Propagation of shape parameterisation for the construction of a statistical shape model of the left ventricle

Matthias Kirschner and Stefan Wesarg

TU Darmstadt, GRIS
Email: {matthias.kirschner,stefan.wesarg}@gris.tu-darmstadt.de

Abstract

Statistical Shape Models (SSMs) have been successfully applied to both segmentation and the description of the dynamic behaviour of the heart. SSMs are learned from a set of training examples, which are represented by vectors of corresponding landmarks. While the construction of a SSM is simple when a landmark representation of the training shapes is available, the extraction of corresponding landmarks from training images or meshes of different sizes is difficult. Optimisation schemes that solve this so-called *correspondence problem* rely on a parameter space representation of the input shapes. These optimisation schemes tend to be sensitive to the initial parameterisation of the input shapes. In this work, we present an algorithm to produce a consistent spherical parameterisation for shapes of the left ventricle. Our algorithm propagates the spherical parameterisation of a root shape within seconds to all other shapes. We demonstrate the effectiveness of our approach by extracting a SSM from the parameterisations generated by our algorithm.

1 Introduction

The dynamic nature of the heart requires 4D imaging (3D volume + time) for analysing its deformation over the cardiac cycle. Over the past decade, a set of parameters have emerged as standard descriptors for the dynamic behavior of the heart [9] – especially the left ventricle (LV). A pre-requisite for such an analysis is the segmentation of the relevant cardiac structures. A family of segmentation algorithms that has been proven to be suitable for that task are approaches that employ statistical shape models [8]. For that, variations of the active shape model (ASM) or active appearance model (AAM) – which have been introduced by Cootes et al. [5, 18] – are widely used [20, 10].

A part of the model building step is the computation of the variability of the vectors that represent the shape (or the texture) of the structure to be modelled. For that purpose, multivariate analysis is the method of choice. Principal component analysis (PCA) is the simplest of these techniques and describes the global variation of the structures of interest over a set of training data. A more localized modelling of the variation can be achieved by means of independent component analysis (ICA) [19]. PCA as well as ICA are not only used for segmenting cardiac structures [11] but also for modelling the deformation of the heart over the cardiac cycle [17, 14, 20].

The shape of an organ – here the LV – can be expressed by a triangular mesh whose vertices are given by the position of a set of landmarks. These landmarks are the constituents of the shape vector for each of the data sets. Before constructing the statistical model, correspondences between the shape vectors have to be established. That means, the *j-th* landmarks on model i and model k should both represent the same anatomical feature.

Many state-of-the-art algorithms for automatic determination of corresponding landmarks treat the problem as one of reparameterisation. For each training shape, a parameter space representation is computed by mapping the shape to a topological primitive. For example, shapes with genus-0 topology can be mapped to the unit sphere. An optimisation algorithm then manipulates the position of sampling points and extracts landmarks by sampling in the parameter space. The process is guided by an objective function which assesses the quality of the extracted landmarks. For example, Davies et al. [6] introduced an objective function based on the Minimum Description Length principle. An illustration of this optimisation scheme is shown in

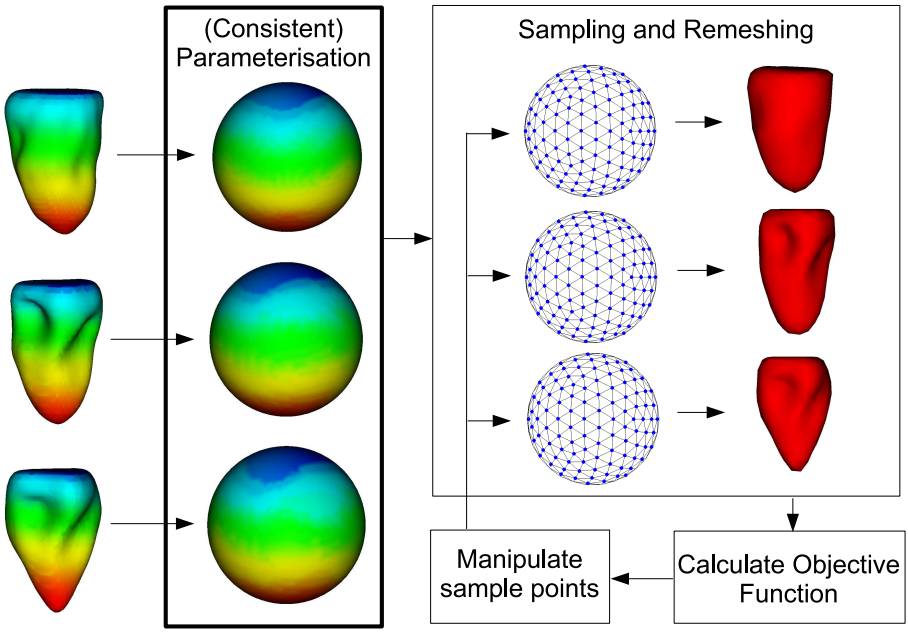

Figure 1: Illustration of a typical optimisation scheme for automatic landmark determination. Each training shape is mapped to a suitable parameter space, for example the unit sphere. The shapes are remeshed using sampling points which are initially uniformly distributed over each parameter space. Correspondence is then optimised by iterative manipulation of sample points and evaluation of correspondences in the remeshed shapes using an appropriate objective function. In this paper, we focus on computing consistent parameterisations of the input shapes.

figure 1.

The scheme described above is computationally very demanding. Even mapping a training shape to a parameter space is time-consuming and typically involves an optimisation process on its own, as the parameter space representation must allow for an easy reconstruction of the original training shape such that no important detail gets lost. For example, when uniform sampling of the parameter space is used to reconstruct a shape, the mapping must be approximately area preserving, that is triangles of the mesh must be mapped to areas of the parameter space with the same relative size. Such area preserving parameterisations can be constructed by various algorithms, which usually takes several minutes or even longer, depending on the size of the shape.

To add to the complexity, the optimisation scheme for landmark determination is sensitive to the initial parameterisations of the shapes. Davies et. al. [7] report that in cases of poor initial correspondence, the convergence time may slow down or the algorithm even fails to converge. Therefore, they propose to produce *groupwise-consistent* parameterisations [7, 8]. With consistent they mean that the same regions in different parameter spaces parameterise approximately the same regions of the training shapes. To obtain such a consistent parameterisation, they firstly select a reference shape and compute an area preserving parameterisation of this shape. For all other shapes, non-optimised parameterisations are computed. Then a point correspondence between points on the new shape and on the reference shape is derived by alignment. Using this point correspondence, the parameterisation of the new shape is manipulated in order to minimise the euclidean distance between corresponding points in the parameter space. One advantage

of this approach is that properties like area preservation are automatically carried over from the reference parameterisation to all other parameterisations. On the other hand, the authors do not give any figures on how much time is spent to compute such a consistent parameterisation. Since their approach is a black-box optimisation approach, it is unlikely that it generates a consistent parameterisation efficiently.

In this paper, we present a method which exploits the similarity of shapes to propagate the parameterisation from a *root mesh* to all other training shapes. In contrast to Davies approach [7, 8], we only compute the parameterisation of this root mesh from scratch and use approximate point correspondences derived from an alignment algorithm in order to quickly propagate the parameterisation to all other shapes. Desireable properties as area preservation are likewise carried over from the parameterisation of the root mesh. In our experiments, propagating the parameterisation to 24 shapes took significantly less time than the parameterisation of the root mesh.

2 Related Work

Beside Davies et. al. [7, 8], several other researchers worked on consistent parameterisations: Alexa [1] computes consistent spherical parameterisations of two polyhedral shapes for shape morphing. Corresponding features on the shapes must be specified by the user. In contrast to our work, both shapes are independently mapped to the sphere and aligned afterwards. Praun et al. [15] establish consistent parameterisations for a group of genus-0 models. They do not use the sphere as a base domain, but a base domain mesh with *feature points*. Both base domain mesh and corresponding features on the input meshes must be specified by the user. The surface patches of the base domain mesh are mapped to the input mesh by a path tracing algorithm. After this mapping, the interior of the patches on the input meshes is parameterised. It is unlikely that their approach easily extends to spherical parameterisation since they choose the base domain mesh to have an abstract geometry similar to the geometry of the meshes they want to parameterise. Asirvatham et al. [2] produce consistent spherical parameterisations. Again, corresponding features must be specified by the user. The running times reported in their paper are by far too long for our purposes as we regard the construction of a consistent parameterisation only as an initialisation step of a larger optimisation scheme.

3 Materials and Methods

3.1 Terminology and notation

A mesh is a tuple $M = (V, T)$, where $V = \{v_j | v_j \in \mathbb{R}^3, 1 \leq j \leq N\}$ are the mesh points and $T = \{\{v_j, v_k, v_l\} | v_j, v_k, v_l \in V\}$ is a triangulation over these points. For a node $v \in V$, the set $\mathcal{N}(v) = \{v' \in V | \exists t \in T : v, v' \in t\}$ is the *neighbourhood* of v. A *spherical parameterisation* $\mathcal{P}$ of the mesh M is a function that assigns to each node $v \in V$ a pair of spherical coordinates $\mathcal{P}(v) = (\theta, \phi)$, where θ is the *colatitude* and ϕ the *longitude*. By $\mathcal{C}(v) = (\mathcal{C}^x(v), \mathcal{C}^y(v), \mathcal{C}^z(v))$, we denote the three-dimensional cartesian representation of the spherical coordinates. A parameterisation is *valid* if there are no triangles on the sphere that are inverted or overlap.

3.2 Algorithm overview

As input, our algorithm receives a labeled $4D$ volume which contains the segmentation of a left ventricle over the whole cardiac cycle. From each 3D subvolume representing one of the K different time phases, we extract a triangle mesh M_i using the Marching Cubes Algorithm [13]. We denote the set of all meshes by $\mathcal{M} = \{M_1, \ldots, M_K\}$. All extracted meshes have genus-0 topology, which is achieved by processing the volumes with standard image processing techniques prior to the surface extraction step in order to fill possible holes. Our goal is to compute a parameterisation $\mathcal{P}_i$ for each mesh, such that all parameterisations are *consistent*, that is the same regions in the parameter spaces describe approximately the same features in the meshes.

Our algorithm works as follows:

1. Determine a rooted *propagation tree* $\mathcal{T} = (\mathcal{M}, \mathcal{A}), \mathcal{A} \subset \mathcal{M} \times \mathcal{M}$. The tree nodes are the input meshes and the arcs determine the order of propagation. (section 3.3)
2. Compute an area preserving parameterisation for the mesh $M_t \in \mathcal{M}$ that corresponds to the root of the propagation tree. (section 3.4).
3. Traverse through the whole tree, starting from the root M_t. For each arc $(M_{\text{ref}}, M_i) \in \mathcal{A}$,

propagate the parameterisation of M_{ref} to M_i. The details of such a propagation step are given in section 3.5.

3.3 Determination of propagation order

We follow the intuitive idea that similar genus-0 shapes have a similar mapping to the sphere. As our data sets represent the beating left ventricle at different points in time, we can exploit the fact that meshes extracted from adjacent time frames are similar to each other. We first select a mesh M_t as tree root, which we name *root mesh*. Assume without loss of generality that $t = \lfloor \frac{K}{2} \rfloor$. By defining the set of arcs of our propagation tree as

$$\begin{aligned} \mathcal{A} &= \{(M_{i+1}, M_i) | 1 \leq i < t\} \cup \\ &\quad \{(M_i, M_{i+1}) | t \leq i < K\}, \end{aligned}$$

we accomplish that the initial parameterisation of the root mesh is propagated in both time directions along similar shapes. It is tempting to select t to be at end-systole, as the mesh extracted from the compressed heart has the smallest size, which reduces the time to compute an initial parameterisation. However, we got slightly better results by selecting t to be at end-diastole. It remains an open question if this observation can be attributed to the greater size or the lower curvature of the mesh extracted at end-diastole.

The determination of the propagation order is the only step of our algorithm which is application specific. In other applications, a natural ordering of similar shapes may not be given a-priori. We postpone the discussion of this point to section 5.

3.4 Parameterisation of the root mesh

To compute parameterisation of the root mesh, we apply a method proposed by Brechbühler et al. [4] to compute an initial parameterisation. In this method, systems of linear equations are constructed for both colatitude and longitude parameters. The equations are derived from the triangulation of the mesh. To resolve ambiguities, additional constraints are obtained by selecting two nodes as poles and computing a meridian that connects these poles. Since this initial parameterisation suffers from large area distortions, we apply an algorithm we presented in [12] to manipulate the parameterisation in order to minimise area distortion. After this optimisation step, the mesh can be reconstructed from the parameter space representation in high detail.

3.5 Propagation of shape parameterisation

Given a mesh M_i, a reference mesh M_{ref} and a spherical parameterisation $\mathcal{P}_{\text{ref}}$ of the reference mesh, we want to compute a spherical parameterisation $\mathcal{P}_i$ of M_i. We do this as follows:

1. Align the shapes in order to derive a relation R of corresponding points of the two meshes. (section 3.5.1)
2. Propagate the parametersation using the relation R and interpolate parameter space coordinates for all nodes that do not occur in R. (section 3.5.2)
3. Correct and refine the parameterisation to obtain a valid and more regular parameterisation. (section 3.5.3)

3.5.1 Deriving corresponding points

Firstly, we align M_i and M_{ref} using the ICP algorithm [3]. After aligning the shapes, we derive a relation $R \subseteq V_i \times V_{\text{ref}}$ of corresponding points such that $(v, w) \in R$ if $w \in \text{argmin}_{w_j \in V_{\text{ref}}} \|v - w_j\|_2$ and $v \in \text{argmin}_{v_j \in V_i} \|v_j - w\|_2$, where $\| \cdot \|_2$ denotes the euclidean norm. We force R to be left- and right-unique, that is for every point $v \in V_i$ there is either one or no $w \in V_{\text{ref}}$ such that $(v, w) \in R$, and vice versa. Possible ties are broken arbitrarily. The relation induces a partition of V_i into the set of *fixed* nodes

$$V_i^f = \{v \in V_i | \exists w \in V_{\text{ref}} : (v, w) \in R\}$$

and the set of *moving* nodes $V_i^m = V - V_i^f$.

To improve the robustness of the propagation step, it is useful to scale the meshes prior to alignment such that they have roughly the same size. We scale the mesh M_i anisotropically, such that it has the same extensions along its principal axes as M_{ref}. Principal axes and the respective scaling factors are computed using PCA. Note that this anisotropic scaling is only temporary and only affects the correspondence relation.

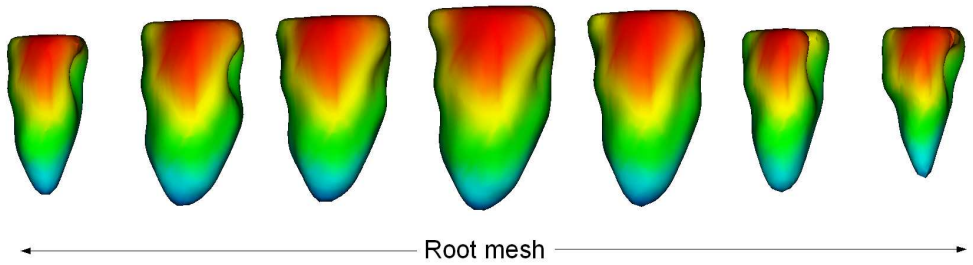

Figure 2: Seven reconstructed shapes from a sequence of 25 shapes. Each shape was remeshed using 1002 sample points. Colour denotes corresponding regions. The shape in the middle is the ventricle at end-diastole and was chosen as root mesh. The arrows show the direction of propagation of parameterisation. The outermost shapes in the figure have the highest distance to the root mesh in the propagation tree (edge distance 12).

3.5.2 Propagation

To generate the parameterisation $\mathcal{P}_i$, we first assign to each fixed nodes $v \in V_i^f$ the spherical coordinates of corresponding node from M_{ref}, that is we set

$$\mathcal{P}_i(v) = \mathcal{P}_{\text{ref}}(w) \tag{1}$$

for each $(v, w) \in R$. The parameter space coordinates of the moving nodes are then interpolated by solving a system of linear equations. Similar to Brechbühler's method [4], we build a $|V_i^m| \times |V_i^m|$-matrix $\mathbf{A}$ from the mesh connectivity. The individual entries of the matrix $\mathbf{A}$ are given by

$$\mathbf{A}_{vw} = \begin{cases} |\mathcal{N}(v)| & v = w \\ -1 & w \in \mathcal{N}(v) \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

Here, we index the matrix entries by the moving nodes to which they correspond. We then use the cartesian representation of the fixed nodes to construct three vectors $\mathbf{b}^x, \mathbf{b}^y, \mathbf{b}^z$, whose components are defined as

$$\mathbf{b}_v^a = \sum_{w \in \mathcal{N}(v) \cap V_i^f} \mathcal{C}_i^a(w) \tag{3}$$

for $a \in \{x, y, z\}$. By solving the three linear systems $\mathbf{A}\mathbf{x}^a = \mathbf{b}^a$, $a \in \{x, y, z\}$, we can assign the cartesian parameter space coordinates

$$\mathcal{C}_i(v) = (\mathbf{x}_v^x, \mathbf{x}_v^y, \mathbf{x}_v^z) \tag{4}$$

to a moving node $v \in V_i^m$. These coordinates are then projected to the unit sphere by normalising them such that $\|\mathcal{C}_i(v)\|_2 = 1$.

3.5.3 Correction and refinement

The method described above produces in many cases invalid parameterisations, that is some triangles on the sphere overlap or are inverted. The occurence of these artifacts is not suprising as we completely ignore the triangulation T_i during the determination of the correspondence relation R. Since the system of linear equations adheres to the triangle structure, artifacts can be removed to some extent by including problematic fixed nodes in the linear systems. That is, after having assigned the coordinates of the fixed nodes (Eq. 1), we identify triangles on the sphere whose nodes only consists of fixed nodes and compute their surface normals. If the normal of such a triangle points to the inside of the sphere, its nodes are removed from the set of fixed nodes and added to the moving nodes.

The remaining artifacts are removed in a subsequent Laplacian-like smoothing step [16]. We iterate several times through the mesh nodes and set the parameter space coordinates of each node v to a convex combination of the coordinates of its neighbours, that is

$$\mathcal{C}_i(v) = \frac{1}{\sum_{w \in \mathcal{N}(v)} a_{vw}} \sum_{w \in \mathcal{N}(v)} a_{vw} \mathcal{C}_i(w).$$

We use $a_{vw} = \|u - w\|^{-1}$. This choice aims at reducing local length distortions, because the nodes are drawn towards those neighbours which are also closest in the original mesh. By normalizing the coordinates $\mathcal{C}_i(v)$, they are projected back to the surface of the unit sphere.

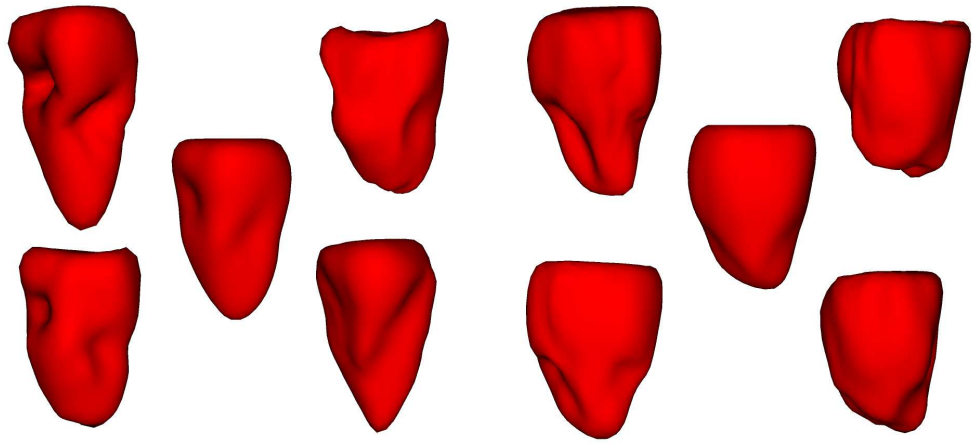

(a) First data set (b) Second data set

Figure 3: The first two modes of variation of the shape models derived from the data sets using our method of consistent parameterisation. In each subfigure, the shape in the middle is the mean shape $\bar{x}$. The shapes on the left and on the right are obtained by adding the eigenvectors p_i, $i \in \{1, 2\}$, to the mean shape. The eigenmodes are weighted by $\pm 3\sqrt{\lambda_i}$, where λ_i is the eigenvalue corresponding to the eigenmode p_i.

Note that the smoothing operations are localized in their effect, which means that they do not destroy the correspondence of regions in parameter space.

On our ventricle data sets, 50 smoothing iterations always produced valid parameterisations. However, we stress that our correction and refinement step is a heuristic approach, so there is no guarantee that it converges to a valid parameterisation.

4 Results

We illustrate the results produced by our algorithm on the basis of two representative data sets. The first one shows the LV from a patient with a myocardial infarction, whereas the second comes from a healthy patient. Figure 2 shows an extract of seven reconstructed ventricles from the first data set. The presented shapes range over the whole cardiac cycle. An approximate uniform sampling of the unit sphere was obtained by projecting the vertices of an subdivided unit icosahedron to the unit sphere and using them as sample points [8]. We mapped the indices of the 1002 sample points to hue in HSV colour space to illustrate corresponding regions. As one can see in the figure, the same colour was indeed mapped to corresponding regions on different shapes. Some sharp edges occur in regions of high curvature on the outermost shapes. These artifacts arise from undersampling in parameter space, since the area preservation property was not perfectly propagated in some local regions.

We computed SSMs based on PCA [5] to illustrate that the quality of correspondence determined by our method is high. Figure 3 presents the first two eigenmodes of the first and the second data set, respectively. In both models, the modes of variation show plausible shape deformations, like the contraction of the ventricle during the cardiac cycle. The folding visible in the first mode of the first data set corresponds to the infarcted region – confirmed by late enhancement imaging. However, our algorithm is only intended as an initialisation phase of a complex optimisation scheme. We are convinced that we can capture even more meaningful statitistics by solving the full correspondence problem.

Table 1 lists the sizes of the data sets and the running times of our algorithm. Most time is spent on computing an area-preserving spherical parameterisation of the root mesh. It takes roughly one minute to compute this first parameterisation. The computation of all remaining parameterisations is done within ten seconds. The times given in the table include loading the images from hard disk, process-

Dataset	K	root mesh size # nodes	root mesh size # triangles	average mesh size # nodes	average mesh size # triangles	Running Times root paramet.	Running Times propagation
1	25	4351	8698	3033	6062	$\approx$ 54 s	$\approx$ 9 s
2	20	5721	11438	4035	8066	$\approx$ 68 s	$\approx$ 10 s

Table 1: Sizes of the test data sets and running times. The test hardware was an Intel Quad Core with 2.4 GHz and 3 GB RAM. The libraries used for image processing and surface extraction exploit multi-threading facilities of the processor, our own algorithms are single-threaded.

ing the images, surface extraction, alignment with the ICP algorithm and propagation of shape parameterisation.

5 Discussion and Future work

In this paper, we presented an algorithm that propagates the parameterisation of a single shape to all other shapes in a training set. Our new method is much faster than computing parameterisations of the individual training shapes independently from scratch, though it handles consistency as an additional requirement. Properties like area preservation can also be propagated to all other training shapes. To the best of our knowledge, we are the first who introduced and applied the idea of parameter space propagation.

Using our method, all training shapes could be reconstructed sufficiently well. Only in rare cases, small sampling artifacts could be observed in shapes which are far away from the root mesh in the propagation tree, which are caused by small local area distortions. The consistency of the produced parameterisations was evaluated by two SSMs learned from our test data sets. Both of them modelled the variability of the beating heart in a plausible way, although correspondence still needs to be refined, as discussed above. Nevertheless, we were able to capture meaningful statistics of the data even without additional optimisation. We therefore conclude that the parameterisations generated by our method are indeed consistent and close to an optimal solution.

The application considered in the paper allows us to derive a natural ordering of shapes in which similar shapes are close together. This ordering is desirable, since the propagation step produces parameterisations of higher quality if the shapes are similar. In applications in which such an ordering is not given a-priori, the determination of the propagation order may be specified by the user. As we want to avoid user intervention in the learning process, we will explore the use of feature extraction algorithms to derive such an ordering automatically.

Currently, our method uses solely the euclidean distance to derive approximate correspondences for the propagation step. As the euclidean distance is often a suboptimal measure for point correspondence, the relation of corresponding points is locally incoherent, which means that some triangles in the parameter space are inverted. In case of the left ventricle, this incoherence is not severe and can be handled by removal and reinsertion of nodes belonging to inverted triangles, as mentioned in section 3.5.3. However, our method lacks robustness when it is applied to classes of shapes with high intra-class-variability. Nevertheless, we are convinced that our general approach can be extended to handle organs with large shape variations as well. To achieve this, we will explore two complementary approaches in order to improve the robustness of our method: Firstly, we want to extend the propagation model to overcome the deficiency of the euclidean distance as a correspondence measure. This may be achieved by replacing the ICP algorithm by a more sophisticated point set alignment algorithm, which also allows for non-rigid alignment. Secondly, we will explore whether the mesh connectivity can be used to improve coherence during the identification of corresponding points.

Furthermore, we want to evaluate our approach in the context of the full correspondence pipeline.

References

[1] M. Alexa. Merging polyhedral shapes with scattered features. *The Visual Computer*, 16(1):26–37, 2000.

[2] A. Asirvatham, E. Praun, and H. Hoppe. Consistent spherical parameterization. In *Interna-*

tional Conference on Computational Science (2), pages 265–272, 2005.

[3] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[4] C. Brechbühler, G. Gerig, and O. Kübler. Parametrization of closed surfaces for 3-D shape description. *Computer Vision and Image Understanding*, 61(2):154–170, 1995.

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

[6] 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 Transactions on Medical Imaging*, 21(5):525–537, 2002.

[7] R. H. Davies, C. J. Twining, and C. J. Taylor. Consistent spherical parameterisation for statistical shape modelling. In *IEEE Symposium on Biomedical Imaging*, pages 1388–1391, 2006.

[8] R. H. Davies, C. J. Twining, and C. J. Taylor. *Statistical Models of Shape - Optimization and Evaluation*. Springer, 2008.

[9] A. F.Frangi, W. J. Niessen, and M. A. Viergever. Three-Dimensional Modeling for Functional Analysis of Cardiac Images: A Review. *IEEE Transactions on Medical Imaging*, 20(1):2–25, 2001.

[10] K. Fritscher and R. Schubert. 4D Endocardial Segmentation using Spatio-temporal Appearance Models and Level Sets. In T. Tolxdorff, editor, *Bildverarbeitung für die Medizin*, Informatik aktuell, pages 1–5. Springer, 2008.

[11] D. Fritz, D. Rinck, R. Dillmann, and M. Scheuering. Segmentation of the left and right cardiac ventricle using a combined bi-temporal statistical model. In K. R. Cleary, R. L. Galloway, and Jr., editors, *Proc. of SPIE Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, volume 6141, 2006.

[12] M. Kirschner and S. Wesarg. Area preserving parameterisation of shapes with spherical topology. In *Informatik 2009: Workshop Medizinische Bildverarbeitung und Mustererkennung*, 2009.

[13] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH '87*, pages 163–169, 1987.

[14] R. Pilgram, K. D. Fritscher, R.-H. Zwick, M. F. Schocke, T. Trieb, O. Pachinger, and R. Schubert. Shape analysis of healthy and diseased cardiac ventricles using PCA. In H. U. Lemke, editor, *Proc. of the 19th CARS 2005*, pages 357–362. Elsevier, 2005.

[15] E. Praun, W. Sweldens, and P. Schröder. Consistent mesh parameterizations. In *SIGGRAPH '01*, pages 179–184, 2001.

[16] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit*. Kitware Inc., 3 edition, 2002.

[17] A. Suinesiaputra, M. Üzümcü, A. F. Frangi, T. A. M. Kaandorp, J. H. C. Reiber, and B. P. F. Lelieveldt. Detecting regional abnormal cardiac contraction in short-axis mr images using independent component analysis. In Christian Barillot, editor, *Proc. of MICCAI 2004*, volume 3216, pages 737–744. Springer, 2004.

[18] G. J. Edwards T. F. Cootes and C. J. Taylor. Active appearance models. In *5th European Conference on Computer Vision*, volume 2, pages 484–498. Springer, 1998.

[19] M. Üzümcü, A. F. Frangi, J. H.C. Reiber, and B. P.F. Lelieveldt. Independent component analysis in statistical shape models. In *Proceedings of SPIE Medical Imaging 2003: Image Processing,*, volume 5032, 2003.

[20] H. Zhang, M. T. Thomas, N. E. Walker, A. H. Stolpen, A. Wahle, T. D. Scholz, and M. Sonka. Four-dimensional functional analysis of left and right ventricles using MR images and active appearance models. In A. Manduca and X. P. Hu, editors, *Proc. of SPIE Medical Imaging 2007: Physiology, Function, and Structure from Medical Images*, volume 6511, 2007.