Fusing Correspondenceless 3D Point Distribution Models

Marco Pereañez[1], Karim Lekadir[1], Constantine Butakoff[2], Corné Hoogendoorn[1], and Alejandro Frangi[3]

[1] CISTIB, Universitat Pompeu Fabra and CIBER-BBN, Barcelona, Spain
[2] PhySense, Universitat Pompeu Fabra, Barcelona, Spain
[3] CISTIB, University of Sheffield, Sheffield, UK

Abstract. This paper presents a framework for the fusion of multiple point distribution models (PDMs) with unknown point correspondences. With this work, models built from distinct patient groups and imaging modalities can be merged, with the aim to obtain a PDM that encodes a wider range of anatomical variability. To achieve this, two technical challenges are addressed in this work. Firstly, the model fusion must be carried out directly on the corresponding means and eigenvectors as the original data is not always available and cannot be freely exchanged across centers for various legal and practical reasons. Secondly, the PDMs need to be normalized before fusion as the point correspondence is unknown. The proposed framework is validated by integrating statistical models of the left and right ventricles of the heart constructed from different imaging modalities (MRI and CT) and with different landmark representations of the data. The results show that the integration is statistically and anatomically meaningful and that the quality of the resulting model is significantly improved.

1 Introduction

Despite their widespread use in medical imaging, point distribution models (PDMs) [2] often suffer from over-fitting, or lack of sufficient generalization ability. This is due to the fact that typically too few training examples are available compared to the dimensionality of the shapes themselves. Extensive work has been dedicated to improving the generalization ability of statistical shape models. These works can be roughly categorized into: a) methods that *add artificial variation to a model*, b) methods that *combine statistical and physical models*, and c) methods that *model object ensembles, and/or object hierarchies*. However, little attention has been placed on methods that aggregate preexisting statistical models into a single more robust model. Yet, this would allow to unify PDMs from different research centers, patient groups and imaging modalities, without the need for the original data, which is not always available and cannot be freely exchanged across centers for various legal and practical reasons.

In the recent years, a few techniques have been published to fuse means and eigenvectors [4,11,1]. However, they require point correspondence between models, a constraint which is rarely valid in practice as the PDMs are derived from

K. Mori et al. (Eds.): MICCAI 2013, Part I, LNCS 8149, pp. 251–258, 2013.

datasets delineated at different clinical centers, using distinct delineation protocols and observers. Unlike for shapes [6], the problem of establishing point correspondence between eigenspaces has not been yet investigated.

In this paper we present a method for the normalization and fusion of statistical models of anatomical shape. To achieve this, an inter-model barycentric mapping is estimated by combining rigid and nonrigid registrations of the means, with the aim to transform the PDMs into a common landmark representation. Subsequently, a model fusion technique is applied which, unlike existing methods [5], takes into account the weights of each model in the final result, thus obtaining a model that is statistically and anatomically meaningful. The proposed framework is validated by integrating anatomical models of various geometrical and statistical complexity, based on left and right ventricular MRI and CT datasets.

2 Normalization Algorithm

The goal of the proposed algorithm is the creation of a single fused model from a set of models built from shapes that do not share point correspondence and hence, have different eigenspace representations, and different dimensionality. Since we do not have the original datasets, we need to solve the point correspondence problem among the different model means, and find a transformation that re-parametrizes all eigenvectors into a common space.

For the remainder of the paper, let $\Omega_i = (\bar{\mathbf{x}}_i, \mathbf{\Phi}_i, \mathbf{\Lambda}_i, K_i, V_i)$, be the i-th model in the set $\{\Omega_i : i = 1, \ldots, N\}$ of pre-built models of the same object with different landmark placement strategy. Here $\bar{\mathbf{x}}_i$ is a vector of concatenated vertex coordinates representing the mean shape, $\mathbf{\Phi}_i$ and $\mathbf{\Lambda}_i$ are the eigenvector and eigenvalue matrices, respectively, and K_i and V_i are the number of observations used to build the model, and their triangulation. Now, let $\mathcal{M}_i = (\bar{\mathbf{x}}_i, V_i)$ be the model mean surfaces defined by the vertices $\bar{\mathbf{x}}_i$ and triangulations V_i.

2.1 Computing Surface Correspondence

Selecting a reference surface mesh is the first step in computing surface correspondence across all models. Here, we choose the highest resolution mesh (*i.e.* CT) as reference, to retain as much shape information – high frequency detail – as possible from the different models. The aim is to establish surface correspondence between each of the models' mean surfaces $\mathcal{M}_i$, and the chosen reference mean surface $\mathcal{M}_{ref}$. This should produce a mapping that linearly transforms the shape representation of each PDM to match that of the reference PDM. To achieve this, we first perform an initial rigid registration using the iterative closest point (ICP) algorithm in order to resolve pose differences between the means. Subsequently we apply the currents-based diffeomorphic registration method [3] to obtain a more accurate surface correspondence. In other words, we estimate a combined transformation $\varphi_i : \mathcal{M}_i \rightarrow \mathcal{M}_{ref}$ such that $\varphi_i(\mathcal{M}_i)$ is as close as possible to $\mathcal{M}_{ref}$. Note that neither ICP nor the currents approach rely on point correspondences, thus making them fit for this purpose.

2.2 Computing Inter-model Landmark Mapping

Once surface correspondence has been obtained, we proceed to find a suitable transformation that re-parametrizes all meshes to match the parametrization found on the reference mesh.

Specifically, we seek a transformation matrix $\mathbf{T}_i$ that expresses every point on $\mathcal{M}_{ref} = (\bar{\mathbf{x}}_{ref}, V_{ref})$ as a linear combination of the vertices on the previously registered mean meshes $\mathcal{M}_i = (\varphi_i(\bar{\mathbf{x}}_i), V_i)$.

Let $\varphi_i(\bar{\mathbf{x}}_i)$, for $i = 1, \ldots, N$, be the points of the i-th mean shape on the reference surface $\mathcal{M}_{ref}$. Without any loss of generality we can use the same triangulation that is already present in the surface $\mathcal{M}_{ref}$. Let $\mathbf{T}_i$ be a linear transformation matrix such that $\bar{\mathbf{x}}_{ref} \approx \mathbf{T}_i\varphi_i(\bar{\mathbf{x}}_i)$. It follows naturally that this is achieved when $\mathbf{T}_i$ contains the barycentric coordinates of the points $\bar{\mathbf{x}}_{ref}$ expressed in terms of points $\varphi_i(\bar{\mathbf{x}}_i)$ and triangulation V_i. $\mathbf{T}_i\bar{\mathbf{x}}_i$ will change the triangulation of the surface $\mathcal{M}_i$ to match that of $\mathcal{M}_{ref}$ and therefore it can be represented by the pair $\mathcal{M}'_i = (\mathbf{T}_i\bar{\mathbf{x}}_i, V_{ref})$.

For every point in the reference shape we compute the normal to its mesh surface, and the point of intersection with a plane described by a triangle on the target mesh. We then compute barycentric coefficients $c_{i,1}$, $c_{i,2}$, $c_{i,3}$ for each point in the triangle, and construct a point transformation matrix $\mathbf{P}_i$ as

$$\mathbf{P}_i = \begin{bmatrix} 0 & c_{i,1} & 0 & 0 & 0 & 0 & c_{i,2} & 0 & 0 & 0 & 0 & c_{i,3} & 0 & 0 & 0 & \\ \cdots 0 & 0 & c_{i,1} & 0 & 0 & \cdots 0 & 0 & c_{i,2} & 0 & 0 & \cdots 0 & 0 & c_{i,3} & 0 & 0 & \cdots \\ 0 & 0 & 0 & c_{i,1} & 0 & 0 & 0 & 0 & c_{i,2} & 0 & 0 & 0 & 0 & c_{i,3} & 0 & \end{bmatrix}_{3\times 3m},$$

where the column-wise position of every 3×3 diagonal matrix $\mathbf{C_i}$ encodes the indexing of a simplex on the target mesh, thus preserving point correspondence between the reference and target meshes. m is the number of points in the target mesh. Finally the complete transformation matrix $\mathbf{T}_i$ is a row-wise concatenation of point transformation matrices $\mathbf{P}_i$,

$$\mathbf{T}_i = \begin{bmatrix} \mathbf{P}_i^{(1)} \cdots \mathbf{P}_i^{(n)} \end{bmatrix}^T_{3n\times 3m}$$

where n is the number of points in the reference mesh. Due to the linearity of $\mathbf{T}_i$, any shape $\mathbf{x}_i \approx \bar{\mathbf{x}}_i + \mathbf{\Phi}_i\mathbf{b}_i$ can also be re-parametrized as $\mathbf{T}_i\mathbf{x}_i \approx \mathbf{T}_i\bar{\mathbf{x}}_i + \mathbf{T}_i\mathbf{\Phi}_i\mathbf{b}_i$. Thus, the models to be fused are $\Omega'_i = (\mathbf{T}_i\bar{\mathbf{x}}_i, \mathbf{T}_i\mathbf{\Phi}_i, \mathbf{\Lambda}_i, K_i, V_{ref})$.

3 Fusion Algorithm

Once all eigenspaces have been transformed to having a common triangulation, the models $\Omega'_i = (\mathbf{T}_i\bar{\mathbf{x}}_i, \mathbf{T}_i\mathbf{\Phi}_i, \mathbf{\Lambda}_i, K_i, V_{ref}), i = 1, \ldots, N$ are fused. To simplify notation, we denote the newly transformed mean shapes and eigenvector matrices as $\bar{\mathbf{x}}'_i = \mathbf{T}_i\bar{\mathbf{x}}_i$ and $\mathbf{\Phi}'_i = \mathbf{T}_i\mathbf{\Phi}_i$, respectively.

The goal of fusion is to compute an eigenspace $\Omega = (\bar{\mathbf{x}}, \mathbf{\Phi}, \mathbf{\Lambda}, K)$, using only information from $\Omega'_i = (\bar{\mathbf{x}}'_i, \mathbf{\Phi}'_i, \mathbf{\Lambda}_i, K_i)$, for $i = 1, \ldots, N$ as described in the previous section. The resulting model should be equivalent to one built from the full set of original observations.

3.1 Global Alignment

Since every PDM has a different mean shape and a different pose, the first step is to align the means $\bar{\mathbf{x}}'_i$ of all PDMs using Procrustes Analysis [8]. Then the aligned means $\bar{\mathbf{x}}'_i$ are used to estimate the fused mean $\bar{\mathbf{x}}$.

During alignment of the means, shapes are centered and rescaled to unit size and a $d \times d$ matrix accounting for rotation is estimated, where d is the dimensionality of landmarks (*i.e.*, $d = 3$). Let $\mathbf{S}_i$ be the $d \times d$ rotation matrix from the transformation that aligns the shape $\bar{\mathbf{x}}'_i$ to the mean $\bar{\mathbf{x}}$. Let $\boldsymbol{\Xi}_i$ be a $dn \times dn$ block-diagonal matrix with repeating $\mathbf{S}_i$ along its diagonal. We use transformations $\boldsymbol{\Xi}_i$ to reorient eigenvectors $\boldsymbol{\Xi}_i\boldsymbol{\Phi}_i$, such that we work with eigenspaces $\Omega''_i = (\boldsymbol{\Xi}_i\bar{\mathbf{x}}'_i, \boldsymbol{\Xi}_i\boldsymbol{\Phi}'_i, \boldsymbol{\Lambda}_i, K_i)$, assuming the mean shapes are already translation and scale-independent.

3.2 Determining the Basis of the Fused Space

The fused model eigenvalues and eigenvectors should satisfy $\mathbf{D} = \boldsymbol{\Phi}\boldsymbol{\Lambda}\boldsymbol{\Phi}^T$, where $\mathbf{D}$ is a combined covariance matrix that can be obtained solely from information in models Ω'_i, and where each covariance matrix $\mathbf{D}_i = \boldsymbol{\Phi}_i\boldsymbol{\Lambda}_i\boldsymbol{\Phi}_i^T$ is weighted according to the number of contributed shapes of each of the models. For the complete mathematical derivation of $\mathbf{D}$ we refer to [1].

The fused space is obtained as follows:

1. *Find an eigenspace that spans all models' eigenspaces:* Construct an orthonormal basis set $\boldsymbol{\Gamma}$, that spans all the eigenspaces $\boldsymbol{\Xi}_i\boldsymbol{\Phi}_i$ of the input models. This is done by orthonormalization of the matrix $\mathbf{H}$, defined as $\mathbf{H} = [\boldsymbol{\Xi}_1\boldsymbol{\Phi}'_1, \ldots, \boldsymbol{\Xi}_N\boldsymbol{\Phi}'_N, \delta_{1,2}, \ldots, \delta_{1,N}, \ldots, \delta_{N-1,N}]$, where $\delta_{i,j} = (\boldsymbol{\Xi}_i\bar{\mathbf{x}}'_i - \boldsymbol{\Xi}_j\bar{\mathbf{x}}'_j)$, for $i, j = 1, \ldots, N$, and $j > i$.
2. *Determine an intermediate eigenproblem:* Use $\boldsymbol{\Gamma}$ to derive an intermediate eigenproblem: $\boldsymbol{\Gamma}\mathbf{D}\boldsymbol{\Gamma}^T = \mathbf{R}\boldsymbol{\Lambda}\mathbf{R}^T$, whose solution provides the fused model eigenvalues $\boldsymbol{\Lambda}$ and eigenvectors $\mathbf{R}$ needed to correctly orient basis $\boldsymbol{\Gamma}$.
3. *Compute the fused eigenvector matrix:* Finally, the eigenvectors $\boldsymbol{\Phi}$ of the fused model are obtained from $\boldsymbol{\Phi} = \boldsymbol{\Gamma}\mathbf{R}$.

4 Results

4.1 Quantitative Evaluation

Numerical assessment of the proposed fusion is carried out using a set of 42 cardiac MRI datasets acquired using a GE Signa CVi-HDx 1.5T scanner (General Electric, Milwaukee, USA). The images were manually delineated by an expert clinician using 880 landmarks for the LV and 1200 landmarks for the RV. The following experiment was then carried out:

1. For each $\mathbf{x}_i = \mathbf{x}_{\text{test}}$, define a ground truth PDM Ω_{OM} (Original Model) constructed from $S_{OM} = S_{\text{all}} \setminus \mathbf{x}_{\text{test}}$ (leave-one-out).

2. Partition S_{OM} into randomly selected subsets S_A, S_B, S_C such that $S_A \cap S_B \cap S_C = \emptyset$, $S_A \cup S_B \cup S_C = S_{OM}$ (repeated 100 times).
3. Without loss of generality, choose subsets S_B and S_C and resample all shapes to obtain S'_B and S'_C such that point correspondence between sets S_A, S_B and S_C is lost.
4. Construct the models Ω_{S_A}, $\Omega_{S'_B}$, and $\Omega_{S'_C}$ from S_A, S'_B, and S'_C, respectively.
5. Apply the normalization and fusion algorithm to PDMs Ω_{S_A}, $\Omega_{S'_B}$, and $\Omega_{S'_C}$ to obtain the fused model Ω_{NF}.
6. Compare the accuracy of shape reconstruction achieved by the Ω_{OM}, and Ω_{NF} PDMs by computing the mean point-to-surface distance between $\tilde{\mathbf{x}}_{\text{test},OM}$ and $\tilde{\mathbf{x}}_{\text{test},NF}$, the reconstructions obtained Ω_{OM} and Ω_{NF}, respectively.

Note that all models were built preserving 98% of total variance, and allowing ± 3 standard deviations from the mean. Resampling of S_B and S_C in step 3 was done through two iterations of Loop-subdivision [9] of the original meshes, followed by a centroidal Voronoi diagram algorithm [10] to remove the correspondence.

We performed the experiments on both the LV and the RV shapes. Table 1 summarizes the results as Hausdorff distances, where it can seen that a low reconstruction error is found, which shows that the unified model is similar to the one obtained directly from the original data. The small reconstruction errors are due to the change in surface parameterization in step 3 which inevitably leads to a loss of geometrical information. However, this is kept under 0.5 mm and with a low maximum error for both the LV and RV. This also shows that the proposed fusion can handle models corresponding to different geometrical complexities and anatomical variabilities (*e.g.*, LV and RV).

In order to demonstrate the capability of the technique to fuse multiple models (more than two), Figs. 1(a) and 1(c) illustrate the first mode of variation of three models obtained through random partition of the original dataset (columns 1-3), their fusion (col. 4), and the original shapes model (col. 5). It can be seen visually that the modes obtained for the Ω_{OM} and Ω_{NF} models are almost identical. Also, Figs. 1(b) and 1(d) show the eigenvalues for both the Ω_{OM} and Ω_{NF} of the same models (Figs. 1(a) and 1(c)), organized by descending variance. It can be seen from the eigenvalue curves that our method is able to accurately merge the statistical qualities of the models, regardless of shape parametrization, and in absence of the original data.

Table 1. Point-to-surface distance statistics between reconstructions (Ω_{OM} and Ω_{NF})

Structure	*Mean* $\pm$ *Stdev (mm)*	*Max (mm)*
Left Ventricle	0.319 ± 0.107	0.674
Right Ventricle	0.497 ± 0.157	0.723

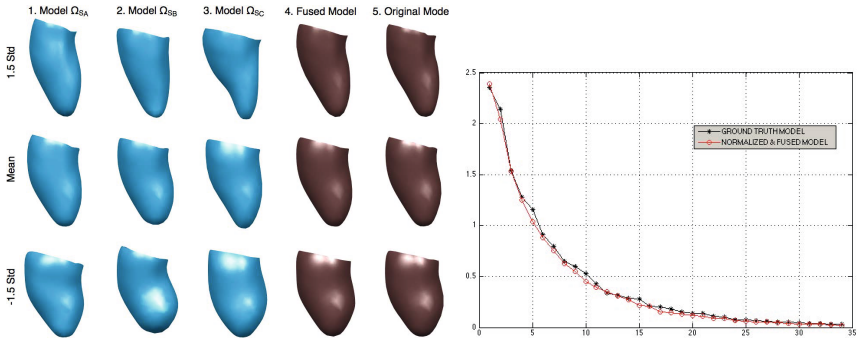

(a) LV first mode: $\mu \pm 1.5\sigma$ (b) LV eigenvalues Ω_{OM}, Ω_{NF}

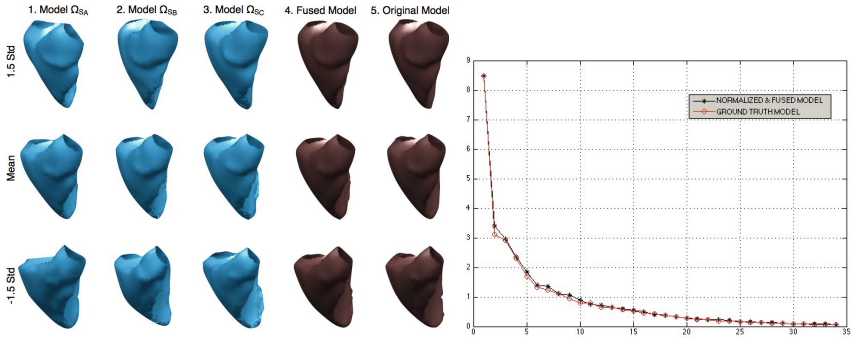

(c) RV first mode: $\mu \pm 1.5\sigma$ (d) RV eigenvalues Ω_{OM}, Ω_{NF}

Fig. 1. First modes of variation of three models Ω_{S_A}, $\Omega_{S'_B}$ and $\Omega_{S'_C}$ (cols. 1, 2 and 3), the fused model Ω_{NF} (col. 4) and original model Ω_{OM} (col. 5); for (a) the LV and (c) the RV. Also shown are the eigenvalues of Ω_{OM} and Ω_{NF} for (b) LV and (d) RV.

4.2 Fusion of MR and CT Models

In this section we evaluate the applicability of the technique by fusing the MR model as described in the previous section with an online PDM constructed from a population of 134 CT datasets[1]. The LV and RV shapes were segmented using an atlas-based approach [7], and described by 1347 and 1748 landmarks, respectively. The original CT datasets are not publicly available and therefore only the obtained means and eigenvectors are used in this validation. The first mode of variation of the MR and CT models are shown in Fig. 2(a) for the LV, and 2(b) for the RV (cols. 1, and 3), respectively. It can be seen that the variability captured by the MR and CT models differ, as can be expected, since they were constructed from different populations. In particular, more localized variation is encoded by the CT statistical model for both the LV and the RV.

We then apply the normalization algorithm to the MR model in order to establish point correspondence, and as observed from comparing Fig. 2 (cols. 1

[1] `http://www.cistib.upf.edu/cistib/index.php/downloads/Statistical_Cardiac_Atl`

and 2), the shape variability is well preserved. This suggests that no significant loss of geometrical and statistical information results from the proposed eigenvector transformation despite the complex geometries involved. We then apply the fusion stage to merge the MR and CT PDMs and the result is displayed in Fig. 2 (col. 4). Visually, it is evident that features of variation found in both the MR and CT models are incorporated into the unified model. In particular, for both the LV and RV, the first mode of variation for the fused model incorporates a pattern describing a transition from the high resolution CT model to the smoother representation found in MRI. In other words, the unified model integrates detailed, as well as smoother surface representations. Inter-modality differences in resolution are encoded as patterns of shape variability into the fused model. In the case of the RV, the regional variation at the valvular level is mostly related to the MR model, while the CT model contributes with the global variation in morphology.

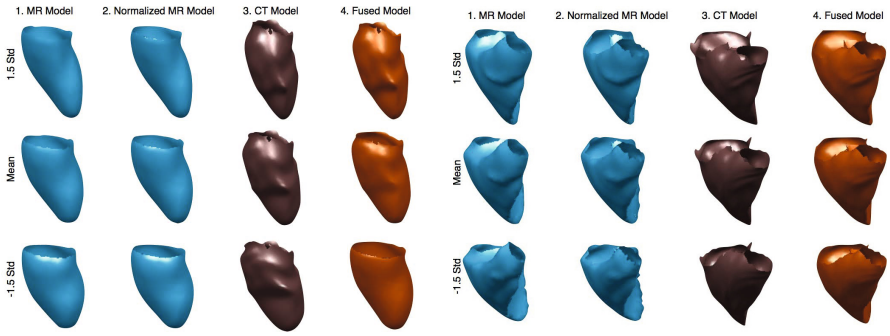

(a) Left Ventricle (b) Right Ventricle

Fig. 2. First modes of variation of two models constructed from imaging modalities MR (col. 1) and CT (col. 3). Also shown are the fused model (col. 4) and the normalized (pre-fusion) MR model (col. 2).

In order to show quantitatively how well the fusion improves shape representation over the individual models, we performed reconstruction of the MR datasets following a leave-one-out scheme based on the MR model alone and then on the fused MR-CT model. The results of the reconstruction are summarized in Table 2. The mean error obtained with the MR-only model for the LV was of 0.953 mm, whereas with the Ω_{NF} it was reduced to 0.465 mm. This is a significant improvement of 51% over the MR model. Also, for the RV, the mean error with the MR-only model is equal to 1.315 mm, while the fusion reduced the reconstruction error to 0.805 mm, which represents a 38% improvement over the MR model. On average the total improvement is 44.5%. These results demonstrate the benefit of the proposed approach to obtain models that generalize better to new instances without the need for any additional datasets but merely by taking advantage of pre-existing PDMs.

Table 2. Reconstruction statistics for the MR dataset (LV, RV), comparing the performance of Ω_{MR}, and a Ω_{MR-CT}

Structure	*Model*	*Mean ± Stdev (mm)*	*Max (mm)*
Left Ventricle	MR model	0.953 ± 0.259	1.800
	MR+CT model	0.465 ± 0.153	0.006
Right Ventricle	MR model	1.315 ± 0.310	2.225
	MR+CT model	0.805 ± 0.197	1.246

5 Conclusion

We present a framework to construct unified multi-modal statistical shape models by normalizing and merging pre-existing PDMs, without the need for the original or any additional datasets. The fused model can be obtained despite differences in the clinical population, delineation protocols and imaging modalities. With the proposed technique, a unique statistical shape model with good specificity and generalization ability is obtained for improved interpretation of medical images.

References

1. Butakoff, C., Frangi, A.F.: A framework for weighted fusion of multiple statistical models of shape and appearance. IEEE Trans. Pattern Anal. Machine Intell. 28(11), 1847–1857 (2006)
2. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models—their training and application. Comput. Vis. Image Understand. 61(1), 38–59 (1995)
3. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. Med. Image Anal. 13(5), 793–808 (2009)
4. Franco, A., Lumini, A., Maio, D.: Eigenspace merging for model updating. In: Proc. Int. Conf. Pattern Recognition (ICPR), vol. 2, pp. 156–159 (2002)
5. Hall, P., Marshall, D., Martin, R.: Merging and splitting eigenspace models. IEEE Trans. Pattern Anal. Machine Intell. 22(9), 1042–1049 (2000)
6. Heimann, T., Meinzer, H.P.: Statistical shape models for 3D medical image segmentation: A review. Med. Image Anal. 13(4), 543–563 (2009)
7. Hoogendoorn, C., Duchateau, N., Sánchez-Quintana, D., Whitmarsh, T., De Craene, M., Sukno, F.M., Lekadir, K., Frangi, A.F.: A high-resolution atlas and statistical model of the human heart from multislice CT. IEEE Trans. Med. Imaging 32(1), 28–44 (2013)
8. Horn, B.K.P.: Closed-form solution of absolute orientation using unit quaternions. J. Opt. Soc. Am. A Optic. Image Sci. Vis. 4(4), 629–642 (1987)
9. Loop, C.: Smooth subdivision surfaces based on triangles. Master's thesis, Department of mathematics, University of Utah, Utah, USA (1987)
10. Valette, S., Chassery, J.M.: Approximated centroidal Voronoi diagrams for uniform polygonal mesh coarsening. Comput. Graph. Forum 24(3), 381–389 (2004)
11. Zhou, X.S., Gupta, A., Comaniciu, D.: An information fusion framework for robust shape tracking. IEEE Trans. Pattern Anal. Machine Intell. 27(1), 115–129 (2005)