Design and development of an ultrasound calibration phantom and system

Alexis Cheng[a], Martin K. Ackerman[b], Gregory S. Chirikjian[b], Emad M. Boctor[a,c,d]

[a]Dept. of Computer Science, Johns Hopkins University, Baltimore, MD, USA
[b]Dept. of Mechanical Engineering, Johns Hopkins University, Baltimore, MD, USA
[c]Dept. of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, USA
[d]Dept. of Radiology, Johns Hopkins University, Baltimore, MD, USA

Abstract

Image-guided surgery systems are often used to provide surgeons with informational support. Due to several unique advantages such as ease of use, real-time image acquisition, and no ionizing radiation, ultrasound is a common medical imaging modality used in image-guided surgery systems. To perform advanced forms of guidance with ultrasound, such as virtual image overlays or automated robotic actuation, an ultrasound calibration process must be performed. This process recovers the rigid body transformation between a tracked marker attached to the ultrasound transducer and the ultrasound image. A phantom or model with known geometry is also required. In this work, we design and test an ultrasound calibration phantom and software. The two main considerations in this work are utilizing our knowledge of ultrasound physics to design the phantom and delivering an easy to use calibration process to the user. We explore the use of a three-dimensional printer to create the phantom in its entirety without need for user assembly. We have also developed software to automatically segment the three-dimensional printed rods from the ultrasound image by leveraging knowledge about the shape and scale of the phantom. In this work, we present preliminary results from using this phantom to perform ultrasound calibration. To test the efficacy of our method, we match the projection of the points segmented from the image to the known model and calculate a sum squared difference between each point for several combinations of motion generation and filtering methods. The best performing combination of motion and filtering techniques had an error of 1.56 mm and a standard deviation of 1.02 mm.

Introduction

Image-guided surgery (IGS) systems are often used in modern surgeries to provide surgeons with additional information support and guidance resulting in less trauma for the patient. Specific benefits to the patient can include cost reduction of the procedure, reduced morbidity rates, and shorter recovery times. In IGS systems, a medical imaging modality is often used to provide a visualization of underlying tissue structures or anatomy that cannot be seen with the naked eye. Many different types of medical imaging modalities have been used in this manner, including fluoroscopy, computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound (US).

In this work, we focus on the use of US in IGS, which provides many unique advantages in a surgical setting. Unlike CT or fluoroscopy, US does not produce or use ionizing radiation, so there is no concern of applied radiation dose to patients or healthcare personnel. Coupled with the fact that US images can be acquired in real-time, a continuous stream of US images can be presented during surgery, providing real-time feedback. Other modalities such as CT or MRI are generally only able to provide feedback to the surgeon before or after the surgery as the imaging equipment is immobile and technical personnel is necessary for its operation. In contrast, US systems are mobile and can be operated with minimal training. Also, whereas MRI imposes the requirement that metal tools cannot be used in its vicinity, US does not impose requirements on the tools used in its imaging field. The final advantage, and possibly the main reasons for ultrasound's widespread adoption, is its low cost in comparison to the other medical imaging modalities. Despite its advantages, the main

drawback for US is its poor image quality relative to the other imaging modalities. This drawback often makes it difficult to automatically segment regions of interest within the images.

If one wants to use US to perform more advanced forms of guidance, an US calibration process must be performed. The first requirement for this process is that a tracked rigid body must be attached to the US transducer. This can be either an optical marker or an electromagnetic (EM) sensor. The external tracker, either optical or EM, provides the tracked rigid body's pose, its orientation and position, within the external tracker's frame of reference. The US calibration process finds the rigid body transformation relating the tracked rigid body to the US image, allowing an US image to be positioned in the external tracker space. The US image is now registered with any other tools or devices that are also being tracked in this external tracker space. Once calibrated, more advanced uses of the US system are possible, such as overlaying the US image onto a video stream or actuating a tracked robotic actuator to target a region within the US image.

To find the tracked body to image transformation in the US calibration, a specialized model or phantom with known configurations or shape is required. There have been many different types of phantoms or models used for US calibration including wall [1], cross-wire [2], Z-fiducial [3], and $AX = XB$ [4] phantoms. Calibrating with a wall model uses many US images of a plane, such as the walls of a water tank. Along with the tracked pose information, the lines in the US images resulting from imaging the planes are reconstructed into a plane in the external tracker's frame of reference. A cross-wire phantom is built such that two wires cross at a single point. US images of this point are then accumulated in a number of different tracked poses. One then uses these tracked poses and the points in the US images resulting from imaging the cross-wire to reconstruct a point in the external tracker's space. This method is considered sufficiently accurate, but it requires many images, resulting in a long acquisition time.

The US calibration phantom that we develop in this work is simultaneously in the Z-fiducial and $AX = XB$ category. Z-fiducial phantoms are those where three wires are oriented in a plane to form a Z or N shape. Given a plane intersecting all three wires of multiple Z-fiducials, the unique pose of the plane can be determined. There are two methods for using Z-fiducial phantoms in US calibration. In the first method, only a single image is required, but the phantom must be externally tracked in addition to the probe fiducial. In the second method, the phantom is not externally tracked, but a moderate number of images, from differing poses, are required. $AX = XB$ phantoms are those that allow the relative rigid body transformation between two images to be recovered based on each image's content. Z-fiducial phantoms are a subset of $AX = XB$ phantoms and therefore also fit this requirement.

The $AX = XB$ problem is also referred to as the hand-eye calibration problem. As seen in figure 1, A^{ij} and B^{ij} are relative motions related by the rigid body transformation X. A tracker provides the homogeneous transformation B_i, representing the pose of the externally tracked marker or sensor. By using an $AX = XB$ phantom, the absolute pose of the image, A_i, relating each image to the phantom's coordinate system, can be computed based on the image content. The $AX = XB$ framework ultimately only requires relative motions, and is therefore advantageous because it does not require the US transducer to be fixed at specific locations or the calibration phantom to be tracked by the external tracker. The calibration phantom must be fixed during the data collection process, but the US transducer can freely move and data can be collected in real-time.

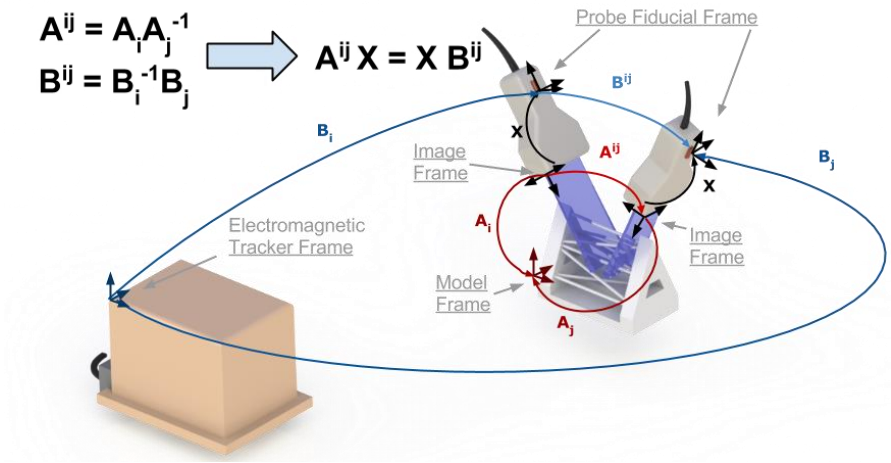

Figure 1. $AX = XB$ formulation with labeled coordinate frames

In this work we design and test a new calibration phantom. This calibration phantom improves on generic Z-fiducial phantoms by having a design that is more compatible with US physics. We will describe the phantom design and the factors behind these decisions. We will also describe the segmentation and registration software to support this calibration phantom. Finally, we will present experimental results using this calibration phantom.

Apparatus

In addition to the calibration phantom that we will describe in the following sections, we use a Sonix RP (Ultrasonix Corp.) US machine, a L12-5 US transducer, an EM tracking system (Ascension Technology Corp.). The EM tracking system provides B_i, which is synchronized with the US images using the MUSiiC Toolkit [5]. This apparatus allows us to collect a continuous real-time stream of data.

Design of Calibration Phantom

The new calibration phantom that is presented in figure 2 where the phantom's computer aided design in figure 2A and the three-dimensional printed phantom in figure 2B are shown. There were several design factors that distinguish this phantom from previous phantoms. The first constraint was that the phantom be created without any need for additional assembly, adjustment, or treatment. Previous phantoms that are printed solely have the frame with holes, identifying the position and orientation of wires. The user would then wire off-the-shelf fishing line [6] into the phantom. In contrast, our entire phantom, including the rods, is printed, allowing our phantom to be "plug and play". The user can download the appropriate computer-aided design (CAD) file, print it from a three-dimensional printer, and be ready to perform US calibration. There are, however, limits to the specifications of a three-dimensional printer. It may be difficult to print thin wires as they can become brittle. For this reason, the printed wires are 2mm in diameter, thicker than off-the-shelf fishing line. Thick wires affect the US images in negative ways, making it more difficult to automatically segment. We will discuss this in the design of the US calibration software. The terms "rods" and "wires" will be used interchangeably in this manuscript.

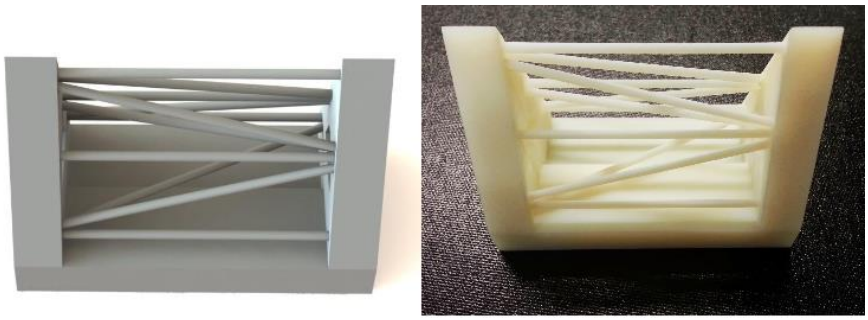

A B

Figure 2. Calibration phantom model A) CAD B) real

Since this phantom is an extension of Z-fiducial phantoms, there were many considerations regarding the construction of the Z-fiducial. First, the length and angle of the Z-fiducial were chosen such that a submillimeter translation could be resolved. In standard phantoms of this type, the Z-fiducials all lie on parallel planes. This geometry does not take advantage of US physics, as the axial dimension has higher resolution than the lateral dimension. In our phantom, a portion of the Z-fiducial change is reflected in the axial dimension, in varying orientations. The new geometry causes rods that are not perpendicular to the imaging plane to have a non-optimal acoustic response, leading to low rod intensities in the US image. To compensate for this, the Z-fiducials are oriented such that in cases where one Z-fiducial becomes difficult to visualize, another Z-fiducial will have a rod that becomes increasingly perpendicular to the image plane. A simulated situation is shown in figure 3A-B. The phantom also has many redundant rods, so that the subset of rods with the highest acoustic response, given the imaging orientation, can be chosen.

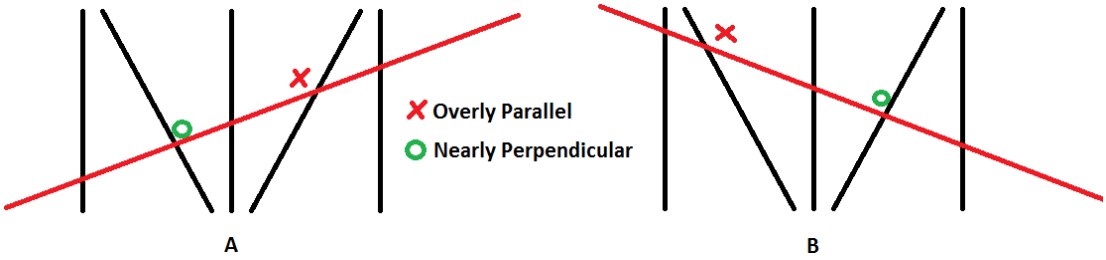

Figure 3. Z-fiducial orientations A) Case 1 B) Case 2

Another consideration was the shadowing effect of the structures within an US image. To avoid the case where many of the rods do not have a clear acoustic path to the transducer, the Z-fiducials are oriented in the shape of a triangle such that when one face of the triangle is experiencing severe shadowing effects, the other faces will be unaffected. This phantom was also designed for use with probes of multiple lateral lengths. For example, with probes of shorter lateral lengths, one could use the six upper most rods and perform the calibration. Finally, we desired the phantom to allow for a large range of motion of at least 3cm for each translational degree of freedom and 45 degrees for each rotational degree of freedom.

Design of Calibration Software

The calibration software has three essential components. Its general workflow can be found in figure 4A. The first is an automatic segmentation algorithm to detect the rods present in the US image. The second component computes the homogeneous transformation, A_i, relating each image with the phantom. The final component solves for the sensor calibration rigid body transformation, X.

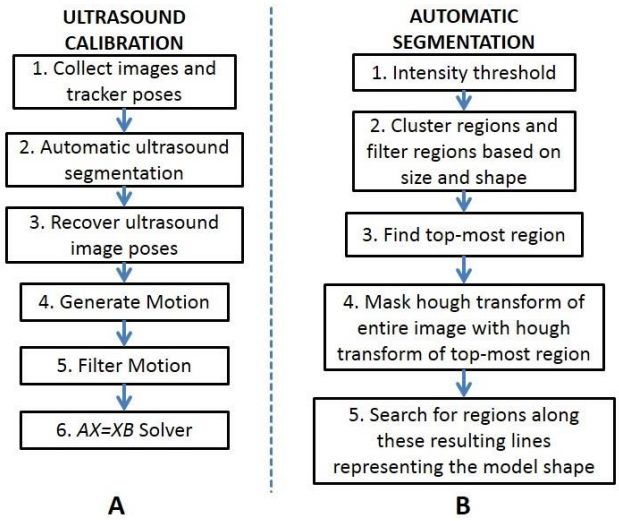

Figure 4. Workflows for A) the overall US calibration and B) the automatic segmentation algorithm

The first component is an automatic segmentation algorithm. Its overall workflow can be seen in figure 4B. The first step of the segmentation algorithm applies an intensity threshold, chosen by Otsu's method [7], to the image. This step allows the algorithm to get a first order estimate of which pixels should be considered the background and which pixels should be considered the signal. A connected regions algorithm is then used to cluster signal pixels together. Afterwards, a filter is applied where only regions containing a certain range of pixels are retained. These steps allow us to remove noise and extract the rods from the US image. After these steps, there will be multiple regions of pixels, indicating rod candidates. An example US image and the segmented rod candidates can be seen in figure 5. The next steps must find correspondence between these wire candidates and the model.

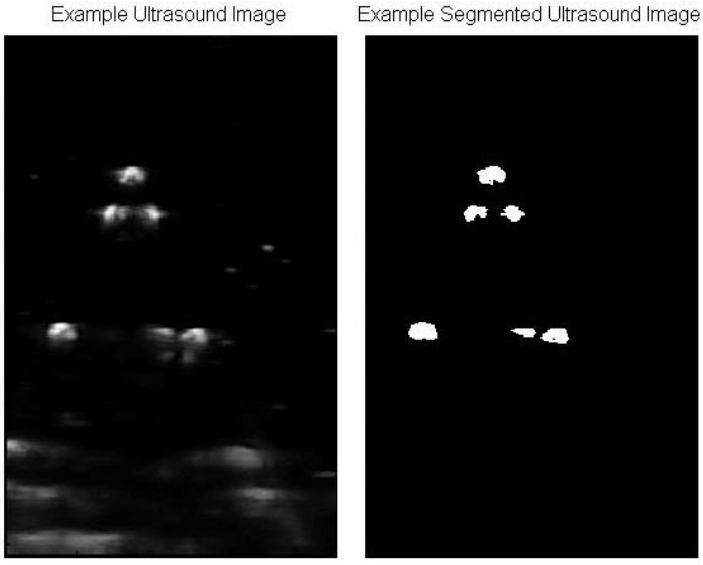

Figure 5. Example Ultrasound Image and Example Segmented Ultrasound Image

The region closest to the US transducer face is selected as the top rod in our model. This is a valid assumption as cases where the phantom is placed in a water bath, it is not practical to place the transducer such that the top rod is not closest to the US transducer face. Given our model, the remaining connected regions in the US

image will exhibit a triangular shape. Thus, the standard Hough transform [8] can be applied to find the edges of this triangular pattern. We use our knowledge of the location of the top rod and the edges of the triangular pattern to establish correspondences between the triangular pattern and the model. With knowledge of these lines, we can search the original image to find rods that had low intensities. This step ensures that even if a rod was only faintly seen in the image, it will still be detected. Finally, for each region, we select the points lying closest to the transducer face from the centroid because the top of the rods are more accurately represented in the US image.

We examined three registration methods to compute the homogeneous transformation, A_i. First, we can directly register these points to a discretized model using coherent point drift [9]. This algorithm does not require known correspondences and also allows for missing points. However, this algorithm is prone to local minima during its iterative procedure. Thus, it may be necessary to initialize the segmented points at different locations and hope a majority converge to the correct answer. The second method requires known correspondences and is based on the Z-fiducials. Given a plane that intersects all three rods of the Z-fiducial, the intersection of the plane and the diagonal rods can be localized in the phantom space. Thus, if a minimum of three Z-fiducials are visible in the US image, the transformation between the image space and the phantom space can be found. However, if certain rods are not visible, there may not be enough Z-fiducials to use this method. A third method uses the inter-point distances of the segmented points and fits this distance vector to the model. This method also requires correspondence. In the case when not all of the points exist in the image, a more complex fitting method, such as the Hungarian method [10], can be used. The fitting process can be accomplished through either an optimization procedure, which is also prone to local minima, or through a look-up table, which could require significant system memory. The preferred approach for this fitting process is to use a hybrid approach: a coarse look-up table is used to find an initial solution to the optimization procedure. The optimization procedure would then be used to refine the solution towards its global minima.

Before we can solve the $AX = XB$ problem, we must obtain the relative A and B transformations from their respective absolute poses. We can either take every consecutive pair where $A^i = A_i^{-1}A_{i+1}$ or every pair combination where $A^{ij} = A_i^{-1}A_j$. In general, every pair will have better results than consecutive pairs at the expense of computation time, as there will be significantly more input data to the chosen $AX = XB$ solver. In addition, we applied a filtering step using the Euclidean invariants [11] of the homogeneous transformations in $AX = XB$, filtering small motions more affected by noise. We also filtered motions where the invariants were not sufficiently close. These two filters can be respectively seen in Equations 1 and 2. θ represents the magnitude of the rotation of the motion. With relative A and B pairs, we can solve the $AX = XB$ problem using various solvers such as, the dual quaternion method [12] or the kronecker product solver [13]. For this manuscript, we used the kronecker product solver.

$$\theta_B < \delta_1 \quad (1)$$

$$|\theta_A - \theta_B| < \delta_2 \quad (2)$$

Experimental Results

To test the calibration software, we calibrated a L12-5 linear US transducer. Since this transducer has a lateral length of 4cm, we only segmented the top six rods. The results shown below use the third registration method of fitting inter-point distances. Figure 6A-B show the results of plotting $B_i X^{-1} p_i$ where p_i is the set of segmented points from each image.

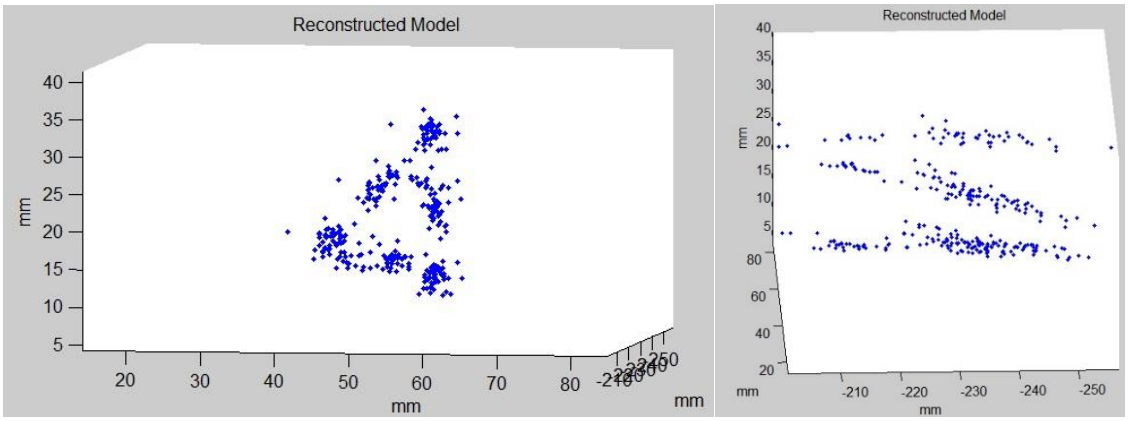

Figure 6. Views of Reconstructed Model

To quantify the difference, we fitted the model to the reconstructed points and computed a normalized metric. For each point, we found its closest point on the model and computed the sum squared difference between them. We tested combinations of A^i or A^{ij} and with and without filtering. A set of test data not used in the computation of X is used to compute the results seen in table 1.

Table 1. Normalized error metric for different combinations of motion generation and filtering

Motion Generation	Filtering	Error Metric (mm)
A^i	Without	2.75 ±1.67
A^i	With	1.74 ±1.00
A^{ij}	Without	2.36 ±1.49
A^{ij}	With	1.56 ±1.02

Discussion

We can see from the results shown in figure 6 and table 1 that this phantom is a feasible US calibration phantom. A closer look at table 1 shows that the third and fourth rows respectively have lower errors than the first and second rows. This reinforces our initial hypothesis that motion generation from all pairs generally provides better results than motion generation from consecutive pairs. However, more experiments may be necessary to reach conclusive results. It can also be seen in table 1 that the second and fourth rows respectively have lower errors than the first and third rows. This corresponds with the notion that filtering using the Euclidean invariants [11] decreases the error.

There is certainly still room for improvement. US images are difficult to segment and there are cases where the US image quality will cause our automatic segmentation algorithm to fail. Since we have a continuous stream of data, there will eventually be enough data that can be automatically segmented. However, since the software is not currently real-time, it is impossible to know during data collection if there is enough data or not. Therefore, we aim to implement the software in real-time in order that a user could receive real-time feedback as they are collecting data.

Conclusions

In this work, we presented the development of an US calibration phantom that takes advantage of US physics and can be easily printed without user modifications. In addition, we describe a segmentation algorithm that leverages the shape and structure of our phantom. Finally, we show results of using this phantom in various

combinations of motion generation and filtering methods. Future work will compare this phantom with other US calibration phantoms, develop a real-time US calibration framework using this phantom, and use this phantom as a test bed for other novel US calibration methods.

Acknowledgement

Financial support was provided by Johns Hopkins University internal funds, NIBIB-NIH grant EB015638, and NSF grant IIS-1162095. We thank Justin Opfermann from the Children's National Medical Center for his help with printing the phantom.

References

[1] Prager R.W., Rohling R.N., Gee A.H., and Berman L.H., "Rapid Calibration for 3-D freehand Ultrasound," Ultrasound in Medicine and Biology, 24 (6), 855-869, 1998.
[2] Detmer P.R., Bashein G., Hodges T., Beach K.W., Filer E.P., Burns D.H., and Strandness Jr D.E., "3D ultrasonic image feature localization based on magnetic scanhead tracking: in vitro calibration and validation," Ultrasound in Medicine and Biology, 20 (9), 923-936, 1994.
[3] Pagoulatos N., Haynor D.R., and Kim Y., "A Fast Calibration Method for 3-D Tracking of Ultrasound Images using a Spatial Localizer," Ultrasound in Medicine and Biology, 27(9), 1219-1229, 2001.
[4] Boctor, E.M., Viswanathan, A., Choti, M.A., Taylor, R.H., Fichtinger, G., Hager, G.D., "A Novel Closed Form Solution for Ultrasound Calibration," IEEE Int Symp. On Biomedical Imaging, 527-530, 2004.
[5] Kang, H. J., Cheng, A., and Boctor, E. M., "MUSiiC ToolKit 2.0: Bidirectional Real-time Software Framework for Advanced Interventional Ultrasound Research," MICCAI 2012, International Workshop on System and Architectures for Computer Assisted Interventions, 2012.
[6] Hsu P.W., Prager R.W., Gee A.H., and Treece G.M., "Real-time Freehand 3-D Ultrasound Calibration," Ultrasound in Medicine and Biology, 34 (2), 239-251, 2008.
[7] Otsu N., "A threshold selection method from gray-level histograms," *IEEE Trans. Sys., Man., Cyber*. 9 (1): 62–66, 1979.
[8] Hough P.V.C., "Machine Analysis of Bubble Chamber Pictures," Proc. Int. Conf. High Energy Accelerators and Instrumentation, 1959.
[9] Myorenko A., and Song X., "Point-set registration: coherent point drift," IEEE Trans. Pattern Anal. Mach. Intell., 32 (12), 2262-2275, 2010.
[10] Kuhn H.W., "The Hungarian Method for the assignment problem," Naval Research Logistics Quarterly, 2, 83-97, 1955.
[11] Ackerman M.K., Cheng A., Boctor E., and Chirikjian G.S., "Online Ultrasound Sensor Calibration Using Gradient Descent on the Euclidean Group," Accepted to International Conference on Robotics and Automation, 2014.
[12] Daniilidis, K., "Hand-Eye Calibration Using Dual Quaternions," The International Journal of Robotics Research Vol 18, pp. 286-298, 1999.
[13] Andreff, N., Horaud, R., Espiau, B., "Robot Hand-Eye Calibration Using Structure-from-Motion" The International Journal of Robotics Research, 2001.