Significant Acceleration of 2D-3D Registration-based Fusion of Ultrasound and X-ray Images by Mesh-based DRR Rendering

Markus Kaiser[ab], Matthias John[a], Anja Borsdorf[a], Peter Mountney[c], Razvan Ionasec[c], Alois Nöttling[a], Philipp Kiefer[d] and Thomas Neumuth[b]
[a]Siemens AG, Healthcare Sector, Forchheim, Germany
[b]Innovation Center Computer Assisted Surgery (ICCAS), University of Leipzig, Germany
[c]Siemens Corporate Research, Princeton, USA
[d]Department of Cardiac Surgery, Heart Center, University of Leipzig, Germany

ABSTRACT

For transcatheter-based minimally invasive procedures in structural heart disease ultrasound and X-ray are the two enabling imaging modalities. A live fusion of both real-time modalities can potentially improve the workflow and the catheter navigation by combining the excellent instrument imaging of X-ray with the high-quality soft tissue imaging of ultrasound. A recently published approach to fuse X-ray fluoroscopy with trans-esophageal echo (TEE) registers the ultrasound probe to X-ray images by a 2D-3D registration method which inherently provides a registration of ultrasound images to X-ray images. In this paper, we significantly accelerate the 2D-3D registration method in this context. The main novelty is to generate the projection images (DRR) of the 3D object not via volume ray-casting but instead via a fast rendering of triangular meshes. This is possible, because in the setting for TEE/X-ray fusion the 3D geometry of the ultrasound probe is known in advance and their main components can be described by triangular meshes. We show that the new approach can achieve a speed up factor of 18 and does not affect the registration accuracy when used in conjunction with the gradient correlation similarity measure. The improvement is independent of the underlying registration optimizer. Based on the results, a TEE/X-ray fusion can be performed with a higher frame rate and a shorter time lag towards real-time registration performance. The approach could potentially accelerate other applications of 2D-3D registrations, e.g. the registration of implant models with X-ray images.

Keywords: Mesh-based 2D3D-Registration, Ultrasound/X-ray Image Fusion, Image Guided Surgery

1. DESCRIPTION OF PURPOSE

More and more procedures in the field of structural heart disease become minimally invasive and catheter-based. This includes for instance trans-catheter aortic valve implantation, trans-catheter mitral valve repair, closure of atrial septal defects and left atrial appendage occlusion. The drivers for this trend from open-heart surgery to trans-catheter procedures are the availability of new catheter devices and the intra-procedural imaging. Usually, these procedures are performed under fluoroscopic X-ray and trans-esophageal echo (TEE) guidance. Intra-operatively these modalities are mainly used independently. An image fusion of both systems could yield a better mutual understanding of the image contents and potentially even allow new kinds of procedures.

Usually the images move relatively to each other, because their position is changed by the operator, but also because of patient-, heart- and breathing motion. Therefore there is a demand of an almost real-time update of the relative position of both images.

An approach for the fusion of ultrasound with fluoroscopic X-ray that does not require additional tracking hardware was recently suggested by Gao et al. [1]. A TEE probe is detected in the X-ray image and thereby derives the 3D position of the TEE probe relatively to the X-ray detector, which inherently provides a registration of the ultrasound image to the X-ray image. To estimate the 3D position, a model of the TEE probe is registered to the X-ray image via a 2D-3D registration algorithm [2]. The method does not need a modification of the ultrasound device and no specific set-up of the system for each procedure. But its main limitation is the runtime of a registration step which currently does not allow interactive registration updates for the image fusion. Therefore, there is a need to accelerate the generation of digital reconstructed radiographs (DRR) which is the most time-consuming part of the overall process.

2. METHODS

The task of an image registration of TEE and X-ray can be stated as a 2D-3D registration problem which is well-known in medical imaging [2]: The goal of a 2D-3D image-based registration system is to estimate the 3D-position of an object only with the available 2D information. Here, the 3D object is the rigid tip of the TEE probe and the 2D images are fluoroscopic X-rays. The final 3D-position of the object is computed by an iterative process starting with an initial 3D-position. In each iteration the similarity of the X-ray image and the projected image of the 3D object is computed based on a similarity measure. Here we use Gradient Correlation (*GC*) which is most suitable for our purpose, because it is sensitive to high image contrasts, e.g. the metal of a probe [3]. For the registration, a 2D projection of the 3D TEE probe model is obtained by simulating a 2D X-ray imaging process, called digitally reconstructed radiograph (DRR). The direction and step size for the three translational (t_x, t_y, t_z) and three rotational $(\theta_x, \theta_y, \theta_z)$ parameters are changed from iteration to iteration and are determined by an optimization algorithm. If the algorithm converges correctly, the correct registration is the final result. Practically, the most time consuming part of the registration process is the generation of the required DRRs which is usually done with ray-casting algorithms.

In contrast to this, we propose a new method for fast generation of DRR-like images based on triangle meshes. The underlying 3D volume of TEE probe was recorded with C-arm CT with a resolution of 512^3 voxel. From this volume we created two triangle meshes T_A and T_B by an isosurface-extraction algorithm. We used a high isovalue for T_A (representing structures with high contrast) and a low isovalue for T_B (which includes low contrast structures like the covering hull of the ultrasound probe). To improve the mesh quality, the generated meshes were preprocessed manually, e.g. by closing holes and removing metal artifacts. Additionally, the number of triangle faces is reduced significantly to 2500 faces per mesh to achieve a higher rendering performance.

To generate a DRR-like image, the two meshes are rendered with the projection geometry of the X-ray image. The X-ray-like translucent effect is obtained by alpha blending with different alpha values. The use of higher alpha values for T_A results in an enhancement of high-contrast structures in the final image. Comparative examples are given in Figure 1 for a ray-casting DRR and for a mesh-rendered DRR.

The *GC* similarity measure first computes the horizontal and vertical gradient images of X-ray and DRR images and then calculates the Normalized Cross Correlation (*NCC*) between the resulting vertical and horizontal gradient images. Determine that G_x and G_y denote the horizontal and vertical gradient images for the X-ray (I_a) and DRR (I_b) images, *GC* is then defined as:

$$GC(I_a, I_b) = NCC\big(G_x(I_a), G_x(I_b)\big)/2 + NCC\big(G_y(I_a), G_y(I_b)\big)/2$$

We use OpenGL to render the meshes and CUDA for fast calculation of the image gradients and the evaluation of *GC*. We also used the OpenGL interoperability of CUDA to provide the OpenGL rendered meshes directly as input of the following CUDA processing. The CUDA implementation was optimized specifically to the Gradient Correlation measure by parallelizing the gradient image computation and the following NCC with utilization of efficient texture and shared memory lookups and fast GPU reduction algorithms.

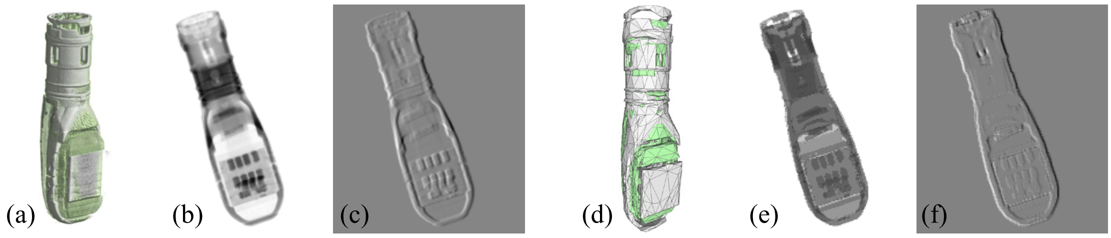

Figure 1: Comparison of volume-based DRR (a-c) and mesh-based renderings (d-f). (a) volume rendering, (b) ray-casting DRR, (c) vertical gradient image of the DRR, (d) mesh rendering, (e) mesh based DRR simulation, (f) vertical gradient image of the simulated DRR.

We compared our method with a conventional ray-casting approach regarding performance and accuracy of the DRR generation and similarity measure evaluation. The test datasets contained various sequences from real clinical cases and animal experiments. The X-ray images have a resolution of *512×512* pixels and were processed by full image edge detection and similarity measure calculation. The DRRs were generated on *512×512*, too. All experiments were performed on a mobile computer with 2.2 GHz Intel CPU (8 cores) and an NVIDIA QUADRO 1000M graphics cards (96 CUDA kernels). The volume-based DRRs were generated with the help of a ray-casting algorithm implemented in CUDA. Here, we calculated the *GC* similarity measure with the help of the open source computer vision library OpenCV.

We evaluated the accuracy of our method by using 1D plots of the similarity measure (see Figure 2). For all six degrees of freedom, the 3D model was stepwise translated and rotated away from the expected ground truth position of our test datasets.

3. RESULTS

The results of the comparison of both approaches for one of the images from the evaluation data set are illustrated in Figure 2. It shows the similarity measure values around the ground truth registration for the variation of two parameters for both the mesh-based DRR simulation approach and the volume-based ray-casting DRR approach. It can be seen that the local maximum for both approaches is close to the ground truth and differs just slightly between the two approaches. This holds for the in-plane parameters, like the translation along the horizontal axis (*X-translation*), as well as for the out-of-plane parameters, such as an out-of plane rotation (*X-rotation*).

For the runtime of the DRR generation based on volume ray-casting we measured a mean value of *45 ms* and for the standard implementation of the similarity measure a mean value of *20 ms*. The DRR simulation by mesh rendering takes in average *1.8 ms*, the CUDA-based evaluation of the GC similarity measure needs in average *1.7 ms*. Therefore, the overall evaluation time for our implementation of the mesh-based approach is *3.5 ms* compared to *65 ms* for the standard implementation of the DRR-based approach.

The results show that the mesh-based approach can yield the same registration accuracy but has a significantly faster runtime.

4. NEW OR BREAKTHROUGH WORK TO BE PRESENTED

We present a new algorithm to significantly accelerate the runtime of a 2D-3D registration algorithm to fuse ultrasound images and X-ray images towards real-time performance. It replaces the usual volume image based ray-casting DRR generation by a simulation of the DRRs by rendering of a triangular mesh model of the ultrasound probe.

We evaluated the runtime and registration accuracy performances for our new approach on clinical data sets. We show that the runtime is significantly accelerated compared to a volume based DRR approach. At the same time our mesh based approach shows the same accuracy when performed with the Gradient Correlation similarity measure.

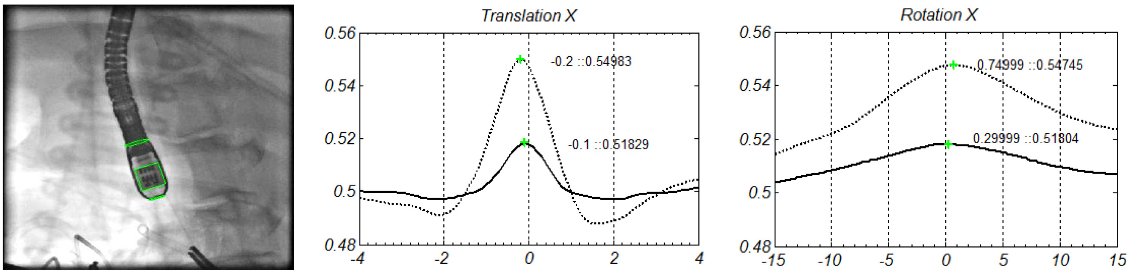

Figure 2: Left: Exemplary dataset with the final registration (indicated by the green structures). Right: Exemplary evaluation of the similarity measure for this image. The Y-axis shows the translation (in mm) / rotation (in degree) of the 3D model relative to the ground truth position (at 0). The X-axis shows the computed result of the *GC* similarity measure. Solid line: mesh rendered DRR, dotted line: ray-casting DRR.

The new algorithm is flexible to be used with any optimization method in the 2D-3D registration pipeline to finally compute a fusion of the images.

5. CONCLUSION

Real-time fusion of ultrasound and X-ray images for trans-catheter based could become a useful clinical tool for trans-catheter structural heart disease procedures. We propose a mesh based DRR simulation algorithm together with an optimized implementation on graphics hardware to accelerate the 2D-3D registration approach for this fusion problem. The algorithm achieves comparable evaluation results on the Gradient Correlation similarity measure and can therefore yield comparable registration accuracies.

Additional runtime optimization might be feasible to allow interactive frame rates also on lower-cost hardware. With an improved mesh generation our approach could even improve the registration accuracy, because the triangular mesh accuracy is independent of any voxel resolution. The algorithm could also be applied to other registration problems if a model of the 3D object is known before, e.g. for the registration of implants in trauma surgery.

DISCLAIMER

This work is nowhere being published. The concepts and information presented in this paper are based on research and are not commercially available.

REFERENCES

[1] G. Gao, G. Penney, Y. Ma, N. Gogin, P. Cathier, A. Arujana, G. Morton, D. Caufield, J. Gill, C. Aldo Rinaldi, J. Hancock, S. Redwood, M. Thomas, R. Razavi, G. Gijsbers, K. Rhode: Registration of 3D trans-esophageal echocardiography to X-ray fluoroscopy using image-based probe tracking. *Medical Image Analysis*, 16(1), Elsevier (2011)

[2] P. Markelj, D. Tomaževič, B. Likar, F. Pernuš, A review of 3D/2D registration methods for image-guided interventions, *Medical Image Analysis*, Volume 16 (3) (2012)

[3] G.P. Penney, J. Weese, J.A. Little, P. Desmedt, D.L. Hill, D.J. Hawkes: A comparison of similarity measures for use in 2-D-3-D medical image registration. *IEEE Transactions on Medical Imaging*, 17(4) (1998)