

Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models

Gábor Székely*, András Kelemen, Christian Brechbühler and Guido Gerig

Communication Technology Laboratory, ETH-Zentrum, CH-8092 Zurich, Switzerland

Abstract

This paper describes a new model-based segmentation technique combining desirable properties of physical models (snakes), shape representation by Fourier parametrization, and modelling of natural shape variability. *Flexible parametric shape models* are represented by a parameter vector describing the mean contour and by a set of *eigenmodes of the parameters* characterizing the shape variation. Usually the segmentation process is divided into an initial placement of the mean model and an elastic deformation restricted to the model variability. This, however, leads to a separation of biological variation due to a global similarity transform from small-scale shape changes originating from elastic deformations of the *normalized model contours only*. The performance can be considerably improved by building shape models normalized with respect to a small set of stable landmarks (AC-PC in our application) and by explaining the remaining variability among a series of images with the model flexibility. This way the image interpretation is solved by a *new coarse-to-fine segmentation procedure* based on the set of deformation eigenmodes, making a separate initialization step unnecessary. Although straightforward, the extension to 3-D is severely impeded by difficulties arising during the generation of a proper surface parametrization for arbitrary objects with spherical topology. We apply a newly developed surface parametrization which achieves a uniform mapping between object surface and parameter space. The 3-D procedure is demonstrated by segmenting deep structures of the human brain from MR volume data.

Keywords: 3-D deformable models, 3-D shape analysis, segmentation of multidimensional images, statistical analysis of anatomical objects

Received October 3, 1995; revised January 12, 1996; accepted January 19, 1996

1. INTRODUCTION

Segmentation of anatomical objects from large 3-D medical data sets, which result for example from routine magnetic resonance imaging (MRI) examinations, represents one of the basic problems of medical image analysis. In some limited applications, segmentation could be achieved by minimal user interaction, providing procedures for the interpretation of medical scenes, which can be applied routinely (Gerig *et al.*, 1991). For general applications, however, adequate segmentation cannot be obtained without expert knowledge,

requiring tedious manual interaction by a human specialist.

Elastically deformable models (snakes) (Kass *et al.*, 1988) have been proposed as tools for supporting manual object delineation. While such procedures can be extended to 3-D (Terzopoulos *et al.*, 1988; Cohen *et al.*, 1992), their initialization becomes difficult. Most often, the initial guess must be very close to the sought contour to guarantee a successful result (Neuenschwander *et al.*, 1994). While a careful and time-consuming analysis is acceptable for outlining complex pathological objects, no real justification for such a procedure can be found for the delineation of normal, healthy organs, as needed in radiation treatment planning, for example.

*Corresponding author
(e-mail: szekely@vision.ee.ethz.ch)

The primary reason for the need of a precise snake initialization is the presence of disturbing attractors in the image, which do not belong to the actual object contour but force the snake into local energy minima. If the deformation of a snake could be limited to shapes within the normal anatomic variation of organs, such local minima could be avoided.

Elastically deformable parametric models offer a straightforward way for the inclusion of prior knowledge in the image interpretation process by incorporating prior distributions on the elastic model parameters to be estimated. Such a procedure has been implemented, for example, by Vemuri and Radisavljevic (1994) using a hybrid primitive called deformable superquadrics constructed in a multi-resolution wavelet base or by Staib and Duncan (1992a) for deformable Fourier models.

For complex shapes described by a large number of possibly highly correlated parameters, the usage of such priors may become tedious. The modal analysis as proposed by Pentland and Sclaroff (1991) offers a promising alternative by changing the basis from the original modelling functions to the eigenmodes of the deformation matrix. This way the dominant part of the deformations can be characterized by the few largest eigenmodes, reducing the dimensionality of the object descriptor space substantially. Such modal analyses have been successfully applied to medical image analysis, for example, by Sclaroff and Pentland (1994) or Nastar and Ayache (1994).

Cootes *et al.* (1993) combined the power of parametric deformable shape descriptors with statistical modal analysis. They use active shape models, which strictly restrict their possible deformations according to the statistics of training samples. Object shapes are described by the point distribution model (PDM) (Cootes and Taylor, 1992), which represents the object outline by a subset of boundary points. There must be a one-to-one correspondence between these points in the different outlines of the training set. After normalization to size, orientation and position they provide the basis for the statistical analysis of the object shape deformations. The mean point positions and their modes of variation (i.e. the eigenvectors corresponding to the largest eigenvalues of their covariance matrix) are used for delimiting the object deformations to a reasonable linear subspace of the complete parameter space. They propose to solve the 3-D shape analysis problem by a slice-by-slice approach (Hill *et al.*, 1992) or by extending the PDM concept for the description of 3-D shapes (Hill *et al.*, 1993). Similar parametrization based on point-by-point correspondence was proposed for the 3-D shape analysis of brain structures by Martin *et al.* (1994). By using the free vibration modes of the nodal displacement matrix of an elastic body, they even separate physical deformation modes which can be modelled by elastic deformation models from experimental

modes, evaluated from statistical analysis of a training sample.

While the idea of restricted elastic deformation of an average surface model is very promising, the parametrization of shapes by displacement of corresponding points on their surfaces is not a convenient technique. For a large training set containing many different anatomic structures, the generation of this parametrization seems to be very tedious and, because of the lack of a reasonable automatization, can be a source of errors. A similar modal analysis, however, can be performed for other contour parametrization techniques, as for example for the 2-D Fourier descriptors which were originally proposed by Persoon and Fu (1977) and Kuhl and Giardina (1982). Staib and Duncan have demonstrated segmentation by parametrically deformable elastic models for 2-D outlines (Staib and Duncan, 1992a) and 3-D object surfaces (Staib and Duncan, 1992b); their 3-D model was, however, limited to star-shaped surfaces. Here, we propose a novel technique based on modal analysis of the parameter vector of object contours, providing the desirable restriction of elastic deformations. The method uses automatic shape parametrization for any surfaces with spherical topology, thus avoiding the problem of finding corresponding points among different boundaries.

The paper is organized as follows. Section 2 discusses the new 2-D modal analysis procedure and the model-based segmentation by restricted deformation. The procedure is illustrated with the segmentation of the corpus callosum structure on mid-sagittal MR images of the human brain. Section 3 proposes several improvements of the procedure. Section 4 generalizes the method to 3-D, addressing the increased complexity requiring new mathematical and numerical solutions. Section 5 finally gives the conclusions and future directions for the development of automated, model-based segmentation techniques.

2. MODAL ANALYSIS OF 2-D FOURIER MODELS

In this section we first summarize the mathematics of Fourier descriptors for the parametrization of simple closed curves, and their use in snake-like elastic matching procedures. We summarize how to apply the modal analysis to 30 normalized outlines of the corpus callosum extracted from 2-D mid-sagittal MR images. The proposed procedure is a combination of an initialization using an average model based on template matching technique followed by an elastic deformation restricted to the major eigenmodes. We demonstrate that this separation of the similarity transform (translation, rotation and scaling) from the elastic deformation often leads to unsatisfactory segmentation results, calling for a modified segmentation process that provides a unified framework for initialization and elastic model matching.

2.1. Parametrization of 2-D contours with Fourier descriptors

The contour of a simply connected object (without holes) is represented by a closed curve with coordinates $(x(t), y(t))$, with t ranging from 0 to 2π . The coordinate functions can be expanded in a Fourier series. Restricting the series to degree K results in the parametric description

$$\begin{aligned} \mathbf{r}(t, \mathbf{p}) &= \begin{pmatrix} x(t, \mathbf{p}) \\ y(t, \mathbf{p}) \end{pmatrix} \\ &= \begin{pmatrix} a_0 \\ c_0 \end{pmatrix} + \sum_{k=1}^K \begin{pmatrix} a_k & b_k \\ c_k & d_k \end{pmatrix} \cdot \begin{pmatrix} \cos(kt) \\ \sin(kt) \end{pmatrix}. \end{aligned}$$

The outline is now completely described by the parameter-vector

$$\mathbf{p} = (a_0 \dots a_K, b_1 \dots b_K, c_0 \dots c_K, d_1 \dots d_K)^\top.$$

The parameters can be easily calculated from the sampling points of the outline $\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_P$ with $\mathbf{q}_0 = \mathbf{q}_P$ (we use maximally dense sampling of the boundary, as provided by the image raster). The resulting parametric shape description can be made invariant under similarity transformation by shifting, rotating and scaling according to the actual displacement, orientation and size of the ellipse determined by the first degree terms of the Fourier series. Similarly, the starting point is moved to a canonical position.

2.2. Fourier snakes

The snake technique as proposed by Kass *et al.* (1988) tries to find the position of a curve $\mathbf{r}(t, \mathbf{p})$, which minimizes the energy

$$E(\mathbf{p}) = E(\mathbf{r}(t, \mathbf{p})) = E_I(\mathbf{r}(t, \mathbf{p})) + E_D(\mathbf{r}(t, \mathbf{p})).$$

By varying \mathbf{p} , the curve deforms itself to minimize the image energy

$$E_I(\mathbf{r}) = \int_0^{2\pi} P(\mathbf{r}(t, \mathbf{p})) dt,$$

searching for an optimal position in the image, described by the potential $P(\mathbf{r}(t, \mathbf{p}))$. A typical choice takes $P(\mathbf{r}(t, \mathbf{p}))$ equal to the negative magnitude of the image gradient:

$$P(\mathbf{r}(t, \mathbf{p})) = -|\nabla I(\mathbf{r}(t, \mathbf{p}))|,$$

where I is the smoothed image. The deformation term $E_D(\mathbf{r})$ is called the internal energy of the snake and serves as a regularization force. It restricts the elongation and bending of the snakes, and normally depends on the first and second derivatives of the curve $\mathbf{r}(t, \mathbf{p})$.

Staib and Duncan (1992a) propose a different energy model, which makes use of the normal direction of the parametrized curve and of the image gradient to achieve a higher selectivity. Normalizing the image potential by the contour length allows contraction and dilation of curves without affecting the energy function.

$$E_I(\mathbf{r}(t, \mathbf{p})) = \pm \left(\int_0^{2\pi} \nabla I(\mathbf{r}(t, \mathbf{p})) \cdot \dot{\mathbf{r}}_\perp(t, \mathbf{p}) dt \right) \cdot \left(\int_0^{2\pi} \|\dot{\mathbf{r}}(t, \mathbf{p})\| dt \right)^{-1}.$$

The sign of the energy $E_I(\mathbf{r})$ will decide between segmenting bright objects on a dark background or vice versa. The polarity of the boundary can be neglected by using the absolute value of the dot product in the integration term.

One has to realize that cutting the Fourier expansion at a finite degree already serves as a regularization, leaving out high frequency variations of the coordinate functions. However, the internal energy cannot be neglected, basically for two reasons:

- even Fourier models of low degree can generate sharp cusps, and
- the parametrized contour can cross itself.

Whereas self-crossings of the outline are expensive to detect, discontinuities of the tangent can be evaluated from the curve parametrization. At such locations the curvature $\kappa(t_0, \mathbf{p})$ and its derivative $\dot{\kappa}(t_0, \mathbf{p})$ both become infinite. While high curvature of the boundary should not be excluded *a priori*, the curvature derivative is chosen to indicate discontinuities (see also Staib and Duncan, 1992), providing the following expression for the internal energy of the Fourier snake:

$$E_D(\mathbf{r}(t, \mathbf{p})) = \int_0^{2\pi} \dot{\kappa}^2(t, \mathbf{p}) \cdot \|\dot{\mathbf{r}}(t, \mathbf{p})\|^2 dt.$$

For the minimization of the total energy function we used the E04JBF routine of the NAGTM library (1988), using a quasi-Newton algorithm for finding an unconstrained minimum of a function of several variables.

2.3. Model-based segmentation of the corpus callosum

We tested the proposed procedure on a collection of 30 mid-sagittal slices of MRI brain images. A few images from the training set are shown in the upper row of Figure 1. Our goal was the automatic segmentation of the outline of the corpus callosum from grey scale images.

2.3.1. Model building

The corpus callosum outline had to be segmented and parametrized for each image of the training set. We used




Figure 1. Training set for the analysis of the corpus callosum on mid-sagittal slices of MRI brain data sets. The upper row shows a region of interest of the original slice, the lower row contains the segmented outline in a standardized invariant configuration (Fourier coefficients up to degree 100).




Figure 2. Template-matching example to find a similarity transform between the mean model curve and the gradient magnitude image. (a) Shows the original image, (b) its Gaussian smoothed version ($\sigma = 5$ pixels) and (c) the corresponding Canny gradient magnitude image, respectively. In (d) the optimal fit found by template matching is overlaid with the edge image.




Figure 3. Segmentation examples (a–c) illustrating segmentation failures. Images in the top row illustrate the initial placement of the model curves (using template matching), the bottom row shows the segmentation result. The *same* initial model curve is used for the different images.

the Fourier snake program as described above to perform this task, with a manual initialization of the optimization. The segmentation results can be seen in the lower row of Figure 1. Fourier coefficients up to degree 100 have been used for the description. The resulting contours of the training set have been normalized to be invariant to translation, rotation, scaling and the starting point of the parametrization.

The mean model has been determined from the normalized outlines by simply averaging their parameters. In order to determine the major deformation modes defined by the above training set, we performed a principal component analysis of the covariance matrix of the normalized Fourier coefficients. As after the first few eigenvectors the variance becomes very small, the few largest eigenmodes have been taken to build a flexible model that explains the biological variability of the shape of the corpus callosum outline.

2.3.2. Initial placement of model contour by template matching

A segmentation of the corpus callosum from grey-valued images based on the deformation of the mean model requires a suitable initialization. Due to the normalization of the Fourier coefficients, the average model only expresses shape deformations up to a similarity transformation. Therefore, the initial placement has to provide a sufficiently good match between the model and the edges in the grey-valued image. A standard template-matching procedure was chosen to solve this first optimization problem by equidistantly sampling the possible parameters of the similarity transformation within a reasonable range. The goodness of fit was calculated for a

Canny edge map on a relatively large scale, since the rigid transformation does not allow for elastic shape deformations. Figure 2 illustrates the result of the initialization by template matching. Figures 2a–c show the original image, the Gaussian smoothed image ($\sigma = 5$ pixels), and the Canny gradient magnitude, respectively. Figure 2d presents an overlay of the optimal fit between model curve and gradient image.

2.3.3. Segmentation by restricted elastic deformation

After initialization, we apply a modified version of the Fourier snakes program. Instead of optimizing in the complete space of the normalized Fourier coefficients, the optimization has been restricted in order to allow only deformations which are prominently represented in the training sample. In the first step the eigenvectors of the parameter covariance matrix are selected as a different set of basis functions in place of the harmonics. The restricted variation is achieved by choosing a subset of eigenmodes, usually the n largest ones, and calculating the optimization in this linear subspace. Starting with the result of the template matching, the model is elastically deformed until it fits (in a local optimum) the edges along the object contour.

The procedure described above provided satisfactory segmentation results in many cases. In some cases, however, we found it difficult to find the correct contour. Figure 3 shows a few examples of unsuccessful segmentations. Analysis of the results has shown that the separation of the similarity transform (in the initialization step) from the elastic deformation is mainly responsible for this failure. As the models have been normalized before the statistical analysis, translation,

rotation and scaling have actually been excluded from the elastic deformation step, forbidding minor corrections of the similarity transform determined by the template matching in the initialization step. Restricting the deformations to the subspace of the dominant eigenmodes made this situation even worse by not allowing corrections through free elastic deformations. This has led to the large segmentation errors demonstrated by the above examples.

The following section describes how these problems can be solved by incorporating the similarity transform into the analysis of biological variability. This way all covariation of the pose and elastic deformation parameters can be handled in a consistent way, too. The selected solution also avoids the somewhat artificial separation of the image analysis process into an initialization and elastic matching, leading to a two-step coarse-to-fine segmentation procedure.

3. IMPROVING FOURIER SNAKE SEGMENTATION

3.1. Alternative initialization techniques

One could try different ways to avoid the segmentation problems analyzed in the previous section.

- The desirable mixing of the effects of similarity transform and elastic deformation could be achieved by the incorporation of eigenvalues belonging to the the largest eigenvectors into the initialization (template matching) procedure. However, the dimensionality of the sampled search space would become far too high and would create, especially in view of a generalization to 3-D, an unacceptable computational burden.
- One could also try to extend the parameters of the elastic deformation with translation, rotation and scaling. This would lead to a higher dimensional local optimization problem without allowing the handling of the optimization of the parameters in one coherent framework.
- The selected optimization procedure finds the nearest local optimum on a high-dimensional strongly non-convex goal function. This makes the procedure highly sensitive to non-desired local optima. We are currently experimenting with alternative global optimization techniques as for example genetic algorithms as proposed by Hill and Taylor (1992) and the taboo search (Reeves, 1993) borrowed from discrete optimization, which may provide some improvement. However, one should realize that *the major problem does not lie in the initialization procedure*. The template matching usually provides a reasonable first approximation using the mean model, which still leads to unsatisfactory results due to the separation of the similarity transform and the elastic deformation described by the principal modes of the Fourier coefficients. This problem is addressed in the next paragraph.




Figure 4. Position of the anterior commissure (AC) and the posterior commissure (PC) on a mid-sagittal MRI image slice. The connecting line represents the unit vector \mathbf{x} of the external anatomical coordinate system.




Figure 5. The 50 largest eigenvalues of the covariance matrix calculated from the Fourier parameters of the 30 curves in the training set.

3.2. Model incorporating full biological variability

We expect for images representing anatomy that the *relative position, rotation and size* of healthy organs is restricted in a similar way as their *elastic deformation*. If we could define a coordinate system fixed to the anatomy, there would be no reason for an unrestricted similarity transform which precedes the elastic matching. The Fourier descriptors of the organ outlines *originally contain this information*, but we suppress it by normalizing the coefficients. In the case of the corpus callosum, the AC/PC line is a generally accepted, well-detectable geometric feature of the mid-sagittal images, which represents such a standard coordinate system. The line from the anterior to the posterior commissure (AC/PC line) as illustrated in Figure 4 has been manually extracted for each image of the training set. After determination of the Fourier




Figure 6. The upper row shows the segmented outlines of the training set with (left) and without (right) normalization of the pose parameters. The lower row shows the corresponding mean models.

coefficients, we apply normalization only to fix the starting point of the curve parametrization. The standardization of the images, necessary for the determination of the deformation modes, is based on a normalization of the AC/PC line to an e_x unit vector. After that, the same statistical analysis of the test set can be performed as previously explained, providing a mean model (now including its relative position and size to the AC/PC line), and the deformation modes which incorporate the parameters of the similarity transform, too. The resulting

eigenvalues are plotted in decreasing order in Figure 5. One can see that the remaining variation after the 12th eigenvalue becomes negligible. Accordingly, all deformations have been restricted to the 12 largest eigenvectors.

Figure 6 illustrates the determination of the mean model by comparing the results using the normalized outlines as described in the previous section (left side) and the average model resulting without normalization presented in the anatomical reference frame (right side). One can see the variation of the

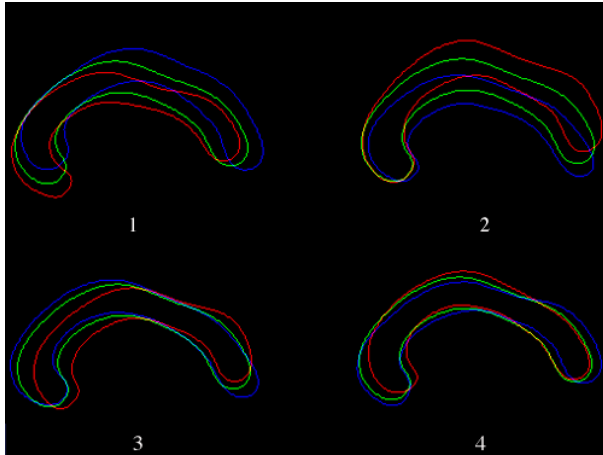


Figure 7. The first four eigenmodes of the deformations of the 30 objects in the training set. The calculations are based on contours represented by Fourier descriptors, which are normalized only with respect to the choice of the starting point. The deformation range amounts to $\pm\sqrt{2}$ eigenvalues.

pose of the training samples in the external coordinate system. At the same time, much more characteristic shape features are retained in the average outline if no normalization is performed before the training.

Figure 7 illustrates the deformations according to the first four eigenmodes, with deformations in the range of $\pm\sqrt{2}$ eigenvalues. It shows how the pose and elastic deformation parameters are mixed together in the dominant eigendeformations.

3.3. Segmentation

The determination of the AC/PC line now becomes part of the segmentation, since the model is built based on a normalization to these landmarks. The AC/PC line is determined manually for the image under analysis. The flexible model, characterized by the mean contour and the eigenmodes, now incorporates *changes of the position and local deformations* of the generic model, which makes the initialization step obsolete. In order to make the optimization robust against local extrema we applied a *two-step coarse to fine* strategy.

In the first step on the coarse level we use a procedure similar to template matching, as previously described. However, we now use a small set of a few dominant modes and calculate the best match in this linear subspace of major deformations by equidistantly sampling the eigenvalues of the selected deformation modes in a reasonable range. This way the complete parameter space is explored, and the result of this coarse segmentation can be used as a reliable initialization

for the following local optimization step. One should realize that the larger the number of eigenmodes participating in this coarse step, the better the initial fit will be in the following step. On the other hand, the computational burden grows exponentially with the dimensionality of the sampled parameter space. The number of eigenmodes considered in this step can be selected by trading the fit quality and the computational requirements.

On the fine level the segmentation is performed in the homogeneous parameter space of the eigenvectors corresponding to the largest eigenvalues, now choosing a larger number of eigenvectors than on the coarse level. Due to the full incorporation of the parameters of the similarity transform into the eigendeformations, the fine tuning of the model pose can be performed within the variability determined by the training set. Figure 8 shows the results of the modified segmentation procedure, which achieved perfect results in almost all cases.

One should mention that in some applications, the identification of the small individual variations of the contour outline, which is not represented even by the whole set of the eigenmodes, is desirable. In this case our strategy can be extended with an additional finer level by applying an unrestricted deformation of the parametrized snake using the result of the previous two levels as initialization.

4. 3-D FOURIER MODELS OF HUMAN BRAIN STRUCTURES

Our goal is the generalization of the improved 2-D procedure to 3-D. The new tool should allow a *well reproducible* segmentation of objects in volume data with *minimal human interaction*. This section describes the first steps in this direction, the generation of 3-D Fourier models from manually segmented training data, and the use of unrestricted 3-D Fourier snakes for elastic matching in grey-valued volume images. As in the previous section, we first summarize the basic mathematics of the 3-D Fourier snakes, and then show first results on some MRI brain data. The major problem, finding a *homogeneous parametrization* of surfaces of *arbitrarily shaped objects*, is solved by applying a very recently developed new parametrization technique. It overcomes limitations given by other surface parametrization schemes, e.g. the torus topology presented in Staib and Duncan (1992b). The following description is guided by the example of segmenting deep grey matter structures of the human brain from MR volume data.

4.1. Description of surfaces by spherical harmonic functions

The description of the surfaces of simply connected 3-D objects in an arbitrary basis can be performed similarly to the




Figure 8. Segmentation of the corpus callosum using the manually determined AC/PC line as a reference coordinate system for model building and segmentation. (a) Shows the initialization, resulting from a coarse match in the subspace of the largest four deformation modes. (b–d) Illustrate single steps of the optimization in the subspace of the largest 12 deformation modes (allowing fine adjustments) and the final result.

2-D case. The surface will be parametrized by two variables, the θ and ϕ polar parameters, and will be defined by three explicit functions

$$\mathbf{r}(\theta, \phi) = \begin{pmatrix} x(\theta, \phi) \\ y(\theta, \phi) \\ z(\theta, \phi) \end{pmatrix}.$$

We emphasize that this is not a radial function. If we select the spherical harmonic functions (Y_l^m denotes the function of

degree l and order m , see Greiner and Diehl (1964)) as a basis, the coordinate functions can be written as

$$\mathbf{r}(\theta, \phi, \mathbf{p}) = \sum_{k=0}^K \sum_{m=-k}^k \mathbf{c}_k^m Y_k^m(\theta, \phi),$$

where

$$\mathbf{c}_k^m = \begin{pmatrix} c_{xk}^m \\ c_{yk}^m \\ c_{zk}^m \end{pmatrix}.$$

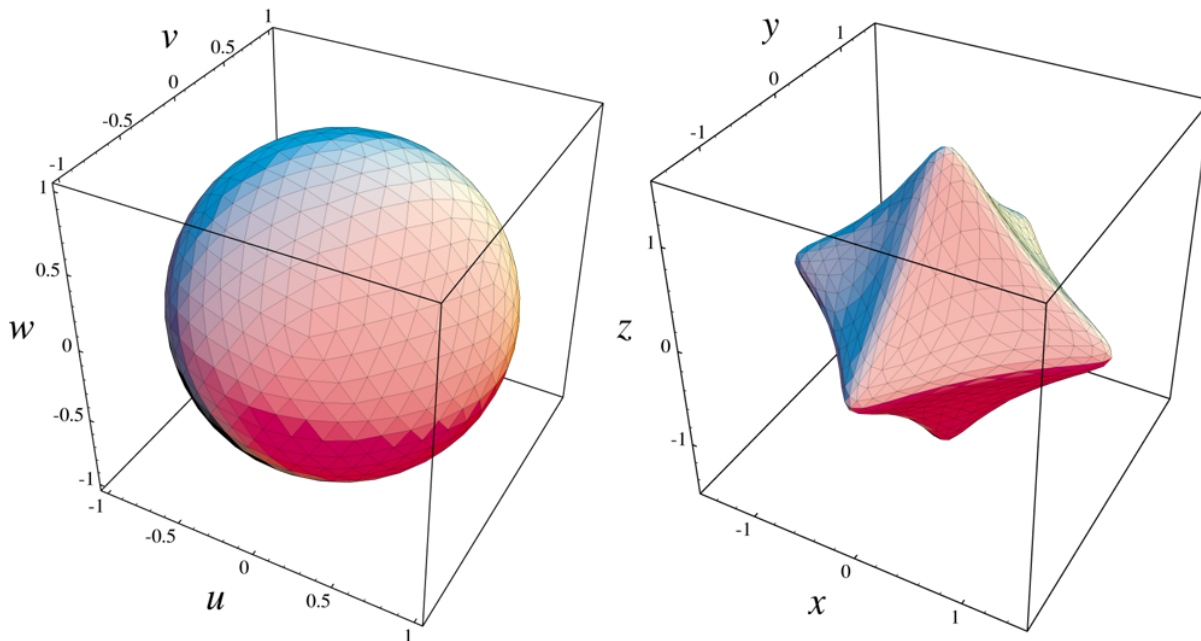


Figure 9. A simple function $U \mapsto \mathbb{R}^3$, visualized by faceting the surfaces in triangles.

We again restricted the expansion to the first $K + 1$ terms.

When the free variables θ and ϕ run over the whole sphere (e.g. $\theta = 0 \dots \pi$, $\phi = 0 \dots 2\pi$ —see Figure 9, left), the point $\mathbf{r}(\theta, \phi)$ runs over the whole surface of our object (Figure 9, right). The sphere Ω_3 is considered a perfectly symmetric surface without any singular points or preferred directions.

The surface is then described by the parameters

$$\mathbf{p} = (c_{x_0}^0, c_{y_0}^0, c_{z_0}^0, c_{x_1}^{-1}, c_{x_1}^0, c_{x_1}^1, c_{y_1}^{-1}, c_{y_1}^0, c_{y_1}^1, c_{z_1}^{-1}, c_{z_1}^0, c_{z_1}^1, \dots, c_{x_K}^{-K}, \dots, c_{z_K}^K)^\top.$$

One has to realize that the Fourier parametrization is just one possibility for the parametric description of contours. Alternative methods, for example deformable superquadrics, have also been proposed in the literature (Terzopoulos and Metaxas, 1991). We would like to emphasize that any reasonable parametric shape model can be used within the presented formalism. We preferred to use Fourier parametrization as it, in contrast to superquadrics, imposes no predetermined symmetries upon the object surface and has no preferred directions in space. It is complete; no additional formalisms are needed to rotate, translate, taper, bend, twist or locally deform the model, as the surface of any simply connected object can be represented to any degree of detail in the simple yet comprehensive formalism.

In the following we summarize the surface parametrization procedure which is described in detail in Brechbühler *et al.* (1992, 1995).

4.2. The surface data structure

Medical CT or MRI images are examples of volumetric data. For each cuboidal cell (volume element or voxel) in a certain volume we have one or more measurements. When segmentation succeeds, one anatomical unit can be characterized by a binary data volume, in which every voxel contains either 1, which means it belongs to the unit, or 0, meaning it is in the background. The object is then the set of '1' voxels and can be pictured as a collection of small cubes, adopting the *cube* notion (Hermann *et al.*, 1979). The surface of a voxel object is a set of unit squares, all parallel to one of the three coordinate planes yz , xz or xy . The edges and vertices that bound the faces are also parts of the surface, which is represented as a data structure that reflects geometry as well as neighbourhood relations.

4.3. Parametrization of closed surfaces

A key step in the shape description of a surface is its mapping to the parameter space, the sphere Ω_3 . Any point on the surface must map to exactly one point on the sphere, and vice versa. The location on the sphere corresponding to a surface point

defines the *parameters* of the point. It can be represented as two polar or three Cartesian coordinates, related through the bijection

$$\begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix}.$$

Mapping a surface to the sphere assigns parameters to every surface point; therefore we call it *parametrization*. The mapping must be continuous, i.e. neighbouring points in one space must map to neighbours in the other space. It is possible and desirable to construct a mapping that preserves areas. Narrowing to the cuberille notion, Figure 10 symbolically illustrates this mapping of a selected facet from the object surface to a portion of Ω_3 . It is not possible in general to map every surface facet to a spherical square: distortions cannot be avoided, but they should be minimal.

It emerges that the parametrization, i.e. the embedding of the object surface graph into the surface of the unit sphere, is a constrained optimization problem. The following paragraphs define the meaning of *variables*, *objective* (goal function), *constraints* and *starting values* in this context.

4.3.1. Variables

The coordinates of all vertices vary in the optimization. Using two (e.g. spherical) coordinates per vertex would be the most economic representation with respect to storage space, but it would make the equal treatment of all spatial directions difficult and pose the problem of discontinuity and singularities in the parameter space. We prefer Cartesian coordinates (u, v, w) for representing a location on the sphere, introducing one virtual degree of freedom per vertex. The number of variables is three times the number of vertices.

4.3.2. Constraints

Two kinds of equalities and one kind of inequality constrain the values that the variables can take.

- (i) Every vertex must lie on the unit sphere in parameter space, i.e. $u^2 + v^2 + w^2 = 1$. This constraint compensates for the virtual degree of freedom and forces.
- (ii) We ask for *area preservation*; any object surface region must map to a region of proportional area on the sphere. We include one constraint for each elementary facet: the area of the spherical quadrilateral must be 4π divided by the total number of faces.
- (iii) All quadrilaterals on the sphere must remain convex; no angle α_k may become negative or exceed π .

4.3.3. Objective function

The *goal* is to minimize the distortion of the surface net in the mapping. It must tend to make the shape of all the mapped faces as similar to their original square form as possible. To fulfil this goal perfectly, a facet should map to a ‘spherical square’ (see Figure 10). This can in general not be reached for all faces, and we need to trade off between the distortions made at different vertices. We observe that the ideal shape of any face, a spherical square, minimizes the circumference $\sum_{i=0}^3 s_i$ of any spherical quadrilateral with a given area. At the same time it maximizes $\sum_{i=0}^3 \cos s_i$. These two measures are similar, but not equivalent if summed over the whole net, as they trade off distortions differently. The second measure punishes too-long sides more and honours too-short sides less than the first one, which is a desirable effect. It is also simpler to calculate; the cosine of a side is the dot product of its endpoints.

4.3.4. Starting values

The variables in our optimization are the positions on the unit sphere to which the vertices are mapped, and *initial values* mean a first rough mapping of the object’s surface to the sphere. It is important for the optimization algorithm that the sphere be completely covered with faces and none of them overlap, even in the beginning.

Arbitrarily, two vertices are chosen as poles and a surface path connecting them as a date line. Discretizing and solving the Poisson equations $\Delta\theta = 0$ and $\Delta\phi = 0$ —with appropriate boundary conditions for the poles and the date line—yields values (θ, ϕ) for each vertex. Figure 11 shows the initial parametrization thus defined for the small object from Figure 10.

The optimization moves the vertices around on the sphere. The poles lose their special meanings, and all faces get equal area. The mapping converges to a state minimizing overall distortions. The same result is reached from all pole assignments (modulo rotations). Figure 12 presents the parametrization result for the same small object.

4.4. Elastically deformable Fourier surface models

Using the parametric surface description previously presented, the parametrized Fourier snake concept can be generalized from 2-D to 3-D. The concept is similar to the technique proposed by Staib and Duncan (1992b). However, a surface can be arbitrarily sampled based on the variables θ and ϕ (there is no distinguished orientation or position of poles), and the surface energy function can similarly be defined and evaluated as in the 2-D case. The image potential can still use the complete (vector-valued) gradient information

$$E_1(\mathbf{r}(\theta, \phi, \mathbf{p})) = \pm \int \int_A \nabla I(\mathbf{r}(\theta, \phi, \mathbf{p})) \cdot \mathbf{n}(\mathbf{r}(\theta, \phi, \mathbf{p})) \, dA$$




Figure 10. Every facet on the object surface maps to a spherical quadrilateral. Its sides are geodesic arcs on the sphere. On the unit sphere, a side s_i is equal to the corresponding centre angle. The quadrilateral in this illustration is special in that its sides $s_0 \dots s_3$ are equal and its angles $\alpha_0 \dots \alpha_3$ are equal: it is a spherical square.

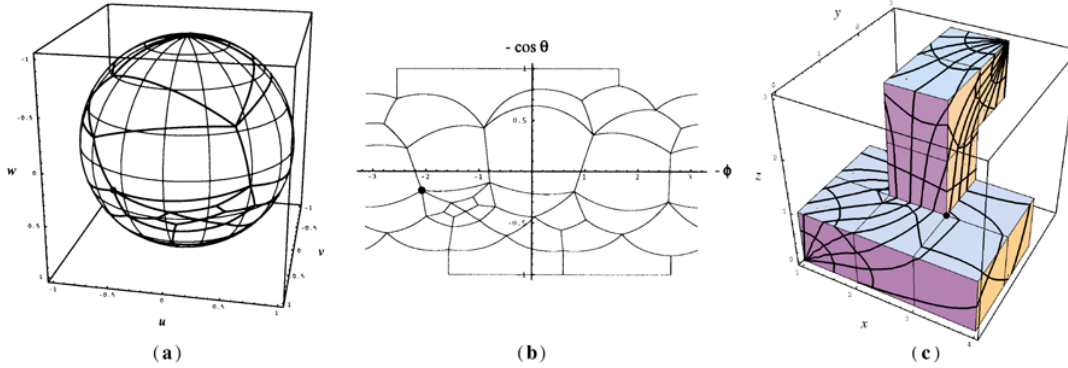


Figure 11. Optimization starts in this state, which is plotted in three ways. (a) The surface net is plotted in thick lines on the spherical parameter space. (b) ϕ and $\cos \theta$ are interpreted as Cartesian coordinates, giving a true-area cylinder projection. The horizontal lines at ± 1 are the poles. (c) Conversely, the globe coordinate grid is drawn over the object. For comparison, one vertex is marked with a black dot in all diagrams.

where dA is the surface element and \mathbf{n} is the surface normal vector. Using $\mathbf{a} = \partial \mathbf{r} / \partial \phi \times \partial \mathbf{r} / \partial \theta = \mathbf{n} dA / d\phi d\theta$, we derive

$$E_1(\mathbf{r}) = \pm \frac{\int_0^\pi \int_0^{2\pi} \nabla I(\mathbf{r}(\theta, \phi, \mathbf{p})) \cdot \mathbf{a}(\mathbf{r}(\theta, \phi, \mathbf{p})) d\phi d\theta}{\int_0^\pi \int_0^{2\pi} \|\mathbf{a}(\mathbf{r}(\theta, \phi, \mathbf{p}))\| d\phi d\theta}.$$

We use the internal energy term $E_D(\mathbf{r}(\theta, \phi, \mathbf{p}))$ to avoid sharp discontinuities in the surface normals. The curvature of the surface can be described by the principal curvatures κ_1 and κ_2 , which are combined to

$$\kappa = \sqrt{\kappa_1^2 + \kappa_2^2},$$

creating a measure for the curvature at every point on the surface. As in the 2-D case we want to limit its derivative,

resulting in the regularization term

$$E_D(\mathbf{r}) = \int_0^\pi \int_0^{2\pi} \|\nabla \kappa(\phi, \theta, \mathbf{p})\|^2 \cdot \|\mathbf{a}(\phi, \theta, \mathbf{p})\|^2 d\phi d\theta.$$

After the definition of the total energy, the problem is completely analogous to the 2-D case: we have to determine the parameters that generate a surface which minimizes this energy. We used exactly the same minimization procedures as in the 2-D case.

For the average surface before and after the deformation, the mapping will not meet the constraint of constant density exactly, nor will it exactly minimize distortions. However, the mapping has been individually adapted to each of the training objects, which are chosen to span the whole range of expected biological variability. A smooth interpolation between them can be expected to reproduce close to uniform and optimal

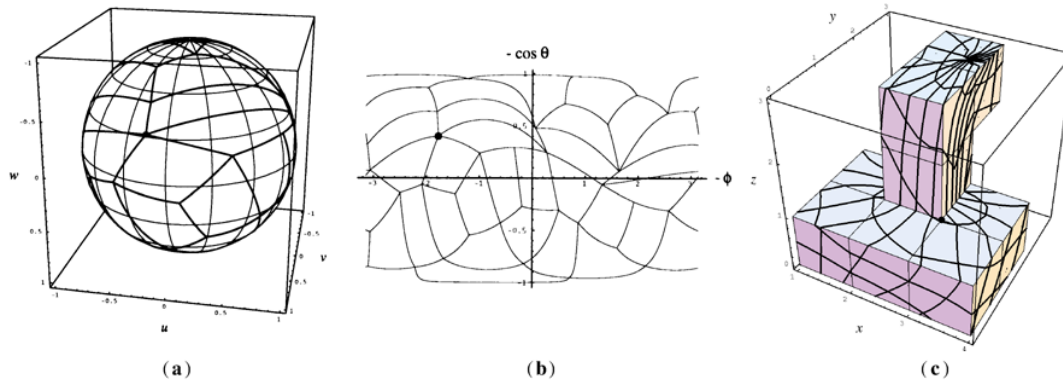


Figure 12. The parametrization achieved by the optimization, visualized as in Figure 11.




Figure 13. Manual segmentation of the putamen, the caudate nucleus and the globus pallidus. The images show a coronal slice from a 3-D data set (a) with overlay of the contours of segmented objects (b).

parametrization for the fitted surface. Exact uniformity and optimality is a means to establish correspondence among the objects of the training set, but it is not crucial for the success of the segmentation.

4.5. 3-D segmentation of deep grey matter structures

The training set consists of a collection of 30 3-D MRI data volumes of the human brain, where deep grey matter structures (putamen, caudate nucleus and globus pallidus) have been manually segmented. Figure 13 shows a coronal slice through one of the volumes and the outlines of the manually segmented objects. In the case of the putamen and globus pallidus one can see that there is practically no grey-level evidence to separate the two objects. Only *a priori* anatomical knowledge will

allow their segmentation, which clearly demonstrates the need for model-based 3-D segmentation procedures.

The modules for 3-D surface parametrization and Fourier description, for the calculation of eigenmodes and for 3-D segmentation by restricted elastic deformations, are implemented and ready for tests and validation. Preliminary results demonstrate two different procedures of the complete segmentation system:

- Figure 14 illustrates the generation of parametric surface descriptions from binary segmentations. The surfaces of the putamen (left) and the caudate nucleus (right) use a spherical harmonic approximation up to degree 6 and 8, respectively.




Figure 14. Parametrized description by spherical harmonics up to degree 8 of the surface of deep grey matter organs. The putamen is shown on the left, the caudate nucleus on the right. The original voxel object is overlaid as a wire-frame structure of the voxel edges.




Figure 15. Segmentation of the left caudate nucleus by 3-D Fourier snakes. Images in panel **a** represent the initial placement of the 3-D model as axial (top) and coronal (bottom) cuts (with spherical harmonics up to degree 3 with 48 parameters). Panel **b** shows the final segmentation result. A graphical display of the elastically deformed model representing the result of the 3-D segmentation is shown to the right (**c**). The final optimization was based on spherical harmonics up to degree 5 (108 parameters).

- Figure 15 demonstrates the segmentation of the caudate nucleus from original grey-valued volume data by the 3-D Fourier snake procedure. Here, the initial placement of the model surface was performed automatically by calculating an initial placement of a spherical harmonic surface of low degree (degrees 0–3). The images in panels **a** and **b** show the initialization and final result after elastic deformation in axial and coronal cuts. For the elastic fit we used spherical harmonics up to degree 5. Panel **c** illustrates a surface rendering of the segmentation result. 15–20 minutes were needed for the elastic fit of the models on a Sparc 10/41 processor.

We are currently calculating the surface parametrization and the description by spherical harmonics of all the 30 segmentations. Based on these results, the deformation eigenmodes for these objects can be determined. The segmentation itself, i.e. the initial placement and restricted elastic deformation, will be carried out with the Fourier snake procedure as shown in Figure 15. However, optimization is performed in the subspace of a small set of eigenmodes rather than the original space of the parameters, as discussed in section 3 for the 2-D case.

5. CONCLUSIONS

Automated, robust segmentation of medical images most often needs *a priori* anatomical knowledge. Typical cases are the segmentation of healthy organs, which present restricted anatomical variability, or the segmentation of organs if only incomplete evidence for boundaries is given by the grey-valued images, requiring an ‘intelligent guess’ about the position of the object boundary.

Geometric organ models and the statistics of their *normal (expected) variation* seem to offer a promising solution to this problem. We proposed the use of the Fourier parametrization of our models, followed by a statistical analysis of a training set, providing mean organ models and their eigendeformations. Elastic fit of the mean model in the subspace of eigenmodes restricts possible deformations and finds an optimal match between the model surface and boundary candidates.

We demonstrated the complete procedure with the determination of the outline of the corpus callosum on a 2-D MRI slice of the human brain. We have shown that the selection of an anatomically defined reference coordinate system allows the inclusion of the variation of the spatial position and orientation into the description of the model variation. This procedure provides a homogeneous framework for the complete analysis, while adding only minimal, well defined user interaction to the procedure. The initial placement of the average model is performed by a coarse elastic deformation defined by only the largest eigenmodes, helping to drive the following optimization into the ‘correct’ local optimum.

In this paper we presented preliminary results of a full 3-D generalization of the Fourier snake procedure. While the modules are implemented and ready for analysis, we still need to spend effort into the 3-D model building. We demonstrated that manually segmented deep grey matter organs of the human brain can be described by spherical harmonic descriptors. We further demonstrated the applicability of an unrestricted 3-D Fourier snake procedure for the segmentation of these organs. In contrast to similar research projects, we could overcome the bottleneck of surface parametrization by applying a new parametrization technique which provides a homogenous parametrization of an arbitrarily shaped simply connected object. Further, the novel method represents a true 3-D extension of the 2-D procedure, which significantly differs from combining results of a stack of 2-D slices.

Our ultimate goal is to provide an automated 3-D segmentation procedure that needs only minimal user interaction. This manual interaction would consist of the selection of a few clearly defined landmarks and could therefore be carried out also by non-experts. The segmentation itself would then run fully automatically. Such an elastic deformation procedure

restricted by prior knowledge about the deformation range would also find applications in *tracking* problems, where objects once defined with a relatively high expense could be automatically tracked in dynamic image sequences.

ACKNOWLEDGEMENTS

The authors thank Ron Kikinis (Brigham and Women’s Hospital, Harvard Medical School, Boston) for organizing the MR acquisitions and for discussions in regard to the clinical needs for robust 3-D image segmentation. Martha Shenton (Brain Imaging Laboratory, Dept. of Psychiatry, Harvard Medical School, Boston) is kindly acknowledged for providing the MR data and the manually segmented 3-D data.

REFERENCES

- Brechbühler, C., Gerig, G. and Kübler, O. (1992) Surface parametrization and shape description. In *Proc. VBC’92 Conf.*, pp. 80–89.
- Brechbühler, C., Gerig, G. and Kübler, O. (1995) Parametrization of closed surfaces for 3-D shape description. *CVGIP: Image Understanding*, 61, 154–170.
- Cohen, I., Cohen, L. D. and Ayache, N. (1992) Using deformable surfaces to segment 3-D images and infer differential structures. *CVGIP: Image Understanding*, 56, 242–263.
- Cootes, T. F. and Taylor, C. J. (1992) Active shape models – ‘smart snakes’. In *Proc. British Mach. Vision Conf. ’92*, pp. 266–275, Springer.
- Cootes, T. F., Hill, A., Taylor, C. J. and Haslam, J. (1993) The Use of Active Shape Models For Locating Structures in Medical Imaging. In *Proc. IPMI’93*, pp. 33–47, Flagstaff USA.
- Gerig, G., Martin, J., Kikinis, R., Kübler, O., Shenton, M. and Jolesz, F. (1991) Automatic Segmentation of Dual-Echo MR Head Data. In *Proc. IPMI’91 (Wye)*, pp. 175–187.
- Greiner, W. and Diehl, H. (1964) Theoretische Physik – Ein Lehr- und Übungsbuch für Anfangssemester, Band 3: Elektrodynamik, Verlag Harri Deutsch, Zürich und Frankfurt am Main, pp. 61–65.
- Hermann, G. T. and Liu, H. K. (1979) Three-dimensional display of human organs from computer tomograms. *CGIP*, 9, 1–21.
- Hill, A. and Taylor, C. J. (1992) Model-based image interpretation using genetic algorithms. *Image Vision Comp.*, 10, 295–300.
- Hill, A., Cootes, T. F. and Taylor, C. J. (1992) A Generic System for Image Interpretation Using Flexible Templates. In *Proc. British Mach. Vision Conf.’92*, pp. 276–285, Springer.
- Hill, A., Thornham, A. and Taylor, C. J. (1993) Model-Based Interpretation of 3-D Medical Images. In *Proc. British Mach. Vision Conf.’93*, pp. 239–348, BMVC Press.
- Kass, M., Witkin, A. and Terzopoulos, D. (1988) Snakes: active contour models. *Int. J. Comp. Vision*, 1, 321–331.
- Kuhl, F. P. and Giardina, C. R. (1982) Elliptic Fourier features of a closed contour. *CVGIP*, 18, 236–258.

- Martin, J., Pentland, A. and Kikinis, R. (1994) Shape Analysis of Brain Structures Using Physical and Experimental Modes. In *Proc. CVPR'94*, pp. 752–755.
- NAG Fortran Library Manual – Mark 13, (1988) 1st edn.
- Nastar, C. and Ayache, N. (1994) Classification of Nonrigid Motion in 3-D Images using Physics-Based Vibration Analysis. In *Proc. IEEE Workshop on Biomed. Image Anal.*, pp. 61–69, Seattle, WA.
- Neuenschwander, N., Fua, P., Székely, G. and Kübler, O. (1994) Initializing Snakes. In *Proc. CVPR'94, IEEE Computer Society Conf. on Computer Vision and Pattern Recognition*, June 21–23, Seattle, WA. pp. 658–663.
- Pentland, A. and Sclaroff, A. P. (1991) Closed-form solutions for physically based shape modelling and recognition. *IEEE PAMI*, 13, 715–729
- Persoon, E. and Fu, K. S. (1977) Shape discrimination using Fourier descriptors. *IEEE Trans. SMC*, 7, 170–179.
- Reeves, C. R. (ed.) (1993) *Modern Heuristic Techniques for Combinatorial Problems*. Blackwell, Oxford.
- Sclaroff, S. and Pentland, A. P. (1994) On Modal Modelling for Medical Images: Underconstrained Shape Description and Data Compression. In *Proc. IEEE Workshop on Biomed. Image Anal.*, pp. 70–79, Seattle, WA.
- Staib, L. H. and Duncan, J. S. (1992a) Boundary finding with parametrically deformable models. *IEEE PAMI* 14, 1061–1075.
- Staib, L. H. and Duncan, J. S. (1992b) Deformable Fourier models for surface finding in 3-D images. In *Proc. VBC'92 Conf.*, pp. 90–194.
- Terzopoulos, D. and Metaxas, D. (1991) Dynamic 3-D models with local and global deformations: deformable superquadrics. *IEEE PAMI*, 13, 703–714.
- Terzopoulos, D., Witkin, A. and Kass, M. (1988) Symmetry-seeking models and 3-D object reconstruction. *Int. J. Comp. Vision*, 1, 211–221.
- Vemuri, B. C. and Radisavljevic, A. (1994) Multiresolution stochastic hybrid shape models with fractal priors. *ACM Trans. Graphics*, 13, 177–20.