Real-time Left Ventricular Speckle-Tracking in 3D Echocardiography With Deformable Subdivision Surfaces

Fredrik Orderud[1], Gabriel Kiss[1], Stian Langeland[2], Espen W. Remme[3], Hans Torp[1], and Stein I. Rabben[2]

[1] Norwegian University of Science and Technology (NTNU), Norway
[2] GE Vingmed Ultrasound, Norway
[3] Department of Cardiology, Rikshospitalet, Norway

Abstract. In this paper, we extend a computationally efficient framework for real-time tracking of deformable subdivision surfaces in 3D echocardiography with speckle-tracking measurements to track material points. Tracking is performed in a sequential state-estimation fashion, using an extended Kalman filter to update the subdivision surface based on displacement vectors from 3D block-matching in the left ventricular wall. Fully automatic tracking is demonstrated in two simulations of an infarcted ventricle, as well as in a set of 21 in-vivo 3D echocardiograms. Credible tracking results were achieved in all cases, with an average drift ratio of $12.08 \pm 2.09\%$ (2.7 ± 1.0mm). The infarcted regions were also correctly identified in both of the simulations. Due to the high computational efficiency of the method, it is capable of operating in real-time.

1 Introduction

With the introduction of 3D echocardiography, rapid and low-cost acquisition of volumetric images of the left ventricle (LV) has become feasible. Tools for assessment of global function, based on semi-automatic shape segmentation of the endocardial boundary, have appeared over the last few years [1]. However, in order to evaluate regional function of the LV, methods that also estimate myocardial deformation by tracking material points are required.

The distinctive speckle pattern found in ultrasound images has often been considered an undesirable image artifact, since it reduces the apparent image quality. However, this pattern have the fortunate property that it, despite being gradually decorrelated, moves in the same manner as the underlying tissue being imaged [2]. This property can be exploited to track the LV myocardial deformation field by means of speckle-tracking techniques. Existing approaches for 2D and 3D speckle-tracking include: Horn-Schunck optical flow speckle-tracking [3], band-pass Gabor filtering prior to block-matching [4], tracking single speckle points with motion coherence regularization of the velocity field [5] and elastic volume registration using B-splines coupled with a mutual-information metric [6]. To our knowledge, none of these approaches have been demonstrated to achieve real-time processing times when applied to volumetric data.

This paper describes a fully automatic, real-time method for LV tracking of material points in 3D echocardiography. It extends the subdivision model based Kalman-filter method of [7], with speckle-tracking to capture the full myocardial deformation pattern, and not only shape changes. A combination of integer voxel displacement estimation using block matching, coupled with optical flow correction of the best match, is used to achieve sub-voxel displacement estimation.

2 Methods

The tracking framework is based on a deformable subdivision surface, consisting of control vertices $\mathbf{q}_i$ for $i \in \{1 \dots N_q\}$ that are allowed to move to alter the shape and parameter-space density of the surface. Unlike in [7], where shape segmentation was the objective, we allow the control vertices to move freely in any direction, and not just in the surface-normal direction. In addition to the control vertices, the topological relationships between the control vertices have to be defined in a list $C(c)$, that maps surface patches $c \in \{1 \dots N_c\}$ to enumerated lists of control vertex indices that define the control vertices influencing each surface patch.

We denote the local deformations $\mathbf{T}_l(\mathbf{x}_l)$ to our deformable model as the deformations obtainable by moving the control vertices of the subdivision model. These local deformations are combined with a global transform $\mathbf{T}_g(\mathbf{x}_g, \mathbf{p}_l)$ to position, scale and orient the model. This leads to a composite state vector $\mathbf{x} = [\mathbf{x}_g^T,\ \mathbf{x}_l^T]^T$ consisting of N_g global and N_l local deformation parameters.

A manually constructed Doo-Sabin subdivision surface [8] consisting of 20 control vertices is used to represent the LV. A distribution of approximately 450 surface points, spread evenly across the surface as in [7], are used as shown in Fig. 1(a). These points are used as a basis for both edge-detection and speckle-tracking measurements, and consists of parametric coordinates (including patch number) for each of the surface points.

The tracking framework is decomposed into the 5 separate steps shown in Fig. 1(b). Most of the steps are similar to [7], with the exception of the measurement step, where edge-detection is replaced with 3D speckle-tracking to update the model. Edge-detection is instead used solely to automatically initialize the model to the endocardial boundary, prior to speckle-tracking. This initialization leads to shape deformations, by moving control vertices inwards and outwards in the direction perpendicular to the surface, but does not impose any deformations along the surface to alter the parameter-space density, as speckle-tracking does.

The Kalman prediction step and Kalman update steps are identical as in [7], and therefore not covered in this paper. Instead this paper focuses on what is new, namely evaluation of the deformable model, speckle-tracking measurements and assimilation of 3D displacement vectors:

2.1 Evaluation of Deformable Model

Calculation of Local Surface Points: The Kalman filter framework requires the creation of a set of surface points $\mathbf{p}_l$ and Jacobi matrices $\mathbf{J}_l$, based on a pre-

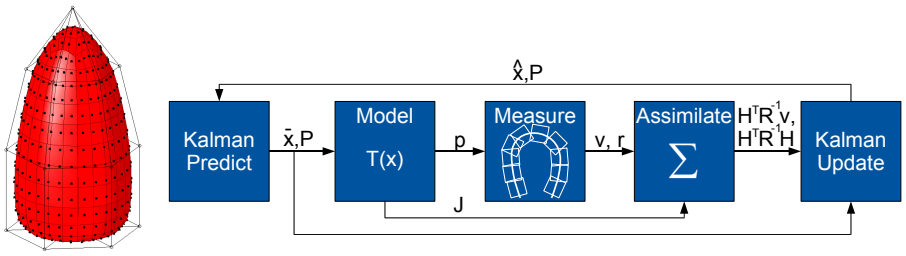

(a) (b)

Fig. 1. (a) The Doo-Sabin subdivision surface used for tracking, which consists of 20 control vertices shown in the encapsulating wire-frame mesh. The speckle-tracking distribution is illustrated with black dots on the surface. (b) Overview over the processing chain for each new frame in the Kalman filter tracking framework.

dicted state vector $\bar{\mathbf{x}}_l$. The creation of these objects can be performed efficiently following the steps below:

1. Update position of control vertices $\mathbf{q}_i$ based on the state vector: $\mathbf{q}_i = \bar{\mathbf{q}}_i + x_{(3i)}\mathbf{v}_x + x_{(3i+1)}\mathbf{v}_y + x_{(3i+2)}\mathbf{v}_z$, where $\bar{\mathbf{q}}_i$ is the initial position of the control vertex, $\mathbf{v}_x,\ \mathbf{v}_y,\ \mathbf{v}_z$ are unit vectors along the x, y and z axis repectively, and $x_{(3i)},\ x_{(3i+1)},\ x_{(3i+2)}$ are the parameters in the state vector corresponding to this control vertex. The full state vector for the model then becomes the concatenation of the state parameters for all control vertices $\mathbf{x}_l = \left[x_0,\ x_1,\ \ldots,\ x_{(3N_q-1)}\right]^T$.
2. Calculate surface points $\mathbf{p}_l$ as a sum of control vertices weighted with their respective basis functions within the surface patch of each surface point: $\mathbf{p}_l = \sum_{i \in C(c_l)} \mathbf{b}_i \mathbf{q}_i$.
3. Calculate Jacobian matrices for the local deformations $\mathbf{J}_l$ by concatenating the unit vectors multiplied with their respective basis functions: $\mathbf{J}_l = \left[\mathbf{b}_{i_1}\mathbf{v}_x, \mathbf{b}_{i_1}\mathbf{v}_y, \mathbf{b}_{i_1}\mathbf{v}_z,\ \mathbf{b}_{i_2}\mathbf{v}_x,\ \ldots\right]_{i \in C(c_l)}$. The Jacobian matrix will here be padded with zeros for columns corresponding to control vertices outside the region of support for the surface patch of each surface point.

Basis functions for these points can be precomputed during initialization if we restrict the parametric coordinate distribution of the surface points to be constant throughout the tracking. This allows the above operations to be performed very quickly, which is crucial for enabling real-time implementations.

Global Transform: We denote $\mathbf{p}_l$ and $\mathbf{J}_l$ for the surface points created from the subdivision surface with local deformations $\mathbf{T}_l(\mathbf{x}_l)$. These points are subsequently transformed by means of a global *pose transform* $\mathbf{T}_g$, that translates, rotates and scales the model to align it correctly within the image volume:

$$\mathbf{p}_g = \mathbf{T}_g(\mathbf{p}_l, \mathbf{x}_g)\ . \tag{1}$$

The Jacobian matrices for the composite deformations then becomes the concatenation of both global and local state-space derivatives. The local part is created by multiplying the 3×3 spatial Jacobian matrix for the global transform with the $3 \times N_l$ local Jacobian matrix for the deformable model, as follows from the chain-rule of multivariate calculus:

$$\mathbf{J}_g = \left[\frac{\partial \mathbf{T}_g(\mathbf{p}_l, \mathbf{x}_g)}{\partial \mathbf{x}_g}, \quad \frac{\partial \mathbf{T}_g(\mathbf{p}_l, \mathbf{x}_g)}{\partial \mathbf{p}_l} \, \mathbf{J}_l \right] . \tag{2}$$

2.2 Speckle-tracking Measurements

Speckle-tracking measurements are performed 1mm outside the endocardial surface, in order to track deformations inside the myocardium. There, 3D displacement vectors $\mathbf{v} = [v_x, \, v_y, \, v_z]^T$ for local motion are inferred by matching predicted surface points $\mathbf{p}$ in the current frame to associated surface point from the updated model in the previous frame. The measurements are computed by first performing 3D block-matching using a sum of absolute differences (SAD) metric to determine integer displacements. This is followed by translative Lucas-Kanade optical flow estimation [9] on the best integer voxel match to correct for sub-sample displacements, as was done for 2D tracking in [10].

Implementation of this matching can be done efficiently on modern processors by using vector instructions and multi-core parallelization of the SAD operations. Furthermore, preprocessing of the data is avoided by doing the tracking directly on the "raw" grayscale ultrasound data acquired in spherical coordinates. Tracking is performed in data decimated in the beam propagation direction by a factor of four to reduce window sizes, since ultrasound image resolution is significantly higher in this direction compared to the two lateral directions. A kernel size of 4x4x4 voxels is used, while the search window has an adaptive size, based on image depth, to make its cartesian dimensions approximately constant regardless of image depth.

Associated measurement noise values r, for the spatial uncertainty of the displacement measurement, are computed based on the ratio between the best and average SAD matching value for each point. After computing the measurements, simple outlier rejection is performed, based on the measurement noise values and a comparison with neighboring displacement vectors in a local search area. Parameters were adjusted by trial and error.

2.3 Measurement Assimilation

The measurements can be efficiently assimilated in information space if we assume that they are uncorrelated [11], since uncorrelated measurements lead to a diagonal measurement covariance matrix $\mathbf{R}$. All measurement information can then be summed into an information vector and matrix of dimensions invariant to the number of measurements:

$$\mathbf{H}^T \mathbf{R}^{-1} \mathbf{v} = \textstyle\sum_i \mathbf{H}_i^T \mathbf{v}_i r_i^{-1} \tag{3}$$

$$\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} = \textstyle\sum_i \mathbf{H}_i^T \mathbf{H}_i r_i^{-1} \, . \tag{4}$$

Usage of unit vectors in x, y and z-direction for displacing control vertices enables direct usage of the Jacobian matrices as measurement matrices in the Kalman filter, since $\mathbf{H}^T = [\mathbf{v}_x,\ \mathbf{v}_y,\ \mathbf{v}_z]^T\mathbf{J} = \mathbf{J}$. A covariance matrix for the measurement can also be used instead of scalar measurement noise values if one desires to capture any non-isotropy in the spatial uncertainty of the displacement, but this is not done in this paper.

3 Experimental Validation & Results

In order to show the feasibility of the method and validate its performance, the method was applied to both simulated and in-vivo data sets:

3.1 Data Description

Two volumetric ultrasound datasets was generated for the experiment, based on a finite element simulation of a left ventricle with an anterio-apical infarction. The first simulation used an ellipsoidal shape for the myocardium, while the second used the average shape of five canine ventricles. The motion and deformation of the ventricle was determined by modeling internal systolic contraction forces and external forces from the cavity pressure, and the infarcted area was modeled by abolishing contractile forces in the anterio-apical region, as in [12]. A *k-space* ultrasound simulator, described in [13], was then used to create realistic 3D ultrasound simulations based on scatter positions extracted from the finite element model.

In addition to the simulations, a collection of 21 apical 3D echocardiography recordings of adult patients, of which half were diagnosed with heart diseases, were used for in-vivo validation of the method. These recordings were acquired with a Vivid 7 scanner (GE Vingmed Ultrasound, Norway) using a matrix array transducer (3V). The exact same configuration was used to initialize tracking in all in-vivo recordings.

3.2 Simulated Data Results

After tracking, the subdivision surfaces were re-meshed into a grid-structured mesh in a manner that preserves material points. Area strain $\epsilon = (a - a_0)/a_0$ values were then computed locally across the surface by comparing the area of each quadrilateral during tracking with associated end-diastolic (ED) areas. These strain values represents variations in the parameter-space density of the subdivision surface, which should correspond to the total muscle contraction locally in a manner that combines the effect of longitudinal and circumferential strain.

Figure 2(b) shows end-systolic (ES) area strain values across the tracked surface, together with ground truth strain values from both simulations. The infarcted regions exhibit small contraction or stretching, and therefore show up as blue and green, while healthy myocardium is contracting, and therefore shows

up as red. One can clearly see that the infarcted regions are correctly identified in the tracked meshes, although the strain values in the infarcted regions are underestimated compared to the ground truth, especially in the second simulation which has a more complex geometry.

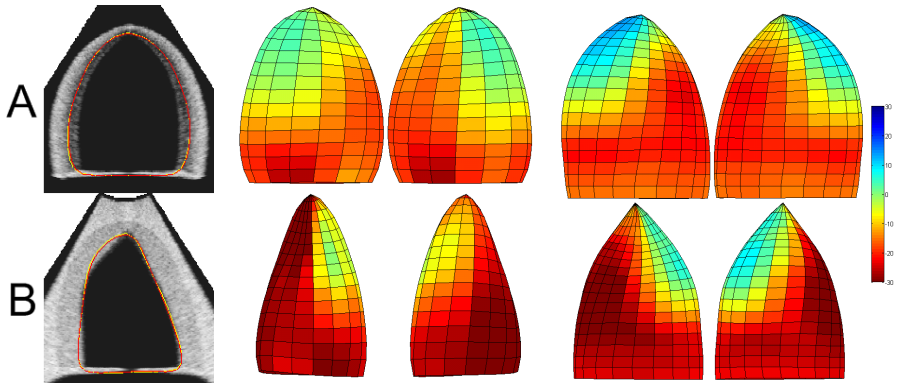

(a) (b) (c)

Fig. 2. Results from tracking in the simulated data, showing (a) intersection slices through the simulations, as well as front & back views of color-coded area strain meshes at ES from (b) the tracked strain meshes and (c) ground truth values. Results from the infarcted ellipsoid (simulation 'A') are shown in the top row, and from the infarcted dog-heart (simulation 'B') in the bottom row.

3.3 In-vivo Results

For the in-vivo data, no ground truth was available, so tracking was instead evaluated by computing the average ratio between the drift of surface points after tracking an entire cardiac cycle, and the walked path distance for the same point.

Tracking in the 21 in-vivo recordings yielded an average drift ratio of $12.08 \pm 2.09\%$, which in absolute values corresponds to 2.7 ± 1.0mm. As a comparison, the drift ratio in the simulated recordings were 8.58% and 10.59%, with absolute drift values of 0.58 and 0.70mm. Fig. 3 shows orthogonal intersection slices of the tracking results at ED and ES in two of the recordings.

Tracking in both the simulated and in-vivo recordings consumed approximately 37ms of processing time per frame on a 2.2GHz Intel Core 2 duo processor. This makes the method capable of operating in real-time, given a typical frame-rate of 25fps.

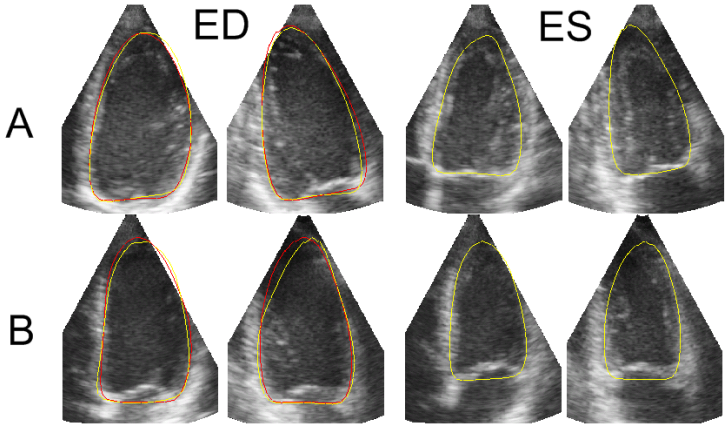

Fig. 3. Orthogonal image slices of the tracked mesh in two of the in-vivo recordings (patient 'A' and 'B'), at both ED and ES. The slices show the initialized mesh in red (only at ED), and tracked meshes in yellow (both at ED and ES). The difference between the tracked and initialized mesh at ED constitutes the drift after tracking in an entire cardiac cycle.

4 Discussion

We have presented a new approach for LV tracking of material points in 3D echocardiography, using a Kalman filter to fit a deformable subdivision model to 3D speckle-tracking measurements. The method is automatically initialized using endocardial edge-detection, and is capable of operating in real-time due to its strong computational efficiency. Usage of this method might therefore enable rapid analysis of regional myocardial function.

Automatic myocardial speckle-tracking was feasible in all tested recordings, both simulated and in-vivo. Based on visual assesment of the tracking and computed drift values, the tracking was found to behave robustly. The obvious discrepancy between absolute drift and drift ratio in the simulations is believed to stem from the fact that apex moves very little in the simulations, which leads to high drift ratios in the apical region, even though the absolute drift is quite small. Furthermore, the method was able to identify the infarcted regions in two different ultrasound simulations of a left ventricle, although the strain in the infarcted areas appeared underestimated compared to ground truth.

Tracking accuracy is dependent on the resolution of the subdivision surface used. Low-resolution surfaces, like the one used in this paper, will exhibit high robustness, due to the inherent regularization of having fewer parameters to estimate. This does, however, come at the expense of stronger spatial smoothing across the surface, compared to more high-resolution surfaces. This might be some of the reason for the underestimated infarcted strain. Usage of a higher resolution model, consisting of more control vertices, might therefore enable more accurate tracking with less spatial smearing of the deformation field. Usage of

more surface points for block-matching might also increase tracking accuracy at the expense of computational efficiency.

With feasibility of the method demonstrated, the next step will be to compare its accuracy to alternative approaches in a more quantitative way. Approaches to improve its accuracy should also be investigated. Currently, tracking is performed sequentially, from one frame to the next. This can, however, be extended with bidirectional tracking, using both a forward and backward Kalman-filter to improve tracking accuracy and reduce drift. Speckle-tracking measurements might also be combined more directly with edge-detection to reduce the surface-normal component of the inherent drift associated with sequential block-matching.

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** (November 2006) 460–468
2. D'hooge, J., Bijnens, B., Thoen, J., Van de Werf, F., Sutherland, G., Suetens, P.: Echocardiographic strain and strain-rate imaging: a new tool to study regional myocardial function. Medical Imaging, IEEE Transactions on **21**(9) (Sep 2002) 1022–1030
3. Meunier, J.: Tissue motion assessment from 3D echographic speckle tracking. Physics in Medicine and Biology **43**(5) (1998) 1241–1254
4. Yu, W., Lin, N., Yan, P., Purushothaman, K., Sinusas, A., Thiele, K., Duncan, J.S.: Motion analysis of 3D ultrasound texture patterns. Functional Imaging and Modeling of the Heart (2003) 1006–1006
5. Song, X., Myronenko, A., Sahn, D.J.: Speckle tracking in 3D echocardiography with motion coherence. Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on (17-22 June 2007) 1–7
6. Elen, A., Loeckx, D., Choi, H.F., Gao, H., Claus, P., Maes, F., Suetens, P., D'hooge, J.: P4a-5 3D cardiac strain estimation using spatio-temporal elastic registration: In silico validation. Ultrasonics Symposium, 2007. IEEE (28-31 Oct. 2007) 1945–1948
7. Orderud, F., Rabben, S.I.: Real-time 3D segmentation of the left ventricle using deformable subdivision surfaces. Computer Vision and Pattern Recognition, 2008. CVPR '08. IEEE Conference on
8. Doo, D., Sabin, M.: Behaviour of recursive division surfaces near extraordinary points. Computer-Aided Design **10**(6) (November 1978) 356–360
9. Lucas, B., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: IJCAI81. (1981) 674–679
10. Behar, V., Adam, D., Lysyansky, P., Friedman, Z.: Improving motion estimation by accounting for local image distortion. Ultrasonics **43**(1) (October 2004) 57–65
11. Bar-Shalom, Y., Li, X.R., Kirubarajan, T.: Estimation with Applications to Tracking and Navigation. Wiley-Interscience (2001)
12. Remme, E., Smiseth, O.: Characteristic strain pattern of moderately ischemic myocardium investigated in a finite element simulation model. In: Functional Imaging and Modeling of the Heart. (2007) 330–339
13. Hergum, T., Crosby, J., Langhammer, M., Torp, H.: The effect of including fiber orientation in simulated 3D ultrasound images of the heart. Ultrasonics Symposium, 2006. IEEE (2-6 Oct. 2006) 1991–1994