Localizing spherical fiducials in C-arm based cone-beam CT

Ziv Yaniv[a)]
Department of Radiology, Imaging Science and Information Systems (ISIS) Center, Georgetown University Medical Center, Washington, DC 2007

(Received 16 March 2009; revised 28 July 2009; accepted for publication 1 September 2009; published 2 October 2009)

Purpose: C-arm based cone-beam CT (CBCT) imaging enables the *in situ* acquisition of three dimensional images. In the context of image-guided interventions, this technology potentially reduces the complexity of a procedure's workflow. Instead of acquiring the preoperative volumetric images in a separate location and transferring the patient to the interventional suite, both imaging and intervention are carried out in the same location. A key component in image-guided interventions is image to patient registration. The most common registration approach, in clinical use, is based on fiducial markers placed on the patient's skin which are then localized in the volumetric image and in the interventional environment. When using C-arm CBCT, this registration approach is challenging as in many cases the small size of the volumetric reconstruction cannot include both the skin fiducials and the organ of interest.
Methods: In this article the author shows that fiducial localization outside the reconstructed volume is possible if the projection images from which the reconstruction was obtained are available. By replacing direct fiducial localization in the volumetric images with localization in the projection images, the author obtains the fiducial coordinates in the volume's coordinate system even when the fiducials are outside the reconstructed region.
Results: The approach was evaluated using two types of spherical fiducials, clinically used 4 mm diameter markers and a custom phantom embedded with 6 mm diameter markers that is part of a commercial navigation system. In all cases, the method localized all fiducials, including those that were outside the reconstructed volume. The method's mean (std) localization error as evaluated using fiducials that were directly localized in the CBCT reconstruction was 0.55(0.22) mm for the 4 mm markers and 0.51(0.18) mm for the 6 mm markers.
Conclusions: Based on the evaluations the author concludes that the proposed localization approach is sufficiently accurate to augment or replace direct volumetric fiducial localization for thoracic-abdominal interventions. This allows the physician to position fiducials in a more flexible manner, relaxing the requirement that both the organ of interest and skin surface be contained in the volumetric reconstruction.
[DOI: 10.1118/1.3233684]

Key words: image-guided therapy, registration, fiducial localization, cone-beam CT

I. INTRODUCTION

C-arm based cone-beam CT (CBCT) is a three dimensional (3D) *in situ* imaging modality. It provides images that have a spatial resolution similar to that obtained with diagnostic CT, but with lower discrimination between tissue types and a smaller spatial extent. Figure 1 illustrates these differences.

More recently, improvements in flat panel detector technology have resulted in improved tissue type discrimination, enabling its introduction into clinical interventions dealing with soft tissue.[1–3] For interventions using an image-guided navigation system, this modality can potentially simplify the procedure workflow with 3D imaging and intervention carried out at the same location.

A typical image-guided navigation workflow consists of the following steps: (1) Place fiducials on patient, (2) acquire a 3D image in an imaging suite, (3) transfer patient to interventional suite, (4) register image and patient space, and (5) navigate.

As CBCT provides *in situ* imaging, it can potentially be used for registrationless navigation,[4–6] streamlining the clinical workflow to only two steps: (1) Acquire a 3D image and (2) navigate. This is possible as the 3D image is acquired *in situ*, allowing us to replace the intraoperative registration step with a one-time calibration. This calibration relates image space to a fixed reference frame that can be detected by the navigation system's tracking device. Figure 2(a) shows the known and estimated transformations relevant for this approach. In this manner, the navigation system transfers tool locations from patient space to image space via the known transformation between the fixed reference frame and the image space. The only requirement is that the fixed reference frame be positioned such that it can be detected by the tracker.

In the clinical environment this requirement is not trivial. The location of the fixed reference frame places constraints on the positioning of the tracker. In addition, when using wired tracking systems the reference frame must be physically connected to the tracker control unit. Given that the

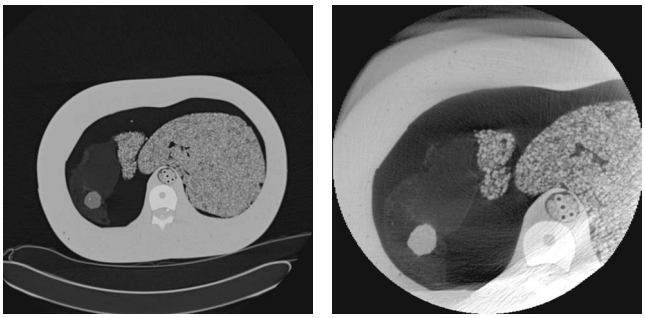

FIG. 1. Similar axial slices from an anthropomorphic phantom, (a) diagnostic CT (Siemens Somatom volume zoom) and (b) C-arm CBCT (Siemens Axiom Artis dFA). Images are displayed using the same window (2600) and level (150). Note that the CBCT reconstructed volume includes the phantom's internal structures but not its skin surface.

clinical environment is most often already cramped, physically positioning the tracking system can become an arduous task. Currently, using electromagnetic trackers for registrationless navigation is all but impossible. These systems are wired, they track in a relatively small spatial region, and they exhibit a variable degree of measurement distortions due to the presence of metallic objects.[7] Thus, navigation systems employing electromagnetic tracking and C-arm based CBCT imaging have a slightly less ideal workflow: (1) Place fiducials on patient, (2) acquire a 3D image, (3) register image and patient space, and (4) navigate. Figure 2(b) shows the known and estimated transformations relevant for this approach.

We have developed a navigation system that employs electromagnetic tracking and C-arm based CBCT for thoracic-abdominal interventions. These include biopsy of liver lesions, radio frequency ablation of liver tumors, and lung biopsies. The system follows the four step workflow described above, with images acquired at breath-hold and registration performed with the widely used paired-point method.[8] That is, for registration, the corresponding fiducials are manually localized in the volumetric image and in the

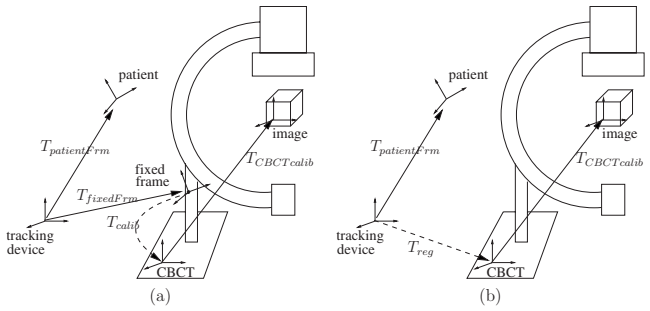

FIG. 2. Coordinate systems used by CBCT based navigation systems. The solid (dashed) arrows denote known (unknown) transformations. (a) Registrationless navigation and (b) registration-based navigation. In both cases, the desired transformation is between the patient mounted coordinate system and the image coordinate system. T_{fixedFrm}, $T_{\text{patientFrm}}$—estimated continuously by tracking device. $T_{\text{CBCTcalib}}$—estimated once when CBCT system is configured. T_{calib}—estimated once when reference frame is fixed onto CBCT. T_{reg}—estimated for every patient.

interventional environment. We have found that this is a challenging task when using C-arm CBCT data. The challenge arises from the small size of the reconstructed spatial region which often does not encompass both fiducials and anatomy of interest.

A recently introduced commercial electromagnetic navigation system, iGuide CAPPA (Siemens AG Healthcare, Erlangen, Germany), is also dependent upon fiducial localization in the reconstructed volume. Unlike our manual fiducial localization approach, this system provides automatic registration. A custom electromagnetically tracked phantom embedded with spherical markers is imaged alongside the patient. The spherical markers are automatically localized in the reconstructed volume enabling image to patient registration. If for some reason automatic fiducial localization fails, then the user can manually localize them in the images. In cases where the fiducials are outside the reconstructed region, the registration phantom must be repositioned and a new scan is acquired.

For CBCT systems, the maximal size of the reconstruction volume is limited by the size of the flat panel detector. It has been recognized as a frequently encountered limitation for hepatic interventional radiology procedures where the liver is not visible in its entirety.[2] Thus, the limited size of the reconstructed region requires careful positioning of the patient so that the anatomical structures of interest are visible.

For navigation purposes, we require that both the anatomical structures of interest and the fiducials placed on the patient's skin be inside the reconstructed region. This is a challenging requirement as the reconstruction of the clinical region of interest (ROI) will most often not include all of the patient's skin surface, as seen in Fig. 1(b). Fiducial placement is thus constrained to regions of the skin surface that are expected to be included in the reconstructed region, or in some cases (i.e., larger patients) may not be possible as the anatomy of interest is too far from the skin surface. These constraints can potentially increase target registration errors or even preclude navigation. The former is a result of missing fiducials, or tight fiducial configurations that ensure visibility but are at odds with the recommended wide spread fiducial placement strategies employed to reduce target registration errors.[9] The latter is a result of an insufficient number of fiducials visible in the reconstruction, with at least three fiducials required for rigid registration.

In this work, we show that even when fiducials are outside the reconstructed region and cannot be directly localized in the volume, they can be localized in the volumetric coordinate system. Our approach is based on the observation that fiducials that are outside the reconstructed region are most often visible in several projection images, the input for the reconstruction algorithm. While the number of these images may not be sufficient for accurate estimation of attenuation coefficients, volumetric reconstruction, they are sufficient for spatial localization of the fiducials.

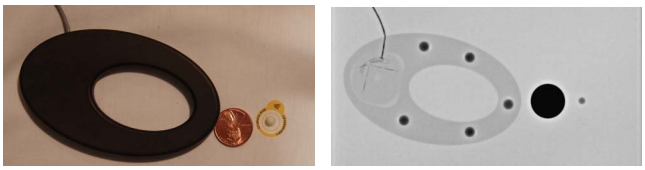

FIG. 3. Picture and x ray of the spherical fiducials used in this work. iGuide CAPPA registration phantom with 6 mm embedded markers and the 4 mm CT-SPOTS markers.

II. EQUIPMENT AND IMAGE ACQUISITION

Our approach is based on the detection of circular markers in the projection images, it is thus directly applicable to all spherical fiducials. In our case we evaluated two types of fiducials, 4 mm CT-SPOTS markers (Beekley Corp., Bristol, CT) and 6 mm markers that are embedded in the custom iGuide CAPPA registration phantom. The former allow us to vary the number and position of fiducials, while the latter consists of five fiducials in a rigid configuration. Figure 3 shows these two types of markers.

All images used in this work were acquired with the Axiom Artis dFA (Siemens AG Healthcare, Erlangen, Germany) CBCT system. The CBCT is modeled as a distortionless pinhole camera. Projection images are acquired using a 20 s rotation in which 543 images are uniformly acquired over a rotation of 220°. Image dimensions are 960×1024 pixels. The resulting volumetric reconstructions are $512\times512\times440$ with an isotropic 0.45 mm voxel size.

A premise of our approach is that the camera's intrinsic and extrinsic, pose, parameters are known and that the pose is given with regard to the CBCT coordinate system. In our case, we have access to the calibration matrices that are used to perform the volumetric reconstruction from the projections. Thus, the matrices encode the camera's pose with regard to the CBCT coordinate system. If these matrices are not accessible, our approach requires additional camera calibration and registration.

In this case a calibration object containing easily identifiable markers is imaged and a volumetric reconstruction is performed. The most widely used CBCT calibration phantom design consists of spherical fiducials with varying sizes positioned in a helical pattern. Other various point fiducial configurations are also applicable.[10] By identifying corresponding fiducials in the reconstructed volume and model coordinate systems, we obtain the transformation from the model to the volumetric coordinate system using paired-point registration. The inputs to the camera calibration are the original projection images and the fiducial coordinates in the CBCT coordinate system instead of in the model coordinate system. Thus, the estimated calibration matrices encode the camera's pose with regard to the CBCT coordinate system, satisfying our assumption.

III. LOCALIZATION APPROACH

To localize a fiducial in the CBCT coordinate system, it is first localized in multiple projection images. Its spatial location is then estimated as the intersection point of the back-projected rays emanating from the camera locations and going through the corresponding image locations. This is an established approach that provides an estimate of the fiducial's spatial location using two or more images. Accurate estimation of a fiducial's spatial location is thus dependent on accurate localization in the projection images. In practice, localization is based on the results of fiducial segmentation in the projection images.

TABLE I. Fiducial localization in volumetric coordinate system.

Input: Projection matrices and corresponding images.
Output: Fiducial location in volumetric coordinate system.
1. Initial localization
(a) Manually define a rectangular ROI in one of the projection images.
(b) Perform edge detection in the ROI and fit a circle to the edge data using the RANSAC algorithm.
2. Constrained localization (adjacent projection images, ±8° C-arm rotation)
(a) Perform edge detection in a ROI defined by the epipolar line corresponding to the fiducial location estimated in the previous phase and fit a circle to the edge data using the RANSAC algorithm.
(b) Compute the backprojected ray emanating from the x-ray source and going through the fiducial location.
(c) Estimate the point that minimizes the distance to all backprojected rays using the RANSAC algorithm.
3. Reprojection localization
(a) Reproject the estimated 3D location onto every nth ($n=10$) image.
(b) Perform edge detection in a ROI surrounding the projected point and fit a circle to the edge data using the RANSAC algorithm.
(c) Compute the backprojected ray emanating from the x-ray source and going through the fiducial location.
(d) Estimate the point that minimizes the distance to all backprojected rays using the RANSAC algorithm.

Following Ref. 11 we view segmentation as two processes. A high level process, recognition, determining the presence and rough location of the object of interest in the image, and a low level process, delineation, accurate estimation of the spatial extent occupied by the object. Humans are more adept at the former and computer algorithms are more adept at the latter. We thus use a semiautomatic fiducial localization approach. Segmentation is manually initialized in a single image followed by three phases that enlarge the search space, from a single image to several adjacent images along the C-arm trajectory, and finally localization in images sampled throughout the trajectory. Table I and Fig. 4 summarize this fiducial localization approach.

Initially, fiducial localization is performed semiautomatically in a single projection image. The user indicates a rectangular ROI containing a single fiducial. We then perform edge detection in that region using the Canny edge detector.[12] This is followed by circle estimation using the random sample consensus (RANSAC) algorithm.[13] Figure 5 shows sample user specified ROIs. To utilize the RANSAC algorithm, we require a method for parameter estimation using the minimal number of data elements, a method for least

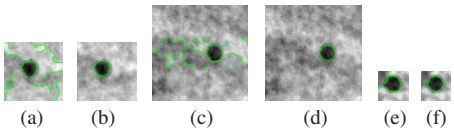

FIG. 4. Graphical representation of localization approach: Phase 1, initial localization in user specified ROI (a) edges and (b) localized fiducial; phase 2, constrained localization (c) edges near epipolar line and (d) localized fiducial; phase 3, reprojection localization (e) edges near reprojection point and (f) localized fiducial.

squares parameter estimation and a distance function. Appendix A describes these elements for spheres of arbitrary dimension.

To improve the probability of a correct estimate, we incorporate an additional constraint into the RANSAC framework. Given that our fiducial is located inside the rectangular region, we constrain the circle's radius to be less than half of the maximal edge length. It is assumed that the fiducial is localized, though it need not be an accurate localization.

Once the fiducial is localized in the initial image, we search for it in adjacent images acquired during the C-arm rotation. In our case we attempt to localize it in eight additional images, corresponding to a uniform sampling of the C-arm orientations in an arc of $\pm 8°$. Given that we are imaging a static scene, acquisition at breath-hold, we can use the epipolar geometry to reduce our search space to a single dimension, the epipolar line.[14] Given point ${}^{l}\mathbf{p}$ in the first image, its location in the second image is on the epipolar line $F{}^{l}\mathbf{p}$, where F is the fundamental matrix which is readily available from the known calibration matrices ${}^{l}P, {}^{r}P$,

$$F = [{}^{r}P\,{}^{l}C]_{\times}\,{}^{r}P\,{}^{l}P^{+},$$

where ${}^{l}P^{+}$ is the pseudoinverse of ${}^{l}P$, ${}^{r}P\,{}^{l}C$ is the projection of the first camera's location in the second image, the epipole, and $[\ ]_{\times}$ is the operator that creates a skew symmetric matrix from the given vector.

Similar to the localization in the initial image, we again perform edge detection in a rectangular region, in this case it is centered on the fiducial location in the initial image. We incorporate the epipolar constraint by ignoring edges that are far from the epipolar line. Again, we localize the fiducial using the RANSAC framework. Once the fiducial is localized in the adjacent images, we estimate an initial location in the volumetric coordinate system as the intersection point of all backprojected rays. As some of the localizations can potentially be erroneous, we estimate the intersection point using the RANSAC framework. Appendix B describes the relevant elements required for using RANSAC to estimate the rays intersection.

Finally, we use the initial estimate of the fiducial location in the volumetric coordinate system to localize it in images sampled throughout the C-arm rotation. In our implementation we use every tenth image, corresponding to a rotation of approximately 4°. It should be noted that the fiducial does not necessarily appear in all images. The initial location is then projected onto all images and the fiducial is localized in each of the images using the same approach as in the initial step, with the region of interest centered on the projected point. The fiducial location in the volumetric coordinate system is then obtained as the intersection point of all backprojected rays using the same robust approach as in the previous phase.

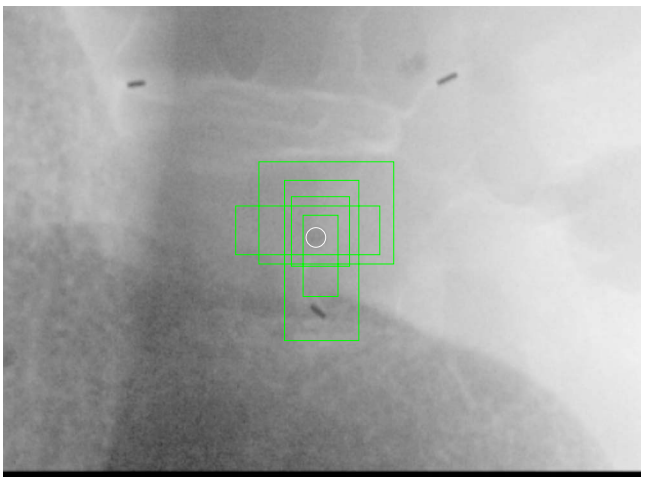

FIG. 5. Five examples of user defined rectangular ROI's, in all cases fiducial localization was successful.

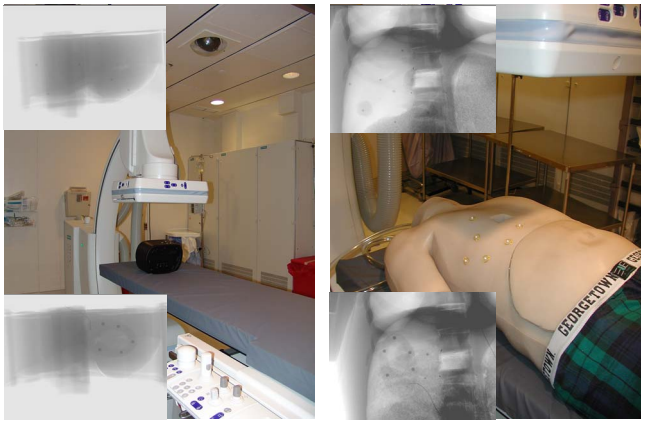

FIG. 6. Experimental setup with our two anthropomorphic phantoms: (a) CIRS abdominal phantom and (b) custom torso phantom. The insets are of sample projection images with CT-SPOTS (top) and iGuide CAPPA registration phantom (bottom).

IV. EXPERIMENTAL EVALUATION

The approach described above was implemented in MATLAB version 7.0.1 (The Mathworks Inc., Natick, MA). All evaluations were performed using a PC, Intel core 2 2.4 GHz processor, 2 Gbytes RAM, running windows XP professional edition.

In our experiments we use two types of anthropomorphic phantoms, the CIRS interventional abdominal phantom (CIRS, Norfolk, VA) and a custom anthropomorphic torso model based on the visible human data.[15] Figure 6 shows the experimental setup and sample projection images showing the various phantom-fiducial combinations.

TABLE II. Summary of qualitative evaluation. Columns denote the number of fiducials that could be localized directly in the volumetric reconstruction, or using the projection images, our approach. Our approach successfully localized all available fiducials in all experiments.

Data set	Fiducials from reconstruction	Fiducials from projection images
	CT-SPOTS 4 mm fiducials	
1	2	5
2	5	6
3	6	6
4	0	6
5	6	6
6	0	4
7	0	6
All	19	39
	iGuide CAPPA 6 mm fiducials	
8	0	5
9	5	5
10	5	5
11	5	5
12	5	5
13	0	5
14	0	5
All	20	35

To perform our evaluations we acquired 14 data sets, seven with CT-SPOTS fiducials and seven with the iGuide CAPPA registration phantom. We evaluate our approach both qualitatively, the number of localized fiducials, and quantitatively, precision, accuracy, and timing.

IV.A. Qualitative evaluation

We assessed the ability of our approach to localize the fiducials based on visual inspection of the localization in the two dimensional projection images and using a three dimensional display of the results. For fiducials that are outside the reconstructed region, we lack a ground truth and thus cannot directly evaluate the localization accuracy. For these fiducials, we infer the expected accuracy as described in the quantitative evaluation section.

In all experiments, our method successfully localized all fiducials. Table II summarizes our qualitative evaluation with Fig. 7 visually illustrating the results for two of the data sets.

IV.B. Quantitative evaluation

We first evaluate the precision of our approach, as it does not require the use of a known ground truth. We then evaluate the method's accuracy and finally its running time.

IV.B.1. Precision

As described above, our method is initialized using a manually delineated rectangular ROI. As a result, the estimated fiducial location may vary based on the user selected image and delineated ROI. To quantify this variation, we performed the following experiment.

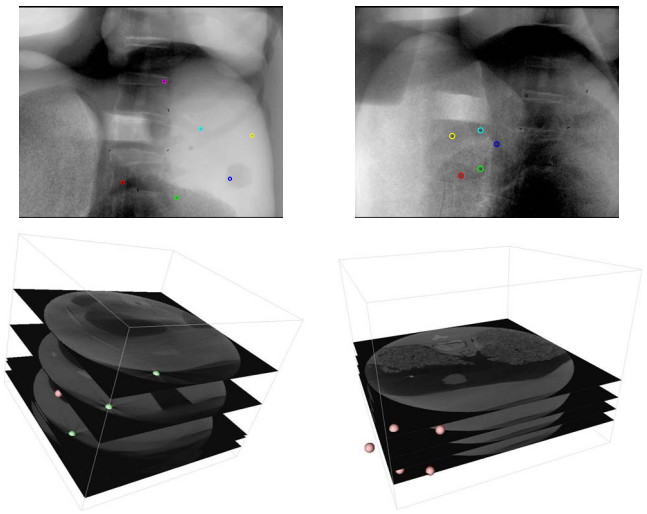

FIG. 7. Fiducial localization in images and corresponding three dimensional locations. Left column corresponds to data set two from Table II and right column to data set eight.

Five 4 mm fiducials and five 6 mm fiducials were repeatedly localized ten times each. Each localization was initialized using a different image and user selected ROI (experiment performed by single user). We then analyzed the distances between the estimates for the same 3D point. For the 4 mm fiducials the mean (std) distances were 0.15 (0.12) mm, with all estimates contained in a sphere with a diameter of 0.64 mm. For the 6 mm fiducials, the mean (std) distances were 0.17 (0.12) mm, with all estimates contained in a sphere with a diameter of 0.69 mm.

IV.B.2. Accuracy

To evaluate the accuracy of our method, we use the fiducials visible in the reconstructed volume as a ground truth, 19 4 mm fiducials and 20 6 mm fiducials. Even though the true location of the fiducials remains unknown, we compare our results to direct fiducial localization in the volumetric data set as it is the standard localization approach used in clinical practice. In our case, the fiducials are localized using the three dimensional equivalent of the first phase in our approach. That is, the user provides a three dimensional ROI, a rectangular box, and the spherical fiducial is detected within it using the RANSAC framework.

We evaluate the accuracy of localization as the distance between the ground truth location and the location estimated using the projection images. We assess both the second and the third phases of our approach. In addition, we evaluate the accuracy of intersecting the backprojected rays using the minimum number of projection images, two images with an angular difference of approximately 40°. In this case we ensure that there are no outliers as the user will discard erroneous fiducial localizations in the projection images. The first phase of our approach is used to localize the fiducial in each of the two images.

Using our proposed approach the mean (std) localization error for the 4 mm CT-SPOTS fiducials was 0.49 (0.21) mm

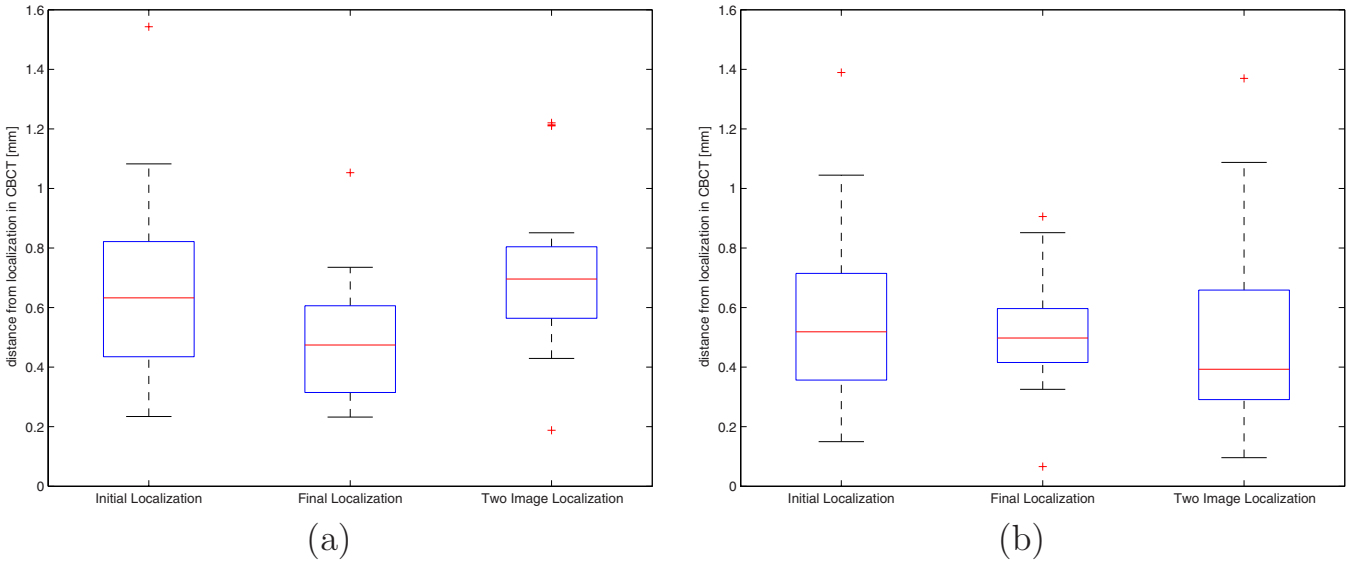

FIG. 8. Accuracy assessment of the initial 3D localization (Table I, step 2) and final localization (Table I, step 3) used by our approach and that of semiautomatic localization using two images with an angular difference of approximately 40°: (a) 4 mm fiducials and (b) 6 mm fiducials. Standard values used in plot construction, box spans interquartile range, median marked inside the box, and maximal whisker length is set to 1.5 times interquartile range.

and for the 6 mm iGuide CAPPA fiduicals it was 0.51 (0.18) mm. Figure 8 shows the results for all three backprojection variations we evaluated. From these plots we see that all three variants provide similar accuracy. We also see that the final phase of our localization method provides more stable results than those obtained by the initial phase and when using only two images.

We next evaluated the effect of the image sampling choice on the localization accuracy. Up to this point, the evaluation used our default sampling for the third phase, every tenth image. We now assess the accuracy when using every 1, 5, 10, 20, 40, and 80th image. These correspond to using 543, 109, 55, 28, 14, and 7 images. For the 4 mm fiducials the mean (std) localization error was 0.53 (0.22) mm when using every image and 0.55 (0.28) mm when using every 80th image. For the 6 mm fiducials, the mean (std) localization error was 0.55 (0.18) mm when using every image and 0.58 (0.2) mm when using every 80th image. We observe that there is only a slight degradation in accuracy and stability as the number of images used decreases. Figure 9 summarizes this evaluation.

We infer expected localization accuracy for fiducials that are outside the reconstructed region, visible in some projection images over a limited angular range, based on the accuracy analysis described above for fiducials that are inside the reconstructed region. We consider the effect of the number of

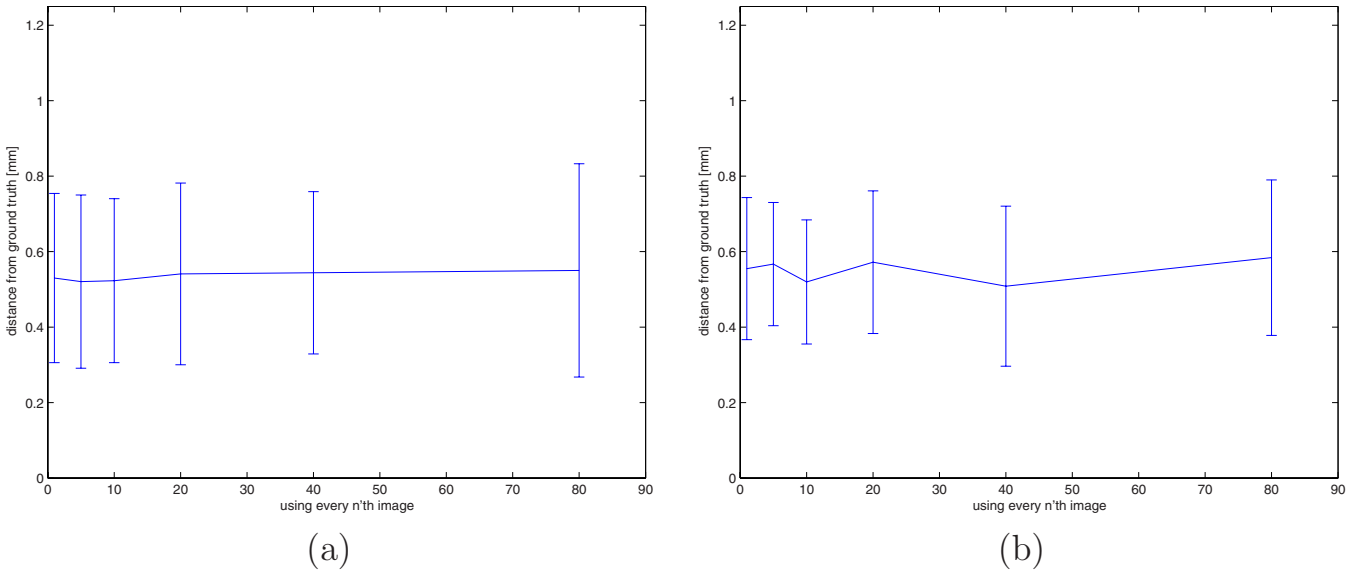

FIG. 9. Assessment of the effect of the image sampling in the third phase of approach on accuracy. The plot shows the mean value for each sampling with error bars of ±1 standard deviation. (a) 4 mm fiducials and (b) 6 mm fiducials.

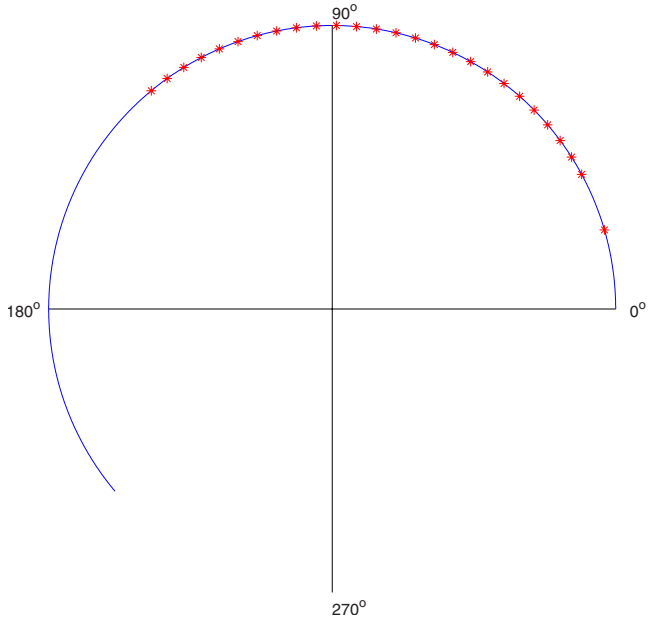

FIG. 10. Example of image set used for localization outside the reconstruction region. C-arm orbit of 220°, with images used for fiducial localization marked. Images correspond to the use of our default sampling rate of every tenth image.

images used in the computations and the angular range spanned by these images. As seen in Fig. 9 even when the number of images was considerably reduced, accuracy was only slightly degraded when the images spanned the same angular range. We also observe in Fig. 8 that the *initial localization*, eight images spanning 16°, and *two image localization*, two images spanning 40°, provided similar accuracy of approximately 0.5 mm. For the fiducials that are outside the reconstructed region, the median (std) number of images used for localization was 27 (8) with a minimum of 13 images spanning 48°. Figure 10 is an example image set used to localize a fiducial outside the reconstructed region.

We can thus expect a mean accuracy of approximately 0.5 mm for these localizations, similar to the results described above for fiducials that are inside the reconstructed region.

IV.B.3. Running time

To assess the time required by our approach, we measured the end to end run time and the time used for processing the images. The former is the time from the moment the user indicates the ROI until the fiducial is localized in the CBCT coordinate system, including both computation time and the time it takes to load images from hard drive. The latter is the net computation time. The end to end times for localizing a single fiducial had a mean (std) of 54.4 (1.4) s and a maximum of 56 s. The time spent on computation had a mean (std) of 4.5 (1.4) s with a maximum of 6.4 s. The order of magnitude difference is due to the time spent loading images during localization. This is due to the limitations imposed by the hardware and software we used. Loading all images into memory before starting the process was not possible using MATLAB on our computer. For systems that can load all images beforehand, the net computation time is the relevant quantity.

V. DISCUSSION AND CONCLUSIONS

We have presented a method for localization of fiducials in the volumetric coordinate system even when their spatial location is outside the reconstructed volume. The approach requires minimal user interaction, indicating a rectangular region of interest surrounding the fiducial in a single image. The method is based on the observation that fiducial localization requires less information than volumetric reconstruction. More specifically, two x-ray images are sufficient for localization but not for reconstruction.

Recently, a similar approach to the one presented in this paper was proposed by Hamming *et al.*,[16] dealing with the small size of the reconstructed region in the context of cranial CBCT-based interventions. In that work, fiducial localization was automatic. Initial localization in several projection images was based on the use of an empirically determined threshold, defining regions of interest, followed by template matching with circular templates having a radius of 2–5 pixels. An initial three dimensional estimate is obtained using a reprojection localization approach, similar to that proposed in the current work (Table I, step 3). False positives are rejected using the known number of fiducials and the expectation that only true fiducials will appear in a consistent manner in the sinogram. Finally, localization in the CBCT coordinate system is performed using the least squares ray intersection estimate.

For our thoracic-abdominal images, this approach is not directly applicable. The use of a threshold to define high attenuation regions that include all fiducials will retain large parts of the image. This is due to different overlaps between anatomical structures and the fiducials in the projections. In projection images of the thoracic-abdominal region the same fiducial may overlap with the ribs, with the spinal column, or only with soft tissue. In addition, clinical images most often include high attenuation regions corresponding to wires from the patient monitoring system. These can potentially result in a large number of false positive template matches. Given the characteristics of our data, we chose to initialize our localization with a user defined ROI. This approach takes advantage of the human operator's ability to correctly identify the fiducials in a cluttered scene, implicitly ignoring distractions such as wires or other medical apparatus visible in the projection images.

The primary difference between the approach described by Hamming *et al.* and our own is in the framework used to obtain robustness. In Ref. 16, robustness is obtained using a model based approach, a set of expected conditions. In our approach, we obtain robustness via user initialization and the use of a hypothesize and test method, RANSAC, providing a more generic fiducial localization framework albeit requiring manual initialization.

In Ref. 16, two conditions are expected to hold, all false positive fiducial localizations are correctly rejected and all fiducials are detected. These two conditions cannot be en-

sured in our settings. When false positives are due to additional apparatus visible in the field of view the rejection criterion, consistency of presentation in the sinogram will not work. In our approach, these false detections will not happen due to the manual initialization. The second condition, detection of all fiducials, assumes that there are no false negatives and that all fiducials are physically positioned so that they appear in the projection images. While both our method and that of Hamming *et al.* increase the size of the regions where fiducials can be placed, they are still limited by the size of the detector. In the case of cranial interventions, the spatial region is relatively small and all fiducials can be expected to be visible in the projection images. For thoracic-abdominal regions, the assumption is easily invalidated, with some fiducials placed outside the extent covered by the detector. This is not an issue for our approach as we localize each fiducial separately and do not take into account the total number of fiducials.

It should be noted that both approaches rely on the accurate calibration of the CBCT system. In our case we are ensured that the system calibration is maintained as it is a clinical system utilized for three dimensional imaging on a routine basis.

In future development of our approach, we will investigate the use of template based localization as proposed by Hamming *et al.* This intensity based approach has the potential to improve the localization accuracy in the projection images as compared to our current edge based method, albeit at a higher computational cost.

We have evaluated the proposed approach using 14 data sets and two types of spherical fiducials, 6 mm diameter spheres that are part of a custom registration target and 4 mm diameter CT-SPOTS. The resulting localizations were sufficiently accurate for our targeted thoracic-abdominal interventional radiology procedures. For both types of fiducials, we obtained similar results even though more edge pixels are available when using the larger fiducials. This is most likely due to the fact that both fiducial sizes provide a sufficient number of edge points. In addition, the RANSAC method effectively removes outlying edge elements, ensuring that only correct edge elements are used in the localization.

The proposed approach provides physicians with additional flexibility in placing fiducial markers while maintaining the accuracy currently obtained in clinical practice using direct fiducial localization in the volumetric data. Also, in cases where the proposed approach localizes fiducials that are outside the reconstructed region, there is no need for rescanning the patient, effectively reducing the patient's radiation exposure.

ACKNOWLEDGMENTS

The author thanks Siemens AG Healthcare for the loan of the iGuide CAPPA navigation system and Mr. Moaaz Ali for his help in acquiring the images used in this work. This work was funded by US Army Grant No. W81XWH-04-1-0078. The content of this manuscript does not necessarily reflect the position or policy of the U.S. Government.

APPENDIX A: *N-D* SPHERE

An *n-D* sphere is defined as the set of points that satisfy the following equation:

$$(\mathbf{p}-\mathbf{c})^T(\mathbf{p}-\mathbf{c}) = r^2, \tag{A1}$$

where $\mathbf{c}$ is the center and r is the radius.

We next describe the methods used by the RANSAC algorithm for estimating the sphere from the minimal number of points and two least squares methods.

1. Estimation using *n*+1, minimal number, of points

Assuming all the given points, $\mathbf{p_i} \in \mathbb{R}^n$, are on the sphere we have

$$\forall i,j \quad (\mathbf{p_i}-\mathbf{c})^T(\mathbf{p_i}-\mathbf{c}) = (\mathbf{p_j}-\mathbf{c})^T(\mathbf{p_j}-\mathbf{c}) = r^2.$$

Each pair of points provides us with one linear equation in $n+1$ unknowns

$$(\mathbf{p_i}-\mathbf{p_j})^T\mathbf{c} = 0.5(\mathbf{p_i}^T\mathbf{p_i} - \mathbf{p_j}^T\mathbf{p_j}).$$

Given $n+1$ linearly independent points, we obtain a regular matrix and solve the equation system to get $\mathbf{c}$. The radius is computed as $\|\mathbf{p_i}-\mathbf{c}\|$, where $\mathbf{p_i}$ is arbitrarily chosen from the points we were given.

2. Algebraic least squares

Following Ref. 17, we derive an algebraic least squares estimate. Note that this formulation does not minimize the geometric point to sphere distance which is the quantity we want to minimize.

Given m points in $\mathbb{R}^n$, $m>n+1$, we want to fit them to a sphere such that the sum of the squared algebraic distances is minimized. The algebraic distance is

$$\delta_i = \mathbf{p_i}^T\mathbf{p_i} - 2\mathbf{p_i}^T\mathbf{c} + \mathbf{c}^T\mathbf{c} - r^2.$$

The optimal sphere parameters are computed as

$$[\mathbf{c}^*, r^*] = \arg\min_{\mathbf{c},r} \sum_{i=1}^{m} \delta_i^2,$$

setting $k=\mathbf{c}^T\mathbf{c}-r^2$ we obtain the following linear equation system ($Ax=b$):

$$\begin{bmatrix} -2\mathbf{p_1}^T & 1 \\ \vdots & \vdots \\ -2\mathbf{p_m}^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ k \end{bmatrix} = \begin{bmatrix} -\mathbf{p_1}^T\mathbf{p_1} \\ \vdots \\ -\mathbf{p_m}^T\mathbf{p_m} \end{bmatrix}.$$

The solution of this equation system minimizes $\Sigma_{i=1}^m \delta_i^2 = \|Ax-b\|^2$.

Note that the equation system admits solutions where $k \geqslant \mathbf{c}^T\mathbf{c}$. That is, we have a solution that does not represent a valid sphere, as $r^2 <= 0$. This situation can arise in the presence of outliers.

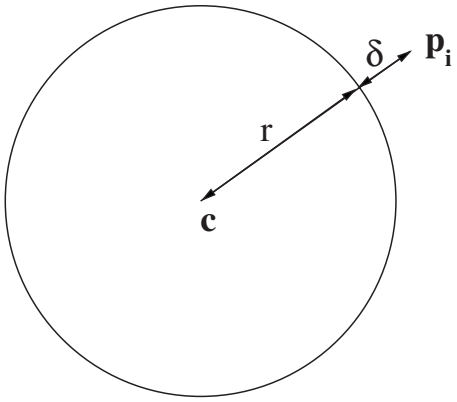

FIG. 11. Orthogonal least squares, point-sphere distance.

3. Geometric least squares

Given m points in $\mathbb{R}^n$, $m>n+1$, we want to fit them to a sphere such that the squared geometric distance is minimized.

The signed geometric distance is (Fig. 11)

$$\delta_i = \|\mathbf{p_i}-\mathbf{c}\| - r.$$

The optimal sphere parameters are computed as

$$[\mathbf{c}^*, r^*] = \arg\min_{\mathbf{c},r} \sum_{i=1}^{m} \delta_i^2 = \arg\min_{\mathbf{c},r} \sum_{i=1}^{m} (\sqrt{(\mathbf{p_i}-\mathbf{c})^T(\mathbf{p_i}-\mathbf{c})} - r)^2.$$

This nonlinear optimization problem is solved using the Levenberg–Marquardt method which requires the computation of δ_i and its partial derivatives with respect to the $n+1$ unknowns,

$$\frac{\partial \delta_i}{\partial c_i} = \frac{c_i - p_i}{\sqrt{(\mathbf{p_i}-\mathbf{c})^T(\mathbf{p_i}-\mathbf{c})}},$$

$$\frac{\partial \delta_i}{\partial r} = -1.$$

APPENDIX B: 3D RAY INTERSECTION

A ray is defined by the equation

$$\mathbf{r}(t) = \mathbf{a} + t\mathbf{n} \text{ s.t., } \quad t \in [0,\infty), \tag{B1}$$

where $\mathbf{a}$, $\mathbf{n} \in \mathbb{R}^3$ are the ray's origin and direction.

To compute the intersection of rays, we ignore the constraint $t_i \in [0,\infty)$, estimate the intersection of lines, and validate that the constraint is satisfied after the fact. We next describe the methods used by the RANSAC algorithm for estimating the ray intersection from the minimal number of rays, and a least squares method.

1. Estimation using two, minimal number, rays

Following Ref. 18, we derive the intersection of two nonparallel rays. Given two such rays we have

$$\mathbf{r_1}(t1) = \mathbf{a_1} + t_1\mathbf{n_1},$$

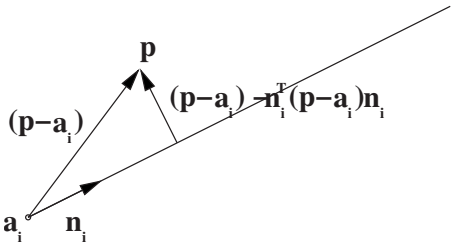

FIG. 12. Distance between point and ray. Note that $\mathbf{n_i}$, ($\|\mathbf{n_i}\|=1$), is the ray direction and not the normal to the ray.

$$\mathbf{r_2}(t2) = \mathbf{a_2} + t_2\mathbf{n_2}.$$

At the intersection point, we have $\mathbf{r_1}(t1)=\mathbf{r_2}(t2)$. Isolating the parameters, t_i, yields the intersection point, or the two points of closest approach on the respective rays. The parameter values are thus

$$t_1 = \frac{((\mathbf{a_2}-\mathbf{a_1})\times\mathbf{n_2})\cdot(\mathbf{n_1}\times\mathbf{n_2})}{\|\mathbf{n_1}\times\mathbf{n_2}\|^2},$$

$$t_2 = \frac{((\mathbf{a_2}-\mathbf{a_1})\times\mathbf{n_1})\cdot(\mathbf{n_1}\times\mathbf{n_2})}{\|\mathbf{n_1}\times\mathbf{n_2}\|^2},$$

with the point of closest approach $[\mathbf{r_1}(t1)+\mathbf{r_2}(t2)]/2$. If the angular deviation between the rays is small, we have a singular configuration as $\|\mathbf{n_1}\times\mathbf{n_2}\|^2 \to 0$.

2. Geometric least squares

The distance between a point $\mathbf{p}$ and a ray $\mathbf{r}(t_i)=\mathbf{a_i}+t_i\mathbf{n_i}$ is (Fig. 12)

$$\delta_i = \|(\mathbf{p}-\mathbf{a_i}) - \mathbf{n_i}^T(\mathbf{p}-\mathbf{a_i})\mathbf{n_i}\|, \quad \|\mathbf{n_i}\| = 1.$$

The intersection is defined as the point that minimizes the sum of squared distances from all rays,

$$\mathbf{p}^* = \arg\min_{\mathbf{p}} \sum_{i=1}^{m} \delta_i^2$$

and in explicit form

$$\Delta = 0.5\sum_{i=1}^{m} \|(\mathbf{p}-\mathbf{a_i}) - \mathbf{n_i}^T(\mathbf{p}-\mathbf{a_i})\mathbf{n_i}\|^2$$

$$= 0.5\sum_{i=1}^{m} [(\mathbf{p}-\mathbf{a_i})^T(\mathbf{p}-\mathbf{a_i}) - \mathbf{n_i}^T(\mathbf{p}-\mathbf{a_i})(\mathbf{p}-\mathbf{a_i})^T\mathbf{n_i}].$$

Deriving Δ with respect to $\mathbf{p}$, we get

$$\frac{\partial \Delta}{\partial \mathbf{p}} = 0.5\sum_{i=1}^{m} [2(\mathbf{p}-\mathbf{a_i}) - 2\mathbf{n_i}^T(\mathbf{p}-\mathbf{a_i})\mathbf{n_i}]$$

$$= m\mathbf{p} - \sum_{i=1}^{m}\mathbf{a_i} - \sum_{i=1}^{m}\mathbf{n_i}^T\mathbf{p}\mathbf{n_i} + \sum_{i=1}^{m}\mathbf{n_i}^T\mathbf{a_i}\mathbf{n_i}.$$

At the optimum the necessary condition, $\partial\Delta/\partial\mathbf{p}=0$, gives us

$$m\mathbf{p} - \left[\sum_{i=1}^{m} \mathbf{n_i}\mathbf{n_i}^T\right]\mathbf{p} = \sum_{i=1}^{m} \mathbf{a_i} - \sum_{i=1}^{m} \mathbf{n_i}^T\mathbf{a_i}\mathbf{n_i}$$

and in matrix notation

$$\begin{bmatrix} m - \sum n_x^2 & -\sum n_x n_y & -\sum n_x n_z \\ -\sum n_x n_y & m - \sum n_y^2 & -\sum n_y n_z \\ -\sum n_x n_z & -\sum n_y n_z & m - \sum n_z^2 \end{bmatrix} \mathbf{p} = \sum_{i=1}^{m} (\mathbf{a_i} - \mathbf{n_i}^T\mathbf{a_i}\mathbf{n_i}).$$

If the angular deviation between all rays is small, the matrix becomes singular.

[a]Electronic mail: zivy@isis.georgetown.edu

[1]D. Ritter, J. Orman, C. Schmidgunst, and R. Graumann, "3D soft tissue imaging with a mobile C-arm," Comput. Med. Imaging Graph. **31**, 91–102 (2007).
[2]M. J. Wallace *et al.*, "Three-dimensional C-arm cone-beam CT: Applications in the interventional suite," J. Vasc. Interv. Radiol. **19**, 799–813 (2008).
[3]K. Wiesent, K. Barth, N. Navab, P. Durlak, T. Brunner, O. Schütz, and W. Seissler, "Enhanced 3D-reconstruction algorithms for C-arm based interventional procedures," IEEE Trans. Med. Imaging **19**, 391–403 (2000).
[4]D. Ritter, M. Mitschke, and R. Graumann, "Markerless navigation with the intra-operative imaging modality SIREMOBIL iso-C3D," Electromedica **70**, 31–36 (2002).
[5]E. B. van de Kraats, B. Carelsen, W. J. Fokkens, S. N. Boon, N. Noordhoek, W. J. Niessen, and T. van Walsum, "Direct navigation on 3D rotational x-ray data acquired with a mobile propeller C-arm: Accuracy and application in functional endoscopic sinus surgery," Phys. Med. Biol. **50**, 5769–5781 (2005).
[6]E. B. van de Kraats, T. van Walsum, L. Kendrick, N. J. Noordhoek, and W. J. Niessen, "Accuracy evaluation of direct navigation with an isocentric 3D rotational x-ray system," Med. Image Anal. **10**, 113–124 (2006).
[7]Z. Yaniv, E. Wilson, D. Lindisch, and K. Cleary, "Electromagnetic tracking in the clinical environment," Med. Phys. **36**, 876–892 (2009).
[8]B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," J. Opt. Soc. Am. A **4**, 629–642 (1987).
[9]J. B. West, J. M. Fitzpatrick, S. A. Toms, C. R. MaurerJr., and R. J. Maciunas, "Fiducial point placement and the accuracy of point-based, rigid body registration," Neurosurgery **48**, 810–817 (2001).
[10]B. E. H. Claus, "Geometry calibration phantom design for 3D imaging," *SPIE Medical Imaging: Physics of Medical Imaging*, edited by M. J. Flynn and J. Hsieh (SPIE, San Diego, CA, 2006), p. 61422E.
[11]J. K. Udupa *et al.*, "A framework for evaluating image segmentation algorithms," Comput. Med. Imaging Graph. **30**, 75–87 (2006).
[12]J. F. Canny, "A computational approach to edge detection," IEEE Trans. Pattern Anal. Mach. Intell. **8**, 679–698 (1986).
[13]M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," Commun. ACM **24**, 381–395 (1981).
[14]R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision* (Cambridge University Press, Cambridge, UK, 2000).
[15]D. Mazilu *et al.*, "Synthetic torso for training in and evaluation of urologic laparoscopic skills," J. Endourol **20**, 340–345 (2006).
[16]N. M. Hamming, M. J. Daly, J. C. Irish, and J. H. Siewerdsen, "Automatic image-to-world registration based on x-ray projections in cone-beam ct-guided interventions," Med. Phys. **36**, 1800–1812 (2009).
[17]W. Gander, G. H. Golub, and R. Strebel, "Least-squares fitting of circles and ellipses," BIT **34**, 558–578 (1994).
[18]R. Goldman, "Intersection of two lines in three-space," *Graphics Gems* (Academic, Cambridge, MA, 1990), p. 304.