

Model-based segmentation

Tobias Heimann, Hervé Delingette

► **To cite this version:**

Tobias Heimann, Hervé Delingette. Model-based segmentation. Deserno, Thomas Martin. Biomedical Image Processing, Springer, pp.279-303, 2011, 978-3-642-15815-5. <10.1007/978-3-642-15816-2_11>. <inria-00616063>

HAL Id: inria-00616063

<https://hal.inria.fr/inria-00616063>

Submitted on 16 Jul 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Model-based Segmentation

Tobias Heimann and Hervé Delingette

Asclepios Project, INRIA Sophia Antipolis, France
tobias.heimann@inria.fr

Abstract. This chapter starts with a brief introduction into model-based segmentation, explaining the basic concepts and different approaches. Subsequently, two segmentation approaches are presented in more detail: First, the method of deformable simplex meshes is described, explaining the special properties of the simplex mesh and the formulation of the internal forces. Common choices for image forces are presented, and how to evolve the mesh to adapt to certain structures. Second, the method of point-based statistical shape models is described. The model construction process is explained and the point correspondence problem is treated in detail. Several approaches of how gray level appearance can be modeled are presented, and search algorithms that use this knowledge to segment the modeled structures in new images are described.

12.1 Introduction

Automatic medical image segmentation is such a challenging task because it involves highly variable objects that have to be extracted from images of very low quality. Often, lack of contrast or artifacts lead to missing data in the image to be segmented, and the boundary of the object of interest cannot be determined from local information as edges or region homogeneity alone. Medical experts are still able to delineate the object because they know what it is supposed to look like: They have a model of the object in their mind, they have an *a-priori* knowledge about its shape and appearance. Model-based segmentation methods strive to translate this knowledge into smart algorithms that have a prior knowledge about the structures of interest. Those methods can be qualified as *top down* and usually consists of two stages. The former initializes the location and appearance of the model. It is based on user input (mouse clicks in an image), the registration of an atlas or machine learning approaches. The latter optimizes the shape and appearance of the model such that it closely matches the ones measured in the images.

As the shape of a model is optimized during the segmentation process, one important property of model-based segmentation lies in the choice of

the deformation strategy. There are two basic ways of deforming a shape: by deforming the embedding space of a shape and by modifying their parameters or Degree of Freedom (DOF). For instance, applying a rigid transform (i.e., rotation and translation) lies in the first category while modifying the vertex position of a mesh lies in the second category. Of course, those two mechanisms are often combined but their formalism and approaches are fairly different. Deforming a shape by deforming the embedding space can be qualified as a *registration approach* whereas modifying the parameters of a shape can be qualified as a *deformable model approach*. The geometric representation of the models is specifically important in the latter case. In Fig. 12.1, we show a coarse taxonomy of possible geometric representations. Each representation having pros and cons, the choice of a representation is often problem dependent. In Sec. 12.2, we detail the deformable model formulation based on simplex meshes. Among others, common representations are level-sets, triangulated meshes and spline curves or surfaces.

If many segmentations of a specific object are available, there are further possibilities to make model more specific. In addition to knowledge about an average template shape, Statistical Shape Models (SSMs) have also knowledge about the principal modes of variation. The SSMs presented in Sec. 12.3 are based on the Point Distribution Model (PDM). To perform statistical analysis on these shapes, point correspondences have to be known: a challenging problem, for which just recently automatic methods have been proposed [?]. The visual appearance of the object can also be captured by statistics on training samples, using techniques from computer vision and machine learning. Finally,




Fig. 12.1. Taxonomy of deformable models. The classification scheme of deformable models for medical image segmentation is based on their geometric representation.

there exist specific algorithms to employ the statistical information to drive a model towards an instance of the object in an image.

To conclude the chapter, we compare the concepts of a deformable simplex mesh and the SSM and give hints which technique to use under which circumstances.

12.2 Deformable Simplex Meshes

In a deformable model approach, the shape of a model is optimized in order to match that of a structure of interest in an image. This technique has been pioneered in 1987 by Terzopoulos et al. [1] with the introduction of active contours or *snakes* [2]. This has been later generalized to active surfaces [3] but one difficulty arises when dealing with three-dimensional (3D) surfaces: the continuous parameterization of surfaces. Indeed, even for surfaces having the same topology as the sphere, it is not straightforward to define a C^1 parametric surface.

In order to avoid this issue, authors have proposed to work with implicit parameterization (e.g., using level-sets function) or with C^0 continuous meshes (e.g., using triangulated or simplex meshes). In the latter case, shape smoothing or regularization cannot be defined in the continuous domain (since the representation is not even C^1) but must be defined in a discrete fashion.

A simplex mesh is an original mesh representation that has been introduced in 1992 [4] as a model-based segmentation method for range data and volumetric images. A k -simplex meshes can represent a manifold surface of dimension k ($k = 2$ for 3D surfaces) and can be defined as a $k + 1$ -cell. We provide below a recursive definition of a cell (Fig. 12.2).

Definition 1. A 0-cell is defined as a vertex P and a 1-cell as an edge, i.e., an unordered pair of distinct vertices (P, M) . We then recursively define a p -cell ($p \geq 2$) \mathcal{C} as the union of c ($p - 1$)-cells, $c \in \mathbb{N}$, such that:

1. Each vertex belonging to \mathcal{C} also belongs to p distinct ($p - 1$)-cells.
2. A ($p - 2$)-cell belongs to 2 and only 2 ($p - 1$)-cells.
3. The intersection of 2 ($p - 1$)-cells is empty or is a ($p - 2$)-cell.




Fig. 12.2. Examples of p -cells. A collection of p -cells is given for $0 \leq p \leq 3$. An edge is a 0-simplex mesh, a face is a 1-simplex mesh, and a 3-cell a 2-simplex mesh.

4. A p -cell is simply connected, i.e. that given 2 vertices of that cell, there exists at least one set of edges that connect those 2 vertices.

Two fundamental topological properties result from this definition. First, a k -simplex mesh has a fixed vertex connectivity: every vertex is adjacent to $k + 1$ distinct vertices. Second, a k -simplex mesh is a topological dual of a k -triangulation. In this duality relationship, a triangle is associated with a vertex of a 2-simplex mesh, a triangulation edge with an edge and a vertex with a face (a 2-cell).

In a deformable 2-simplex mesh for image segmentation, the vertex position \mathbf{P}_i evolves over time under the influence of internal $\mathbf{f}_i^{\text{int}}$ and external forces $\mathbf{f}_i^{\text{ext}}$ capturing the influence of prior knowledge on shape and appearance, respectively

$$\mathbf{P}_i^{t+1} = \mathbf{P}_i^t + (1 - \delta)(\mathbf{P}_i^t - \mathbf{P}_i^{t-1}) + \alpha \mathbf{f}_i^{\text{int}} + \beta \mathbf{f}_i^{\text{ext}} \quad (12.1)$$

A description of those forces are the topic of the next two sections.

12.2.1 Internal Forces on Simplex Meshes

Local Internal Forces

There are two types of prior knowledge that can be applied. A *weak* shape prior consists in assuming that the structure of interest is smooth, which is a reasonable assumption for many anatomical structures (liver, heart, bones...) but not all of them (e.g., infiltrating tumor, brain gray matter). With a *strong* shape prior, it is assumed that the structure of interest has a typical shape, i.e. that its shape varies to a “small extent” around a reference shape.

Simplex meshes are well suited to enforce both types of priors. Indeed, one can define at each vertex of a 2-simplex mesh, discrete geometric quantities that locally control the shape of the mesh. At vertex \mathbf{P}_i those quantities are the three metric parameters $(\epsilon_{1i}, \epsilon_{2i}, \epsilon_{3i})$ and the simplex angle ϕ_i (Fig. 12.3, left). Metric parameters control the relative spacing of vertices through a tangential internal force while the simplex angle controls the local surface smoothness through a normal internal force. By combining those two tangential and normal components of the internal force, it is possible to define weak and strong shape priors in a simple way.

For instance, Fig. 12.3 (right) shows the effect of a specific tangential force that leads to a concentration of vertices at parts of high curvatures. In Fig. 12.4, the tangential component impose metrics parameters equal to $1/3$, while the normal component of the internal forces smooths the surface with normal and mean curvature continuity. In some restricted cases, discrete mean curvature on a simplex mesh has been proved [5] to converge towards its continuous value.

The straightforward definition of forces imposing mean curvature continuity is a major advantage of simplex meshes. Indeed, with such regularizing

constraint, no shrinking effect is observed unlike the mean curvature flow [6] widely used in the level-sets method [7]. With those flows minimizing surface areas, it is often necessary to limit the shrinking effect by adding a “balloon force”, which may lead to additional problems [8].

To provide a strong shape prior, a specific internal force has been proposed that constrains the metrics parameters and simplex angle to be equal to some user-defined values corresponding to a reference shape. Under the influence of this shape-memory force, a simplex mesh is attracted towards a reference shape up to a translation, rotation and scale. This local representation of shape is also a unique feature of simplex meshes compared to other representations.




Fig. 12.3. Metric parameters. Left: point P and its three neighbors on a 2-simplex mesh define a circumscribed sphere. The simplex angle in P is related to the height $\|FP\|$ of point P with respect to triangle A, B, C such that the angle stays constant when P lies on the circumsphere. The metric parameters $(\epsilon_1, \epsilon_2, \epsilon_3)$ are simply the barycentric coordinates of F in the triangle A, B, C ; Middle: cube simplex mesh with metric parameters equal to $1/3$; Right: mesh with optimized metric parameters so that vertices are concentrated at parts of high curvature.




Fig. 12.4. Regularization of a simplex mesh. Left: simplex mesh with zero mean curvature; Middle: the red part is smoothed with a mean curvature continuity to the four cylinders; Right: normal continuity between the surface and 3D contours can also be enforced.

12.2.2 Image Forces

Image forces drive a deformable simplex mesh towards the apparent boundary of a structure of interest and rely on some prior knowledge about the appearance of that structure. In all cases, the external force is formulated as displacement in the normal direction towards the closest boundary point \mathbf{C}_i of the structure: $\mathbf{f}_i^{\text{ext}} = ((\mathbf{C}_i - P_i) \cdot \mathbf{n}_i) \mathbf{n}_i$. Again one can rely on a weak appearance prior that assumes generic boundary features (such as gradient maxima) or on a strong appearance prior that relies on machine learning trained on segmented dataset. In fact, there is a continuum of priors between those two extremes:

- *Contour prior*: The closest boundary points are searched among the maxima of the gradient norm in the gradient direction. Additional criteria may be used to select the proper boundary point such as intensity values or the dot product between the surface normal \mathbf{n}_i and the gradient vector at that point [9].
- *Contour and region prior*: The principle is to characterize the inner (or outer) region by a range of intensity. The boundary is then determined as the voxel located next to this region whose gradient norm is greater than a threshold [10].
- *Region histogram prior*: From a collection of segmented structure, one can build the probability density function that a pixel belongs to the structure from the region histogram (using Parzen windowing for instance). The closest point is then determined as the one which maximize (resp. minimize) the probability for the inner (resp. outer) voxels to belong to the inner region [11].
- *Intensity profile matching prior*: At each mesh vertex, an intensity profile along the normal direction is extracted. The external force is computed by searching along the normal direction the displacement which maximize the correlation with a reference intensity profile. That reference profile may result from a Principal Component Analysis (PCA) as described in Sec. 12.3.4. Several criteria may be used such as the sum of the square differences, the linear correlation or the Mahalanobis distance.
- *Block matching prior*: This approach is similar to the intensity profile approach but based on the matching of image blocks rather than intensity profiles. The main difference is that the search for the closest point is no longer constrained to be along the normal direction but may be in all directions. This is of interest when homologous points can be extracted from images (such as points of high curvatures or anatomically relevant features).
- *Texture classification prior*: The image force can be applied after performing some image classification based on some machine learning method. The output of the classification is a probability map that a voxel belongs to the structure of interest. The closest point along the normal direction

whose probability is equal to 50% may then be used to compute the image forces. Linear classifiers, support vector machines [12] or even neural networks [13] can be trained on segmented datasets to classify the image.

12.2.3 Globally Constrained Deformation

As mentioned previously, the first stage of model-based segmentation consists in positioning the deformable model in the image through various methods, manual or automatic. Due to the convergence to a local minimum of the functional, the second stage corresponding to the optimization of the model parameters may not give good segmentation results if the initialization is not close enough to the final shape. One common strategy to decrease the sensitivity of the optimal shape to the initialization position is to use a coarse-to-fine strategy. In this approach, the complexity (in terms of DOF) of the model is slowly increased in order to start the optimization with a small number of local minima and then track the global optimum during the increase of complexity.

A first method to control the DOF consists in using a multi-resolution approach where a coarse mesh is deformed on a given image and replaced with a finer mesh when convergence is reached. This technique is fairly straightforward to implement but requires to create meshes of increasing resolution which is not trivial on unstructured meshes [14]. On simplex meshes, the $\sqrt{3}$ subdivision scheme is well suited to globally or even locally refine the mesh [15]. One limitation of multi-resolution schemes is that the coarser mesh




Fig. 12.5. Deformable simplex meshes. The meshes are driven by the matching of intensity profiles; Upper row: reference CT image scan (left) and head segmented with a simplex mesh (right). Lower row: an intensity profile is stored at each vertex as a reference and this mesh is used to register the head of the same patient on an MRI: mesh position before (left) and after rigid registration (right).

Table 12.1. Geometric transforms and constrains. The table shows geometric transforms that can constrain the deformation of a simplex mesh.

Transform	Number of DOF	Transformation type
Rigid	6	global
Similarity	7	global
Affine	12	global
B-spline	$3 \times$ number of control points	restricted to a bounding box
Axis-symmetric	$3 \times$ number of points on the central line	restricted to the mesh
PCA	$6 +$ number of modes	restricted to the mesh

may still have too many DOF to prevent falling into a local minimum of the functional.

An alternative is to use globally constrained deformations [16], where global spatial transforms are used in the early stages of segmentation in order to constrain the DOF. Table 12.1 lists some transforms that may be applied with a growing number of DOF. Starting from a rigid transform, it is possible to use only six DOF which greatly constrain the optimization problem.

However, it is often difficult to smoothly increase the number of DOFs, except when using the PCA (Sec. 12.3.3). To this end, the concept of globally constrained deformation has been introduced [16] that can tune the complexity of the deformation with a locality parameter λ . More precisely, the motion equation (12.1) is refined by adding a global force $\mathbf{f}_i^{\text{global}} = \mathcal{T}(\mathbf{P}_i^t) - \mathbf{P}_i^t$ computed from the application of a global transform \mathcal{T}

$$\mathbf{P}_i^{t+1} = \mathbf{P}_i^t + (1 - \delta)(\mathbf{P}_i^t - \mathbf{P}_i^{t-1}) + \lambda(\alpha\mathbf{f}_i^{\text{int}} + \beta\mathbf{f}_i^{\text{ext}}) + (1 - \lambda)\mathbf{f}_i^{\text{global}} \quad (12.2)$$

When $\lambda = 0$ the simplex mesh deforms according to the application of a global transformation (with few DOFs) whereas when $\lambda = 1$, the deformation is only driven by local internal and external forces. When $0 < \lambda < 1$ the deformation is a trade-off between global and local behaviors corresponding to an intermediate number of DOFs. The transformation \mathcal{T} is computed such that it minimizes the discrepancy between the vertex position \mathbf{P}_i^t and its closest point on the image \mathbf{C}_i^t .

In many cases, it is sufficient to apply global transforms from rigid to affine and then to slowly increase the locality parameter from 0 to 1 or to a smaller value (Fig. 12.6). The concept of globally constrained deformation may also be successfully applied to enforce to some degree the axis-symmetry around a deformable axis (for instance to segment vessels) or to take into account the statistical shape variation around a mean shape.

12.2.4 3D+T Deformable Simplex Meshes

For the analysis of time series of medical images, e.g., gated cardiac imaging, it is important to not only to segment a structure but also to track its deformation over time. A common approach consists in using the result of the segmentation at time t as the initialization of the segmentation at time $t + 1$.

However this approach is prone to the propagation of errors if one image is especially challenging and tends to bias the result noticeably by underestimating the motion of the structure.

Rather than processing the images one after the other, 3D+T deformable models [17] can be used to process all images of the time series at once (if they can fit in the memory of the computer). In this framework, a family of simplex meshes having the same topology (Fig. 12.7) are deformed on each image of the time series, the mesh deformation at time t being coupled with that of time $t - 1$ and $t + 1$.

The temporal coupling forces act as additional internal force and can be formulated in a weak or a strong way. The weak prior consists in minimizing the kinetic energy of the system and thus tries to make vertex \mathbf{P}_i^t close to both points \mathbf{P}_i^{t-1} and \mathbf{P}_i^{t+1} . With a strong motion prior, the shape of the trajectories of each vertex are constrained to be similar to some reference trajectory curves, described with geometric quantities that are invariant to translation, rotation and scale. Examples of segmentation of time-series of echocardiographic images are shown in Fig. 12.8.





Fig. 12.6. Coarse-to-fine strategy of 3D mesh reconstruction. The coarse to fine strategy runs from clouds of points: the different stages of the deformations with increasing deformation complexity are shown ending with a globally constrained deformation ($\lambda = 0.3$ in panel f).

Fig. 12.7. 3D+T simplex mesh. The array of meshes has the same topology but different geometries.



12.2.5 Advanced Segmentation Strategies

The segmentation of anatomical structures from medical images is often a difficult problem due to their low contrast with surrounding structures, the inter-subject variability, the presence of pathologies, etc. Problem specific segmentation strategies must be defined in order to achieve a robust and accurate delineation.

Atlas-Based Initialization

The first stage of model-based segmentation consists in finding the location of the structure of interest in the image. This can be done manually by the user but the inter-user variability in providing this information may be very large and may impair the reproducibility of the method. An interesting alternative is to use non-rigid registration [12] to find a geometric transformation between




Fig. 12.8. Segmentation of the left ventricle endocardium. Segmentation from 3D+T echo-cardiographic images by 3D+T simplex meshes.

a reference image (or an atlas) and the target image. In the reference image, meshes of the structures of interest can then be mapped by this transformation in the space of the target image and finally serve as initial shape for model-based segmentation. This approach has the advantage of being automatic (and therefore reproducible) and usually provides good model initializations. However, when the structure location varies greatly, it is required to provide a suitable initialization of the non-rigid registration algorithm, which can be achieved by performing some coarse exhaustive search [18].

Multi-Structure Segmentation

It is often a good strategy to first segment the most contrasted structures and then proceed with more challenging ones (least contrasted). Indeed, already delineated structures can restrict the possible shape of neighboring structures during the model-based segmentation of other structures. Furthermore, when boundaries between structures are hardly visible, it is best to segment both structures at the same time, the shape of the former providing inter-penetration constraints [19] for the latter. Those constraints are based on the computation of distance maps measuring the distance of a given points to the two surface meshes. Fig. 12.9 shows an example of the joint segmentation of bladder and prostate.

Rule-based segmentation

Ideally, all model-based segmentation algorithms should proceed in two stages: initialization followed by an optimization stage. However, most segmentation algorithms are meant to be applied on patients with some pathologies (e.g., tumor, metastasis) or some rare anatomical configurations (e.g., abnormal branching of vessels, already ablated tissue). In such case, it is difficult to




Fig. 12.9. Joint segmentation of bladder and prostate. The bladder (red) and prostate (green) are shown in a CT image. Left: no inter-penetration constraints are used which causes a leakage; Middle: with inter-penetration constraint and some statistical shape constraint on the prostate, the segmentation is improved despite the low contrast between structures; Right: 3D view of the two organs.

build a shape and appearance model that can cope with such variation of cases. Therefore, to improve the robustness of the segmentation, some alert mechanisms should be included in order to detect that the segmentation is not performing well. For instance, for segmentation of hippocampi in brain MRI, a projection on the main PCA mode was used [20] to detect the leakage of a deformable model thus signaling that the locality parameter λ should be decreased. Additional rules and meta-rules can be used to test the validity of hypothesis and eventually reset the optimization stage with new parameters.

12.2.6 Geometric Representations for Model-Based Segmentation

The choice of a geometric representation of a model has a significant impact on the performance of a model-based segmentation algorithm. Many components of those models can be implemented independently of the representation, however some implementations may be far easier than others. For instance, PCA of the shape variability may be implemented for both point-set distributions [21] (Sec. 12.3.1) and level-set functions [22] but are far simpler for the former.

Simplex Meshes vs. Triangulations

Triangulations are widely used representations for and computational purposes. They can represent manifold and non-manifold surfaces while simplex meshes are only suitable for manifold surfaces. Furthermore, it is straightforward to compute the closest distance between a point to a triangulation because of its planar faces (triangles). On simplex meshes, the faces are curved and therefore computing accurately the closest point is more complex (no closed form solution). However, due to its fixed vertex connectivity, it is much easier to control the spreading of vertices on simplex meshes for instance to concentrate vertices at parts of high curvature. Also one can define on simplex meshes shape memory regularization as well as curvature continuous smoothing without shrinkage. On triangulations, the most common way to smooth a shape is to use the mean-curvature flow [23] which entails shrinkage.

Simplex Meshes vs. Parametric Surfaces

Parametric surfaces, such as B-splines, the Non-Uniform Rational B-Spline (NURBS), or Finite Element Model (FEM), can provide C^k continuity ($k > 0$) everywhere while for simplex meshes only C^0 continuity is enforced. Therefore, they describe smoother surfaces which can be important for some applications (e.g., rapid prototyping, physical simulation). Also, differential quantities such as surface normals or curvatures can be described everywhere, which is important for shape analysis (e.g., extraction of feature points). However, using those parametric surfaces brings additional complexity especially when the

surface cannot be parameterized by a single planar patch (i.e., surface of planar, cylindrical and toroidal topology). Indeed, even for surfaces of spherical topology, parametric surfaces may be complex to handle for instance with the creation of poles (deformable super-quadrics [3]). Simplex meshes as well as triangulated surfaces can cope with surfaces of any topology. Furthermore in many image segmentation problems, only the position, volume and distance of the segmented structure are of interest. Therefore, describing surfaces with C^0 continuity is a good trade-off between complexity and accuracy.

Simplex Meshes vs. Level-Sets

Level-sets level-set rely on an implicit description of the geometry. They can describe manifolds of co-dimension one (i.e., hyper-surface) of any topology but cannot represent surfaces with borders (except if the border is on the edge of the image) such as a planar patch or a cylinder. Simplex meshes can represent a manifold of any dimension and co-dimension. In level-sets, differential quantities (normal, curvatures) may be computed everywhere by filtering the embedding image while in simplex meshes they are only available at vertices. Level-sets can adapt their topologies easily while it is much more difficult for simplex meshes [24] or triangulations [25]. Also, the coupled segmentation of structures is easier to implement with level-sets since the distance function is already computed to update their position. Since they are based on an Eulerian scheme, level-sets are not suited to transport *a-priori* information from a model to the segmented structure. Also, the regularization of level-sets is usually local and based on the mean curvature flow which entails shrinkage.

12.3 Statistical Models of Shape and Appearance

As presented in Sec. 12.2.1, shape priors can be employed as stabilizing forces during the evolution of a deformable model. With strong shape priors, it is possible to enforce smooth shape changes around a defined template. Apart from the smoothness constraint, vertices generally move independently from one another, driven by their individual image forces. For many structures in the human body, it is possible to constrain the possible variations further, creating more specific models while still maintaining a good ability to adapt to individual instances. The higher the specificity of a model, i.e. the fewer non-plausible shapes it produces, the more robust it can segment images in the presence of noise and artifacts.

To determine plausible and non-plausible shapes for a certain structure requires either an artificial parametric model (which has to be hand-crafted each time), or a large amount of examples from which this information can be extracted automatically. The SSM employs the latter approach; it generally uses second order statistics to extract the principal modes of variation from a set of example shapes. These modes can be used instead of the internal

forces of deformable models. In addition, it is also possible to employ statistics to generate more specific image forces. As for shape modeling, the required information is extracted from a set of example images. Using both techniques together allows to develop very specific and accurate segmentation algorithms for medical applications. A more detailed discussion of the techniques and approaches is given in the comprehensive review by Heimann and Meinzer [26].

12.3.1 Shape Representation

As mentioned in Sec. 12.2.6, there exist several different representations for shapes. In the following, we assume the most common and simplest representation: a cloud of points densely sampled on the surface of each shape. Thus, each training shape can be represented as a single vector \mathbf{x} that concatenates coordinates from all k points on the surface

$$\mathbf{x} = (x_1, y_1, z_1, \dots, x_k, y_k, z_k)^T \quad (12.3)$$

Connecting edges between neighboring points leads to a mesh structure. Although these connections are not required for a Point Distribution Model (PDM) [21], they are essential to define inside, outside, and normal vectors for segmentation applications.

12.3.2 Point Correspondence

To be able to perform statistical analysis on a set of training shapes, the individual points representing each training shape have to correspond over the entire set of examples (Fig. 12.10). This means that if a point is placed on a certain anatomic landmark in one example, it should be placed on the same landmark in all other shapes (and the number of sample points is necessarily equal for all training shapes). Thus, all points do represent a kind of landmark, and the two terms are often used synonymously in shape modeling.

Setting landmarks manually on 3D shapes is an extremely tedious and time-consuming process. Additionally, results are not reproducible, as different experts will label different landmarks, especially in areas without discernible features. For this reason, methods to automatically determine correspondences are of paramount importance to shape modeling, and the quality of extracted correspondences has a large impact on the quality of the resulting shape model. The variety of approaches to determine correspondences over a set of shapes can be classified according to the type of registration used to match one shape to the other. In the following sections, we briefly present the most important approaches.

Mesh-to-Mesh Registration

Assuming that the original training shapes are available as meshes, a straightforward solution to the point correspondence problem is to register the different surfaces to each other. One training shape is selected as template and

matched to all others by a generic surface matching method as Iterative Closest Point (ICP) [27] or Softassign Procrustes [28]. For each point of the template, these algorithms deliver the closest point in the target mesh after both have been aligned by an optimal similarity transform. However, this solution has a number of draw-backs. Firstly, the selection of the template induces a bias in the process; different templates usually lead to different correspondences. Secondly, possible corresponding points are limited by the resolution of the target mesh: If this mesh is not sampled densely enough, the approach can introduce considerable artificial variation in the resulting SSM. Finally, the most serious problem is that the employed similarity transform may be insufficient to match the template to the other training shapes if the training set shows large geometric variations. In this case, the approach leads to invalid correspondences and the resulting statistical model is not representative of the training data.

An alternative approach is to use non-rigid registration to match the training surfaces. For this, a limited number of matching landmarks is determined in advance, either by manual labeling or by automatic feature matching. These sparse point sets can then be used to initialize a non-rigid registration as the thin-plate-spline deformation, which produces exact matches for the pre-defined landmarks. Correspondences for the remaining points can be determined simply by Euclidean distance. However, to avoid topological problems it is safer to regularize this matching, e.g., by mesh relaxation [29].




Fig. 12.10. Correspondences for a set of liver shapes. Corresponding areas are marked by colors and a sparse coordinate grid.

Mesh-to-Volume Registration

Most training shapes for medical imaging applications originate from volumetric image modalities as CT or MRI. In these cases, it is possible to acquire point correspondences by registering a deformable mesh to each volumetric training image. The basic idea is that the positions to which vertices of the deformable mesh converge are corresponding over the entire training set. Obviously, the deformable mesh has to adapt accurately to each shape, which is difficult when using one of the typical image forces (Sec. 12.2.2) on intensity images. Therefore, the deformable mesh is mostly used with the segmented, binary volumes, which provide a hard feature for the image force. The geometry for the deformable mesh can be extracted from an arbitrary instance of the training set. However, this procedure can introduce a bias in the corresponding detection. A technique to minimize this bias has been proposed in [30], where intermediate templates are generated for difficult cases.

Volume-to-Volume Registration

Similar to registering a deformable mesh to all training volumes, it is also possible to employ a volumetric atlas for that purpose. The atlas can either consist of a single volume or can be built from a collection of images. After registering the atlas to an instance of the training set, the resulting deformation field is used to propagate landmarks on the atlas to the respective training shape. As with mesh-to-volume registration, the atlas registration is usually conducted on the segmented volumes to increase accuracy. To provide the required flexibility, non-rigid transforms are essential for the registration; popular methods include B-splines or thin-plate-splines deformations. As usual, a multi-resolution approach increases robustness of the matching.

Parameterization-to-Parameterization Registration

A parameterization is a bijective mapping between a mesh and an appropriate base domain [31]. For closed 2D contours, this base domain is the circle, and determining correspondences by parameterization is equivalent to matching points by relative arc-length. Depending on the topology of training shapes the respective base domains can change, which makes registration between them impossible. This means that a prerequisite for using this approach is that all training shapes have to feature the same topology. Please note that this is also true for the previously presented registration approaches, with the difference that this constraint is not verified there. For 3D shapes, most parameterization approaches are geared to S^2 , i.e. 2-manifolds without holes or self-intersections. The corresponding base-domain is the sphere. When using unconstrained parameterizations as spherical harmonics [32], the obtained correspondences are often arbitrary. Therefore, a common approach is to determine a limited set of matching landmarks in advance (as for the mesh-to-mesh registration) and employ to match the remaining regions of the shape.

Population-Based Optimization

In all the correspondence approaches presented above, matching points are determined in a pair-wise comparison or without a comparison at all (for the parameterization-based approaches). Population-based correspondence methods analyze the complete set of training shapes at once and select corresponding points based on a global cost function. Starting from initially arbitrary landmarks, an optimization algorithm iteratively modifies correspondences until the cost function converges at its minimum. Although this procedure sounds straight-forward and should lead to optimal results, there are a number of challenges involved in the design.

First, as it is difficult to define what *good correspondences* are, coming up with a practically usable and theoretically sound cost function is not evident. Practically all current approaches rely on a form of Minimum Description Length (MDL) as cost function, which was pioneered for correspondence optimization by Davies [33]. MDL implements the principle of Occam’s razor that a simpler model is better than a complex one. Translated to SSMs, this means that a model is better when it features fewer and more compact modes of variation. As the cost function for the full MDL is computationally very expensive, a simplified version [34] is often used. Here, costs F are directly calculated from the eigenvalues λ of the shape model

$$F = \sum_m \mathcal{L}_m \quad \text{with} \quad \mathcal{L}_m = \begin{cases} 1 + \log(\lambda_m/\lambda_{\text{cut}}) & \text{for } \lambda_m \geq \lambda_{\text{cut}} \\ \lambda_m/\lambda_{\text{cut}} & \text{for } \lambda_m < \lambda_{\text{cut}} \end{cases} \quad (12.4)$$

where λ_{cut} describes the amount of noise in the dataset.

A second challenge is to adjust the landmark positions in the course of optimization. The difficulty lies in the fact that the optimization problem is heavily constrained, as individual landmarks cannot be moved independently from their neighbors (to avoid flipping triangles in the model mesh). Additionally, all 3D landmarks have to remain on the 2D manifolds of the surfaces of the training shapes. This is either accomplished by employing parameterizations to modify landmark positions [33, 35] or by particle approaches [36].

12.3.3 Construction of Statistical Shape Models

Once point correspondences are known, each training shape i can be represented as coordinate vector \mathbf{x}_i in (12.3). Before applying the dimensionality reduction that delivers the predominant modes of variation, the individual shapes first have to be aligned into the same coordinate system.

Alignment

The commonly used procedure to align a set of training shapes is the Generalized Procrustes Analysis (GPA) [37]. The underlying Procrustes analysis

determines the optimal similarity transform to minimize the sum of squared distances between two corresponding point sets. GPA employs this method iteratively to match a whole set of shapes to their unknown mean. An additional step for constructing a shape model is a rescaling of the thus aligned shapes. This is required since the minimization of squared distances introduces non-linearities in the shape space, which are difficult to capture by commonly used techniques. A straight-forward solution is to scale each training shape i with $1/(\mathbf{x}_i \cdot \bar{\mathbf{x}})$, where $\bar{\mathbf{x}}$ is the mean as determined by GPA.

Dimensionality Reduction

Statistical shape modeling is essentially a problem of dimensionality reduction: The input data of s training shapes \mathbf{x}_i (with the dimension of \mathbf{x}_i generally in the thousands) should be reduced to a limited number of modes of variation. This task is usually solved by Principal Component Analysis (PCA) [38]. The first step is to calculate the mean shape $\bar{\mathbf{x}}$ of the training set

$$\bar{\mathbf{x}} = \frac{1}{s} \sum_{i=1}^s \mathbf{x}_i \quad (12.5)$$

Then, the covariance matrix S of the input data is computed by:

$$S = \frac{1}{s-1} \sum_{i=1}^s (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \quad (12.6)$$

An eigendecomposition of S yields the eigenvectors ϕ_m (representing the principal modes of variation) and the corresponding eigenvalues λ_m (indicating the variance per mode). Sorting all modes from largest to smallest variance, the first c modes are employed to model the distribution, while the remaining modes are discarded. Thus, all valid shapes can be approximated by a linear combination

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{m=1}^c b_m \phi_m \quad (12.7)$$

with \mathbf{b} as the vector of shape parameters. To constrain variation of the model, \mathbf{b} is usually bound to certain limits, either by constraining $|b_m| < 3\lambda_m$ for each mode individually or by constraining b globally to lie within an appropriate ellipsoid. Typically, c is chosen so that the accumulated variance $\sum_{m=1}^c \lambda_m$ reaches a certain percentage (e.g. 95%) of the total variance. Fig. 12.11 shows the first three principal modes of variation as extracted from a collection of liver shapes.

12.3.4 Modeling Object Appearance

In Sec. 12.2.2, we already described possible image forces for deformable models. As we assume to have training data available for statistical modeling, we now focus on strong appearance priors.

Boundary-Based Features

The most popular appearance models for SSMs are related to the above-mentioned intensity profile. At each landmark of the shape model, a profile is sampled in normal direction and compared to a model built from training data at the same landmark. This implies that one shape model commonly features a large number of appearance models. The employed profiles are not limited to plain intensity alone; gradient profiles and normalized intensity or gradient profiles are also popular choices. Different features can also be combined in the same model. Each appearance model is built similar to the shape model, i.e. mean profile and covariance matrix are generated. Using the covariance matrix, the Mahalanobis distance to profiles sampled during the search can be determined. For Gaussian distributions, this represents a robust measure of the goodness of fit [39].

In cases where the distribution of profiles is essentially non-Gaussian (e.g., because the modeled object is adjacent to several tissues of different appear-

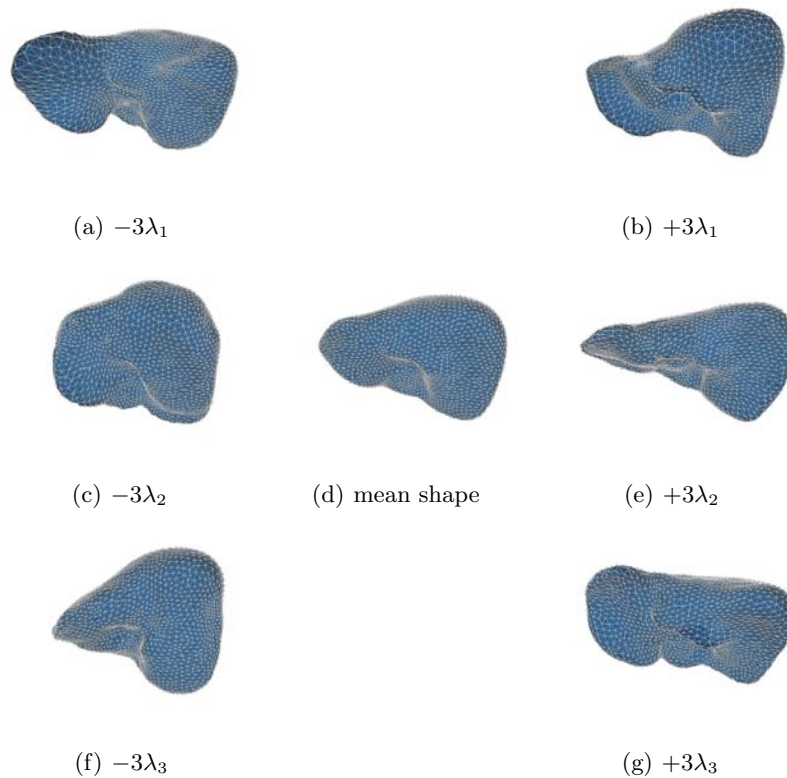


Fig. 12.11. Statistical shape model of the liver. Panels (a,b), (c,d), and (f,g) show the first three modes of variation for a SSM of the liver with mean shape (d).

ance), non-linear appearance model can deliver better results. A straight-forward method to employ a non-linear model is using a k-Nearest-Neighbor (k-NN) classifier [40]. To train it, both true and false examples of boundary profiles must be used. The required false examples can be generated by shifting profiles from their known positions towards the inside and outside of the Object of Interest (OOI).

Region-Based Features

As mentioned in Sec. 12.2.2, region histogram priors can be used to evolve a deformable model in order to match a certain distribution of intensities. For statistical models, it is straight-forward to model this distribution from training data. However, there are a number of much more specific features that can be used in this case. The most consequent approach for learning object appearance uses the entire texture of the OOI [41]. To combine textures from all training images, the respective objects are morphed to the same shape (the mean shape of the SSM). Subsequently, the mean texture and its covariance matrix can be calculated as usual. A drawback of this approach is the high memory footprint of the appearance model, especially for 3D models where a volumetric texture is used. One solution is to model the texture only in certain parts of the object, e.g. around the boundary.

Modeling local regions inside and outside of the object of interest opens up more possibilities regarding the search algorithm for model evolution (Sec. 12.3.5), as the goodness of fit can be determined locally (like when using profiles). In addition to the above mentioned techniques of histogram matching and texture matching, local appearance can also be modeled by multi-scale derivatives, Gabor filterbanks, wavelets, etc. In comparison to boundary-based features, region-based features are often more robust as they use more image data to determine the goodness of fit. At the same time and for the same reason, they are also computationally more demanding.

12.3.5 Local Search Algorithms

As mentioned in Sec. 12.2.5, the first step to employ a deformable model for segmentation is to initialize it for the respective image. This initialization places the model close to the object of interest (e.g., by atlas registration), so that the model can iteratively evolve towards the correct contour. In the following, we present several techniques to conduct this subsequent evolution (the local search) for statistical models of shape and appearance.

Active Shape Model

The Active Shape Model (ASM) was introduced by Cootes et al. [21]. It is one of the most popular model-based approaches for segmentation of medical

images. An instance \mathbf{y} of the model in an image is defined by a similarity transform T and the shape parameter vector \mathbf{b}

$$\mathbf{y} = T(\bar{\mathbf{x}} + \Phi\mathbf{b}) \quad (12.8)$$

where $\Phi = (\phi_1 \dots \phi_c)$ is the matrix of eigenvectors (12.7). To evolve the shape over time, the optimal displacement \mathbf{dy}_p is determined for each landmark point separately by querying the goodness of fit for the respective appearance model at various positions along the . The update then consists of two steps: First, transform T is updated by the result of the Procrustes match of the model to $\mathbf{y} + \mathbf{dy}_p$. This leads to new residual displacements \mathbf{dy}_s . Second, shape parameters \mathbf{b} are updated by transforming \mathbf{dy}_s into model space and then applying a projection into shape space

$$\mathbf{db} = \Phi^T \tilde{T}^{-1}(\mathbf{dy}_s) \quad (12.9)$$

where \tilde{T} is equal to T without the translational part. The resulting adjustments \mathbf{db} are added to \mathbf{b} . Usually, \mathbf{b} is restricted to certain limits to maintain a valid shape (Sec. 12.3.3). The presented two steps are repeated iteratively until the model converges at the best local fit in the image.

To increase , especially against incorrect initialization, the search is commonly conducted in a multi-resolution fashion. Contrary to the multi-resolution approach presented for simplex meshes (Sec. 12.2.3), the DOF usually remains unchanged. Instead, appearance models are constructed for a number of down-sampled versions of all training images, and the search is started on a down-sampled version of the original image. Here the search radius is larger, since the goodness of fit is evaluated at points at a greater distance from the original contour. When the model evolution converges on the rough scale, appearance model and search radius are switched to the next (more detailed) versions, until the original resolution is reached.

Another approach to increase robustness is to decrease the effect of outliers, i.e. wrongly detected optimal displacements for specific landmarks. There are two basic techniques for this problem. Firstly, outliers can be detected (e.g. by comparing suggested landmark positions to collected statistics) and corrected before the shape is updated. Secondly, the update of transform and shape parameters can be weighted and each suggested landmark position receives a reliability weighting in each step.

Active Appearance Model

The Active Appearance Model (AAM) [41] belongs to the class of generative models, i.e., they can generate realistic images of the modeled data. This is accomplished by storing a complete texture model, consisting of the mean and the principal modes of variation, in addition to the shape model. However, AAMs are more than a shape model with region-based features, they employ a specific search method which is completely different from the ASM

approach. Moreover, shape and appearance parameters are combined into one linear system, where shape \mathbf{x} and appearance \mathbf{g} are described by a common parameter vector \mathbf{c}

$$\begin{aligned}\mathbf{x} &= \bar{\mathbf{x}} + \Phi_s W_s Q_s \mathbf{c} \\ \mathbf{g} &= \bar{\mathbf{g}} + \Phi_g Q_g \mathbf{c}\end{aligned}\tag{12.10}$$

Here, Φ_s and Φ_g are the independent eigenvector matrices of shape and appearance model, respectively, and W_s is a diagonal weight matrix for the shape parameters. $Q = \begin{pmatrix} Q_s \\ Q_g \end{pmatrix}$ is the eigenvector matrix of the combined shape and appearance parameters. It is the result of a PCA on the independent parameters $\mathbf{b} = \begin{pmatrix} W_s \mathbf{b}_s \\ \mathbf{b}_g \end{pmatrix}$. An instance of the model in an image is defined by a similarity transformation T and the combined shape-appearance parameters \mathbf{c} . In the following, we name the latter ones simply parameters \mathbf{p} . To evaluate the goodness of fit, the image texture is warped to the mean shape and normalized, resulting in \mathbf{g}_s . With the modeled appearance $\mathbf{g}_m = \mathbf{g}$ from (12.10), the residuals are given by $\mathbf{r}(\mathbf{p}) = \mathbf{g}_s - \mathbf{g}_m$, and the error by $E = \mathbf{r}^2$.

The key idea of AAM search is to assume a constant relationship between texture residuals $\mathbf{r}(\mathbf{p})$ and parameter updates \mathbf{dp} over the entire search

$$\mathbf{dp} = -R\mathbf{r}(\mathbf{p})\tag{12.11}$$

The success of this optimization scheme largely depends on the derivative matrix R . In the first presentation of AAMs, R was computed using multivariate linear regression on a large number of simulated disturbances of the training images. Later, regression was replaced by numeric differentiation, claimed to be both faster and more reliable.

The main challenge when employing AAMs for medical image segmentation is the enormous amount of data that has to be captured by the model. For 3D models, appearance is generally modeled as a 3D texture, which quickly leads to prohibitively large equation systems. Therefore, texture resolution generally has to be scaled down. It is noteworthy that the AAM search procedure is not limited to region-based features. Indeed, the general idea of determining parameter updates from residuals can also be used with alternative appearance features as described in Sec. 12.3.4.

12.4 Conclusion

Model-based approaches to segmentation are arguably the most robust methods when image data is noisy or includes artifact. Therefore, they are a prime choice for many applications in medical image analysis. The deformable simplex meshes and point-based statistical models presented in this chapter represent state-of-the-art techniques for this purpose.

Comparing both methods, an obvious difference is that building a shape model requires a large collection of training images (typically 20–50 as a minimum for 3D models), while a simplex mesh with strong shape prior can be

based on a single example. Another point is that the additional constraints of SSMs generally result in a higher robustness, but at the same time limit the accuracy of the final result. Statistical models are therefore often applied as part of a coarse-to-fine strategy that features an additional refinement step afterwards. As already hinted to in Sec. 12.2.3, it is also possible to combine both techniques in one framework. That way, the strict shape constraints can gradually be lifted and the final free deformation ensures an optimal fit to the data.

References

1. Terzopoulos D. Multiresolution computation of visible-surface representation. MIT. Cambridge, MA, USA; 1984.
2. Kass M, Witkin A, Terzopoulos D. Snakes: active contour models. *Int J Comput Vis.* 1988;1:321–31.
3. Metaxas D, Terzopoulos D. Constrained deformable superquadrics and nonrigid motion tracking. In: *Proc CVPR.* Maui, Hawaii; 1991. p. 337–43.
4. Delingette H, Hébert M, Ikeuchi K. Shape representation and image segmentation using deformable surfaces. *Image Vis Comput.* 1992;10(3):132–44.
5. Delingette H. Modélisation, déformation et Reconnaissance d'objets tridimensionnels a l'aide de maillages simplexes [Thèse de sciences]. Ecole Centrale de Paris; 1994.
6. Chopp DL, Sethian JA. Computing minimal surfaces via level set curvature flow. *J Comp Phys.* 1993;106:77–91.
7. Malladi R, Sethian JA, Vemuri BC. Shape modeling with front propagation: a level set approach. *IEEE Trans Pattern Anal Mach Intell.* 1995;17(2):158–75.
8. Cohen LD, Cohen I. Finite element methods for active contour models and balloons from 2-D to 3-D. *IEEE Trans Pattern Anal Mach Intell.* 1993;15(11).
9. Delingette H. General object reconstruction based on simplex meshes. *Int J Comput Vis.* 1999;32(2):111–46.
10. Montagnat J, Delingette H, Malandain G. Cylindrical echocardiographic images segmentation based on 3D deformable models. *Lect Notes Computer Sci.* 1999;1679:168–75.
11. Herbulot A, Besson SJ, Duffner S, et al. Segmentation of vectorial image features using shape gradients and information Measures. *J Math Imaging Vis.* 2006;25(3):365–86.
12. Pitiot A, Delingette H, Thompson PM, et al. Expert knowledge guided segmentation system for brain MRI. *NeuroImage.* 2004;23(Suppl 1):S85–S96.
13. Pitiot A, Toga A, Ayache N, et al. Texture based MRI segmentation with a two-stage hybrid neural classifier. In: *World Congr Comput Intell;* 2002. .
14. Ciampalini A, Cignoni P, Montani C, et al. Multiresolution decimation based on global error. *Vis Comput.* 1997;13(5):228–46.
15. Kobbelt L. Sqrt(3) subdivision. In: *Proc SIGGRAPH;* 2000. p. 103–12.
16. Montagnat J, Delingette H. Globally constrained deformable models for 3D object reconstruction. *Signal Process.* 1998;71(2):173–86.
17. Montagnat J, Delingette H. 4D Deformable models with temporal constraints: application to 4D cardiac image segmentation. *Med Image Anal.* 2005;9(1):87–100.

18. Montagnat J, Delingette H. Volumetric medical image segmentation using shape constrained deformable models. *Lect Notes Computer Sci.* 1997;1205:13–22.
19. Costa J, Delingette H, Novellas S, et al. Automatic segmentation of bladder and prostate using coupled 3D deformable models. *Lect Notes Comp Sci.* 2007;4791:252–60.
20. Pitiot A, Delingette H, Ayache N, et al. Expert-Knowledge-Guided Segmentation System for Brain MRI. *Lect Notes Comp Sci.* 2003;2879:644–52.
21. Cootes TF, Taylor CJ, Cooper DH, et al. Active shape models – their training and application. *Comput Vis Image Underst.* 1995;61(1):38–59.
22. Leventon ME, Grimson WEL, Faugeras OD. Statistical shape influence in geodesic active contours. In: *CVPR; 2000.* p. I: 316–23.
23. Chopp DL, Sethian JA. Flow under curvature: singularity formation, minimal surfaces, and geodesics. *J Exp Math.* 1993;2(4):235–55.
24. Delingette H, Montagnat J. New algorithms for controlling active contours shape and topology. *Lect Notes Comp Sci.* 2000;(1843):381–95.
25. McInerney T, Terzopoulos D. T-snakes: topology adaptive snakes. *Med Image Anal.* 2000;4:73–91.
26. Heimann T, Meinzer HP. Statistical shape models for 3D medical image segmentation: a review. *Med Image Anal.* 2009;13(4):543–63.
27. Besl PJ, McKay ND. A method for registration of 3-D shapes. *IEEE Trans Pattern Anal Mach Intell.* 1992;14(2):239–56.
28. Rangarajan A, Chui H, Bookstein FL. The Softassign Procrustes matching algorithm. *Lect Notes Comp Sci.* 1997;1230:29–42.
29. Lorenz C, Krahnstöver N. Generation of point-based 3D statistical shape models for anatomical objects. *Comput Vis Image Underst.* 2000;77(2):175–91.
30. Zhao Z, Teoh EK. A novel framework for automated 3D PDM construction using deformable models. In: *Proc SPIE.* vol. 5747; 2005. p. 303–14.
31. Floater MS, Hormann K. Surface Parameterization: a Tutorial and Survey. In: *Dodgson NA, Floater MS, Sabin MA, editors. Advances in multiresolution for geometric modelling.* Springer; 2005. p. 157–86.
32. Kelemen A, Székely G, Gerig G. Elastic model-based segmentation of 3-D neuroradiological data sets. *IEEE Trans Med Imaging.* 1999;18(10):828–39.
33. Davies RH, Twining CJ, Cootes TF, et al. 3D statistical shape models using direct optimization of description length. *Lect Notes Comp Sci.* 2002;2352:3–20.
34. Thodberg HH. Minimum description length shape and appearance models. *Lect Notes Comp Sci.* 2003;2732:51–62.
35. Heimann T, Wolf I, Williams TG, et al. 3D Active shape models using gradient descent optimization of description length. *Lect Notes Comp Sci.* 2005;3565:566–77.
36. Cates J, Fletcher PT, Styner MA, et al. Shape modeling and analysis with entropy-based particle systems. *Lect Notes Comp Sci.* 2007;4584:333–45.
37. Goodall C. Procrustes methods in the statistical analysis of shape. *J Roy Stat Soc B.* 1991;53(2):285–339.
38. Jolliffe IT. *Principal Component Analysis.* 2nd ed. Springer; 2002.
39. Cootes TF, Taylor CJ. Using grey-level models to improve active shape model search. In: *Proc ICPR.* vol. 1; 1994. p. 63–7.
40. de Bruijne M, van Ginneken B, Viergever MA, et al. Adapting active shape models for 3D segmentation of tubular structures in medical images. *Lect Notes Comp Sci.* 2003;2732:136–47.

41. Cootes TF, Edwards GJ, Taylor CJ. Active appearance models. *IEEE Trans Pattern Anal Mach Intell.* 2001;23(6):681–85.

List of Acronyms

AAM	Active Appearance Model
ASM	Active Shape Model
DOF	Degree of Freedom
FEM	Finite Element Model
GPA	Generalized Procrustes Analysis
ICP	Iterative Closest Point
k-NN	k-Nearest-Neighbor
MDL	Minimum Description Length
NURBS	Non-Uniform Rational B-Spline
OOI	Object of Interest
PCA	Principal Component Analysis,
PDM	Point Distribution Model,
SSM	Statistical Shape Model

Index

- 3D model, 22
- 3D texture, 22

- AAM, 21, 22
- active contour, 3
- active surface, 3
- appearance model, 20, 21
- appearance prior, 6, 18
- ASM, 20
- atlas, 16
- atlas registration, 20
- atlas-based segmentation, 10

- B-spline, 8, 12, 16
- balloon force, 5
- block prior, 6
- boundary-based feature, 20
- bounding box, 8
- brain MRI, 12

- cell, 3
- coarse-to-fine strategy, 7, 9, 23
- computer vision, 2
- contour prior, 6
- covariance matrix, 19, 20
- CT, 16

- deformable super-quadric, 13
- DOF, 2, 8

- edge, 3
- eigendecomposition, 18
- eigenvalue, 18
- eigenvector, 18, 21

- Euclidean distance, 15
- Eulerian scheme, 13

- face, 3, 4
- FEM, 12

- Gabor filterbank, 20
- Gaussian distribution, 19
- generative model, 21
- genus zero shape, 16
- global transform, 8
- goodness of fit, 19–22
- GPA, 17

- histogram matching, 20
- histogram prior, 6, 20
- hyper-surface, 13

- ICP, 15
- initialization, 11
- inter-subject variability, 10
- inter-user variability, 10
- internal force, 9

- k-NN, 20

- landmark, 15, 16
- landmark point, 21
- level-set, 2, 3, 5, 13
- level-set function, 12
- linear classifier, 7
- linear system, 22
- local appearance, 20
- local search, 20

- locality parameter, 8
- machine learning, 2
- Mahalanobis distance, 6, 19
- manifold surface, 12
- matching prior, 6, 19
- MDL, 17
- mean curvature, 13
- mean shape, 8
- mean-curvature flow, 12
- mesh relaxation, 15
- mesh-to-mesh registration, 16
- mesh-to-volume registration, 16
- meta-rule, 12
- model-based approach, 20
- model-based segmentation, 3
- MRI, 16
- multi-resolution approach, 7, 16, 21
- multi-scale derivative, 20
- multi-structure segmentation, 11
- neural network, 7
- non-rigid registration, 11, 15
- normal vector, 21
- NURBS, 12
- object appearance, 20
- Occam's razor, 17
- OOI, 20
- optimization, 11, 17
- param-to-param registration, 16
- parameter vector, 22
- parameterization, 16
- parametric surface, 12
- Parzen windowing, 6
- PCA, 6, 8, 12, 18
- PDM, 2, 14
- point-set distribution, 12
- population-based optimization, 17
- principal mode, 13
- Procrustes algorithm, 15
- Procrustes analysis, 17
- range data, 3
- region prior, 6
- region-based feature, 20, 22
- registration
 - mesh-to-volume, 16
 - param-to-param, 16
 - volume-to-volume, 16
- regularization, 3, 13
- regularizing constraint, 5
- residual, 22
- rigid registration, 15
- rigid transform, 2, 8
- robustness, 21
- rule-based segmentation, 11
- segmentation
 - atlas-based, 10
 - multi-structure, 11
 - rule-based, 11
- self-intersection, 16
- shape, 20
- shape model, 17
- shape modeling, 14
- shape parameter, 21
- shape prior, 4, 13
- shape space, 21
- shape-memory force, 5
- shrinking effect, 5
- simplex mesh, 3, 12
- snake, 3
- Softassign Procrustes, 15
- spline curve, 2
- SSM, 2, 3, 13, 15, 17, 20
- statistical model, 20
- statistical shape variation, 8
- support vector machine, 7
- surface, 12
- taxonomy, 2
- texture, 20
- texture matching, 20
- texture model, 21
- texture prior, 6
- thin-plate-spline, 15, 16
- topology, 13
- training data, 20
- training sample, 2
- training set, 18
- transform
 - affine, 8
 - axis-symmetric, 8
 - B-spline, 8
 - rigid, 8
 - similarity, 8

triangulated mesh, 2
triangulation, 12
vertex connectivity, 12
visual appearance, 2

visualization, 12
volume-to-volume registration, 16
wavelet, 20

