a very different, computationally more efficient implementation. The scheme then takes a different approach to solve the refinement step.

Initial grouping step : Initially the grouping algorithm puts all faces into a *list* of ungrouped faces. It then selects and removes an arbitrary face f_0 from the list of ungrouped faces as a starting point for a new, current group. This face f_0 is put into a new *list of unprocessed faces*. The list of unprocessed faces contains all faces that will be grouped into the current group but they are not yet processed by the algorithm. The list of ungrouped faces contains all faces of the whole skeleton that are not yet grouped. From the list of unprocessed faces a face f_i is selected. It receives the identification label of the current group and is removed from the list of ungrouped faces. For all edges of f_i , which bound exactly one other face f_{adj} with f_i , the face f_{adj} is added to the list of unprocessed faces only if f_{adj} is in the list of ungrouped faces. This is repeated until there are no more faces in the list of unprocessed faces. Then a new face f_0 is selected from the list of ungrouped faces for a new group and the whole procedure is repeated until the list of ungrouped faces is empty. A typical progression of the grouping algorithm is displayed in Fig. 4.10.

Some undesired situations can arise as shown in Fig. 4.11. The arrows show a possible sequence in which the faces could have been processed by the grouping al-



Figure 4.10: Progression of the initial grouping algorithm (left to right, top to bottom). In this example, no sheet will contain sheet-internal branches. Pictures by M. Naef.

gorithm. The resulting group no longer represents a non-branching sheet because it consists of edges where 3 faces of this group meet. In the following, I will call such edges sheet-internal branches. An additional refinement step is necessary to remove these sheet-internal branches.



Figure 4.11: Creation of sheet-internal branches in initial grouping algorithm. In this example, the medial sheet generated by the initial grouping will contain sheet-internal branches (bold edges). The arrows indicate a possible progression of the grouping algorithm. Pictures by M. Naef.

Refinement step : The refinement step assigns three new groups along the edges of a sheet-internal branch. The new groups are then propagated inside the original group. This is repeated until there are no more sheet-internal branches. Naef proposed to use a graph-breadth-first propagation. My algorithm uses a weighted breadth-first propagation based on a geometric continuity criterion. It disallows propagation to adjacent faces whose orientation differs by more than a predefined angle. This results in patches that are impossible to reach. These patches are detected by inspecting whether the whole original group has been reassigned to one of the three new groups. If this is not the case and thus there are unreachable patches, the predefined angle is doubled and the propagation is continued at the boundary of these patches. The current implementation of the algorithm starts with an angular difference constraint of 15 degrees, which was determined empirically. This propagation algorithm ensures that the partitioning of a branching sheet into sub-sheets satisfies a geometric continuity. This grouping scheme produces visually good results. Nevertheless, an additional criterion to ensure the continuity of the radius function along the sub-sheets might further improve the algorithm. I did not implement such an extension of the propagation algorithm.

The need for a computationally efficient implementation is evident. The rather simple hippocampus-amygdala object is described by ~ 4000 sampling points on the boundary. The skeleton generation results in $\sim 30,000$ inner Voronoi vertices. The grouping algorithm computes ~ 1000 sheets on the unpruned skeleton. A first, straightforward implementation of the grouping algorithm needs a computation time for the *initial* grouping step of more than 1 hour on a Sun Sparc Ultra 10 workstation. My new implementation that uses set of precomputed look-up-tables to store all relationships between the Voronoi diagram and the groups is much less memory efficient but computationally performs much better. A computation time of about 5 seconds is needed to set up the tables, compute the initial grouping, and compute the refinement steps for the same object on the same machine.

4.3.1.4 Pruning Voronoi sheets

A large number of Voronoi vertices are generated due to noise artifacts or due to the dense sampling of the surface. These vertices often contribute only an insignificant amount to the reconstruction of the object. By pruning, which means removing insignificant vertices, the skeleton properties are significantly less complex. My proposed pruning does not remove single Voronoi vertices, but rather remove medial sheets that comprise a group Voronoi faces. Pruning is also used to create a coarse-scale representation of the skeleton. A sampling of this coarse-scale skeleton leads to the coarse-scale m-rep description as the final step of my scheme. The pruning algorithm proposed in this dissertation is based on a topology preserving pruning, which removes vertices only if the topology of the skeleton is preserved. The topology preserving pruning method was proposed by Attali [3] and is not discussed in this document.

Pruning algorithms can be divided into 2 groups: algorithms based on local significance criteria and algorithms based on global criteria. A local criterion is only influenced by its neighborhood, and it is independent of the rest of the skeleton. Local criteria are, for example, the size of the local thickness or the local bisector angle. A global criterion is influenced by all parts of the skeleton, for example the areal coverage of a part of the skeleton in relation to the whole skeleton.

Most research by others has proposed 3D skeleton pruning based on local significance criteria. Such pruning methods proposed by Naef [51] and Attali [5, 3] were influential to this dissertation but the reader is referred to the original publications. Local significance criteria seem to be unsuited for pruning because there is no direct link between the local significance and the object shape. Thus heuristics are often used to create such a link in local pruning criteria. To my knowledge, no heuristic has been found that results in an appropriate pruning. My scheme is purely based on global significance criteria.

The next two sections discuss two different global significance criteria that prune whole medial sheets of Voronoi faces criterion is based in the areal contribution of the sheet to the manifold and the other is based on the volumetric contribution of the reconstruction to the object. For these criteria a preliminary step that groups the Voronoi faces into sheets is essential. The sheets are being treated as independent of each other and thus a simple thresholding is applied in order to mark the sheets that are to be pruned. Only those sheets that can be pruned without changing the topology of the skeleton are actually pruned. The threshold levels used in this dissertation are suitable for the studied objects but other, more complex objects might need a different set of thresholds.

The pruning of medial sheets usually changes the branching topology of the skeleton. The pruning can create new sheets or merge existing sheets. Therefore, a complete sheetpruning scheme includes an additional grouping step that is performed directly after the pruning step. Since sheets were possibly created, the skeleton needs to be pruned again with the same criterion. This possibly changes the branching topology again. Thus, the sheet-pruning scheme applies a loop consisting of a grouping step followed by pruning step until no more sheets can be pruned.



Full Display

Wire frame Display

Figure 4.12: Final pruning step that prunes sheets consisting only of one Voronoi face. The sheets are labeled with side-wise written Roman numbers I,II and III (in white). The single-face sheet (III, only visible in Wire frame display), which is present before the application of this pruning step (top row), is removed afterwards (bottom row).

As a consequence of the discussions below I implemented the following pruning scheme. The scheme starts with a preprocessing step of the Voronoi skeleton. Then the scheme applies a conservative threshold for the areal contribution criterion. This step removes a large number of tiny sheets. The remaining medial sheets are pruned using the volumetric contribution criterion. After this pruning step a few medial sheets consisting only of a single Voronoi face are often present. These sheets cannot be pruned since this would alter the topology of the skeleton. The contribution of these sheets is minimal, and therefore they are merged in a final pruning step with one of their neighbors (see Fig. 4.12).

4.3.1.5 Pruning by areal contribution to the medial manifold

Naef proposed a global significance criterion based on the areal contribution of a sheet to the whole medial manifold. He assumes that the areal contribution of a sheet is proportional to the number of Voronoi vertices of that sheet. This criterion prunes a sheet s_i if it has fewer Voronoi vertices than a predefined threshold:

$$C_{area} = \frac{n_{vertices,s_i}}{n_{vertices,skel}} \tag{4.5}$$

The areal contribution does not sufficiently correlate with the contribution of a sheet to the object shape. August [6, 7] described several generic situations that generate large medial sheets with little significant contribution to the object shape. Nevertheless, choosing a conservative pruning threshold, this pruning criterion can be used to remove tiny sheets that are unlikely to have a high significance to the object shape. If a non-conservative threshold is chosen, the criterion prunes sheets that a human observer considers significant to the object shape. This criterion can be implemented in a computationally efficient way. As a first pruning step in my pruning scheme tiny sheets are efficiently removed from the skeleton using the areal contribution criterion with a threshold of 0.1% of the overall number of skeleton vertices. Since only tiny sheets are removed, the skeleton still has a large number of sheets (see Fig. 4.14).

This pruning step will not remove medial sheets that are made of incorrect vertices, if these are large enough. For example, in section 4.3.1.2, the computed Voronoi skeleton has a topology violation due to an inclusion. This topology violation was generated by Voronoi vertices that are not part of the exact skeleton. The preprocessing step punches a hole into the inclusion so that it *can* be pruned (see Fig. 4.8). This modified bubble forms a large sheet that *cannot* be pruned using the areal contribution criterion.

4.3.1.6 Pruning by volumetric contribution to the whole object

The second sheet-pruning algorithm is based on the volumetric contribution of a sheet. I define the volumetric contribution of a sheet s_i as the relative volumetric difference of the object reconstruction from the skeleton with and without the sheet s_i :

$$C_{volume} = \frac{V_{skel} - V_{skel-s_i}}{V_{skel}} \tag{4.6}$$

The volumetric contribution criterion proved to be far superior to the areal contribution criterion since the volumetric contribution correlates directly with the significance of a sheet to the object shape.

Parts of the 3D Voronoi skeleton that do not belong to the exact skeleton are removed by the volumetric contribution criterion because the volumetric contribution of these parts is close to zero. The areal contribution criterion is not able to remove most of these parts.



A: Boundary | B: After preprocess | C: After pruning | D: Reconstruction

Figure 4.13: Application of the pruning scheme to a human left hippocampus-amygdala. A: Boundary representation. B: Voronoi skeleton after preprocessing. C: Voronoi skeleton after pruning. D: Reconstructions of B (transp. red) and C (blue). The skeleton is pruned from ~ 1200 to 3 sheets.

My pruning scheme first applies a conservative threshold, pruning sheets with a volumetric contribution smaller than 0.1% (see Fig. 4.14). In a next step, an additional merging step is applied that tries to merge neighboring sheets (see section 4.3.1.7). Following the merging step, a less conservative pruning step with a threshold of 1.0% is

applied. Some results are presented in Figs. 4.13- 4.17.

The computational efficiency of this pruning criterion largely depends on the computation of the volume of the object reconstruction from the skeleton. The approach in the current implementation computes the volume by counting voxels in a fine-scale isotropic 3D voxel-based object reconstruction. The reconstruction is based on drawing a sphere with corresponding radius for every Voronoi vertex. Considering the previously mentioned example of a hippocampus-amygdala with \sim 30,000 Voronoi vertices, it is evident that such a sphere drawing algorithm needs to be computationally efficient. As I did not find an existing voxel-drawing algorithm for spheres, I developed a filled-sphere drawing algorithm based on similar principles as Bresenham's line drawing algorithm.

The results presented here show that a considerable reduction of the number of medial sheets is possible with sacrificing only little accuracy of the reconstruction. In fact, so far, the pruned skeletons of *all* objects had a volumetric overlap with the original object of more than 98%.

4.3.1.7 Merging Voronoi sheets

The grouping algorithm described in section 4.3.1.3 handles a sheet-internal branch by partitioning the sheet into 3 parts along the branch. It is always possible to merge 2 of these 3 parts without creating new branches on the merged sheet. Thus, the partition into 3 parts unnecessarily fragments the medial sheets. In order to avoid this fragmentation, neighboring sheets should be merged. The result of the merge strongly depends on the choice of the merging order, which ranks all possible merges of the skeleton. My scheme first establishes an appropriate merging order. It then merges pairs of neighboring sheets according to that order if no new sheet-internal branches are created on the merged sheet.

My scheme computes the order of merging using a heuristic. The first developed heuristic aims to create maximally largest sheets. This heuristic orders neighboring



Figure 4.14: Steps of my pruning scheme. The resulting skeleton of a human hippocampus-amygdala is shown after each step of my scheme. The skeleton is pruned from 1226 to 3 sheets. The refconstructions of the preprocessed skeleton (in transparent red) and the pruned skeleton in blue) is shown at each step. Since the pruned skeleton is a subset of the preprocessed skeleton, no blue can be seen, but rather only violet (=blue+red) parts where both agree and tiny red parts that are only represented in the preprocessed skeleton.


Figure 4.15: Pruning scheme applied to the average hippocampus-amygdala objects of the left and right brain hemisphere (side views). A: Boundary. B: Skeleton after preprocessing. C: Skeleton after pruning. D: Reconstruction of B (transp. red) and C (blue). Skeletons are pruned from ~ 1200 to 2 sheets (both left and right).



Figure 4.16: Pruning scheme applied to a lateral ventricle pair (side views). A: Boundary. B: Skeleton after preprocessing. C: Skeleton after pruning. D: Reconstruction of B (transp. red) and C (blue). Skeletons are pruned from ~ 1600 to 3 sheets (both).

sheets according to the number of Voronoi vertices per sheet. The results are visually not satisfying because sheets are merged that clearly do not belong to each other. Inspecting these cases, I observed a break in the geometric or the radial consistency of the merged sheets. This means that a discontinuity of the orientation and/or of the radius next to the merged edges was created by the merge. The second developed heuristic handles this problem. This heuristic orders pairs of neighboring sheets (s_i, s_j) according to a combined radial and orientational continuity criterion at the common border:

$$C_{continuity}(s_{i}, s_{j}) = \sum_{x \in border(s_{i}, s_{j})} \frac{C_{rad}(x, s_{i}, s_{j})^{2} + C_{ori}(x, s_{i}, s_{j})^{2}}{2}$$
(4.7)

$$C_{rad}(x, s_{i}, s_{j}) = \frac{|r(x, s_{i}) - r(x, s_{j})|}{\max_{s_{i}, s_{j}}(r) - \min_{s_{i}, s_{j}}(r)}$$

$$C_{ori}(x, s_{i}, s_{j}) = \frac{\angle(\mathbf{N}(x, s_{i}), \mathbf{N}(x, s_{j}))}{\pi}$$

The ordering of the merging is determined by the ranking pairs by the continuity value $C_{continuity}$. Pairs that exceed a certain threshold for the continuity criterion are not merged at all. This means that at the end of the merging process, there are still sheets that could be but are not merged. I observed only very few cases in my studies where I would merge such non-merged sheets, since a clear discontinuity of the orientation or the radius is otherwise created. The second heuristic produced considerably better results.

In my pruning scheme, a merging step precedes every non-conservative pruning step. Thus, the non-conservative pruning is less likely to prune parts that are below the threshold but are significant to the object shape. Nevertheless, both the conservative and the non-conservative pruning might remove parts of the skeleton that a human observer would judge significant. According to my experience, the probability of such an event is lower if a merging step is applied.



Figure 4.17: Application of the pruning scheme to various brain structures. A: Pallidate globe, B: Thalamus. C: Putamen. For each structure the initial sampled boundary (top) and the computed pruned Voronoi skeleton (bottom) is shown. Skeletons are pruned from $\sim 1500-2000$ to 2 sheets.



Figure 4.18: The pruning scheme applied to the object set of a left hippocampusamygdala population. First row: Average hippocampus-amygdala. Other rows: Objects at shape space positions -2, -1, +1, +2 along axes of eigenmodes $\lambda_1 \dots \lambda_6$.

4.3.2 Common branching topology via common spatial frame

In the preceding sections, my scheme to compute a set of medial sheets from a Voronoi skeleton was described. This set of medial sheets is then used in a description of the branching topology of that single object. In this section, the branching topologies of a group of sample objects are combined into a common branching topology that unifies the whole set of branching topologies described in the object set. The object set comprises quite a large number of objects (e.g. 25 objects in Fig. 4.18). The details and properties of the object set were described in the previous section 4.2. The common branching topology is computed via spatially matching medial sheets of the Voronoi skeletons between different objects using the correspondence on the boundary. In Fig. 4.19, the branching topology of selected members of an example object set are displayed next to the computed common branching topology.



Figure 4.19: Topology matching scheme applied to the object set (see also Fig. 4.18). A,B,C: Objects from the object set that contribute to the common branching topology. F: Final branching topology with sheets as clouds of Voronoi vertices.

The problem of comparing branching topologies has been adressed in 2D by Siddiqi [60] and others mainly via matching medial graph structures. To my knowledge, there has been no work reported in 3D to date. In 2D, The results of August [6] and others have shown that the medial branching topology is quite unstable. In 3D, the medial branching topology is even more unstable and ambiguous than in 2D. The best matching algorithms developed for 2D all use optimization methods to solve the NP-hard problem of matching trees in an acceptable time. These algorithms would be computationally less efficient in 3D. Also, they cannot be extended straightforwardly to 3D since in the graph of the 3D branching topology of an object of spherical topology is no longer a tree as in 2D but a general graph. Thus, I developed a matching algorithm that is not based on graph matching but on spatial correspondence. The branching topology is thereby not represented as an abstract graph but rather via its spatial distributions. These spatial distributions are defined by the Voronoi vertices of the medial sheets. After defining a common spatial frame (see section 4.3.2.1), a spatial correspondence can be computed using a distance measure between the sheets (see section 4.3.2.2). The common branching topology is then computed in iterative procedure that matches the branching topologies in the object set with the current common branching topology (see section 4.3.2.3). All non-matching medial sheets are thereby incorporated into the common branching topology so that the common branching topology incorporates all medial sheets necessary to describe the shape space.

4.3.2.1 Common frame for spatial correspondence between medial sheets

Spatial comparisons between medial sheets of objects in the shape space can be done after the objects have been mapped into one template object in a common spatial frame. The mapping is applied to both the boundary and the Voronoi skeleton of the objects. The common spatial frame is chosen as the boundary of the average object according to the SPHARM shape space, in order to minimize the mapping distortions. Every member of the object set is mapped into the common frame based on the correspondence of the individual boundaries with the common frame boundary, which is the average case boundary. The objects are already in rigid registration according to the first order ellipsoid of the SPHARM description. Since shape variability in the object set suggests that a warped registration technique is necessary, I chose to warp all objects into the common frame. The SPHARM description is used to create correspondences (see section 3.1.4) on the PDM's of the boundary between each object and the template object in the common frame. The PDM boundary correspondence is interpolated in the whole 3D space via thin plate splines (TPS) [13, 64]. Thus, the PDM's perfectly overlay for all objects in the common frame.

In summary, all branching topologies are mapped into the common frame by a TPS warp of the Voronoi skeletons into the average object, where their PDM, which are the Voronoi skeleton's generating points, perfectly overlay. A schematic overview of this scheme is displayed in Fig. 4.20.



Match medial sheets - Extract common topology

Figure 4.20: Schematic overview of matching procedure

4.3.2.2 Identifying corresponding sheets in a common spatial frame

Given that all medial sheets of the object set are mapped into a common spatial frame, a matching criterion can be defined to assess how well two different sheets spatially correspond. Visually, a high degree of overlap between matching sheets in the common frame can be observed. The centers of the medial sheets match better than the edges, which are quite sensitive to boundary noise. I defined a quite robust matching criterion that takes into account the non-isotropic spatial distribution of the Voronoi vertices of the medial sheets. Specifically, for every sheet s_i the covariance matrix of its Voronoi vertices and the average Voronoi vertices' locations μ_i (= sheet center) is computed. This covariance matrix Σ_i can be seen as an ellipsoid approximating the medial sheet s_i . The matching criterion is then computed as the paired Mahalanobis distance between the sheet centers (4.8).

$$d_{Maha}(s_i, x) = d_{Maha}(\mu_i, \Sigma_i, x) = (x - \mu_i)' \cdot \Sigma_i^{-1} \cdot (x - \mu_i)$$

$$crit_{Maha}((\mu_i, \Sigma_i), (\mu_j, \Sigma_j)) = \frac{d_{Maha}(\mu_i, \Sigma_i, \mu_j) + d_{Maha}(\mu_j, \Sigma_j, \mu_i)}{2}$$

$$(4.8)$$

$$if \ crit_{Maha}((\mu_i, \Sigma_i), (\mu_j, \Sigma_j)) > threshold = 2 \Rightarrow no \ match$$

$$if \ crit_{Maha}((\mu_i, \Sigma_i), (\mu_j, \Sigma_j)) \leq threshold = 2 \Rightarrow s_i \ and \ s_j \ match$$

$$(4.9)$$

I performed several tests to determine an empirical match-threshold. The determined threshold represents the rejection of a match if the sheet centers are further away than twice the paired Mahalanobis distance. This empirical threshold is a parameter of the matching procedure and produced good results with the datasets studied so far. For objects of different complexity a different threshold might be more appropriate.

4.3.2.3 Computation of the common branching topology

The common branching topology is computed iteratively. First, the topology of the average object is chosen as the initial guess for the common branching topology. For each object set member, the algorithm compares its branching topology with the current common branching topology until the object set is fully processed. The comparison is performed by computing the paired Mahalanobis distance between the medial sheets to identify corresponding sheets. Those sheets that do not correspond to any sheet in the current common branching topology are added to it. This means that every sheet of the whole object set is matched by at least one of the sheets of the final common branching topology. The common topology is a set of medial sheets originating from various members of the object set mapped into the common spatial frame.

The algorithm of the common topology computation is shown in Fig. 4.3.2.3 as pseudocode. The algorithm allows comparing medial sheets in an one-to-many fashion. This is desired since due to skeletal instabilities a single medial sheet in one object can have two or more corresponding sheets in another object. The matching procedure Match(sheet, commonTopo) computes the minimum of matching criterion (eq. 4.8) between this sheet and each of the sheets in the common branching topology.

For this algorithm, the object set processing order in the shape space is chosen as follows: The next object to be matched is the unprocessed object set member whose shape space location is the closest to the center of the shape space. If the shape space location of multiple members are equally close, the shape space location of the lowest eigenmode is taken.

In our experiments presented in chapter 5, the computed common branching topologies comprised only a few medial sheets. The common branching topology of the hippocampus-amygdala structure comprises 4 medial sheets. For the lateral ventricle and hippocampus structure, the common branching topologies are each a single medial sheet.

```
ComputeCommonBranchTopo(shapeSet)
{
  // Warp all shapes into common frame
   commonFrame = averageObject(shapeSet)
  warpShapeSet = warp(shapeSet, commonFrame)
  // Initialize common branching topology
   commonTopo = Empty
  foreach sheet in commonFrame
      Add(sheet,commonTopo)
   do until (warpShapeSet == Empty)
      object = getClosestToCenter(warpShapeSet)
      Remove(object,warpShapeSet)
      foreach sheet in object
         if ( NOT( Match(sheet, commonTopo))
             Add(sheet,commonTopo)
      end
  end
  return commonTopo
```

}

Figure 4.21: Schematic algorithm for computing the common branching topology.

4.4 Computation of the grid sampling for a common medial model

Figure 4.22: Computation of common m-rep model with minimal grid dimensions: The common m-rep model is computed as the common branching topology sampled by the minimal grid given a maximal approximation error in the shape space.

In the previous section, the scheme to compute a common branching topology from a population was presented. This section describes how the medial sheets of this common branching topology are sampled to create the common m-rep model with the minimal sampling.

M-reps sample the medial manifold of a sheet by a grid of medial atoms. The set of medial sheets and a set of grid parameters determine an m-rep model. I propose a grid sampling algorithm that is based on the medial axis of a medial sheet. This sampling algorithm is described in section 4.4.1. The computation of the m-rep model associated with the grid sampling is described in section 4.4.2. The computed m-rep is a good *initial* estimate to the m-rep description. An additional step is needed in order to get the appropriate m-rep description: the medial atoms are deformed to optimally fit the object boundary. This deformation is described in the section 4.5.

The grid dimensions are optimized to be minimal while the corresponding m-rep has a predefined maximal approximation error in the shape space. The approximation error is defined as the Mean Absolute Distance (*MAD*) of the m-rep implied boundary and the original boundary. The *MAD* error is normalized to make it independent of the object size (4.10). The normalization of the error can be done either relative to the individual object or relative to the population. E_i is the error relative to the individual object using the average radius r_{avg} of the object's skeleton. E_{pop} is the error relative to the population using the average radius over all skeletons of the population $r_{avg,pop}$.

$$E_{i} = MAD_{norm} = \frac{MAD}{r_{avg}}$$

$$E_{pop} = \frac{MAD}{r_{avg,pop}}$$

$$(4.10)$$

Section 4.4.3 describes the algorithm to compute the minimal grid sampling of the medial sheets of a single object. This algorithm is used to compute the m-rep model of the common branching topology in the common frame. This m-rep model is then checked whether it describes sufficiently the shape space. If this is not the case, the sampling is subsequently adjusted, as described in section 4.4.4.

4.4.1 Grid sampling of a single medial sheet

The grid sampling algorithm solves following problem: Given a medial sheet from the Voronoi skeleton and a set of m-rep grid dimensions n, m, how can we determine the grid samples for a most uniform grid on the medial sheet? The algorithm that I propose computes this sampling on the volumetric reconstruction from the medial manifold rather than on the Voronoi skeleton since efficient and well-tested algorithms exist for a wide range of image operations on volumetric representations. In this dissertation I chose to compute a most uniform grid in Euclidean space, whereas in the general m-rep theory m-rep atoms are sampled proportional to the local thickness.

Figure 4.23: Computing the sampled medial sheet axis. A: Rendering of the sheet. B: Overlay of the sheet boundary (purple) with the smoothed sheet (blue). C: Overlay of the sheet and the 1D skeleton of the smoothed sheet. D: Overlay of the sheet and the sampled sheet axis extracted from the skeleton.

In overview, the algorithm first smoothes the voxel sampling of the medial sheet at its boundary to form the *sheet image*. Then, the medial axis of the sheet is computed from the sheet image, and this axis is uniformly sampled. Next, the m-rep grid samples on the grid-edge are computed using the sampled axis and the curve forming the voxel sampling of the sheet boundary (the *sheet-boundary image*). Finally, the remaining grid samples are interpolated. These steps of the algorithm are visualized in the Figs. 4.23 and 4.24. The rest of this section describes these steps in more detail.

The first step of the sampling algorithm is a smoothing step that erodes the the initial sheet image with a spherical structuring element only at the locations of the sheet boundary. Erosion is not discussed here, and the reader is referred to [39] for details. The radius of the structuring element is proportional to the second largest axis of the covariance matrix of the sheet's Voronoi vertices (see section 4.3.2.2). This generates appropriate results for flat sheets. If a sheet isn't flat but shaped like a screw, for example, the radius has been chosen too large and important parts of the skeleton are removed. Such a case can be detected by visual inspection of the medial sheets and demands an adaptation of the radius of the structuring element.

From the smoothed sheet, I compute the 1D skeleton using an isotropic thinning procedure. The thinning algorithm is an extended version of the original parallel thinning algorithm described by Fu and Tsao [34]. After the parallel thinning step, the algorithm continues sequentially while keeping the line-ends. The thinning-skeleton is translated into a graph via a straightforward graph-compilation. The longest path is then computed from this graph. This longest path is divided uniformly into n samples and forms the sampled medial axis of the sheet.

The sampling algorithm determines the samples at the m-rep grid-edge using the sampled medial axis and the sheet-boundary image as follows. For every medial axis sample, the algorithm computes the 3D direction that is normal to the axis and that lies in the plane tangential to the medial surface. This direction is scaled by the distance of the axis sample to the sheet-boundary to compute the estimated locations of the grid-edge. The distance of the axis samples to the sheet-boundary is obtained from a distance map that is computed from the sheet-boundary image. The estimated locations of the grid-edge do not lie on the sheet-boundary. Thus, the algorithm projects the estimated grid-edge locations to the sheet-boundary using the sheet-boundary distance map another time. In summary, the grid-edge samples are computed as the closest sheet-boundary points of estimated locations on the directions normal to the medial axis.

The sampling algorithm is finished at this stage if the second grid dimension $m \leq 3$. Otherwise, intermediate samples need to be calculated. These samples are linearly interpolated along the lines connecting medial axis samples and grid-edge samples. The interpolated samples lie on the medial sheet only if it is fully flat. Thus, the interpolated samples are projected to the closest sheet locations using a distance map from the sheet image. The samples that lie on the grid-edge are projected to the sheet-boundary. As a final step, the medial sheet axis is discarded if m is even. The medial sheet axis is kept as part of the grid if m is odd.

Figure 4.24: Visualization of the sampling method. Starting from the sampled axis (top left, boundary in black, eroded boundary in purple, axis in red), the grid-edge (top right, blue) is estimated. The grid-edge is projected to the sheet-boundary (bottom right) and the remaining samples (violet) are interpolated.

The algorithm presented in this section computes the grid sampling from the voxel sampling of the medial sheet. Thus, the medial samples do not lie at the locations of the Voronoi vertices. These medial samples by themselves are inadequate for an m-rep description since they only posses the location property of an m-rep atom. No information about the radius, the frame or the angulation is known for these medial samples. The next section describes how all the m-rep atom properties can be computed for the grid samples via the closest Voronoi vertices and its generating points.

4.4.2 M-rep computation from a sampled medial sheet

The previous section described how the grid sampling of a medial sheet is computed, and this section determines the m-rep description from this sampling. First, the grid samples are bijectively projected to the closest Voronoi vertices of the medial sheet. Since the medial manifold is densely sampled with Voronoi vertices, this projection affects the sample locations only slightly. At the Voronoi vertice \mathbf{x}_{skel} , the additional information from the generating points ($\mathbf{p}_{gen_1}, \mathbf{p}_{gen_2}$) and the Voronoi neighborhood ($\mathbf{x}_{v_1} \dots \mathbf{x}_{v_4}$) can be used to estimate the m-rep atom properties (see eq. 4.11): position \mathbf{x}_{atom} , radius r_{atom} , frame \underline{F}_{atom} and angulation θ_{atom} . The frame and angulation properties are checked for consistency along both grid dimensions and are adjusted via linear interpolation if needed.

The properties of the m-rep atom $m_{i,j}$ with grid index i, j are computed as follows:

$$\begin{aligned} \mathbf{x}_{atom} &= \mathbf{x}_{skel} = \mathbf{x}_{skel}^{i,j}; \end{aligned} \tag{4.11} \\ r_{atom} &= \frac{||\mathbf{x}_{skel} - \mathbf{p}_{gen1}|| + ||\mathbf{x}_{skel} - \mathbf{p}_{gen2}||}{2} \\ \underline{F}_{atom} &= \{\mathbf{b}_{atom}, \mathbf{b}^{\perp}_{atom}, \mathbf{n}_{atom}\}, \text{ where} \\ \mathbf{b}_{atom} &= \frac{\{\mathbf{x}_{skel}^{i,j} - \mathbf{x}_{skel}^{i+1,j}\} + (\mathbf{x}_{skel}^{i-1,j} - \mathbf{x}_{skel}^{i,j})}{2}(\mathbf{n}_{atom}) \\ \mathbf{n}_{atom} &= \frac{1}{4} \sum_{k}^{4} (\mathbf{x}_{skel} - \mathbf{x}_{v_k}) \times (\mathbf{x}_{skel} - \mathbf{x}_{v_{(k+1)} \mod 4}) \\ \mathbf{b}^{\perp}_{atom} &= \mathbf{n}_{atom} \times \mathbf{b}_{atom} \\ \theta_{atom} &= \frac{1}{2} \sum_{k}^{2} \arccos \frac{||\mathbf{b}_{atom}||}{||\mathbf{x}_{skel} - \mathbf{p}_{gen_k}||} \end{aligned}$$

The equations in 4.11 are valid for all internal atoms. Medial atoms at the gridedge are computed differently since the generating points of these Voronoi vertices are generally ill-behaved. This usually results in a small angulation θ_{atom} and an unstable orientation \mathbf{b}_{atom} . These properties are adjusted as follows: \mathbf{b}_{atom} is directed outward orthogonal to the grid-edge in the tangential plane of the medial sheet; θ_{atom} is minimally $\frac{2\pi}{3}$.

4.4.3 Minimal grid sampling for a single object

I propose an extended version of a nonlinear-optimization algorithm that uses Evolutionary Schemes (ES) to find the optimal sampling for a set of medial sheets of a single object. The implemented (1+1)-ES algorithm handles the discrete nature of the optimization space and keeps a history of computed parameters since revisiting parameters is frequently done in a discrete optimization space. Further details of the (1+1)-ES algorithm are discussed in [67].

The sheets are treated by the optimization as being independent of each other, and the grid samplings of the sheets are optimized simultaneously. The value of the goalfunction for a set of grid dimensions incorporates the total number of medial samples n_{atoms} and the degree of approximation to the original boundary. The goal-function first computes the m-rep description as described in the previous sections. This m-rep is only an initial description. It is deformed to fit optimally to the object boundary (see section 4.5). The approximation error E_i of the deformed m-rep is the MAD_{norm} error (4.10). If this error is larger than a predefined maximal approximation error E_{max} then the goal-function value f_{opt} is penalized. Formally, f_{opt} is computed as follows:

if
$$E_i \leq E_{max} \rightarrow f_{opt} = n_{atoms} + MAD_{norm}$$
 (4.12)
otherwise $\rightarrow f_{opt} = (n_{figures} + MAD_{norm}) * C_{pen}$

The penalty constant C_{pen} was chosen to be 100, which is an appropriate value for all objects with a minimal grid sampling of less than 10x10.

The output of the optimization procedure is the minimal grid sampling whose associ-

Grid	Initial position			Final position		
	Reconstruction	MAD	E_{vol}	Reconstruction	MAD	E_{vol}
2x6		0.156	69%		0.068	89%
2x7		0.134	73%		0.054	91%
3x6		0.113	78%		0.043	94%
3x7		0.100	83%		0.045	94%
4x7		0.082	85%		0.039	95%
3x12		0.070	88%		0.040	96%
4x12		0.050	91%		0.033	97%

Figure 4.25: M-rep approximation errors of a hippocampus in the initial position (left) and in the final position after deformation (right). The reconstruction shows the m-rep (red lines), points of the implied boundary (dark blue dots) and the original boundary (light blue transparent). The Mean Absolute Distance MAD_{norm} and the volumetric overlap error are shown for both. Average radius r_{avg} is 2.6 mm. The approximation errors are low even for m-rep's with coarse grid samplings.

ated m-rep fits the object at least with the approximation error E_{max} . Any grid sampling with a smaller number of samples fits the object less well than E_{max} . Empirical values for E_{max} were determined in the range of 5% to 10% of the average radius through tests. For every object population, an adjustment of this value is recommended. The three examples presented in chapter 5 have slightly different E_{max} values.

4.4.4 Minimal grid sampling for an object population

The final m-rep must generate a minimally sampled m-rep model that fits every object in the object set at least with a maximal approximation error E_{max} . This is achieved in two steps. First, the m-rep model of the minimal sampling for the average object s_{avg} is computed as described in the previous section. Second, this m-rep model is checked whether it appropriately fits into all objects from the object set. If an object o_i of the object set has a larger E_i than E_{max} , the current m-rep model is not appropriate for the whole shape space and has to be adjusted. In this case, the algorithm computes a new m-rep model with a minimal sampling for the object o_i . This m-rep model becomes the current m-rep model, which has to be checked to appropriately fit the whole object set. After all members of the object set have been handled by the algorithm, then the current m-rep model is the common m-rep model sought.

4.5 Fitting an m-rep into a SPHARM object

In the previous sections I described how to compute a common m-rep model. Once such a common m-rep model is computed, I want to fit it to individual objects. This fit is done in two steps. First a good initial estimate is obtained, which is then refined in an optimization step to fit to the boundary. The initial estimate is computed by warping the m-rep model from the common frame into the frame of the individual object o_i using the SPHARM correspondence on the boundary. This is the inverse process of mapping o_i into the common frame, which is done when the common topology is computed (see section 4.3.2.1). The mapping is defined via a TPS warp of the boundary points. The warped m-rep model is located quite close to the final position, which allows a wellbehaved optimization.

Starting from the initial position, an optimization procedure changes the features of the m-rep atoms to improve the fit to the boundary. Local similarity transformations and rotations of the local angulation are applied to the medial atoms. The optimization is constrained by an additional prior on neighboring atoms. This prior guarantees the smoothness of the medial manifold. The weight of the prior is a parameter of the optimization and is empirically determined. The optimization procedure is reinitialized every time it slows down.

With the exception of a few contributions of mine, the optimization was designed and developed by other members of the MIDAG group under the guidance of Stephen Pizer. My main contributions were in the development of a calling scheme with automatic re-initialization and adjustment of the prior strength. The scheme uses the approximation error to the original boundary E_i to detect a slowing down due to an excessive contribution of the prior. Convergence is detected if, despite re-initialization, E_i stays constant. Also, the change in the approximation error ΔE_i is constrained to be maximally ΔE_i^{max} , which is chosen empirically at about 20% E_{max} . In the evolutionary optimization described in detail below, this results in a moderate upper limit for accepting an optimization step with worse E_i . For the case that $\Delta E_i > \Delta E_i^{max}$, the optimization is reinitialized at the previous step with an increased penalty weight. This calling scheme is not applicable in the general case of fitting an m-rep because it is based upon an exact description of the boundary to compute E_i .

The following paragraphs describe the details of the optimization procedure after Joshi [41]:

The optimization applies to the medial atoms $m_{i,j} = \{\mathbf{x}, r, \underline{F}, \theta\}$ local similarity transformations as well as rotations of the local angulation, $S_{i,j} = (\alpha, O, \mathbf{t}, \beta)_{i,j} \in [(\mathbb{R}^+ \times SO(3)) \times \mathbb{R}^3] \times [-\frac{\pi}{2}, \frac{\pi}{2}]$. The transformed medial atom is computed as follows,

$$m'_{i,j} = S_{i,j} \circ m_{i,j} = (\alpha_{i,j}O_{i,j}\mathbf{x}_{i,j} + \mathbf{t}_{i,j}, \alpha_{i,j}r_{i,j}, O_{i,j} \circ \underline{F}_{i,j}, \theta_{i,j} + \beta_{i,j})$$
(4.13)

A prior on the local atom transformations $S_{i,j}$ is induced based on the displacement of the implied boundary with an additional neighborhood dependent prior on the translations, guaranteeing the smoothness of the medial manifold. In keeping with the level of locality let B be the portion of the implied boundary affected by the atom $m_{i,j}$. The prior energy on the local transformation $S_{i,j}$ of the atom $m_{i,j}$ depends on corresponding points on the figural boundary y and on the deformed boundary y'. The prior energy becomes

$$\log P(S) = \left[-\int_{B_{i,j}} \frac{||y - y'||^2}{(\sigma r(y))^2} dy - \sum_{i,j} \sum_{n,m=-1}^{n,m=1} \frac{||\mathbf{t}_{i,j} - \mathbf{t}_{i+n,j+m}||^2}{||\mathbf{x}_{i,j} - \mathbf{x}_{i+n,j+m}||} \right],$$

where $\mathbf{t}_{i,j}$ is the translation component of the local transformation $S_{i,j}$. Association between points on the boundary y and the deformed boundary y' is made using the figural coordinate system (u, v, t) described below. The point y' is the point on the deformed model having the same coordinates as that of the original point y. The integral in the above prior is implemented as a discrete sum over a set of boundary points by defining a sampling of the (u, v, t) coordinate space and calculating the associated



Figure 4.26: Different stages of the m-rep fit for an individual object o_i . Two examples from the object set of a lateral ventricle population. On the left, the object is smaller and on the right it is larger than the common m-rep model. A: Common m-rep model in common frame. B: Common m-rep model in frame o_i . C: Warped m-rep model in frame o_i . D: Fitted m-rep model in frame o_i .

implied boundary before and after an atom deformation.

The continuous medial manifold of a figure, defined via a spline interpolation, is parameterized by (u, v), with u and v taking the atom index numbers at the discrete mesh positions. A parameter $t \in \{-1, 1\}$ designates the side of the medial manifold on which an implied boundary point lies. t varies continuously between -1 and 1 as the implied boundary point moves around the crest of the object from one side of the medial axis to another. For single figures, boundary correspondences are defined via the common parameterization (u, v, t). Positions in the neighborhood of the implied boundary are indexed by (u, v, t, \hat{d}) , where (u, v, t) is the parameterization of the closest point on the medially implied boundary and \hat{d} is the signed distance (interior = negative, exterior = positive) from the boundary in multiples of the local radius r of the medial point at (u, v).

Chapter 5

Clinical applications

5.1 Example 1: Hippocampus structures from a schizophrenia study

This section presents two cases of hippocampal shape that illustrate the intuitive representation of shape changes inherent in the m-rep description. In medial descriptions the local shape properties of location, orientation and thickness are separated. In this section, I primarily investigate the thickness property. Two different cases exploring shape asymmetry between the hippocampi of the left and the right brain hemisphere are presented. The hippocampus is a sub-cortical structure in the limbic system of the human brain. It is involved in the laying down and retrieval of long-term memory via interconnections with cortical regions of the brain. Asymmetry is defined via the interhemispheric plane. Volume measurements and medial axis length measurements demonstrate interhemispheric hippocampal asymmetry. These measurements do not provide localization of the detected asymmetry. A localized asymmetry analysis, however, can be computed using the m-rep description.

For the asymmetry analysis, the following steps were taken on a set of hippocampi. The right side hippocampi were mirrored at the interhemispheric plane. The SPHARM coefficients were determined and normalized for rotation and translation using the first order ellipsoid. Scaling normalization was not applied. I determined the average SPHARM object for each case. Since the number of objects was 2 for both cases, the object set comprises the average structure and the two individual objects, a left side and a mirrored right side object. All objects have a medial branching topology of a single medial sheet with a volumetric overlap error of more than 98%. Thus, the common topology is a single sheet. The computed minimal grid sampling has an E_i of less than 5% for all objects after the fitting procedure.

As will be demonstrated in the results, these hippocampal studies provide evidence that the observed asymmetry can be better understood via the medial description than by morphologic measurements like the volume or the medial axis length. A coarse scale description does not lessen the power of the analysis, and the asymmetry can be reliably detected and localized.

5.1.1 First hippocampus case

In Fig. 5.1, the objects of the first hippocampus case are displayed and asymmetry is clearly visible. Volume and medial axis length measurements indicate that the right hippocampus is larger than the left: $vol_{right} = 2184mm^3$, $vol_{left} = 2023mm^3$; $axis_{right} = 65.7mm$, $axis_{left} = 64.5mm$. These measurements, however, do not provide a localization of the detected asymmetry. Locality of asymmetry can be determined and visualized using the m-rep description.

I first present the medial axis of the medial sheet, which is a single column of m-rep atoms from the full m-rep model presented later. The right hippocampus is thicker over the full length of the axis, and the difference is pronounced in the middle part of the axis as one can see in Fig. 5.2. In order to relate this thickness information with the appropriate location, I chose to visualize it on the m-rep itself. Each medial atom is displayed by a sphere of size and color that is proportional to its thickness information. This kind of display can also be used to display the difference in thickness between



Figure 5.1: Visualization of the first hippocampus case. From left to right: Boundary, pruned Voronoi Skeleton with thickness coloring (same range for all objects), m-rep description as single axis and grid. Top to bottom: Average object, right (mirrored) and left hippocampus.

corresponding locations of the right and left hippocampus. The radius of the sphere is proportional to the absolute thickness difference and the color to the real thickness difference: $r_{diff} \sim |R - L|$; $col_{diff} \sim (R - L)$.



Figure 5.2: Visualization of thickness asymmetry along the medial axis (tail to head) for the first case: A: Thickness plot $\{\mathbf{r}_{right}, \mathbf{r}_{left}\}$. B: Plot of difference between right and left: $\mathbf{r}_{right} - \mathbf{r}_{left}$. C: M-rep with thickness information displayed as spheres. Radius and color are proportional to the corresponding thickness. D: Difference between right and left thickness (R - L) at medial atoms. Radius and color are proportional to the difference.

As a next step, I take into account the full grid of medial atoms and perform the same analysis as for the axis (see Figs. 5.4 and 5.3). One can clearly see that the right hippocampus is bigger over most parts of the object but the difference is pronounced in the middle part.



Figure 5.3: Visualization of thickness asymmetry in the first case: Plot of Difference between right and left thickness (R - L) at medial atoms. Thickness plotted along the 3 rows of the longitudinal grid direction (tail to head).



Figure 5.4: Visualization of thickness asymmetry in the first case: A,B: Thickness values displayed as spheres in the m-rep grid (A: right, B: left). Radius and color are proportional to the corresponding thickness. C: Difference between right and left thickness (R - L) at medial atoms. Radius and color are proportional to the difference.

5.1.2 Second hippocampus case

In Fig. 5.5, the hippocampi of the second case are displayed. It is very hard to see an asymmetry between the left and right side. Inspecting the volume measurement, we see that there is an asymmetry: $vol_{right} = 3318mm^3$, $vol_{left} = 3214mm^3$.



Figure 5.5: Visualization of second hippocampus case. From left to right: SPHARM boundary, pruned Voronoi Skeleton with thickness coloring (same range for all objects), m-rep description. Top to bottom: Average object, right (mirrored) and left hippocampus.

I performed the same thickness analysis on the second case as I have done for the first case. Looking at the 'asymmetry' m-rep in Fig. 5.6, we realize that the right hippocampus is thicker in the upper head and in the medial part. The left hippocampus is slightly thicker in the lower head and the tail. Further, we observe that the magnitude of the thickness asymmetry is about the same as in the first case. The length of the

medial axis also indicates a slight asymmetry: $axis_{right} = 72.8mm$, $axis_{left} = 72.2mm$.



Figure 5.6: Visualization of thickness asymmetry in the second case: A,B: Radius values displayed as spheres in the m-rep grid(A: right, B: left). C: Difference between right and left thickness (R - L) at medial atoms. Radius and color are proportional to the difference.

The two cases presented in this section show a similar locality of left/right thickness difference in the middle part. The second case also represents differences in the tail region. Due to the small number of studied cases, we can't draw any conclusions from these observations. A full study with a much larger number of cases would be needed to explore the statistical significance of these findings. However, the two cases serve well as an example of the methodology of my scheme and of the shape analysis.

5.2 Example 2: Hippocampus-amygdala complex from a schizophrenia study

This section presents my shape description scheme applied to hippocampus-amygdala objects from a pooled control-schizophrenia population (30 subjects). There is a left and a right hippocampus-amygdala for each subject (60 datasets). The data was segmented manually via outlining. The original MR and segmented data were provided by Ron Kikinis and Martha Shenton, Brigham and Women's Hospital, Harvard Medical School, Boston. The data was used to detect if the brain morphometry of schizophrenic patients is different from the morphometry of healthy subjects. In this dissertation, no shape analysis was performed on this dataset. It was purely used to test the capabilities of the scheme to deal with branching topologies comprising multiple medial sheets.

I decided to build the m-rep model on the left hippocampus-amygdala population. The training population and the study population are the same for the left hippocampusamygdala. Then the m-rep model's validity to represent the right hippocampus-amygdala was tested. Also, the m-rep model with the right hippocampus-amygdala as training population was computed and compared with the model from the left hippocampusamygdala.

The SPHARM coefficients were normalized regarding rotation and translation by the first order ellipsoid. The size of the objects was normalized with the individual volume. The first 6 eigenmodes λ_i of the PCA shape space contain 97% of the variability in the left population, so the shape space was defined as { $\bar{\mathbf{c}} \pm 2 \cdot \sqrt{\lambda_i}; i = 1...6$ }. The object set was computed by uniformly sampling 5 objects along each eigenmode axis (see Fig. 4.3).

The 3D Voronoi skeletons were computed from the PDM's and pruned. No member of the shape space object set had a branching topology comprising more than 5 medial sheets. In Fig. 4.19, the individual medial branching topology of several members of the object set are displayed together with the common medial branching topology, which consists of only four sheets. Every member of the object set had more than 98% volumetric overlap between its reconstruction and the original object. The minimal sampling was computed with a maximal $E_i = MAD_{norm} \leq 0.09$ in the shape space.

The m-rep description of all individual cases of the left population have an E_i error lower than 0.19. Despite the fact that only a small error was visually present, the MAD_{norm} error was high (> 0.15) when the structure was very thin. The population approximation error E_{pop} is the range between 0.11 and 0.035 for all individual cases. Fig. 5.7 displays the m-rep description of some sample objects with the E_i and E_{pop} errors. Also, the pruned Voronoi skeleton is shown for a visual comparison of the fine scale skeleton and the coarse scale m-rep.

The m-rep representation of the right hippocampus-amygdala objects was computed using the model computed from the left population. In order to do so, I first mirrored the right objects at the interhemispheric plane. The approximation error E_i is maximally 0.143 for the right population, and E_{pop} is maximally 0.10. The fit of the left m-rep model is slightly better for the individual right population than for the left population. This suggests that the 'left' m-rep model appropriately describes the left objects and the right objects.

I also computed the m-rep model from the right hippocampus-amygdala population with $E_i \leq 0.09$. As expected, the model is not the same but it is indeed similar (see Fig. 5.8). The following properties are different :

- Branching topology: The right model consists of 3 medial sheets, one sheet less than the left model. The 3 sheets of the right model each have a matching sheet in the left model. The right model is a subset of the left model in regard to the branching topology.
- Grid sampling dimensions: The sampling dimensions are equal for some of the sheets that are common to both the left and the right model. For all sheets the



Figure 5.7: M-rep description of left hippocampus-amygdala population: The common m-rep model applied to the average object (left), a member of the shape space object set (middle) and an individual case from the training population (right). The pruned Voronoi skeletons and the volumetric overlap between original and reconstructed object is shown. In the middle row, the warped m-rep and its E_i , E_{pop} errors are displayed ($r_{avg,pop} = 3.6mm$). In the bottom row, the final m-rep and corresponding E_i , E_{pop} errors are displayed.

grid sampling is similar as expected. The grid dimensions are (left vs right): sheet 1: 9x3 vs. 7x3; sheet2: 4x3 vs. 4x3; sheet3: 3x2 vs. 4x2; sheet 4: 3x3 vs. unmatched.

• M-rep atom properties: The m-rep atom properties are similar for the matching sheets. The properties are more similar for m-rep atoms in the center than for those at the grid edge.



Figure 5.8: Application of the scheme to the populations of the left and right hippocampus-amygdala. In the top row, the two different common branching topologies are visualized. In the bottom row, the resulting common m-rep models are displayed. The models are shown in the frame of each side's average object.

The m-rep description of all individuals from the left and right population using the right m-rep model was computed next. The approximation errors were compared with those of the previously computed left model. Table 5.1 shows the range of the approximation errors for each model and population. These error ranges and the differences in the m-rep models suggest that the two populations are not the same. A more detailed analysis of the two populations is necessary to confirm this. This analysis has not been done in this dissertation.

	E_i range	E_{pop} range
left model on left pop.	0.042 - 0.190	0.035 - 0.112
left model on right pop.	0.044 - 0.143	0.038 - 0.100
right model on left pop.	0.043 - 0.202	0.043 - 0.158
right model on right pop.	0.048 - 0.104	0.046 - 0.073

Table 5.1: Error ranges for E_i and E_{pop} for the models of the left and right hippocampusamygdala population.

The results in this section show that my proposed scheme can handle populations of objects with multi-sheet branching topology. The scheme computed for the right and left hippocampi structures similar m-rep models, but the few differences between the models suggest that the right and left hippocampi are not part of the same population.

5.3 Example 3: Lateral ventricles from a mono/dizygotic twin study

This section presents my scheme applied to a population of lateral ventricles, a fluid filled structure in the center of the human brain. The image data is part of a mono/dizygotic twin study and consists of 20 twin subjects (10 pairs, 5 monozygotic and 5 dizygotic pairs). The original brain images were provided by Daniel Weinberger, NIMH Neuroscience in Bethesda, Maryland. The segmentation of the ventricles was performed at the Neuro Image Analysis Lab, University of North Carolina. The segmenting method used a single gradient-echo channel with manual seeding for Parzen-window based nonparametric statistical classification. The segmented structures were preprocessed using a closing operation with a spherical structuring element of radius of two voxels. Some objects needed further manual preprocessing with the IRIS software to fill up holes in the structure in order to establish simply-connected objects with spherical topology.



Figure 5.9: Three-dimensional rendering of the skin surface (transparent) and the lateral ventricles.

I mirrored the left objects at the interhemispheric plane to perform analysis of shape asymmetry and shape similarity. The SPHARM coefficients were determined and normalized with respect to rotation and translation using the first order ellipsoid. Fig. 5.10 displays the lateral ventricles of all twin pairs. The first 8 eigenmodes λ_i of the shape space hold 96% of the variability in the population and thus the shape space is defined as { $\bar{\mathbf{c}} \pm 2 \cdot \sqrt{\lambda_i}$; i = 1...8}. The medial branching topologies in the object set varied between one to three medial sheets with an volumetric overlap of more than 98% for each object. The single medial sheet topology of the average object matched all sheets in the common frame since the matching algorithm allows one-to-many matches. Thus, the common medial topology was computed to be a single sheet.

The minimal sampling of the medial topology was computed with a maximal $E_i \leq 0.22$ and $E_{pop} \leq 0.10$ in the shape space. The average radius in the population $r_{avg,pop}$ is 2.26. The deformation in the first two principal components thinned the structure non-uniformly, which leads to a very small average radius for the corresponding object set members (see Fig. 5.11).

Finally, the application of the medial model to all individual cases of the population did not produce an E_i larger than 0.35 and an E_{pop} larger than 0.15, which is about 0.35mm. Most objects had a smaller approximation error, and only for the extremely 'shrunken' ventricles was such an error present. This is a very small error for a coarse scale description considering that the lateral ventricle is a thin, long structure and its medial axis length is about 115mm. Further, the original images used for segmentation of the individual structures have a voxel-size of $0.9375 \times 0.9375 \times 1.5$ mm. This means that, in regard to the original sampling, the individual m-reps are computed with subvoxel accuracy.

One twin pair is presented in more detail in the next section, and a difference in asymmetry between the twins of this pair is shown. Further, I present a group analysis comparing monozygotic (MZ) twins and dizygotic (DZ) twins.



Figure 5.10: Visualization of the lateral ventricles of all twin pairs (same color for pairs) scaled with the individual volume (correct relative size). Top row: Left side. Bottom row: Right side. Top row: MZ twins. Bottom Row: DZ twins.


Figure 5.11: Approximation error for the object set members along the first two eigenmodes. E_i is high for some objects despite a good visual approximation. E_{pop} (with $r_{avg,pop} = 2.26$) seems to be a better estimate of the approximation error.



Figure 5.12: Two examples of fitting the common m-rep model to the study population. Left: Case with worst approximation error in population. Right: Case with low approximation error. The implied m-rep boundary (dark blue) is overlayed with the original boundary (light blue).

5.3.1 Asymmetry and similarity analysis of a monozygotic twin pair

I examined a monozygotic twin pair more closely for an asymmetry and similarity analysis. The objects in this analysis were not normalized with respect to size. Fig. 5.13 clearly shows asymmetry, which is also confirmed by the volume measurements: twin 1 $L/R = 9365/12442mm^3$; twin 2 $L/R = 8724/6929mm^3$. An analysis of the thickness is shown in Fig. 5.14. It is clear that the asymmetry in the twin 1 is larger than in twin 2. Moreover, the similarity between the two twins is much higher on the left than on the right side. The location where the structures differ and how strong they do so is clearly visible from the visualization of the m-rep descriptions: the location of strongest difference is located in the atrium of the lateral ventricle. As for the hippocampus cases presented in chapter 5.1, we can't draw any conclusions from these results due to the extremely low number of studied cases. However, the analysis serves well as an example for the asymmetry and similarity analysis.



Figure 5.13: Visualization of the lateral ventricles of two MZ twins and its corresponding m-reps. Radius and color is proportional to the thickness. A clear asymmetry of the boundary display is shown in both cases.



Figure 5.14: Asymmetry (top) and similarity (bottom) analysis of the thickness: difference of thickness at medial atoms. Radius and color is proportional to the difference (same range for all objects). Maximal absolute differences: a) 1.5mm, b) 0.28mm, c) 1.7mm, d) 0.42mm. The first twin has a higher degree of asymmetry. The similarity of the left ventricles is considerably higher than the right ones.

5.3.2 Monozygotic (MZ) vs. dizygotic (DZ) twin pairs

This section describes a statistical group difference analysis to distinguish monozygotic (MZ) twins from dizygotic (DZ) twins. In this study, there are 5 pairs for each population. The population size is very small, so the observed effect must be quite large for the analysis to yield a significant result. Another study was performed by Bartley et al [10] on the same datasets with the goal of distinguishing the populations. They compared cortical gyral patterns and the total brain volumes. Both measures show significant differences between the MZ and DZ populations. In this section, I will show that my ventricular volume analysis did not yield any significant findings. However, using my new shape description scheme, I get significant difference of ventricle shape between MZ and DZ twins.

5.3.2.1 Volume analysis



Figure 5.15: Volume difference between lateral ventricles of twin pairs (sum of left and right volume) show a strong overlap, p-value = 0.15.

The lateral ventricles are fluid filled connected structures that allow fluid to flow from one side to the other. Thus, I studied in this experiment the sum of the log normalized volumes, assuming that the sum is an important descriptor: $vol_i = \log vol_{right_i} +$ $\log vol_{left_i}$. In order to study the twin pairs, the absolute volume difference is computed and analyzed: $\Delta vol_{T_{1,2}} = |vol_{T_1} - vol_{T_2}|$.

As shown in Fig. 5.15, there is hardly any difference or trend visible between the two populations since the volume measurements are overlapping. This confirms what Bartley measured on the same original images though she was using a different segmentation technique. The *p*-value for discriminating the two population is at 0.15 suggesting a trend that the two populations are different, but this finding is non-significant at 5% significance level.

5.3.2.2 Shape analysis via SPHARM

In the earlier section 3.1.5, I described the calculation of the Mean Squared Distance (MSD) between SPHARM objects. I assumed that similar objects that are aligned and volume-scaled have a lower MSD than dissimilar objects since their surfaces are much closer.

As a prerequisite for any shape similarity calculation, shapes have to be normalized with respect to a reference coordinate frame. Since I am interested in measuring shape differences, a normalization is needed to eliminate differences that are due to rotation, translation and magnification. Normalization of translation and rotation is accomplished by aligning the SPHARM objects via the first order ellipsoid. In order to normalize for magnification, an appropriate scaling method has to be defined. The choice of the scaling method depends on the task and the type of objects. The three possibilities considered were the following:

- A No scaling correction: The computation of shape differences without any scale normalization reveals differences between small and large objects even though they might have the same shape properties. Thus, the differences will reflect mixed values of both the shape differences and the size differences.
- B Longest ellipsoid axis: The longest axis of the first order ellipsoid is appropriate

for objects that have their main size difference along a dominant axis. Normalizing the elongation of such objects removes differences along that axis and emphasizes differences that are orthogonal to it.

C Object volume: The object volume is a measure that captures the whole object. In this example we found that there are volume differences but that they were not significant. Thus, creating a shape difference measure that is orthogonal in its nature to the volume measure has the potential to reveal information additional to size. The volume measurements can be incorporated later into a multivariate statistical analysis as an additional orthogonal feature.

I computed the \sqrt{MSD} difference between twins for each of the different scaling normalizations methods mentioned above. As in the volume analysis, the values of the left and right ventricles were combined by adding the MSD's. A *t*-test analysis was applied to test whether the two populations differ significantly. This means that the hypothesis was tested whether the populations have the same mean.

If the objects are not normalized for scaling, my analysis yields no significant difference between the two populations as shown in Fig. 5.16. A significant difference is observed if the objects are scaling normalized. The best p-value can be obtained if the objects are normalized with their individual volumes. The p-value is at 0.012, which suggests significance. This result demonstrates that the lateral ventricle of MZ twins are more similarly shaped than those of DZ twins.

The SPHARM analysis revealed that the two populations differ in their shape similarity despite their volume similarity. This result might be of clinical importance. For example, it might be of interest in the analysis of discordant MZ twin studies. A twin pair is said to be discordant for a disease, if one of the twins is sick while the other twin is healthy. If the discordant twins have less similarly shaped ventricles, this might indicate that the studied disease manifests itself also in brain shape changes. This hypothesis can't easily be tested in a non-related population as the biological variability would hide



Figure 5.16: Statistics for \sqrt{MSD} shape similarity between lateral ventricles of twin pairs (sum of left and right differences): The two populations show a better separation than in the volume analysis. Top left: No scaling. Top right: Scaling with longest axis of first order ellipsoid. Bottom left: Scaling with the individual volume. Bottom right: *p*-values for group differences using the 3 scaling methods.

the expected subtle differences and only reveal large group effects.

The SPHARM analysis does not yield the locations where the object populations differ. The m-rep shape analysis has the potential to provide locality as described in the next section.

5.3.2.3 Shape analysis via m-rep

Based on the encouraging results of the SPHARM analysis, the m-rep descriptions were computed on objects scaled with the individual volume. Additionally, I computed the statistics on the unscaled objects for comparison. The goal of the m-rep shape analysis is to pinpoint the location of maximal difference between the populations. These differences can manifest in different ways, each of them potentially at different locations: thickness changes as the result of local uniform growth; location changes as the result of bending, twisting and local non-uniform growth; combined thickness and location changes as an overall measure of change.

In order to compute a *p*-value for the case of the combined thickness and location measure, the Fisher linear discriminant axis was determined and the values were projected onto this axis. The Fisher axis is the most discriminating line of the two populations, assuming Gaussian distributions. The computed *p*-value will no longer be unbiased since the Fisher axis is computed and applied on the same data. This is called data snooping and creates a bias on the resulting statistics. An unbiased Fisher axis cannot be computed due to the small sample size for each population. Unbiased multivariate statistics would be more representative and stable, but this has not been done in this dissertation.

All three measures detect a higher level of similarity between MZ twins than between DZ twins at a significant p-value. The similarity is most pronounced in the combined thickness and location measure with a p-value of 0.0233. Using the joint MZ and DZ population, I also extracted an age and gender matched population of 10 non-related

pairs. The shape analysis shows that the similarity in MZ twins significantly differs from the similarity of non-related pairs. This is not the case for DZ twins and non-related pairs. The statistics are visualized in Figs. 5.17 and 5.18 with the corresponding p-values in table 5.2. The individual volume scaling procedure is necessary for a significant result, which is evident from the p-values and statistical plots of the unscaled objects shown in table 5.3 and Figs. 5.19 and 5.20.

Volume scaled objects	MZ/DZ	MZ/Other	DZ/Other
Pure thickness	0.0282	0.0510	0.8980
Pure location	0.0270	0.0175	0.9089
Fisher axis on thickness / location	0.0233	0.0144	0.9659

Table 5.2: Class statistics of the m-rep properties and its associated p-values. The objects were individually scaled to their volume.

Unscaled objects	MZ/DZ	MZ/Other	DZ/Other
Pure thickness	0.5279	0.2088	0.3474
Pure location	0.2046	0.0394	0.3875
Fisher axis on thickness / location	0.1968	0.0377	0.5342

Table 5.3: Class statistics of the m-rep properties and its associated *p*-values. No scaling was performed.

The statistics and *p*-values shown in the tables and plots above are computed in a global analysis using cumulative values. These values were computed by integrating the individual local differences over the whole object. A local statistical analysis is also possible in order to compute the locations, where the two populations differ the most. For these statistics, the atoms are assumed to be independent of each other. Thus the local analysis is performed for each medial atom individually. This viewpoint is not fully correct but reasonable in a preliminary analysis. A more thorough analysis should be done by computing the statistics using a regional kernel instead of doing it atom by atom. The local analysis was applied individually for the thickness and location difference measure. Additionally, the Fisher linear discriminant axis on the combined



Scaling by individual volume

Figure 5.17: Class statistics of the pure m-rep properties for objects individually scaled to their *volume*. The statistical plot is shown on the left and the corresponding estimated Gaussian distribution on the right (MZ = red, DZ = blue, non-related = black). Top row: statistics for changes in local thickness. Bottom row: statistics for changes in location. The difference between the MZ twins and the DZ twins or the non-related subjects is significant. There is no significant difference between DZ twins and non-related subjects (see corresponding table).



Scaling by individual volume

Figure 5.18: Class statistics of the combined m-rep properties for objects individually scaled to their *volume*. The combined location and thickness measures are displayed in a 2D feature space (top row) with quartile ellipsoids (MZ = red, DZ = blue, non-related = black). The Fisher linear discriminant axis is displayed on the right side as a purple line. The 1D feature space (bottom row) from the projection to the Fisher linear discriminant axis is shown on the bottom row. The statistical plot is shown on the left and the corresponding estimated Gaussian distributions on the right. The difference between the MZ twins and the DZ twins or the non-related subjects is significant. There is no significant difference between DZ twins and non-related subjects (see corresponding table).





Figure 5.19: Class statistics of the pure m-rep properties for *unscaled* objects. The statistical plot is shown on the left and the corresponding estimated Gaussian distribution on the right (MZ = red, DZ = blue, non-related = black). Top row: statistics for changes in local thickness. Bottom row: statistics for changes in location. The difference between the MZ twins and the non-related subjects is significant. There is no significant difference between MZ and DZ twins or between DZ twins and non-related subjects (see corresponding table).

No scaling



Combined thickness and location

Figure 5.20: Class statistics of the combined m-rep properties for *unscaled* objects. The combined location and thickness measures are displayed in a 2D feature space (top row) with quartile ellipsoids (MZ = red, DZ = blue, non-related = black). The Fisher linear discriminant axis is displayed on the right side as a purple line. The 1D feature space (bottom row) from the projection to the Fisher linear discriminant axis is shown on the bottom row. The statistical plot is shown on the left and the corresponding estimated Gaussian distributions on the right. The difference between the MZ twins and the non-related subjects is significant. There is no significant difference between MZ and DZ twins or between DZ twins and non-related subjects (see corresponding table).

measure was computed for each atom and the values were projected onto the Fisher axes.

The local analysis is visualized in Fig. 5.21. The locations of significant shape difference are not the same for the thickness and location feature. I also observe that the combined feature doesn't capture additional significant locations that are not captured by the two pure features. The view of quasi-orthogonal features for thickness and location is thus justified in this example.



Figure 5.21: Locations of significant difference between MZ and DZ twins. The differences are shown in the frame of the common m-rep model. The radius and color is inversely proportional to the *p*-value of the shape difference between the two population: $r, col \sim 1/p$. Top row: top view. Bottom row: side view.

Chapter 6

Discussion and conclusions

In this chapter, I first discuss the computational efficiency of the different steps of the model building implementation. Next, I discuss the stability of these steps. Then, I discuss the homology of the computed m-rep model. The conclusion of this dissertation are summarized next, followed by a list of items for future work.

6.1 Computation time

The computation time for the different algorithms were computed in several experiments with different populations of anatomical objects: single figure model for a lateral ventricle population; single figure model for a hippocampus population; four figure model for a hippocampus-amygdala complex population. The computation time depends on the processor type and the amount of available memory. The later is quite important in our computations since the methods are implemented to be computationally efficient by sacrificing memory efficiency. The values presented in Table 6.1 are computed on a processor of type SUN Ultra-Sparc II and enough memory to fit the whole process at all time. A memory size of 750 MB on a Sparc Ultra 10 station satisfies these conditions.

The tools 'IRIS' and 'AVS' were used for the preprocessing the data. The 'param' tool developed by C. Brechbühler is used for the parameter optimization. The remaining steps of the SPHARM computation are performed using Mathematica. All other computations are performed with the 'VSkelTool' program that I have developed in this dissertation. For the m-rep fitting process 'VSkelTool' calls library procedures in the 'pablo' tool, which is being developed within the MIDAG group at UNC.

	Step	Tool	Auto.	Approx. costs
Α	Preprocessing of segmented object	IRIS, AVS	No	30 min/object
В	Parameter optimization	param	Yes	30 min/object
С	SPHARM description	Mathematica	Yes	$10 \min/object$
D	A+B+C for population of 30 ob-	-	-	35 hrs/pop
	jects			
E	Shape space and object set	Mathematica	Yes	30 min/pop
F	PDM description	Mathematica	Yes	5 min/object
G	Inner Voronoi skeleton	VSkelTool	Yes	5 min/object
Η	Pruned/grouped Voronoi skeleton	VSkelTool	Yes	40 min/object
Ι	Warping object set	VSkelTool	Yes	5 min/pop
J	Common topology extraction	VSkelTool	Yes	$5 \min/pop$
K	E+F+G+H+I+J for 25 object set	-	-	21 hrs/pop
	members			
L	M-rep sampling of single sheet	VSkelTool	Yes	1 min/object
М	Fit of a single figure M-rep	VSkelTool	Yes	2 hrs/object
Ν	Optimal sampling for a single ob-	VSkelTool	Yes	90 hrs/pop
	ject			
Ο	Optimal sampling of object set	VSkelTool	Yes	300 hrs/pop
Р	D+K+O Total estimated time cost	_	-	350 hrs/pop

Table 6.1: Estimated computation time for all stages in the m-rep model computation scheme, measured per object or per population (pop). The costs are computed for a single figure m-rep model. The computation time is independent of the number of figures with the exceptions of steps L, M, N and O, which are scaled linearly with the number of figures. The computation platform is a SUN Sparc Ultra 10 workstation with 750 MB of memory (i.e., large enough to fit the process)

The high cost of computing the minimal m-rep sampling can be approached by using multiple workstations. The algorithms involved in computing the sampling continuously store intermediate values to the disk. This allows the computation to be continued at the point of the last stored intermediate value if the algorithm or the computer crashes. This can happen in case of electricity failure or memory problems. Since these intermediate values can also be accessed by other 'VSkelTool' processes through network transparent file systems, the computation time of the optimal m-rep sampling can be reduced linearly with the number of 'VSkelTool' processes running on different machines.

6.2 Stability

Shape space - All computations of the common m-rep model are based on the shape space determined via PCA. In my experiments PCA has shown to be a quite stable procedure that produces good results in leave-one-out experiments (see also section 7). The PCA shape space stabilizes the computation of the common m-rep model by removing shape variations due to noise. A leave-one-out analysis should be performed on the PCA with following computation of the m-rep model, which would take several weeks of computations. This has not been performed in this dissertation.

Common medial branching topology - In all my tests, the stability of the common medial branching topology was very good. The tested anatomical brain structures were the following ones: hippocampus-amygdala, hippocampus, lateral ventricles, putamen, globus pallidus and thalamus. The common branching topology depends strongly on the boundary correspondence. Although I have not experienced problems with the quality of the boundary correspondence, it is evident that for objects with a high degree of rotational symmetry the first order ellipsoid correspondence is not appropriate. In this case, the common branching topology couldn't be extracted reliably.

In my tests, the procedure has shown to be robust to the ordering of the object set. Changing the ordering results in the same graph properties of the branching topology. The originating objects of the medial sheets might change but the sheets and their spatial distributions remain similar. Suggested by my experiments, an ordering is likely to generate the same number of sheets and similar spatial distributions as other, different orderings (see also Fig. 6.1).

Minimal sampling computation - The borders of the Voronoi skeleton are sensi-



Figure 6.1: Computed common branching topology for two different orderings of the matching procedure. Top row: Display of the Voronoi vertices (sheets are color coded). Bottom row: List of objects in the shape space from which the medial sheets are originating. Left: Branching topology from the implemented ordering. Right: Branching topology from a (different) random ordering. The graph properties are the same in both cases. The spatial distributions of the sheets are similar.

tive to small perturbations on the boundary, unlike the center part of the skeleton, which is quite stable. Thus the properties of the sampled m-rep atoms are quite stable in the grid center but not as stable at the grid edges. The computation the grid dimensions for a single population is stable. My experiments suggest that the grid dimensions for a similar population are unlikely to be the same, but they are likely to be close. Thus the grid dimension computation can be considered to be at least not unstable.

M-rep fit procedure - The m-rep fit procedure starts from an position that is close to the final position because the boundary correspondence is used to compute the initial position. The change of the m-rep properties during the fit procedure is strongly constrained by a prior on neighboring atoms. This prevents the m-rep atoms to move freely on the medial sheet if the radial function is constant along all directions. The fit procedure is non-deterministic but due to the strong prior, the stability and reliability can be considered good.

Conclusion - A quantitative analysis on the stability has not yet been performed. If the common medial model for 2 similar populations (e.g., 2 different studies of the same structure regarding the same disease) is computed, I expect to compute very similar medial branching topologies, similar grid parameters, very similar m-rep properties in Figure 6.2: Schematic 2D visualization of the homology for the medial description (red) and its implied boundary (blue) in two cases (I + II). The common model (left) gets warped into every individual case (middle) and then deformed into the fine scale boundary (green). The skeleton implied by the boundary is shown in cyan. The final positions of the m-rep atoms do not have to be on the skeleton since the m-rep is a coarse scale description and the skeleton is a fine scale description. The homology strongly depends on the boundary correspondence.

the center and less similar m-rep properties at the edge of the grid. The model is constructed for a population only once, and its extraction is mainly deterministic and repeatable. Thus, in contrast to the stability of the extraction of a medial model from a single object with noise, I consider the stability of a medial model for a *single* population as stable.

6.3 Homology

The main influencing factors of establishing homologous properties for the m-rep atoms are schematically visualized in Fig. 6.2. First, the common model is highly influenced by the Voronoi skeletons and the common frame, which is based on the boundary correspondence. The warp of the common m-rep model into every individual case is purely based on the boundary correspondence. Thus, the properties of the m-rep atoms prior to the fit procedure is mainly determined by the 'warp'-influence of the boundary correspondence transferred to the Voronoi skeletons. In the m-rep fit procedure, the m-rep is adjusted to maximize a boundary match at its implied boundary while it is constrained by the initial position. A strong neighborhood dependent prior prevents the medial atoms from moving freely on the medial sheet. In summary, the established m-rep homology strongly depends on the boundary homology.

6.4 Conclusions

In this dissertation I presented a new approach to the description and analysis of shape in the presence of shape variability. The proposed description is based on the boundary SPHARM and the medial m-rep description. The generation of the m-rep description takes into account the shape variability of a set of training objects, which is a novel concept and a step towards a shape representation for natural objects. Using the mrep description, locally computed shape features can quantitate and visually illustrate asymmetry or similarity. Since a correspondence is given on both the boundary and the medial manifold, a statistical analysis can directly be applied.

The SPHARM description and thus also the derived m-rep is constrained to objects of sphere topology. Also, the SPHARM boundary correspondence has shown to be a good approach in the general case, but it has inherent problems in special cases presenting rotational symmetry.

The choice of a *fixed* topology for the m-rep description has the advantages of enabling an implicit correspondence for statistical analysis. However, a fixed topology m-rep cannot precisely capture the topology of an individual object. The determined individual m-rep is therefore always an approximation, which emphasizes my decision to provide a coarse scale m-rep description.

The medial representation is constrained by the assumption that the shape variability can be captured based on the shape space spanned by the principal component analysis of SPHARM. It is thus not guaranteed that I am able to describe pathological objects not represented in the shape space. However, such pathological objects can be detected by inspecting the approximation error. Also, the m-rep model computation is designed for objects whose major deformation eigenmodes of the fine-scale boundary incorporate the coarse scale deformations. The computed m-rep model will not be appropriate if the deformation eigenmodes comprise a significant component of fine-scale deformation.

The results on the hippocampus-amygdala population in chapter 5.2 show that the

choice of the training population is important. The subjects of the training population have to incorporate most of the shape variability in the desired population. If that is not the case, then the computed m-rep model will not appropriately describe the population. In the presented applications, I showed that the m-rep model computation is able to incorporate a large shape variability. Thus, an m-rep model can be computed for a population that in fact is a collection of sub-populations. For example, an m-rep model can be computed that incorporates the patient and control population instead of computing two separate m-rep models. The disadvantage of such a super-population is the increased dimensionality of the shape space because the super-population's variability is likely to increase the number of principle modes to capture 95% variability.

All parts of the presented scheme have been implemented, applied and tested. The scheme has been applied to populations of several structures of neurological interest: hippocampus-amygdala(60 cases), hippocampus(20), thalamus(56), globus pallidus(56), putamen(56) and lateral ventricles(40). However, I expect that highly complex objects like the cortex of a human brain would be hard to handle without further adaptation of the algorithms.

6.4.1 Scientific contributions of this dissertation

This section briefly summarizes the new developments and findings that this dissertation contributes to different fields of this multi-disciplinary research project.

- 1. I have developed new shape description scheme that incorporates prior statistical knowledge about the shape variability. The scheme is presented in chapter 3, and the methods are described in chapter 4. The shape description scheme is suitable for shape analysis as demonstrated in the presented applications in chapter 5.
- 2. This work is the first to *compute* a common medial branching topology for a population of objects. This common medial branching is necessary to deal with one

of the major disadvantages of medial descriptions: the sensitivity of the branching topology to even small shape variations. In section section 4.3 the methods for the branching topology computation are presented, and the complexity of the common branching topology is shown to be of the same magnitude as the individual branching topology. The computation of the common branching topology is shown in section 6.2 to be stable.

- 3. In regard to computational geometry issues of Voronoi skeletons, this dissertation presents in sections 4.3.1.3 4.3.1.7 a novel scheme that automatically prunes 3D skeletons with results superior to those published elsewhere. My experiments showed that a small number of skeletal sheets are necessary to describe even quite complex objects, a surprising and encouraging finding.
- 4. Medial representations are known to be sensitive to small boundary perturbations. This work presents in section 4.4 a medial sampling technique that together with the m-rep fit procedure allows dealing with this sensitivity as discussed in section 6.2.
- 5. I presented a new shape description scheme for shape analysis via incorporating prior knowledge about the shape variability. This description scheme allows new insights and paths of exploration in various fields of morphological research as demonstrated by the applications in chapter 5. The shape features thereby are meaningful and allow to answer questions that could not be answered before.
- 6. I have shown an example in which shape information carries information that is superior to volume measurements. This example, which is presented in section 5.3, showed that the lateral ventricles of monozygotic twins are significantly more similarly shaped than those of dizygotic twins. I was able to describe and measure the shape similarity. Further, I was even able to describe locality and type of the shape differences, features not accessible previously.

6.5 Future work

Although the scheme has been tested on a large set of images, more validation of the scheme needs to be done. A possible validation would have to include synthetic data of two or more populations that can be separated with respect to a local effect of variable size. Also, real datasets can further be used for validation. For example, the left and right hippocampus-amygdala populations, presented in section 5.2, can be used to build one joint population and compare the m-rep model of the joint model to the individual models.

The first application presented in section 5.1 can be extended to include not only two single subjects but rather the whole set of hippocampi in this schizophrenia study. The full dataset includes 82 datasets of which 26 are controls, 28 are treatment responsive patients and 28 are treatment non-responsive patients. The m-rep model should be build on the joint population of all 82 subjects. The computed individual m-rep descriptions can be used for shape analysis to investigate the differences in the three populations.

In the second application chapter 5.2 further analysis can be done on the hippocampusamygdala objects. An earlier analysis showed a significant difference between schizophrenia patients and controls in regard to the asymmetry index |L - R|/(L + R). This difference was detected in both the volume measurements and the SPHARM \sqrt{MSD} measure. Using the individual m-rep descriptions, we could investigate the locations of most significant differences.

It is evident that the statistical analysis that was presented in the application chapters can be considerably improved. Instead of using the raw m-rep properties, a set of new features can be computed that separate uniform/non-uniform growth from twisting and bending. This has been proposed by Yushkevich [79] on 2D m-reps. He also introduces methods of doing a local PCA do detect locations that mostly discriminate 2 populations. The extension of these methods to 3D m-reps is probably the most important future work. The computed m-rep models incorporates shape variability in the common branching topology and the medial grid dimensions. This is not the case for the medial atom properties, which are determined from the average object. The m-rep description of the individual objects could be improved if the atom properties would additionally incorporate shape variability. This would lead to a fully statistical m-rep model. Since the m-rep model has been fitted to all object set members in the shape space, a distribution for the medial atom properties is directly accessible.

Chapter 7

APPENDIX 1: Properties of the Principal Component Analysis (PCA)

Principal component analysis (PCA) is a computation of an efficient basis for a distribution that is described by a set of multivariate samples. The computation of the basis is such that the variability is concentrated in a small subset of the basis. This has 2 main effects. First, the basis directions are decorrelated given that the samples are from a single Gaussian population. Secondly, the distribution basis can be reduced by removing directions of small variability contribution.



Figure 7.1: Principal component analysis as basis rotation. q_1 , q_2 is the original basis, and q'_1 , q'_2 the rotated basis after decorrelation with PCA. Clearly a decorrelation is visible as well as an ordering of the basis directions by variability contribution.

Fig. 7.1 shows the PCA principle for a two-dimensional example of a Gaussian distribution. Obviously the axes q_1 and q_2 are correlated. PCA computes the new decorrelated axes q'_1 and q'_2 , which are orthogonal linear combinations of q_1 and q_2 . As much variability as possible is thereby represented in q'_1 . The data could be approximated by only regarding the q'_1 axis and thus reducing the dimensionality.

The principal components are computed from the empirical covariance matrix S of the training set, which is defined as

$$\boldsymbol{S} = \frac{1}{N-1} \sum_{i=1}^{N} d\mathbf{x}_i d\mathbf{x}_i^{T}, \qquad (7.1)$$

where $d\mathbf{x}_i = \mathbf{x}_i - \bar{\mathbf{x}}$ is the deviation of the *L*-dimensional sample \mathbf{x}_i from the arithmetic mean $\bar{\mathbf{x}} = \sum_{i=1}^{N} \mathbf{x}_i$.

The modes of variation $\mathbf{v}_{\mathbf{k}}$, $k = 1 \dots \min(L, N)$ are the unit eigenvectors of the matrix \boldsymbol{S} and defined by

$$\boldsymbol{S}\,\mathbf{v}_k = \lambda_k\,\mathbf{v}_k \tag{7.2}$$

and

$$\mathbf{v}_k^T \, \mathbf{v}_k = 1, \tag{7.3}$$

where the λ_k are the eigenvalues of the matrix S ordered so that $\lambda_k \geq \lambda_{k+1}$. Rewriting equation (7.2) to

$$(\boldsymbol{S} - \lambda_k \boldsymbol{I}) \, \mathbf{v}_k = 0 \,, \tag{7.4}$$

where I is the identity matrix, makes clear that a set of linear equations has to be solved. Note that the number of eigenmodes $(\lambda_k, \mathbf{v_k})$ of a matrix is equal to its rank. Accordingly, S has min(L, N) eigenmodes. The variance described by an eigenvector \mathbf{v}_k is equal to the corresponding eigenvalue λ_k . So the eigenvectors belonging to the largest eigenvalues describe the most significant modes of variability.

In shape analysis applications, the number of samples N is often smaller than the number of parameters L gathered in the parameter vectors \mathbf{x}_i . The eigenmodes of the empirical covariance matrix $\mathbf{S} = \frac{1}{N-1} \sum_{i=1}^{N} d\mathbf{x}_i d\mathbf{x}_i^T = \Delta \mathbf{X} \Delta \mathbf{X}^T$ can be obtained by

solving for the eigenmodes of a reduced covariance matrix. As the mean was subtracted from the \mathbf{x}_i , the matrix S has at most rank N - 1 and a maximum of N - 1 non-zero eigenvectors results from the PCA.

As a first step, one of the vectors in ΔX is dropped so that the remaining are linearly independent. Then, an orthonormal basis M of ΔX is obtained by applying a Gram-Schmidt procedure that takes an arbitrary basis and generates an orthonormal one.

A reduced covariance matrix s is then given by

$$\boldsymbol{s} = \frac{1}{N-1} \boldsymbol{M}^T \Delta \boldsymbol{X} \Delta \boldsymbol{X}^T \boldsymbol{M} = \frac{1}{N-1} \boldsymbol{M}^T \boldsymbol{S} \boldsymbol{M} .$$
(7.5)

Solving the reduced eigensystem for each eigenvector \mathbf{v}_{si} ,

$$\mathbf{s}\,\mathbf{v}_{\mathbf{s}i} = \lambda\,\mathbf{p}_{\mathbf{s}i},\tag{7.6}$$

the (N-1)-dimensional \mathbf{v}_{si} yield the *L*-dimensional eigenvectors through

$$\mathbf{v}_i = \boldsymbol{M} \mathbf{v}_{\boldsymbol{s}i},\tag{7.7}$$

which are gathered in the eigenvector matrix $V = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{N-1})$.

Most of the variation can usually be explained by a relatively small number of eigenmodes t. When the number of training samples N is smaller than the number of parameters, then all variation in the training set can be described by t = N - 1 eigenmodes. Considering a smaller number t of eigenmodes, they describe a proportion λ_t of the total variance of all variables

$$\lambda_t = \sum_{k=1}^t \lambda_k . \tag{7.8}$$

The number t of selected eigenmodes is thereby chosen considering the overall proportion of variance reflected in the selected eigenmodes or by selecting eigenmodes with eigenvalues above a given minimum. Having chosen t, any object in the training set can be approximated by a weighted sum of the first t eigenmodes and the mean object $\bar{\mathbf{x}}$.

$$\mathbf{x} = \bar{\mathbf{x}} + \boldsymbol{V}_t \mathbf{b}_t \,, \tag{7.9}$$

where $\mathbf{b}_t = (b_1, b_2, \dots, b_{t-1}, b_t)^T$ is the weight vector, and $\mathbf{V}_t = (\mathbf{v}_1 \dots \mathbf{v}_t)$ is the eigenvector matrix.

Bibliography

- N. Amenta, M. Bern, and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. In SIGGRAPH 1998, Computer Graphics Proceedings, pages 415–421, 1998.
- [2] N.C. Andreasen, L. Flashmann, M. Flaum, S. Arndt, V. Swayze II, D. O'Leary, J.C. Ehrhardt, and W.T.C. Yuh. Regional Brain Abnormalities in Schizophrenia Measured with Magnetic Resonance Imaging. JAMA, 272(22):1763–1792, December 1994.
- [3] D. Attali and A.Montanvert. Computing and simplifying 2d and 3d continuous skeletons. Comp. Vision and Image Under., 67(3):261–273, September 1997.
- [4] D. Attali and A. Montanvert. Semi-continuous skeletons of 2d and 3d shapes. In Aspects of Visual Form Processing, World Scientific, pages 32–41, 1994.
- [5] D. Attali, G. Sanniti di Baja, and E. Thiel. Skeleton simplification through non significant branch removal. *Image Proc. and Comm.*, 3(3-4):63–72, 1997.
- [6] J. August, K. Siddiqi, and S. Zucker. Ligature instabilities in the perceptual organization of shape. In *IEEE Comp. Vision and Pattern Rec.*, 1999.
- [7] J. August, A. Tannenbaum, and S. Zucker. On the evolution of the skeleton. In Int. Conf. on Comp. Vision, pages 315–322, 1999.
- [8] R. Bajcsy and F. Solina. Three dimensional object representation revisited. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 231–239, 1987.
- [9] D. H. Ballard and C. M. Brown. Computer Vision. Prentice-Hall Inc., Englewood Cliffs, New Jersey, 1981.
- [10] A. Bartley, D. Jones, and D. Weinberger. Genetic variability of human brain size and cortical patterns. *Brain*, 120:257–269, 1997.
- [11] T.O. Blum. A transformation for extracting new descriptors of shape. In *Models* for the Perception of Speech and Visual Form. MIT Press, 1967.
- [12] J.D. Boissonnat and P. Kofakis. Use of the delaunay triangulation for the identification and the localization of objects. In *IEEE Comp. Vision and Pattern Rec.*, pages 398–401, 1985.
- [13] F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pat. Anal. and Machine Intel.*, 1(6), June 1989.
- [14] F.L. Bookstein. Morphometric Tools for Landmark Data: Geometry and Biology. Cambridhe University Press, 1991.

- [15] F.L. Bookstein. How to produce a landmark point: the statistical geometry of incompletely registered images. In SPIE, volume 2573, 1995.
- [16] F.L. Bookstein. Landmark methods for forms without landmarks: Localizing group differences in outline shape. In Math. Methods in Biomed. Image Analysis, 1996.
- [17] F.L. Bookstein. Shape and the information in medical images: A decade of the morphometric synthesis. In *Math. Methods in Biomed. Image Analysis*, 1996.
- [18] G. Borgefors, I. Ragnemalm, and G. di. The Euclidean distance transform: Finding the local maxima and reconstructing the shape. In SCIA Conf, pages 974–981, 1991.
- [19] J.W. Brandt and V.R. Algazi. Continous skeleton computation by Voronoi diagram. Comp. Vision, Graphics, Image Proc.: Image Under., 55(3):329–338, 1992.
- [20] C. Brechbühler. Description and Analysis of 3-D Shapes by Parametrization of Closed Surfaces. 1995. Diss., IKT/BIWI, ETH Zürich, ISBN 3-89649-007-9.
- [21] C. Brechbühler, G. Gerig, and O. Kübler. Parametrization of closed surfaces for 3-D shape description. CVGIP: Image Under., 61:154–170, 1995.
- [22] C.A. Burbeck, S.M. Pizer, B.S. Morse, D. Ariely, G. Zauberman, and J Rolland. Linking object boundaries at scale: a common mechanism for size and shape judgements. *Vision Research*, 36:361–372, 1996.
- [23] A. Caunce and C.J. Taylor. 3D Point Distribution Models of the Cortical Sulci. In Sixth International Proceedings on Computer Vision, pages 402–407, Bombay, India, January 1998. IEEE, Narosa Publishing House.
- [24] G.E. Christensen, R.D. Rabbitt, and M.I. Miller. 3d brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39:209–618, 1994.
- [25] T. Cootes, C. Beeston, G.J. Edwards, and C.J. Taylor. A unified framework for atlas matching using active appearance models. In *Proc. of Information Processing* in Medical Imaging, IPMI'99, pages 322–333, 1999.
- [26] T.F. Cootes and C.J. Taylor. Active Shape Models 'Smart Snakes'. In British Mach. Vision Conf., pages 266–275. Springer-Verlag, 1992.
- [27] J.G. Csernansky, S. Joshi, L.E. Wang, J. Haller, M. Gado, J.P. Miller, U. Grenander, and M.I. Miller. Hippocampal morphometry in schizophrenia via high dimensional brain mapping. *Proc. Natl. Acad. Sci. USA*, 95:11406–11411, September 1998.
- [28] C. Davatzikos. Spatial transformation and registration of brain images using elastically deformable models. *Computer Vision and Image Understanding*, 66(2):207– 222, 1997.

- [29] C. Davatzikos, M. Vaillant, S. Resnick, J.L Prince, S. Letovsky, and R.N. Bryan. A computerized method for morphological analysis of the corpus callosum. J. of Comp. Ass. Tomography., 20:88–97, Jan./Feb 1996.
- [30] H. Delingette. General Object Reconstruction based on Simplex Meshes. Technical Report 3111, INRIA, 1997.
- [31] L DeLisi, T William, X. Shu-hong, A. Hoff, M Sakuma, M. Kushner, G. Lee, K. Shedlack, A. Smith, and R. Grimson. A prospective follow-up study of brain morphology and cognition in first episode schizophrnic patients: Preliminary findings. *Biological Psychiatry*, 38:349–360, 1995.
- [32] I. Dryden. Highly resistant shape analysis. Technical report, Dep. Statistics, Univ. of Chicago, June 1997.
- [33] T. Fletcher, S. Pizer, A. Thall, and G Gash. Shape modeling and image visualization in 3d with m-rep object models. Technical report, UNC, Computer Science, 2000.
- [34] K. Fu and Y. Tsao. A parallel thinning algorithm for 3-d pictures. Comp. Graphics and Image Proc., 17:315–331, 1981.
- [35] Y. Ge, C. Maurer, and J. Fitzpatrick. image registration using the iterative closest point algorithm with a closestpoint transform. In *SPIE*, volume 2710, pages 358– 367, 1996.
- [36] P. Giblin and B. Kimia. A formal classification of 3d medial axis points and their local geometry. In *IEEE Comp. Vision and Pattern Rec.*, pages 566–573, 2000.
- [37] P. Golland and W.E.L. Grimson. Fixed topology skeletons. In Int. Conf. on Comp. Vision, 1999.
- [38] P. Golland, W.E.L. Grimson, and R. Kikinis. Statistical shape analysis using fixed topology skeletons: Corpus callosum study. In *Inf. Proc. in Med. Imaging*, pages 382–388, 1999.
- [39] R.M. Haralick. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):532–550, 1987.
- [40] S. Joshi, M Miller, and U. Grenander. On the geometry and shape of brain submanifolds. *Pattern Recognition and Artificial Intelligence*, 11:1317–1343, 1997.
- [41] S. Joshi, S. Pizer, T. Fletcher, A. Thall, and G. Tracton. Multi-scale deformable model segmentation based on medial description. In *Inf. Proc. in Med. Imaging*, 2001. accepted for publication.
- [42] G. Kelemen, A.and Székely and G. Gerig. Elastic model-based segmentation of 3d neuroradiological data sets. *IEEE Trans. Med. Imaging*, 18:828–839, October 1999.

- [43] B. Kimia, A. Tannenbaum, and S. Zucker. Shape, shocks, and deformations i: The components of two-dimensional shape and the reaction-diffusion space. Int. J. of Comp. Vision, 15:189–224, 1995.
- [44] A. Kotcheff and C. Taylor. Automatic construction of eigenshape models by genetic algorithm. In *Inf. Proc. in Med. Imaging*, pages 1–14, 1997.
- [45] L. Lam, S. Lee, and C. Suen. Thinning methodologies: A comprehensive survey. *IEEE Trans. Pat. Anal. and Machine Intel.*, 14:869–885, 1992.
- [46] S. Loncaric. A survey of shape analysis techniques. Pattern Recognition, 31(8):983– 1001, 1998.
- [47] S. G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pat. Anal. and Machine Intel.*, 11(7):674–693, 1989.
- [48] R. McCarley, C. Wible, M. Frumin, Y. Hirayasu, J. Levitt, I. Fischer, and M. Shenton. Mri anatomy of schizophrenia. *Biological psychiatry*, 1999.
- [49] M. Miller. Algorithms and codes for refining and linking of brain data. EnVision, 15(3):4–5, July-September 1999.
- [50] M. Miller. Brain Warping, chapter Large deformation fluid diffeomorphisms for landmark and image matching, pages 115–131. Academic Press, 1999.
- [51] M. Näf. Voronoi Skeletons: a semicontinous implementation of the 'Symmetric Axis Transform' in 3D space. PhD thesis, ETH Zürich, Communication Technology Institue IKT/BIWI, 1996.
- [52] A. Pitiot, P.M. Thompson, and A.W. Toga. Spatially and temporally adaptive elastic template matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1999.
- [53] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Trans. Med. Imaging*, 18:851–865, Oct. 1999.
- [54] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. Numerical Recipes in C. Cambridge Univ. Press, 2 edition, 1993. http://www.ulib.org/ webRoot/Books/Numerical_Recipes/bookcpdf.html.
- [55] A. Rangarajan, H. Chui, and F. Bookstein. The softassign procrustes matching algorithm. In *Inf. Proc. in Med. Imaging*, number 1230 in LNCS, pages 29–42, 1997.
- [56] M. Schmitt. LNCS:Geometry and robotics, volume 391, chapter Some Examples of Algorithms Analysis in Computational Geometry by Means of Mathematical Morphological Techniques, pages 225–246. Springer Verlag, 1989.

- [57] M.E. Shenton, C.G. Wible, and RW McCarley. Brain Imaging in Clinical Psychiatry, chapter MRI studies in schizophrenia, pages 297–380. Marcel Decker Inc., 1997.
- [58] F. Y. Shih and C. C. Pu. A skeletonization algorithm by maxima tracking on Euclidean distance transform. *Pattern Recognition*, 28:331–341, 1995.
- [59] K. Siddiqi, A. Ahokoufandeh, S. Dickinson, and S. Zucker. Shock graphs and shape matching. In Int. Conf. on Comp. Vision, January 1998.
- [60] K. Siddiqi, A. Ahokoufandeh, S. Dickinson, and S. Zucker. Shock graphs and shape matching. Int. J. Computer Vision, 1(35):13–32, 1999.
- [61] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. Zucker. The Hamilton-Jacobi skeleton. In Int. Conf. on Comp. Vision, pages 828–834, 1999.
- [62] K. Siddiqi, B. Kimia, S. Zucker, and A. Tannenbaum. Shape, shocks and wiggles. *Image and Vision Computing Journal*, 17:365–373, 1997.
- [63] K. Siddiqi, A. Tannenbaum, and S. Zucker. Hyperbolic smoothing of shapes. In Int. Conf. on Comp. Vision, 1998.
- [64] R. Sprengel, K. Rohr, and H. Stiehl. Thin-plate spline approximation for image registration. In Internat. Conf. of the IEEE Engineering in Medicine and Biology Society, 1996.
- [65] L.H. Staib and J.S. Duncan. Boundary Finding with Parametrically Deformable Models. *IEEE Transactions on Pattern Analysis and Machine Intelligece*, 14(11):1061–1075, November 1992.
- [66] L.H. Staib and J.S. Duncan. Model-based Deformable Surface Finding for Medical Images. *IEEE Trans. Med. Imaging*, 15(5):1–12, 1996.
- [67] M. Styner and G. Gerig. Evaluation of 2d/3d bias correction with 1+1esoptimization. Technical Report 179, Image Science Lab, ETH Zürich, 1997.
- [68] M. Styner and G. Gerig. Hybrid boundary-medial shape description for biologically variable shapes. In Proc. of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA) 2000, pages 235–242, June 2000.
- [69] M. Styner and G. Gerig. Medial models incorporating shape variability. In *Inf. Proc. in Med. Imaging*, 2001. accepted for publication.
- [70] G. Subsol, J.P. Thirion, and N. Ayache. A scheme for automatically building threedimensional morphometric anatomical atlases: application to a skull atla. *Med. Image Analysis*, 2(1):37–60, 1998.
- [71] G. Székely. Shape characterization by local symmetries, 1998. Habilitation, ETH Zürich.

- [72] G. Székely, A. Kelemen, Ch. Brechbühler, and G. Gerig. Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models. *Medical Image Analysis*, 1(1):19–34, 1996.
- [73] H. Tagare. Non-rigid curve correspondence for estimating heart motion. In Inf. Proc. in Med. Imaging, pages 489–494, 1997.
- [74] H. Tagare, D. O'Shea, and A. Rangarajan. A geometric criterion for shape-based non-rigid correspondence. In *Int. Conf. on Comp. Vision*, pages 434–439, 1995.
- [75] H. Tek and B. Kimia. Curve evolution, wave propagation, and mathematical morphology. In Int. Conf. on Comp. Vision, June 1998.
- [76] H. Tek and B. Kimia. Symmetry maps of free-form curve segments via wave propagation. In Int. Conf. on Comp. Vision, 1999.
- [77] D. Thomson. On Growth and Form. Cambridge University Press, second edition, 1942.
- [78] A.W. Toga. Brain warping. Academic Press, 1999.
- [79] P. Yushkevich and S. Pizer. Coarse to fine shape analysis via medial models. In Inf. Proc. in Med. Imaging, 2001. accepted for publication.