



Swansea University
Prifysgol Abertawe



Cronfa - Swansea University Open Access Repository

This is an author produced version of a paper published in :
British Machine Vision Conference

Cronfa URL for this paper:

<http://cronfa.swan.ac.uk/Record/cronfa27022>

Conference contribution :

Palmer, R., Xie, X. & Tam, G. (2015). *Automatic Aortic Root Segmentation with Shape Constraints and Mesh Regularisation*. Xianghua Xie, Mark W. Jones, and Gary K. L. Tam (Ed.), *British Machine Vision Conference*, (pp. 83.1-83.11).

<http://dx.doi.org/10.5244/C.29.83>

This article is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Authors are personally responsible for adhering to publisher restrictions or conditions. When uploading content they are required to comply with their publisher agreement and the SHERPA RoMEO database to judge whether or not it is copyright safe to add this version of the paper to this repository.

<http://www.swansea.ac.uk/iss/researchsupport/cronfa-support/>

Automatic Aortic Root Segmentation with Shape Constraints and Mesh Regularisation

Robert Ieuan Palmer

Department of Computer Science
Swansea University
Swansea, UK

Xianghua Xie

csvision.swan.ac.uk

Gary Tam

Abstract

Fully automated 3D segmentation is not only challenging due to, for instance, ambiguities in appearance, but it is also computationally demanding. We present a fully-automatic, learning-based deformable modelling method for segmenting the aortic root in CT images using a two-stage mesh deformation: a non-iterative boundary segmentation with a statistical shape model for shape constraint, followed by an iterative boundary refinement process. At both stages, we introduce a B-spline mesh regularisation technique to avoid mesh entanglement during deformation. The initialisation of the deformable model is achieved through efficient detection and localisation of the aortic root using marginal space learning, which carries out similarity parameter estimation in an incremental fashion. Quantitative comparisons are carried out against a state-of-the-art deformable model-based approach and an active shape model based segmentation. The proposed method achieves both a lower average mesh error of $1.39 \pm 0.29mm$, and Hausdorff distance of $6.75 \pm 2.05mm$. Compared to these two approaches, it results in much more regularised mesh surfaces with no tangled mesh faces.

1 Introduction

Aortic stenosis is a common heart disease among elderly people where the aortic valve does not fully open. In many cases the disease requires surgical treatment, and a non-invasive procedure called Transcatheter Aortic Valve Implantation (TAVI) has emerged as an alternative procedure for patients that cannot undergo standard open-heart surgery. As with many non-invasive surgical procedures, the role of medical imaging is very important to the success of the procedure, and a full segmentation of the aortic root is essential for patient selection, procedural planning, and post-evaluation [5].

A number of studies have proposed segmenting aortic structures in medical images, but to our knowledge, very few studies have concentrated on segmenting the aortic root. The majority of studies in the literature have segmented the aortic root as part of a larger segmentation procedure, [3, 4, 6, 13, 15], and most have proposed model-based approaches for this purpose, e.g. [6, 13, 15]. There are non-model-based techniques, but they may easily fail to detect boundary between the aortic root and the left ventricular output tract (LVOT) [15]

due to image noise and calcifications. Here, we focus on the discussion of model-based approaches. Waechter *et al.* [13] proposed an automatic segmentation method that first detects cardiac objects using a generalised Hough transform. This was followed with deforming an initial mesh by formulating an energy function that included a data term based on image appearance and a curvature cost. However, detecting and localising complex anatomical structures, such as the aortic root, with poor image quality using Hough transform in 3D is not only error prone but also computationally expensive. In [3], the authors adapted the popular active shape model (ASM) to segment tubular aortic structures. Although the reported mesh error was low, the aortic root was not part of the segmented area in this study. Furthermore, the pose of the initial shape had to be manually placed before segmentation could take place. Elattar *et al.* [4] applied a 3D normalised cut to segment the aortic root and ascending aorta. Their method produced a low mesh error, however the method was evaluated on a small number of images and 3D graph-cut based methods are computationally expensive.

Notably, a recent work by Zheng *et al.* [15] introduces an aortic root segmentation method that is automatic and efficient. The authors proposed a marginal space learning (MSL) technique to drastically reduce the computational burden of 3D exhaustive search of similarity transform with 9-degree freedom. The position, orientation, and anisotropic scale estimations are carried out incrementally, which leads to sub-optimal solutions but allows efficient search in a much denser computational grid that is otherwise not possible with exhaustive search. The initialised mesh template is then deformed using a learning-based boundary detector and generic mesh smoothing in an iterative process. Grbic *et al.* [6] adapted this work to segment valvular heart apparatus in 4D. MSL was used for localisation and a multi-linear active shape model, which decomposes the shape space into temporal and spatial components, was used to carry out non-rigid deformation. However, the deformation method in [15] is simplistic and requires the mesh template to be very close to the object of interest after rigid deformation. The movements of individual vertices on the mesh is independent at the non-rigid deformation stage. Active shape model defined on a coarse computational mesh on the other hand, as we demonstrate in this work, is often too restrictive to allow local deformations that are necessary to achieve accurate segmentation.

In this paper, we propose a fully-automatic method to segment the aortic root in 3D cardiac CT images. The MSL method is adopted to efficiently detect and localise the aortic root at natural voxel resolution. However, we employ more robust features and a different classifier to those used in [15] so that the feature vector length is dramatically reduced but with improved discriminative power. Learning-based boundary descriptors are used to drive a free-form-deformation (FFD) with B-spline mesh regularisation. This allows large mesh deformations which is carried out non-iteratively whilst maintaining its mesh regularity, followed by a prior based regularisation using a statistical shape model with thin-plate-spline warping. The proposed method also allows refined local adjustment to provide flexibility in shape modelling. We show that the proposed method is not only more efficient but also more accurate in segmenting complex anatomical structures with ambiguous image appearance.

The remainder of this paper is organised as follows. Section 2 presents our fully automatic segmentation method. Section 3 presents the results of our proposed method and comparison to the state of the art. The paper concludes in Section 4.




Figure 1: Overview of the proposed automatic aortic root segmentation at the testing stage.

2 Method

The proposed method consists of aligning an initial mesh with an initial aortic root pose estimation, before deforming the mesh towards the object boundary in the image. The initial mesh is the median aortic root mesh from the training set. The estimation of the aortic root pose can be considered as an object detection problem, where the position, orientation, and anisotropic scale parameters need to be determined. Once the initial mesh is aligned with the estimated pose, we implement a two-stage mesh deformation method: non-iterative boundary segmentation followed by iterative boundary refinement with mesh smoothing. The initial non-iterative process consists of deformable segmentation with mesh regularisation followed by an SSM based regularisation with thin-plate-spline warping. This process is capable of large vertex deformations without introducing mesh irregularities. Following this, a boundary refinement iterating a local deformable segmentation with mesh regularisation process is applied. This process again combines the learning-based boundary detector and B-spline based mesh regularisation. Finally, generic mesh smoothing is used to remove any faceted mesh faces. Figure 1 outlines the steps taken at the testing stage for our automatic segmentation method.

The remainder of this section is organised as follows. Section 2.1 describes our MSL based similarity transformation estimation. Section 2.2 describes our initial non-iterative deformable segmentation, consisting of deformable segmentation with mesh regularisation (Section 2.2.1) and an SSM with TPS warping (Section 2.2.2). Finally, Section 2.3 describes the iterative local refinement and final segmentation.

2.1 MSL Based Similarity Transformation Estimation

To initialise the segmentation process, we need a robust method to detect the aortic root position, orientation, and scale. Conventional exhaustive techniques search every possible combination of position, orientation and scale hypotheses for the desired object, which is computationally expensive and impractical in our case (e.g. 9^n number of hypothesis need to be evaluated, where n is the number positions to be tested in a 256^3 volume). MSL significantly speeds up the parameter search process by marginalising the search space and reducing the number of hypotheses. Instead of a single unified exhaustive search, we use three separate pose classifiers. First, all position hypotheses are tested, with only a few (~ 100) of the highest scoring hypotheses being tested further. The highest scoring position-orientation hypotheses (~ 50) are then tested for scale, and the highest scoring hypothesis of pose-orientation-scale is taken as the detection result. The last stage is an exhaustive search but there are only very few hypotheses to evaluate.

Different to [15], where an extensive set of Haar wavelets [12] and a large bank of steerable filters [14] are used to extract features for position estimation and position-orientation estimation, we employ a much smaller feature set that are more discriminative in marginal space classification. For efficiency and accuracy, we train a random forest (RF) classifier, instead of Adaboost in [15], to derive the similarity parameters $(x, y, z, \psi, \phi, \theta, Sx^l, Sy^l, Sz^l)$, where x, y, z denote position, ψ, ϕ, θ denote orientation, and Sx^l, Sy^l, Sz^l are the anisotropic scale parameters in local coordinates. Two sets of features are extracted for the estimators: Haar features and 3D spherical histogram of oriented gradients (SHOG) features [8, 10]. The Haar wavelets are used to highlight the tubular structure of the aorta and thus their features are extracted at rather coarse scales. Haar features ranging in size between 0.25 to 3 times the mean global scale of the training meshes $(\overline{Sx_t^g}, \overline{Sy_t^g}, \overline{Sz_t^g})$, and centred at the position hypothesis are extracted. These features incorporate large appearance differences around the root, as well as local appearance. To detect detailed features at the aortic root, 3D SHOG features are extracted, which highlight areas of high curvature. These areas are heavily present around the aortic hinges in the root, and are not present along the tubular structures of the ascending aorta and aortic arch. This makes these complementary features substantially more discriminative with regards to the aortic root. The Haar feature extraction in position-orientation estimation is aligned locally to a positive hypothesis so that the features are orientation dependent, whereas for the position estimator the features are extracted irrespective of aortic root orientation.

To train the RF classifier, positive position hypotheses are extracted based on the condition that $\max\{|x - x_t|, |y - y_t|, |z - z_t|\} \leq 1$, where (x, y, z) is a position hypothesis and (x_t, y_t, z_t) is the positional ground truth. A positive position-orientation hypothesis satisfies both positive position hypothesis and orientation requirement such that $\max\{|\psi - \psi_t|, |\phi - \phi_t|, |\theta - \theta_t|\} \leq 10^\circ$, where (ψ, ϕ, θ) is an orientation hypothesis and $(\psi_t, \phi_t, \theta_t)$ are the ground truth Euler angles. As for scale, we use the mean scale from the training set for simplicity and efficiency, i.e. $(Sx^l, Sy^l, Sz^l) = (\overline{Sx_t^l}, \overline{Sy_t^l}, \overline{Sz_t^l})$. As shown later, such assumption does not affect the performance as our proposed model can undergo large deformations without losing mesh regularity. The position estimator is applied to all voxels in the image, and the highest scoring 100 hypotheses are tested with the position-orientation estimator. Finally, the highest scoring position-orientation hypothesis is combined with the mean scale to obtain the final nine-element similarity estimation vector $(x, y, z, \psi, \phi, \theta, Sx^l, Sy^l, Sz^l)$.

2.2 Non-Iterative Deformable Segmentation

2.2.1 Deformable Segmentation with Mesh Regularisation

Here, we apply learning-based boundary detectors to localise the aortic root boundary. Once the initial mesh is aligned with the estimated pose, the mesh vertices must be adjusted to fit the object boundary. To avoid vertex entanglement, a search path is defined along the surface normal direction for each vertex. The path coordinate that has the strongest boundary response is then taken as the new vertex position. To this end, we train an RF classifier with two efficient Haar features.

The vertex normals can be straightforwardly determined from the mesh, therefore the Haar templates can be aligned with the normal direction. As a positive sample is a ground truth vertex sitting on the object boundary, one rectangle box will be inside the object whilst the other will be outside, yielding an optimal boundary response. This feature, along with the mean intensity of each rectangle yields a total of just three features.

After boundary detection, we have a set of original mesh vertices V , and a new set of vertex positions V' that are on the boundary. Due to the length of the search path and image noise, it is possible that a vertex may have found an optimal boundary response further down the search path than its neighbouring vertex. This can lead to vertices crossing over, resulting in tangled and overlapping mesh faces. To address this problem, we propose a *B-spline based mesh regularisation* method which estimates a non-rigid transformation between V and V' before performing a free-form-deformation (FFD) on V to fit V' .

The FFD is estimated by warping an underlying voxel lattice controlled by a set of control points which act as parameters of a B-spline. The control points are defined as $\phi_{i,j,k}^h$ of size $[n_x \times n_y \times n_z]$, which are separated by δ , and the FFD is formulated as follows

$$T(x, y, z) = \sum_{l=0}^3 \sum_{m=0}^3 \sum_{n=0}^3 B_l(u)B_m(v)B_n(w)\phi_{i+l,j+m,k+n} \quad (1)$$

where B_l represents the l th basis function of the B-spline, and $i = \lfloor x/n_x \rfloor - 1$, $j = \lfloor y/n_y \rfloor - 1$, $k = \lfloor z/n_z \rfloor - 1$, $u = x/n_x - \lfloor x/n_x \rfloor$, $v = y/n_y - \lfloor y/n_y \rfloor$, and $w = z/n_z - \lfloor z/n_z \rfloor$ [7, 9].

In addition, the non-rigid transformation is estimated in a multi-resolution procedure which is expressed as a summation of FFDs at multiple resolutions H :

$$T^H(x, y, z) = \sum_{h=1}^H T^h(x, y, z) \quad (2)$$

At each mesh resolution h , the voxel lattice is warped by moving the set of control points $\phi_{i,j,k}^h$ which is consequential of δ_h . If δ_0 is the original control point spacing, then at each resolution h , the spacing is defined as $\delta_h = \delta_0/2^h$. The B-spline parameters $\phi_{i,j,k}^h$ are optimised using gradient descent, formulated as

$$E(\phi) = E_s(V', V) + \lambda E_r(T), \quad (3)$$

where E_r is a smoothness cost and λ is a constant that defines the contribution of the smoothness term. E_s is a similarity metric, which in our case is a sum-of-squared-difference (SSD) metric between V and V' . In our experiments the number of mesh resolutions during transformation estimation is set as $H = 3$, and the FFD control point spacing $\delta_0 = 10$.

2.2.2 Statistical-Shape-Model with TPS Warping

Having performed mesh deformations that are capable of large movements, the shape space needs to be constrained so that any deformation is consistent with the shapes in the training set. Applying constraints based on statistical-shape-models (SSMs) is a popular approach to provide prior shape knowledge to a deforming model, and was first proposed by Cootes *et al.* with active-shape-models [2]. It is believed that an SSM built with a complete set of vertices may be too restrictive during mesh deformation, however too few number of correspondences may not adequately describe the shape. We propose building an SSM with a corresponding subset of mesh vertices, called *fiducial points*, across the training set. It is worth noting that this SSM is not applied in an iterative fashion in the proposed method.

Each training mesh is manually labelled with 10 corresponding primary landmark points on distinctive anatomical features. Three of these are aortic valve *hinge* points, three are aortic *commissure* points, three are *sinotubular junction* points (directly beneath the hinge

points), and one point at the centre of the aortic root. Using these primary landmarks, a denser set of *fiducial points* is interpolated automatically. Pairs of primary landmarks are defined and the surface paths between them are determined using Dijkstra’s shortest path algorithm. Fifteen paths are determined in all: six between the hinge points and their nearest commissure points, three between the hinge points and the centre point, three between the hinge points and their corresponding sinotubular junction points, and three between the sinotubular junction points. Six evenly spaced points are interpolated along each path, leaving a total of 90 *fiducial points*. An example of the *fiducial points* can be seen on the resulting meshes in Figure 2.

Principal component analysis (PCA) is then applied to the corresponding landmark points, and an SSM is generated. After applying the SSM, only the *fiducial points* are curved, therefore a further process is required to adjust the remaining vertices appropriately. Thin-plate-spline (TPS) warping [1] allows surface interpolation depending on the physical bending energy of a thin metal plate, whilst keeping certain nodes fixed in location. This makes it ideal for fixing the adjusted *fiducial points* in place while interpolating the remaining vertices around them.

2.3 Local Refinement and Final Segmentation

Following initial segmentation, the mesh vertices are locally refined towards the boundary in an iterative fashion to obtain final segmentation. This process involves repeating our deformable segmentation with mesh regularisation technique (Section 2.2.1) until deformation convergence or the maximum number of iterations is reached. As this process is only finely adjusting the mesh vertices, the boundary detector search path is much smaller here, and decreases with every iteration.

Finally, a generic mesh smoothing technique [11] is used to smooth the mesh surface following the iterative local refinement. The technique proposed by Taubin [11], uses a low-pass filter which removes high curvature variations in the mesh, and is well suited to smoothing faceted mesh faces.

3 Results and Discussion

Experiments were conducted comparing our proposed method with that proposed by Zheng *et al.* [15]. We attempted to recreate as close as we could what was outlined in the original paper. This involved an iterative process of a learning-based boundary detector trained with steerable features and an AdaBoost-based algorithm, followed by a generic mesh smoothing technique [11]. In addition we also compared our method with an active-shape-modelling-based method. Applying ASM in a conventional manner (i.e. an iterative process of a boundary detector followed by SSM application [2]) led to severe entanglement, due to the complexity of those meshes. As a result we decided to include our B-spline mesh regularisation and TPS warping in our implementation, which essentially became an iterative process of our deformable segmentation with mesh regularisation technique (Section 2.2.1). Generic mesh smoothing [11] was also applied after the iterative process. In all cases the initial pose was estimated using our MSL-based approach for fair comparison.

All experiments were conducted on 36 aortic root volumes in a threefold cross-validation fashion. For every training image, the aortic root, including the ascending aorta and aortic arch were manually labelled in CT TAVI images, of size $512 \times 512 \times (500 \sim 800)$, and the




Figure 2: Results of the proposed method. Each row is a result of a different test image. The first column shows the resulting mesh with its corresponding *fiducial points*. The following columns show a selection of image slices with the segmentation results. Green contours are the ground truth, while red contours are the segmentation results.

voxel size was $0.48mm \times 0.48mm \times 0.62mm$. The advanced marching-cube algorithm was used for mesh generation, and the 10 anatomical landmarks were manually labelled. The ascending aorta and aortic arch were then discarded from the mesh below the sinotubular junction to leave a mesh of the aortic root, and a ground truth vector including the root's position, orientation and scale were determined. The average local mesh size was $35mm \times 34mm \times 24mm$.

The segmentation accuracy was quantitatively evaluated using a symmetric point-to-mesh distance error. This was calculated as the mean of the Euclidean distances between the nearest neighbours of the ground truth mesh vertices and the resulting mesh vertices, and vice-versa to make the measurement symmetric. This seems to be the most popular evaluation metric used in the literature [3, 4, 6, 13, 15]. The symmetrical Hausdorff distance was also used to give an idea of any outlying regions in the resulting mesh compared to the ground truth. This was a measurement of the maximum Euclidean distance between nearest neighbours. In a few cases, a small percentage of vertices were stretched below the plane of the sinotubular junction after mesh deformation. In order to ensure that we quantitatively

Table 1: Quantitative comparison.

Method	Mesh Error		Hausdorff Distance	
	(mm space)	(voxel space)	(mm space)	(voxel space)
Zheng <i>et al.</i> [15] Method	1.47 ± 0.51	2.76 ± 0.95	10.79 ± 2.66	19.56 ± 4.37
ASM-based Method	1.86 ± 0.46	3.53 ± 0.87	8.87 ± 2.79	16.77 ± 4.65
Proposed Method	1.39 ± 0.29	2.62 ± 0.54	6.75 ± 2.05	12.72 ± 3.79




Figure 3: Comparison of results, showing mesh appearance and segmentation output. a) Zheng *et al.* [15] method; b) ASM-based method; c) our method; d) ground truth.




Figure 4: Further example segmentation results with our proposed method.

evaluated the segmentation of the aortic root and not parts of the ascending aorta, these vertices were not considered. Qualitative results produced by the proposed method are shown in Figure 2, which also shows the appearance of the resulting meshes. Table 1 compares the quantitative results of our method to that of [15], and the ASM-based method. Figure 3 compares the qualitative segmentations and resulting mesh appearances.

Our method produced highly comparable results to manual segmentations, and we obtained an average point-to-mesh error of $1.39 \pm 0.29\text{mm}$, which translates to ~ 3 voxels. The average Hausdorff distance was $6.75 \pm 2.05\text{mm}$. Figure 2 also highlights that our resulting meshes have no compressed or stretched mesh faces and no tangled mesh edges, leaving a smooth regularised mesh surface.

The boundary refinement approach proposed in [15] produced an average mesh error of 1.47mm. The average Hausdorff distance of 10.79mm, which was notably higher than that produced by our method, suggests that large regions of the resulting meshes deviated significantly from the ground truth. Furthermore, Figure 3 shows that the appearance of the resulting meshes were severely disordered compared to our method, which was due to the absence of any mesh regularisation after vertex boundary detection. Generic mesh smoothing [11] was applied after boundary detection, which was suitable to smooth faceted mesh faces, but could not cope with untangling crossed mesh edges. In reality we attempted to run 15 iterations of this boundary refinement approach, however in all 36 cases the mesh became tangled after three or less iterations. The results presented here for this method were the best ones.

The mesh appearance of the ASM-based results was much better, with all meshes having a reasonably smooth surface. However, this approach produced a significantly higher mesh error than our proposed method. This suggests that applying strong shape constraints on every boundary refinement iteration was too restrictive, and did not allow enough freedom for the mesh to deform towards the boundary. The failure of this method is also due to the fact that vertex-level correspondence in smooth and complex anatomical structures is very difficult to obtain and local deformation can be non-linear.

A further experiment was conducted for the proposed method where the ground truth pose was used coupled with our proposed boundary refinement method, making the segmentation semi-automatic. This produced a point-to-mesh error of $1.19 \pm 0.21\text{mm}$ and a Hausdorff distance of $5.79 \pm 1.76\text{mm}$, suggesting that improvement to the pose estimation can lead to even better segmentation.

The images in Figures 2 and 3 clearly show the complexity of the data. The proposed method achieved much better delineation of the aortic roots. The variations in both geometry and appearance are evident in the further 3D examples shown in Figure 4. The proposed method achieved robust segmentation results.

4 Conclusion

In this paper, we presented a fully-automatic, deformable modelling-based approach for the 3D segmentation of the aortic root in cardiac CT images. The proposed method was very efficient in pose estimation and deformable segmentation. The learning based features were efficient and robust in defining aortic roots and the proposed intrinsic mesh regularisation allowed large deformations to take place. We showed that statistical shape regularisation was effective at the coarser scale and nonlinear local refinement was necessary to archive accurate segmentation. Both qualitative and quantitative results showed that the proposed method outperformed the state of the art, as well as the classical ASM.

References

- [1] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *Trans. PAMI*, 11(6):567–585, 1989.
- [2] T. F. Cootes, D. H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
- [3] M. de Bruijne, B. van Ginneken, M A. Viergever, and W. J. Niessen. Adapting active shape models for 3d segmentation of tubular structures in medical images. In *Proc. Information Processing in Medical Imaging*, pages 136–147, 2003.
- [4] A. Elattar, M. E. M. Wiegerinck, R. N. Planken, E. vanbavel, H. C. van Assen, and J. Baan Jr. Automatic segmentation of the aortic root in ct angiography of candidate patients for transcatheter aortic valve implantation. *Med. Biol. Eng. Comput.*, 52:611–618, 2014.
- [5] M. Gessat, T. Frauenfelder, L. Altwegg, J. Grunenfelder, and V. Falk. Transcatheter aortic valve implantation: Role of imaging. *Aswan Heart Centre Science and Practice Series*, 1(3), 2011.
- [6] S. Grbic, R. Ionasec, D. Vitanovski, I. Voigt, B. Georgescu, N. Navab, and D. Comaniciu. Complete valvular heart apparatus model from 4d cardiac ct. *Med. Image Anal.*, 16:1003–1014, 2012.
- [7] S. Lee, G. Wolberg, and S. Y. Shin. Scattered data interpolation with multilevel b-splines. *Trans. Visualization and Computer Graphics*, 183(3), 1997.
- [8] K. Liu, H. Skibbe, T. Schmidt, T. Belin, K. Palme, and O. Ronneberger. 3d rotation-invariant description from tensor operation on spherical hog field. In *Proc. of the BMVC*, 2011.
- [9] D. Rueckert, L. I. Sonoda, C. Hayes, D. Mill, O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: Application to breast mr images. *Trans. Med. Imaging*, 18(8):712–721, 1999.
- [10] H. Skibbe, M. Reisert, and H. Burkhardt. Shog - spherical hog descriptors for rotation invariant 3d object detection. In *Proc. DAGM*, 2011.

- [11] G. Taubin. A signal processing approach to fair surface design. In *Proc. SIGGRAPH '95*, pages 351–358, 1995.
- [12] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 511–518, 2001.
- [13] I. Waechter, R. Kneser, G. Korosoglou, J. Peters, N. H. Bakker, R. v. d. Boomen, and J. Weese. Patient specific models for planning and guidance of minimally invasive aortic valve implantation. In *Proc. Int'l Conf. Medical Image Computing and Computer Assisted Intervention*, pages 526–533, 2010.
- [14] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu. Four-chamber heart modeling and automatic segmentation of for 3d cardiac ct volumes using marginal space learning and steerable features. *Trans. Med. Imaging*, 27(11):1668–1681, 2008.
- [15] Y. Zheng, M. John, R. Liao, A. Nottling, J. Boese, J. Kempfert, T. Walther, G. Brockmann, and D. Comaniciu. Automatic aorta segmentation and valve landmark detection in c-arm ct for transcatheter aortic valve implantation. *Trans. Med. Imaging*, 31:2307–2321, 2012.