
Radon Transform Based Instrument Shaft Detection in Three-Dimensional Ultrasound

Paul M. Novotny, Todd E. Zickler, Robert D. Howe
Division of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
novotny@fas.harvard.edu
1-617-496-9098

***Abstract.** Real-time three-dimensional ultrasound enables new intra-cardiac 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 modified Radon transform to search for long straight objects in the ultrasound image, a feature that is characteristic of instruments and not found in cardiac tissue. Experimental validation shows mean accuracy is within half a degree in orientation and within 2 mm in position. In vivo trials demonstrate effective detection with instruments inserted into a beating heart. In addition, the algorithm is shown to be amenable to parallel architectures such as PC graphics hardware resulting in near real-time instrument detection in three-dimensional ultrasound.*

Introduction

Real-time three-dimensional ultrasound has been demonstrated as a viable tool for guiding surgical procedures [1]. This imaging technique enables a range of new minimally invasive cardiac procedures. For example, beating heart intracardiac procedures are now possible with the use of three-dimensional ultrasound and minimally invasive instruments [2]. These procedures eliminate the need for cardio-pulmonary bypass and its well documented adverse effects, including delay of neuro development in children, mechanical damage produced by inserting tubing into the major vessels, stroke risk, and significant decline in cognitive performance [3-5]. Cannon *et al.* [1] found complicated 3D tasks, such as navigation, anchoring, and grasping, are possible with 3D ultrasound. However, initial animal trials highlighted many challenges to the goal of ultrasound guided intracardiac surgery [2]. Most notable is the appearance of surgical instruments in ultrasound. Tools look incomplete and distorted (Figure 1), making it difficult to identify and orient instruments during procedures.

To address this issue, researchers have developed techniques to detect instruments in ultrasound. By detecting instruments as they move within the ultrasound image it is possible to highlight the position of the instrument for the surgeons. Eventually, this will allow surgeons to more accurately control the instruments as they perform a surgical task. In addition, real-time tracking of instruments in conjunction with a surgical robot opens the door for a range possibilities, such as surgical macros, virtual fixtures, and other visual servoing techniques. Previous work in instrument detection can be broadly separated into two categories: external tracking systems such as electromagnetic and optical tracking [6,7] and image-based detection algorithms [8-11]. External tracking systems have suffered from the limitations of the surgical environment. Electromagnetic tracking is difficult to implement due to the abundance of ferro-magnetic objects in the operating room. In addition, optical tracking of

instruments is complicated by line-of-sight issues. 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 have been developed to track instruments within the ultrasound image [8-11]. However, these methods have been restricted to two-dimensional ultrasound images or are only appropriate for finding bright objects such as needles in an ultrasound images that stand out amongst relatively homogeneous tissue. These techniques are not appropriate for the dynamic, heterogeneous environment of intracardiac surgery.

In this work we present a technique capable of detecting instruments used in minimally invasive procedures, such as endoscopic graspers, staplers, and cutting devices. Most instrument shafts used in minimally invasive procedures are fundamentally cylindrical in shape, typically 3-10 mm in diameter. Our technique identifies instruments in 3D ultrasound by looking for these long straight cylinders, a feature that is not found in cardiac ultrasound images. We use a modified form of the Radon transform, a common imaging technique, to identify these instruments within the ultrasound volumes. In the following sections we derive a form of the Radon transform that is appropriate for identifying instrument shafts in 3D ultrasound volumes. We also introduce a search schema used to quickly distinguish instruments from cardiac tissue. 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 near real-time. The accuracy of the proposed method is examined in both tank studies and an *in vivo* animal trial.

Methods and Materials

1 The Radon Transform

When looking for linear objects in 2D images, the Radon transform is a powerful tool [12]. However, the traditional Radon transform is not capable of detecting lines in three dimensions as it

reduces to a plane detector and is not well suited for this application. As a result, a modified 2D Radon transform is needed for the detection of linear objects in 3D ultrasound volumes.

As described in [12], it is possible to generalize the 2D Radon transform into three dimensions, thereby creating a modified three dimensional Radon transform for line detection. To this end, we start with the parametric form of the 2D Radon transform where \check{g} is the Radon transform of the 2D image g and is computed using

$$\check{g}(\theta, \rho) = \int g(s\tau(\theta) + \rho\alpha(\theta)) ds. \quad (1)$$

Where θ and ρ are a slope and perpendicular distance from the origin that describe a line. τ and α are functions of θ denoting the direction of the line of integration and the perpendicular distance of that line from the origin, respectively.

$$\tau(\theta) = \begin{pmatrix} \cos \theta \\ \sin \theta \end{pmatrix}, \quad \alpha(\theta) = \begin{pmatrix} -\sin \theta \\ \cos \theta \end{pmatrix} \quad (2)$$

In three dimensions, the direction of the line of integration is a function of two parameters, θ and ϕ , and we define

$$\tau(\theta, \phi) = \begin{pmatrix} \cos \theta \cos \phi \\ \sin \theta \cos \phi \\ \sin \phi \end{pmatrix}. \quad (3)$$

In addition, two orthogonal vectors are needed to describe the offset of the line from the origin in three dimensions.

$$\alpha(\theta, \phi) = \begin{pmatrix} -\sin \theta \\ \cos \theta \\ 0 \end{pmatrix}, \quad \beta(\theta, \phi) = \begin{pmatrix} -\cos \theta \sin \phi \\ -\sin \theta \sin \phi \\ \cos \phi \end{pmatrix} \quad (4)$$

As shown in Figure 2, τ , α , and β form an orthonormal basis whereby a line in three dimensions is defined. Using this framework, it is possible to generalize the 2D Radon transform (Equation 1) into three dimensions.

$$\check{g}(\theta, \phi, u, v) = \int_{-\infty}^{\infty} g(s\tau(\theta, \phi) + u\alpha(\theta, \phi) + v\beta(\theta, \phi)) ds \quad (5)$$

Equation 5 describes a transform in which each point in $\check{g}(\theta, \phi, u, v)$ corresponds to the integral of $g(r)$ along a three dimensional line defined by the four parameters θ , ϕ , u , and v . It is not strictly speaking a Radon transform; however, it closely relates to Equation 1, the 2D Radon transform, and can be seen as a generalization of this transform into three dimensions.

Identifying lines in the 3D volume now becomes a problem of finding local maximums of $\check{g}(\theta_m, \phi_m, u_m, v_m)$, from Equation 5, where $(\theta_m, \phi_m, u_m, v_m)$ denotes a local maximum or likely instrument position. In other words, we integrate the volume, $g(r)$, along a direction defined by (θ, ϕ) and identify maximums in the integrated plane (u, v) . This scheme is illustrated in Figure 3 where integrations are illustrated for multiple directions. Figure 3c contains points with a high integral, or simply, the image is brighter than the other four images (Figure 3a, 3b, 3d, and 3e). This is a result of the correspondence of τ_c and the object's axis. In addition, the position of the maximum in the integrated plane gives two positional coordinates of the axis of the rod. As a result, by finding the maximum value of \check{g} , the axis of the instrument in 3D space is implicitly defined by the parameters $(\theta_m, \phi_m, u_m, v_m)$.

2 Implementation

The modified 3D Radon transform outlined above provides a basis for detection of instruments in 3D ultrasound. Calculating this transform in its basic form is computationally costly. However, the

algorithm is especially well suited to many performance improvements that enable near real-time instrument detection. In this implementation, we took the following three steps to harness characteristics of the algorithm that enable high performance instrument detection. The first step is to break up the large ultrasound volume into more manageable subvolumes. Secondly, a course-to-fine search is used to broadly sample the \checkmark space, and then zero-in on areas that contain a maximum. Finally, because the algorithm is especially well suited for parallel computing, we harnessed the parallel architecture found on the latest generation graphics processing units.

2.1 Subvolumes

The first step is to break up the large 3D ultrasound volumes into smaller regions for analysis. The reason is two-fold. First, subvolumes that do not contain an instrument are disregarded while promising subvolumes are analyzed in greater detail. Second, smaller subregions are more likely to contain an instrument in blood without large sections of tissue that would otherwise obscure the instrument. Our graphics card implementation requires the sub-volume size to be a power of two; however, we empirically found that subvolumes of 32x32x32 were sufficient because instrument shafts used in this study were 22 voxels in diameter. As a result, the ultrasound volume, typically 128x64x204 voxels, is divided into spherical subvolumes with a diameter of 32. The volume was sampled by selecting subvolumes in 16 voxel intervals in all three directions. A spacing of 16 voxels guarantees complete sampling of the volume, as well as ensuring an instrument axis lies near the center of at least one sphere.

2.2 Coarse to Fine Search

The performance of the algorithm was further improved with a coarse-to-fine search that samples the Radon space (Equation 5). Each subvolume is first searched coarsely by projecting every 10

degrees (0 to 170 degrees) for both θ and ϕ , resulting in an 18x18 search space. Subvolumes are then ranked by the maximum \check{g} to determine which subvolume is likely to contain an instrument.

It is necessary to eliminate erroneous subvolumes from this ranking that may produce bright projections, but in fact are not instruments. Large segments of tissue, or even instrument artifacts, introduce areas that could be misidentified. However, the consistent appearance of the cross-section of instruments (Figure 3c) is used to eliminate non-instruments in our algorithm. By determining the shape of this cross-section *a priori*, we are able to eliminate candidates that are not consistent with the instrument appearance. In this implementation, we use a simple template matching scheme to calculate the correspondence of the candidate projection with the template. The normalized cross correlation (NCC) provides a measure of the similarity between the template and the candidate projections

$$NCC = \left(\frac{1}{\sigma_P \sigma_Q} \right) (P - \bar{P}) \cdot (Q - \bar{Q}). \quad (6)$$

Here P and Q are defined as the candidate and template images, with mean \bar{P} and \bar{Q} , and variance σ_P and σ_Q . Empirically we found that a $NCC > 0.85$ was sufficient to distinguish instruments from non-instruments. In addition, this threshold was found to be insensitive to different imaging conditions as it was used throughout the study.

The highest ranked subvolume that passed the template matching test is chosen for a finer sampling of the Radon space. To this end, the space centered on $(\theta_m, \phi_m, u_m, v_m)$ found in the coarse search is then more finely sampled at 1 degree increments within $\pm 10^\circ$. As a result, the axis of the instrument is found at the maximum $\check{g}(\theta, \phi, u, v)$, thereby implicitly identifying the instruments axis $(\theta_m, \phi_m, u_m, v_m)$ within 1 degree for θ and ϕ .

2.3 Graphics Card Implementation

One of the most promising features the modified Radon transform is the potential for real-time computation. Calculation of Equation 5 can be independently calculated for each (θ, ϕ, u, v) and is ideally suited for computation on parallel architectures. Modern graphics cards are built with features well suited for this application as they were designed to meet the high demands of the consumer gaming industry. Many researchers have shown that highly parallel calculations, when implemented on these graphics cards, show significant performance advantages over CPU based implementations [13-15]. Using a similar approach, we were able to program these graphics processors to calculate the necessary integrations, \check{g} , and thereby harness the available speed of this hardware.

In this application, a 7800GTX PC graphics card (nVidia Corp., Santa Clara, CA) was used. The three-dimensional ultrasound data was loaded onto the graphics card as a three-dimensional texture. Sixteen parallel pipelines on the graphics card (programmable pixel shaders) calculate the integral defined in Equation 5 by stepping through the volume for sixteen (θ, ϕ, u, v) simultaneously. Trilinear interpolation, implemented in hardware, is used as the algorithm discretely sums the intensity values along the line.

3 *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. Both trials used a Sonos 7500 ultrasound system (Philips Medical, Andover, MA) to produce the ultrasound images.

3.1 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 three-axis translational stage with 1 μm resolution, and 2 rotational stages with a resolution of 5 minutes. Registration to the ultrasound coordinates was done by using a flat-plate phantom and two-wire phantom [16]. Registration accuracy was 0.4 mm for positions and 0.6 degrees for angular measurements.

The instruments used in the procedure were chosen to mimic minimally invasive instruments. Multiple material types and coatings were used in order to compare their effects on the appearance of the instruments in ultrasound, but a focused investigation into suitable coatings is left for future work. In this study, we chose material types and coatings representative of those found by this ongoing area of research. The instruments consisted of cylindrical rods 5-6 mm in diameter with five different material types and coatings. Shown in Table 1, a stainless steel tube was tested with 3 different coatings: no coating, covered in polymer coated fabric gaffer's tape, and coated in an acrylic based dip (0.55 mm thickness). In addition to the stainless steel, cylinders made from solid acetyl and wood were used.

Table 1: Materials used in tank and in vivo study.

Material	Coating	Diameter	Inside Diameter
Stainless Steel	-	5.2mm	4.4mm
Stainless Steel	tape	5.2mm	4.4mm
Stainless Steel	acrylic	5.2mm	4.4mm
Acetyl	-	6.3mm	-
Wood	-	6.3mm	-

Two series of ultrasound images were taken by varying the two orientation angles of the instrument (θ_{Inst} and ϕ_{Inst}) independently. As shown in Figure 4, these angles refer to rotations about

the z and x axis of the ultrasound image, identical to the θ and ϕ defined in the modified Radon transform (Equation 5). ϕ_{Inst} ranged from 0 to 90 degrees in 10 degree increments. In the case of θ_{Inst} , symmetry allowed for limiting the study to 0 to 45 degree in 5 degree increments. 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 direction lateral directions (Figure 4). As a result, 5 images were taken for each θ and ϕ .

3.2 In Vivo Animal Study

An animal study tested the algorithm under *in vivo* conditions. In this study, the above 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 consistency of the algorithm in both detecting the instrument and reporting its orientation was studied. The instruments were held statically within 0.22 degrees and 0.04mm by a small robot (Omni Phantom, Sensible Technologies, Woburn, MA) during two cardiac cycles, about 50 volumes. This data set provided images of an instrument as the heart moved around it, testing the algorithm with varying configurations of cardiac tissue. In the absence of absolute calibration between ultrasound and robot coordinates, relative measurements were obtained.

Results

In the tank trial, the algorithm had varying degrees of success depending on the instrument material and coating. In the best case, shown in Figure 5, tape coated stainless steel was detected

accurately for a wide range of angles. For $\phi_{Inst} \leq 50$ degrees, the mean fell within 0.5 degrees of the measured angle, with a standard deviation of 1 degree. However, beyond 50 degrees the algorithm was unable to accurately find the instrument within the ultrasound image with errors over 20 degrees. For θ_{Inst} , the algorithm did not fare as well. The detected angle was on average only 3 degrees of the measured angle with standard deviation of the sample generally around 5 degrees.

Acetyl, wood, and acrylic-coated stainless steel instruments (Figures 6, 7, and 8) also were successfully detected, but for a much narrower range of orientations than the tape-coated stainless steel instrument. For acetyl, ϕ_{Inst} estimates began to deteriorate for angles greater than 45 degrees, while wood and acrylic coated stainless steel suffered above 30 and 20 degrees, respectively. As with tape-coated stainless steel, θ_{Inst} displayed a much larger standard deviation; however, the errors were largely independent of angle.

The positional accuracy of the algorithm was also measured by comparing the distance between the position (u, v) found by the algorithm and actual location of the axis measured with the five-axis positioning system. The error was calculated by finding the perpendicular distance between this calculated point (u, v) in the projection plane and the measured instrument axis. With the tape coating, mean errors were minimal, between 0.4 and 3.1 mm depending on angle (Figure 9). When ϕ_{Inst} was 20 degrees or less the mean positional error remained below 1 mm. As this angle increased, positional accuracy degraded as was also seen with angular accuracy. The estimated positional accuracy of acetyl, wood, and acrylic coated stainless steel instruments were found to be less than tape-coated stainless steel. These materials also showed a similar dependence on ϕ_{Inst} , as this angle increased accuracy decreased. However, for 20° or less all materials were accurate to within 2 mm and for 60° or less the accuracy remained below 6mm.

For the *in vivo* experiment in a beating porcine heart, Figure 10 shows the standard deviation of the ϕ_{Inst} and θ_{Inst} measurements given by the algorithm. In addition, the fraction of the 50 images in which the algorithm identified the incorrect subvolume is shown. Similar to results found in the tank studies, the performance of the algorithm varied depending on the material and coating used. The stainless steel with a tape covering performed well: the algorithm correctly distinguished the instrument amongst the dynamic cardiac tissue for all 50 images. The standard deviation of ϕ_{Inst} and θ_{Inst} estimates for the fixed instrument was 1.4 and 2.4 degrees, respectively. Wood and acetyl did not fare as well. The algorithm misidentified segments of tissue as the instrument, resulting in failure for wood and acetyl in 15% and 54% of the acquired images, respectively. For wood, the algorithm consistently calculated the orientation of the instrument incorrectly. The standard deviation of ϕ_{Inst} and θ_{Inst} was 7.8 degrees and 5.2 degrees. Acetyl did much better in its measurements, with ϕ_{Inst} and θ_{Inst} standard deviations of 5.2 and 2.8. However, this includes only the 54% of the images where the instrument was detected.

The performance of the algorithm was found to be well within the range of a real-time implementation. On a 3GHz Pentium 4 with 2 gigabytes of RAM and the nVidia 7800GTX graphics card, the algorithm took 0.6 seconds to search an entire 128x48x204 ultrasound volume.

Discussion

These results confirm the utility of the algorithm for detecting instruments in three-dimensional ultrasound images. With appropriate instrument coatings it is possible to detect rods made of the same material (stainless steel tubes) and same shaft geometry (5-6 mm cylinders) as instruments commonly used in minimally invasive procedures. This was demonstrated by first estimating the accuracy of the

method in tank studies, and then by demonstrating its ability to correctly detect the instrument in an *in vivo* study. These two studies taken together demonstrate the effectiveness of the algorithm for detecting instruments inside a beating heart using 3D ultrasound.

In addition, we presented an implementation that performed this technique in near real-time speeds. This was accomplished by using a simple search technique combined with implementing the algorithm on a powerful and inexpensive graphics processor. The current implementation was able to search an entire ultrasound volume in 0.6 seconds, which is very near real-time speeds. However, the search technique is not optimal and further study is necessary to find more efficient and effective search schemes. Future work will focus on improving performance with the following strategies: incorporating a more sophisticated search technique to find the maximum of \check{g} ; restricting the search volume within a radius of the instrument location found in the previous frame; and upgrading the hardware implementation. With these improvements, it will be possible to detect the instrument at the speeds necessary to handle the 25 volumes per second produced by the three-dimensional ultrasound machine.

Despite its success, the algorithm was not without its drawbacks. As shown in all of the tank studies, regardless of material or coating, the algorithm was unable to calculate θ_{Instr} to the degree of accuracy of ϕ_{Instr} . This is a result of the shape of the instrument when imaged by ultrasound. As shown in Figure 3c, the cross-section is larger in the lateral (x) direction than the axial (z) direction. As a result, local maximums of \check{g} are more sensitive to changes in ϕ_{Instr} which leads to increased accuracy in identifying the correct orientation. With the right material or coating, however, this effect is greatly reduced as was seen with tape covered stainless steel.

The algorithm also highlighted the importance of coatings and material selection. Coatings had a large effect on misidentification and accuracy, as seen in Figure 10. This occurred because the

assumption that the cross-section of the instruments is consistent in ultrasound does not hold for certain materials and coatings. The tape covering on stainless steel produced a near uniform angular scattering of the acoustic energy and therefore had a consistent cross-section in ultrasound and was identified accurately. The other materials' cross-sections were much more variable. As a result, the template matching portion of the technique was unable to distinguish the instrument from the surrounding cardiac tissue. To address this issue, future work will explore the use of coatings and materials similar to the tape used in this experiment. Emphasis will be placed on material types and coatings that are medically viable.

The largest limitation of this technique is the inability to identify all 6 degrees of freedom of the instrument. Consequently, we were unable to fully characterize the performance of the technique in detecting the instruments' tip location and roll angle about the instrument shaft, an important requirement of instrument detection. However, we are working to add a second complimentary detection technique proposed by Stoll *et al.* [17]. This method introduces a series of passive markers to the instrument shaft that are used to determine the final two degrees of freedom. By combining both methods a full detection of the six degrees of freedom of an instrument will be possible.

Real-time instrument tracking promises to be a great aid to surgical procedures performed with three-dimensional ultrasound. This paper presents a necessary first step to this goal by developing a high performance algorithm capable of accurately identifying the axis of the instrument. It demonstrates that detecting instruments in near real-time is possible and achieving a real-time implementation is not far off. With this tool, there are many new possibilities to aid surgical procedures, such as highlighting instruments for surgeons and tracking instruments for controlling a surgical robot.

Figure 1:




Figure 2:




Figure 3:




Figure 4:




Figure 5:

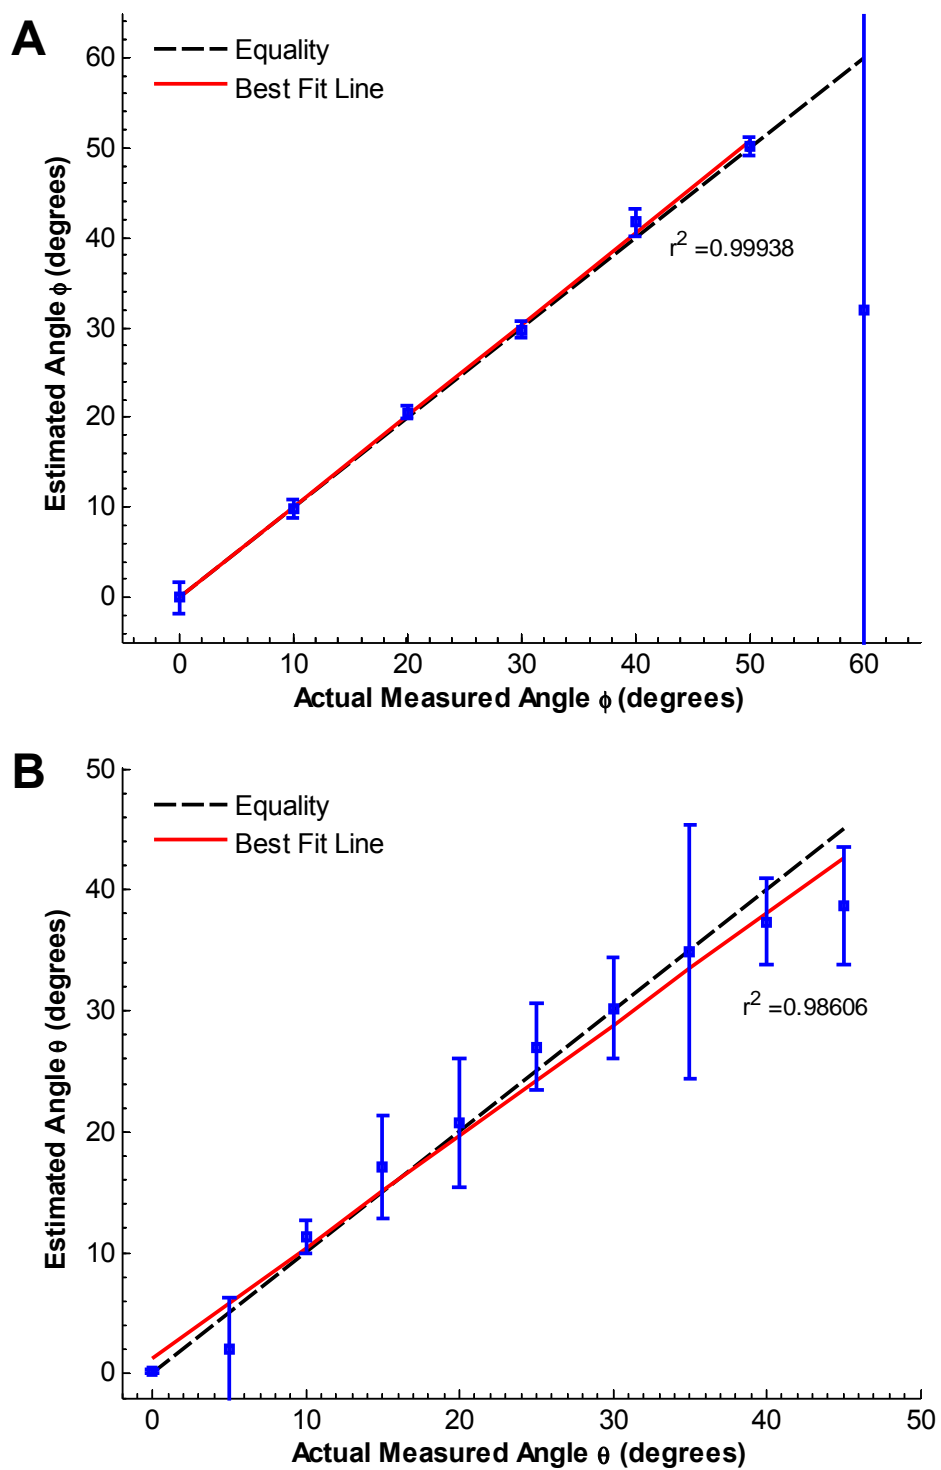


Figure 6:

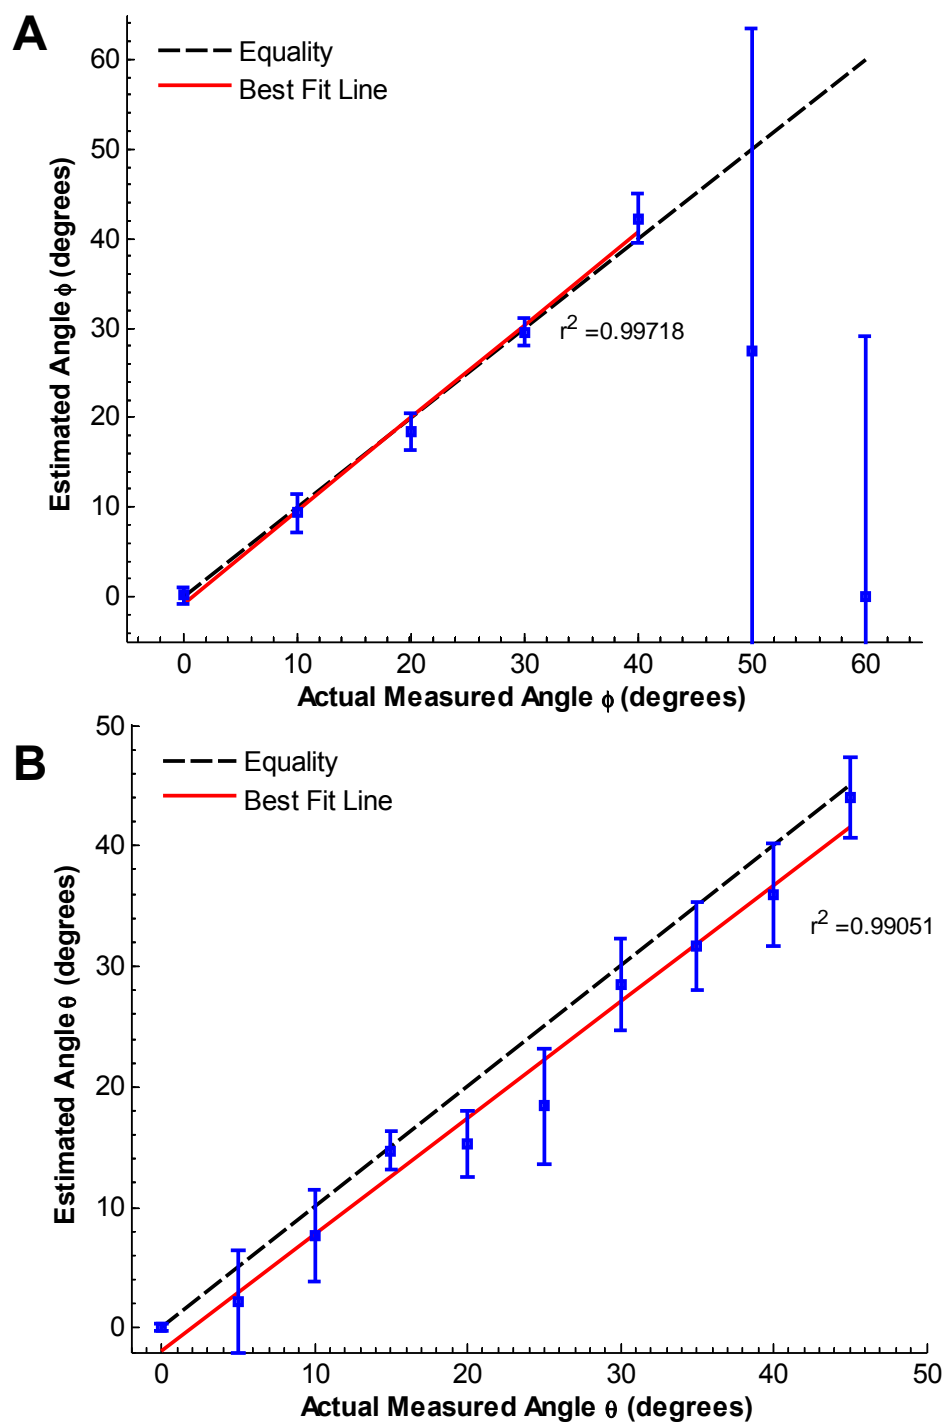


Figure 7:

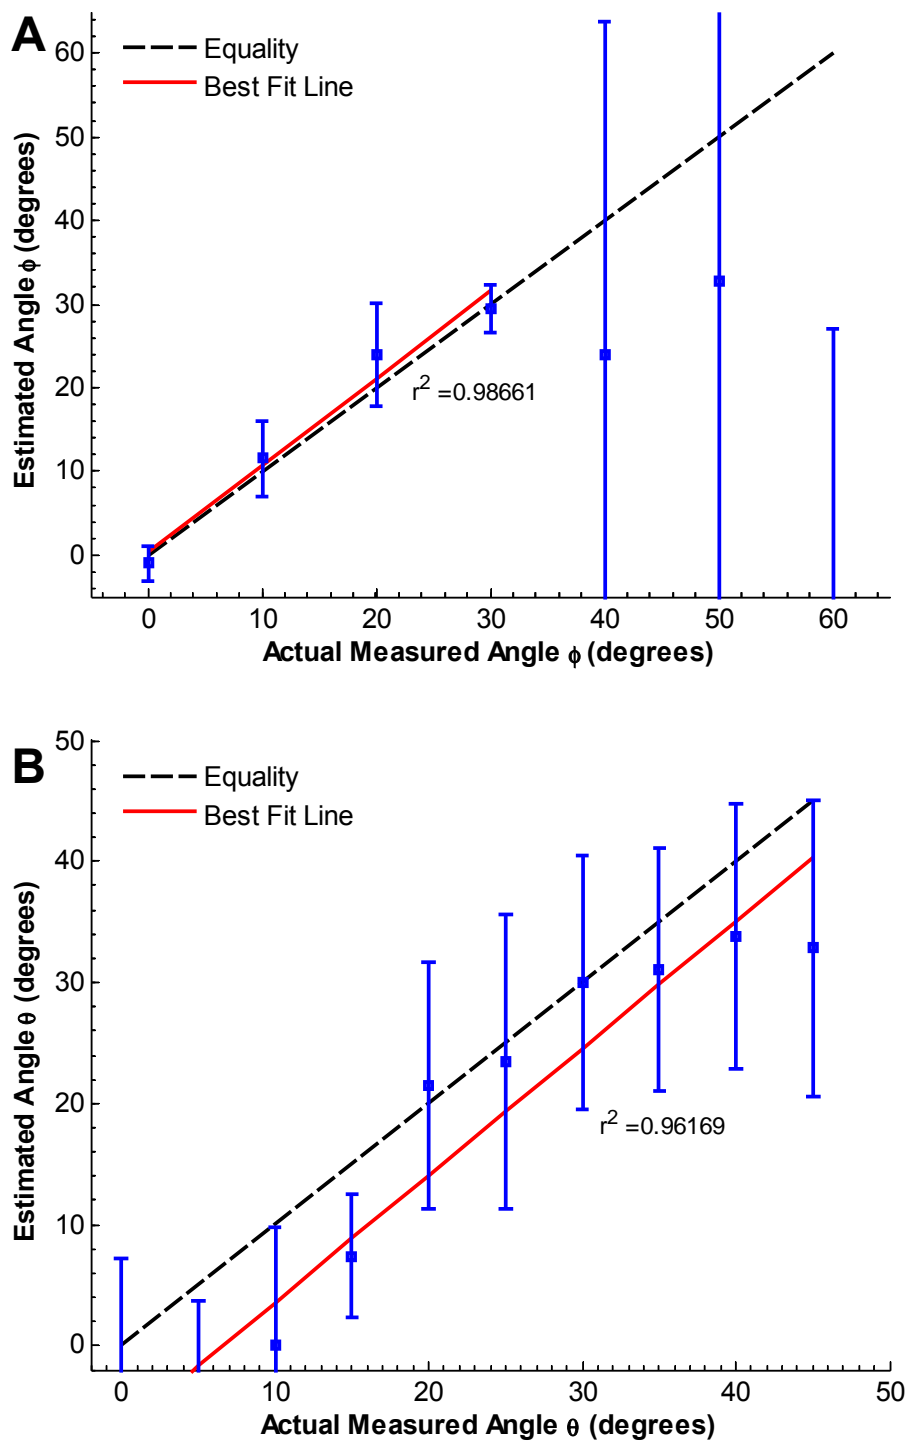


Figure 8:

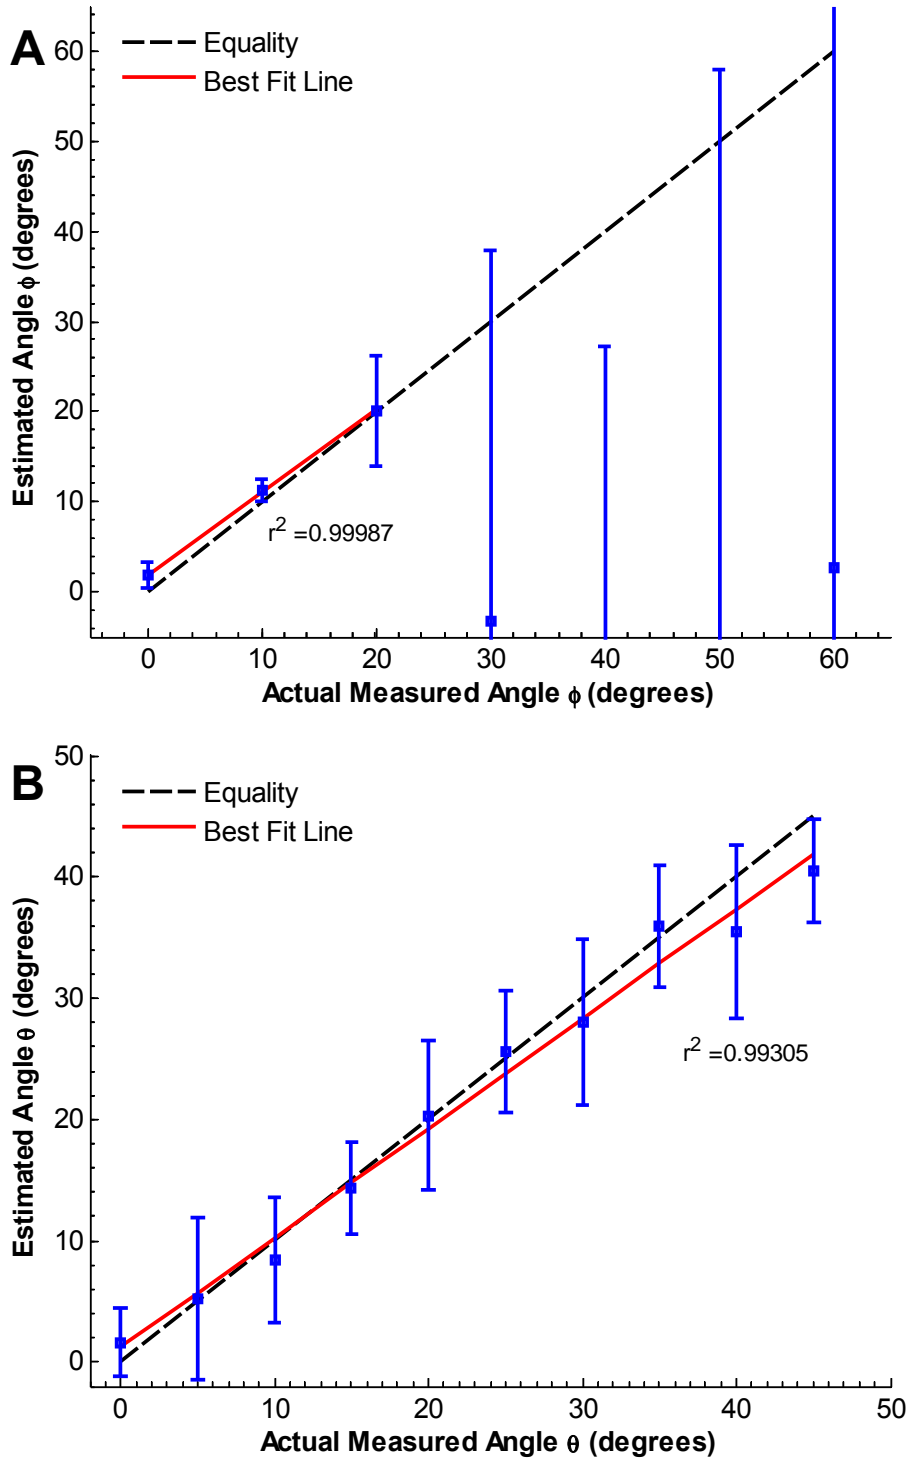


Figure 9:

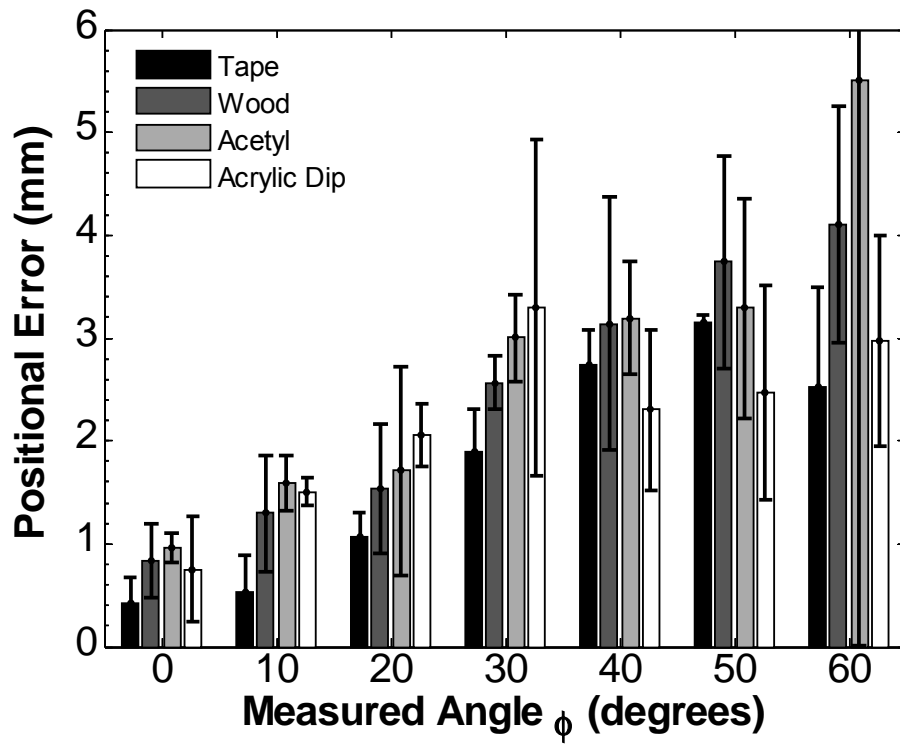


Figure 10:

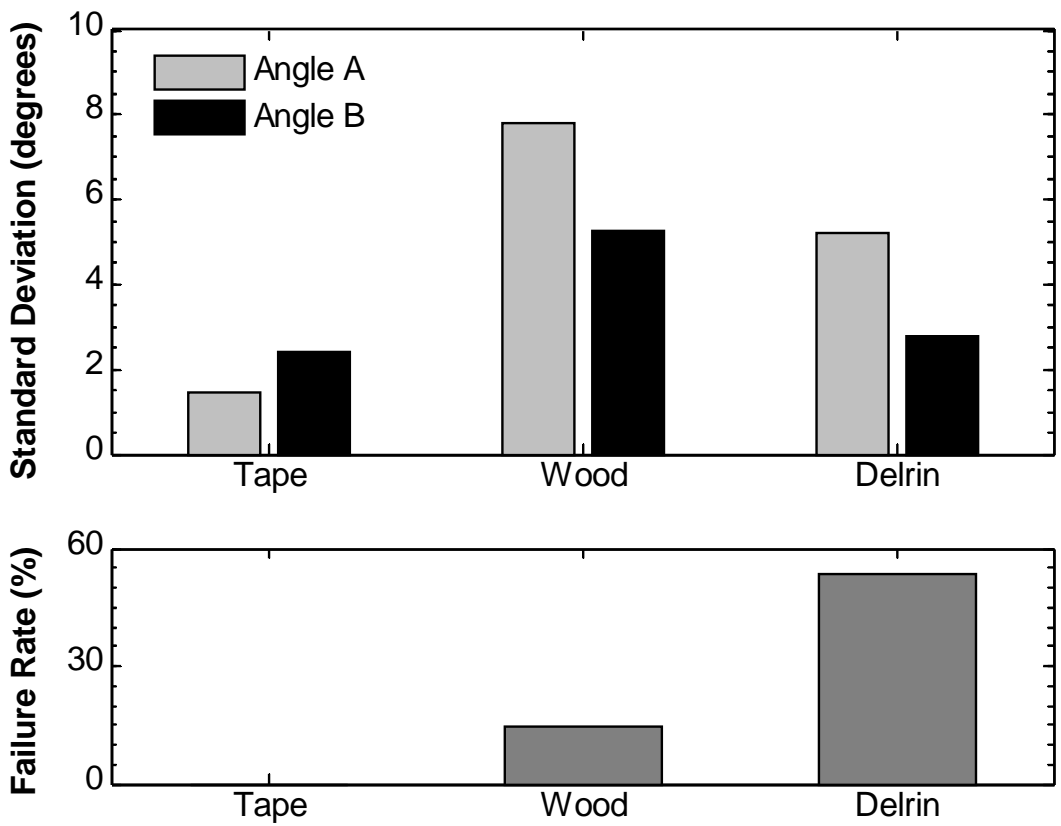


Figure Captions

FIGURE 1

Ultrasound images of a tape-coated stainless steel rod. (A) Rod immersed in water tank; note large tip artifact. (B) Same rod inserted into a porcine heart. Arrows indicate instrument shaft. Ultrasound probe is located at the top of the image.

FIGURE 2

Schematic of the modified Radon transform in 3D. The transform integrates along a line (shown as dashed), defined by the vectors β , α , and τ .

FIGURE 3

Example of the modified Radon transform performed on the ultrasound image shown in Figure 1A. Each image (A-E) is a projection of the ultrasound image along the corresponding direction shown in the schematic. The projection is 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.

FIGURE 4

In the tank study instruments were imaged at different θ_{Inst} and ϕ_{Inst} . Images were taken of the instrument at five positions for each orientation.

FIGURE 5

Tank results for stainless steel coated in tape for angles ϕ (A) and θ (B). For ϕ the fit does not include points beyond 50 degrees because the algorithm fails past this point.

FIGURE 6

Tank results for acetyl for angles ϕ (A) and θ (B). For ϕ the fit does not include points beyond 40 degrees because the algorithm fails past this point.

FIGURE 7

Tank results for wood for angles ϕ (A) and θ (B). For ϕ the fit does not include points beyond 30 degrees because the algorithm fails past this point.

FIGURE 8

Tank results for stainless steel coated in acrylic dip for angles ϕ (A) and θ (B). For ϕ the fit does not include points beyond 20 degrees because the algorithm fails past this point.

FIGURE 9

Positional error of algorithm for four materials: stainless steel coated in tape, wood, acetyl, and stainless steel coated in acrylic. Error bars are standard deviation.

FIGURE 10

In vivo results of instruments detection within a beating porcine heart. The standard deviation of the calculated angles is given for 50 images. In addition, the rate at which the algorithm misidentified the instrument is shown.

References

- [1] Cannon JW, Stoll JA, Salgo IS, Knowles HB, Howe RD, Dupont PE, Marx GR, del Nido PJ, "Real-time three-dimensional ultrasound for guiding surgical tasks," *Comput Aided Surg*, vol. 8, no. 2, pp. 82-90, 2003.
- [2] Suematsu Y, Marx GR, Stoll JA, DuPont PE, Cleveland RO, Howe RD, Triedman JK, Mihaljevic T, Mora BN, Savord BJ, Salgo IS, del Nido PJ, "Three-dimensional echocardiography-guided beating-heart surgery without cardiopulmonary bypass: a feasibility study," *J Thorac Cardiovasc Surg*, vol. 128, no. 4, pp. 579-87, 2004.
- [3] Murkin JM, Boyd WD, Ganapathy S, Adams SJ, Peterson RC, "Beating heart surgery: why expect less central nervous system morbidity?," *Ann Thorac Surg*, vol. 68, no. 4, pp. 1498-501, 1999.
- [4] Zeitlhofer J, Asenbaum S, Spiss C, Wimmer A, Mayr N, Wolner E, Deecke L, "Central nervous system function after cardiopulmonary bypass," *Eur Heart J*, vol. 14, no. 7, pp. 885-90, 1993.
- [5] Bellinger DC, Wypij D, Kuban KC, Rappaport LA, Hickey PR, Wernovsky G, Jonas RA, Newburger JW, "Developmental and neurological status of children at 4 years of age after heart surgery with hypothermic circulatory arrest or low-flow cardiopulmonary bypass," *Circulation*, vol. 100, no. 5, pp. 526-32, 1999.
- [6] Leotta DF, "An efficient calibration method for freehand 3-D ultrasound imaging systems," *Ultrasound Med Biol*, vol. 30, no. 7, pp. 999-1008, 2004.
- [7] Lindseth F, Tangen GA, Lango T, Bang J, "Probe calibration for freehand 3-D ultrasound," *Ultrasound Med Biol*, vol. 29, no. 11, pp. 1607-23, 2003.
- [8] Ortmaier T, Vitrani M, Morel G, Pinault S, "Robust real-time instrument tracking in ultrasound images for visual servoing," *Proceedings of International Conference on Robotics and Automation*, 2005.
- [9] Draper K, Blake C, Gowman L, Downey D, Fenster A, "An algorithm for automatic needle localization in ultrasound-guided breast biopsies," *Medical Physics*, vol. 27, no. 8, pp. 1971-1979, 2000.
- [10] Novotny P, Cannon J, Howe R, "Tool localization in 3D ultrasound images," *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003. 6th International Conference. Proceedings. Part II*, pp. 969-970, 2003.
- [11] Zhouping W, Gardi L, Downey D, Fenster A, "Oblique needle segmentation for 3D TRUS-guided robot-aided transperineal prostate brachytherapy," *2004 2nd IEEE International Symposium on Biomedical Imaging: Macro to Nano*, vol. 1, pp. 960, 2004.
- [12] Toft P, "The Radon Transform - Theory and Implementation," PhD thesis, Department of Mathematical Modelling, Technical University of Denmark, 1996.
- [13] Enhua Wu, Youquan Liu, Xuehui Liu, "An improved study of real-time fluid simulation on GPU," *Computer Animation & Virtual Worlds*, vol. 15, no. 3-4, pp. 139-146, 2004.
- [14] Kyoung-Su Oh, Keechul Jung, "GPU implementation of neural networks," *Pattern Recognition*, vol. 37, no. 6, pp. 1311-1314, 2004.

-
- [15] Kruger J, Westermann R, "Acceleration techniques for GPU-based volume rendering," *IEEE Visualization 2003*, pp. 287-292, 2003.
- [16] Prager RW, Rohling RN, Gee AH, Berman L, "Rapid calibration for 3-D freehand ultrasound," *Ultrasound Med Biol*, vol. 24, no. 6, pp. 855-69, 1998.
- [17] Stoll J, Dupont P., "Passive Markers for Ultrasound Tracking of Surgical Instruments," *Medical Image Computing and Computer-Assisted Intervention*, vol. 3750, pp. 41-48, 2005.