

Automatic Construction of Correspondences for Tubular Surfaces

Toon Huysmans, Jan Sijbers, and Brigitte Verdonk

Abstract—Statistical shape modeling is an established technique and is used for a variety of tasks in medical image processing, such as image segmentation and analysis. A challenging task in the construction of a shape model is establishing a good correspondence across the set of training shapes. Especially for shapes of cylindrical topology, very little work has been done. This paper describes an automatic method to obtain a correspondence for a set of cylindrical shapes. The method starts from an initial correspondence which is provided by cylindrical parameterization. The quality of the obtained correspondence, measured in terms of the description length, is then improved by deforming the parameterizations using cylindrical b-spline deformations and by optimization of the spatial alignment of the shapes. In order to allow efficient gradient-guided optimization, an analytic expression is provided for the gradient of this quality measure with respect to the parameters of the parameterization deformation and the spatial alignment. A comparison is made between models obtained from the correspondences before and after the optimization. The results show that, in comparison with parameterization-based correspondences, this new method establishes correspondences that generate models with significantly increased performance in terms of reconstruction error, generalization ability, and specificity.

Index Terms—Point correspondence problem, statistical shape models, tubular structures, minimum description length, image segmentation, image shape analysis.

1 INTRODUCTION

SHAPE correspondences and the derived statistical shape models have a wide range of applications in medical image computing [1], [2]. They have been used to analyze shape differences between different classes of objects, for example, the lateral ventricles of schizophrenics versus a healthy group [3]. They have also been employed to gain more knowledge about the anatomical variability of certain organs or bones as, for example, the human ear canal [4]. Such knowledge can in turn be used to reconstruct malformed, missing, or fractured bone structures [5]. A widespread application of statistical shape models is their use as prior knowledge in automatic image segmentation [6], [7], [8]. The probability density function of the shape is estimated from a set of manual segmentations. This knowledge is then used to guide the segmentation process of an unseen instance and to restrict the segmentation result to the class of plausible shapes.

The major hurdle in the construction of a statistical shape model is establishing a dense correspondence over the surfaces of a large set of training shapes. These correspondences should be of high quality, i.e., the correspondence

should match anatomically equivalent points over the surfaces. If this requirement is not met, artificial modes of variation are introduced into the shape model and this has a negative effect on the performance of the model when used for image segmentation or interpretation [9].

For 2D shapes, a correspondence between the boundary curves of the individual shapes is often defined by manual landmarking [10]. Although this approach is feasible, it turns out to be a time-consuming and error-prone task. In principle, the approach can be extended to 3D, but it becomes highly impractical due to the large amount of landmarks that need to be located and the increased level of difficulty in pinpointing them. Several approaches have been proposed to automate this labor-intensive procedure.

A relatively simple but effective approach is to establish the sought-after correspondence by means of surface parameterization [11], [12], [13], [14], [15], [16], [17]. Thereby, a one-to-one map is constructed between each surface of the set and some common, predefined, and usually mathematically simple parameter domain, such as a planar disc or a sphere. For each surface, the obtained one-to-one map associates a 2D parameter coordinate with each point of the 3D surface. The surface-to-surface correspondence is then defined by the assigned parameter values, i.e., points between the surfaces correspond when they share the same parameters.

Although the parameterization approach produces valid correspondences, there is still room for improvement. The parameterization of each shape in the training set is done independently of the other shapes; thus, correlations between the shapes are not taken into account. When an approach takes this extra information into account, it can obtain better correspondences. Surface parameterization can provide a good initial correspondence for such a method. Davies et al. developed such a correspondence method in [18]. They use the description length of the

• T. Huysmans and J. Sijbers are with the Department of Physics, University of Antwerp (CDE), Universiteitsplein 1, B-2610 Antwerp, Belgium.
E-mail: {toon.huysmans, jan.sijbers}@ua.ac.be.

• B. Verdonk is with Emerging Computational Techniques, Department of Mathematics and Computer Science, University of Antwerp (CMI), Middelheimlaan 1 (M.G.103a), B-2020 Antwerpen, Belgium.
E-mail: brigitte.verdonk@ua.ac.be.

Manuscript received 28 May 2008; revised 17 Dec. 2008; accepted 2 Apr. 2009; published online 22 Apr. 2009.

Recommended for acceptance by A. Srivastava, J.N. Damon, I.L. Dryden, and I.H. Jermyn.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMISI-2008-05-0308.

Digital Object Identifier no. 10.1109/TPAMI.2009.93.

derived shape model as a measure of correspondence quality and optimize the correspondences with respect to this criterion. To date, their minimum description length (MDL) approach is considered the method of choice for the construction of correspondences [9].

Most of the aforementioned techniques focus on sets of surfaces of *spherical* topology since these are prevalent in biomedical image processing, e.g., kidneys, liver, and brain ventricles. Nonetheless, the developed principles can be translated to surfaces of other topologies. More specifically, this paper deals with the translation of these principles to sets of surfaces of *cylindrical* topology, such as the trachea, cochlear channels, aortic aneurysms, and the rectum. For such surfaces, the cylinder is the natural choice for the parameter domain. An initial correspondence is obtained by specialized cylindrical parameterization techniques. Following this, the correspondence is improved by an optimization framework that adopts the MDL criterion as a measure of correspondence quality.

The remainder of this paper is organized as follows: Related work is discussed in Section 1.1. An overview of the correspondence method is given in Section 2. Parameterization of surfaces of cylindrical topology is discussed in Section 3. The measure of correspondence quality, namely, description length, is treated in Section 4. Section 5 details the procedure of how an initial correspondence is obtained, and Section 6 elaborates on the correspondence optimization. Experimental results of applying the method to a number of phantoms and real data sets are presented and discussed in Section 7. Section 8 concludes the paper.

1.1 Related Work

Parameterization of surfaces with disc-like topology onto a planar convex region has been addressed in [11]. A parameterization technique to parameterize surfaces of spherical topology onto the sphere has been developed by Brechbühler et al. [12] and is known as SPHARM. This was later used to model the shape of brain structures for segmentation [6] and analysis [3]. Conformal parameterization techniques for surfaces of spherical topology were introduced in [19], [20], [21]. A more efficient alternative to SPHARM, utilizing a progressive surface representation, was provided by Praun and Hoppe [13]. Based on this spherical parameterization technique, Huysmans et al. have developed a cylindrical parameterization technique to parameterize tubular surfaces onto the cylinder in a progressive way [14]. Conformal parameterization of surfaces of cylindrical topology was addressed in [22], [23], and [24]. Algorithms for more complex topologies (genus- n) have also been proposed [15], [16], [17], but, from a correspondence optimization point of view, these algorithms are suboptimal since they rely on a heuristic or manually defined surface chartification. In [25], Ricci flow is used to construct a seamless periodic tiling of a genus- n surface in the plane, sphere, or hyperbolic space and it is known as the universal covering space. The obtained map is angle preserving and can be used to construct harmonic maps between surfaces of the same topology [32].

In [26], Kottchef and Taylor used the determinant of the landmark covariance matrix as an optimization objective. Based on their ideas, Davies et al. developed an MDL formulation for the assessment of correspondence quality

for 2D curves [18], which was later simplified by Thodberg [27]. The method for building MDL correspondences has been extended to surfaces of spherical topology [28], which was later improved by Heimann et al. [29] in terms of computational efficiency. Horkeaw and Yang applied the MDL principle to surfaces of disc-like topology [30], and an extension to more complex topologies was obtained by cutting surfaces into topological discs prior to optimization [31]. This, of course, constrains the optimization along the cuts. In the recent work of Li et al. [32], a globally optimal map, in terms of harmonic energy, is obtained between two surfaces sharing arbitrary complex topology. An extension of their method to populations of surfaces and to other correspondence metrics, e.g., MDL, is, however, not addressed. Recently introduced point-based correspondence techniques can also handle arbitrary topology. In [33], Ferrarini et al. use self-organizing maps to obtain a pairwise correspondence between each population member and a template. In [34], Cates et al. use particle systems to optimize geometry sampling and groupwise correspondence with respect to an information-theoretic measure comparable to MDL. Point-based correspondence techniques are promising but problems occurring with highly convoluted surfaces still need to be addressed.

Little work has been done in shape modeling for cylindrical surfaces. In [35], de Bruijne et al. proposed an improved modeling scheme for tubular objects. They used a manually determined correspondence. Some of the previously discussed techniques could be employed to build a correspondence for a set of cylindrical shapes. For example, the spherical MDL framework can be applied when the holes of the cylindrical shapes are closed but this can result in an invalid correspondence at the boundaries and this approach will not perform well for elongated surfaces [36]. The correspondence methods that rely on chartification [16], [31] could also be applied to cylindrical surfaces, but, obviously, the performance of the chartification heuristic will influence the final quality. The method proposed in this paper does not suffer from these drawbacks since, here, the description length minimization method is translated specifically to treat sets of surfaces of cylindrical topology.

2 METHOD OUTLINE

The method treated in this paper constructs a dense surface correspondence together with an optimal spatial alignment for an arbitrary population of cylindrical surfaces. It proceeds in two steps. First, a correspondence is derived from the surface parameterizations by alignment of the surfaces and parameterizations. The result is referred to as the *rigid correspondence*. Then, the rigid correspondence is improved by applying local, nonrigid, deformations to the parameterizations while keeping the surfaces optimally aligned. The finally obtained correspondence is referred to as the *nonrigid correspondence*. An overview of these two steps can be found in Fig. 1.

The input to the method is a set of n_s triangle surfaces $\{\mathcal{M}_1, \dots, \mathcal{M}_{n_s}\}$ of cylindrical topology. Each surface \mathcal{M}_i is defined by a tuple (V_i, T_i) , where V_i is the set of n_{V_i} vertices $\{v_1^i, \dots, v_{n_{V_i}}^i\}$ with $v_j^i \in \mathbb{R}^3$ and T_i is the set of n_{T_i} triangles $\{t_1^i, \dots, t_{n_{T_i}}^i\}$.



Fig. 1. (a) Flow chart visualization of the rigid correspondence construction. First, each surface is equipped with a cylindrical parameterization and approximated with a b-spline surface. Then, the surfaces are brought into a reference coordinate system by alignment of their principal axes. This is followed by the alignment of the parameter spaces of the surfaces by optimization w.r.t. the correspondence quality, i.e., model description length. Finally, the rigid correspondence is obtained after a second optimization where both the spatial and parameterization alignment parameters are set free. (b) Flow chart visualization of the nonrigid correspondence improvement. The correspondence is improved by simultaneously optimizing the parameters of the spatial alignments and the b-spline parameterization deformations w.r.t. the correspondence quality. The optimizations are performed at multiple resolution levels sequentially, starting at the lowest level with coarse deformations and gradually adding more detail with each new level. Finally, the nonrigid correspondence is obtained.

The construction of the rigid correspondence is covered by the flow chart in Fig. 1a and presented in detail in Section 5. The rigid correspondence, denoted as

$\{x_1, \dots, x_{n_s}\}$, is obtained by parameterization of the surfaces \mathcal{M}_i , followed by rigid alignment of the surfaces and their parameter spaces. The spatial transformations for

the alignment of the surfaces are denoted as $\tau^\circ(\bullet | \Phi_i^{\tau^\circ})$ and the parameterization transformations for the parameter space alignments are denoted as $\rho^\circ(\bullet | \Phi_i^{\rho^\circ})$. In order to avoid convergence to local minima, the optimal parameters, $\hat{\Phi}_i^{\tau^\circ}$ and $\hat{\Phi}_i^{\rho^\circ}$, of these transformations are obtained by a number of consecutive optimizations with respect to the MDL correspondence quality criterion, denoted as μ . The construction pipeline for the rigid correspondence comprises following steps:

1. A cylindrical parameterization x_i° is constructed for each surface \mathcal{M}_i . It constitutes a one-to-one map between the surface \mathcal{M}_i and the open-ended cylinder \mathcal{C}_h^2 of height h .
2. A smooth b-spline approximation \tilde{x}_i° is constructed for each surface. The b-spline representation results in smooth optimization objectives and provides a multiresolution representation of the surface.
3. An initial spatial alignment of the surfaces is obtained by alignment of their principal axes.
4. While keeping the spatial alignment fixed, an alignment of the parameterizations is determined by optimization with respect to MDL.
5. Both the spatial and parameter space alignments are improved by simultaneous optimization with respect to MDL.

In order to obtain the rigid correspondence, the optimal spatial transformations and parameterization transformations are applied to the b-spline approximations of each of the parameterized surfaces:

$$x_i = \tau^\circ(\bullet | \hat{\Phi}_i^{\tau^\circ}) \circ \tilde{x}_i^\circ \circ \rho^\circ(\bullet | \hat{\Phi}_i^{\rho^\circ}). \quad (1)$$

The construction of nonrigid correspondence is covered by the flow chart in Fig. 1b and is presented in detail in Section 6. The nonrigid correspondence is obtained as an improvement of the rigid correspondence by applying a nonrigid b-spline transformation to the parameter space of each surface and simultaneously optimizing their spatial alignment. Again, the MDL correspondence criterion was used as the optimization objective. For reasons of efficiency and in order to avoid local minima, the optimization is done successively at a number of resolution levels. With each resolution level L , the grid size $m_{u^{(0)}}^L \times m_{u^{(1)}}^L$ for the b-spline approximation \tilde{x}_i^L of x_i and the grid size $n_{u^{(0)}}^L \times n_{u^{(1)}}^L$ for the b-spline parameterization deformation $\rho^L(\bullet | \Phi_i^{\rho^L})$ are doubled, allowing for a more detailed correspondence improvement. Also, the number of landmarks n_p^L used to calculate the shape model is increased and the convergence tolerance becomes more strict. The spatial transformation $\tau(\bullet | \Phi_i^{\tau^L})$ is rigid and has optimal parameters $\hat{\Phi}_i^{\tau^L}$. The optimization of a resolution level L is initialized with the result from the previous resolution level $L - 1$. The final correspondence is obtained from the optimal transformation parameters, $\hat{\Phi}_i^{\tau^L}$ and $\hat{\Phi}_i^{\rho^L}$, of the last resolution level:

$$\hat{x}_i^L = \tau(\bullet | \hat{\Phi}_i^{\tau^L}) \circ \tilde{x}_i^L \circ \rho^L(\bullet | \hat{\Phi}_i^{\rho^L}). \quad (2)$$

From the final correspondence $\{\hat{x}_1, \dots, \hat{x}_n\}$, a map from surface \mathcal{M}_i to surface \mathcal{M}_j can be obtained by composition of the inverse of parameterization \hat{x}_i with the

parameterization \hat{x}_j , i.e., $q = \hat{x}_j \circ \hat{x}_i^{-1}(p)$, where $p \in \mathcal{M}_i$ and the corresponding point $q \in \mathcal{M}_j$.

3 SURFACE REPRESENTATION

3.1 Parameterization

Starting from a cylindrical surface \mathcal{M} defined by its vertices V and triangles T , a parameterized version x of \mathcal{M} is obtained by assigning a unique pair of cylindrical coordinates to each point of the surface \mathcal{M} . Usually, the parameter coordinates are only defined explicitly at the vertices of \mathcal{M} and the extension over the triangles is implied by barycentric interpolation of the parameter coordinates at the vertices. To be more precise, for surfaces of cylindrical topology, the parameterization x is a homeomorphic function from \mathcal{C}_h^2 to the surface \mathcal{M} , i.e.,

$$\begin{aligned} x : [0, 2\pi] \times [0, h] &\rightarrow \mathcal{M} \subset \mathbb{R}^3 \\ u &\rightarrow x(u), \end{aligned}$$

where \mathcal{C}_h^2 denotes the open-ended 2D cylinder of length h with unit radius, which is parameterized by an angular coordinate $u^{(0)}$ and an axial coordinate $u^{(1)}$, i.e., $u = (u^{(0)}, u^{(1)})$. For x to be a homeomorphism, it must be a bijective, continuous function, and have a continuous inverse. If the topology of \mathcal{M} is consistent, such a homeomorphism can always be obtained although a solution is not unique. This can be seen from the fact that the composition $x \circ \rho$ of any automorphism ρ of the cylinder with the parameterization x of the surface \mathcal{M} again is a valid parameterization of the same surface. The particular solution that a parameterization technique will propose is usually the result of the minimization of an energy functional. Different functionals result in different parameterizations, the quality of which depends on the actual application.

In order to obtain good correspondences between surfaces, a parameterization technique should create similar maps for similar surfaces. In addition, it is also desirable that it retains relative areas and angles as much as possible (i.e., distortion). When the parameterizations systematically suffer from large area distortions, undersampling of parts of the surface can occur in the final correspondence. For a tubular surface with an approximately constant width, the use of an harmonic parameterization technique is recommended, e.g., [22], [23]. The harmonic cylindrical parameterization is uniquely defined as the solution of a system of linear equations. As a consequence, it is computationally very efficient. However, it fails to keep area distortions within acceptable bounds when a large variation in cross-sectional diameter is present. For these cases, the progressive nonlinear cylindrical parameterization technique from [14] is a better alternative. It allows control over the trade-off between angle and area distortions, at the cost of increased computation time. See Fig. 2 for a comparison of the two methods.

3.2 B-Spline Representation

Since the surface \mathcal{M} is a piecewise linear surface, the partial derivatives of the surface coordinates with respect to the parameter coordinates ($\frac{\partial x}{\partial u^{(0)}}$ and $\frac{\partial x}{\partial u^{(1)}}$) are discontinuous at the triangle boundaries, excluding boundaries between

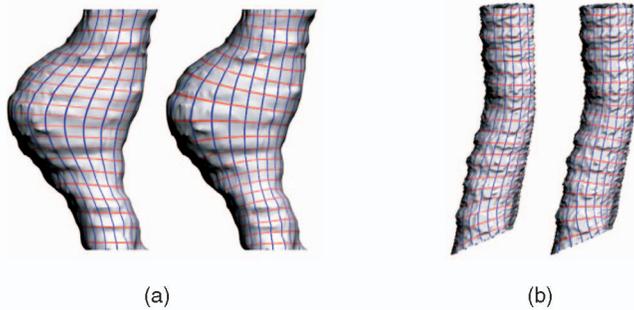


Fig. 2. A qualitative comparison of the harmonic parameterization method with the nonlinear parameterization method. The parameterization of each surface is visualized by the blue and red lines, which correspond to the $iso-u^{(0)}$ and $iso-u^{(1)}$ lines of the cylinder, respectively. (a) Two parameterizations of a surface with a large variation in cross-sectional diameter are shown. It is clearly visible that the harmonic parameterization (right) has large-area distortions at the bulge compared to the nonlinear method (left), so the nonlinear method is preferred for this kind of surface. (b) Two parameterizations of a surface with an approximately constant cross-sectional diameter are shown. For this surface, both methods can be used since they perform approximately equally well.

coplanar triangles. This is undesirable since these partial derivatives are utilized in the gradient-guided optimization of the correspondences. Using cubic b-splines, approximation of x will result in a surface \tilde{x} that is C^2 continuous within the b-spline patches and C^1 continuous at the patch boundaries. By varying the number of control points used for the approximation, a multiresolution representation of the surface can be obtained. Moreover, evaluation of the b-spline representation \tilde{x} is much faster than evaluation of the triangle-based representation x . This is because point location in a regular grid is far more efficient than point location in a triangulation. A good b-spline approximation of x can be achieved with a number of control points much lower than the original number of vertices that x comprises, resulting also in better memory efficiency.

The approximation uses a 2D tensor product b-spline surface. The cubic b-spline kernel, denoted by β , is defined as in [37]:

$$\beta(u) = \begin{cases} \frac{1}{6}(3|u|^3 - 6u^2 + 4), & |u| \in [0, 1], \\ \frac{1}{6}(2 - |u|)^3, & |u| \in [1, 2], \\ 0, & |u| \in [2, \infty]. \end{cases} \quad (3)$$

The b-spline surface is defined by a uniform grid of knots $K = \{k_{ij}\}$ positioned on the cylinder C_h^2 and a corresponding grid of control points $P = \{p_{ij}\}$ in \mathbb{R}^3 . An example knot grid is shown in Fig. 4a. The b-spline surface then has the following form:

$$\tilde{x}(u|K, P) = \sum_{i=-1}^{m_{u^{(0)}}+1} \sum_{j=-1}^{m_{u^{(1)}}} \beta\left(\frac{u - k_{ij}}{\Delta}\right) p_{ij}, \quad (4)$$

where $m_{u^{(0)}}$ is the number of knots in the $u^{(0)}$ -direction and $m_{u^{(1)}}$ is the number of knots in $u^{(1)}$ -direction that are within the range $[0, 2\pi] \times [0, h]$. The 2D cubic b-spline kernel $\beta(u)$ is separable, i.e., $\beta(u) = \beta(u^{(0)})\beta(u^{(1)})$, and the grid spacing is denoted as $\Delta = (\frac{2\pi}{n_{u^{(0)}}, \frac{h}{n_{u^{(1)}}-1})$. The division in the argument of the b-spline kernel is executed element-wise. In order to obtain a closed and smooth surface at the

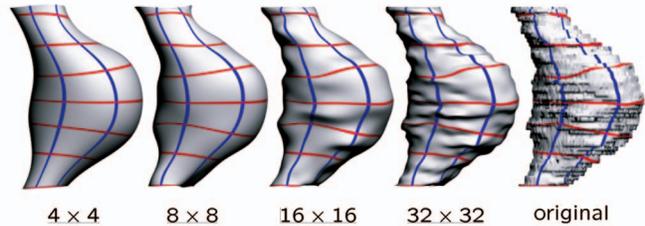


Fig. 3. Approximation of a parameterized surface using a b-spline surface with an increasing number of control points (indicated as $m_{u^{(0)}} \times m_{u^{(1)}}$).

parameter boundary $u^{(0)} = 2\pi$, the control points satisfy the following conditions:

$$\begin{cases} p_{-1,j} & = & p_{m_{u^{(0)}}-1,j}, \\ p_{m_{u^{(0)},j} & = & p_{0,j}, \\ p_{m_{u^{(0)}+1,j} & = & p_{1,j}. \end{cases} \quad (5)$$

Least-squares fitting is used in order to find a set of $m_{u^{(0)}}(m_{u^{(1)}} + 2)$ control points that provides a good approximation to the surface x . For this purpose, a set of m_p uniformly distributed parameter locations is chosen on the cylinder: $U^{m_p} = \{u_1, \dots, u_{m_p}\} \subset C_h^2$. Using these parameter locations, the approximation error can be determined as the sum of the squared distances between the points on the original surface and the points on the approximating surface at corresponding parameter locations. The set of control points \hat{P} for which this error is minimal is regarded as the optimal set in a least-squares sense:

$$\hat{P} = \arg \min_P \sum_{i=1}^{m_p} \|\tilde{x}(u_i|K, P) - x(u_i)\|^2. \quad (6)$$

The minimum is found as the solution of the system $BP = X$, where X is an $m_p \times 3$ matrix having surface points $\{x(u_i)\}$ as its rows, P is an $m_{u^{(0)}}(m_{u^{(1)}} + 2) \times 3$ matrix having the control points $\{p_{ij}\}$ as its rows, and B is an $m_p \times m_{u^{(0)}}(m_{u^{(1)}} + 2)$ matrix. Note that B_{kl} is the contribution of the l th control point P_l to the approximation of the k th surface point X_k . The solution is obtained efficiently from the normal equations $B^T B P = B^T X$ [38]. In Fig. 3, a surface is shown together with four cylindrical b-spline approximations obtained using an increasing number of control points.

4 CORRESPONDENCE QUALITY MEASURE

In this work, the quality of the correspondence of a set of surfaces is measured by the description length of the shape model that was built from the correspondence. The actual computation of the description length can be divided into three steps. First, each surface is represented with a set of corresponding landmarks. Then, the point distribution model (PDM) is calculated from these landmarks. This PDM consists of a mean surface, a set of shape modes, and the variance expressed by each of these modes. Finally, the description length is calculated from the obtained shape mode variances. The following three sections expound these steps.

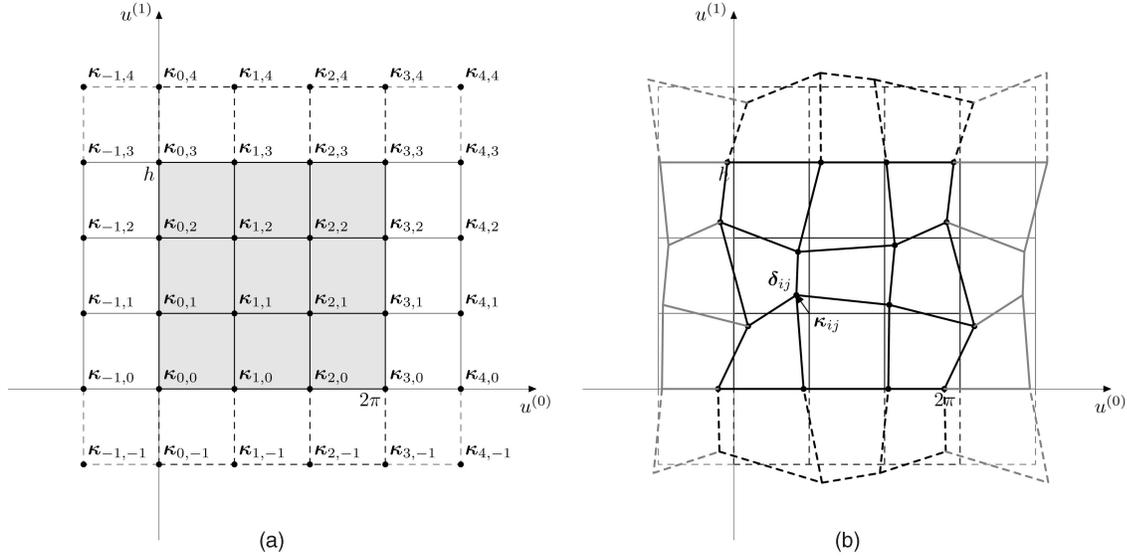


Fig. 4. (a) A 3×4 grid of knots on \mathcal{C}_h^2 . The knots $\kappa_{i,-1}$ and $\kappa_{i,4}$ ensure that the cubic b-spline can be evaluated over the whole domain of \mathcal{C}_h^2 (indicated as the shaded area) and the knots $\kappa_{-1,j}$, $\kappa_{3,j}$, and $\kappa_{4,j}$ make the b-spline periodic in the $u^{(0)}$ -direction. (b) An example grid of b-spline control points overlaid on the knot grid. Together these define a cylindrical parameterization transformation. The shape of the transformation is controlled by the displacements δ_{ij} from the corresponding knots κ_{ij} .

4.1 Statistical Modeling

In order to build a statistical shape model or a point distribution model [1] for a set of n_s surfaces $\{\mathcal{M}_i\}$, a correspondence needs to be established and the surfaces have to be aligned in a common reference coordinate system. Suppose that $\{x_i\}$ are the parameterizations that express this correspondence in the common reference coordinate system. Then, the goal of statistical shape modeling is to capture the shape present in the set of surfaces with a probability density function. Here, a distribution is assumed that is symmetric about its mean, namely, a multivariate Gaussian distribution and the actual distribution parameters are obtained using principal components analysis (PCA) [39].

In this work, the computation of the PCA is done by means of singular value decomposition (SVD). This is more efficient than the traditional method where an eigenvalue decomposition of the large covariance matrix is used. The SVD method also allows the computation of the partial derivatives of the shape mode variances w.r.t. the landmark positions, which is important for the gradient-based optimization. A matrix representation of the set of surfaces $\{x_i\}$ is obtained by sampling each surface at a set of uniformly distributed cylindrical parameter locations $U^{n_p} = \{u_1, \dots, u_{n_p}\}$. For each surface, the coordinates of the n_p landmarks are concatenated and a $3n_p$ row vector \hat{x}_i , representing the surface x_i , is obtained:

$$\hat{x}_i = [x_i(u_1) \dots x_i(u_{n_p})]. \quad (7)$$

The landmark matrix X is then obtained from the n_s shape vectors as $X = [\hat{x}_1^T \dots \hat{x}_{n_s}^T]$, resulting in a matrix of dimensions $3n_p \times n_s$. The mean shape vector \bar{x} is computed as

$$\bar{x} = \frac{1}{n_s} \sum_{i=1}^{n_s} \hat{x}_i, \quad (8)$$

and the row-centered landmark matrix is obtained by subtracting this mean shape from each column of X , i.e.,

$$X_c = [\hat{x}_1^T - \bar{x}^T \dots \hat{x}_{n_s}^T - \bar{x}^T]. \quad (9)$$

Now, let the SVD of the centered landmark matrix be defined as

$$\frac{1}{\sqrt{n_s - 1}} X_c = PSQ^T, \quad (10)$$

where P is a $3n_p \times 3n_p$ orthonormal matrix containing the left singular vectors p_j as its columns, S is a $3n_p \times n_s$ diagonal matrix where the diagonal elements are the singular values σ_j in descending order, and Q is an $n_s \times n_s$ orthonormal matrix with the right singular vectors q_j as its columns. The $n_s \times n_s$ surface covariance matrix D is defined as

$$D = \frac{1}{n_s - 1} X_c^T X_c = QS^2Q^T, \quad (11)$$

and its SVD can be calculated efficiently. The first m columns of P , denoted as P_m , contain the m shape modes p_j and are obtained from the SVD of D as

$$P_m = \frac{1}{\sqrt{n_s - 1}} X_c QS_m^{-1}, \quad (12)$$

where S_m is the matrix that contains the first m rows of S . The first m corresponding shape mode variances λ_j are also obtained from (11) as the first m squared nonzero singular values σ_j^2 .

Using the obtained shape modes p_j and the corresponding mode variances λ_j , a new shape instance \hat{x} can be obtained by adding a linear combination of the principal shape modes to the mean surface:

$$\hat{x} = \bar{x} + \sum_{j=1}^{n_s-1} p_j b_j, \quad (13)$$

where b_j is the contribution of the j th principal shape mode to \hat{x} . Equation (13) defines a shape space spanned by the shape parameters b_j and with the mean shape as the origin. The bounds on the shape parameters of the shape space are usually chosen as a small multiple of the standard deviation of the point cloud along that direction, i.e., $-3\sqrt{\lambda_j} \leq b_j \leq +3\sqrt{\lambda_j}$.

In what follows, Λ will denote the function that maps a set of corresponded surfaces to the mode variances of their derived shape model:

$$(\lambda_1, \dots, \lambda_{n_s-1}) = \Lambda(\mathbf{x}_1, \dots, \mathbf{x}_{n_s} | U^{n_p}). \quad (14)$$

4.2 Description Length

In the work of Davies et al. [18], a correspondence measure for curves and surfaces is introduced which is regarded as the current standard for correspondence optimization. Their measure is adopted here but in a simplified form. The original measure is based on the minimum description length principle: The sampled surfaces are coded in a message where the encoding is determined by the PCA model built from the correspondence. The total message length of the encoded surfaces, together with the encoded model, determines the quality of the model, and therefore, also the quality of the correspondence. In this way, a trade-off is made between model complexity and goodness-of-fit. Over the years, the MDL measure has been tuned and, in this work, the simplified MDL measure, introduced by Thodberg [27], is used. It is a function of the shape mode variances λ_j and is defined as follows:

$$\mu(\lambda_1, \dots, \lambda_{n_s-1}) = \sum_{\lambda_i \geq \lambda_c} \left(1 + \log \frac{\lambda_i}{\lambda_c}\right) + \sum_{\lambda_i < \lambda_c} \frac{\lambda_i}{\lambda_c}, \quad (15)$$

The free parameter λ_c is set to be the expected noise variance in the data. The variation captured by all modes with an eigenvalue (variance) below λ_c is thus considered noise. As can be seen from the first and the second term in (15), respectively, the benefit of decreasing normal modes is logarithmic, while for noise modes, it is constant. Furthermore, the quality measure μ goes to zero when all eigenvalues go to zero, i.e., it favors compact models. Also, both μ and its partial derivatives $\frac{\partial \mu}{\partial \lambda_i}$ are continuous. This is an attractive property for optimization. Note that shorter description length, i.e., a lower value of μ , indicates better quality of correspondence.

4.3 Gradient of Description Length

The L-BFGS minimizer used in this work requires not only the value of the objective μ but also the gradient $\nabla \mu$. In this section, the gradient with respect to the landmark positions x_{ij} is derived. In [40], Ericsson and Åström explain how to obtain the partial derivatives of the description length μ w.r.t. centered the landmark positions x_{ij}^c . Their derivation is based on a result obtained by Papadopoulos and Lourakis in [41]. This result is a simple expression, in function of the elements p_{ik} of P and q_{kj} of Q , for the derivative of the singular values σ_k of X_c w.r.t. the matrix values x_{ij}^c , namely,

$$\frac{\partial \sigma_k}{\partial x_{ij}^c} = p_{ik} q_{kj}. \quad (16)$$

The partial derivatives of the description length w.r.t. the noncentered landmarks x_{ij} can be obtained as:

$$\frac{\partial \mu}{\partial x_{ij}} = \sum_{\lambda_k \geq \lambda_c} \frac{1}{\lambda_k} \frac{\partial \lambda_k}{\partial x_{ij}} + \sum_{\lambda_k < \lambda_c} \frac{1}{\lambda_c} \frac{\partial \lambda_k}{\partial x_{ij}}. \quad (17)$$

The derivatives of the shape mode variances λ_k can be refined into:

$$\frac{\partial \lambda_k}{\partial x_{ij}} = 2\sigma_k p_{ik} q_{kj} - \frac{2}{n_s} \sigma_k p_{ik} \sum_j q_{kj}, \quad (18)$$

where (16) was used, together with the fact that

$$\frac{\partial x_{ij}^c}{\partial x_{ij}} = \begin{cases} 1 - \frac{1}{n_s}, & \text{if } i = i \text{ and } j = j, \\ -\frac{1}{n_s}, & \text{if } i = i \text{ and } j \neq j, \\ 0, & \text{if } i \neq i. \end{cases} \quad (19)$$

Hereby, the gradient of the description length w.r.t. the landmarks is obtained. The gradient with respect to the transformation parameters Φ can be obtained by multiplying the landmark gradient with the Jacobian of the function that maps the transformation parameters to the landmark positions.

5 RIGID CORRESPONDENCE

Establishing an initial correspondence for a set of surfaces $\{\mathcal{M}_1, \dots, \mathcal{M}_{n_s}\}$ starts by parameterizing each surface onto the cylinder. The parameterized piecewise linear surfaces $\{x_1^o, \dots, x_{n_s}^o\}$ are then approximated using b-splines, resulting in the smooth surfaces $\{\tilde{x}_1^o, \dots, \tilde{x}_{n_s}^o\}$. Details on this can be found in Section 3. The approximation uses a grid of $m_{u^{(0)}} \times m_{u^{(1)}}$ control points. It was observed that accurate approximations are achieved with a grid size of the order $32 \times \frac{h}{2\pi} 32$ for all surfaces considered in this work. An initial correspondence, denoted by $\{\hat{x}_1^o, \dots, \hat{x}_{n_s}^o\}$, is obtained from the b-spline surfaces after applying a spatial alignment $\tau^o(\bullet | \hat{\Phi}_i^o)$ and parameter space rotation $\rho^o(\bullet | \hat{\Phi}_i^o)$ to each of the surfaces \tilde{x}_i^o . The optimal transformation parameters $\hat{\Phi}_i^{\tau^o}$ and $\hat{\Phi}_i^{\rho^o}$ are determined by optimization w.r.t. the model description length μ . The objective μ contains multiple local minima, and therefore, suitable initialization for the transformation parameters is required. First, the spatial transformation parameters are initialized by principal axes alignment of the surfaces. This is followed by initialization of the parameterization rotation parameters by running a description length optimization for the parameterization rotations while keeping the spatial transformation parameters fixed. Finally, starting from these initial parameters, the full optimization of the description length, where both spatial transformation parameters and parameterization rotation parameters are set free, is executed. This results in the desired rigid correspondence $\{\hat{x}_1^o, \dots, \hat{x}_{n_s}^o\}$. An overview of establishing a rigid correspondence is given in the flow chart of Fig. 1a.

5.1 Spatial Alignment Initialization

The spatial transformation $\tau^o(\bullet | \Phi^{\tau^o}) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, used for the alignment of the surfaces, is a 3D rigid transformation. It is composed of a 3D rotation around the surface center followed by a 3D translation. The rotation is parameterized by a unit quaternion $\mathbf{q} = (w, q_x, q_y, q_z) \in \mathcal{S}^3$, where \mathcal{S}^3 is the

three-sphere. In order to be of unit length, the quaternion should adhere to the following constraint:

$$\sqrt{qq^T} = 1. \quad (20)$$

Quaternion parameterization for rotation does not suffer from the singularities encountered with Euler angles. The translation is parameterized by a 3D vector $\mathbf{t} = (t_x, t_y, t_z)$. The transformation τ° is thus controlled by seven parameters: $\Phi^{\tau^\circ} = (\mathbf{q}, \mathbf{t})$.

Let

$$\hat{\mathbf{x}}_i^\circ(\bullet | \Phi_i^{\tau^\circ}) = \tau^\circ(\bullet | \Phi_i^{\tau^\circ}) \circ \tilde{\mathbf{x}}_i^\circ \quad (21)$$

be a shorthand notation for the surface obtained after applying the spatial transformation τ° to the surface $\tilde{\mathbf{x}}_i^\circ$, for given parameters $\Phi_i^{\tau^\circ}$. Then, the parameters $\{\Phi_1^{\tau^\circ}, \dots, \Phi_{n_s}^{\tau^\circ}\}$ are chosen so that the surfaces

$$\{\hat{\mathbf{x}}_1^\circ(\bullet | \Phi_1^{\tau^\circ}), \dots, \hat{\mathbf{x}}_{n_s}^\circ(\bullet | \Phi_{n_s}^{\tau^\circ})\}$$

have their principal axes aligned with the axes of the reference coordinate system. The translation vector t_i for surface $\tilde{\mathbf{x}}_i^\circ$ centers the surface at the origin of the coordinate system, i.e.,

$$\mathbf{t}_i = -\bar{\mathbf{v}}^i = -\frac{1}{n_{V_i}} \sum_{j=1}^{n_{V_i}} \mathbf{v}_j^i,$$

where $\bar{\mathbf{v}}^i$ is the average of the vertices \mathbf{v}_j^i of surface $\tilde{\mathbf{x}}_i^\circ$. The rotation for the surface is obtained using singular value decomposition. Let $V_i = [(\mathbf{v}_1^i - \bar{\mathbf{v}}^i)^T \dots (\mathbf{v}_{n_{V_i}}^i - \bar{\mathbf{v}}^i)^T]^T$ be the matrix that has the centered vertices of \mathcal{M}_i as its rows and let the singular value decomposition of the coordinate covariance matrix be defined as $\frac{1}{n_{V_i}-1} V_i^T V_i = U_i S_i^2 U_i^T$. Then, U_i is the rotation matrix that aligns the principal axes of the surface with the reference coordinate axes. Note that the rotation matrix U_i is not uniquely defined since the singular vectors are defined up to their sign. Thus, there are eight possible rotations, from which four can be eliminated since they produce a mirrored surface. From the four remaining rotations, the one is chosen that best matches a reference surface \mathcal{M}_r in terms of the following error:

$$\sum_{j=1}^{n_p} [D(\hat{\mathbf{x}}_i^\circ(\mathbf{u}_j | \Phi_i^{\tau^\circ}), \mathcal{M}_r)]^2, \quad (22)$$

where $D(p, \mathcal{M})$ measures the distance from p to the closest point on \mathcal{M} . The quaternion \mathbf{q}_i that represents the best rotation together with the translation t_i forms the initialization for the spatial transformation parameter set, i.e., $\Phi_i^{\tau^\circ} = (\mathbf{q}_i, t_i)$.

5.2 Parameterization Alignment Initialization

The parameterization transformation $\rho^\circ(\bullet | \Phi^\rho)$, used to align the parameterizations, is a parameter space rotation and it is controlled by a single parameter $\Phi^\rho \in [0, 2\pi]$, defining the angle of rotation. The transformation has the following form:

$$\rho^\circ(\bullet | \Phi^\rho) : \mathcal{C}_h^2 \mapsto \mathcal{C}_h^2, \quad (23)$$

$$: \mathbf{u} \mapsto (\mathbf{u}^{(0)} + \Phi^\rho, \mathbf{u}^{(1)}). \quad (24)$$

Similar to (21), a shorthand notation for the surface obtained after applying the spatial and parameterization transformation for given parameters $\Phi_i^{\rho^\circ}$ and $\Phi_i^{\tau^\circ}$ is as follows:

$$\hat{\mathbf{x}}_i^\circ(\bullet | \Phi_i^{\rho^\circ}, \Phi_i^{\tau^\circ}) = \tau^\circ(\bullet | \Phi_i^{\tau^\circ}) \circ \tilde{\mathbf{x}}_i^\circ \circ \rho^\circ(\bullet | \Phi_i^{\rho^\circ}). \quad (25)$$

The initial values for the parameterization rotation parameters $\{\Phi_1^{\rho^\circ}, \dots, \Phi_{n_s}^{\rho^\circ}\}$ are then obtained by solving the following optimization problem:

$$\arg \min_{\Phi_i^{\rho^\circ}, \forall i} \mu \circ \Lambda(\dots, \hat{\mathbf{x}}_i^\circ(\bullet | \Phi_i^{\rho^\circ}, \Phi_i^{\tau^\circ}), \dots | U^{m_\circ}), \quad (26)$$

where the $\Phi_i^{\tau^\circ}$ are the initial spatial alignment parameters obtained in the previous section.

5.3 Full Alignment

Now, starting from the initial spatial transformation parameters $\{\Phi_1^{\tau^\circ}, \dots, \Phi_{n_s}^{\tau^\circ}\}$ and the initial parameterization rotation parameters $\{\Phi_1^{\rho^\circ}, \dots, \Phi_{n_s}^{\rho^\circ}\}$, obtained in the previous two sections, the full description length minimization with rigid transformations can be solved:

$$\arg \min_{\Phi_i^{\rho^\circ}, \Phi_i^{\tau^\circ}, \forall i} \mu \circ \Lambda(\dots, \hat{\mathbf{x}}_i^\circ(\bullet | \Phi_i^{\rho^\circ}, \Phi_i^{\tau^\circ}), \dots | U^{m_\circ}) + \frac{\alpha^\tau}{n_s} \sum_j \eta^\tau(\Phi_j^{\tau^\circ}). \quad (27)$$

Here, the regularization term $\eta^\tau(\Phi_j^{\tau^\circ})$ penalizes parameter sets with a quaternion that violates (20). The penalty for a quaternion \mathbf{q} is measured as $(1 - \sqrt{qq^T})^2$. The regularization term attains its minimum when all quaternions are in S^3 and smoothly penalizes any deviation from this. In all of the experiments, a regularization factor $\alpha^\tau = 10^6$ was used. The parameters that solve (27) are denoted $\{\hat{\Phi}_1^{\tau^\circ}, \hat{\Phi}_1^{\rho^\circ}, \dots, \hat{\Phi}_{n_s}^{\tau^\circ}, \hat{\Phi}_{n_s}^{\rho^\circ}\}$ and they provide the final rigid correspondence $\{\hat{\mathbf{x}}_1^\circ(\bullet | \hat{\Phi}_1^{\rho^\circ}, \hat{\Phi}_1^{\tau^\circ}), \dots, \hat{\mathbf{x}}_{n_s}^\circ(\bullet | \hat{\Phi}_{n_s}^{\rho^\circ}, \hat{\Phi}_{n_s}^{\tau^\circ})\}$.

6 NONRIGID CORRESPONDENCE

In the previous section, a rigid correspondence $\{x_1, \dots, x_{n_s}\}$ was produced by applying the optimal rigid spatial transformations and rigid parameterization transformations to the parameterized surfaces. An improvement over this rigid correspondence can be obtained by allowing local deformations of the parameterizations. Such local deformations can be realized by a nonrigid parameterization transformation. Here, a cylindrical b-spline parameterization deformation ρ^L is used, where L indicates the level of resolution of the deformation. The spatial transformation τ is the same as in Section 5. The optimal nonrigid parameterization deformations, together with the spatial transformation parameters, are determined by optimization. The optimization is done at increasing levels of resolution, sequentially, to avoid convergence to a local optimum. The flow chart in Fig. 1b gives an overview of the multiresolution correspondence optimization.

At every resolution level L , the optimal transformation parameters for that level are determined in the following manner: First, each surface x_i , obtained from the rigid correspondence procedure, is approximated with a b-spline surface, denoted as \tilde{x}_i^L . The approximation uses

an $m_{u^{(0)}}^L \times m_{u^{(1)}}^L$ grid of control points. Here, $m_{u^{(0)}}^L = 3 \cdot 2^{L-1}$ and $m_{u^{(1)}}^L = \lfloor \frac{h}{2\pi} m_{u^{(0)}}^L \rfloor$ are used; thus, the resolution is approximately isotropic and doubled from one level to the next. Each surface \tilde{x}_i^L is transformed according to the spatial transformation parameters $\Phi_i^{\tau^L}$ and parameterization transformation parameters $\Phi_i^{\rho^L}$ as follows:

$$\hat{x}_i^L(\bullet | \Phi_i^{\rho^L}, \Phi_i^{\tau^L}) = \tau(\bullet | \Phi_i^{\tau^L}) \circ \tilde{x}_i^L \circ \rho^L(\bullet | \Phi_i^{\rho^L}). \quad (28)$$

The level- L b-spline parameterization transformation ρ^L is defined by a grid of $n_{u^{(0)}}^L \times n_{u^{(1)}}^L$ control points. Similar to the approximation b-spline, the transformation b-spline has an isotropic resolution that is doubled at each new level, i.e., $n_{u^{(0)}}^L = 3 \cdot 2^{L-1}$ and $n_{u^{(1)}}^L = \lfloor \frac{h}{2\pi} n_{u^{(0)}}^L \rfloor$. Using the notation from (28), the optimal level- L transformation parameters are determined by solving the following optimization problem:

$$\begin{aligned} \arg \min_{\Phi_i^{\rho^L}, \Phi_i^{\tau^L}, \forall i} \mu \circ \Lambda(\dots, \hat{x}_i^L(\bullet | \Phi_i^{\rho^L}, \Phi_i^{\tau^L}), \dots | U^{n_p^L}) \\ + \frac{\alpha^\tau}{n_s} \sum_j \eta^\tau(\Phi_j^{\tau^L}) + \frac{\alpha^\rho}{n_s} \sum_j \eta^\rho(\Phi_j^{\rho^L}). \end{aligned} \quad (29)$$

Here, η^τ is the regularization for the spatial transformations, as defined in Section 5.3, and η^ρ is the regularization for the parameterization deformations in order to avoid overfitting. The experimentally determined regularization constants are $\alpha^\tau = 10^6$ and $\alpha^\rho = 0.2$. The set of parameter coordinates $U^{n_p^L}$ used to estimate the shape covariance matrix contains $n_p^L = 250 \cdot 4^{L-1}$ parameter locations. Thus, the number of samples per area increases fourfold with every new level. This mirrors the doubling of the resolution of the surface approximation and parameterization transformation. Note that the optimal transformation parameters for the optimization problem of level $L-1$ are used to initialize the parameters of current level L . To initialize the b-spline parameterization transformation parameters $\Phi_j^{\rho^L}$ from $\Phi_j^{\rho^{L-1}}$ the b-spline upsampling technique from [42] is used. The initialization of the spatial transformations is trivial, i.e., $\Phi_j^{\tau^L} = \Phi_j^{\tau^{L-1}}$. In this work, three levels of resolution were used, and thus, the optimal transformation parameters $\{\hat{\Phi}_1^{\tau^3}, \hat{\Phi}_1^{\rho^3}, \dots, \hat{\Phi}_{n_s}^{\tau^3}, \hat{\Phi}_{n_s}^{\rho^3}\}$ of the third-level optimization problem are the final transformation parameters. These provide the final correspondence $\{\hat{x}_1, \dots, \hat{x}_{n_s}\}$.

In the following two sections, the actual form of the parameterization transformation will be detailed and a suitable regularizer is introduced.

6.1 Reparameterization Transformation

The parameterization space \mathcal{C}_h^2 is deformed using a parameterization transformation $\rho^L(u | \Phi^{\rho^L})$, where L denotes the level of resolution. The transformation is an automorphism of the parameter space, i.e., ρ^L constitutes a continuous one-to-one map of \mathcal{C}_h^2 . The space of possible reparameterizations is spanned by the transformation parameters Φ^{ρ^L} . Different bases can be used to represent a reparameterization function [18], [28], [29], [43]. In this work, ρ^L is a 2D cubic b-spline function with knot positions on a regular grid. Such a representation has a number of convenient properties: 1) Cubic b-spline deformations are C^2 continuous with respect to their parameters within the patches and C^1 continuous at the patch boundaries. This is required for

efficient, gradient-guided optimization. 2) It has compact support which makes it fast to evaluate and allows local control. And 3) it can be used in a multiresolution method by refining the grid that controls the shape of the deformation.

The b-spline deformation function ρ^L is defined by a set of knots $K^L = \{\kappa_{ij}^L\}$ and a set of control point displacements $\Phi^{\rho^L} = \{\delta_{ij}^L\}$:

$$\begin{aligned} \rho^L(u | \Phi^{\rho^L}) \\ = \sum_{i=-1}^{n_{u^{(0)}}^L+1} \sum_{j=-1}^{n_{u^{(1)}}^L} \beta\left(\frac{u - \kappa_{ij}^L}{\Delta^L}\right) (\kappa_{ij}^L + \delta_{ij}^L) \pmod{(0)}{2\pi}, \end{aligned} \quad (30)$$

where the modulo operator, $\pmod{(0)}$, acts on the first coordinate of the parameter space and keeps the deformed parameter within the bounds $[0, 2\pi]$. β is the 2D separable cubic b-spline kernel from (3). The knots and the corresponding control points are arranged on a regular $n_{u^{(0)}}^L \times n_{u^{(1)}}^L$ grid on the cylinder \mathcal{C}_h^2 . The spacing between the knots is denoted as

$$\Delta^L = \left(\frac{2\pi}{n_{u^{(0)}}^L}, \frac{h}{n_{u^{(1)}}^L - 1} \right).$$

The following constraints on the control point displacements δ_{ij} make sure that ρ^L is periodic and continuous at the parameter boundary $u^{(0)} = 2\pi$:

$$\begin{cases} \delta_{-1,j} = \delta_{n_{u^{(0)}}-1,j}, \\ \delta_{0,j} = \delta_{n_{u^{(0)}},j}, \\ \delta_{1,j} = \delta_{n_{u^{(0)}}+1,j}, \end{cases} \quad \forall j. \quad (31)$$

In order to make ρ^L one-to-one along the boundaries of \mathcal{C}_h^2 , the following constraints are also enforced:

$$\begin{cases} \delta_{i,-1}^{(1)} = -\delta_{i,1}^{(1)}, \\ \delta_{i,n_{u^{(0)}}}^{(1)} = -\delta_{i,n_{u^{(0)}}-2}^{(1)}, \\ \delta_{i,0}^{(1)} = 0, \\ \delta_{i,n_{u^{(0)}}-1}^{(1)} = 0, \end{cases} \quad \forall i. \quad (32)$$

In this work, the same arrangement of knots K^L is used for the reparameterization of each surface x_k . Only the control point displacements $\Phi_k^{\rho^L} = \{\delta_{ij,k}^L\}$ differ from surface to surface. Fig. 4a shows a 3×4 cylindrical grid of knots in an unfolded view and Fig. 4b shows a grid of control points for that knot grid, where the control points are obtained by adding the control point displacement to the knot location, i.e., $\kappa_{ij}^L + \delta_{ij}^L$. Observe that both the knots and the control points reside on the parameter domain \mathcal{C}_h^2 . The b-spline function ρ^L thus defines a reparameterization function of the cylindrical parameter domain \mathcal{C}_h^2 .

6.2 Reparameterization Regularization

From (30), it can be seen that the reparameterization transformation ρ^L is a linear combination of translated versions of the cubic b-spline kernel β . As a result, the reparameterization transformation is C^2 continuous within the patches and C^1 continuous at the patch boundaries. The inherent smoothness of the b-spline transformation is convenient but it does not avoid overfitting by the b-spline transform. Overfitting is perceived as an irregular local deformation and occurs in regions where μ is insensitive to

local deformations, e.g., when a b-spline kernel is not supported by any landmark samples in the calculation of Λ .

To counter these irregularities, a regularization term is introduced, denoted as η^ρ . For b-splines, a number of regularization terms have been used in the past [44], [45]. Here, a simple regularizer is chosen that measures the Dirichlet energy of the parameterization displacement function. This displacement function, denoted by ϱ , can be easily derived from (30):

$$\varrho(\mathbf{u} \mid \Phi^\rho) = \sum_{i=-1}^{n_{u(0)}+1} \sum_{j=-1}^{n_{u(1)}} \beta \left(\frac{\mathbf{u} - \kappa_{ij}}{\Delta} \right) \delta_{ij}. \quad (33)$$

The regularization term η^ρ for the deformation is then defined as

$$\eta(\Phi^\rho) = \int \int_{C_h^2} \left(\left| \frac{\partial \varrho(\mathbf{u} \mid \Phi^\rho)}{\partial u^{(0)}} \right|^2 + \left| \frac{\partial \varrho(\mathbf{u} \mid \Phi^\rho)}{\partial u^{(1)}} \right|^2 \right) d\mathbf{u}. \quad (34)$$

In what follows, $\beta_{u^{(c)}}^i = \beta(\frac{u^{(c)}}{\Delta} - i)$ is used as a shorthand notation and the prime mark symbol denotes the derivative. Substituting the b-spline transformation into (34) results in:

$$\begin{aligned} \eta(\Phi^\rho) &= \sum_{ij} \sum_{kl} \delta_{ij} \delta_{kl}^\tau \left(\frac{1}{\Delta^{(0)2}} \int_0^{2\pi} \beta_{u^{(0)}}^{i'} \beta_{u^{(0)}}^{k'} du^{(0)} \int_0^h \beta_{u^{(1)}}^j \beta_{u^{(1)}}^l du^{(1)} \right. \\ &\quad \left. + \frac{1}{\Delta^{(1)2}} \int_0^{2\pi} \beta_{u^{(0)}}^i \beta_{u^{(0)}}^k du^{(0)} \int_0^h \beta_{u^{(1)}}^{j'} \beta_{u^{(1)}}^{l'} du^{(1)} \right). \end{aligned} \quad (35)$$

From this, the derivatives of the smoothness energy w.r.t. the control point displacements are easily derived:

$$\begin{aligned} \frac{\partial \eta}{\partial \delta_{kl}^{(0)}} &= 2 \sum_{ij} \delta_{ij}^{(0)} \left(\frac{1}{\Delta^{(0)2}} \int_0^{2\pi} \beta_{u^{(0)}}^{i'} \beta_{u^{(0)}}^{k'} du^{(0)} \int_0^h \beta_{u^{(1)}}^j \beta_{u^{(1)}}^l du^{(1)} \right. \\ &\quad \left. + \frac{1}{\Delta^{(1)2}} \int_0^{2\pi} \beta_{u^{(0)}}^i \beta_{u^{(0)}}^k du^{(0)} \int_0^h \beta_{u^{(1)}}^{j'} \beta_{u^{(1)}}^{l'} du^{(1)} \right). \end{aligned} \quad (36)$$

Equations (35) and (36) can be evaluated analytically, and as a result, they can be used efficiently with a gradient-based optimization method.

7 RESULTS AND DISCUSSION

7.1 Data Sets

In this work, six populations of surfaces were used to test the proposed correspondence method. Three of these are phantom populations, each consisting of 30 surfaces: a population of disks where the position of the disk on the cylinder is variable, a population of beams with varying width and depth, and a population of cylinder-like bent surfaces with an elliptic cross section where the width and height of the cross section are variable together with the amount of bending. A sample surface of each of these populations is shown in Fig. 5a. There are also three populations of real, ct-scanned surfaces: a population of 25 clavicles, a population of 23 tracheas, and a population of

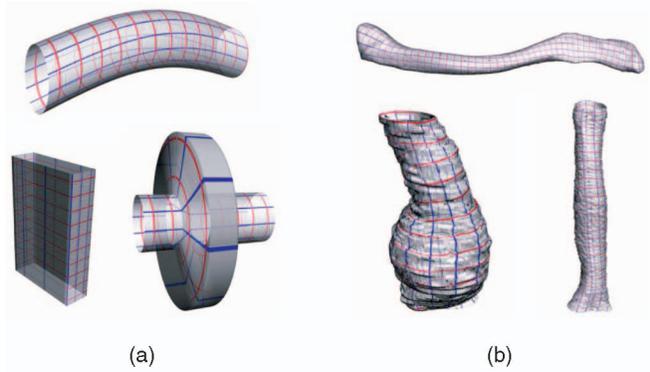


Fig. 5. From each of the six considered surface populations, one random surface is shown here. Each surface is textured with the isocontours of its parameter coordinates obtained after parameterization. (a) Three populations of synthetic surfaces in clockwise order are cylinders, disks, and beams. (b) Three populations of real surfaces in clockwise order are clavicles, tracheas, and thrombi.

50 aortic sections with a thrombus. A sample surface of each of the real populations can be found in Fig. 5b. All surfaces of these populations have cylindrical topology except the clavicles, which are of spherical topology. The clavicles are made cylindrical by puncturing both ends of the clavicle.

7.2 Performance Measures

In what follows, different correspondences will be constructed for the aforementioned populations. The quality of the established correspondences is evaluated by deriving a PCA model from the correspondence and reporting performance measures for the obtained model. The performance of a model is measured here by the compactness, reconstruction ability, generalization ability, and specificity of the model. The performance measures are reported for the full k -mode model and all restricted m -mode, $m < k$, versions of the PCA model. In a comparison, the correspondence having the best performance measures, for its derived model, is considered the best correspondence.

Now, given a set of surfaces $\{\mathcal{M}_1, \dots, \mathcal{M}_{n_s}\}$, let the correspondence be denoted as $\{x_1, \dots, x_{n_s}\}$ and the derived shape model as \hat{x}^m , where m is the number of modes of the shape model. Then, the *compactness* of a model is measured as the cumulative variance:

$$C(m) = \sum_{i=1}^m \lambda_i, \quad (37)$$

where λ_i is the variance of the i th shape mode. The *reconstruction ability* indicates how good a model is able to reconstruct the surfaces that were used to build the model. It is measured as the average approximation error after fitting the model to each of the surfaces $\{\mathcal{M}_1, \dots, \mathcal{M}_{n_s}\}$:

$$R(m) = \frac{1}{n_s} \sum_{i=1}^{n_s} \min_{\Phi^\tau, \mathbf{b}} D(\tau(\Phi^\tau) \circ \hat{x}^m(\mathbf{b}), \mathcal{M}_i), \quad (38)$$

where Φ^τ are the parameters of the rigid transformation τ , \mathbf{b} are the model parameters, and $D(x, y)$ measures the average closest point distance from surface y to surface x . The optimal parameters Φ^τ and \mathbf{b} , resulting in the best model-to-surface fit, are determined iteratively by alternately estimating Φ^τ and \mathbf{b} in a least-squares sense. The *generalization ability* of a

TABLE 1
Ideal, Rigid, and Nonrigid Correspondence Results

population	n_s	h	MDL			time (min.)
			ideal	rigid	non-rigid	
disks	30	2π	6.2	17.2	4.6	66
beams	30	2π	9.6	9.8	9.5	19
cylinders	30	2π	13.6	11.7	10.9	82
clavicles	25	8π		61.9	41.5	157
trachea	23	4π		47.2	33.3	385
thrombi	50	2π		39.1	28.0	411

model determines how well the model generalizes to unseen instances of the modeled class. It is measured as the average approximation error after fitting leave-one-out versions of the model to the left out surfaces:

$$G(m) = \frac{1}{n_s} \sum_{i=1}^{n_s} \min_{\Phi^T, b} D(\tau(\Phi^T) \circ \hat{x}_i^m(b), \mathcal{M}_i), \quad (39)$$

where \hat{x}_i^m is the m -mode model where the i th surface was left out, i.e., it is built from the corresponded surfaces $\{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{n_s}\}$. The model *specificity* measures how much random samples, generated by the model resemble the original surfaces:

$$S(m) = \frac{1}{n_t} \sum_{i=1}^{n_t} \min_{j, \Phi^T} D(\tau(\Phi^T) \circ \hat{x}^m(b_i^m), \mathcal{M}_j), \quad (40)$$

where the b_i^m are the random Gaussian model parameters for the sample of the i th trial and n_t is the number of random

samples used to estimate the specificity. Note that the model performance measures from [18] result in a bias toward the MDL-optimized models. The model performance measures in (38), (39), and (40) do not suffer from this drawback. The compactness measure from (37) is closely related to the MDL-measure, and therefore, biased. It is, however, reported here because it contains important information of how a model captures the variation of a population.

7.3 Rigid versus Nonrigid Correspondence

The surfaces of each of the six phantom and real populations were parameterized using the progressive parameterization technique of [14]. The chosen height h of the cylinder for each of the populations can be found in Table 1. From the parameterized surfaces, a $16 \times 16 \frac{h}{2\pi}$ b-spline representation was computed for each surface using $m_p = 10,000$ points. This was followed by the construction of the rigid correspondence. The number of landmarks to estimate the covariance was $n_p = 4,000$ for all populations. Starting from the rigid correspondence, the nonrigid correspondence was calculated. In this construction, a b-spline regularization factor of $\alpha^p = 0.2$ was used for all populations. Three levels of scale were used for the b-spline surface, the reparameterization transformation, and the number of landmarks. On the coarsest scale, a 4×4 b-spline surface, together with a 4×4 reparameterization transformation, was the covariance matrix that was estimated based on 250 landmarks. At each new resolution label, the values are increased as detailed in Section 6. All optimizations problems were solved with the L-BFGS routine [46]. For the alignment initialization (26) and full

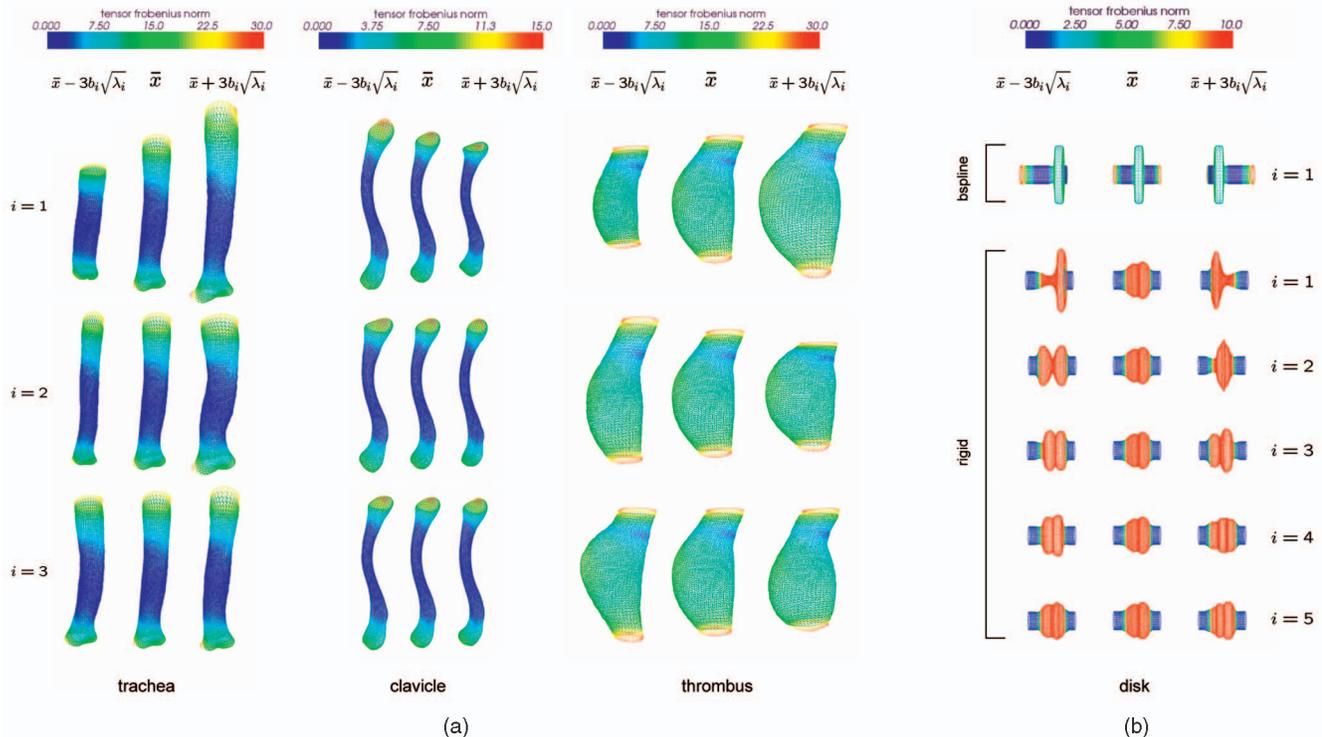


Fig. 6. A visualization of the modes of models obtained with the approach of this paper. The color coding represents the frobenius norm of the landmark covariance matrix and is a measure of the local variability of the surface. On the left, the first three modes are shown for the three CT-scanned populations: trachea, clavicle, and thrombus. For each mode i , the average shape \bar{x} is shown together with the positive and negative offset in the direction of mode i : $\bar{x} \pm 3b_i\sqrt{\lambda_i}$. On the right, the modes of model for the disk population are shown. On the top, the b-spline-optimized model is shown. At the bottom, the model of the rigid correspondence is shown. Clearly, the b-spline optimization dramatically improved the correspondence.

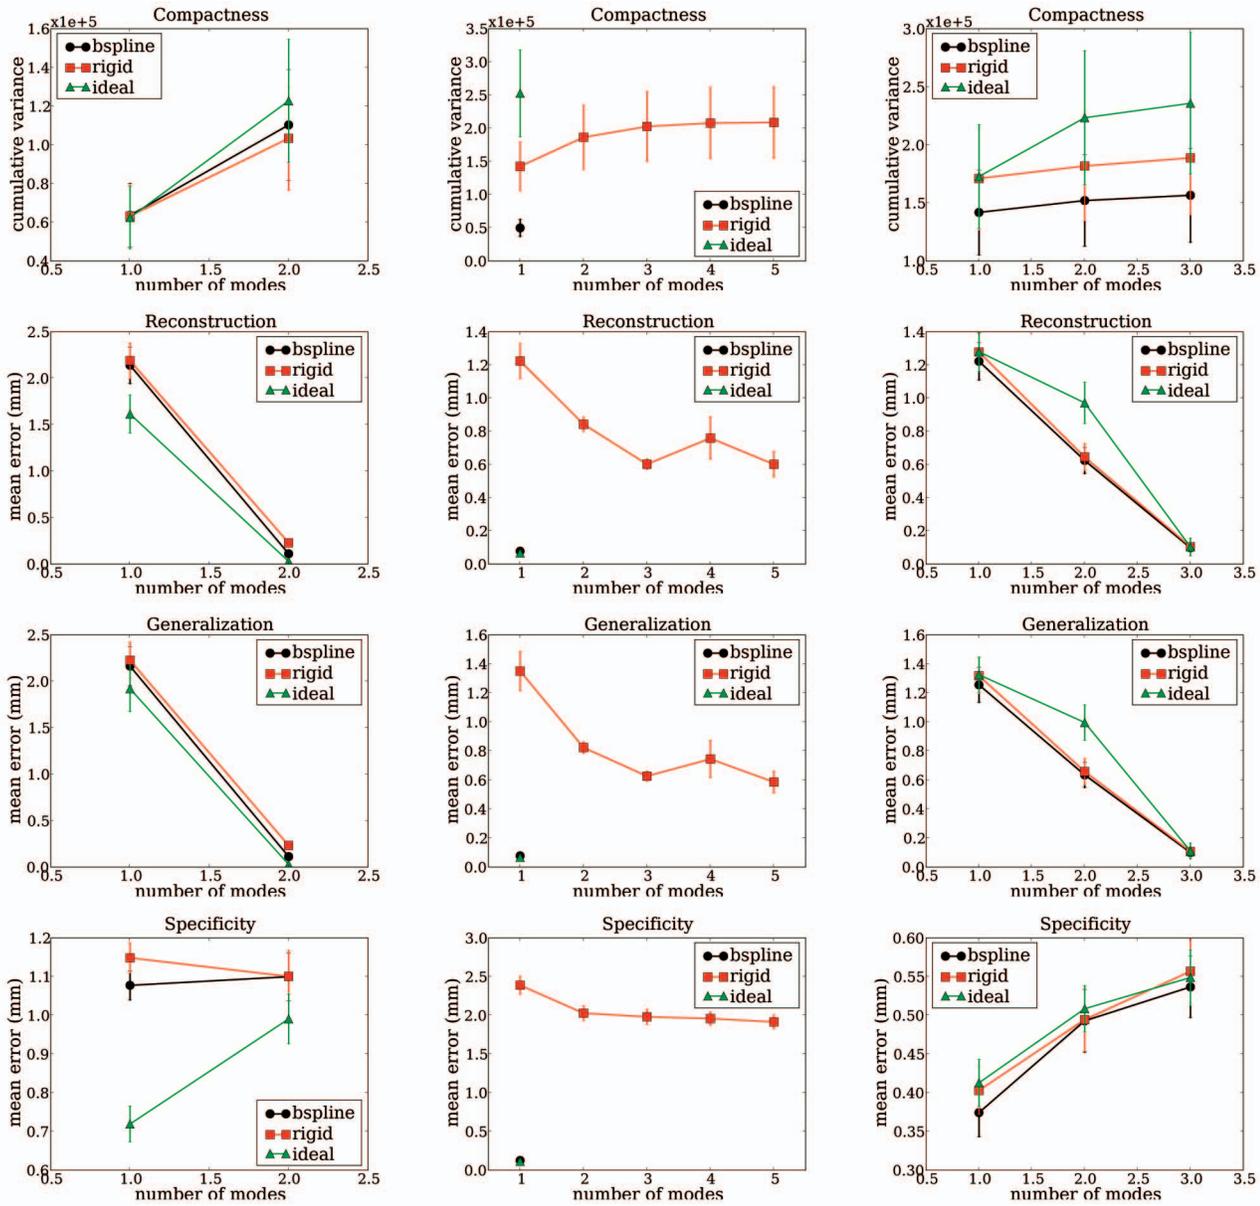


Fig. 7. The model performance measures, together with the standard error, for the phantom populations. From left to right: beams, disks, and cylinders populations. From top to bottom: compactness, reconstruction, generalization, and specificity measure. In each graph, a comparison is made of the ideal correspondence (green), the rigid correspondence (red), and the nonrigid correspondence (black). For each model, the first m modes are shown that capture 99 percent of the total variance in the model.

alignment (27), a gradient tolerance of 0.01 was used to determine convergence. For the multiresolution correspondence optimization (29), a gradient tolerance of 0.01×2^{-L} was used. The quality of the obtained correspondences was assessed using the aforementioned model performance measures. The results are shown in Fig. 7 for the phantom populations and in Fig. 8 for the real populations.

In Fig. 7, the ideal (intuitive), the rigid, and the nonrigid correspondences for the phantom populations are compared. For the disk population, it can be seen that the nonrigid and the ideal correspondence are of comparable quality. The rigid correspondence, on the other hand, is much worse. This is due to the fact that the parameterization technique maps the disk part of the surfaces to a different location in the parameter space. The nonrigid correspondence improvement, on the other hand, moves

the disks to the same part of the parameter space and this results in an optimal correspondence. See Fig. 6 for a visualization of the rigid and b-spline-optimized model. The improved compactness of the nonrigid over the ideal correspondence can be attributed to the reduced area that the disk part of the surface takes in the parameter space. For the beam population, the quality of the correspondences is comparable, with a slight advantage for the ideal, since the parameterization technique already generates a good correspondence. The most notable difference is the improved specificity of the ideal model, which is due to the fact that the other models can generate samples with rounded corners. For the cylinders population, it can be seen that the nonrigid correspondence is an improvement over the ideal, which is mainly due to the improved spatial alignment of the surfaces.

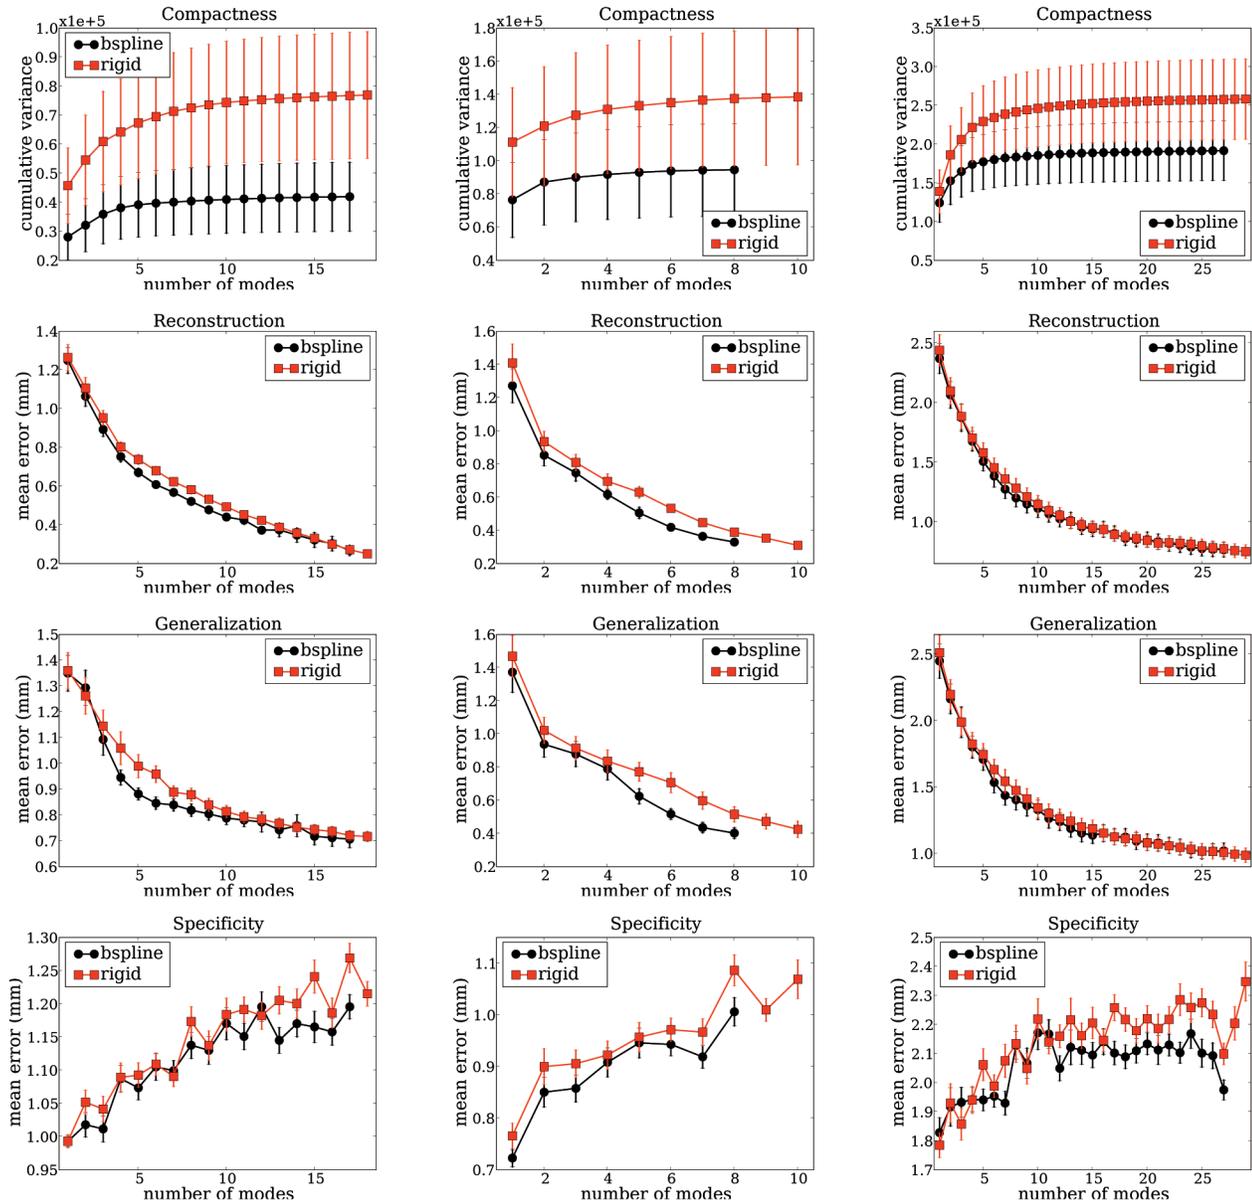


Fig. 8. Different model performance measures, together with the standard error, for the real data sets. From left to right: clavicles, tracheas, and thrombi populations. From top to bottom: compactness, reconstruction, generalization, and specificity measure. In each graph, the rigid correspondence (red) is compared to the nonrigid correspondence (black). For each model, the first m modes are shown that capture 99 percent of the total variance in the model.

In Fig. 6, a visualization of the first few modes of the models of the CT-scanned populations can be found. In Fig. 8, the model performance measures for the rigid and the nonrigid correspondence are shown for the CT-scanned populations. It is more difficult to analyze these results because there is no ideal correspondence available. It can be seen though that the rigid correspondence generates good models for all three populations and that the nonrigid correspondence is a significant improvement in most cases. Compared to the clavicle and trachea populations, the approximation errors for the thrombi population errors are higher. This can be attributed to the large variability that is present in the thrombi population, together with the lower resolution of the ct-scans, that is $0.5 \times 0.5 \times 0.5 \text{ mm}^3$ versus $0.5 \times 0.5 \times 2.0 \text{ mm}^3$.

In Table 1, the MDL values for the ideal (when available), the rigid, and the nonrigid correspondences can be found for all populations. It can be seen that the b-spline correspondence optimization always succeeds in decreasing the MDL-value. Such a decrease indicates that the resulting shape model became less complex. This is most apparent for the disk population: The five modes of the rigid model (MDL-value of 17.2) are reduced to a single mode by the b-spline optimization (MDL-value of 4.6). Table 1 also reports the execution time for the construction of the correspondences once the parameterizations are obtained. About 10 percent of the time is taken by the rigid correspondence construction. The correspondences for the phantom populations are constructed more efficiently compared with the CT-scanned populations. This is because the phantom shape models are relatively simple, i.e., they

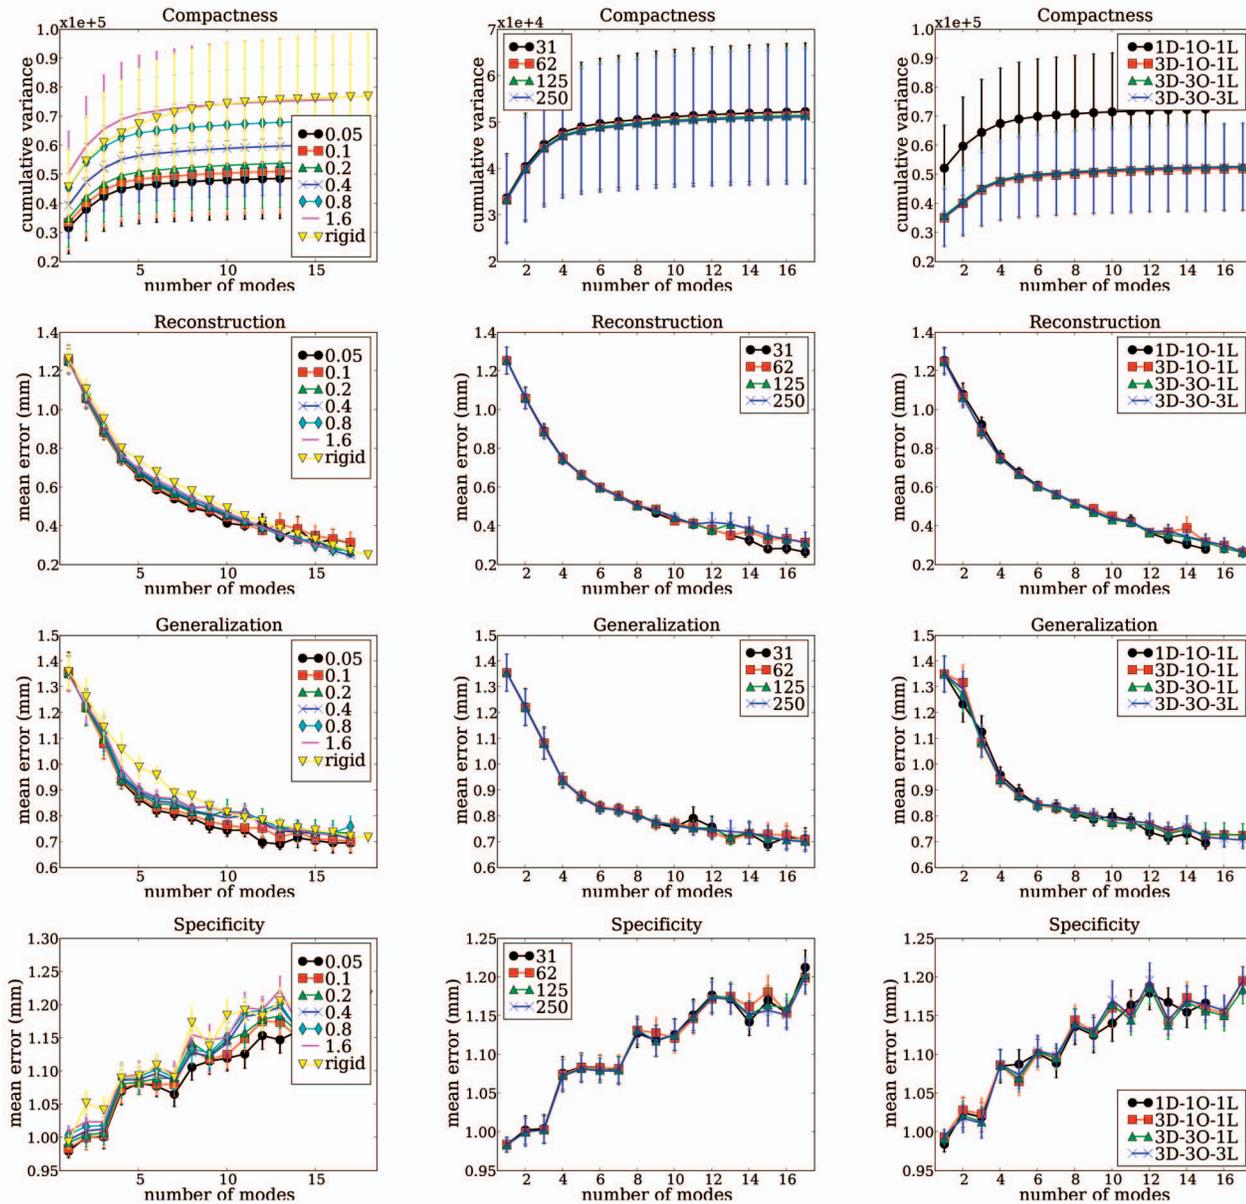


Fig. 9. Influence of the method parameters on the model performance for the clavicle data set. Left: influence of the b-spline deformation regularization factor α^ρ . Seven models with regularizations from $\alpha^\rho = 0.05$ to $\alpha^\rho = 1.6$ are compared. The rigid case, i.e., $\alpha^\rho = \infty$, is provided as a reference. It can be observed that models with less regularization tend to perform better. However, too low values for α^ρ may result in irregularities. Middle: influence of the number of landmarks n_p used in the estimation of the shape covariance matrix. Four models are shown with number of landmarks from $n_p^L = 31 \cdot 4^{L-1}$ to $n_p^L = 250 \cdot 4^{L-1}$. It can be seen that the performance is insensitive to this parameter. However, for certain populations, using too few landmarks may result in degraded models due to undersampling of highly variable surface regions. Right: influence of the multiresolution scheme. The scheme encoded as $aD-bO-cL$ uses a levels for the b-spline deformation, b levels for the surface approximation, and c levels for the number of landmarks. It can be observed that the single resolution scheme 1D-1O-1L generates a degraded model and that the other schemes result in comparable models. The full scheme 3D-3O-3L is preferred because it is computationally less expensive.

have a small number of modes. The construction for the CT-scanned populations takes a couple of hours. Together with the construction of the parameterizations, a correspondence can easily be established overnight.

7.4 Influence of Parameters

In this section, the influence of the method parameters on the resulting correspondence is investigated. Fig. 9 shows the model performance parameters for the clavicle population when using different values for the most important parameters: 1) the b-spline deformation regularization controlled by factor α^ρ in (29); 2) the number of landmarks

to estimate the shape covariance matrix, controlled by n_p in (14); and 3) the number of scale levels L for the b-spline surface approximation, the b-spline parameterization transform, and the landmarks.

In the first column of Fig. 9, the influence of the b-spline regularization factor α^ρ is shown and it can be seen that, as expected, lower regularization, i.e., smaller α^ρ , generates better models. However, the regularization cannot be lowered too much since then irregularities can appear. No irregularities were noted with regularization factors $\alpha^\rho \leq 0.2$ for all considered populations.

TABLE 2
Influence of the Multilevel Scheme for the Clavicle Population

number of levels			MDL	time(h)
reparameterization	surface	landmarks		
1	1	1	42.6	4.5
3	1	1	41.3	7.5
3	3	1	41.3	7.5
3	3	3	41.5	2.5

In the second column of Fig. 9, it can be seen that the influence of the number of landmarks n_p is negligible for the clavicle population. It can, however, happen that using too few landmarks results in undersampling of highly variable parts of the surface, which, in turn, will result in a degraded correspondence. This was observed for the thrombi population, where $n_p^L = 250 \cdot 4^{L-1}$ landmarks generated a significantly better correspondence than $n_p^L = 31 \cdot 4^{L-1}$ landmarks (result not shown). Using fewer landmarks reduces computation time, but it can also result in a degraded correspondences.

In the last column of Fig. 9, the model performance measures for different multiscale schemes are shown, and in Table 2, the corresponding MDL values are listed. Four different schemes are shown: 1D-1O-1L optimizes the correspondence using a single resolution, 3D-1O-1L uses three resolution levels for the deformations, 3D-3O-1L uses three levels for the deformations and the surfaces, and 3D-3O-3L is the full multiresolution scheme using three levels for the deformations, the surfaces, and the landmarks. It can be seen that, for the clavicle, the single-scale scheme generates a degraded correspondence and the three other schemes generate comparable correspondences. However, the full three-scale scheme is considerably faster. For the disk population, the performance degradation of the single-scale method was even more notable since, as opposed to the three other schemes, it did not successfully correspond the disk parts of the surface (result not shown).

8 CONCLUSIONS

In this work, the minimum description length approach for shape modeling was translated to surfaces of cylindrical topology. The proposed method establishes an alignment and a correspondence for a population of surfaces of cylindrical topology. It generates a rigid correspondence based on cylindrical surface parameterizations and an improved correspondence using multilevel b-spline reparameterizations. Care was taken to ensure that the objective functions are differentiable with respect to the alignment and reparameterization parameters and, where necessary, an expression for the gradient was provided. It was shown that the method produces correspondences that agree with the intuitive correspondence and that the derived shape models generate small approximation errors.

The cylindrical correspondence method of this paper, together with the spherical correspondence methods from [28] and [29], and the disc-like correspondence method of [30] already cover a wide range of biomedical surfaces. However, it would be interesting to extend the method to other topologies. The method of this paper can be trivially extended to surfaces of genus-1 topology, where the torus can

be used as the parametric domain. More complex topologies can be treated by decomposing the surfaces in a consistent set of discs and tubes. Populations of tubular structures with bifurcations could be handled well with this approach. However, arbitrary complex surfaces will suffer from the boundary constraints imposed by the surface decomposition. The development of a method that can handle arbitrary topology, without constraints, is a very challenging problem to be solved in the future.

ACKNOWLEDGMENTS

This work was financially supported by the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) and the Fund for Scientific Research (FWO)-Flanders, Belgium. The thrombus data sets were kindly provided by Marleen de Bruijne and the Department of Vascular Surgery, University Medical Center Utrecht.

REFERENCES

- [1] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham, "Active Shape Models: Their Training and Application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38-59, 1995.
- [2] I.L. Dryden and K. Mardia, *Statistical Shape Analysis*. John Wiley Sons, 1998.
- [3] M. Styner, J.A. Lieberman, R.K. McClure, D.R. Weinberger, D.W. Jones, and G. Gerig, "Morphometric Analysis of Lateral Ventricles in Schizophrenia and Healthy Controls Regarding Genetic and Disease-Specific Factors," *Proc. Nat'l Academy of Sciences USA*, vol. 102, no. 13, pp. 4872-4877, Mar. 2005.
- [4] R.R. Paulsen, R. Larsen, S. Laugesen, C. Nielsen, and B.K. Ersbøll, "Building and Testing a Statistical Shape Model of the Human Ear Canal," *Proc. Fifth Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, 2002.
- [5] S. Zachow, H. Lamecker, B. Elsholtz, and M. Stiller, "Reconstruction of Mandibular Dysplasia Using a Statistical 3D Shape Model," *Proc. Computer Assisted Radiology and Surgery*, H. Lemke et al., eds., pp. 1238-1243, 2005.
- [6] A. Kelemen, G. Szekely, and G. Gerig, "Elastic Model-Based Segmentation of 3D Neuroradiological Data Sets," *IEEE Trans. Medical Imaging*, vol. 18, no. 10, pp. 828-839, Oct. 1999.
- [7] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam, "Use of Active Shape Models for Locating Structures in Medical Images," *Image and Vision Computing*, vol. 12, no. 6, pp. 355-365, 1994.
- [8] M. de Bruijne, B. van Ginneken, M.A. Viergever, and W.J. Niessen, "Interactive Segmentation of Abdominal Aortic Aneurysms in cta Images," *Medical Image Analysis*, vol. 8, no. 2, pp. 127-138, 2004.
- [9] M.A. Styner, K.T. Rajamani, L.-P. Nolte, G. Zsemlye, G. Szekely, C.J. Taylor, and R.H. Davies, "Evaluation of 3D Correspondence Methods for Model Building," *Information Processing in Medical Imaging*, vol. 18, pp. 63-75, July 2003.
- [10] F. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology*. Cambridge Univ. Press, 1991.
- [11] M.S. Floater, "Parametrization and Smooth Approximation of Surface Triangulations," *Computer Aided Geometric Design*, vol. 14, no. 3, pp. 231-250, 1997.
- [12] C. Brechbühler, G. Gerig, and O. Kübler, "Parametrization of Closed Surfaces for 3D Shape Description," *Computer Vision and Image Understanding*, vol. 61, no. 2, pp. 154-170, 1995.
- [13] E. Praun and H. Hoppe, "Spherical Parametrization and Remeshing," *ACM Trans. Graphics*, vol. 22, no. 3, pp. 340-349, 2003.
- [14] T. Huysmans, J. Sijbers, and B. Verdonk, "Parameterization of Tubular Surfaces on the Cylinder," *J. WSCG*, vol. 13, no. 3, pp. 97-104, 2005.
- [15] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, "Multiresolution Analysis of Arbitrary Meshes," *Proc. ACM SIGGRAPH '95*, pp. 173-182, 1995.
- [16] H. Lamecker, T. Lange, and M. Seebaß, "A Statistical Shape Model for the Liver," *Proc. Fifth Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, T. Dohi and R. Kikinis, eds., pp. 421-427, 2002.

- [17] H. Lamecker, M. Seebaß, H.-C. Hege, and P. Deuffhard, "A 3D Statistical Shape Model of the Pelvic Bone for Segmentation," *Proc. SPIE—Medical Imaging 2004: Image Processing*, J. Fitzpatrick and M. Sonka, eds., pp. 1341-1351, May 2004.
- [18] R.H. Davies, C.J. Twining, T.F. Cootes, J.C. Waterton, and C.J. Taylor, "A Minimum Description Length Approach to Statistical Shape Modeling," *IEEE Trans. Medical Imaging*, vol. 21, no. 5, pp. 525-537, May 2002.
- [19] M.K. Hurdal, P.L. Bowers, K. Stephenson, D.W.L. Sumners, K. Rehm, K. Schaper, and D.A. Rottenberg, "Quasi-Conformally Flat Mapping the Human Cerebellum," *Proc. Second Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 279-286, 1999.
- [20] C. Gotsman, X. Gu, and A. Sheffer, "Fundamentals of Spherical Parameterization for 3D Meshes," *Proc. ACM SIGGRAPH '03*, pp. 358-363, 2003.
- [21] X. Gu, Y. Wang, T. Chan, P. Thompson, and S.-T. Yau, "Genus Zero Surface Conformal Mapping and Its Application to Brain Surface Mapping," *IEEE Trans. Medical Imaging*, vol. 23, no. 8, pp. 949-958, Aug. 2004.
- [22] S. Haker, S. Angenent, A. Tannenbaum, and R. Kikinis, "Non-distorting Flattening Maps and the 3D Visualization of Colon CT Images," *IEEE Trans. Medical Imaging*, vol. 19, no. 7, pp. 665-670, July 2000.
- [23] W. Hong, X. Gu, F. Qiu, M. Jin, and A. Kaufman, "Conformal Virtual Colon Flattening," *Proc. 2006 ACM Symp. Solid and Physical Modeling*, pp. 85-93, 2006.
- [24] T. Huysmans and J. Sijbers, "Cylindrical Surface Parameterization," to be published.
- [25] M. Jin, J. Kim, F. Luo, and X. Gu, "Discrete Surface Ricci Flow," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 5, pp. 1030-1043, Sept./Oct. 2008.
- [26] A.C.W. Kotcheff and C.J. Taylor, "Automatic Construction of Eigenshape Models by Direct Optimization," *Medical Image Analysis*, vol. 2, no. 12, pp. 303-314, Dec. 1998.
- [27] H.H. Thodberg, "Minimum Description Length Shape and Appearance Models," *Proc. Int'l Conf. Information Processing in Medical Imaging*, 2003.
- [28] R.H. Davies, C.J. Twining, T.F. Cootes, J.C. Waterton, and C.J. Taylor, "3D Statistical Shape Models Using Direct Optimisation of Description Length," *Proc. European Conf. Computer Vision*, pp. 3-20, 2002.
- [29] T. Heimann, I. Wolf, T. Williams, and H. Meinzer, "3D Active Shape Models Using Gradient Descent Optimization of Description Length," *Proc. Int'l Conf. Information Processing in Medical Imaging*, pp. 566-577, 2005.
- [30] P. Horkaew and G.-Z. Yang, "Construction of 3D Dynamic Statistical Deformable Models for Complex Topological Shapes," *Proc. Seventh Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 217-224, Sept. 2004.
- [31] P. Horkaew and G.Z. Yang, "Optimal Deformable Surface Models for 3d Medical Image Analysis," *Proc. Int'l Conf. Information Processing in Medical Imaging*, pp. 13-24, 2003.
- [32] X. Li, Y. Bao, X. Guo, M. Jin, X. Gu, and H. Qin, "Globally Optimal Surface Mapping for Surfaces with Arbitrary Topology," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 4, pp. 805-819, July/Aug. 2008.
- [33] L. Ferrarini, H. Olofsen, W.M. Palm, M.A. van Buchem, J.H. Reiber, and F. Admiraal-Behloul, "Games: Growing and Adaptive Meshes for Fully Automatic Shape Modeling and Analysis," *Medical Image Analysis*, vol. 11, no. 3, pp. 302-314, 2007.
- [34] J. Cates, P.T. Fletcher, M. Styner, M. Shenton, and R. Whitaker, "Shape Modeling and Analysis with Entropy-Based Particle Systems," *Information Processing in Medical Imaging*, vol. 20, pp. 333-345, 2007.
- [35] M. de Bruijne, B. van Ginneken, W. Niessen, and M. Viergever, "Adapting Active Shape Models for 3D Segmentation of Tubular Structures in Medical Images," *Proc. Int'l Conf. Information Processing in Medical Imaging*, C. Taylor and J. Noble, eds., pp. 136-147, 2003.
- [36] T. Huysmans, J. Sijbers, F. Vanpoucke, and B. Verdonk, "Improved Shape Modeling of Tubular Objects Using Cylindrical Parameterization," *Lecture Notes in Computer Science*, pp. 84-91, Springer, 2006.
- [37] R.H. Bartels, J.C. Beatty, and B.A. Barsky, *An Introduction to Splines for Use in Computer Graphics & Geometric Modeling*. Morgan Kaufmann Publishers, Inc., 1987.
- [38] I. Lloyd, N. Trefethen, and D. Bau, *Numerical Linear Algebra*. SIAM, 1997.
- [39] I.T. Jolliffe, *Principal Component Analysis*. Springer, Oct. 2002.
- [40] A. Ericsson and K. Åström, "Minimizing the Description Length Using Steepest Descent," *Proc. British Machine Vision Conf.*, vol. 2, pp. 93-102, 2003.
- [41] T. Papadopoulos and M.I.A. Lourakis, "Estimating the Jacobian of the Singular Value Decomposition: Theory and Applications," *Proc. European Conf. Computer Vision*, D. Vernon, ed., pp. 554-570, 2000.
- [42] E. Catmull and J. Clark, "Recursively Generated B-Spline Surfaces on Arbitrary Topological Meshes," *Computer-Aided Design*, vol. 10, no. 6, pp. 350-355, Nov. 1978.
- [43] P. Horkaew and G.-Z. Yang, "Optimal Deformable Surface Models for 3D Medical Image Analysis," *Proc. Int'l Conf. Information Processing in Medical Imaging*, pp. 13-24, 2003.
- [44] D. Loeckx, "Automated Nonrigid Intra-Patient Image Registration Using B-Splines," PhD dissertation, Katholieke Univ. Leuven, May 2006.
- [45] D. Rueckert, L.I. Sonoda, C. Hayes, D.L. Hill, M.O. Leach, and D.J. Hawkes, "Nonrigid Registration Using Free-Form Deformations: Application to Breast mr Images," *IEEE Trans. Medical Imaging*, vol. 18, no. 8, pp. 712-721, Aug. 1999.
- [46] D.C. Liu and J. Nocedal, "On the Limited Memory BFGS Method for Large Scale Optimization," *Math. Programming*, vol. 45, no. 3, pp. 503-528, 1989.



shape modeling, analysis, and visualization.



processing and analysis.



Toon Huysmans received the MSc degree in computer science in 2003 from the University of Antwerp, Belgium, where he is currently working toward the PhD degree in the Vision Lab. His MSc thesis was on optimization in diamond cut studies. From 2004 to 2008, he was a doctoral fellow of the IWT (Institute for the Promotion of Innovation through Science and Technology in Flanders). His main interests are image and geometry processing with focus on biomedical

Jan Sijbers received the MSc degree in physics and the PhD degree in sciences from the University of Antwerp, Belgium, in 1993 and 1998, respectively. From 2002 to 2008, he was a postdoctoral fellow of the FWO (Fund for Scientific Research—Flanders). Since 2005, he has been a lecturer at the University of Antwerp. His main interests are in biomedical image processing with focus on magnetic resonance and computer tomography image

Brigitte Verdonk received the MSc degree in computer science from Stanford University, California, in 1984, and the PhD degree from the Universiteit Antwerpen, Belgium, in 1988. After a few years in industry, she returned to academic life. She is currently a full professor in the Department of Mathematics and Computer Science, Universiteit Antwerpen. She has authored or coauthored some 50 refereed publications in international journals and proceedings.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.