

Shape Spaces via Medial Axis Transforms for Segmentation of Complex Geometry in 3D Voxel Data

Jochen Abhau ^{*} Oswin Aichholzer [†] Sebastian Colutto [‡]
Bernhard Kornberger [§] Otmar Scherzer ^{*¶}

February 21, 2013

Abstract

In this paper we construct a shape space of medial ball representations from given shape training data using methods of Computational Geometry and Statistics. The ultimate goal is to employ the shape space as prior information in supervised segmentation algorithms for complex geometries in 3D voxel data. For this purpose, a novel representation of the shape space (i.e., medial ball representation) is worked out and its implications on the whole segmentation pipeline are studied. Such algorithms have wide applications for industrial processes and medical imaging, when data are recorded under varying illumination conditions, are corrupted with high noise or are occluded.

Keywords: medial ball representation, medial axis transform, image segmentation, Procrustes analysis, skin surfaces, EM algorithm

1 Introduction

Segmentation and detection of objects in image data is of major importance in several applications, such as for instance medical imaging, video surveillance and motion tracking. There are basically two different categories of segmentation algorithms,

^{*}Computational Science Center, University of Vienna, Nordbergstraße 15, A-1090 Wien, Austria, {jochen.abhau, otmar.scherzer}@univie.ac.at; present address of Jochen Abhau: Institute of Mathematics, University of Innsbruck, Technikerstraße 21a, A-6020 Innsbruck, Austria

[†]Institute for Software Technology, Graz University of Technology, Austria, oaich@ist.tugraz.at

[‡]Institute of Mathematics, University of Innsbruck, Technikerstraße 21a, A-6020 Innsbruck, Austria, Sebastian.Colutto@uibk.ac.at; present address of Sebastian Colutto: Department of Digitization and Digital Preservation, University of Innsbruck, Innrain 52, A-6020 Innsbruck, Austria

[§]Geom e.U. Softwareentwicklung, Brockmanngasse 15, A-8010 Graz, Austria office@geom.at

[¶]Radon Institute of Computational and Applied Mathematics, Austrian Academy of Sciences, Altenberger Straße 69, A-4040 Linz, Austria

namely supervised and unsupervised methods. Unsupervised algorithms do not make benefit out of object models or statistics on object models; this models will not be studied in this paper. Supervised methods and algorithms on the other hand make use of statistics on object models to be segmented. This is the type of algorithm which we investigate in this paper. Historically, unsupervised segmentation algorithms have been introduced first. However, there the segmentation results are poor if the data is noisy, has little contrast, or if the object is partially occluded. To cope with these difficult situations supervised segmentation has to be used.

In this paper we focus on supervised segmentation methods based on minimization of energy functionals, which have been proven to be very efficient. The shape prior model is a medial ball representation of low dimensionality. In this context it is common to differ between edge based (see [51]) and region based energy functionals (see [41]).

In the following we review some energy minimization segmentation approaches which incorporate shape prior information:

In [19] (this work is actually in 2D) contours are represented as simple, closed B-spline curves. The shape prior and statistic are computed as the mean and the principal components of the spline control points of the training shapes. For segmentation a Mumford-Shah like energy functional is supplemented by the Mahalanobis distance to the shape prior, which is calculated from the statistics. This approach is based on the assumption that the shapes are distributed according to a multivariate Gaussian distribution. Using kernel space techniques, a generalization to non-Gaussian distributed shapes is given in [18]. Applications to medical MR image segmentation have been studied in [15].

A level set approach for segmentation with shape priors has been studied in [38]. There shapes are associated with the associated signed distance functions. On the training distance functions of the shapes the actual statistics, which is used in the segmentation algorithm, is calculated. Shape statistics are implemented in a geodesic active contour evolution as maximum a posteriori estimator of the shape. By this approach geometrically and topologically complex shape priors can be considered. However, this representation requires large computational resources, and efficient implementations with narrow bands tend to create inaccuracies and artifacts in the segmentation results. In [23], the work of [38] has been improved concerning efficiency. The same shape prior strategy as in [38] has been used in combination with the unsupervised Chan-Vese energy minimization [13] in [50]. Further generalizations of [38] are given in [45] and [46].

In [16], shape statistical prior information is incorporated in a variational segmentation functional with an additional regularization term. A generalization of this work is presented in [27]. More recently, the paper [12] proposes a variational level set framework that takes into account shape prior information which combines gradient and region based segmentation.

A compact representation of shapes, called *discrete m-rep* (Medial atom representation) has been proposed in [25] and [43] (see also [47]). A discrete m-rep consists of *medial atoms* aligned on a regular grid structure, that approximates the medial axis defined by Blum in [9]. Medial atoms consist of a coordinate position in space, a radius and two boundary vectors. The actual surface associated with the m-rep can be

constructed using an enfolding B-spline of the surface that interpolates the boundary points and normals defined by the medial atoms. In contrast to discrete m-reps as defined above, *continuous* m-reps were introduced in [52]. There, the boundary of the object *and* the medial axis itself are approximated using B-splines. Discrete m-reps can be regarded as elements of a Riemannian shape manifold, which defines distances between m-reps by lengths of geodesics amongst them. An m-rep shape space is suitable for modeling shapes of very similar geometry and topology. Image segmentation with m-reps has been considered in [30, 42]. Mumford-Shah energy segmentation on an m-rep shape space has been presented in [17]. A limitation of using m-reps defined on a fixed grid structure is that the geometry of the object to be segmented has to be determined in advance.

Therefore, in this work, we use shapes which are defined via a fixed number of *medial balls* which are defined by its position and radius only. This shape representation does *not* require an underlying grid-structure as m-reps do. In fact this flexibility in combination with surface generation using skin surfaces (cf. Section 5.1) allows us to consider objects of complex geometry. To clarify and highlight the notation we give the following comparison of the two different approaches:

medial balls:	uses (position, radius)	[our work]
medial atom:	uses (position, radius, two boundary vectors)	[47]

Note that m-reps are defined by medial atoms on a regular grid that approximates the medial axis, while our representation is a collection of medial balls that sample the medial axis. Medial balls are a basic tool in Computational Geometry. Since in our approach we rely on medial balls rather than atoms, it is natural to work with established Computational Geometry algorithms for surface creation rather than algorithms usually applied for m-reps. Such combinations of algorithms – as for instance skin surfaces – are outlined below.

In this paper, we consider the construction of *medial ball shape spaces* and the computation of statistics of training shapes, which are then used for supervised energy minimization segmentation. We are given training shapes, which are represented as triangular surface meshes. They can either be provided by expert segmentation, or have been obtained from unsupervised segmentation (see e.g. [2, 17]). For the latter one can also imagine to use data of other imaging modalities, which have higher contrast or are less affected by noise.

The following scheme, which can as well serve as an *outline* of our work, is studied:

Scheme 1.1

Preprocessing (Shape Statistics)

- For each mesh a discrete medial axis transform

$$M = (x_1, \dots, x_k; r_1, \dots, r_k) \tag{1}$$

is computed. We refer to $c(M) = (x_1, \dots, x_k) \in \mathbb{R}^{3 \times k}$ as the centers of the k medial balls of M , and r_1, \dots, r_k are their radii. It is important for our

application that each discrete medial axis transform M_j consists of the same number k of medial balls $(x_i, r_i) \in \mathbb{R}^3 \times \mathbb{R}^+$. The algorithm developed in here is capable to do so, which is a novel aspect in the literature. For sake of simplicity, the discrete medial axis transform computed in this step will be called ball representation according to its definition (Section 2).

- A labeling of the ball representation M_j is computed by an EM-algorithm and the Kuhn-Munkres algorithm (Section 3).
- Using Procrustes analysis, a mean ball representation is computed from the labeled ball representations, and a Mahalanobis distance between ball representations is defined. These are the building blocks of our medial ball shape space. (Section 4).

Segmentation

- For segmentation the shape space is implemented as shape prior information in a simplified Mumford-Shah functional. The implied surface of the ball representations is constructed by a skin surface [21] (Section 5).

In Section 6 the full pipeline of statistics and segmentation is applied to synthetic data, and Section 7 concludes the paper.

2 Medial Axis Transform

In this section we consider the problem of approximating a number of bounded open sets $\Omega_1, \dots, \Omega_n \subset \mathbb{R}^3$ by sets B_1, \dots, B_n of approximate medial balls, where all sets B_i have the same cardinality. Stability of the approximation, described below, is of vital importance for the statistical analysis of medial axis representations. Our novel approach for calculating sets of approximate medial balls of the same cardinality for a class of objects combines three well known methods. First we utilize Voronoi diagrams to approximate objects by sets of (approximate) medial balls [7, 8]. As these sets will usually have a rather huge cardinality we then use set covering methods to obtain sufficiently small and stable subsets [5, 4]. Finally, we use a k -means clustering algorithm [31] to control the cardinality of the approximating sets in order to obtain sets of approximate medial balls of uniform cardinality.

We start by reviewing some basic facts of the medial axis transform.

2.1 Discrete medial axis transform

The following definitions of the medial axis and medial axis transform are standard and can for instance be found in [10].

Definition 2.1 ([10]) *The inner Medial Axis $M(\Omega)$ of a bounded open set $\Omega \in \mathbb{R}^3$ is the set of points $x \in \Omega$ which have at least two closest points in $\partial\Omega$, where $\partial\Omega$ denotes the boundary of Ω . The inner Medial Axis Transform $MT^{\text{in}}(\Omega)$ is the collection of maximal (with respect to inclusion) open balls centered at $M(\Omega)$ and included in Ω .*

Note that the cardinality of the inner Medial Axis Transform is in general infinite. Moreover the medial axis $M(\Omega)$ behaves unstable with respect to perturbations of high curvature of $\partial\Omega$. That is, the medial axis might contain branches that are far from being intuitive. The corresponding inner medial axis transform $MT^{\text{in}}(\Omega)$ expresses this unstable behavior by small balls near the surface of Ω . For these reasons we seek for an approximation of $MT^{\text{in}}(\Omega)$, which we define next.

Definition 2.2 ([10]) *The inner Discrete Medial Axis Transform $DMT(\Omega)^{\text{in}}$ is a finite set of open, approximate medial balls, which approximates the inner Medial Axis Transform $MT^{\text{in}}(\Omega)$.*

In the following we assume that the input object Ω is given by a set S_Ω of sample points of the boundary $\partial\Omega$ of Ω and a triangular mesh $T(S_\Omega)$, representing $\partial\Omega$. We first show how to approximate **one** three-dimensional object Ω by the union of its discrete medial axis transform $DMT(\Omega)^{\text{in}}$ in a stable way. We start with the well known Voronoi approach [7, 8].

Definition 2.3 *The Voronoi Diagram of a set of points S_Ω in \mathbb{R}^3 is a partition of \mathbb{R}^3 into (possibly unbounded) convex polyhedral regions, called Voronoi cells, such that each point $s_i \in S_\Omega$ has an associated Voronoi cell $v(s_i)$ with*

$$v(s_i) := \{x \in \mathbb{R}^3 : \|x - s_i\| \leq \|x - s_j\|, \forall i \neq j, s_i, s_j \in S_\Omega\}.$$

We first compute the Voronoi diagram of S_Ω . Then we extract for each sample point $s_i \in S_\Omega$ the inner pole point p_i , which is the vertex of the Voronoi cell $v(s_i)$ being farthest away from s_i and inside Ω . Finally, we construct for each inner pole point p_i a so called polar ball B_{p_i, ρ_i} centered at p_i with radius $\rho_i = \|s_i - p_i\|$. The set of all polar balls created in this way is the inner discrete medial axis transform $DMT(\Omega)^{\text{in}}$, and S_Ω is contained in the boundary of its union [7, 8]. See Figure 1(a) for an example with more than 10000 balls.

In the original approach a dense sampling S_Ω of a smooth surface $\partial\Omega$ of Ω is required in order to be able to distinguish inner Voronoi vertices from outer ones [7, 8]. To overcome this restriction we use, as mentioned above, the triangular surface mesh $T(S_\Omega)$ as additional input. This allows us to easily distinguish between inner and outer Voronoi vertices. Thus noise - always present in real-world data sets - and poor sampling quality do not affect the correctness of our inner/outer labeling, and correct operation of this step is ensured.

The centers of $DMT(\Omega)^{\text{in}}$ are close to $M(\Omega)$, see [6] for a quantitative analysis and a precise statement of that fact. This implies that $DMT(\Omega)^{\text{in}}$ might include small, surface near balls corresponding to unwanted features of $MT^{\text{in}}(\Omega)$. Moreover, $DMT(\Omega)^{\text{in}}$ approximates Ω by up to $|S_\Omega|$ balls. For our purposes the instability and the high cardinality of $DMT(\Omega)^{\text{in}}$ prevent its direct usage.

To avoid these disadvantages we apply a pruning algorithm to $DMT(\Omega)^{\text{in}}$ in order to extract a proper subset of $DMT(\Omega)^{\text{in}}$. The result will be a stable (we remove surface near balls which result from instability) and compact representation of Ω , cf.

Figure 1(d) for an example. Moreover the pruning step will give us control to obtain sets of approximate medial balls of predefined cardinality, see Subsection 2.2 for details. In the following we briefly describe the pruning approach of [4], see there for further details.

By the above construction each ball in $\text{DMT}(\Omega)^{\text{in}}$ has four sample points on its boundary but none in its interior. We enlarge each ball of $\text{DMT}(\Omega)^{\text{in}}$ by a sufficiently small constant $\varepsilon > 0$. Then we use a spatial search structure to find for each enlarged ball b_i all sample points from S_Ω which are now covered by b_i (typically tens or even hundreds of sample points are contained). Finally we use a set-covering algorithm to find an (almost) minimal subset $\text{DMT}(\Omega)_*^{\text{in}}$ of the enlarged balls whose union covers all sample points S_Ω . This set $\text{DMT}(\Omega)_*^{\text{in}}$ is the output of the pruning step.

Let us stress the fact that the applied set-covering algorithm to find a minimal subset of balls will favor balls that cover a large fraction of S_Ω and thus large areas of $\partial\Omega$ in order to make the cardinality of $\text{DMT}(\Omega)_*^{\text{in}}$, $|\text{DMT}(\Omega)_*^{\text{in}}|$, as small as possible. These balls are centered near stable parts of $\text{M}(\Omega)$. This implies that surface-near balls, which originate from small perturbations of high curvature of $\partial\Omega$ (recall the above discussion) are avoided. Thus stability of our approach is significantly improved by the set-covering algorithm, see also [4]. In addition the one sided Hausdorff distance from the union of the obtained compact discrete medial axis transform $\text{DMT}(\Omega)_*^{\text{in}}$ to the original object Ω is bounded by $O(\varepsilon)$.

2.2 Sets of approximate medial balls of uniform cardinality

As described in Section 1 our goal is to obtain sets of approximate medial balls of uniform cardinality for different input objects, $\Omega_i, i = 1, \dots, n$. So far we have shown how to compute a stable representation $\text{DMT}(\Omega_i)_*^{\text{in}}$ of the enlarged discrete medial axis transform for one fixed object Ω_i . We can control $|\text{DMT}(\Omega_i)_*^{\text{in}}|$ to some extent by the value ε by which we enlarge the polar balls.

For each single object Ω_i , the value ε is chosen such that the resulting set $|\text{DMT}(\Omega_i)_*^{\text{in}}|$ is at least as large as the desired uniform cardinality k , that is, $k \leq \min_i (|\text{DMT}(\Omega_i)_*^{\text{in}}|)$. This approach is justified by the fact, that for piecewise linear $\partial\Omega$, the quantity $|\text{DMT}(\Omega_i)_*^{\text{in}}|$ tends to infinity for $\varepsilon \rightarrow 0$. We then apply the k -means clustering algorithm [31] to the centers of each $\text{DMT}(\Omega_i)_*^{\text{in}}$ and get for every set $\text{DMT}(\Omega_i)_*^{\text{in}}$ exactly k clusters. For each cluster we choose as an representative ball the one whose center is closest to the center of that cluster. This results in n stable representations of $\Omega_1, \dots, \Omega_n$ by sets of balls B_1, \dots, B_n with uniform cardinality k .

For the remainder of the paper we will use the term *ball representation* to refer to the pruned inner discrete medial axis transform of uniform cardinality, denoted by M_1, \dots, M_n .

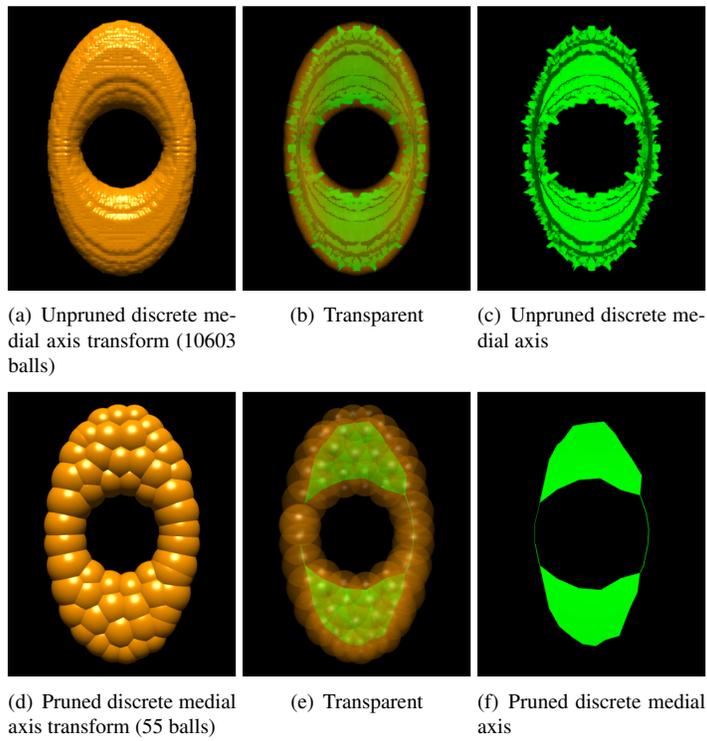


Figure 1: Discrete Medial Axis Transforms and Discrete Medial Axes

3 Labeling of the Ball Representations

In this section we use statistical algorithms for labeling of the ball representations. Given two ball representations $M = (x_1, \dots, x_k; r_1, \dots, r_k)$ and $\tilde{M} = (\tilde{x}_1, \dots, \tilde{x}_k; \tilde{r}_1, \dots, \tilde{r}_k)$, the labeling problem consists in determining a rearrangement of indices such that the two ball representations optimally match. Every matching involves a distance measure between ball representations. In view of the used definition of an Euclidean invariant shape space of ball representations considered below (see Section 4), we consider the following minimization problem for determining the optimal alinement:

$$\min_{\rho > 0, A \in SO(3), b \in \mathbb{R}^3, \pi \in \mathcal{S}^k} \sum_{i=1}^k \|\rho A x_{\pi(i)} - b - \tilde{x}_i\|^2 + \alpha (\rho r_{\pi(i)} - \tilde{r}_i)^2. \quad (2)$$

Here $SO(3)$ is the special orthogonal group of degree 3 (that is the group of rotations in \mathbb{R}^3 , which can be represented as unitary matrices with determinant one), \mathcal{S}^k is the symmetric group of degree k , and $\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^3 . The parameter α is a weighting parameter, providing a trade-off between matching of balls and radii. Since it is difficult to compute global minimizers of (2) efficiently, we use a two-step algorithm from [35], to compute approximate minimizers. This algorithm consists of the following two steps:

- An Expectation Maximization (EM) algorithm is used to compute a similarity transformation, i.e., a translation $b \in \mathbb{R}^3$, a rotation $A \in SO(3)$, and a scaling $\rho > 0$, which is close to an optimal one in (2). The algorithm is explained in Subsection 3.1.
- Keeping the similarity transform fixed, the Kuhn-Munkres algorithm is used to compute an optimal labeling $\pi \in \mathcal{S}^k$ in (2). This algorithm is explained in Subsection 3.2.

The algorithm outlined below also applies to variants of the minimization problem (2), such as

$$\min_{\rho > 0, A \in SO(3), b \in \mathbb{R}^3, \pi \in \mathcal{S}^k} \sum_{i=1}^k \|\rho A x_{\pi(i)} - b - \tilde{x}_i\|^2 + \alpha (\log(1 + (\rho r_{\pi(i)})^2) - \log(1 + \tilde{r}_i^2))^2.$$

For the sake of simplicity of notation we restrict attention to (2) and omit a discussion on general minimization problems.

3.1 EM-algorithm to compute the similarity transformation

EM-algorithms have been used successfully for solving a wide range of labeling problems [34]. In general, EM algorithms [40] consist in maximizing a likelihood function, which depends on parameters and hidden variables. There two successive steps are performed iteratively:

- An **Expectation** step, which consists in computation of the expected values of the hidden variables and
- a **Maximization** step, which is to optimize the likelihood function for the parameters.

We propose an algorithm which is closely related to [35]. However, the difference is that we take into account information on the radii of the medial balls in addition. The key idea of [35] is to compute the optimal similarity transformation of (2) relaxing the constraint $\pi \in \mathcal{S}^k$. This allows to apply the EM algorithm.

1. Relaxation step.

A mapping $\pi \in \mathcal{S}^n$ can be represented as a matrix $P \in \{0, 1\}^{k \times k}$ with entries

$$P_{j,i} = \begin{cases} 1 & \text{if } \pi(j) = i \\ 0 & \text{otherwise} \end{cases} .$$

The relaxation consists in assuming instead of \mathcal{S}^k *stochastic matrices*

$$\{P \in [0, 1]^{k \times k} : \sum_{i=1}^k P_{j,i} = 1 \text{ for all } j\} .$$

It is common to interpret $P_{j,i}$ as the probability that j is mapped to i . For medial ball representation this means that the ball (x_j, r_j) is mapped onto $(\tilde{x}_i, \tilde{r}_i)$.

Furthermore, a distribution $p = (p_1, \dots, p_k)$ is assumed over $\{1, \dots, k\}$. The number p_i represents the probability, that an arbitrary element of $\{1, \dots, k\}$ is mapped to i .

2. Distances of balls.

Under the condition $\pi(j) = i$ and for a given similarity transformation (ρ, A, b) , we assume that the transformed ball (x_j, r_j) is normally distributed around $(\tilde{x}_i, \tilde{r}_i)$, i.e., $\rho Ax_j + b$ is normally distributed with mean \tilde{x}_i and ρr_j is normally distributed with mean \tilde{r}_i . Concerning the radii, this assumption is slightly inaccurate since radii cannot become negative; we overcome this problem by choosing the variance σ sufficiently small, such that the 0.01 quantile is positive. Hence we define (for $i, j = 1, \dots, k$)

$$g_i(x_j, r_j) = (2\pi\sigma^2)^{-\frac{1}{2}} \exp\left(-\frac{1}{2\sigma^2} (\|\rho Ax_j + b - \tilde{x}_i\|_2^2 + \alpha(\rho r_j - \tilde{r}_i)^2)\right) . \quad (3)$$

The function $g_i(x_j, r_j)$ is large if (x_j, r_j) is mapped closely to $(\tilde{x}_i, \tilde{r}_i)$ and small otherwise.

The likelihood function for the parameters of the similarity transform (ρ, A, b) and the relaxed labeling (P, p) for given ball representations M and \tilde{M} is then given by the product probability

$$L(\rho, A, b, P, p) = \prod_{i,j=1}^k (p_i g_i(x_j, r_j))^{P_{j,i}} . \quad (4)$$

Taking the logarithm on both sides defines the log-likelihood function

$$l(\rho, A, b, P, p) = \sum_{i,j=1}^k P_{j,i} (\log p_i + \log g_i(x_j, r_j)) . \quad (5)$$

It is important to note, that if P is a permutation, then $p_i = 1/k$. Therefore in this case, maximizers of (5) are minimizers in (2) and vice versa. Following [35], enlarging the class of labelings from \mathcal{S}^k to stochastic matrices has little effect on the optimal similarity transform.

The advantage of the formulation (5), compared to (2), is, that (5) can be minimized efficiently with an EM-algorithm, where the variables ρ, A, b, p are regarded as parameters, and the entries of P are regarded as hidden variables of the likelihood function l . The EM algorithm for minimization of (5) reads as follows:

Algorithm 1 EM Maximization of l

Initialize $\rho = 1$, $A = id$ and $b = 0 \in \mathbb{R}^3$. Choose values for σ, α .

while l still increases **do**

E-Step: Compute the expected value of the relaxed permutation matrix P ,

$$P_{j,i} \leftarrow \frac{p_i g_i(x_j, r_j)}{\sum_{s=1}^k p_s g_s(x_j, r_j)} \quad (6)$$

By this assignment it is guaranteed that $\sum_{i=1}^k P_{j,i} = 1$ for all j .

M-Step: Maximize

$$\sum_{i,j=1}^k P_{j,i} (\log p_i + \log g_i(x_j, r_j)) \quad (7)$$

 over A, ρ, b and p_1, \dots, p_k .

 Compute the new value of l .

end while

Compared to different algorithms for minimizing (2) such as Markov Chain Monte Carlo Methods (MCMC), the output of the EM Algorithm is more sensitive to initial values of the parameters. This is no drawback in our case, since the ball representations are often already quite well aligned. However, the provided algorithm is assumed to be more efficient than MCMC.

3.2 Labeling

This subsection is concerned with computing an optimal labeling $\pi \in \mathcal{S}^k$ in (2).

It is instructive to reformulate the labeling problem as a combinatorial optimization problem.

1. A weighted graph is constructed of the following data:

- Nodes are the $2k$ balls $(x_1, r_1) \dots, (x_k, r_k)$ and $(\tilde{x}_1, \tilde{r}_1), \dots, (\tilde{x}_k, \tilde{r}_k)$.
- Let α be as in (2). The weights of the edges between the nodes are defined by

$$w_{j,i} = \|\rho Ax_j - b - \tilde{x}_i\|^2 + \alpha(\rho r_j - \tilde{r}_i)^2, \quad \text{with } \alpha \in [0, 1].$$

Here, (ρ, A, b) is the similarity transform computed in Subsection 3.1.

2. A permutation $\pi \in \mathcal{S}^k$ minimizing

$$\sum_{i=1}^k w_{\pi(i),i} \tag{8}$$

is computed.

For the fixed similarity transform transformation (ρ, A, b) , a minimizer $\pi \in \mathcal{S}^k$ of (2) is a minimizer of (8) and vice versa. For the solution of the matching problem (8) we apply the Kuhn-Munkres algorithm, which is a special case of the primal-dual algorithm in linear programming [3]. According to [49] the complexity of the algorithm is $O(k^3)$. In our applications, the high order of complexity of the algorithm is not an issue since it is only applied for medial ball representations of dimension k , which is in the range of 10 to 50.

4 Medial Ball Shape Space by Procrustes Analysis

In the following we perform a statistics of the ball representations M_1, \dots, M_n and establish a shape space, which we call *medial ball shape space*. Moreover, a distance between elements of the medial ball shape space is employed to compare them quantitatively.

4.1 Elements of the medial ball shape space

A shape is all the geometrical information that remains when location, scale and rotational effects are filtered out from an object (taken from [20, Chapter 1]). We consider this definition when constructing the medial ball shape space as the factor space of ball representations $M = (x_1, \dots, x_k; r_1, \dots, r_k)$ modulo similarity transformations. That is

$$\Sigma = \mathbb{R}^{3 \times k} \times (\mathbb{R}^{>0})^k / \sim_{ST} . \tag{9}$$

In other words, two medial ball representations $M = (x_1, \dots, x_k; r_1, \dots, r_k)$ and $\tilde{M} = (\tilde{x}_1, \dots, \tilde{x}_k; \tilde{r}_1, \dots, \tilde{r}_k)$ represent the same shape (or in other words are equivalent with respect to \sim_{ST}) if there exists a similarity transformation (ρ, A, b) such that

$$\rho A \tilde{x}_i + b = x_i \text{ and } \rho \tilde{r}_i = r_i \text{ for } i = 1, \dots, k \tag{10}$$

We emphasize that in the segmentation literature (see e.g. [17, 42, 38, 19]) frequently shape spaces are not defined as factor spaces. As a consequence the according shape space allows for different representations of the same shape.

Below we derive a segmentation model taking into account the original shape concept and coordinates of Kendall [33]. Compared to Bookstein [11], Kendall coordinates do not depend on the choice of a baseline i.e., two medial balls in our case. Therefore statistical analysis like PCA is not distorted. Since it is bulky to define a distance between equivalence classes we systematically choose representatives, for which we define a shape distance. The process is similar to [20, Chapter 4] where mid axis representations $M = (x_1, \dots, x_k)$ (without radii) have been considered.

1. A *mean* ball representation μ of the (labeled) ball representations M_1, \dots, M_n is defined by

$$\mu = \arg \min_{\tilde{\mu} = \{(\tilde{\mu}_1, \tilde{r}_1), \dots, (\tilde{\mu}_k, \tilde{r}_k)\}} \sum_{l=1}^n d(M_l, \tilde{\mu})^2 \quad (11)$$

where

$$d(M_l, \tilde{\mu})^2 = \min_{\rho > 0, A \in SO(3), b \in \mathbb{R}^3} \sum_{i=1}^k \|\rho A x_i^{(l)} - b - \tilde{\mu}_i\|^2 + \alpha (\rho r_i^{(l)} - \tilde{r}_i)^2 \quad (12)$$

The parameter α is the same as in Equation (2). The minimization problem (11) can be solved with an iterative algorithm, which is along the lines of [20, Chapter 5].

2. M is normalized with respect to *translation* by multiplying the ball centers of M with the $(k-1) \times k$ *Helmert submatrix* H (see [20, p.34] for the definition). The advantage in using the Helmert matrix for this task (compared to for instance centering the coordinates) is that the arising representation only consists of $k-1$ centers and that the Helmert matrices can be computed efficiently inductively. The resulting representation $T(M)$ is given by

$$T(M) = ((H[c(M)^T])^T; r_1, \dots, r_k) \in \mathbb{R}^{3 \times (k-1)} \setminus \{0\} \times (\mathbb{R}^{>0})^k. \quad (13)$$

3. $T(M)$ is normalized with respect to *scaling* by dividing the coordinates and radii by the Frobenius norm $\|c(T(M))\|$. The resulting scaled representation is denoted by $S(T(M))$.
4. $S(T(M))$ is normalized with respect to *Rotation* by minimizing the Frobenius norm

$$\min_{A \in SO(3)} \|c(S(T(\mu)))^T - c(S(T(M)))^T A\|. \quad (14)$$

Rotating $S(T(M))$ by the optimal $A \in SO(3)$ gives

$$R_{\mu}(S(T(M))) = \left(c(S(T(M)))A; \frac{r_1}{\|c(T(M))\|}, \dots, \frac{r_k}{\|c(T(M))\|} \right), \quad (15)$$

which is the representative of the class.

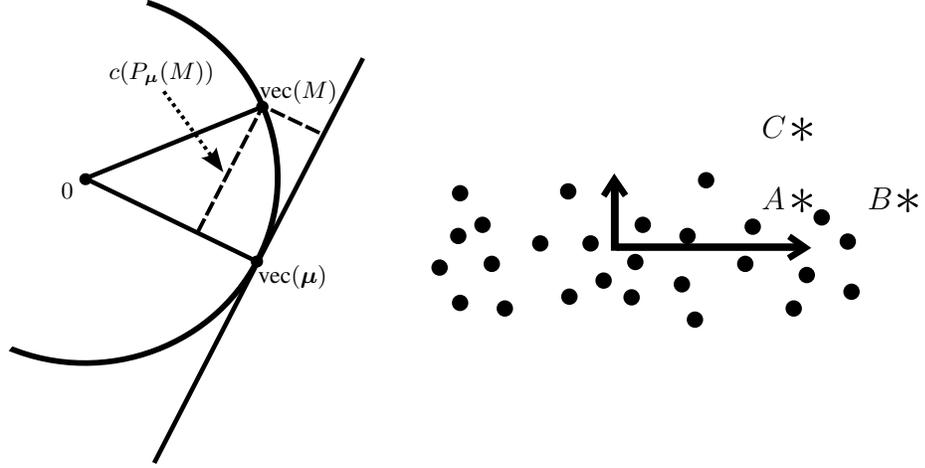


Figure 2: On the left hand side, Procrustes tangent coordinates of the centers of a ball representation are shown. Translated, scaled and rotated ball representations μ , M are on the circle. $\text{vec}(M)$ is orthogonally projected to the subspace spanned by $\text{vec}(\mu)$. The difference between the projection of $\text{vec}(M)$ and $\text{vec}(\mu)$ is $c(P_\mu(M))$. PCA and Mahalanobis distance are visualized on the right hand side. The data points vary in horizontal direction much more than in vertical direction. The Mahalanobis distance between A and B is smaller than the distance between A and C , which is the opposite measured in the Euclidean distance.

As shown in [33], the optimal rotation A in (14) can be calculated as follows: If

$$(c(S(T(\mu))))c(S(T(M)))^T = V\Lambda U^T \quad (16)$$

is the signed singular value decomposition of $(c(S(T(\mu))))c(S(T(M)))^T$, then

$$A = UV^T. \quad (17)$$

The mean shape μ and the representative are utilized to compute distances in Σ in the next Subsection 4.2.

We emphasize that in our approach we consider ball representations of the shapes. This should not be confused with m-reps, which have an additional grid structure. In the latter a principal geodesic analysis seems preferable to a PCA (see [47]).

4.2 Variability in the medial ball shape space

Standard techniques for capturing the variability of data are principal component analysis (PCA) and the Mahalanobis distance. Historically, these techniques have been defined in Euclidean space. A standard way to transfer these concepts to Riemannian manifolds is to project the data to a tangent space (which is Euclidean) and perform a variability analysis there. Recently, techniques have been established, which consider

the whole tangent bundle for parallel transport of tangent vectors along geodesics and perform a *principal geodesic analysis*, [24]. In our case, the medial ball shape space Σ is a Riemannian manifold with singularities. As is detailed in [48] for shape spaces of point configurations and carries over to our situation, Σ has singularities at equivalence classes of ball representations under translation and scaling, where the rotation group does not act free. These shapes arise from ball representations consisting of balls of the same radius lying on a straight line. We can exclude these ball representations from our further considerations since they are not of interest to us, or apply the simulation of simplicity concept [22] to ball representations M_1, \dots, M_n and avoid this case. For sake of simplicity, we present only details on the projection technique, although the whole shape analysis and segmentation pipeline is generalizable to a nonlinear framework.

The tangent space of Σ at $[\boldsymbol{\mu}]_{\sim_{ST}}$.

For our medial ball shape space Σ , we compute projections of shapes on the tangent space at $[\boldsymbol{\mu}]_{\sim_{ST}}$. Let $M = (x_1, \dots, x_{k-1}; r_1, \dots, r_k)$ be a ball representation, in normalized form (15), and $\text{vec}(M) = (x_1^T, \dots, x_{k-1}^T)^T \in \mathbb{R}^{3k-3}$. Then the orthogonal projection of M to the tangent space of Σ at $\boldsymbol{\mu}$ is given by

$$P_{\boldsymbol{\mu}}(M) = (\text{vec}(M) - \langle \text{vec}(M), \text{vec}(\boldsymbol{\mu}) \rangle \text{vec}(\boldsymbol{\mu}); r_1, \dots, r_k) \in \mathbb{R}^{3k-3} \times (\mathbb{R}^{>0})^k \quad (18)$$

Concerning the center points $c(M)$ of M , these are the Procrustes tangent coordinates as in [20, p.76]. The tangent projection is illustrated in Figure 2, left part. Tangent coordinates represent the linear deviation of M relative to $\boldsymbol{\mu}$ and will be used for statistical analysis in the sequel.

PCA and Mahalanobis distance for ball representations in Σ .

First we recall the concepts of *principal component analysis (PCA)* and *Mahalanobis distance* in \mathbb{R}^d . For data points $m_1, \dots, m_n \in \mathbb{R}^d$, a PCA is given by the normed eigenvectors of the covariance matrix

$$\text{Cov}(m_1, \dots, m_n) = \frac{1}{n} \sum_{i=1}^n (m_i - \bar{m})(m_i - \bar{m})^T \quad \text{with} \quad \bar{m} = \frac{1}{n} \sum_{i=1}^n m_i. \quad (19)$$

The eigenvectors e_1, \dots, e_d are sorted by the size of the corresponding eigenvalues $\lambda_1, \dots, \lambda_d$ and give an orthonormal basis of \mathbb{R}^d . This basis represents the main directions of variability of the data vectors m_1, \dots, m_n .

A distance which takes into account the variability of the data m_1, \dots, m_n is given by the Mahalanobis distance. This distance is defined by

$$d_{\mathbb{R}^d}(\tilde{m}_1, \tilde{m}_2) = (\tilde{m}_1 - \tilde{m}_2)^T (\text{Cov}(m_1, \dots, m_n))^{-1} (\tilde{m}_1 - \tilde{m}_2) \quad (20)$$

for \tilde{m}_1 and $\tilde{m}_2 \in \mathbb{R}^d$. Compared to the Euclidean distance, the Mahalanobis distance of two data points is small, if they differ along the main directions of variability (which are the principal components of m_1, \dots, m_n), and it is large, if they differ along other

directions. If d is large it might be appropriate to use a PCA first, and taking only the \bar{d} most significant components $e_1, \dots, e_{\bar{d}}$ to eigenvalues $\lambda_1, \dots, \lambda_{\bar{d}}$. This results in a simplified Mahalanobis distance

$$\tilde{d}_{\mathbb{R}^d}(\tilde{m}_1, \tilde{m}_2) = (\tilde{m}_1 - \tilde{m}_2)^T \tilde{U} \tilde{D}^{-1} \tilde{U}^T (\tilde{m}_1 - \tilde{m}_2) \quad (21)$$

with $\tilde{U} = (e_1, \dots, e_{\bar{d}})$ and \tilde{D} the diagonal matrix with entries $\lambda_1, \dots, \lambda_{\bar{d}}$.

Applying these concepts to ball representations, let $P_{\mu}(M_i) = (c_{\mu}^i; r_{\mu}^i) \in \mathbb{R}^{3k-3} \times (\mathbb{R}^{>0})^k$ be the tangent coordinates of M_i , split into centers c_{μ}^i and radii r_{μ}^i . A PCA of ball representations M_1, \dots, M_n , is given by the eigenvectors and eigenvalues of the covariance matrices

$$\text{Cov}(c_{\mu}^1, \dots, c_{\mu}^n) \text{ and } \text{Cov}(r_{\mu}^1, \dots, r_{\mu}^n). \quad (22)$$

The Mahalanobis distance d_{Σ} between M and \tilde{M} in Σ is then given by (with $P_{\mu}(M) = (c_{\mu}; r_{\mu})$ and $P_{\mu}(\tilde{M}) = (\tilde{c}_{\mu}; \tilde{r}_{\mu})$)

$$\begin{aligned} d_{\Sigma}(M, \tilde{M}) &= (c_{\mu} - \tilde{c}_{\mu})^T \text{Cov}(c_{\mu}^1, \dots, c_{\mu}^n)^{-1} (c_{\mu} - \tilde{c}_{\mu}) \\ &\quad + \alpha (r_{\mu} - \tilde{r}_{\mu})^T \text{Cov}(r_{\mu}^1, \dots, r_{\mu}^n)^{-1} (r_{\mu} - \tilde{r}_{\mu}). \end{aligned} \quad (23)$$

In case that the ball representations contain many balls, a PCA is performed first and a Mahalanobis distance as in (21) is computed.

PCA and Mahalanobis distance are illustrated in Figure 2, right part.

As a result of this section, a mean shape and distances between shapes have been defined and this statistics can be used as shape prior in segmentation in the sequel (cf. (32) and (33)). Generally spoken, the Mahalanobis distance is used to make a point cloud uniformly distributed. Therefore one can use quadratic regularization functionals. For a comprehensive work on the problem of defining suitable approximations of shape metrics in image processing we refer to Charpiat et al. [14] where the authors also perform warping of two shapes onto another by gradient descent that minimizes the corresponding distance.

5 Segmentation

In this section we consider object segmentation in 3D voxel data. Thereby we aim for taking into account shape prior information given in the form that the object to be recovered belongs to the medial ball shape space. The distance on the shape space is given by the Mahalanobis distance (see Sections 2-4). The medial ball shape space is computed from boundaries of training objects as discussed in Section 2.

The boundary of a ball representation is defined as the *skin surface* [21]. Below we give a definition and some properties of skin surfaces which predestine them as boundary surfaces of medial ball representations.

Finally we describe the whole pipeline of the proposed segmentation process.

5.1 Skin Surfaces

In this subsection, we review the definition and basic properties of skin surfaces.

Properties of skin surfaces.

We employ the concept of a skin surface later on for definition of a surface from a ball representation. We believe that skin surfaces are very suitable for this task, since they have several attractive properties which we discuss in the sequel.

- *Information efficiency:* A skin surface only requires center points, weights and a shrinking factor for its complete description. Center points and weights are exactly the information available to us after (mean) ball representation computation in Sections 2 and 3. The shrinking factor can be used as a tuning parameter depending on the application. Other popular surface construction methods like NURBS (or other splines), used for instance in [19], or Gregory patches (see [44, 17]) suffer from the drawback, that an additional grid structure of the control points is required for surface generation, which is difficult to obtain in an automatic way.
- *Computational efficiency:* There exists an efficient algorithm for computing skin surfaces (see [37]). For k weighted input points, the algorithm produces a triangular mesh consisting of $O(k^2)$ vertices.
- *Ball representation fidelity:* If M is a ball representation of some surface in \mathbb{R}^3 , and $\text{skn}(M)$ is the skin surface defined by M , then M is contained in the ball representation of $\text{skn}(M)$. This follows easily from the definition of the skin surface as a convex hull.
- *Economy:* A small number of input points can generate complicated skin surfaces [21].
- *Universality:* Every orientable closed surface has a skin representation [21].

An example of a skin surface, meshed with the algorithm from [37], is given in Figure 3.

Definition of a skin surface.

Basic building blocks of skin surfaces are *weighted points* $p = (x_p, w_p) \in \mathbb{R}^3 \times \mathbb{R}$. Weighted points can for instance represent balls with center x and radius r , setting $x_p = x$ and $w_p = r^2$, and therefore a ball representation can be regarded as a set of weighted points. Addition and scalar multiplication of weighted points p, q are defined by

$$p + q = (x_p + x_q, w_p + w_q + 2\langle x_p, x_q \rangle), \quad (24)$$

and

$$sp = (sx_p, sw_p + (s^2 - s)\|x_p\|^2) \quad \text{for } s > 0. \quad (25)$$

Here, $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ denote the Euclidean scalar product and its induced norm. As usual, but with these algebraic operations, the convex hull of a set of weighted points

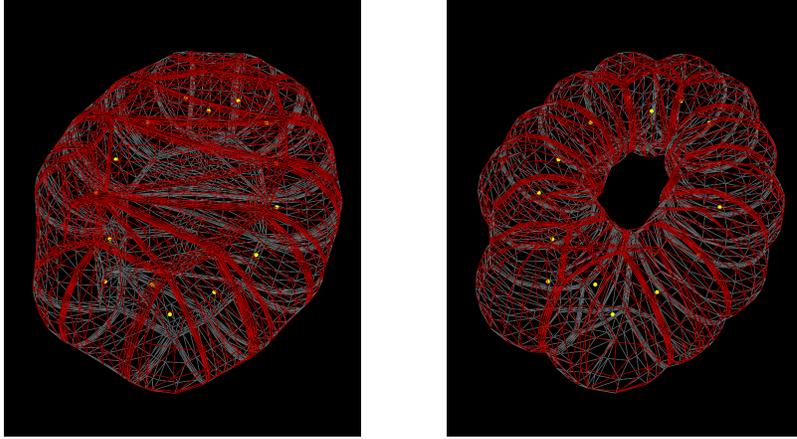


Figure 3: Examples of a skin surface of 15 balls with shrink factor $s = 0.1$ (left) and $s = 0.9$ (right). The dots indicate the centers of the medial balls. Note the topology changing effect of increasing the shrink factor, which can be utilized for constructing shapes of different topology.

$P = \{p_1, \dots, p_k\}$ is defined by $\text{conv}(P) = \left\{ \sum_{i=1}^k \lambda_i p_i \mid \lambda_i \geq 0, \sum_{i=1}^k \lambda_i = 1 \right\}$. The *shrinkage* p^s of a weighted point $p = (x_p, w_p)$ by a factor $s \geq 0$ is defined by $p^s = (x_p, s w_p)$.

Assume now we are given a set P of weighted points $p = (x, w)$ with $w > 0$. Shrinkage applied to P by $s \geq 0$ is defined by $P^s = \{p^s \mid p \in P\}$. Let $B_w(x) = \{y \in \mathbb{R}^3 \mid \|x - y\| \leq w\}$ be a ball with radius w and center x . We denote by ∂ denote the boundary of a set in Euclidean space.

Let $s \geq 0$. Then the s -*skin* of P is defined by

$$\text{skn}^s(P) = \partial \left(\bigcup_{q=(x_q, w_q) \in (\text{conv} P)^s} B_{\sqrt{w_q}}(x_q) \right) \quad (26)$$

Note that the square root of a negative number is imaginary and a ball with an imaginary radius is the empty set.

In the special case $s = 0$, the skin surface of p_1, \dots, p_k reduces to the convex hull of x_{p_1}, \dots, x_{p_k} , and in case $s = 1$, the skin surface is the boundary of the union of the input weighted points.

Since weighted points are shrunk by a factor of $s \in [0, 1]$ in the construction of the skin surface, the skin surface of a surface represented by medial balls is smaller than the original surface. A way to circumvent this problem is by prescaling the input weights w_{p_1}, \dots, w_{p_k} by $1/s$.

The implied skin surface of a shape.

For a fixed ball representation μ , and some given similarity transformation (ρ, A, b) , a shape in Σ implies a surface. For its computation, it is necessary to reverse the transformations in Subsection 4.1. To give some details, assume that the shape in Σ is given by its tangent coordinates $P_\mu = (c_\mu, r_\mu)$. First note that the tangent projection (18) is injective, with inverse

$$(c_\mu; r_\mu) \mapsto \left(\text{vec}^{-1} \left(\sqrt{1 - c_\mu^T c_\mu} \text{vec}(c_\mu(\bar{\mu})) + c_\mu \right); r_\mu \right) \quad (27)$$

mapping tangent coordinates to normalized ball representations. Normalized ball representations are then aligned in \mathbb{R}^3 by applying A^{-1} and scaling with $1/\rho$. When reversing translation, note that the Helmert matrix H has a right inverse $\tilde{H} \in \mathbb{R}^{(k-1) \times (k-1)}$.

Applying \tilde{H} and translating the result by b , an (aligned) ball representation M consisting of k balls in \mathbb{R}^3 is obtained, to which the skin surface

$$\gamma(M) = \gamma(P_\mu, \rho, A, b) \quad (28)$$

is the implied boundary.

5.2 Region and edge based segmentation

For segmentation of voxel images, there are essentially two types of segmentation methods. *Region based* segmentation is applied to images, if the mean intensity of voxels inside the object to segment differs significantly from the mean intensity outside the object. If contrasts are low and objects are only separated by curves, *gradient based* segmentation is used. Here, we briefly discuss both approaches. Let $\Omega \subset \mathbb{R}^3$ be a bounded domain, and $u : \Omega \rightarrow \mathbb{R}$ an image intensity function.

Region based segmentation.

For a closed surface $\gamma \subset \Omega$, let $\mathcal{I}(\gamma) \subset \Omega$ be the inner part of γ and $\mathcal{O}(\gamma) \subset \Omega$ the outer part of γ . The mean values of an image intensity function u on $\mathcal{I}(\gamma)$ and $\mathcal{O}(\gamma)$ are then given by

$$\bar{u}_{\mathcal{I}(\gamma)} = \frac{1}{|\mathcal{I}(\gamma)|} \int_{\mathcal{I}(\gamma)} u \, dx \quad \text{and} \quad \bar{u}_{\mathcal{O}(\gamma)} = \frac{1}{|\mathcal{O}(\gamma)|} \int_{\mathcal{O}(\gamma)} u \, dx \quad (29)$$

A simplified version of the Mumford-Shah model introduced in [13] consists in minimization of the functional

$$I_{SMS}(\gamma) = \int_{\mathcal{I}(\gamma)} (\bar{u}_{\mathcal{I}(\gamma)} - u)^2 \, dx + \int_{\mathcal{O}(\gamma)} (\bar{u}_{\mathcal{O}(\gamma)} - u)^2 \, dx. \quad (30)$$

If we assume that the mean intensity of an object in image u differs significantly from the mean intensity of the region outside of this object, the functional I_{SMS} is minimized by a submanifold γ which represents the boundary of the object to segment, and $\mathcal{I}(\gamma)$ is the object itself.

Edge based segmentation.

In edge based segmentation, an object and its closed boundary surface γ are implied by high gradients. In this case, it is common to minimize a functional which penalizes small gradients, as for example the Snakes energy introduced in [32]:

$$I_{KWT}(\gamma) = - \int_{\gamma} |\nabla f(\gamma)| dx dy + \alpha \text{Area}(\gamma) \quad (31)$$

Note that edge based energy functionals are usually minimized by active contour methods. They consist of evolving a small initial contour mesh towards the boundary of the object, which is a minimizer of the energy functional. An advantage of active contour methods is that the evolution converges to a minimizer close to the initial contour. However, since the evolution equations only contain local information, but shape prior is a global information, active contour methods cannot be used efficiently in this case. In order to retain control over local minimizers of (31), we use the term $\alpha \text{Area}(\gamma)$ in (31) instead, thus forcing the contour closer to the initial contour.

In the presence of noisy data that e.g. arises in MRI or ultrasound imaging, minimization of both functionals (30) and (31) over γ is not well-posed. Furthermore, if the object to segment is partially occluded, minimizers of (30) and (31) might not only segment the object, but also parts of the background, and knowledge about the true shape of the object to segment has to be considered.

Here we use the mean ball representation μ of Equation (11) as a priori information and the Mahalanobis distance d_{Σ} as defined in Subsection 4.2, formula (23) to measure the distance between two ball representations. We recall that with a given fixed ball representation μ , an equivalence class $[M]_{\sim_{ST}}$ of ball representations can be uniquely represented by tangent coordinates $P_{\mu}(M) = (c_{\mu}(M), r_{\mu}(M))$, see (18). Tangent coordinates and a similarity transformation $(\rho, A, b) \in \mathbb{R}^{>0} \times SO(3) + \times \mathbb{R}^3$ uniquely determine a skin surface $\gamma = \gamma(P_{\mu}(M), f)$ as defined in Subsection 5.1, formula (28). We obtain the following functionals as a regularization of (30) resp. (31) for some $\beta > 0$:

$$\begin{aligned} I_{\beta} : \Sigma \times \text{SO}(3) \times \mathbb{R}^3 \times \mathbb{R}^{\geq 0} &\longrightarrow \mathbb{R}, \\ (P_{\mu}, f) &\mapsto I_{SMS}(\gamma(P_{\mu}, f)) + \beta d_{\Sigma}(0, P_{\mu}) \quad (32) \\ (P_{\mu}, f) &\mapsto I_{KWT}(\gamma(P_{\mu}, f)) + \beta d_{\Sigma}(0, P_{\mu}) \quad (33) \end{aligned}$$

Note that the regularization term actually measures the Mahalanobis distance between M and μ , since $P_{\mu}(\mu) = 0$.

6 Results

In this section we discuss the results obtained by applying our approach to test images. A dataset of synthetic torus voxel images was segmented using functional (32) in Section 6.1. Additionally, we conducted an experiment on real-world data by segmenting the human hippocampus in Section 6.2 using functional (33). Both functionals were

minimized using the evolutionary algorithm CMA-ES (see [28, 36]). This choice is mainly motivated by the fact that the calculation of the gradient for the functional is tough, since it requires finding the derivative with respect to the skin surface, which is given by the coordinates of the medial ball representation. Thus, instead of analytical differentiation, numerical differentiation should be used for this purpose. However, the skin surface meshing process we used for surface construction ([37]) does not guarantee that the resulting polyhedra always have the same structure or even the same number of vertices and faces, as it is for example the case by using spline based interpolation surfaces. Those irregularities in the resulting surface meshes complicate the numerical differentiation process and thus we decided to use gradient-free evolutionary minimization algorithms. We chose the evolutionary algorithm CMA-ES for the minimization of (32), because it has been successfully applied to a similar task in [17], where its superior behavior over two other common evolutionary minimization techniques was demonstrated. For more details on the minimization process, especially also concerning the choice of the minimization parameters, we also refer to [17].

6.1 Synthetic Examples

We now consider the segmentation of a synthetically created and corrupted torus voxel image. By this example we demonstrate not only the influence of the shape prior on the segmentation result, but in particular also the capability of the method to segment surfaces of complex geometry, i.e., a torus in this case. We generated 15 binary training voxel images containing elliptic tori with differing radii and positions (see Figure 4). Using the software ITK-SNAP [53] we computed the according input meshes $\Omega_1, \dots, \Omega_{15}$ which are the actual input data for the shape statistics.

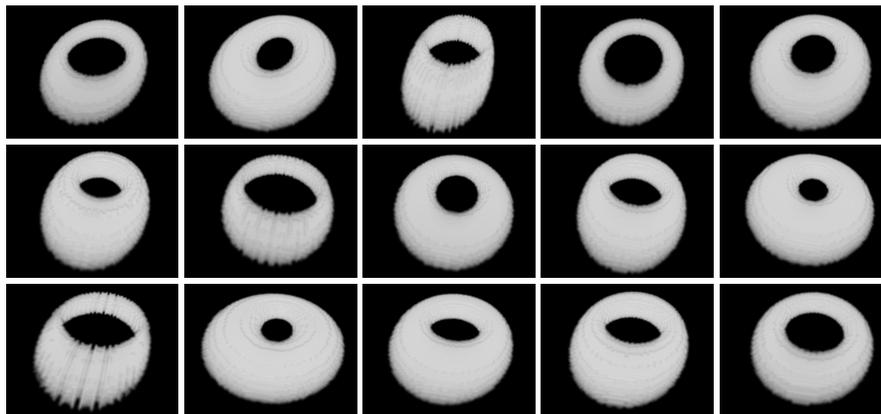


Figure 4: 15 different voxel images of tori used as training data

Using the surface meshes, Scheme 1.1 has been used to compute a shape space and a Mahalanobis distance in it. Figure 3 shows the mean medial ball representation of the shape space. This shape information was then used to segment a torus corrupted

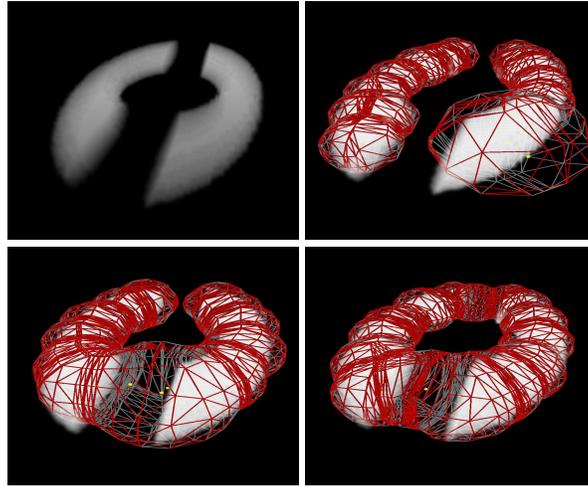


Figure 5: Segmentation results for the corrupted torus image for different values of the regularization parameter β . From upper left to lower right: (a) Voxel image of corrupted torus data. (b) Segmentation result with regularization parameter $\beta = 10^{-3}$ (c) $\beta = 10^{-2}$ (d) $\beta = 5 * 10^{-2}$.

by a large black strip as shown in Figure 5(a). Because high contrasts occur in the images, we chose (32) as the segmentation functional, which we minimized using the CMA-ES. The segmentation algorithm was run with $k = 15$ balls. The shrink factor s was included in the minimization process. As initial value for s it proved suitable to take high values if long and thin structures were expected and smaller values in case of more compact objects. The results of the segmentation process with different values for the regularization parameter β are shown in Figure 5(b)-(d). As expected, small values of β , i.e., segmentation with low influence of the regularization, results in improper segmentation results (Figure 5(b)-(c)), while the right choice of β yields a good segmentation, as seen in Figure 5(d).

Note that the slight inaccuracies in the segmentation are a consequence of using skin surfaces as surface generation method, which often results in a surface of ball-like structure by nature. This effect can be minimized by using more medial balls to represent the shape. In principle, the medial ball representation presented above is able to describe complex shapes with bended or flat surfaces and sharp corners provided that the number of medial balls is sufficiently high. It has to be mentioned though that the time complexity of the segmentation substantially rises when using more input balls. This is mainly due to the skin surface meshing algorithm which has to be computed for every new instantiation of a medial ball representation. Minimization using evolutionary optimization algorithms like the CMA-ES involves multiple evaluations of the functional with different models at each step of the segmentation and thus the time complexity for each step rises significantly. However, each evaluation of the functional is performed independently of each other and thus parallelization can be applied which is especially boosting the performance on modern multi-core processor environments.

For the synthetic images above, 15 balls were used to describe the torus shape and the CMA-ES algorithm was applied with $\lambda = 100$ functional evaluations at each step. Considering 100 iterations for one segmentation, the average time for one step in this setting was 20 seconds on a standard desktop PC with an Intel Core 2 Q9550 2.83 Ghz Quad Core CPU and 4 GB of RAM. For the hippocampus model with 50 medial balls (cf. Section 6.2 below), an average time of 60 seconds per step was measured.

6.2 Real-World Example - Hippocampus Segmentation

In this section we test the presented algorithm on a real-world example from medical imaging. Our object of interest to segment is the human hippocampus. To this end, we use a database of MRI images of the human brain presented in [29] which includes 25 labelled samples of the right and left hippocampi respectively. Using the method described above we constructed a shape space out of the set of right hippocampi for all training images. Here, $k = 50$ balls were used to model the hippocampus. Afterwards we segmented a test image utilizing the shape prior with the edge based functional (33). This is due to the fact that region based segmentation is not applicable in MRI images where many regions with different contrast exist. To construct a gradient image of the hippocampus, we thresholded the test image according to the mean contrast values gained from the training images and computed a morphological gradient which is shown in Figure 7. The computation of the gradient image thus depends on the given training image data. We want to point out however that the presented methodology is not intended as a fully automatic medical imaging segmentation workflow but could prove to be effective in this setting when used in conjunction with standard medical imaging preprocessing steps.

To evaluate our segmentation results we used two common similarity coefficients in the image segmentation community: Jaccard and Dice’s coefficient. Interpreting two binary segmentation images as sets A and B Jaccard and Dice’s coefficient, J_C and D_C , are defined as follows:

$$J_C = \frac{|A \cap B|}{|A \cup B|}, \quad (34)$$

$$D_C = \frac{2|A \cap B|}{|A| + |B|}. \quad (35)$$

Additionally, to compare our results against a similar approach based on classical m-reps (cf. [30, 42]), we *manually* created a shape space of m-reps for the training images using a mesh of size 5×10 , i.e. 50 atoms as with the medial ball model. This shape prior was then used to perform image segmentation with a similar edge-based functional as (33). For a detailed outline of the segmentation algorithm using m-reps and a-priori shape knowledge we refer to the work of [17]. Figure 6 shows the mean medial models constructed for the m-rep and medial ball representation respectively.

Figures 8 and 9 show the segmentation results using the m-rep and medial ball model. In Table 1 the similarity coefficients for the resulting segmentations are summarized. Both segmentations yield good results, however the medial ball model is able to outperform the m-rep based method by a small margin. We want to note however,

that the m-rep approach does not include *automatic* shape space generation as it is the case with our method.

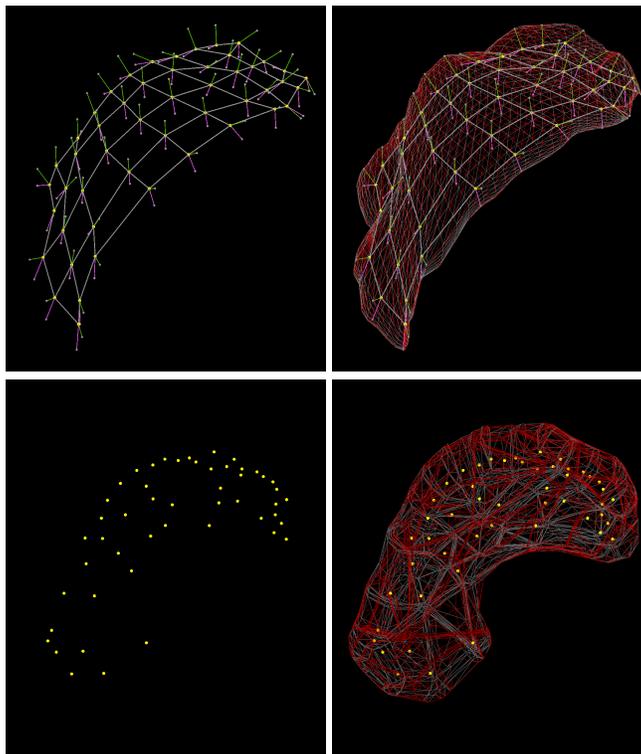


Figure 6: Mean models for the 25 training hippocampi. On the top the mean for the m-rep model due to Pizer et al. [30, 42] with and without surface is shown. The bottom row corresponds to the mean medial ball representation of the hippocampus as presented in this work.

7 Conclusion

In this paper we presented a new approach for segmentation of 3D voxel images taking into account statistical shape information. The algorithm requires minimal user interactions for segmentation which is achieved by implementing a pipeline of algorithms. The pipeline involves the automatic generation of training shapes represented as medial ball representations that were automatically generated from a given set of input meshes. Afterwards the training shapes are labeled, and a shape space based on Procrustes statistics is established. The resulting shape space and the Mahalanobis distance are used as a prior in region- and edge based segmentation algorithms.

We did a proof-of-concept for our algorithm by evaluating its performance on a synthetic dataset of 15 randomly generated torus images. By utilizing a segmentation

<i>Model</i>	J_C	D_C
m-rep	0.7380	0.8492
medial balls	0.7572	0.8618

Table 1: Segmentation evaluation results for the test image. The suggested medial ball algorithm performs better than the m -reps algorithm with respect to the two distance measures J_C and D_C .

energy which includes statistical regularization using the PCA of the training data, we were able to segment a distorted torus image correctly. Additionally, this example shows the benefit of using only a set of medial balls as object model in combination with skin surfaces as surface generation method, which enables us to segment images with complex geometry.

Additionally, the algorithm was applied to a real-world dataset of human hippocampi in MRI images. A comparison with the segmentation using shape priors on the traditional m-rep model showed that our methodology is able to outperform m-rep based segmentation by a small margin.

The automatic generation of a medial ball object model from a set of training images is another advantage of our algorithm in contrast to the work of Pizer et al. (e.g. [42]), where expert segmentations have to be constructed by using a predefined medial ball model itself.

We emphasize that Siddiqi and Pizer [47] work with **medial atoms** while we work with **medial balls**, which allows for more flexibility in the objects to be represented.

Acknowledgement

This work has been supported by the Austrian Science Fund (FWF) within the research networks NFNs *Industrial Geometry*, Projects S09203 and S09205, and Photoacoustic Imaging in Biology and Medicine, Project S10505-N20. The authors thank the referees for stimulating comments which helped to improve the quality of the manuscript.

References

- [1] Medical image computing and computer assisted intervention. <http://www.miccai2009.org>.
- [2] J. Abhau, W. Hinterberger, and O. Scherzer. Segmenting surfaces of arbitrary topology: A two-step approach. In *Ultrasonic Imaging and Signal Processing*. Proceedings of SPIE – Volume 6437, 2007.
- [3] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [4] O. Aichholzer, F. Aurenhammer, T. Hackl, M. Kornberger, B. M. Peternell, and H. Pottmann. Approximating boundary-triangulated objects with balls. In *Proc.*

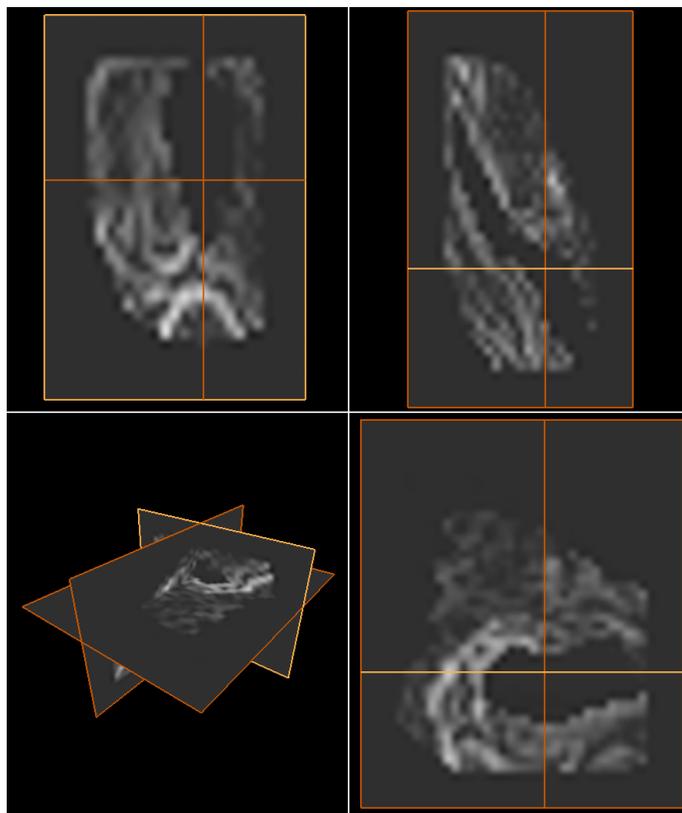


Figure 7: The absolute gradient for the test image of the hippocampus segmentation. The top left, top right and bottom right images show the 2D slices of the volume while the bottom left image displays all 3 slices in 3D.

23rd European Workshop on Computational Geometry, pages 130–133, Graz, 2007.

- [5] O. Aichholzer, F. Aurenhammer, and B. Kornberger. Stable piecewise-linear approximations of 3d medial axes. Manuscript.
- [6] O. Aichholzer, F. Aurenhammer, B. Kornberger, S. Plantinga, G. Rote, A. Sturm, and G. Vegter. Recovering structure from r -sampled objects. In *Proc. Eurographics Symposium on Geometry Processing*, volume 28, pages 1349–1360, Berlin, 2009.
- [7] N. Amenta and M. Bern. Surface reconstruction by “Voronoi” filtering. *Discrete & Computational Geometry*, 22:481–504, 1999.
- [8] N. Amenta and R. Kolluri. Accurate and efficient unions of balls. In *Proc. 16th Ann. Symp. Computational Geometry*, pages 119–128, Hong Kong, 2000. ACM.

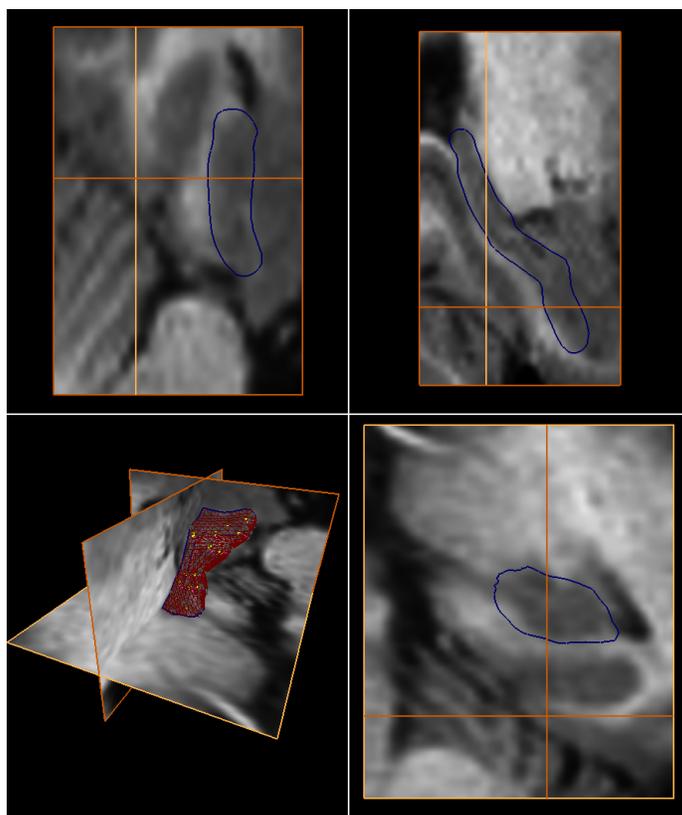


Figure 8: Segmentation result for the m-rep model. On the bottom left the resulting 3D model is shown, while the blue contours on the other images indicate the segmentation in the 2D slice images.

- [9] H. Blum and R.N. Nagel. Shape description using weighted symmetric axis features. *Pattern Recognition*, 10:167–180, 1978.
- [10] J.-D. Boissonnat and M. Teillaud, editors. *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006.
- [11] F. L. Bookstein. *Morphometric tools for landmark data: geometry and biology*. Cambridge University Press, Cambridge, 1997.
- [12] X. Bresson, P. Vandergheynst, and J.P. Thiran. A variational model for object segmentation using boundary information and shape prior driven by the mumford-shah functional. *Int. J. Comput. Vision*, 28(2):145–162, 2006.
- [13] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Process.*, 10(2):266–277, 2001.

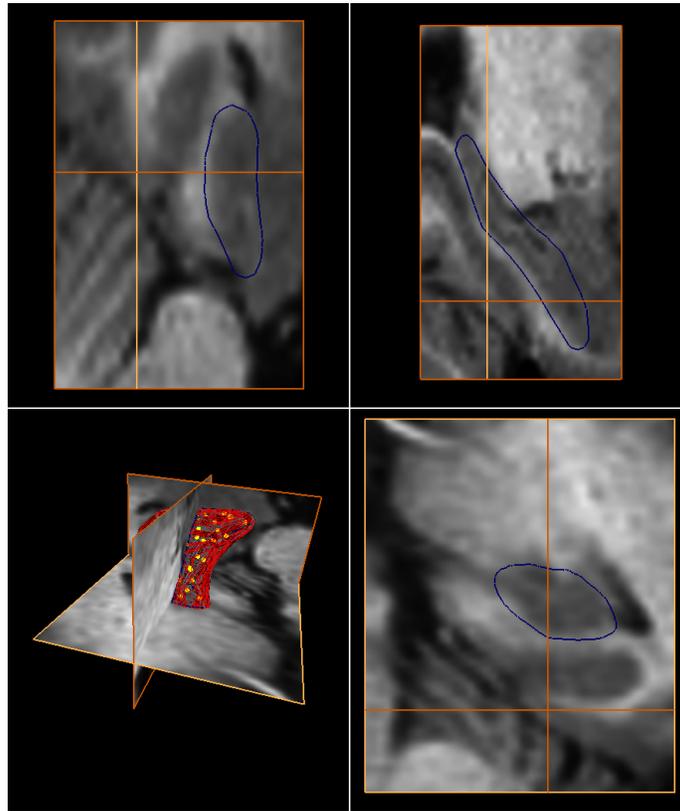


Figure 9: Segmentation result for the medial ball representation. On the bottom left the resulting 3D model is shown, while the blue contours on the other images indicate the segmentation in the 2D slice images.

- [14] G. Charpiat, O. Faugeras, and R. Keriven. Approximations of Shape Metrics and Application to Shape Warping and Empirical Shape Statistics. *FOUNDATIONS OF COMPUTATIONAL MATHEMATICS*, 5:1–58, 2004.
- [15] Q. Chen, Z.M. Zhou, M. Tang, P.A. Heng, and D.S. Xia. Shape statistics variational approach for the outer contour segmentation of left ventricle mr images. *IEEE Trans. Inf. Technol. Biomed.*, 10(3):588–597, 2006.
- [16] Y. Chen, H.D. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K.S. Gopinath, R.W. Briggs, and E.A. Geiser. Using prior shapes in geometric active contours in a variational framework. *Int. J. Comput. Vision*, 50:315–328, 2002.
- [17] S. Colutto, F. Frühauf, M Fuchs, and O. Scherzer. The CMA-ES on Riemannian manifolds to reconstruct shapes in 3D voxel images. *IEEE Transactions on Evolutionary Computation*, 14(2): 227 - 245 2010.

- [18] D. Cremers, T. Kohlberger, and C. Schnoerr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36(9):1929–1943, 2003.
- [19] D. Cremers, F. Tischhäuser, J. Weickert, and Ch. Schnörr. Diffusion snakes: Introducing statistical shape knowledge into the Mumford-Shah functional. *Int. J. Comput. Vision*, 50:295–313, 2002.
- [20] I.L. Dryden and K.V. Mardia. *Statistical Shape Analysis*. John Wiley, Chichester, 1998.
- [21] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comp. Geom.*, 21:87–115, 1999.
- [22] H. Edelsbrunner and E.P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, 1990.
- [23] W. Fang and K.L. Chan. Incorporating shape prior into geodesic active contours for detecting partially occluded object. *Pattern Recognition*, 40(7):2163–2172, 2007.
- [24] P.T. Fletcher, S. Joshi, C. Ju, and S.M. Pizer. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Med. Imag.*, 23:995–1005, 2004.
- [25] D.S. Fritsch, S.M. Pizer, L. Yu, V. Johnson, and E.L. Chaney. Localization and segmentation of medical image objects using deformable shape loci. *Lecture Notes in Computer Science*, 1230:127–140, 1997. IPMI 1997.
- [26] M. Fuchs and S. Gerber. Variational shape detection in microscope images based on joint shape and image feature statistics. In *CVPR Workshops 2008. IEEE Computer Society Conference on*, pages 1–8, 2008.
- [27] M. Gastaud, M. Barlaud, and G. Aubert. Combining shape prior and statistical features for active contour segmentation. *IEEE Trans. Circuits and Systems*, 14(5):726–734, 2004.
- [28] N. Hansen, S. D. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 2003.
- [29] K. Jafari-Khouzani, K. Elisevich, S. Patel, and H. Soltanian-Zadeh. Dataset of magnetic resonance images of nonepileptic subjects and temporal lobe epilepsy patients for validation of hippocampal segmentation techniques. *Neuroinformatics*, 2011.
- [30] S. Joshi, S. Pizer, P.T. Fletcher, A. Thall, and G. Tracton. Multi-scale 3-d deformable model segmentation based on medical description. In *Proc. International Conference on Information Processing in Medical Imaging (IPMI)*, pages 64–77, 2001.

- [31] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. An efficient k-means clustering algorithm: Analysis and implementation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(7):881–892, 2002.
- [32] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Comput. Vision*, 1(4):321–331, 1987.
- [33] D.G. Kendall. Shape manifolds, procrustean metrics and complex projective spaces. *Bull. Lond. Math. Soc.*, 16:81–121, 1984.
- [34] J. T. Kent, K.V. Mardia, and C.C. Taylor. Matching problems for unlabelled configurations. *Bioinf. Images Wavelets*, pages 33–36, 2004.
- [35] J.T. Kent, K.V. Mardia, and C.C. Taylor. Matching unlabelled configurations and protein bioinformatics. *to appear*.
- [36] S. Kern and N. Hansen. Evaluating the cma evolution strategy on multimodal test functions. In *Eighth International Conference on Parallel Problem Solving from Nature PPSN VIII*, volume 3242, pages 282–291. Springer, 2004.
- [37] N. Kruithof and G. Vegter. Meshing skin surfaces with certified topology. *Comput. Geom.*, 36:166–182, 2007.
- [38] M.E. Leventon, W.E.L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR-00)*, pages 316–323, Los Alamitos, 2000. IEEE.
- [39] K.V. Mardia, J.T. Kent, and J.M. Bibby. *Multivariate Analysis*. Academic Press, London, 1977.
- [40] G. McLachlan and T. Krishnan. *The EM algorithm and extensions*, volume 2. Wiley-Interscience, Hoboken, NJ, 2008.
- [41] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42(5):577–685, 1989.
- [42] S.M. Pizer, P.T. Fletcher, S. Joshi, A. Thall, J.Z. Chen, Y. Fridman, D.S. Fritsch, A.G. Gash, J.M. Glotzer, M.R. Jiroutek, C. Lu, K.E. Muller, G. Tracton, P. Yushkevich, and E.L. Chaney. Deformable m-reps for 3d medical image segmentation. *Int. J. Comput. Vision*, 55(2):85–106, 2003.
- [43] S.M. Pizer, A.L. Thall, and D.T. Chen. M-reps: A new object representation for graphics. Technical Report TR99-030, Department of Computer Science, University of North Carolina — Chapel Hill, 1999.
- [44] M. A. Puso and T. A. Laursen. A 3-d contact smoothing algorithm method using gregory patches. *Numer. Meth. in Engineering*, 2002.

- [45] M. Rousson and N. Paragios. Shape priors for level set representations. In *European Conference in Computer Vision (ECCV)*, pages 78–93, 2002.
- [46] M. Rousson and N. Paragios. Prior knowledge, level set representations & visual grouping. *Int. J. Comput. Vision*, 76(3):231–243, 2007.
- [47] K. Siddiqi and S. M., Pizer. *Medial representations*, Computational Imaging and Vision Series Vol. 37, Springer, New York, 2008.
- [48] C. G. Small. *The Statistical Theory of Shape*. Springer Series in Statistics. Springer, 1996.
- [49] S. Suri. Bipartite matching and the hungarian method. www.cse.ust.hk/golin/COMP572/Notes/Matching.pdf.
- [50] A. Tsai, A. Yezzi, C. Tempny, D. Tucker, A. Fan, W.E.L. Grimson, and A. Will-sky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Trans. Med. Imag.*, 22(2):137–154, 2003.
- [51] A. Witkin, M. Kass, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Comput. Vision*, 1, No 4:321–331, 1987.
- [52] P. Yushkevich, P. T. Fletcher, S. C. Joshi, A. Thall, S. M. Pizer. Continuous medial representations for geometric object modeling in 2D and 3D, *Image and Vision Computing*, 21, No.1:17–27, 2003.
- [53] P.A. Yushkevich, J. Piven, C. Hazlett, H. and G. Smith, R. Ho, S., J.C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006.