

Augmenting CT Cardiac Roadmaps with Segmented Streaming Ultrasound

Qi Duan ^{a,b}, Guy Shechter ^a, Luis F. Gutiérrez ^a, Douglas Stanton ^a, Lyubomir Zagorchev ^a,
Andrew F. Laine ^b, Daniel R. Elgort ^{*a}

^a Philips Research North America, 345 Scarborough Rd., Briarcliff Manor, NY 10510

^b Department of Biomedical Engineering, Columbia University, ET351, 1210 Amsterdam Avenue,
New York, NY, 10025

ABSTRACT

Static X-ray computed tomography (CT) volumes are often used as anatomic roadmaps during catheter-based cardiac interventions performed under X-ray fluoroscopy guidance. These CT volumes provide a high-resolution depiction of soft-tissue structures, but at only a single point within the cardiac and respiratory cycles. Augmenting these static CT roadmaps with segmented myocardial borders extracted from live ultrasound (US) provides intra-operative access to real-time dynamic information about the cardiac anatomy. In this work, using a customized segmentation method based on a 3D active mesh, endocardial borders of the left ventricle were extracted from US image streams (4D data sets) at a frame rate of approximately 5 frames per second. The coordinate systems for CT and US modalities were registered using rigid body registration based on manually selected landmarks, and the segmented endocardial surfaces were overlaid onto the CT volume. The root-mean squared fiducial registration error was 3.80 mm. The accuracy of the segmentation was quantitatively evaluated in phantom and human volunteer studies via comparison with manual tracings on 9 randomly selected frames using a finite-element model (the US image resolutions of the phantom and volunteer data were 1.3 x 1.1 x 1.3 mm and 0.70 x 0.82 x 0.77 mm, respectively). This comparison yielded 3.70±2.5 mm (approximately 3 pixels) root-mean squared error (RMSE) in a phantom study and 2.58±1.58 mm (approximately 3 pixels) RMSE in a clinical study. The combination of static anatomical roadmap volumes and dynamic intra-operative anatomic information will enable better guidance and feedback for image-guided minimally invasive cardiac interventions.

Keywords: Abdominal Procedures, Cardiac Procedures, Multimodality Display, Registration, Segmentation and Rendering, Ultrasound Guidance

1. INTRODUCTION

In cardiovascular minimally invasive interventional procedures such as catheter-based radiofrequency ablation¹, pre-procedural roadmaps can be valuable for surgical planning and intra-procedural guidance by complementing the intra-procedural imaging modality. Though multiple imaging modalities have been proposed for acquiring such roadmaps, cardiovascular X-ray computed tomography (CT) is widely used since it can provide three-dimensional high-resolution depiction of soft-tissues² not usually attainable with intra-procedural imaging modalities. The utility of pre-procedural roadmaps is usually limited in that they are static volumes acquired at a single point within cardiac and respiratory cycles, whereas during the interventional procedure, the geometry and position of the heart changes over time due to cardiac motion and breathing. Several modalities, such as X-ray fluoroscopy³, interventional magnetic resonance imaging (iMRI)⁴, and 2D/3D ultrasound (US) have been used for intra-procedural imaging to acquire dynamic information about the targeted tissues, though each modality has its own limitation. For example, X-ray based methods do not depict soft tissue structures well, are limited to 2D projection perspectives, and introduce additional ionizing radiation to the patient and the operators; MR-based methods require an MR-compatible operating room and instruments. On the other hand, real-time 3D (RT3D) US is used for intra-procedural live monitoring since it is portable, does not generate ionizing radiation, and does not require any special operating environments.

* daniel.elgort@philips.com; phone: 1 914 945-6092; fax: 1 914 945-6580; www.philips.com

Development of RT3D echocardiography started in the late 1990s by Volumetrics⁵ based on matrix phased arrays transducers. Current systems can acquire roughly one quadrant of the heart in a single 3D field-of-view in real-time, or can stitch together acquisitions over four cardiac cycles for a fully sampled cardiac volume. This quadrant-wise acquisition design enables a dramatic increase in spatial resolution and image quality compared to those reconstructed 3D ultrasound from 2D B-mode slices, which makes such 3D ultrasound techniques increasingly attractive for daily cardiac clinical diagnoses. Since RT3D ultrasound acquires volumetric ultrasound sequences with fairly high temporal resolution (about 20 fps per quadrant) and a stationary transducer, complex 3D cardiac motion can be captured with high fidelity. RT3D US data from one quadrant can roughly cover the left ventricle and can be recorded in true real-time fashion, with a temporal resolution of about 20 frames per second. This dynamic 3D imaging modality provides new opportunities for non-invasive monitoring of intra-operative tissue dynamics and functional information.

The purpose of this work was to explore the feasibility of augmenting static cardiac CT roadmaps with segmented myocardial borders extracted from RT3D US streams. Emphasis was placed on the integration of information based on the inherent strengths of each imaging modality. CT provides a static high-resolution depiction of cardiac soft-tissue structures; RT3D US provides intra-operative access to real-time dynamic information about the cardiac tissue without exposing the patient to additional ionizing radiation. Integration and fusion of images and information from multiple modalities is a critical enabling technology that will improve the quality of intra-operative guidance for many minimally invasive cardiac interventions.

2. METHODS

There are different ways of incorporating information from streamed US into a CT roadmap. One direct way is to overlay ultrasound images onto the CT volume. However, even though the multimodality registration step and the volume rendering step can be implemented with real-time performance, overlaying all of the image data from the high speed volumetric streaming ultrasound on top of the CT volume can be overwhelming to operators. In this work, we extract the relevant dynamic cardiac information by online segmentation of endocardial borders from RT3D US images; then, the extracted endocardium is overlaid on the static CT roadmap in real-time to augment the intra-procedural roadmap display. This work involves multi-modality registration and online segmentation of RT3D US image data. In the following sections, the hardware and imaging system will be introduced; then the registration framework will be explained; finally, the design and performance of a real-time segmentation method will be analyzed.

2.1. Hardware and Imaging

A multi-modality imaging phantom was constructed to facilitate technology development and validation. The custom-built phantom was tuned for realistic soft-tissue contrast under both CT and US imaging and has geometric similarity to a human heart (Fig. 1a).

X-ray CT imaging was performed on a Philips Brilliance 16 slice CT scanner. The CT image data was acquired with a spatial resolution of $0.59 \times 0.59 \times 1$ mm and field-of-view of $300 \times 300 \times 200$ mm. This CT volume was used as the static pre-procedural roadmap (Fig 1b). Ultrasound imaging was performed on a Philips iE33 system, using the X3-1 3D imaging probe (Fig 1c). For development purposes, RT3D US data sets were saved to disk for repeated use. In our segmentation experiments, streaming of the data was simulated by loading a single 3D frame at a time from disk and applying our processing steps before loading the next 3D frame. During the US scanning, the cardiac phantom was manually deformed to simulate cardiac motion. The spatial resolution of the RT3D US image data of the phantom was $1.3 \times 1.1 \times 1.3$ mm. The temporal resolution was 18 fps.

Cardiac US image data was collected in two human volunteers. The clinical protocols were approved by the institutional review board and informed consent was secured. The RT3D US images were acquired with the same hardware that was used in the phantom experiments and with a spatial resolution of $0.70 \times 0.82 \times 0.77$ mm, and a temporal resolution of 23 fps. CT imaging was not performed on the human volunteers.




Fig. 1. (a) A photograph of the multi-modality phantom under experimental settings; Three orthogonal planes through (b) the corresponding CT roadmap and (c) one sample volume taken from the 4D ultrasound stream.

2.2. Registration

The CT coordinate system was chosen as the reference coordinate system, and the ultrasound data was registered to it. In this initial work, the ultrasound probe was left in a fixed position relative to the cardiac anatomy, so a one-time coordinate registration was sufficient.

$$X_{CT} = R_{US}^{CT} X_{US} \quad (1)$$

Equation (1) describes the relationship between the CT and US coordinate systems. X_{US} refers to image coordinates in the US data space. When multiplied by the registration matrix, R_{US}^{CT} , US coordinates are transformed to corresponding image coordinates in CT data space, X_{CT} . The registration matrix, R_{US}^{CT} , describes a rigid body transform between the two coordinate systems. Anatomical features such as the mitral valve annulus, tip of the apex, and cross-sectional points of the artificial coronary artery, which were visible in both imaging modalities, were used to define the corresponding points in the CT image and in the first frame of the ultrasound image data. The registration matrix, R_{US}^{CT} , was calculated analytically using a least squares technique for eight manually selected point landmarks in the two image datasets.

2.3. Segmentation

In ultrasound image analysis, including segmentation, a preprocessing step is usually applied to suppress the speckle noise in the ultrasound images⁶. However, such processing usually needs full volume linear or non-linear processing, which may require substantial computational power. This is acceptable for offline processes but this requirement is incompatible with this online or “real-time” segmentation application.

In cardiac ultrasound segmentation, besides the classical methods based on thresholding^{7, 8} and morphological operations^{9, 10}, parametric active surface models^{11, 12, 13, 14}, level set frameworks^{15, 16}, and active shape/appearance models (ASM/AAM)^{17, 18, 19} have been proposed as alternate methods. Since the general level set implementation is relatively slow, and the requirement of a large training database for ASM/AAM is non-realistic in the cardiac intervention context, we chose to develop a segmentation algorithm using a customized 3D active mesh model^{11, 20, 21}.

2.3.1. Surface Representation

In order to efficiently capture the shape of endocardium, instead of using a traditional triangulated mesh with a linear surface within each patch, 3rd order Hermite polynomials were used as a surface descriptor^{11, 22}.

In 1D, there are four cubic Hermite basis functions:

$$\begin{aligned}
H_0^0(\xi) &= 1-3\xi^2+2\xi^3 \\
H_0^1(\xi) &= \xi(\xi-1)^2 \\
H_1^0(\xi) &= \xi^2(3-2\xi) \\
H_1^1(\xi) &= \xi^2(\xi-1)
\end{aligned} \tag{2}$$

On a 2D finite element patch, as shown in Fig. 2, there will be 4 basis functions associated with each local coordinate direction (ξ_1 or ξ_2), which generates a total of 16 2D basis functions. Fig. 2b shows an example of using 3rd order Hermite polynomials to represent a convoluted 3D surface in 3D space using a single 2D finite element patch. For comparison, if linear quadrilateral patches are used, about 100 linear patches are required to represent the same surface with comparable accuracy. For this reason, Hermite polynomials are widely used in cardiac biomechanics studies for surface representation^{22, 23, 24}. A simple 8x8 finite element model (FEM) with intrinsic C^1 continuity can sufficiently represent the geometry of the endocardium^{23, 24}. In our implementation, this 8x8 convention was followed, i.e. an 8x8 FEM using cubic Hermite polynomials as surface descriptors was used to represent the endocardium. The model was driven to segment the endocardium by the forces derived from an energy function described in the following section.

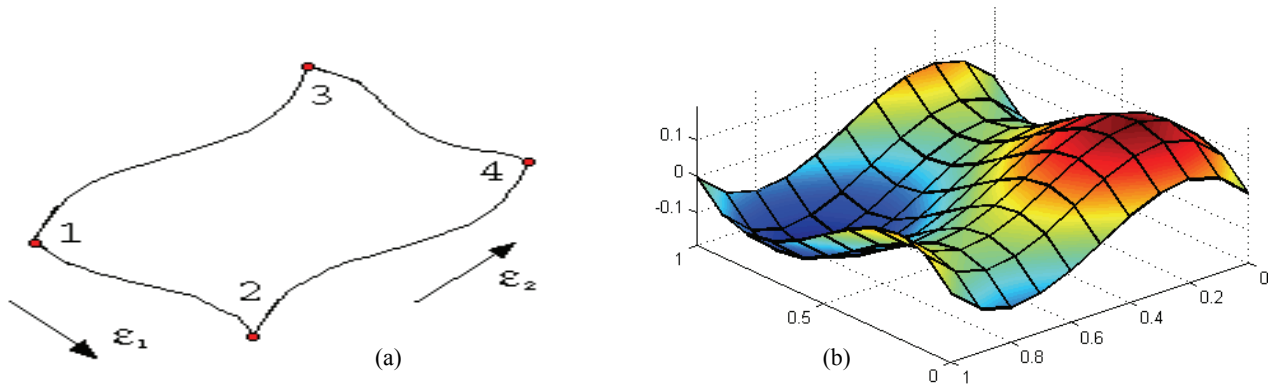


Fig. 2. Illustration of cubic Hermite surface representation: (a) a 2D surface finite element patch used in our model; (b) a single patch with cubic Hermite representation can efficiently characterize a convoluted surface as shown in color, whereas it takes about 100 linear quadrilateral patches to achieve similar accuracy.

2.3.2. Segmentation Energy Functional

Since it would be too computationally expensive to pre-process or smooth the US images in our real-time application, traditional segmentation approaches that use an image gradient based energy minimization function are not robust enough due to the high noise level of ultrasound images. In order to overcome this problem, we borrowed an idea from the “active contour without edges” concept, which was originally proposed by Chan and Vese²⁵ for a level set framework.

The 3D active mesh model is deformed under external forces derived from a minimization of the Mumford-Shah energy functional [5]:

$$E(f, \bar{C}) = \beta \int_{\Omega} (f - g)^2 dV + \alpha \int_{\Omega \setminus \bar{C}} |\nabla f|^2 dV + \gamma \oint_{\bar{C}} ds \tag{3}$$

in which \bar{C} denotes the smoothed and closed segmentation, g represents the observed US data, f is a piecewise smoothed approximation to g with discontinuities only along \bar{C} , and Ω denotes the image domain. The first integral enforces the similarity between f and g , which is equivalent to homogeneity constraint since f is piecewise smoothed; the second integral controls the smoothness of f ; and the last integral is actually the length of the segmented boundary, which, acts as internal elasticity to prevent leaking at weak boundaries. The external forces driving the 3D active mesh

were formulated using the same image homogeneity based rationale proposed by Chan and Vese²⁵, using homogeneity measures inside and outside regions based on the current segmentation. Specifically, the optimum segmentation corresponding to endocardium divides the local region of the image into two relatively homogeneous regions. In this application, these regions correspond to blood pool and myocardium.

The Mumford-Shah equation (Equation (3)) was only evaluated at discrete super-sampled sub-node points. To ensure that the system was well-defined, each surface patch was super-sampled by a 4x4 sub-node grid. In this way, the whole surface model was over-constrained with continuity constraints between each adjacent surface patch since the surface basis functions are cubic. The Mumford-Shah equation (Equation (3)) was minimized using a Newton Downhill method, chosen for its computational efficiency.

2.3.3. Repositioning and Reorientation of the Active Surface

One common drawback of segmentation using parametric active surface models is that the capture range of the method is usually small compared to other methods. In ultrasound segmentation applications, an initial contour is usually required to be fairly close to the actual boundary²⁶. To accomplish this, manual initialization is usually required, which is not preferred in our interventional application. Automated and semi-automated methods have been proposed to solve this problem, which use optical flow tracking²⁶, Hough transform²⁷, or a multi-scale implementation²⁸. Since temporal performance is critically important for our real-time interventional application, a computationally efficient approach was needed.

In order to reduce the dependence on the initial position, after the segmentation converges for the current frame, the model repositions itself so its center corresponds to the centroid of the current segmentation and reorients itself so its long axis corresponds to the long axis of the current segmentation's long axis. The repositioned and reoriented model is used to initialize the segmentation of the next frame. This extra step reduces the dependence on the model's initial position, speeds up the convergence, and keeps the Hermite coefficients at each node as small as possible, which increases the numerical stability of the optimization procedure.

2.3.4. Implementation

The segmentation software was implemented in C++ using an ITK²⁹/VTK³⁰ compatible framework. These open source libraries were used for rendering, visualization, and interaction with the image data.

2.4. Validation

To quantitatively evaluate the segmentation's performance, nine 3D image volumes from a 4D phantom data set, (which contained a total of 500 3D frames), and four image volumes from a 4D clinical data set, (which contained a total of 87 3D frames), were randomly selected. The endocardial surfaces were manually traced using a customized tracing interface implemented using Matlab (The MathWorks, Inc, Natick, Massachusetts); roughly 200-300 boundary points were selected for each volume. To quantitatively compare the manual segmentation to the automatic segmentation, the output of the automatic segmentation for the selected frame, i.e. the converged finite element surface model (FEM), was fitted to the corresponding manual tracing point cloud. The FEM which has been re-fitted to the manual tracing points now has updated Hermite parameters at each node. The endocardium achieved by manual tracing and automated segmentation can be efficiently represented by two sets of Hermite parameters at each node of a single finite element surface model, which enables point-wise quantitative surface comparison. In this study, surface distance at each corresponding node point from the FEM for each surface was used as the criterion for surface discrepancy. A similar approach for segmentation comparison can be found in previous studies³¹.

3. RESULTS

3.1. Results on Phantom Data

The CT and US coordinate systems were rigidly registered using 8 manually selected landmarks, and the RMS fiducial registration error was 3.80 mm.

The segmentation algorithm was applied to each ultrasound frame (500 in total) in the sequence, requiring 4.12 ± 0.73 (mean \pm stdev) iterations per frame to converge on the endocardial border. The processing for each image frame took 150-200 ms on a Pentium 4 (2.80GHz desktop with 1.00GB RAM running Microsoft Windows XP), enabling a 5 frame/second application update. The actual time required for the segmentation was 50 ± 8 ms; the remaining roughly 130 ms were required for data loading from the hard disk and rendering.




Fig. 3. (a) Representative ultrasound frame from an ultrasound stream, with chamber surface (red) being segmented at 5 frames per second. The ultrasound derived surface is registered and overlaid onto a static “un-deformed” pre-procedural CT image volume. The surface in (b) is derived from an ultrasound frame acquired in an “un-deformed” state, while the surface in (c) was derived from an ultrasound frame acquired in a “deformed” state.

Fig. 3(a) shows a snapshot of the endocardial border segmentation (red surface) in the space of the ultrasound images. After applying the rigid body transformation, the segmented endocardial surface can be mapped into the pre-procedural CT dataset shown in (b) using the transform R_{US}^{CT} . As the ultrasound images capture dynamic changes in the shape of the phantom, the segmentation is updated, and continuously projected into the static CT dataset (Fig. 3c).

Quantitative evaluation of the automatic segmentation compared to the manual tracings, as described in Section 2.4, yielded an RMS error of 3.70 ± 2.5 mm which represents less than 3 pixels. Sample comparison frames are shown in Fig. 4, where the red surface is the endocardium segmented by our automatic segmentation method and the green surface is the segmentation from the manual tracings.




Fig. 4. Comparison of phantom data between the automated method (red) and manual tracing (green) on one frame at three different view angles.

3.2. Results on Clinical Data

The segmentation algorithm was also tested on each frame of a clinical patient data set (87 frames in total). The clinical data set had better image resolution than the phantom data. Quantitative evaluation of the segmentation via comparison with manual tracings, as described above, yields an RMSE of 2.58 ± 1.58 mm which corresponds to approximately 3 pixels. The actual time required for segmentation of each frame of clinical US data was 75 ± 12 ms (this time does not include loading the data from the disk or rendering). Sample comparison frames are shown in Fig. 5.

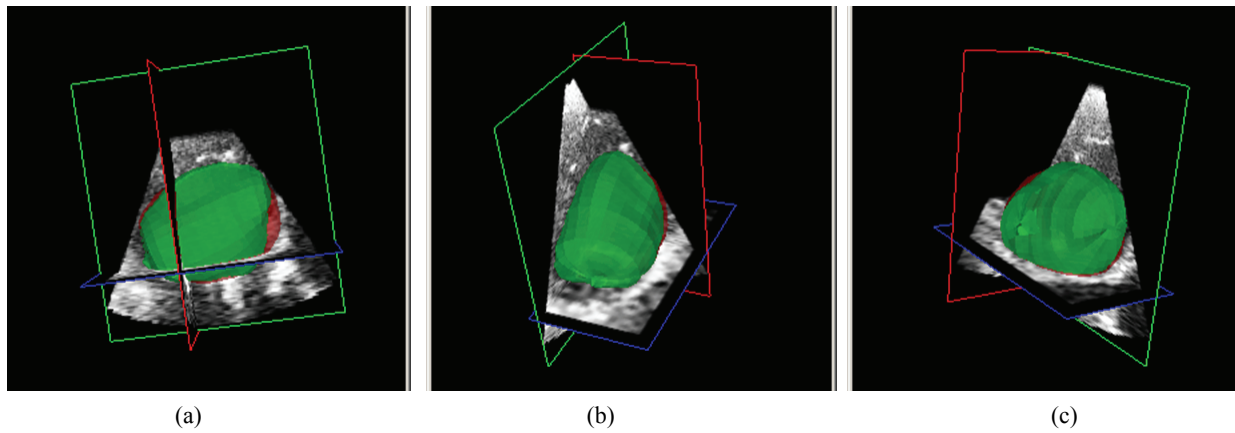


Fig. 5. Comparison of clinical data between the automated method (red) and manual tracing (green) on one frame at three different view angles.

4. DISCUSSIONS

This work presents a method for extracting endocardial surfaces in near real-time from a RT3D cardiac US stream and using this intra-procedural information to augment static CT roadmaps. This work demonstrates the feasibility of using streaming RT3D cardiac ultrasound as a real-time non-invasive tool for qualitative assessment of dynamic tissue morphology and proposes a framework for incorporating this information into an intra-procedural display for image guidance.

4.1 Registration and Segmentation Accuracy

The RMS segmentation error of the phantom data was 3.70 ± 2.5 mm. This is comparable to the findings in other studies using active contours for 2D cardiac ultrasound segmentation²⁶. The segmentation errors in the clinical data set were smaller (2.58 ± 1.58 mm RMSE). The RMSE values reported in this study were calculated via a comparison between two finite element surface models (FEM), one fit by the automatic segmentation algorithm to the endocardial surface within

the image data, and one fit to a manual tracing. Inaccurate manual tracings and/or a less-than-perfect fit of the FEM surface to the manual tracing will also contribute to this RMSE value. The improved segmentation accuracy with the clinical cardiac images is likely due to the data set's higher spatial resolution. The increased image resolution provides more support points along the endocardial surface. Additionally, higher image resolution may facilitate more accurate manual tracings.

The multi-modal RMS registration error was 3.80 mm. This registration error is possibly due to the variation in manually picked land marks, the limited resolution of the ultrasound images, and deformation of the phantom between the CT and US scans. This error is comparable to the RMS error in other similar studies³² (5mm RMS surface distance between CT and MR co-registration).

In the multi-modality intra-operative display, the segmented endocardial surface from RT3D US will be registered to the CT volume. The registration errors and the segmentation errors will both contribute to the overall error in the augmented roadmap display. The results of the phantom registration and segmentation experiments in this study indicate that this overall error could be as large as 7.5 mm. It is important to note that in the context of this application, i.e. augmentation of a static CT roadmap, the segmentation is not intended to provide precise quantitative functional measurements, such as ejection fraction. Therefore the implementation emphasizes temporal performance at the cost of optimal segmentation accuracy.

4.2 Sensitivity to Initialization

In order to minimize the amount of manual interaction that is required to initialize the segmentation model, the algorithm was simply initialized as a small sphere in the center of the ultrasound volume. A sample screenshot of the initial state is shown in Fig. 6a. Only a small portion of the surface was inside the left ventricle at the time of initialization. As shown in Fig. 6a, the surface is initialized with a major part of the model residing within the myocardium; only a small portion of model is actually inside the ventricle. Several frames later, as shown in Fig. 6b, the model begins to grow into the ventricle. The segmentation depicted in Fig. 6b is the result after the model converged. The incorrect segmentation in Fig. 6b is a transient result due to the poor initialization. Repositioning and reorienting the model based on the current segmentation are key algorithmic steps that enable the model to adjust itself and overcome poor initialization. In our experiments, the model only required about 1-2 seconds (about 25 frames) for it to properly align with the left ventricle as long as some of the active surface was within the ventricle at initialization. This 1-2 second overhead at the start of US imaging is acceptable in the context of most interventional applications. This analysis regarding the segmentation method's relative insensitivity to initialization conditions is based on qualitative observations in a small number of experiments. A rigorous investigation of capture range will need to be performed in future studies.

In the current implementation, repositioning and reorientation accounts for an average of 66% of the overall processing time for each frame. This could be a bottleneck in the algorithm and a candidate for further optimization to achieve higher frame rates. Fortunately, it has been found in cardiac biomechanics²³ that the long axis of the left ventricle is relative stable through the cardiac cycle, therefore, it might be possible to turn off the reorientation/reposition step after the initial several seconds which are required for the model to adapt to actual image from its initial positions. This modification to the algorithm could provide a three-fold gain in the update rate.




Fig. 6. Segmented model at (a) initial state, (b) 0.5 seconds later, and (c) 2 seconds later. The arrow is indicating the rough orientation of the model's long axis in each frame.

4.3 Segmentation Speed

The update rate is currently about 5 frames per second, which is well below the maximum frame rate at which the RT3D US images can be acquired. The temporal performance of the segmentation can be increased by several means. In the current implementation, the 4D data stream was stored on the system hard disk. Therefore, the algorithm needed to load the data from the hard disk frame by frame, an effort that takes about 50 ms per frame. In a real clinical setting, the US data will be streaming into the system memory and will be processed in place; loading from the hard disk will not be required. Additionally, the time spent on visualization can also be decreased by using high performance graphics cards and rendering algorithms.

As pointed out in section 4.2, since repositioning and reorientation accounts for an average of 66% of the overall processing time for each frame, the computational load can potentially be cut by a factor of 3 by reducing the frequency of these steps once the model is roughly aligned with the anatomy. Also, since most of the operations in the current implementation are node-wise, the algorithm can be parallelized in order to utilize the efficiency of multi-threaded and parallel computing. Finally, running the whole framework on a more powerful hardware platform will reduce the overall processing time.

By taking all these factors into account, the update rate of the overall framework may be able to reach 20 frames per second or higher for true real-time performance.

4.4 Future Development

This work represents preliminary work toward extracting dynamic information from a 3D intra-procedural US stream. In addition to the endocardium's position and morphology, further valuable information may also be extracted using the current framework. Since the segmentation model is a parametric surface model based on finite elements, the segmentation process is actually a LV parameterization process. Once the segmentation converges, a patient specific finite element LV surface model is automatically built. Based on this model, model-based LV motion analysis²³ can be easily adapted into this framework. More detailed cardiac functional metrics, such as mechanical dyssynchrony, ejection fraction, local 3D fractional shortening, etc, may be directly computed or estimated from the model.

This framework can be also applied to other applications, such as tumor tracking for liver ablation guidance or tracking the atria or right ventricle in minimally invasive cardiac interventions. Ultimately, this segmentation method could be applied in the context of a multimodality guidance system using electromagnetic or optical tracking to localize the ultrasound probe for automatically and dynamically computing the registration with datasets from other imaging modalities.

5. CONCLUSIONS

This paper presents preliminary results and demonstrates the technical feasibility of augmenting static pre-procedural CT roadmap volumes with information extracted in near-real time from streamed RT3D cardiac ultrasound images acquired intra-procedurally. In the phantom experiments, the registration between the CT and US coordinate systems was accomplished with an RMSE of 3.8 mm. Segmentation of the phantom RT3D US image data was accomplished with an RMSE of 3.70 ± 2.5 mm. Experiments using clinical cardiac image data yielded 2.58 ± 1.58 mm RMS segmentation errors. The segmentation performed with an update rate of approximately 5 fps, but the majority of the processing time was used for loading the US data from disk and rendering the images within the display. The actual segmentation of each image frame required only 50 ± 8 ms and 75 ± 12 ms in the phantom and clinical images, respectively. This combination of static anatomical roadmap volumes and dynamic intra-operative anatomic information will enable better guidance and feedback for many image-guided minimally invasive interventions.

REFERENCES

1. M. R. M. Jongbloed, M. S. Dirksen, J. J. Bax, E. Boersma, K. Geleijns, H. J. Lamb, E. E. van der Wall, A. de Roos, and M. J. Schalij, "Atrial Fibrillation: Multi-Detector Row CT of Pulmonary Vein Anatomy prior to Radiofrequency Catheter Ablation--Initial Experience," vol. 234, 2005, pp. 702-709.
2. P. Schoenhagen, "Three-dimensional imaging for the guidance of coronary interventional procedures: impact on clinical decision making?," *The International Journal of Cardiovascular Imaging*, vol. 20, pp. 531-532, 2004.
3. O. Wink, R. Kemkers, S.-Y. J. Chen, and J. D. Carroll, "Intra-procedural coronary intervention planning using hybrid 3-dimensional reconstruction techniques," *Academic Radiology*, vol. 10, pp. 1433-1441, 2003.
4. C. M. Hillenbrand, D. R. Elgort, E. Y. Wong, A. Reykowski, F. K. Wacker, J. S. Lewin, and J. L. Duerk, "Active device tracking and high-resolution intravascular MRI using a novel catheter-based, opposed-solenoid phased array coil," *Magnetic Resonance in Medicine*, vol. 51, pp. 668-675, 2004.
5. O. T. V. Ramm and S. W. Smith, "Real time volumetric ultrasound imaging system," *Journal of Digital Imaging*, vol. 3, pp. 261-266, 1990.
6. Q. Duan, E. D. Angelini, and A. Laine, "Assessment of visual quality and spatial accuracy of fast anisotropic diffusion and scan conversion algorithms for real-time three-dimensional spherical ultrasound," in *SPIE International Symposium Medical Imaging*, 5373, 331-342, San Diego, CA, USA, 2004.
7. M. Mulet-Parada and J. A. Noble, "2D+T acoustic boundary detection in echocardiography," in *Medical Image Computing and Computer-Assisted Intervention-MICCAI'98*, 806-813, Cambridge, MA, 1998.
8. D. Boukerroui, O. Basset, A. Baskurt, and A. Noble, "Segmentation of echocardiographic data. Multiresolution 2D and 3D algorithm based on grey level statistics," in *MICCAI*, 516-523, Cambridge, UK, 1999.
9. J. G. Thomas, R. A. Peters, and P. Jeanty, "Automatic segmentation of ultrasound images using morphological operators," *IEEE Transactions on Medical Imaging*, vol. 10, pp. 180-186, 1991.
10. W. L. Jiang, Q. Liu, T. Wang, and Y. Zheng, "Edge detection method of Doppler echocardiography based on morphological transformation," in *Multispectral Image Processing and Pattern Recognition*, v 5286, 827-830, Beijing, China, 2003.
11. B. Vallet, E. D. Angelini, and A. F. Laine, "Variational segmentation framework in prolate spheroidal coordinates for 3D real-time echocardiography " in *Medical Imaging 2006: Image Processing*, 6144, 1370-1380, San Diego, CA, USA, 2006.
12. V. Chalana, D. T. Linker, D. R. Haynor, and Y. Kim, "A multiple active contour model for cardiac boundary detection on echocardiographic sequences," *IEEE Transactions on Medical Imaging*, vol. 15, pp. 290-298, 1996.
13. R. Drezek, G. D. Stetten, T. Ota, C. Fleishman, E. Lily, C. Lewis, C. J. Ohazama, T. Ryan, D. Glower, J. Kisslo, and O. T. v. Ramm, "Active contour based on the elliptical Fourier series, applied to matrix-array ultrasound of the heart," in *25th AIPR Workshop: Emerging Applications of Computer Vision*, 2962, 26-34, 1997.
14. Z. Tao and H. D. Tagare, "Stopping rules for active contour segmentation of ultrasound cardiac images," in *Medical Imaging 2005: Image Processing*, V 5747, 475-484, San Diego, CA, USA 2005.

15. E. Angelini, A. Laine, S. Takuma, J. Holmes, and S. Homma, "LV volume quantification via spatio-temporal analysis of real-time 3D echocardiography," *IEEE Transactions on Medical Imaging*, vol. 20, pp. 457-469, 2001.
16. N. Paragios, "A level set approach for shape-driven segmentation and tracking of the left ventricle," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 773-776, 2003.
17. R. Medina, A. Bravo, P. Windyga, J. Toro, P. Yan, and G. Onik, "A 2-D active appearance model for prostate segmentation in ultrasound images," in *27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* 3363-3366, 2005.
18. J. G. Bosch, S. C. Mitchell, B. P. F. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. H. C. Reiber, "Fully automated endocardial contour detection in time sequences of echocardiograms by three-dimensional active appearance models," in *Medical Imaging 2002: Image Processing*, v 4684, 452-462, San Diego, CA, USA 2002.
19. S. C. Mitchell, J. G. Bosch, B. P. F. Lelieveldt, R. J. Van der Geest, J. H. C. Reiber, and M. Sonka, "3-D active appearance models: Segmentation of cardiac MR and ultrasound images," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1167-1178, 2002.
20. J. Liang, T. McInerney, and D. Terzopoulos, "United Snakes," in *Seventh International Conference on Computer Vision (ICCV'99)*, 2, 933-940, 1999.
21. A. Montillo, D. Metaxas, and L. Axel, "Automated Segmentation of the Left and Right Ventricles in 4D Cardiac SPAMM Images," in *MICCAI 2002*, 620-633, 2002.
22. G. R. Christie, D. P. Bullivant, S. A. Blackett, and P. J. Hunter, "Modelling and visualising the heart," *Computing and Visualization in Science*, vol. 4, pp. 227-235, 2004.
23. S. Herz, C. Ingrassia, S. Homma, K. Costa, and J. Holmes, "Parameterization of left ventricular wall motion for detection of regional ischemia.," *Annals of Biomedical Engineering*, vol. 33, pp. 912-919, 2005.
24. P. M. Nielsen, I. J. Le Grice, B. H. Smaill, and P. J. Hunter, "Mathematical model of geometry and fibrous structure of the heart," vol. 260, 1991, pp. H1365-1378.
25. T. F. Chan and L. A. Vese, "Active Contours Without Edges," *IEEE Transactions on Image Processing*, vol. 10, pp. 266-277, 2001.
26. I. Mikic, S. Krucinski, and J. D. Thomas, "Segmentation and tracking in echocardiographic sequences: active contours guided by optical flow estimates," *IEEE transactions on medical imaging*, vol. 17, pp. 274-284, 1998.
27. C. Sheng, Y. Xin, Y. Liping, and S. Kun, "Segmentation in echocardiographic sequences using shape-based snake model combined with generalized Hough transformation," *International Journal of Cardiovascular Imaging*, vol. 22, pp. 33-46, 2006.
28. W. Liu, J. A. Zagzebski, T. Varghese, C. R. Dyer, U. Techavipoo, and T. J. Hall, "Segmentation of elastographic images using a coarse-to-fine active contour model," *Ultrasound in Medicine and Biology*, vol. 32, pp. 397-408, 2006.
29. "Insight Segmentation & Registration Toolkit, www.itk.org."
30. "The Visualization Toolkit, www.vtk.org."
31. Q. Duan, E. D. Angelini, S. L. Herz, O. Gerard, P. Allain, C. M. Ingrassia, K. D. Costa, J. W. Holmes, S. Homma, and A. F. Laine, "Tracking of LV Endocardial Surface on Real-Time Three-Dimensional Ultrasound with Optical Flow," in *Third International Conference on Functional Imaging and Modeling of the Heart 2005*, LNCS 3504, 434-445, Barcelona, Spain, 2005.
32. T. O'Donnell, S. Aharon, S. S. Halliburton, A. Gupta, G. Funka-Lea, and R. D. White, "Multi-modality model-based registration in the cardiac domain," in *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000* 2, 790-791, 2000.