Real-Time Active Shape Models for Segmentation of 3D Cardiac Ultrasound

Jøger Hansegård[1], Fredrik Orderud[2], and Stein I. Rabben[3]

[1] University of Oslo, Norway
jogerh@ifi.uio.no
[2] Norwegian University of Science and Technology, Norway
fredrik.orderud@idi.ntnu.no
[3] GE Vingmed Ultrasound, Norway
stein.rabben@med.ge.com

Abstract. We present a fully automatic real-time algorithm for robust and accurate left ventricular segmentation in three-dimensional (3D) cardiac ultrasound. Segmentation is performed in a sequential state estimation fashion using an extended Kalman filter to recursively predict and update the parameters of a 3D Active Shape Model (ASM) in real-time. The ASM was trained by tracing the left ventricle in 31 patients, and provided a compact and physiological realistic shape space. The feasibility of the proposed algorithm was evaluated in 21 patients, and compared to manually verified segmentations from a custom-made semi-automatic segmentation algorithm. Successful segmentation was achieved in all cases. The limits of agreement (mean$\pm$1.96SD) for the point-to-surface distance were 2.2$\pm$1.1 mm. For volumes, the correlation coefficient was 0.95 and the limits of agreement were 3.4$\pm$20 ml. Real-time segmentation of 25 frames per second was achieved with a CPU load of 22%.

1 Introduction

Left ventricular (LV) volumes and ejection fraction (EF) are among the most important parameters in diagnosis and prognosis of heart diseases. Recently, real-time three-dimensional (3D) echocardiography was introduced. Segmentation of the LV in 3D echocardiographic data has become feasible, but due to poor image quality, commercially available tools are based upon a semi-automatic approach [1,2]. Furthermore, most reported methods are using iterative and computationally expensive fitting schemes. These factors make real-time segmentation in 3D cardiac ultrasound challenging.

Prior work by Blake *et al.* [3,4] and Jacob *et al.* [5,6], have shown that a state estimation approach is well suited for real-time segmentation in 2D imagery. They used a Kalman filter, which requires only a single iteration, to track the parameters of a trained deformable model based on principal component analysis (PCA), also known as Active Shape Models (ASMs) [7]. ASMs can be trained on manually traced LV contours, resulting in a sub-space of physiologically probable shapes, effectively exploiting expert knowledge of the LV anatomy

W.G. Kropatsch, M. Kampel, and A. Hanbury (Eds.): CAIP 2007, LNCS 4673, pp. 157–164, 2007.

and function. For segmentation of 3D cardiac data, Van Assen *et al.* [8] introduced the 3D ASM. However, there are to our knowledge no reports of real-time implementations of 3D ASMs.

Based on the work in [3,4,5,6], real-time LV segmentation of 3D cardiac ultrasound was recently introduced by Orderud [9]. He used an extended Kalman filter for robust tracking of a rigid ellipsoid LV model. Later this framework has been extended to use a flexible spline-based LV model coupled with a global pose transform to improve local segmentation accuracy [10]. However, expert knowledge of LV anatomy could not be modeled directly.

To utilize expert knowledge of LV anatomy during segmentation, we propose to use a 3D ASM for real-time segmentation of 3D echocardiograms, by extending the framework described in Orderud [9]. The 3D ASM, trained on LV shapes traced by an expert, gives a compact deformable model which is restricted to physiologically realistic shapes. This model is fitted to the target data in real-time using a Kalman filter. The feasibility of the algorithm is demonstrated in 21 patients, where we achieve real-time segmentation of the LV shape, and instantaneous measurements of LV volumes and EF.

2 Shape Model

A set of 496 triangulated LV training meshes were obtained from 31 patients using a custom-made segmentation tool (GE Vingmed Ultrasound, Norway). The training tool provides manual editing capabilities. When necessary, the user hence did manual editing of the segmentation to make it equivalent to manual tracing.

Building the ASM requires pair-wise point correspondence between shapes from different patients [7,8]. We developed a reparametrization algorithm for converting triangulated LV training shapes into quadrilateral meshes. This algorithm produced meshes with 15 longitudinal and 20 circumferential segments, with vertices approximately identifying unique anatomical positions. The meshes were aligned separately to remove trivial pose variations, such as scaling, translation and rotation.

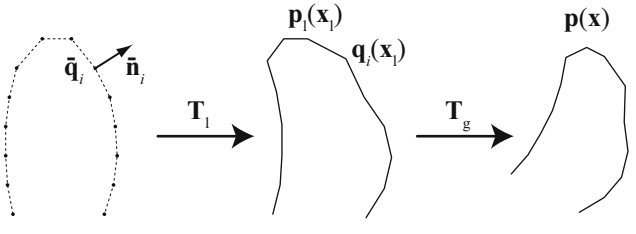

Fig. 1. A point $\mathbf{p}(\mathbf{x})$ on the ASM is generated by first applying a local transformation $\mathbf{T}_l$ described by the ASM state vector $\mathbf{x}_l$ on the mean shape, followed by a global pose transformation $\mathbf{T}_g$ to obtain a shape in real-world coordinates

From the aligned training set, the mean vertex position $\overline{\mathbf{q}}_i$ was computed, and PCA was applied on the vertex distribution to obtain the N_x most dominant eigenvectors. In normalized coordinates, the ASM can be written on the form

$$\mathbf{q}_i(\mathbf{x}_\mathrm{l}) = \overline{\mathbf{q}}_i + \mathbf{A}_i\,\mathbf{x}_\mathrm{l} \ , \tag{1}$$

where the position of a vertex $\mathbf{q}_i$ is expressed as a linear combination of the associated subspace of the N_x most dominating eigenvectors combined into the $3 \times N_\mathrm{x}$ deformation matrix $\mathbf{A}_i$. Here, $\mathbf{x}_\mathrm{l}$ is the local state vector of the ASM. The expression for the ASM can be optimized assuming that the deformation at vertex $\overline{\mathbf{q}}_i$ is primarily directed along the corresponding surface normal $\overline{\mathbf{n}}_i$ of the average mesh. This is done by projecting the deformation matrix $\mathbf{A}_i$ onto the surface normal, giving an N_x-dimensional vector of projected deformation modes $\mathbf{A}_i^\perp = \overline{\mathbf{n}}_i^T\,\mathbf{A}_i$. The optimized expression for the ASM can now be written on the form

$$\mathbf{q}_i(\mathbf{x}_\mathrm{l}) = \overline{\mathbf{q}}_i + \overline{\mathbf{n}}_i\,(\mathbf{A}_i^\perp\,\mathbf{x}_\mathrm{l}) \ , \tag{2}$$

reducing the number of multiplications by a factor of three.

Due to the quadrilateral mesh structure of the ASM, a continuous surface is obtained using a linear tensor product spline interpolant. An arbitrary point on the ASM in normalized coordinates can be expressed as $\mathbf{p}_\mathrm{l}(\mathbf{x}_\mathrm{l}) = \mathbf{T}_\mathrm{l}\big|_{(u,v)}$ where (u, v) represents the parametric position on the surface, and the local transformation $\mathbf{T}_\mathrm{l}$ includes the deformation and interpolation applied to the mean mesh. By coupling this model with a global pose transformation $\mathbf{T}_\mathrm{g}$ with parameters $\mathbf{x}_\mathrm{g}$ including translation, rotation, and scaling, we obtain a surface

$$\mathbf{p}(\mathbf{x}) = \mathbf{T}_\mathrm{g}(\mathbf{p}_\mathrm{l}(\mathbf{x}_\mathrm{l}), \mathbf{x}_\mathrm{g}) \tag{3}$$

in real-world coordinates, with a composite state vector $\mathbf{x}^T \equiv \left[\mathbf{x}_\mathrm{g}^T,\, \mathbf{x}_\mathrm{l}^T\right]$. An illustration showing the steps required to generate the ASM is shown in Fig. 1. In our experiments we used 20 eigenvectors, describing 98% of the total variation within the training set.

3 Tracking Algorithm

The tracking algorithm extends prior work by Orderud [9,10], to enable usage of 3D ASMs in the Kalman filter for real-time segmentation. This is accomplished by using the ASM shape parameters $\mathbf{x}_\mathrm{l}$ directly, in addition to the global pose parameters $\mathbf{x}_\mathrm{g}$, in the Kalman filter state vector.

3.1 Motion Model

Modeling of motion in addition to position can be accomplished in the prediction stage of the Kalman filter by augmenting the state vector to contain the last two

successive state estimates. A *motion model* which predicts the state $\bar{\mathbf{x}}$ at timestep $k+1$, is then expressed as

$$\bar{\mathbf{x}}_{k+1} - \mathbf{x}_0 = \mathbf{A}_1(\hat{\mathbf{x}}_k - \mathbf{x}_0) + \mathbf{A}_2(\hat{\mathbf{x}}_{k-1} - \mathbf{x}_0) \ , \tag{4}$$

where $\hat{\mathbf{x}}_k$ is the estimated state from timestep k. Tuning of properties, like damping and regularization towards the mean state $\mathbf{x}_0$ for all deformation parameters, can then be accomplished by adjusting the coefficients in matrices $\mathbf{A}_1$ and $\mathbf{A}_2$. Prediction uncertainty can similarly be adjusted by manipulating the process noise covariance matrix $\mathbf{B}_0$ used in the associated covariance update equation. The latter will then restrict the change rate of parameter values.

3.2 Measurement Processing

Edge-detection is based on *normal displacement* measurements v_i [4], which are calculated by measuring the radial distance between detected edge-points $\mathbf{p}_{\text{obs},i}$ and the contour surface $\mathbf{p}_i$ along selected search normals $\mathbf{n}_i$. These displacements are coupled with associated measurement noise r_i to weight the importance of each edge, based on a measure of edge confidence. Measurement vectors are calculated by taking the normal projection of the composite state-space Jacobian for the contour points

$$\mathbf{h}_i^T = \mathbf{n}_i^T \left[\frac{\partial \mathbf{T}_{\text{g}}(\mathbf{p}_{\text{l}}, \mathbf{x}_{\text{g}})_i}{\partial \mathbf{x}_{\text{g}}} \quad \frac{\partial \mathbf{T}_{\text{g}}(\mathbf{p}_{\text{l}}, \mathbf{x}_{\text{g}})_i}{\partial \mathbf{x}_{\text{l}}} \right] \ , \tag{5}$$

which is the concatenation of a global and a local state-space Jacobi matrix. The global Jacobian is trivially the state-space derivative of the global pose transformation, while the local Jacobian has to be derived, using the chain-rule for multivariate calculus, to propagate surface points on the spline through mesh vertices, and finally to the ASM shape parameters:

$$\frac{\partial \mathbf{T}_{\text{g}}(\mathbf{p}_{\text{l}}, \mathbf{x}_{\text{g}})}{\partial \mathbf{x}_{\text{l}}} = \sum_{n \in x,y,z} \frac{\partial \mathbf{T}_{\text{g}}(\mathbf{p}_{\text{l}}, \mathbf{x}_{\text{g}})}{\partial \mathbf{p}_{\text{l},n}} \sum_{j \in 1..N_q} \left(\frac{\partial \mathbf{T}_{\text{l}}(\mathbf{x}_{\text{l}})_n}{\partial \mathbf{q}_j} \cdot \bar{\mathbf{n}}_k \right) \mathbf{A}_j^{\perp} \ . \tag{6}$$

Here, $\partial \mathbf{T}_{\text{g}}(\mathbf{p}_{\text{l}}, \mathbf{x}_{\text{g}}) / \partial \mathbf{p}_{\text{l}}$ is the spatial derivative of the global transformation, and $\partial \mathbf{T}_{\text{l}}(\mathbf{x}_{\text{l}}) / \partial \mathbf{q}_j$ is the spatial mesh vertex derivative of the spline interpolant.

3.3 Measurement Assimilation and State Update

All measurements are assimilated in *information space* prior to the state update step. Assumption of independent measurements leads to very efficient processing, allowing summation of all measurement information into an information vector and matrix of dimensions invariant to the number of measurements:

$$\mathbf{H}^T \mathbf{R}^{-1} \mathbf{v} = \textstyle\sum_i \mathbf{h}_i r_i^{-1} v_i \tag{7}$$

$$\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} = \textstyle\sum_i \mathbf{h}_i r_i^{-1} \mathbf{h}_i^T \ . \tag{8}$$

The updated state estimate $\hat{\mathbf{x}}$ at timestep k can then by computed by using the information filter formula for measurement update [11], and the updated error covariance matrix $\hat{\mathbf{P}}$ is calculated directly in information space:

$$\hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + \hat{\mathbf{P}}_k \mathbf{H}^T \mathbf{R}^{-1} \mathbf{v} \tag{9}$$

$$\hat{\mathbf{P}}_k^{-1} = \bar{\mathbf{P}}_k^{-1} + \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} \ . \tag{10}$$

Using this form, we avoid inversion of matrices with dimensions larger than the state dimension.

4 Evaluation

4.1 Data Material

For evaluation of the proposed algorithm, apical 3D echocardiograms of one cardiac cycle from 21 adult patients (11 diagnosed with heart disease) were recorded using a Vivid 7 scanner (GE Vingmed Ultrasound, Norway) with a 3D transducer (3V). In all patients, meshes corresponding to the endocardial boundary were determined using a custom-made semi-automatic segmentation tool (GE Vingmed Ultrasound, Norway). The segmentations were, if needed, manually adjusted by an expert to serve as independent references equivalent to manual tracing.

4.2 Experimental Setup and Analysis

Edge measurements were done perpendicular to the mesh surface within a distance of ± 1.5 cm to the surface at approximately 450 locations, using a simple edge model based on the transition criterion [12]. The ASM was initialized to the mean shape, and positioned in the middle of the volume in the first frame. Segmentation was performed on the evaluation set by running the algorithm for a couple of heartbeats, to give the ASM enough time to lock on to the LV.

The accuracy of the ASM was assessed using the mean of absolute point-to-mesh distances between the ASM and the reference, averaged over one cardiac cycle. Volume differences (bias) between the ASM and the reference were calculated for each frame. End-diastolic volume (EDV), end-systolic volume (ESV), and EF $((\text{EDV} - \text{ESV})/\text{EDV} \cdot 100\%)$ were compared to the manually verified reference (two-tailed t-test assuming zero difference), with 95% limits of agreement (1.96 standard deviations (SD)). EDV and ESV were computed as the maximum and minimum volume within the cardiac cycle respectively.

5 Results

We observed that common challenges with 3D cardiac ultrasound, such as dropouts, shadows, and speckle noise were handled remarkably well, and segmentation was successful in all of the 21 patients. Some examples are shown in Fig. 2(b-d).

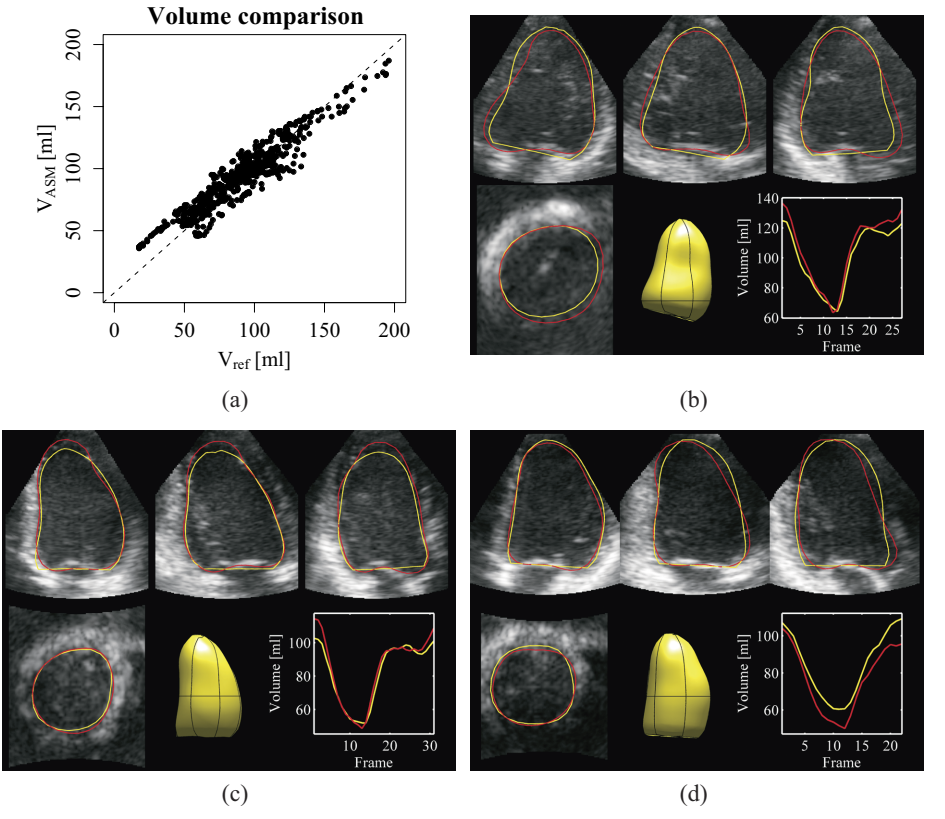

Fig. 2. LV volumes obtained by the ASM (V_{ASM}) in all 21 patients is compared to the reference (V_{ref}) and shown with the identity line (dashed) in (a). In (b-d), the end-diastolic segmentation (yellow) in three patients is compared to the reference (red) and shown along with the volume curve for one cardiac cycle.

The limits of agreement (mean±1.96SD) for the point-to-surface distance were 2.2±1.1 mm, indicating good overall agreement between the ASM and the reference. From Tab. 1, column 2, we see that the limits of agreement for volumes were 3.4±20 ml, with a strong correlation (r=0.95). The volume correspondence between the ASM and the reference is shown in Fig. 2(a).

Table 1. Segmentation results showing ASM versus reference

	Volume [ml]	EDV [ml]	ESV [ml]	EF [%]
Difference (mean±1.96SD)	3.4*±20	-5.9*±21	6.2*±19	-7.7*±12
Correlation coeff. (r)	0.95	0.91	0.91	0.74

* Significantly different from 0, $p<0.05$.

We also found a strong EDV correlation (r=0.91), with a bias and 95% limits of agreement of -5.9±21 ml (Tab. 1, column 3). The correlation in ESV was 0.91, with limits of agreement of 6.2±19 ml, (Tab. 1, column 4), while the correlation in EF was 0.74, with limits of agreement of -7.7±12% (Tab. 1, column 5).

The CPU load required to maintain real-time segmentation at 25 frames per second (fps) was approximately 22% on a 2.16 GHz Intel Core 2 Duo processor.

6 Discussion

We have presented a fully automatic real-time algorithm for robust and accurate LV segmentation in 3D cardiac ultrasound. This was achieved by combining a 3D ASM with a Kalman filter based tracking algorithm. The feasibility of the algorithm was demonstrated in 21 patients.

Computational performance was excellent with a CPU load of 22% at 25 fps. Compared to traditional ASM update schemes, the Kalman filter gives good segmentation in a single iteration, allowing real-time implementations.

Contours detected by the ASM showed good overall agreement with the reference shapes, both with respect to point-to-mesh distances and volumes. There was a significant bias in estimated EDV, ESV, and EF, but with relatively narrow 95% limits of agreement. We speculate if the bias is primarily caused by the simple edge detector used, and better results are expected using a more advanced edge detector. Robustness was high in the evaluation set, with successful segmentation in all patients.

Since no user-interaction is required, the algorithm provides rapid analysis of LV function, and it can potentially provide higher reproducibility than semi-automatic methods. Adding means of manual corrections when segmentation fails will be subject for further studies.

The algorithm was evaluated on a population with varying image quality, but for evaluation of clinical applicability, the algorithm must be tested on a larger population. Also, an inherent challenge when using ASMs for clinical applications, is that care must be taken when assembling the training set to ensure inclusion of a sufficiently wide range of pathologies.

Traditional applications where our algorithm fits well includes rapid analysis of LV volumes, EF, and regional function. With real-time segmentation, we expect new applications to emerge, such as patient monitoring, and automated operator guidance.

7 Conclusion

We have developed a fully automatic algorithm for real-time segmentation of the left ventricle in 3D cardiac ultrasound. Initial evaluation in 21 patients is very promising, suggesting that this method is applicable in a clinical setting.

Acknowledgment. The authors would like to thank Brage Amundsen at the Norwegian University of Science and Technology (NTNU) for providing the 3D echocardiography datasets.

References

1. Jacobs, L.D., Salgo, I.S., Goonewardena, S., Weinert, L., Coon, P., Bardo, D., Gerard, O., Allain, P., Zamorano, J.L., de Isla, L.P., Mor-Avi, V., Lang, R.M.: Rapid online quantification of left ventricular volume from real-time three-dimensional echocardiographic data. European Heart Journal 27, 460–468 (2006)
2. Sugeng, L., Mor-Avi, V., Weinert, L., Niel, J., Ebner, C., Steringer-Mascherbauer, R., Schmidt, F., Galuschky, C., Schummers, G., Lang, R.M., Nesser, H.J.: Quantitative assessment of left ventricular size and function. Side-by-side comparison of real-time three-dimensional echocardiography and computed tomography width magnetic resonance reference. Circulation 114, 654–661 (2006)
3. Blake, A., Curwen, R., Zisserman, A.: A framework for spatiotemporal control in the tracking of visual contours. International Journal of Computer Vision 11(2), 127–145 (1993)
4. Blake, A., Isard, M.: Active Contours: The Application of Techniques from Graphics, Vision, Control Theory and Statistics to Visual Tracking of Shapes in Motion. Springer-Verlag New York, Inc., Secaucus, NJ, USA (1998)
5. Jacob, G., Noble, J.A., Kelion, A.D., Banning, A.P.: Quantitative regional analysis of myocardial wall motion. Ultrasound in Medicine & Biology 27(6), 773–784 (2001)
6. Jacob, G., Noble, J.A., Mulet-Parada, M., Blake, A.: Evaluating a robust contour tracker on echocardiographic sequences. Medical Image Analysis 3(1), 63–75 (1999)
7. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models - Their training and application. Computer Vision and Image Understanding 61(1), 38–59 (1995)
8. van Assen, H.C., Danilouchkine, M.G., Behloul, F., Lamb, H.J., van der Geest, R., Reiber, J.H.C., Lelieveldt, B.P.F.: Cardiac LV segmentation using a 3D active shape model driven by fuzzy inference. In: Ellis, R.E., Peters, T.M. (eds.) MICCAI 2003. LNCS, vol. 2878, pp. 533–540. Springer, Heidelberg (2003)
9. Orderud, F.: A framework for real-time left ventricular tracking in 3D+T echocardiography, using nonlinear deformable contours and kalman filter based tracking. In: Computers in Cardiology (2006)
10. Orderud, F., Hansegård, J., Rabben, S.I.: Real-time tracking of the left ventricle in 3D echocardiography using a state estimation approach: Medical Image Computing and Computer-Assisted Intervention - MICCAI (Submitted 2007)
11. Comaniciu, D., Zhou, X.S., Krishnan, S.: Robust real-time myocardial border tracking for echocardiography: An information fusion approach. Medical Imaging, IEEE Transactions on 23(7), 849–860 (2004)
12. Rabben, S.I., Torp, A.H., Støylen, A., Slørdahl, S., Bjørnstad, K., Haugen, B.O., Angelsen, B.: Semiautomatic contour detection in ultrasound M-mode images. Ultrasound in Med. & Biol. 26(2), 287–296 (2000)