Available online at www.sciencedirect.com
ScienceDirect
Medical Image Analysis 11 (2007) 197–206
MEDICAL IMAGE ANALYSIS
www.elsevier.com/locate/media

Reslicing axially sampled 3D shapes using elliptic Fourier descriptors

Yongwon Jeong, Richard J. Radke *

Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, USA

Received 12 September 2005; received in revised form 9 October 2006; accepted 22 December 2006
Available online 9 January 2007

Abstract

We propose a new method that interpolates between parallel slices from a 3D shape for the purposes of reslicing and putting into correspondence organ shapes acquired from volumetric medical imagery. By interpolating the coefficients of elliptic Fourier descriptors for a set of parallel contours, a new set of slices can be directly generated at desired axial locations. Neither an explicit correspondence between points on adjacent contours nor a 3D interpolating surface needs to be obtained. We apply the proposed reslicing method to experimental datasets of both synthetic 3D shapes and real prostate contours, and demonstrate that it performs as well as a common method based on variational implicit surfaces, for a much lower computational cost. We also show that reslicing and putting into correspondence an ensemble of axially sampled 3D organs enables the construction of shape models for accurate 3D segmentation.

Keywords: Resampling; Surface reconstruction; Shape modeling; Elliptical Fourier descriptors; Variational implicit surfaces

1. Introduction

The statistical analysis of ensembles of organ shapes is an important prerequisite for developing deformable shape models that can be used to anticipate inter- or intra-patient variability or to accurately segment organs from volumetric imagery (McInerney and Terzopoulos, 1996). For example, Fig. 1 shows several examples of manually labeled points on the surface of a patient's prostate obtained during radiotherapy treatment planning, represented as sets of points arranged in parallel planes (or "contours"). However, obtaining a large number of accurate correspondences between multiple hand-labeled organs contoured in parallel axial slices is difficult, since the number of slices in each dataset and the number of sample points around each contour generally vary between datasets, and the spacing of the sample points around each contour is usually nonuniform. A reasonable approach is to resample the ensemble of shapes so that each dataset has the same number of slices with the same number of samples at roughly homologous positions in each slice. This can naturally be posed as an interpolation problem, and is frequently addressed using a variational implicit function approach that (1) estimates a function whose zero-level-set is a 3D surface that interpolates the original points, and (2) evaluates the interpolating function to obtain surface points at each desired axial location. This method of interpolation was recently reintroduced into the computer graphics literature by Turk and O'Brien (1999). While the process produces smooth, continuous interpolating surfaces, it can be quite time-consuming, since a large linea\r system must be solved and numerical evaluation of the interpolating function is required to extract each resampled slice.

In this paper, we present a technique for the fast and accurate *reslicing* of 3D organ shapes specified as points on parallel axial slices, using a technique based on elliptical Fourier descriptors. The technique is advantageous in that neither an explicit correspondence between adjacent contours nor a 3D interpolating surface needs to be obtained. We demonstrate that computationally intensive computer graphics algorithms are not required to reslice organ

* Corresponding author. Tel.: +1 518 276 6483; fax: +1 518 276 8715.
E-mail addresses: jeongy@rpi.edu (Y. Jeong), rjradke@ecse.rpi.edu (R.J. Radke).

1361-8415/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2006.12.003

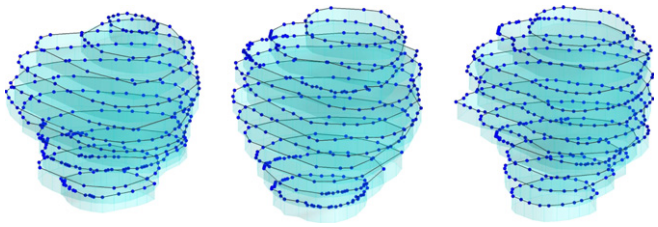

Fig. 1. Axial slices of one patient's prostate acquired on three different days of radiation treatment, contoured from CT imagery. The number of axial slices for each dataset and the number of sample points around each contour generally vary between datasets, and the spacing of the sample points around each contour is usually nonuniform.

shapes typical of prostate radiotherapy applications. We compare the elliptic Fourier descriptor method for reslicing against the implicit function method, and show that the results have the same accuracy obtained at less than 15% of the computational cost. Since time is valuable in a busy radiotherapy clinic, this speed advantage is important.

The paper is organized as follows. In the next section, we discuss related work on 3D surface reconstruction and correspondence. In Section 3, we give a detailed description of the variational implicit function method for interpolation, and in Section 4, we present a new interpolation method based on elliptic Fourier descriptors. In Section 5, we apply both reslicing methods to experimental datasets containing both synthetic shapes and hand-labeled prostate contours, and discuss the results. We also demonstrate how a shape model built from the resliced prostate contours can be used to accurately segment the prostate from 3D computed tomography (CT) images. We conclude the paper in Section 6.

2. Related work

Point distribution models (PDMs) (Cootes et al., 1992), also known as active shape models (Cootes et al., 1995), have been frequently and successfully used for shape modeling and segmentation since their introduction, since they can capture shape variability with only a few parameters, while generating a wide range of shapes consistent with an ensemble of training data. The key requirement for building a PDM is obtaining a fairly dense set of corresponding points (landmarks) between all shapes in the training dataset (e.g. images of the same organ from the same patient on different days of treatment, or images of the same organ from different patients). The model cannot express shape variability accurately if this correspondence is not well established. However, in medical image analysis, manually determining correspondence between organ shapes is time-consuming, subjective, and very challenging in 3D cases. Thus, it is important to automatically obtain unbiased and consistent correspondences, so that a model built using the landmarks can capture the desired statistical characteristics of shape variability. As mentioned above, the ensemble of training data that is available for shape analysis usually comes in the form of points on an organ arranged around parallel contours, and our goal is to consistently reslice the data in the axial direction to obtain corresponding slices across the ensemble of training data. A common approach is to first construct a 3D surface interpolating the contours, and then resample it into the desired set of slices, so most related work is from the surface reconstruction literature. Bolle and Vemuri (1991) gave a good general survey of 3D surface reconstruction methods.

In the context of medical imagery, early work focused on building a 3D triangulated model from cross-sectional polygonal contours. The key sub-problem is "tiling": that is, constructing an optimal (e.g. surface-area-minimizing) triangular mesh between two contours at different levels by constructing "spans", or edges between vertices at adjacent levels. Good surveys of tiling techniques were given by Meyers et al. (1992) and Barequet and Sharir (1996). Such techniques are not particularly well suited to the contour interpolation problem, since the triangular meshes produced by connecting original vertices are coarse instead of smooth.

Higher quality 3D models based on surface reconstruction from unorganized points provide a more natural framework for the contour resampling problem, since the resulting surfaces are smooth, no explicit connections between original points are required to arrive at a solution, and branching structures can be handled gracefully if they are present. Early work by Hoppe et al. (1992) used a surface interpolation method for unorganized points based on estimating samples of the signed distance function to the fitting surface. However, the algorithm results in no convenient functional form for the distance field. Turk and O'Brien (1999) recently reintroduced the variational approach of implicit surface fitting to the computer graphics community, though it had been known for some time previously (e.g. Bookstein, 1989; Grimson, 1982). This method, described in detail in the next section, is commonly applied for smooth interpolation of very large datasets, and is appealing since the implicit surface is parameterized by coefficients on basis functions centered at the original points.

An alternate approach to using an implicit surface is to fit a parametric model that represents the surface as an embedding of a two-dimensional parameter domain (e.g. the plane or the sphere) in $\mathbb{R}^3$. Finely tessellating the 2D domain into triangles, quads, or other polygons (e.g. simplex meshes (Delingette, 1994)) induces a corresponding tessellation of the 3D surface, resulting in a finite-element model. The tessellation may be desired to be as uniform as possible, or to concentrate more vertices in regions of high curvature (i.e. an adaptive tesselation). There are many such techniques; see Bolle and Vemuri (1991) for a review. One extension is the deformable superquadric model proposed by Terzopoulos and Metaxas (1991). The idea is to augment a parametric superquadric mesh with a local displacement field to better approximate surfaces with fine details. Vemuri and Radisavljevic (1993)

applied a wavelet transform to the displacement field to give more control over deformations at different resolutions. In medical image analysis, fitting such models to data usually uses imagery as input as opposed to raw 3D contour points. These techniques are generally more related to segmentation of a single image or interactive simulation (e.g. Delingette and Ayache, 2004) than ensemble shape analysis, since the goal is usually not to establish good correspondences across a large number of different shapes.

Designing methods for determining good correspondences between 3D shapes is in itself a large area of active research interest. Many techniques have been developed for estimating correspondences in specific application domains, such as 3D images of faces (Blanz and Vetter, 1999) or brains (Wang et al., 2000). Davies et al. (2002) approached the correspondence problem based on compact parameterization of a shape ensemble, using the Minimum Description Length (MDL) principle. Since correct correspondences result in a highly compressible model, they hypothesized that minimizing the ensemble description length should produce good correspondences. However, the technique is difficult and time-consuming to implement in 3D. Styner et al. compared this MDL approach with several other correspondence methods in Styner et al. (2003).

A newer family of point-set alignment techniques matches ensembles of shapes to each other without requiring any initial explicit correspondence between them. For example, Chui et al. (2004) described a joint clustering and matching technique that used a deterministic annealing algorithm to compute a mean shape from multiple sample point sets. During the process, correspondences between regions of the original shapes are formed. They qualitatively demonstrated their algorithm on nine 2D corpus callosa and twenty 3D hippocampus shapes. This approach was extended by Wang et al. (2006), who used the Jensen–Shannon divergence between cumulative distribution functions to non-rigidly register multiple unlabeled point sets. They qualitatively demonstrated their algorithm on seven 2D corpus callosa and four 3D hippocampus shapes. These techniques may be useful for the reslicing problem, although they would need to be quantitatively validated on larger datasets and compared against other algorithms.

Finally, we note that the reslicing and correspondence problems are related to the remeshing of 3D models, used in computer graphics applications to improve mesh quality (e.g. vertex regularity, unsharp triangles) for the purposes of mesh editing, animation, simplification, denoising, rendering, and compression. Alliez et al. (2006) gave a good, recent review of remeshing techniques. Of particular relevance to the problem addressed here is remeshing for the purpose of "morphing" between two meshes; see Alexa (2002) for a good review. In particular, "compatible" or "consistent" mesh parameterizations, in which two meshes are desired to have semantically corresponding vertices and identical connectivity, are especially useful for the morphing problem (Kraevoy and Sheffer, 2004; Praun, 2001). However, manually labeled correspondence between well-chosen mesh vertices is still required for such techniques. Also, remeshing techniques generally require fine initial meshes. The input data in our application is not connected in a mesh, and even applying a tiling technique as mentioned above would result in a coarser mesh than remeshing algorithms generally expect.

3. Interpolation using implicit functions

We now describe how the reslicing problem can be solved based on surface interpolation using an implicit function formulation. This implicit function on $\mathbb{R}^3$ is defined by

$$s(p) = q(p) + \sum_{i=1}^{N} \lambda_i \phi(p - p_i), \tag{1}$$

where $\{p_i = (x_i, y_i, z_i),\ i = 1, \ldots, N\}$ are the sample points to be interpolated, $p \in \mathbb{R}^3$ is a test point, $q(p)$ is a polynomial of degree $K-1$, $\phi(p) : \mathbb{R}^3 \to \mathbb{R}$ is a basis function, and the λ_i are unknown weights to be determined such that $s(p_i) = 0,\ i = 1, \ldots, N$ for the given sample points on the surface. The implicit function must also be set to a non-zero value at additional sample points to prevent $s(p) = 0$ from being a viable solution. For example, Turk and O'Brien (1999) set the implicit function to a positive value (e.g. 1) at points located slightly inward from and normal to the surface at each surface sample point.

Smoothing splines (also known as thin-plate splines) and their variations are popularly used in surface reconstruction (Bookstein, 1989). One advantage of using smoothing splines is that they produce a smooth surface interpolating a given set of points, which is generally desirable for organs in medical imagery. Smoothing splines produce fitting surfaces by minimizing a combination of bending energy (related to the smoothness of the fitting surface) and potential energy (related to the closeness of the fitting surface to the set of original points). That is, a smoothing surface is the function f that minimizes

$$E = \alpha \int_{\Omega} f_{uu}^2(u,v) + 2f_{uv}^2(u,v) + f_{vv}^2(u,v)\,\mathrm{d}S + \sum_{i=1}^{N} (p_i - f(u_i, v_i))^2, \tag{2}$$

where Ω is the domain of the fitting surface in $\mathbb{R}^2$, the p_i's are data points in $\mathbb{R}^3$, and α is a weighting factor. In our case, we want the reconstructed surface to exactly interpolate the original data points, so we only minimize the bending energy, namely:

$$E = \int_{\Omega} f_{uu}^2(u,v) + 2f_{uv}^2(u,v) + f_{vv}^2(u,v)\,\mathrm{d}S, \tag{3}$$

which is an integral over the 2D surface embedded in 3D. Duchon (1977) showed that thin-plate spline interpolation in 3D was equivalent to using the basis function $\phi(p) = \|p\|^2 \log\|p\|$ in (1).

There are many unknown parameters in (1) to be determined: the coefficients $\{\lambda_i,\ i=1,\ldots,N\}$ and K coefficients of the polynomial $q(p)$, denoted as a vector a. Since there are more equations than unknowns, additional constraints are required to uniquely define a solution; a usual choice is to force

$$\sum_{i=1}^{N}\lambda_i q_j(p_i)=0,\quad j=1,\ldots,M, \tag{4}$$

where the q_j form a basis for the space of polynomials of degree K in two dimensions (so $M=\binom{K+2}{2}$). In this case, the unknown parameters can be found as the solution to a linear system

$$\begin{bmatrix} A & Q \\ Q^T & \mathbf{0} \end{bmatrix}\begin{bmatrix}\lambda \\ a\end{bmatrix}=\begin{bmatrix}h\\0\end{bmatrix}. \tag{5}$$

Here, $A\in\mathbb{R}^{N\times N}$ is a matrix with coefficients $A_{ij}=\phi(p_i-p_j)$, $Q\in\mathbb{R}^{N\times M}$ is a matrix with coefficients $P_{ij}=q_j(p_i)$, and $\mathbf{0}$ is a $M\times M$ matrix of zeros. On the right-hand side, h is the vector of values of the implicit function at the sample points (in this case, the values are always either 0 or 1).

Once the implicit functional form $s(p)$ has been obtained, the level set of points $s(p)=0$ must be determined to extract the interpolating surface. Typically, the implicit function is evaluated on a grid and the surface is obtained using isosurfacing techniques. Isosurfacing methods fit successively finer surface primitives (e.g. polygons) to find constant-value contours in volumetric data (Elvins, 1992). Since our objective is resampling on axial planes, the isosurface really only needs to be extracted on each of the specified planes.

The overall procedure is outlined in Fig. 2. Fig. 2a shows original sample points arranged in parallel slices. An interpolating surface constructed using thin-plate splines is shown in Fig. 2b. Fig. 2c shows new points obtained by resampling the interpolated surface at new axial planes.

The variational implicit function method has two computationally expensive steps: the solution of the linear system (5) and the numerical evaluation of the interpolating function to find the new contours. The first issue has been addressed using fast multipole methods and far-field approximations (Billings et al., 2002; Cherrie et al., 2002), and there have been efforts to speed up the isosurface extraction using efficient surface-following algorithms (Carr et al., 2001). However, both steps are still time-consuming, and it seems clear that an entire implicit function need not be calculated if the main goal is simply resampling new axial slices. In the next section, we present a simple method for the reslicing problem that produces almost exactly the same result as isosurfacing an interpolated implicit function, for a fraction of the computational cost.

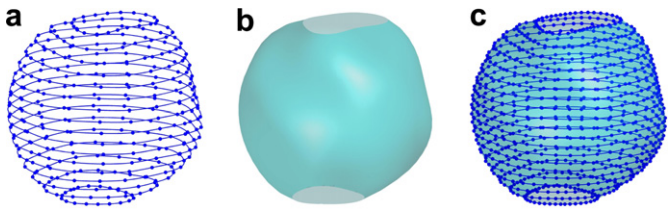

Fig. 2. Reslicing using a variational implicit surface. (a) A dataset of points arranged in 17 parallel slices. (b) A reconstructed surface is estimated that interpolates all the points, by extracting the isosurface of a 3D implicit function. (c) The reconstructed surface is sampled into the target number of parallel slices (e.g. 20 slices).

4. Interpolation using elliptic Fourier descriptors

A closed planar curve can be expressed in parametric form using elliptic Fourier descriptors (EFDs) (Kuhl and Giardina, 1982; Tello, 1996). Specifically, a closed curve in 2D parameterized by $0\leqslant t<2\pi$ can be expressed as a weighted sum of the Fourier basis functions:

$$\begin{bmatrix}x(t)\\y(t)\end{bmatrix}=\begin{bmatrix}a_0\\c_0\end{bmatrix}+\sum_{k=1}^{\infty}\begin{bmatrix}a_k & b_k\\ c_k & d_k\end{bmatrix}\begin{bmatrix}\cos kt\\ \sin kt\end{bmatrix}. \tag{6}$$

The coefficients are given in closed form by

$$\begin{aligned} a_0&=\frac{1}{2\pi}\int_0^{2\pi}x(t)\,\mathrm{d}t, & c_0&=\frac{1}{2\pi}\int_0^{2\pi}y(t)\,\mathrm{d}t,\\ a_k&=\frac{1}{\pi}\int_0^{2\pi}x(t)\cos kt\,\mathrm{d}t, & b_k&=\frac{1}{\pi}\int_0^{2\pi}x(t)\sin kt\,\mathrm{d}t,\\ c_k&=\frac{1}{\pi}\int_0^{2\pi}y(t)\cos kt\,\mathrm{d}t, & d_k&=\frac{1}{\pi}\int_0^{2\pi}y(t)\sin kt\,\mathrm{d}t. \end{aligned} \tag{7}$$

Thus, the closed curve $\{x(t),y(t)\}$ is equivalently represented by $\{a_0,c_0,a_1,b_1,c_1,d_1,\ldots\}$. EFDs were used to model 2D biological shapes in Staib and Duncan (1992) and 3D surfaces in Wu and Sheu (1998). EFDs are particularly suitable for representing biological objects, since such objects usually do not have sharp edges (and hence, few coefficients are required to accurately represent them).

The main idea of this paper is to resample a set of parallel slices by interpolating corresponding EFD coefficients. First, we obtain the EFDs for each original slice. We must obtain the same number of coefficients for each slice, as well as align the phase of all slices. We do this by interpolating the points in each slice with a closed cubic spline (Farin, 2001), and extending R rays from the set's center of gravity at uniform angles (starting at $\theta=0°$) to the slice boundary, as illustrated in Fig. 3. This set is used to compute the EFD coefficients, using the discretized version of (7). That is, for a closed curve $\{x(t_j),y(t_j)\}$ where $t_j=\frac{2\pi j}{R}$, $j=0,1,\ldots,R-1$:

$$\begin{aligned} a_0&=\frac{1}{2\pi}\sum_{j=0}^{R-1}x(t_j), & c_0&=\frac{1}{2\pi}\sum_{j=0}^{R-1}y(t_j),\\ a_k&=\frac{1}{\pi}\sum_{j=0}^{R-1}x(t_j)\cos kt_j, & b_k&=\frac{1}{\pi}\sum_{j=0}^{R-1}x(t_j)\sin kt_j,\\ c_k&=\frac{1}{\pi}\sum_{j=0}^{R-1}y(t_j)\cos kt_j, & d_k&=\frac{1}{\pi}\sum_{j=0}^{R-1}y(t_j)\sin kt_j. \end{aligned} \tag{8}$$

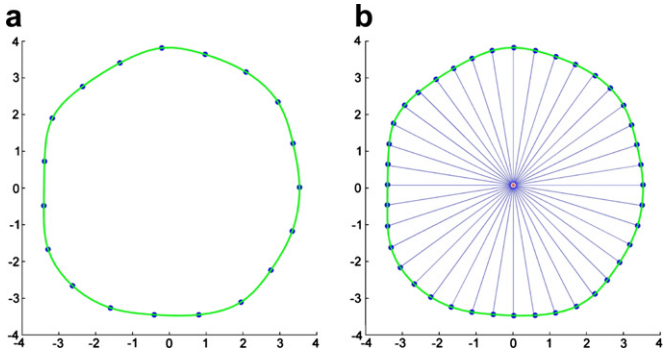

Fig. 3. (a) The original set of points on a slice. (b) The set of points after spline interpolation and resampling.

At this point, each slice at height z_i is represented by a vector of EFD coefficients $[a_0(z_i),c_0(z_i),a_1(z_i),b_1(z_i),c_1(z_i), d_1(z_i),\ldots]$. These vectors are interpolated element-by-element to obtain a_0, c_0, etc. as continuous functions of z. For example, if there are D slices in the dataset, we compute

$$a_0(z) = g(\{a_0(z_i), \quad i = 1,\ldots,D\}), \tag{9}$$

where g is, in our experiment, a piecewise cubic spline interpolant. Then we can evaluate the interpolated functions and plug them into (6) to obtain the surface contour at any value of z; the interpolant coincides with and smoothly interpolates all of the original sample points. The ability of EFDs to accurately represent shape contours depends on the number of coefficients used for approximation. The more coefficients used, the higher spatial frequency (i.e. higher curvature) that can be captured. In the next section, we analyze reslicing accuracy as a function of the number of harmonic coefficients.

In terms of computational complexity, the EFD method has clear advantages over reslicing using implicit functions. While there are several 1D spline interpolation problems to solve, these are inexpensive compared to solving a large linear system or isosurfacing on a fine grid. Furthermore, we show in the next section that for shapes corresponding to organs in medical imagery, the resampling accuracy does not suffer.

5. Experimental results

To validate the proposed method, we conducted two experiments. First, in order to determine the ground truth accuracy of both the implicit function and EFD methods, we generated synthetic 3D shapes with continuous surfaces described in closed form. Each shape was sampled into two sets of slices with different (uniform) spacing in the z-dimension. One set of slices was designated as the reference set, and the other set was used as input to the two resampling methods. The resampled slices from the two methods were compared with the reference slices to determine each method's accuracy. In the second experiment, we acquired sets of slices from hand-contoured CT images of the prostate. Each set of slices was resampled using the two methods and the results were compared with each other. The goal is to show that the EFD method produces results that are virtually indistinguishable from the implicit function method at a considerable computational savings. In the experiments, we followed the implementation of the variational implicit surface method as proposed by Turk and O'Brien (1999) and find the intersection of the fitting surface with each parallel slice by isosurfacing. All experiments were conducted using MATLAB on a Pentium 4-2.6 GHz PC with 1 GB RAM.

To determine the appropriate number of EFD harmonic coefficients for the experiments, we measured reslicing accuracy for the synthetic and clinical datasets as a function of the number of harmonics, as illustrated in Fig. 4. Reslicing accuracy is measured for each dataset as described in the following subsections. As expected, the approximation using EFD becomes more accurate as more harmonics are used. However, there is an obvious "elbow" at $k = 8$ harmonics, after which the reslicing accuracy obtained by using more harmonics is negligible. Therefore, in our experiments, we used 8 harmonics ($\{a_0, c_0, a_1, b_1, c_1, d_1, \ldots, a_8, b_8, c_8, d_8\}$) which produced good surface approximations for both the synthetic and clinical organ shapes of interest. However, we note that there was virtually no additional time consumed by increasing the number of harmonics in our experiments.

5.1. Synthetic shapes

To generate random synthetic 3D shapes whose surfaces could be described in closed form, we used a shape parameterization based on spherical harmonics (Brechbühler et al., 1995; Arfken and Weber, 2005). First, we generate a random radius for each of many altitude and azimuth angle pairs (θ,φ) uniformly spaced in spherical coordinates. Each radius is modeled as an i.i.d. random variable with $r(\theta,\varphi) \sim 3 + \sigma U$, where U is the uniform distribution with support $[0,1]$. Connecting these vertices into faces produces a noisy, jagged object as illustrated in Fig. 5a. The bumpiness of the shape can be controlled by the width of the uniform distribution. Next, the shape is smoothed using a 2D moving average filter over (θ,φ), as illustrated in Fig. 5b. Finally, the smoothed shape is approximated using spherical harmonics up to degree L as in Fig. 5c. That is, we parameterize a surface point $p \in \mathbb{R}^3$ in terms of (θ,φ) as:

$$p(\theta,\varphi) = \sum_{l=0}^{L} \sum_{m=-l}^{l} \mathbf{c}_l^m Y_l^m(\theta,\varphi), \tag{10}$$

where the spherical harmonic $Y_l^m(\theta,\varphi)$ of degree l and order m is given by

$$Y_l^m(\theta,\varphi) = \begin{cases} \sqrt{\frac{2l+1}{4}\frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta)\, e^{im\varphi}, & m \geqslant 0, \\ (-1)^m (Y_l^m)^*(\theta,\varphi), & m < 0, \end{cases} \tag{11}$$

and $P_l^m(x)$ is the associated Legendre polynomial. The coefficients in (10) can be computed as

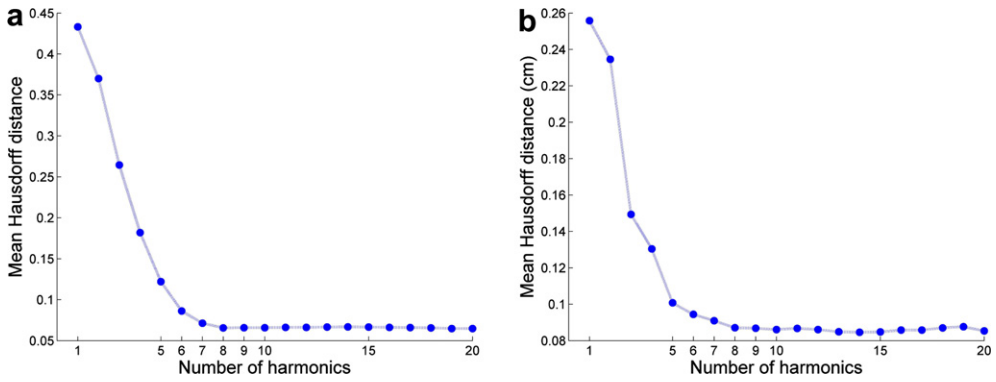

Fig. 4. Approximation error as a function of the number of EFD harmonic coefficients. (a) The mean Hausdorff distance between each set of 17 synthetic shape contours with $\sigma = 2.5$ interpolated using EFDs with the given number of harmonics and the true shape. (b) The mean Hausdorff distance between each set of prostate contours from Patient 4 interpolated using EFDs with the given number of harmonics and the shape estimated using implicit functions.

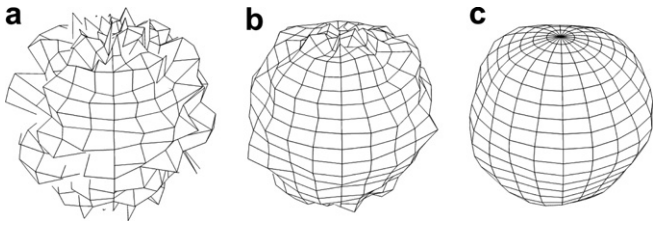

Fig. 5. Generating random shapes using spherical harmonics. (a) A random radius is added to samples of a sphere at uniform azimuth and altitude angles. (b) The shape is smoothed using a moving average filter. (c) The smoothed shape is approximated using spherical harmonics up to degree 7.

$$\mathbf{c}_l^m = \int_0^{\pi} \int_0^{2\pi} p(\theta,\varphi) Y_l^m(\theta,\varphi) \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\varphi. \tag{12}$$

By using a low number of harmonics, an additional smoothing effect is achieved. The objective of using a weighted sum of spherical harmonics is that the resulting shapes do not have regions of high curvature (which generally matches the characteristics of biological objects) but are not as simplistic as ellipsoidal models. Since the spherical harmonics represent the surface in closed form, we can precisely sample the shape into any set of slices, and obtain analytic expressions for the resulting contours. See Brechbühler et al. (1995) for more details about surface parameterization using spherical harmonics.

Several random synthetic shapes generated using this procedure with different widths for the radial distribution are shown in Fig. 6. As the parameter σ increases, the bumpiness increases as well as the overall size of the shape. In our experiment, we generated the radial random variables from the uniform distribution over $[3, 3+\sigma]$ with five different σ values: $\sigma \in \{0.5, 1.0, 1.5, 2.0, 2.5\}$. For each σ, 20 shapes were generated, giving a total of 100 shapes. Each shape was smoothed using a 2×2 moving average filter and approximated by spherical harmonics up to degree 7. Each shape can be fit in a cube with a side length of 10 (there is no corresponding physical unit since the shapes are synthetic).

Each randomly generated shape was then sampled into sets of 20 and 17 equally spaced slices as shown in Fig. 7. The sets of 20 slices were set aside as reference and the sets of 17 slices were resampled using both the implicit function and EFD methods to yield 20 slices at the axial locations of the reference set. For each slice in a shape, we measured the Hausdorff distance between the reference slice and the resampled slice, defined as

$$D_{\mathrm{H}}(A,B) = \max\{d_{\mathrm{H}}(A,B), d_{\mathrm{H}}(B,A)\}, \tag{13}$$

where

$$d_{\mathrm{H}}(A,B) \equiv \max\{\min\{\|a-b\|, b \in B\}, a \in A\}, \tag{14}$$

for two sets of points, A and B. This roughly measures the maximum distance between the two curves. One slice from the experiment is shown in Fig. 8. The Hausdorff distances

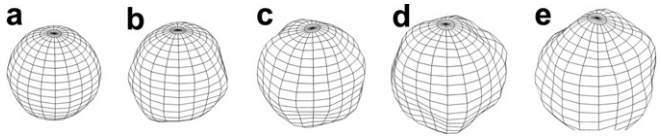

Fig. 6. Examples of synthetic shapes constructed using different widths for the radial distribution. (a) $\sigma = 0.5$. (b) $\sigma = 1.0$. (c) $\sigma = 1.5$. (d) $\sigma = 2.0$. (e) $\sigma = 2.5$. All shapes are in the same scale.

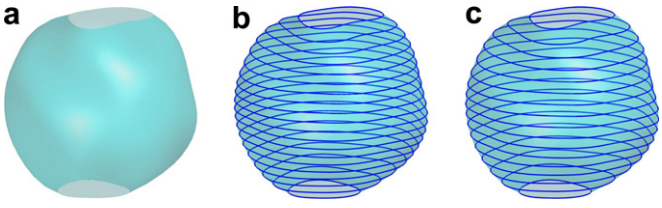

Fig. 7. (a) A synthetic shape generated using spherical harmonic descriptors. (b) The shape is sampled into 20 equally spaced parallel slices, used as reference slices. (c) The shape is sampled into 17 equally spaced parallel slices, used as input to each resampling method to produce 20 parallel slices to compare with the reference slices.

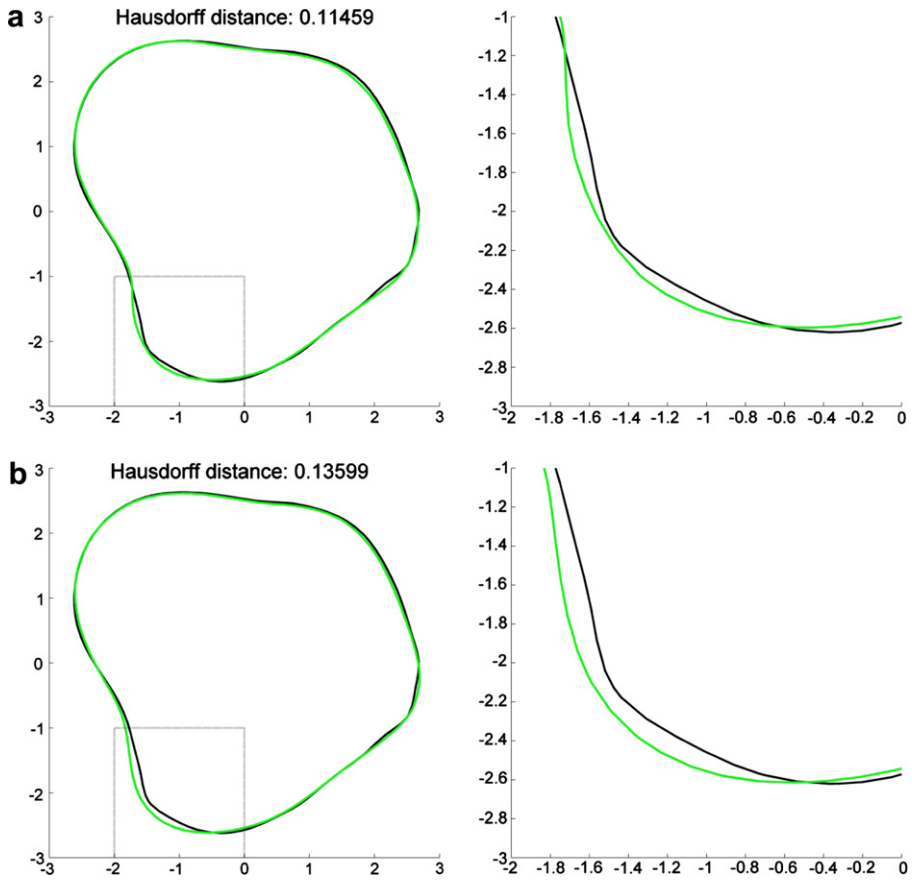

Fig. 8. An example of one synthetic slice interpolated using both the proposed method and the implicit function method. (a) The original slice (black) and the resampled slice using EFD (green) (Hausdorff distance = 0.11). (b) The original slice (black) and the resampled slice using the implicit function method (green) (Hausdorff distance = 0.14). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

from all 20 slices in each shape were then averaged to yield an overall measure for each shape/method combination. The results for all 100 shapes are summarized in Table 1 as a function of the uniform distribution width. As can be seen from the table, while the error grows for "bumpier" shapes, the overall performance for both methods is virtually identical. The average computational time required to resample each shape was 5.3 min using the EFD method and 38.4 min using the implicit function method, a factor of 7.2 difference.

Table 1
Results for the synthetic shape experiment

Radial parameter	EFD		IF	
	Mean	SD	Mean	SD
$\sigma = 0.5$	0.020	0.002	0.020	0.001
$\sigma = 1.0$	0.029	0.004	0.029	0.004
$\sigma = 1.5$	0.039	0.005	0.040	0.005
$\sigma = 2.0$	0.051	0.011	0.052	0.010
$\sigma = 2.5$	0.062	0.011	0.063	0.010

The EFD columns give the mean and standard deviation of the Hausdorff distance for resampling using elliptic Fourier descriptors, while the IF columns give the mean and standard deviation for resampling using the implicit function method. The results are basically identical, but the EFD required about 7.2 times less computational time than the IF method.

5.2. Clinical data

The data used in this experiment is composed of contours of the prostate from serial CT images, drawn by a physician. Several examples of such contours are shown in Fig. 1. All shapes can be fit in a cube with a side length of 8 cm. Unlike the synthetic shape experiment described above, here there are no reference slices to compare with the resampled slices. Hence, we compared the resampled slices obtained using both methods against each other.

Table 2
Results for the prostate shape experiment

Patient	Datasets	Slices per dataset	Mean (mm)	SD (mm)
1	14	9–14	0.70	0.16
2	17	11–13	0.56	0.12
3	17	10–12	0.81	0.19
4	19	9–18	0.64	0.16
5	19	9–15	0.66	0.21
6	13	8–9	0.48	0.08

The columns give the mean and standard deviation of the Hausdorff distance between the resliced EFD and IF contours for each dataset. The EFD required about 6.2 times less computational time than the IF method.

Six datasets were used for the experiment in this section; each dataset corresponds to the same patient's prostate on different days of radiation treatment. The number of examples in each dataset ranges from 13 to 19, giving a total of 99 sets of contours. The number of axial slices per prostate ranges from 8 to 18. Each set of contours was resampled into 20 slices using both the EFD and implicit function methods, and the Hausdorff distance between the two results was measured slice-by-slice and averaged for each set as described above. The results are summarized in Table 2 on a patient-by-patient basis.

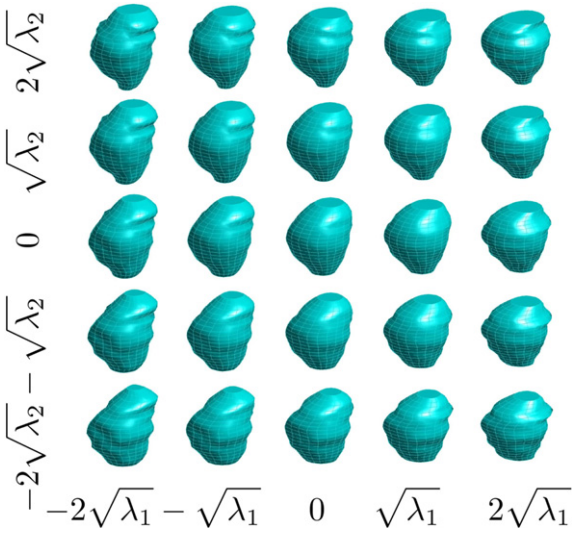

Fig. 9. A 3D active shape model built using the resliced contours can produce a variety of shapes that are consistent with the training data (in this case, a single patient's prostate observed on 13 consecutive days of radiotherapy). The horizontal axis shows the variation along the first mode, and the vertical axis shows the variation along the second mode. The learned mean prostate shape is in the middle of the figure.

The difference between the EFD method and the implicit function method is below 1mm on the average, with a variance of 0.04 mm. This error seems acceptable, since it is less than reported intra-observer variability for contouring the prostate in CT images (which is itself less than reported inter-observer variability) (Fiorino et al., 1998). The average computational time required to resample each shape was 3.2 min using the EFD method and 19.9 min using the implicit function method, a factor of 6.2 difference.

The resampled contours obtained by the method can subsequently be used for 3D shape modeling and analysis. For example, we used the resampled prostate contours described here to build active shape models for the automatic segmentation of the prostate in CT imagery in the context of radiotherapy planning (Freedman et al., 2005). Since axial correspondence is established after resampling the original contours into slices at desired locations, it is easy to establish points at consistent locations on each 2D slice based on arc-length and angle constraints, resulting in automatic selection of 3D landmarks. One resulting shape model is illustrated in Fig. 9, which indicates the mean shape and the first two orthogonal modes of variation for the prostate of a single patient across 13 different radiotherapy fractions. We designed a segmentation technique that evolves the parameters of the active shape model so that the histogram of pixel intensities within the shape's interior matches as well as possible with an appearance model learned from training data. The result of applying the 3D segmentation technique with a 10-mode active shape model to a new volumetric CT scan from the same patient is illustrated in Fig. 10. In this case, the median surface-to-surface distance between the implicit surface interpolating the physician's ground truth contours and the segmentation estimated by the model is 0.34 mm. Constraining the seg-

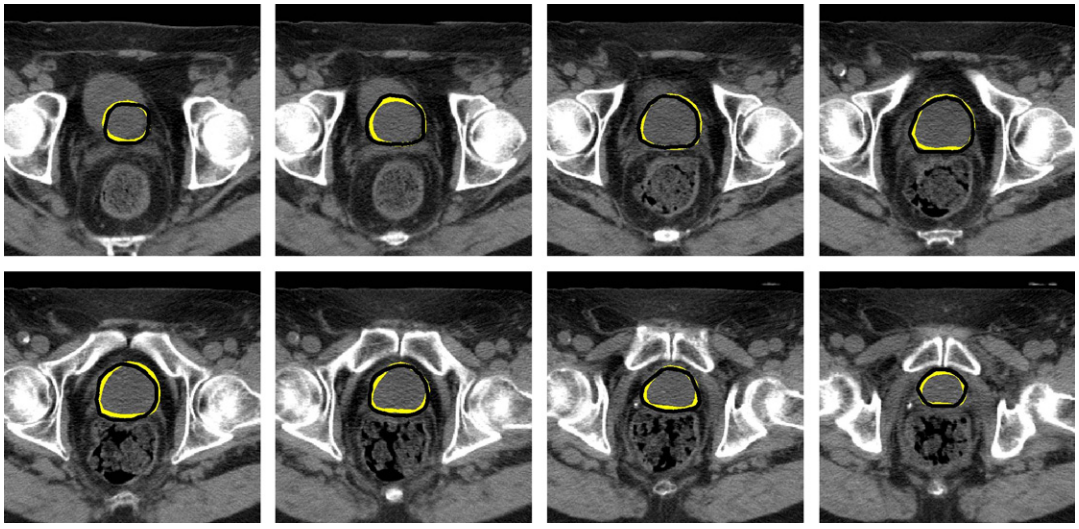

Fig. 10. Using the appearance-histogram-based 3D segmentation method described in Freedman et al. (2005), the 3D active shape model is used to segment the prostate from a new volumetric CT scan from the same patient. The black contours indicate an expert's hand-drawn estimate of the prostate, and the yellow contours indicate the automatic segmentation result. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

mentation with the active shape model learned from expertly contoured training data is important for this problem, since the prostate presents no distinct "edge" in the image itself. A segmentation algorithm based purely on intensity would naturally expand to segment the entire region that has similar intensity to the prostate, including the pelvic floor muscles and seminal vesicles. For more details on the modeling and segmentation techniques, refer to Freedman et al. (2005).

6. Conclusions

We introduced a new algorithm using elliptic Fourier descriptors for reslicing axially sampled 3D shapes acquired from medical imagery. The EFD method shows performance virtually indistinguishable from a method using the common approach of fitting and resampling an implicit function, but has a much lower computational cost. Therefore, we believe there is no compelling reason to use the generic implicit function method over the simple EFD method for the axial resampling application. While the EFD method does not apply to certain types of 3D shapes (e.g. shapes with holes, non-closed shapes), we find it to be quite useful for interpolating soft-tissue organs from tomographic imagery (e.g. the prostate, bladder, and rectum), since these objects generally have simple, blob-like cross-sections without branches or regions of high curvature. Several other organs would have similar blob-like shapes, such as the liver or kidneys. Manually determining a large number of definitive landmarks on such organs across a population of patients or even across scans of the same patient on different days would be difficult even for an expert. As the shapes to be matched become more complex, a point set alignment algorithm such as Chui et al. (2004) might be more appropriate, with a possible penalty in speed. In such cases, it may also be easier to semi-automatically associate meaningful surface landmarks across organs from different patients, compared to the prostate case. Our goal in this paper was to show that for simple shapes, fast, accurate reslicing solutions are possible.

Acknowledgments

This work was supported by CenSSIS, the NSF Center for Subsurface Sensing and Imaging Systems, under the award EEC-9986821. Thanks to D. Michael Lovelock of Memorial Sloan-Kettering Cancer Center for providing the prostate datasets and contours.

References

Alexa, M., 2002. Recent advances in mesh morphing. Computer Graphics Forum 21 (2), 173–196.

Alliez, P., Ucelli, G., Gotsman, C., Attene, M., 2006. Recent advances in remeshing of surfaces. In: de Floriani, L., Spagnuolo, M. (Eds.), Shape Analysis and Structuring. Springer, Berlin. Available from: <ftp://ftp-sop.inria.fr/geometrica/alliez/survey_remeshing.pdf>.

Arfken, G.B., Weber, H.J., 2005. Mathematical Methods for Physicists, sixth ed. Academic Press, London.

Barequet, G., Sharir, M., 1996. Piecewise-linear interpolation between polygonal slices. Computer Vision and Image Understanding 63 (2), 251–272.

Billings, S.D., Beatson, R.K., Newsam, G.N., 2002. Interpolation of geophysical data using continuous global surfaces. Geophysics 67 (6), 1810–1822.

Blanz, V., Vetter, T., 1999. A morphable model for the synthesis of 3D faces. In: SIGGRAPH 1999, pp. 187–194.

Bolle, R.M., Vemuri, B.C., 1991. On three-dimensional surface reconstruction methods. IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (1), 1–13.

Bookstein, F., 1989. Principal warps: thin plate splines and the decomposition of deformations. IEEE Transactions on Pattern Analysis and Machine Intelligence 11, 567–585.

Brechbühler, C., Gerig, G., Kübler, O., 1995. Parameterization of closed surfaces for 3-D shape description. CVGIP: Image Understanding 61, 154–170.

Carr, J.C., Beatson, R.K., Cherrie, J.B., Mitchell, T.J., Fright, W.R., McCallum, B.C., Evans, T.R., 2001. Reconstruction and representation of 3D objects with radial basis functions. In: Computer Graphics (Proceedings of SIGGRAPH), pp. 67–76.

Cherrie, J.B., Beatson, R.K., Newsam, G.N., 2002. Fast evaluation of radial basis functions: methods for generalized multiquadrics in R^n. SIAM Journal on Scientific Computing 23 (5), 1549–1571.

Chui, H., Rangarajan, A., Zhang, J., Leonard, C., 2004. Unsupervised learning of an atlas from unlabeled point-sets. IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (2), 160–172.

Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1992. Training models of shape from sets of examples. In: Proceedings of the 3rd British Machine Vision Conference, pp. 9–18.

Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., 1995. Active shape models: their training and application. Computer Vision and Image Understanding 61 (1), 38–59.

Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J., 2002. A minimum description length approach to statistical shape modeling. IEEE Transactions on Medical Imaging 21 (5), 525–537.

Delingette, H., 1994. Simplex meshes: a general representation for 3D shape reconstruction. In: CVPR 1994, pp. 856–859.

Delingette, H., Ayache, N., 2004. Soft tissue modeling for surgery simulation. In: Ayache, N. (Ed.), Computational Models for the Human Body. In: Ciarlet, P. (Ed.), Handbook of Numerical Analysis. Elsevier, Amsterdam, pp. 453–550.

Duchon, J., 1977. Splines minimizing rotation-invariant semi-norms in Sobolev spaces. In: Schempp, W., Zeller, K. (Eds.), Constructive Theory of Functions of Several Variables, Lecture Notes in Mathematics. Springer, Berlin, pp. 85–100.

Elvins, T.T., 1992. A survey of algorithms for volume visualization. Computer Graphics (ACM) 26 (3), 194–201.

Farin, G., 2001. Curves and surfaces for CAGD: a practical guide, fifth ed. Morgan Kaufmann, Los Altos, CA.

Fiorino, C., Reni, M., Bolognesi, A., Cattaneo, G.M., Calandrino, R., 1998. Intra- and inter-observer variability in contouring prostate and seminal vesicles: implications for conformal treatment planning. Radiotherapy and Oncology 47, 285–292.

Freedman, D., Radke, R.J., Zhang, T., Jeong, Y., Lovelock, D.M., Chen, G.T.Y., 2005. Model-based segmentation of medical imagery by matching distributions. IEEE Transactions on Medical Imaging 24, 281–292.

Grimson, W.E.L., 1982. A computational theory of visual surface interpolation. Philosophical Transactions of the Royal Society of London, Series B 298 (1092), 395–427.

Hoppe, H., DeRose, T., Duchamp, T. 1992. Surface reconstruction from unorganized points. In: Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques, pp. 71–78.

Kraevoy, V., Sheffer, A., 2004. Cross-parameterization and compatible remeshing of 3D models. In: SIGGRAPH 2004, pp. 861–869.

Kuhl, F.P., Giardina, C.R., 1982. Elliptic Fourier features of a closed contour. Computer Graphics and Image Processing 18 (3), 236–258.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Medical Image Analysis 1 (2), 91–108.

Meyers, D., Skinner, S., Sloan, K., 1992. Surfaces from contours. ACM Transactions on Graphics 11 (3), 228–258.

Praun, E., Sweldens, W., Schröder, P., 2001. Consistent mesh parameterizations. In: SIGGRAPH 2001, pp. 179–184.

Staib, L.H., Duncan, J.S., 1992. Boundary finding with parametrically deformable models. IEEE Transactions on Pattern Analysis and Machine Intelligence 14 (11), 1061–1075.

Styner, M.A., Rajamani, K.T., Nolte, L.-P., Zsemlye, G., Szëkely, G., Taylor, C.J., Davies, R.H., 2003. Evaluation of 3D correspondence methods for model building. In: Information Processing in Medical Imaging.

Tello, R., 1996. Fourier descriptors for computer graphics. IEEE Transactions on Systems Man and Cybernetics 25 (5), 861–865.

Terzopoulos, D., Metaxas, D., 1991. Dynamic 3D models with local and global deformations: deformable superquadrics. IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (7), 703–714.

Turk, G., O'Brien, J.F., 1999. Shape transformation using variational implicit functions. In: SIGGRAPH 1999, pp. 335–342.

Vemuri, B.C., Radisavljevic, A., 1993. From global to local: a continuum of shape models with fractal priors. In: CVPR 1993, pp. 307–313.

Wang, Y., Peterson, B., Staib, L., 2000. Shape-based 3D surface correspondence using geodesics and local geometry. In: CVPR 2000, pp. 644–651.

Wang, F., Vemuri, B., Rangarajan, A. 2006. Groupwise point pattern registration using a novel CDF-based Jensen–Shannon divergence. In: IEEE Computer Vision and Pattern Recognition, pp. 1283–1288.

Wu, M.-F., Sheu, H.-T., 1998. Representation of 3D surfaces by two-variable Fourier descriptors. IEEE Transactions on Pattern Analysis and Machine Intelligence 20 (8), 858–863.