



GPU based real-time instrument tracking with three-dimensional ultrasound [☆]

Paul M. Novotny ^a, Jeff A. Stoll ^b, Nikolay V. Vasilyev ^c, Pedro J. del Nido ^c,
Pierre E. Dupont ^b, Todd E. Zickler ^a, Robert D. Howe ^{a,*}

^a *School of Engineering and Applied Sciences, Harvard University, 29 Oxford Street, Cambridge, MA, 02138, USA*

^b *Department of Aerospace and Mechanical Engineering, Boston University, USA*

^c *Department of Cardiovascular Surgery, Children's Hospital Boston, USA*

Received 2 April 2007; accepted 4 June 2007

Available online 5 July 2007

Abstract

Real-time three-dimensional ultrasound enables new intracardiac surgical procedures, but the distorted appearance of instruments in ultrasound poses a challenge to surgeons. This paper presents a detection technique that identifies the position of the instrument within the ultrasound volume. The algorithm uses a form of the generalized Radon transform to search for long straight objects in the ultrasound image, a feature characteristic of instruments and not found in cardiac tissue. When combined with passive markers placed on the instrument shaft, the full position and orientation of the instrument is found in 3D space. This detection technique is amenable to rapid execution on the current generation of personal computer graphics processor units (GPU). Our GPU implementation detected a surgical instrument in 31 ms, sufficient for real-time tracking at the 25 volumes per second rate of the ultrasound machine. A water tank experiment found instrument orientation errors of 1.1° and tip position errors of less than 1.8 mm. Finally, an in vivo study demonstrated successful instrument tracking inside a beating porcine heart.

© 2007 Elsevier B.V. All rights reserved.

Keywords: 3D Ultrasound; Cardiac Surgery; GPU; Instrument tracking

1. Introduction

Real-time three-dimensional ultrasound has been demonstrated as a viable tool for guiding minimally invasive surgery (Cannon et al., 2003). For example, beating heart intracardiac surgery is now possible with the use of three-dimensional ultrasound and minimally invasive instruments (Suematsu et al., 2004). These techniques eliminate the need for cardio-pulmonary bypass and its well-documented adverse effects, including delay of neural development in children, mechanical damage produced by

inserting tubing into the major vessels, stroke risk, and significant decline in cognitive performance (Murkin et al., 1999; Zeitlhofer et al., 1993; Bellinger et al., 1999). Cannon et al. (2003) showed that complex surgical tasks, such as navigation, approximation, and grasping are possible with 3D ultrasound. However, initial animal trials revealed many challenges to the goal of ultrasound guided intracardiac surgery (Suematsu et al., 2004). In a prototypical intracardiac procedure, atrial septal defect closure, an anchor driver was inserted through the cardiac wall to secure a patch covering the defect (Fig. 1). Surgeons found it difficult to navigate instruments with 3D ultrasound in the dynamic, confined intracardiac space. Tools look incomplete and distorted (Fig. 1), making it difficult to distinguish and orient instruments.

To address this issue, researchers are developing techniques to localize instruments in ultrasound. Enhancing

[☆] Supported by United States National Institutes of Health (R01 HL073647-01).

* Corresponding author.

E-mail address: howe@seas.harvard.edu (R.D. Howe).

URL: www.biorobotics.harvard.edu (R.D. Howe).




Fig. 1. Schematic (a) and ultrasound image (b) of an atrial septal defect repair. In this procedure a patch is inserted into the heart to cover the atrial septal defect. An anchor driver is used to attach the patch to the septum.

the displayed position of the instrument allows surgeons to more accurately control the instruments as they perform surgical tasks. In addition, real-time tracking of instruments in conjunction with a surgical robot opens the door for a range of enhancements, such as surgical macros, virtual fixtures, and other visual servoing techniques (Kragic et al., 2005; Shinsuk et al., 2001).

Previous work in instrument detection can be broadly separated into two categories: external tracking systems such as electromagnetic and optical tracking (Leotta, 2004; Lindseth et al., 2003) and image based detection algorithms (Draper et al., 2000; Novotny et al., 2003; Zhouping et al., 2004; Ortmaier et al., 2005). External tracking systems have suffered from the limitations of the surgical environment. Electromagnetic tracking has limited accuracy and is problematic to implement due to the abundance of ferro-magnetic objects in the operating room. Optical tracking of instruments is complicated by line-of-sight requirements. Both of these systems suffer from errors introduced by improper registration of the ultrasound image coordinates to the tracking coordinate frame. To eliminate such errors, image based algorithms are used to track instruments within the ultrasound image. Most of this work focused on tracking needles (Draper et al., 2000; Novotny et al., 2003; Zhouping et al., 2004) and more recently surgical graspers (Ortmaier et al., 2005) in 2D ultrasound images. As 3D ultrasound systems have become widely available, these 2D techniques have been adapted for implementation in 3D. An appealing approach to instrument localization is the Radon or Hough transform. These techniques have wide spread use in 2D image analysis for detection of wide variety of shapes. Applications of these techniques have focused on detecting 2D objects in 2D images, however, Hough and Radon based techniques have shown promise in 3D medical image analysis. Most relevant is a needle tracking technique for prostate biopsy (Ding and Fenster, 2003) that projects the ultrasound volume onto two orthogonal planes. A Hough transform is then performed on the two 2D images to identify the needle.

Beating heart intracardiac procedures pose different challenges and requirements than the 2D breast and pros-

tate biopsy procedures in previous work. For example, the high data rates of 3D ultrasound machines, 30–40 MB/s, require very efficient algorithms for real-time implementation. Previous 2D ultrasound techniques are too computationally costly, or inappropriate for three dimensions. These methods are only appropriate for finding bright objects such as needles in ultrasound images that stand out amongst relatively homogeneous tissue. In cardiac procedures, larger instruments such as anchor drivers and graspers are used that do not stand out amongst the surrounding dynamic, heterogeneous environment. To work in this environment, the algorithm must be efficient for handling the large data rates, and capable of distinguishing instruments from fast-moving cardiac structures of similar intensity.

2. Methods and materials

In this work we present a technique capable of detecting instruments used in minimally invasive procedures, such as endoscopic graspers, staplers, and cutting devices. Instruments used in minimally invasive procedures are fundamentally cylindrical in shape and typically 3–10 mm in diameter, a feature that is not found in cardiac tissue. We use a form of the Radon transform to identify these instruments within the ultrasound volumes. In the following sections we describe a generalization of the Radon transform that is appropriate for identifying instrument shafts in 3D ultrasound volumes. Furthermore, we show that this technique can be implemented on a parallel architecture such as inexpensive PC graphics hardware, enabling the detection of instruments in real-time. The proposed method is examined in both tank studies and an in vivo animal trial.

2.1. The Radon transform

The Radon and Hough transforms are widely used for detecting lines in two dimensional images. In its original formulation (Radon, 1917), the Radon transform maps an image, I , into the Radon space $\mathcal{R}\{I\}$

$$\mathcal{R}\{I\}(d, \phi) = \int I(d \cos(\phi) - s \sin(\phi), d \sin(\phi) + s \cos(\phi)) ds. \quad (1)$$

d and ϕ define a line in two dimensions by its perpendicular distance to the origin and slope, respectively. s is a free parameter that corresponds to a specific point on the shape. This method transforms a difficult image processing problem into simply identifying maximums in the Radon space. In its original formulation, the Radon transform is unsuitable for instrument tracking in 3D, but has been extended for identification of arbitrary shapes. Using the notation introduced by Luengo Hendriks et al. (2005), points on a parametric shape are defined by the function $\mathbf{c}(s, \mathbf{p})$, where \mathbf{p} defines the parameters of the shape. In this framework, the Radon transform is rewritten as

$$\mathcal{R}_{\mathbf{c}(s,\mathbf{p})}\{I\}(\mathbf{p}) = \int I(\mathbf{c}(s,\mathbf{p}))ds. \quad (2)$$

A line segment in 3D is parametrized with six variables. (x_0, y_0, z_0) defines the center of the line segment in 3D. Two angular parameters, (θ, ϕ) , describe its orientation, and L defines its length. More compactly, these six parameters are written as

$$\mathbf{p} = [x_0, y_0, z_0, \theta, \phi, L]. \quad (3)$$

Points lying on a line segment are now defined for \mathbf{p} and s as

$$\mathbf{c}(s,\mathbf{p}) = \begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} + \begin{bmatrix} \cos(\theta) \cos(\phi) \\ \sin(\theta) \cos(\phi) \\ \sin(\phi) \end{bmatrix} \left(\frac{L}{2}\right)s, \quad s \in [-1, 1]. \quad (4)$$

Combining Eqs. (2) and (4) yields a form of the generalized Radon transform for line segments in 3D volumes

$$\mathcal{R}_{\mathbf{c}(s,\mathbf{p})}\{I\}(\mathbf{p}) = \int_{-1}^1 I \left(\begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} + \begin{bmatrix} \cos(\theta) \cos(\phi) \\ \sin(\theta) \cos(\phi) \\ \sin(\phi) \end{bmatrix} \left(\frac{L}{2}\right)s \right) ds. \quad (5)$$

Identifying lines in the 3D volume now becomes a problem of finding local maximums of $\mathcal{R}_{\mathbf{c}(s,\mathbf{p})}\{I\}(\mathbf{p})$, from Eq. (5), where \mathbf{p} denotes a local maximum or likely instrument position. In other words, we integrate the image volume, I , along a direction defined by (θ, ϕ) , through the point (x_0, y_0, z_0) and identify maximums. This scheme is illustrated in Fig. 2 where integrations are illustrated for multiple directions. Fig. 2c contains points with a high integral, or




Fig. 2. Example of the Radon transform detection of the instrument axis. Each image (a–e) is a projection of the ultrasound image along the corresponding direction shown in the schematic. The projection along the axis of the instrument (c) is the brightest. Note that this diagram omits out-of-plane projections that are part of the implementation.

simply, the image is brighter than the other four images (Fig. 2a, b, d, and e). This is a result of the correspondence of the integration direction and the object's axis. As a result, by finding the maximum value of $\mathcal{R}_{\mathbf{c}(s,\mathbf{p})}\{I\}(\mathbf{p})$, the axis of the instrument in 3D space is implicitly defined by the parameters \mathbf{p} .

2.1.1. Graphics card implementation

One of the most promising features of the Radon transform approach described above is the potential for real-time computation. Eq. (5) can be independently calculated for each $[x_0, y_0, z_0, \theta, \phi, L]$ and is ideally suited for computation on parallel architectures; current personal computer graphics cards are built with features well-suited for this application. Many researchers have shown that highly parallel calculations, when implemented on these graphics processor unit (GPU) cards, show significant performance advantages over CPU based implementations (Wu et al., 2004; Oh and Jung, 2004; Kruger and Westermann, 2003). Using a similar approach, we programmed a PC graphics card (7800GTX, nVidia Corp., Santa Clara, CA) to calculate the necessary integrations, $\mathcal{R}_{\mathbf{c}(s,\mathbf{p})}\{I\}(\mathbf{p})$.

Before the algorithm is executed, data are transferred to the GPU: the sequential values of \mathbf{p} to be searched are pre-loaded into five textures in the GPU memory, where each texture corresponds to each parameter $[x_0, y_0, z_0, \theta, \phi]$. The three-dimensional ultrasound data is loaded onto the graphics card into a three-dimensional texture. Once the data is loaded, the algorithm runs by 'rendering' to an output texture. Sixteen parallel pipelines on the graphics card (programmable pixel shaders) calculate the integral defined in Eq. (5) by stepping through the volume simultaneously a length of L . The pixel shaders perform the integrations for each input parameter set defined in the input textures, and output the results to the corresponding position in the output texture. Tri-linear interpolation, implemented in hardware, is used as each pixel shader integrates the image intensity in the volume along the direction defined by (θ, ϕ) through the point (x_0, y_0, z_0) . The output texture is then transferred back to main memory for post-processing by the CPU. On the CPU, the maximum intensity in the output texture is identified as the instrument axis position.

2.1.2. Passive markers

Once the axis of the instruments is found, it is necessary to detect the final 2 degrees of freedom of the instrument (tip position and roll angle) to fully define its position and orientation. To this end, we build on work first introduced by Stoll and Dupont (2005). Here we use a new marker design, shown in Fig. 3. To produce distinct elements, 800 μm polyurethane foam was wrapped around the instrument shaft as shown in Fig. 3. Uncoated metals such as the stainless steel used for surgical instruments are highly reflective in ultrasound. Huang et al. (2006) showed that uncoated metal instruments that were more than 20° from perpendicular to the ultrasound probe reflected almost no ultrasound energy back to the ultrasound probe. To ensure




Fig. 3. Picture (a) and ultrasound image (b) of a minimally invasive anchor driver with passive markers. The instrument tip and roll angle is calculated using the distances x_1 and x_2 .

that the instrument is visible in ultrasound, a more diffusive interaction with the ultrasound pulse is desired. As a result, an 80 μm fiberglass embedded PTFE coating was applied to the instrument in order to improve the appearance.

Finding the markers begins with the image volume already loaded into texture memory from the Radon transform algorithm, then the built-in tri-linear interpolation is used to quickly render a slice through the instrument axis, orientated so that the instrument axis is horizontal. To identify the position of the bumps, a template matching algorithm is used on the ultrasound slice. The algorithm uses the sum of the absolute differences between a candidate region of the slice and a template, shown in Fig. 4a.

The positions of the three best matches found in the slice are used to determine the tip position of the instrument and the roll angle (Fig. 4b). The tip position is found with a known offset of 3 mm between the two closest markers and the instrument tip. To find the roll angle, the ratio of the distances x_1 and x_2 are used (Fig. 3a). Since the third marker is wrapped in a helical pattern around the instrument shaft, the roll angle is a linear function of this ratio.




Fig. 4. (a) Passive marker template used to identify the location of the three markers on the instrument. (b) 3D ultrasound image of a surgical instrument in a water tank. White dots indicate tracked passive markers.

2.1.3. System architecture

For instrument tracking, the ultrasound data is produced by a Sonos 7500 ultrasound system (Philips Medical, Andover, MA). The ultrasound volumes are streamed from the ultrasound machine to a personal computer over a gigabit LAN using TCP/IP. The data stream is captured on the target PC and passed to the instrument tracking algorithm running in a separate asynchronous thread. The tracking algorithm calculates the modified Radon transform on the graphics processing unit (7800GTX, nVidia Corp, Santa Clara, CA) using DirectX 9.0c. The entire system runs on a Pentium 4 3 GHz personal computer with 2 GB of RAM.

For the first ultrasound volume, Eq. (5) is calculated for evenly spaced points throughout the volume space I . Spatially, the volume is sampled at 5 voxel increments in x , y , and z . With an ultrasound resolution of 0.5 mm per voxel, this search space corresponds to a 5 mm cube. For angles θ and ϕ , Eq. (5) is sampled in 10° increments. Due to symmetry, the angles are only sampled from 0 to 180° . This search constitutes the initialization of the instrument tracking, as the entire volume is searched for the instrument.

For subsequent frames, the tracking algorithm confines its search space to an area centered on the location found in the previous frame. Since the ultrasound volumes are updated at 25–28 Hz, this search space can be fairly small. In our trials, we empirically found that limiting the search space to ± 5 voxels spatially in the x , y , and z directions and $\pm 10^\circ$ around the angles θ and ϕ found in the previous frame was sufficient to capture typical surgical movements.

2.2. Experimental validation

To validate the proposed methodology, two sets of experiments were conducted. The first trial measured the accuracy of the algorithm by imaging instruments in a controlled tank environment. While this study carefully characterized the accuracy of the method, it does not reflect the target conditions for the algorithm, detecting instruments within a beating heart. As a result, a second study was necessary to validate the effectiveness of the technique in a surgical setting.

2.2.1. Tank study

The accuracy of the proposed method was measured with a system that precisely positioned and oriented the instruments within the ultrasound field. The instruments were connected to a 3-axis translational stage with 1 μm resolution, and two rotational stages with a resolution of 5 min. Registration to the ultrasound coordinates was done by using a flat-plate phantom and two-wire phantom (Prager et al., 1998). Registration accuracy was 0.4 mm for positions and 0.6° for angular measurements.

A series of ultrasound images were taken by varying the orientation angles of the instrument ϕ_{instr} . As shown in Fig. 5, this angle refers to rotations about the x -axis of

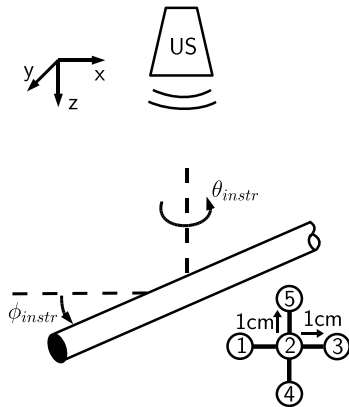


Fig. 5. In the tank study, instruments were imaged at different ϕ_{instr} . Images were taken of the instrument at five positions for each orientation.

the ultrasound image, identical to x defined in the modified Radon transform (Eq. (5)). ϕ_{instr} ranged from 0 to 60° in 5° increments. It was physically not possible to image the instrument beyond 60° because of the size of the field of view and dimensions of the ultrasound probe. This constraint exists in surgical situations also and therefore orientations beyond 60° were not considered in this study. At each angular orientation, the instrument was imaged in five different positions within the ultrasound field. The first position had the instrument tip in the center of the image, and the other four positions were each 1 cm from this initial center position in the axial and lateral directions (Fig. 5). As a result, five images were taken for each ϕ_{instr} .

2.2.2. In vivo animal study

Second, in vivo validation was performed by tracking an instrument within a beating porcine heart. The instruments were imaged inside a porcine heart during an open chest beating heart procedure. The instruments were inserted through ports in the right atrial wall and secured by purse-string sutures. The ultrasound probe was positioned epicardially on the right atrium to give a view of the right and left atrium. The surgeon was instructed to move the instrument toward the atrial septum. During this movement, the instrument tip position calculated by the algorithm was recorded to a data file. Electromagnetic tracking (miniBIRD 800, Ascension Technology, Burlington, VT) was used to simultaneously track the instrument tip position.

3. Results

The results of the tank studies demonstrate the accuracy of the method. In Fig. 6, the angular accuracy of the method is shown for different orientations of the instrument with respect to the ultrasound probe. Fig. 6 shows that for angles from 0 to 60° of ϕ_{instr} , the instrument tracking algorithm accurately determined its orientation. Across all trials, the RMS difference of the angle calculated by the tracking algorithm and the angle measured by the testing

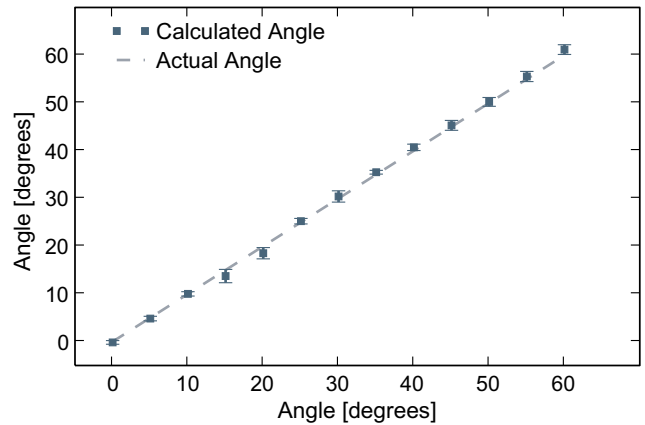


Fig. 6. The plot shows the mean angle ϕ_{instr} calculated by the tracking algorithm for all five position. Errors bar indicate standard deviation.

setup was 1.07°. There was no dependence on the accuracy of the algorithm and the orientation of the instrument.

The tip position was also tracked with the algorithm in the tank study. The tip position was compared to the actual tip position recorded by the 5-axis positioning stage. The results are shown in Fig. 7 for each orientation angle, where each data point corresponds to one of the five positions of the instrument and orientation angle. It was found that the tip position accuracy was extremely dependent on the accuracy of the passive marker identification algorithm. When the algorithm correctly identified the positions of the three passive markers, the tip position had an RMS error of 1.8 mm. However, for eight of the trials, the algorithm mis-identified the marker location. As a result, the tip position error was greater than 5 mm.

Fig. 8 shows the results of the in vivo trials. The figure plots the x , y , and z trajectories of the instrument tip versus time. Both the instrument tip position reported by the instrument tracking algorithm and the electromagnetic trackers is shown. The tracking method correctly tracked the instrument tip as the surgeon moved it for 5 s. The RMS difference between the tip position reported by the electromagnetic tracker and tracking algorithm was




Fig. 7. Tip position error from tank trials for each angle ϕ_{instr} .




Fig. 8. In vivo x , y , and z position of the instrument tip as reported by an electromagnetic tracker (dashed line) and the tracking algorithm (solid line).

1.4 mm, 0.5 mm, and 0.9 mm in x , y , and z directions. Fig. 9 graphically shows the results of the instrument tracking in a beating porcine heart. The figure shows an overlay of the tracked instrument position on the 3D ultrasound volume over 5 s. For each of the images, the overlay correctly matches the position of the instrument as the heart beats around it.

In our experimental setup, the instrument tracking technique required 1.7 s to initially detect the instrument in the

entire ultrasound volume. For subsequent tracking, the algorithm required 32 ms per volume. This speed is well within the 38 ms required for the algorithm to keep pace with the 25 volumes per second generated by the ultrasound machine. This performance is a significant improvement over implement at ion on a CPU. When running the same algorithm on a CPU (Pentium 4, 3 GHz), the algorithm took 11.7 s for the initial detection and 0.54 s for subsequent tracking of the instrument. As a result, the GPU based tracking approach is 14 times faster than a CPU based approach. However, the 32 ms does include the required time to produce and transfer the ultrasound volume by the ultrasound machine. This time was typically 30 ms.

4. Discussion

This paper demonstrates for the first time real-time tracking of surgical instruments in intracardiac procedures with 3DUS. The algorithm was both capable of distinguishing instruments from fast-moving cardiac structures and efficient enough to work in real-time. The generalized Radon transform is effective here because it integrates over the length of the instrument shaft to minimize the effects of noise and spatial distortion in ultrasound images. By taking advantage of the unique shape of surgical instruments, we were able to correctly distinguish instruments from cardiac structures of similar intensity. The Radon transform was also selected because of its amenability to parallel implementation, which was exploited by using GPUs to achieve real-time performance. As a result, it was possible to detect an instrument in the ultrasound volume in 32 ms,




Fig. 9. Ultrasound images of the instrument inside a beating porcine heart. The red dots indicate the instrument position calculated by the tracking algorithm. Each image (a–f) are images taken each second for 6 s.

which is sufficient to handle the 25 volumes per second produced by the ultrasound machine.

The approach demonstrated good performance in both tank trials and an in vivo study. These studies accomplished two separate and complimentary goals. The tank trials permitted rigorous exploration of the accuracy of the technique in a controlled environment with precision that is not possible in a surgical setting. The in vivo trials, on the other hand, tested the technique's ability to correctly distinguish an instrument from tissue in the presence of highly dynamic aberrations and clutter typical of cardiac ultrasound imaging. Rigorously characterizing accuracy is difficult in this setting, due to the limited range of instrument motion and access to the intracardiac space.

The tank experiments showed that the Radon transform algorithm is very accurate. Across all the trials, it correctly identified the orientation of the instrument to within 1° . By adding passive markers, the full 6 degree of freedom tracking was possible. This is of importance to surgical applications because it enables instrument tip tracking. For most cases, the accuracy of the technique was within 1.8 mm. However, as seen in Fig. 7, the marker detection algorithm sometimes incorrectly identifies the position of the passive markers. As a result, the tip position is incorrectly calculated along the shaft axis. While not the focus of this work, the passive marker detection and tracking is being studied by parallel efforts by Stoll and Dupont (2005). Integration of improved passive marker tracking with the Radon based shaft tracking presented here, will be addressed in future work.

In vivo trials provided verification of the effectiveness of the algorithm when instruments are surrounded by blood and highly inhomogeneous and rapidly moving tissue. The electromagnetic tracking used for verification is by no means a "gold standard"; however, the data showed the tracking technique presented here is in fact following the surgical instrument with reasonable accuracy (Fig. 9).

With real-time instrument tracking techniques, it is now possible to introduce guidance enhancements to aid in our target procedures, intracardiac surgery. Real-time tracking can now be used for instrument overlays and navigational aids to help the surgeon deal with the distorted appearance of tissue and instruments in 3DUS images. In addition, tracking of instruments opens a wide range of possibilities for implementing robot control under 3DUS guidance.

References

- Bellinger, D.C., Wypij, D., Kuban, K.C., Rappaport, L.A., Hickey, P.R., Wernovsky, G., Jonas, R.A., Newburger, J.W., 1999. Developmental and neurological status of children at 4 years of age after heart surgery with hypothermic circulatory arrest or low-flow cardiopulmonary bypass. *Circulation* 100, 526–532.
- Cannon, J.W., Stoll, J.A., Salgo, I.S., Knowles, H.B., Howe, R.D., Dupont, P.E., Marx, G.R., del Nido, P.J., 2003. Real-time three-dimensional ultrasound for guiding surgical tasks. *Computer Aided Surgery* 8, 82–90.
- Ding, M., Fenster, A., 2003. A real-time biopsy needle segmentation technique using hough transform. *Medical Physics* 30, 2222–2233.
- Draper, K., Blake, C., Gowman, L., Downey, D., Fenster, A., 2000. An algorithm for automatic needle localization in ultrasound-guided breast biopsies. *Medical Physics* 27, 1971–1979.
- Huang, J., Dupont, P.E., Undurti, A., Triedman, J.K., Cleveland, R.O., 2006. Producing diffuse ultrasound reflections from medical instruments using a quadratic residue diffuser. *Ultrasound in Medicine and Biology* 32, 721–727.
- Kragic, D., Marayong, P., Li, M., Okamura, A., Hager, G., 2005. Human-machine collaborative systems for microsurgical applications. *International Journal of Robotics Research* 24, 731–741.
- Kruger, J., Westermann, R., 2003. Acceleration techniques for GPU-based volume rendering. In: *Proceedings of IEEE Visualization*, pp. 287–292.
- Leotta, D.F., 2004. An efficient calibration method for freehand 3D ultrasound imaging systems. *Ultrasound in Medicine and Biology* 30, 999–1008.
- Lindseth, F., Tangen, G.A., Lango, T., Bang, J., 2003. Probe calibration for freehand 3D ultrasound. *Ultrasound in Medicine and Biology* 29, 1607–1623.
- Luengo Hendriks, C.L., Van Ginkel, M., Verbeek, P.W., Van Vliet, L.J., 2005. The generalized Radon transform: sampling, accuracy and memory considerations. *Pattern Recognition* 38, 2494–2505.
- Murkin, J.M., Boyd, W.D., Ganapathy, S., Adams, S.J., Peterson, R.C., 1999. Beating heart surgery: why expect less central nervous system morbidity? *The Annals of Thoracic Surgery* 68, 1498–1501.
- Novotny, P., Cannon, J., Howe, R., 2003. Tool localization in 3D ultrasound images. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention (MICCAI'03)*, pp. 969–970.
- Oh, K., Jung, K., 2004. GPU implementation of neural networks. *Pattern Recognition* 37, 1311–1314.
- Ortmaier, T., Vitrani, M., Morel, G., Pinault, S., 2005. Robust real-time instrument tracking in ultrasound images for visual servoing. In: *Proceedings of the International Conference on Robotics and Automation (ICRA'05)*, pp. 2167–2172.
- Prager, R.W., Rohling, R.N., Gee, A.H., Berman, L., 1998. Rapid calibration for 3D freehand ultrasound. *Ultrasound in Medicine and Biology* 24, 855–869.
- Radon, J., 1917. Über die bestimmung von funktionen durch ihre integralwerte langs gewisser mannigfaltigkeiten. *Berichte Sachsische Akademie der Wissenschaften, Leipzig, Mathematisch-Physikalische Klasse* 69, 262–277.
- Shinsuk, Park., Robert D, Howe., David F, Torchiana., 2001. Virtual fixtures for robotic cardiac surgery. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention (MICCAI'01)*, pp. 1419–1420.
- Stoll, J., Dupont, P., 2005. Passive markers for ultrasound tracking of surgical instruments. In: *Proceedings of the Medical Image Computing and Computer-Assisted Intervention (MICCAI'05)*, vol. 8, pp. 41–48.
- Suematsu, Y., Marx, G.R., Stoll, J.A., DuPont, P.E., Cleveland, R.O., Howe, R.D., Triedman, J.K., Mihaljevic, T., Mora, B.N., Savord, B.J., Salgo, I.S., del Nido, P.J., 2004. Three-dimensional echocardiography-guided beating-heart surgery without cardiopulmonary bypass: a feasibility study. *The Journal of Thoracic and Cardiovascular Surgery* 128, 579–587.
- Wu, E., Liu, Y., Liu, X., 2004. An improved study of real-time fluid simulation on GPU. *Computer Animation and Virtual Worlds* 15, 139–146.
- Zeitlhofer, J., Asenbaum, S., Spiss, C., Wimmer, A., Mayr, N., Wolner, E., Deecke, L., 1993. Central nervous system function after cardiopulmonary bypass. *European Heart Journal* 14, 885–890.
- Zhouping, W., Gardi, L., Downey, D., Fenster, A., 2004. Oblique needle segmentation for 3D trus-guided robot-aided transperineal prostate brachytherapy. In: *Proceedings of the 2nd IEEE International Symposium on Biomedical Imaging: Macro to Nano*, vol. 1, p. 960.