

Spatiotemporal Segmentation and Modeling of the Mitral Valve in Real-Time 3D Echocardiographic Images

Alison M. Pouch¹(✉), Ahmed H. Aly², Eric K. Lai³,
Natalie Yushkevich¹, Rutger H. Stoffers⁴, Joseph H. Gorman IV³,
Albert T. Cheung⁵, Joseph H. Gorman III³, Robert C. Gorman³,
and Paul A. Yushkevich¹

¹ Department of Radiology, University of Pennsylvania, Philadelphia, PA, USA
pouch@mail.med.upenn.edu

² Department of Bioengineering,
University of Pennsylvania, Philadelphia, PA, USA

³ Gorman Cardiovascular Research Group, University of Pennsylvania,
Philadelphia, PA, USA

⁴ University of Groningen, Groningen, The Netherlands

⁵ Department of Anesthesiology, Perioperative and Pain Medicine,
Stanford University Medical Center, Stanford, CA, USA

Abstract. Transesophageal echocardiography is the primary imaging modality for preoperative assessment of mitral valves with ischemic mitral regurgitation (IMR). While there are well known echocardiographic insights into the 3D morphology of mitral valves with IMR, such as annular dilation and leaflet tethering, less is understood about how quantification of valve *dynamics* can inform surgical treatment of IMR or predict short-term recurrence of the disease. As a step towards filling this knowledge gap, we present a novel framework for 4D segmentation and geometric modeling of the mitral valve in real-time 3D echocardiography (rt-3DE). The framework integrates multi-atlas label fusion and template-based medial modeling to generate quantitatively descriptive models of valve dynamics. The novelty of this work is that temporal consistency in the rt-3DE segmentations is enforced during both the segmentation and modeling stages with the use of groupwise label fusion and Kalman filtering. The algorithm is evaluated on rt-3DE data series from 10 patients: five with normal mitral valve morphology and five with severe IMR. In these 10 data series that total 207 individual 3DE images, each 3DE segmentation is validated against manual tracing and temporal consistency between segmentations is demonstrated. The ultimate goal is to generate accurate and consistent representations of valve dynamics that can both visually and quantitatively provide insight into normal and pathological valve function.

Keywords: Medial axis representation · Deformable modeling · Multi-atlas segmentation · Mitral valve · 3D echocardiography

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-66182-7_85](https://doi.org/10.1007/978-3-319-66182-7_85)) contains supplementary material, which is available to authorized users.

1 Introduction

Ischemic mitral regurgitation (IMR) is a condition in which the mitral valve becomes incompetent due to ischemic disease of the left ventricle. Transesophageal real-time 3D echocardiography (rt-3DE) is the standard of care for preoperative assessment of IMR in many major medical centers. Given the high recurrence rate of moderate to severe IMR after valve repair surgery (33% in the first year [1]), there is a critical need to identify patients at elevated risk of recurrence, for whom alternative surgical approaches (complete valve replacement) may provide better long-term outcomes. Recently, 3D morphological features extracted from intra-operative rt-3DE of the diseased valve were shown to predict post-repair IMR recurrence [2]. Static features such as leaflet tethering angles demonstrate promising predictive value but only take advantage of rt-3DE data at a single frame in the cardiac cycle. Since the etiology of IMR is functional left ventricular disease, it is likely that features of mitral valve *dynamics* could potentially bolster recurrence prediction. Towards this end, we have developed a spatiotemporal image analysis algorithm that can be used to quantitatively describe mitral valve shape and motion in normal and IMR subjects.

Several studies have explored 4D mitral valve segmentation and modeling. In [3], the mitral valve is first segmented in a diastolic 3DE image using graph cuts and the resulting mesh evolves under a set of data-driven and regularization forces to capture valve shape in the other images in the series. Application-specific regularization includes forces that tether the leaflet free edges into the left ventricle, enforce a physiological strain on the leaflets, and prevent collision along the leaflet free margin. In [4], a landmark-based deformable model of the mitral and aortic valves is first initialized in two frames in the cardiac cycle. For all other frames, landmark motion is predicted using manifold learning and clustering and then updated by an optical flow tracker and boundary detection. The result is a full 4D surface model of both valves. A disadvantage of these methods is they do not generate volumetric segmentations that distinguish the atrial and ventricular surfaces or capture locally varying thickness of the leaflets. In this work, we present a 4D segmentation method that combines the benefits of atlas-based segmentation and deformable modeling with medial axis representation, which produces topologically consistent *volumetric models* of the mitral valve. The former uses expert knowledge of valve image appearance to label the target rt-3DE series and the latter establishes geometric correspondences across subjects and facilitates statistical shape analysis. Temporal coherence is enforced in the segmentation stage by employing a groupwise implementation of multi-atlas label fusion, which updates segmentations of individual 3DE frames using information from its temporal neighbors in the same image series. Temporal consistency is also enforced during the modeling phase by using the Kalman filter to smooth the trajectories of deformable models fitted to the label fusion results. In contrast to previous work, this method is validated against label maps of the valve at every frame in the rt-3DE series and volumetric models provide information about locally varying leaflet thickness. Moreover, it demonstrates a comparison of functional valve measurements in subjects with normal mitral valve function and in subjects with severe IMR.

2 Materials and Methods

2.1 Image Data and Manual Segmentation

The iE33 imaging platform (Philips Medical Systems, Andover, MA) was used to acquire EKG-gated rt-3DE images of the mitral valve from 10 patients: 5 with normal mitral valve structure and function and 5 with severe IMR. Images were obtained over four consecutive cardiac cycles with a 2–7 MHz matrix-array transesophageal transducer. Each subject’s rt-3DE image series consisted of 10–37 frames (temporal resolution of 18–40 Hz) showing the mitral valve over one cardiac cycle beginning at early systole. The images were exported in Cartesian format with nearly isotropic resolution ranging from 0.4 to 0.8 mm. All data series, totaling 207 individual 3DE images, were manually segmented in ITK-SNAP, an interactive medical image segmentation tool. An expert observer separately labeled the anterior and posterior leaflets in each 3DE image. The manual segmentations served as atlases for multi-atlas label fusion and as references for cross-validation of automated segmentation.

2.2 Automated Image Analysis

Automated 4D mitral valve segmentation requires a reference atlas set consisting of manually traced images at each phase of the cardiac cycle, and a deformable model of the valve. The only manual interaction required to segment a new unseen target image series is the identification of several temporal and physical landmarks in the target series, described below. First, the series is segmented using a groupwise implementation of multi-atlas label fusion. Then, a mitral valve model in the form of a medial axis representation is deformed to the groupwise segmentation result with enforcement of temporal coherence. The combination of these techniques facilitates standardized automated measurements of mitral valve dynamics in rt-3DE data.

Cardiac Phase Detection and Landmark Initialization. Given a “target” rt-3DE image series to segment, a user first identifies two frames in the cardiac cycle: mid-systole and mitral valve opening. The user then identifies five landmarks in the mid-systolic image: the anterior aortic peak of the annulus, two commissures, the midpoint of the posterior annulus, and the midpoint of the coaptation line.

Groupwise Multi-atlas Segmentation. Multi-atlas label fusion (MALF) uses a set of expert-labeled images, referred to as atlases, to generate individual candidate segmentations of a target image. Each candidate segmentation is created by performing deformable registration between the target and atlas image, and then warping the atlas labels to the target image space. Registration is initialized with the user-identified landmarks described above. Since the candidate segmentations may not be accurate on their own, they are merged into a higher quality segmentation using a consensus-seeking strategy called label fusion, which uses spatially-varying intensity-based weighted voting to combine the candidates into a consensus label map [5]. In order to extend 3D MALF to the spatiotemporal domain, we use a *groupwise* adaptation of MALF [6]. First, each 3DE image in the series is independently segmented with MALF using other subjects’ atlases. Once an initial segmentation is generated for each frame,

the series of segmentations is iteratively updated using both the original reference atlases and the MALF segmentations of neighboring 3DE images as “pseudo” reference atlases. By weighting the “pseudo” atlases more than the original atlas set, the iteratively updated segmentations become more coherent with one another, capturing the similarity in the segmentations of the same structure moving over time.

A challenge of using groupwise MALF (GW-MALF) for 4D valve segmentation is that performing within-subject deformable registration between the valve at systole and diastole often produces poor results because the temporal resolution of image acquisition is not sufficient to provide many intermediate images of the leaflets during rapid valve opening. To overcome this challenge, GW-MALF is performed separately in overlapping cardiac phases: systole, transition, and diastole. Figure 1 illustrates that each image in a sample target series $\{I_n\}_{n=1}^8$ is registered to a set of atlases $\{A_f^s\}$ drawn from the same phase and a neighboring phase of the cardiac cycle. Here, f denotes a set of frame numbers and s denotes a set of subject identifiers. Each image I_n is likewise registered to its neighboring images in the same series. In the first round of MALF, each I_n is segmented using only the reference atlas set $\{A_f^s\}$. In the subsequent groupwise iterations, the segmentation of I_n is updated using the segmentations of I_{n-1} and I_{n+1} as pseudo atlases, along with the original atlas set $\{A_f^s\}$.

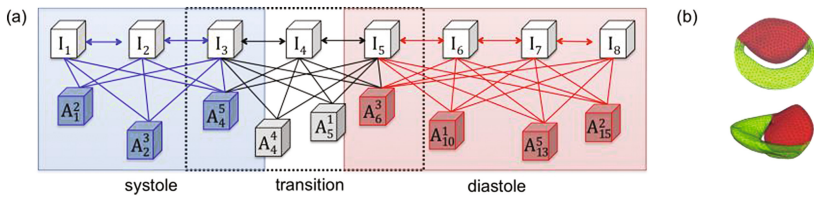


Fig. 1. (a) Atlas assignment based on cardiac phase in GW-MALF. (b) Deformable template fitted to the groupwise segmentation results. Medial mesh (top) and boundary mesh (bottom) shown with the anterior leaflet in red and posterior leaflet in green.

Deformable Medial Modeling. Once a series of valve segmentations is obtained, a deformable model in the form of a continuous medial representation (cm-rep) [7] is warped to the GW-MALF results. This step imposes a fixed topology on the final segmentation, establishes correspondences on different patients’ segmentations, and facilitates automated morphological measurement. Cm-rep describes the valve in terms of its medial axis geometry and is parameterized by $\{\mathbf{m}, R\} \in \mathbb{R}^3 \times \mathbb{R}^+$, where \mathbf{m} is a continuous medial manifold and R is a radial thickness field defined over the medial manifold. Numerically, the medial manifold and leaflet surfaces are represented by triangulated meshes that can be sequentially Loop subdivided to a desired vertex density. The template medial mesh, shown in Fig. 1b, is generated using a method similar to that described in [8] and is deformed to maximize its overlap with the target segmentation while imposing soft regularization constraints to ensure mesh regularity and validity of medial axis geometry as described in [8].

Kalman Filtering. Dynamic measurements derived from cm-reps that are independently fitted to each frame in the output of GW-MALF segmentation are inherently noisy. To impose spatiotemporal smoothness on the model series, a Kalman filter [9] is used to recursively estimate the true state of the system (the valve’s configuration and motion), which cannot be observed directly but can be estimated by combining noisy observations with the predictions of kinematic equations. Let \mathbf{x}_t denote the true system state at time t and let $\hat{\mathbf{x}}_{t|t-1}$ be the *a priori* estimate of the system’s state at time t with variance $\mathbf{P}_{t|t-1}$. The vector $\hat{\mathbf{x}}_{t|t-1}$ is determined by applying kinematic equations to the state estimate at $t - 1$. Let \mathbf{z}_t be observations of the system made at time t . The goal of the Kalman filter is to combine the *a priori* estimate $\hat{\mathbf{x}}_{t|t-1}$ and measurements \mathbf{z}_t to generate an *a posteriori* prediction of the state at time t , denoted by $\hat{\mathbf{x}}_{t|t}$.

In this work, the sought “true” system state is defined by $\mathbf{x}_t = \begin{bmatrix} x_t \\ \dot{x}_t \end{bmatrix}$, where x_t denotes vertex coordinates that comprise the medial surface \mathbf{m} and \dot{x}_t denotes their velocities at time t . We treat the cm-rep models fitted independently to different rt-3DE frames as the noisy observations \mathbf{z}_t . With a medial mesh consisting of n_p nodes, the total number of state variables is $n_v = 6n_p$. The recursive filter begins by initializing $\hat{\mathbf{x}}_{0|0} = \mathbf{z}_0$ as the medial node coordinates and velocities of the medial model fitted to the first frame of the rt-3DE series. Its covariance is initialized as $\mathbf{P}_{0|0} = \mathbf{0}_{n_v}$. For each subsequent frame, the state is predicted to evolve according to $\hat{\mathbf{x}}_{t|t-1} = \mathbf{F}_t \hat{\mathbf{x}}_{t-1|t-1} + \mathbf{w}_t$, where \mathbf{F}_t is a transition matrix that displaces the nodes with constant velocity and $\mathbf{w}_t \sim \mathcal{N}(0, \mathbf{Q}_t)$ is zero-mean Gaussian distributed process noise with covariance \mathbf{Q}_t . The *a priori* covariance is updated as $\mathbf{P}_{t|t-1} = \mathbf{F}_t \mathbf{P}_{t-1|t-1} \mathbf{F}_t^T + \mathbf{Q}_t$. Noisy measurements of the state at time t are modeled as $\mathbf{z}_t = \mathbf{H}_t \mathbf{x}_t + \mathbf{v}_t$, where \mathbf{H}_t is the identity matrix \mathbf{I}_{n_v} and $\mathbf{v}_t \sim \mathcal{N}(0, \mathbf{R}_t)$ is zero-mean Gaussian distributed measurement noise with covariance \mathbf{R}_t . Assuming that the process noise and measurement errors are uncorrelated, the *a posteriori* state estimate is given by:

$$\begin{aligned} \hat{\mathbf{x}}_{t|t} &= \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t (\mathbf{z}_t - \mathbf{H}_t \hat{\mathbf{x}}_{t|t-1}) \\ \mathbf{P}_{t|t} &= \mathbf{P}_{t|t-1} - \mathbf{K}_t \mathbf{H}_t \mathbf{P}_{t|t-1} \end{aligned} \quad (1)$$

where \mathbf{K}_t is the Kalman gain defined by $\mathbf{K}_t = \mathbf{P}_{t|t-1} \mathbf{H}_t^T (\mathbf{H}_t \mathbf{P}_{t|t-1} \mathbf{H}_t^T + \mathbf{R}_t)^{-1}$. In this study, $\mathbf{R}_t = \mathbf{I}_{n_v}$ remains constant over time. Since the state evolution equations assume a constant velocity, which does not accurately reflect the nonlinear motion of the valve, we employ a time-varying process noise covariance that increases in magnitude at t_{trans} , the transition time in the cardiac cycle when the valve opens:

$$\mathbf{Q}_t = \alpha(t) \mathbf{I}_{n_v}, \quad \text{where } \alpha(t) = \begin{cases} 0.1, & t < t_{trans} \\ 1.0, & t \geq t_{trans} \end{cases}. \quad (2)$$

The time-varying \mathbf{Q}_t allows the model fittings to have a stronger influence on the state estimation when the valve accelerates. Once the Kalman filter is applied to the entire series of medial models, measurements of valve dynamics are computed.

3 Results

Valve segmentation and modeling were evaluated in a leave-one-out experiment using manual image segmentation for comparison. For all patients' rt-3DE data, each frame in the target image series was first automatically segmented using an atlas set of 8 randomly selected frames from other subjects in the same group (normal or IMR, a diagnosis that is known prior to preoperative imaging). Six of these atlas frames were from the same phase of the cardiac cycle as the target image and 2 were from the transition phase. For target images in the transitional phase, the atlases consisted of 1 from systole, 6 transitional, and 1 from diastole. The series of segmentations were updated with two groupwise iterations, wherein the intra-subject pseudo atlases were weighted 4 times the atlases from other subjects. A representative segmentation is shown in Fig. 2, and the accuracy of GW-MALF and model fitting relative to manual tracing are given in Table 1.

Table 1. Segmentation accuracy based on the symmetric mean boundary error (BE) metric.

| | Mean BE (mm) | | 95 th percentile BE (mm) | |
|-------------------------|---------------|---------------|-------------------------------------|---------------|
| | Systole | Diastole | Systole | Diastole |
| MALF | 0.5 ± 0.1 | 0.5 ± 0.1 | 2.4 ± 1.6 | 2.5 ± 0.7 |
| GW-MALF | 0.5 ± 0.1 | 0.5 ± 0.1 | 2.5 ± 1.7 | 2.6 ± 0.7 |
| Model fitted to GW-MALF | 0.6 ± 0.2 | 0.6 ± 0.1 | 2.8 ± 1.8 | 2.7 ± 0.8 |

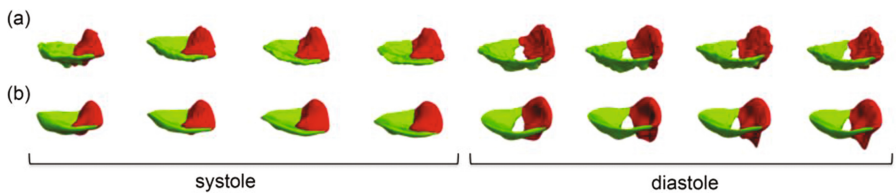


Fig. 2. (a) Manual segmentation and (b) model fitting to the GW-MALF results for a subject with normal mitral valve function. The anterior leaflet is red, posterior leaflet green.

To evaluate the smoothing effect of the groupwise updates and Kalman filter, the dynamic measurements listed in Table 2 were computed from cm-reps fitted to the MALF segmentations and from the Kalman filtered models fitted to the GW-MALF results. Shown in Fig. 3a, a cubic polynomial was fitted to each measurement curve and the residual between the measurements (black) and the cubic polynomial (red) was computed. The mean differences in the measurements before and after groupwise updates and Kalman filtering were calculated to assess whether smoothing introduced a significant bias. An embedded video illustrating the medial mesh before and after groupwise updates and Kalman smoothing is shown in Fig. 3b.

Table 2. Root mean square residual (RMSR) in dynamic measurements relative to their best-fit cubic polynomials before and after GW + Kalman (groupwise iterations and Kalman filtering). Measurement bias in GW + Kalman relative to MALF is given.

| Measurement | RMSR (MALF) | RMSR (GW + Kalman) | <i>p</i> -value | Bias |
|------------------------|---------------|--------------------|----------------------|------------------|
| Annular bending angle | 0.6 ± 0.2 deg | 0.3 ± 0.2 deg | 1 × 10 ⁻³ | -0.17 ± 0.19 deg |
| Septolateral diameter | 0.8 ± 0.3 mm | 0.4 ± 0.2 mm | 4 × 10 ⁻⁶ | -0.17 ± 0.26 mm |
| Intercommissural width | 3.8 ± 2.0 mm | 2.4 ± 1.3 mm | 8 × 10 ⁻³ | -0.23 ± 1.14 mm |

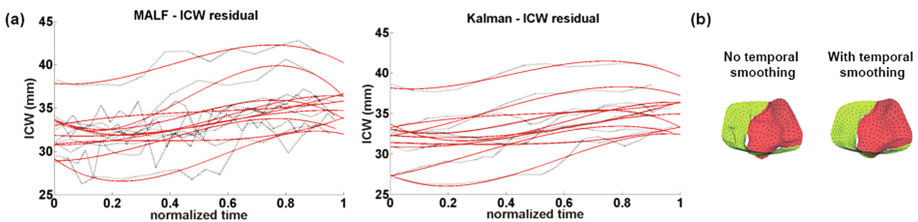


Fig. 3. (a) Cubic polynomials (red) fitted to dynamic intercommissural width measurements (black) before and after groupwise iterations and Kalman filtering. (b) Embedded video of a valve with severe IMR before (left) and after (right) groupwise iterations and Kalman filtering.

For the normal and IMR subgroups, the septolateral diameter, annular bending angle, and intercommissural width were plotted as a function of time, normalized to the cardiac cycle. Bending angle is the angle between planes fitted through the anterior and posterior mitral annuli and is therefore a measure of annular bend, with a larger angle signifying a flatter annulus. Figure 4 shows the IMR measurements in blue and normal valve measurements in red. These measurements were significantly different ($p < 0.01$) over the entire cardiac cycle and underscore differences in annular dynamics in the normal and IMR patients.

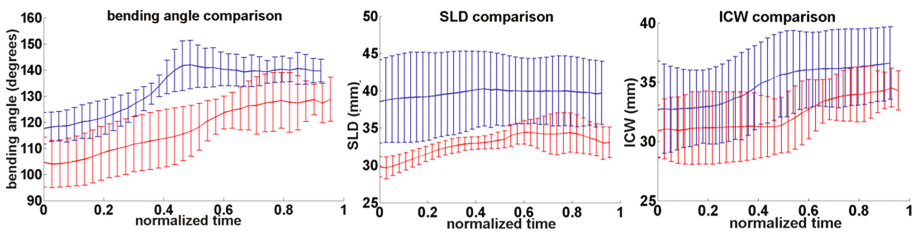


Fig. 4. Comparison of septolateral diameter (SLD), bending angle, and intercommissural width (ICW) in subjects with normal mitral valves (red) and severe IMR (blue) over one cardiac cycle. The error bars represent one standard deviation from the mean.

4 Discussion

The proposed 4D mitral valve segmentation pipeline combines the attractive properties of multi-atlas segmentation and deformable medial modeling to create descriptive, quantifiable representations of mitral valve dynamics from rt-3DE images. Exemplified in Fig. 2, the output models are visually consistent with manual tracing throughout the cardiac cycle. In addition, the accuracy metrics presented in Table 1 demonstrate that in each cardiac phase, automated segmentation performs on par with user-guided segmentation (0.40 ± 0.32 mm [10]), and compares favorably to other automated 4D segmentation algorithms (0.59 ± 0.49 mm [3] and 1.54 ± 1.17 mm [4]). The pipeline's novelty is its extension of MALF and cm-rep into the spatiotemporal domain by enforcing temporal consistency during both the segmentation and modeling stages. The effect of temporal smoothing is illustrated in Table 2 and in Fig. 3, which highlight that Kalman-filtered cm-reps fitted to the GW-MALF results generate dynamic measurements that are less noisy than those fitted to MALF segmentations without temporal smoothing. Another important observation in this study is the identification of differences in valve dynamics in patients with normal mitral valve function and in patients with severe IMR. Figure 4 demonstrates that the normal and IMR valves in this study differ significantly in annular function. Increased SLD, ICW, and bending angle in IMR patients suggest annular dilatation and flattening, which are common features of IMR. Moreover, Fig. 4 suggests that annular dynamics, particularly change in size along on the septolateral diameter, may be less pronounced over the cardiac cycle in IMR patients relative to normal patients.

While the proposed pipeline only requires manual identification of two key frames in the cardiac cycle and five valve landmarks at mid-systole, future work will focus on full automation of the algorithm. In this study, the use of a Kalman filter with a constant velocity model and time-varying process noise provide meaningful time series smoothing. A physical model that better captures the valve's non-linear dynamics could further improve temporal coherence in the output 4D model. Finally, although the dynamic measurements presented in this work suggest significant differences in annular function in IMR patients, studies with a larger patient population are needed to verify these findings. The development of this 4D automated segmentation and modeling tool is a step towards better understanding the functional mechanisms underlying IMR progression and post-operative disease recurrence.

Acknowledgement. This work was funded by the National Institutes of Health: EB017255 and HL103723.

References

1. Acker, M.A., Parides, M.K., Perrault, L.P., Moskowitz, A.J., Gelijns, A.C., Voisine, P., Smith, P.K., Hung, J.W., Blackstone, E.H., Puskas, J.D., Argenziano, M., Gammie, J.S., Mack, M., Ascheim, D.D., Bagiella, E., Moquete, E.G., Ferguson, T.B., Horvath, K.A., Geller, N.L., Miller, M.A., Woo, Y.J., D'Alessandro, D.A., Ailawadi, G., Dagenais, F., Gardner, T.J., O'Gara, P.T., Michler, R.E., Kron, I.L.: Mitral-Valve repair versus replacement for severe ischemic mitral regurgitation. *N. Engl. J. Med.* **370**(1), 23–32 (2013)
2. Bouma, W., Lai, E.K., Levack, M.M., Shang, E.K., Pouch, A.M., Eperjesi, T.J., Plappert, T. J., Yushkevich, P.A., Mariani, M.A., Khabbaz, K.R., Gleason, T.G., Mahmood, F., Acker, M.A., Woo, Y.J., Cheung, A.T., Jackson, B.M., Gorman III, J.H., Gorman, R.C.: Preoperative three-dimensional valve analysis predicts recurrent ischemic mitral regurgitation after mitral annuloplasty. *Ann. Thorac. Surg.* **101**(2), 567–575 (2016)
3. Schneider, R.J., Tenenholtz, N.A., Perrin, D.P., Marx, G.R., del Nido, P.J., Howe, R.D.: Patient-specific mitral leaflet segmentation from 4D ultrasound. *Med. Image Comput. Comput. Assist. Interv.* **14**, 520–527 (2011)
4. Ionasec, R.I., Voigt, I., Georgescu, B., Wang, Y., Houle, H., Vega-Higuera, F., Navab, N., Comaniciu, D.: Patient-specific modeling and quantification of the aortic and mitral valves from 4-D cardiac CT and TEE. *IEEE Trans. Med. Imaging* **29**, 1636–1651 (2010)
5. Wang, H., Suh, J.W., Das, S., Pluta, J., Craige, C., Yushkevich, P.: Multi-atlas segmentation with joint label fusion. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(3), 611–623 (2013)
6. Wang, H., Yushkevich, P.A.: Groupwise segmentation with multi-atlas joint label fusion. *Med. Image Comput. Comput. Assist. Interv.* **16**(1), 711–718 (2013)
7. Yushkevich, P.A., Zhang, H., Gee, J.C.: Continuous medial representation for anatomical structures. *IEEE Trans. Med. Imaging* **25**(12), 1547–1564 (2006)
8. Pouch, A.M., Wang, H., Takabe, M., Jackson, B.M., Gorman 3rd, J.H., Gorman, R.C., Yushkevich, P.A., Sehgal, C.M.: Fully automatic segmentation of the mitral leaflets in 3D transesophageal echocardiographic images using multi-atlas joint label fusion and deformable medial modeling. *Med. Image Anal.* **18**, 118–129 (2014)
9. Kalman, R.E.: A new approach to linear filtering and prediction problems. *J. Basic Eng.* **82** (SeriesD), 35–45 (1960)
10. Siefert, A.W., Icenogle, D.A., Rabbah, J.P., Saikrishnan, N., Rossignac, J., Lerakis, S., Yoganathan, A.P.: Accuracy of a mitral valve segmentation method using J-Splines for real-time 3D echocardiography data. *Ann. Biomed. Eng.* **41**(6), 1258–1268 (2013)