Fiducial localization in C-arm based Cone-Beam CT

Ziv Yaniv

Imaging Science and Information Systems (ISIS) Center, Dept. of Radiology,
Georgetown University Medical Center, Washington, DC, USA

ABSTRACT

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. In this paper we show that fiducial localization outside of 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 we obtain the fiducial coordinates in the volume's coordinate system even when the fiducials are outside of the reconstructed region. The approach was evaluated using two anthropomorphic phantoms. When using the projection images all fiducials were localized, including those that were outside the reconstruction volume. The method's maximal localization error as evaluated using fiducials that could be directly localized in the CBCT reconstruction was 0.67 millimeters.

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

1. INTRODUCTION

C-arm based Cone-Beam CT 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. 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] Figure 1 shows corresponding axial slices from an anthropomorphic phantom obtained with a diagnostic CT and a C-arm CBCT system. 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; (5) navigate.

As CBCT provides in-situ imaging it can potentially be used for registration-less navigation,[4–6] streamlining the clinical workflow to only two steps: (1) acquire a 3D image; (2) navigate. This is possible as the 3D image is acquired in-situ, allowing us to replace the intra-operative 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 tracker. 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 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

E-mail: zivy@isis.georgetown.edu

Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling,
edited by Michael I. Miga, Kenneth H. Wong, Proc. of SPIE Vol. 7261, 72610I
© 2009 SPIE · CCC code: 1605-7422/09/$18 · doi: 10.1117/12.811633

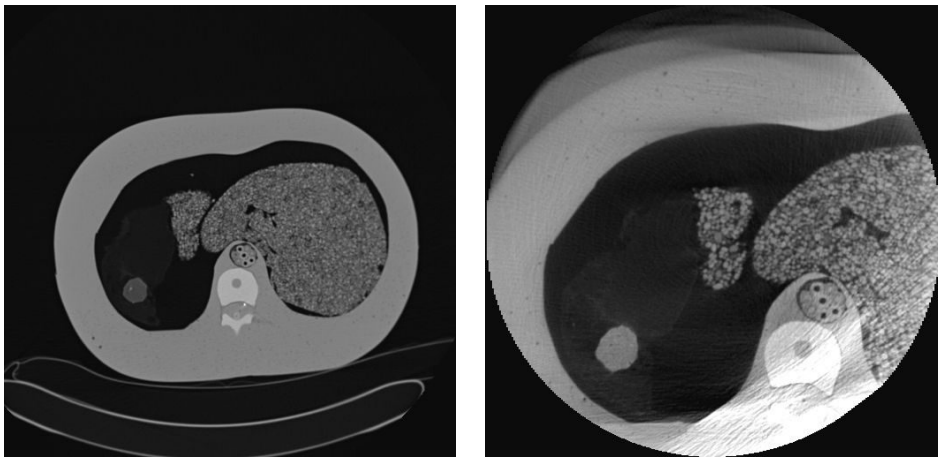

(a) (b)

Figure 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).

frame must be physically connected to the tracker control unit. Given that the clinical environment is most often already cramped, physically positioning the tracking system can become an arduous task. Currently, using electromagnetic trackers for registration-less 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; (4) navigate.

We have developed a navigation system that employs electromagnetic tracking and C-arm based CBCT for thoracic-abdominal interventions. 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, corresponding fiducials are localized in the volumetric image and in the 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.

Figure 2 shows axial and coronal maximum intensity projections of an abdominal phantom acquired with a CT and with a C-arm based CBCT system. The small reconstruction volume of the CBCT system is clearly evident. In many cases, the reconstructed region cannot encompass both the organ of interest and the fiducials placed on the patient's skin. The missing fiducials in the volumetric data can potentially increase the target registration error or may even preclude registration if less than three fiducials are visible.

In this work we show that even when the fiducials cannot be directly localized in the reconstructed volume they can be localized in the volumetric coordinate system, outside of the reconstructed region. Our approach is based on the observation that fiducials that are outside of 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.

We next describe our localization approach in detail.

2. MATERIALS AND METHODS

The CBCT system used in this work is the Axiom Artis dFA (Siemens AG, Medical Solutions, Erlangen, Germany). The system is modelled as a distortion-less pinhole camera. Projection images are acquired using a 20sec rotation in which 543 images are uniformly acquired over a rotation of 220^o. Image dimensions are 960x1024

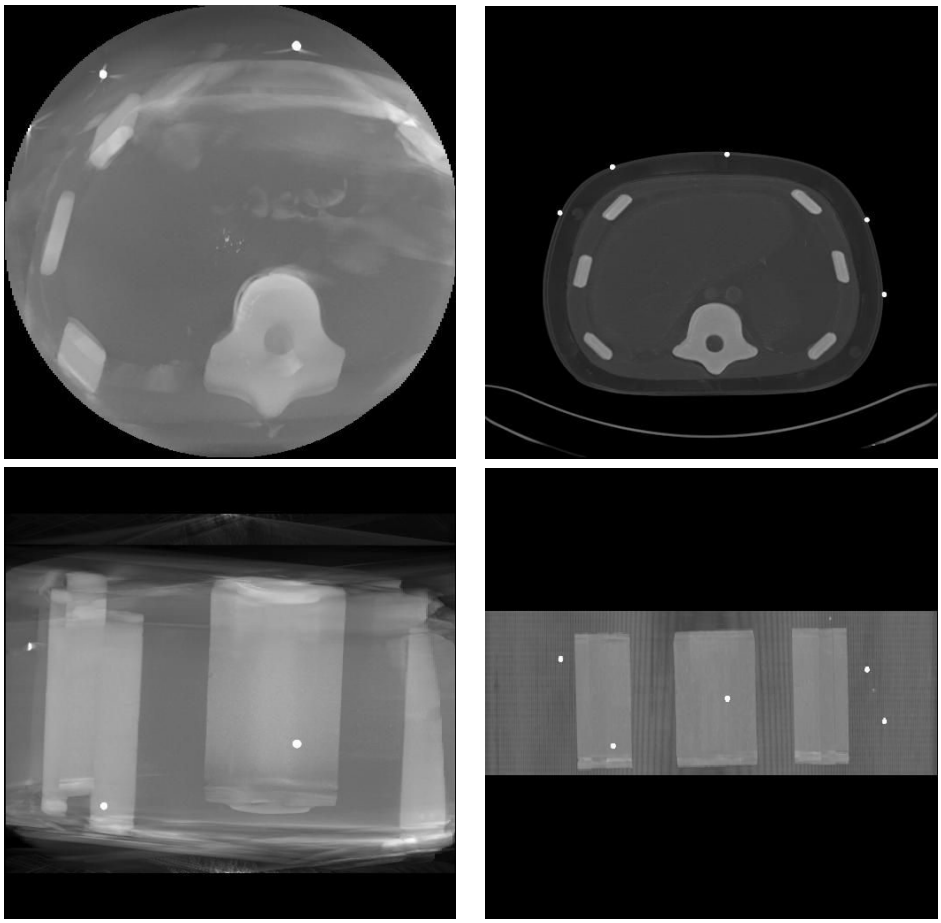

(a) (b)

Figure 2. Maximum intensity projections (MIP) of data acquired using an abdominal phantom with five skin fiducials. Axial and coronal MIPs from (a) CBCT data and (b) CT data. The small reconstruction volume of the CBCT system is clearly evident, with only two fiducials inside the reconstructed region vs. the five visible on CT (white circles).

pixels. Resulting volumetric reconstructions are 512x512x440 with an isotropic $0.45mm$ voxel size. In our experiments we use two anthropomorphic phantoms, the CIRS interventional abdominal phantom (CIRS, Norfolk VA, USA) and a custom anthropomorphic torso model based on the visible human data.[9] The fiducials we use are the 4mm spherical CT-SPOTS markers (Beekley Corp., Bristol CT, USA).

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 volumetric 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 volumetric coordinate system. If these matrices are not accessible, our approach requires additional camera calibration and registration. A calibration object containing easily identifiable markers such as the one used in[10] is imaged and a volumetric reconstruction is performed. The camera's poses with regard to the calibration object are obtained from the projections using standard calibration methods. The transformation from the calibration object coordinate system to the volumetric one can be obtained using paired point registration, identifying corresponding points in the volume and the calibration object. When combined with the camera's pose with regard to the calibration object we obtain the camera pose with regard to the volumetric coordinate system.

To localize a fiducial in the volumetric coordinate system it is first localized in multiple projection images. Its spatial location is then estimated as the intersection point of the backprojeced 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 projective images. In essence this is a segmentation task.

2.1 Fiducial localization

Following[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 later. We thus follow a semi automatic approach to localizing fiducials. 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.

Initially, fiducial localization is performed semi-automatically in a single projection image. The user indicates a rectangular region of interest containing the 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] To utilize the RANSAC algorithm we require a method for parameter estimation using the minimal number of data elements, a method for least 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 correctly 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 $\pm 8^o$. 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 ${}^lP$, ${}^rP$:

$$F = [{}^rP\,{}^lC]_\times\,{}^rP\,{}^lP^+$$

where ${}^lP^+$ is the pseudo inverse of ${}^lP$, ${}^rP\,{}^lC$ 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.

Input: Projection matrices and corresponding images.
Output: Fiduical 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, $\pm 8^0$ 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 back-projected ray emanating from the X-ray source and going through the fiducial location.
   (c) Estimate the point that minimizes the distance to all back-projected rays using the RANSAC algorithm.
3. Reprojection localization
   (a) Reproject the estimated 3D location onto every n'th (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 back-projected ray emanating from the X-ray source and going through the fiducial location.
   (d) Estimate the point that minimizes the distance to all back-projected rays using the RANSAC algorithm.

Table 1. Fiducial localization in volumetric coordinate system.

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 four degrees. 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 fiducials 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.

Table 1 and Figure 3 summarize our fiducial localization approach.

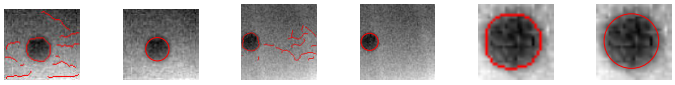

(a) (b) (c) (d) (e) (f)

Figure 3. Graphical representation of localization approach: initial localization (a, b), constrained localization (c, d), reprojection localization (e, f).

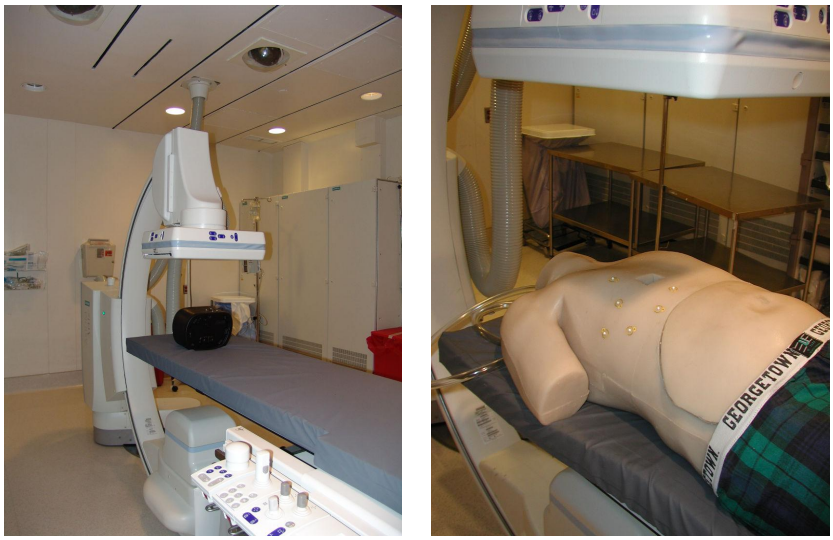

(a) (b)

Figure 4. Experimental setup with our two anthropomorphic phantoms (a) CIRS abdominal phantom and (b) custom torso phantom.

3. EXPERIMENTAL RESULTS

To evaluate our approach we performed the following experiment. Multiple fiducial markers were placed on an anthropomorphic phantom and a 20 second clinical image acquisition scan was performed. The experiment was performed once with the CIRS phantom using five fiducials and once with the custom torso phantom using six fiducials. Figure 4 shows the experimental setup.

We evaluated our approach both qualitatively, the number of localized fiducials, and quantitatively, how accurate is our localization in the volumetric coordinate system. The later evaluation is only applicable to fiducials that are inside the spatial extent of the reconstructed volume. Finally, we compared the quality of our approach to one in which a user semi-automatically localized the fiducial in two projection images with an angular difference of approximately 40^o. The fiducial's volumetric location was again obtained via intersection of the backprojected rays. The choice of this specific angular difference corresponds to the maximal range in which one of the fiducials was actually visible.

For quantitative evaluation purposes we semi-automatically localized all fiducials that were visible in the volumetric reconstruction. The user provided a three dimensional region of interest and the spherical fiducial was detected using the RANSAC framework. This is the three dimensional equivalent to the initial phase of our approach. Example results are shown in Figure 5.

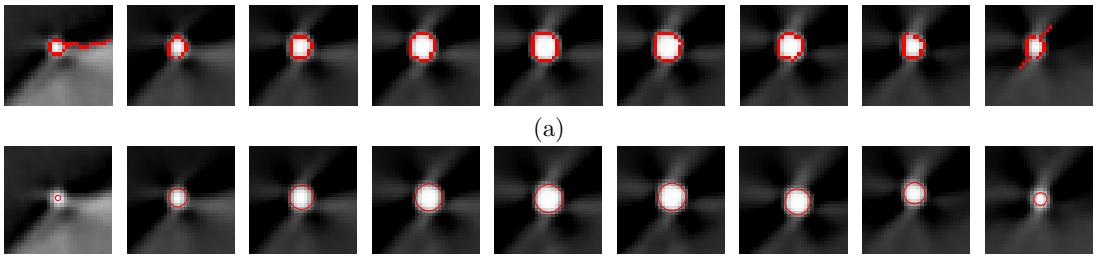

(a)

(b)

Figure 5. Localization of spherical fiducial in CBCT reconstruction using the RANSAC algorithm. Overlaid in red are (a) results of edge detection and (b) intersection of localized fiducial with axial slice.

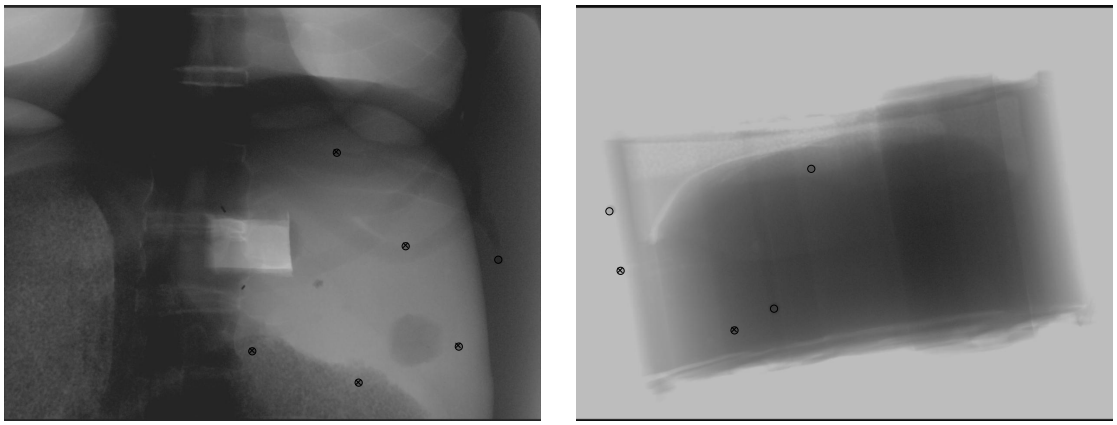

(a) (b)

Figure 6. Fiducial locations in 3D are projected onto the images. Fiducials localized directly in volumetric data are marked with an 'x' and those localized using our approach with an 'o'.

In both experiments our approach successfully localized all fiducials in the volumetric coordinate system. This is in contrast to the direct localization in the volumetric data. In the experiment using the CIRS phantom only two out of the five fiducials were visible in the reconstruction, precluding the use of this data set for navigation. In the experiment using the custom torso phantom five of the six fiducials were visible in the reconstruction. Figure 6 provides a visual summary of the qualitative evaluation.

In both experiments our approach and the method of user localized fiducials in two images provided accurate localizations in the volumetric coordinate system. The maximal difference from direct volumetric localization was less than $0.85mm$ when using two images and less than $0.67mm$ when using our approach. Figure 7 summarizes the quantitative evaluation.

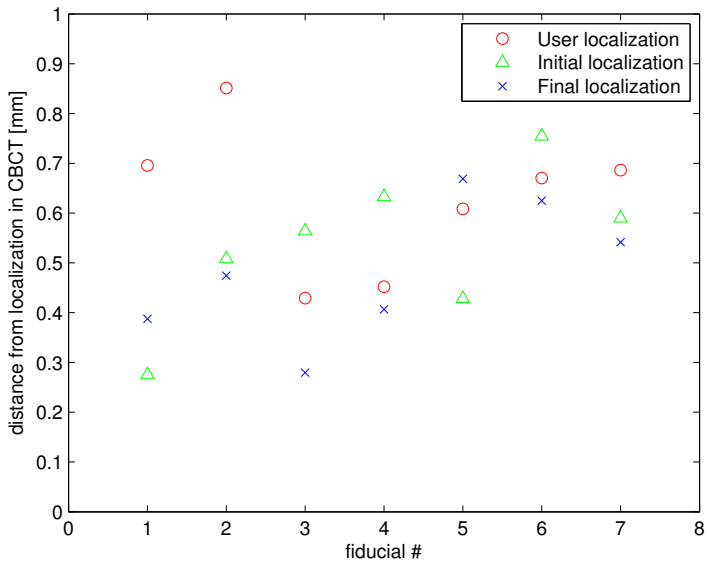

Figure 7. Distances between fiducials that were localized in the CBCT reconstruction and their location as estimated using the proposed approach (initial 2D fiducial subset and final 2D fiducial subset), and the approach in which the user semi-automatically localized the fiducial in two images.

4. DISCUSSION AND CONCLUSIONS

We have presented a method for accurately localizing fiducials in the volumetric coordinate system even when their spatial location is outside of the reconstructed volume. Our approach requires minimal user interaction, indicating a rectangular region of interest surrounding the fiducial in a single image. Our 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.

We evaluated the proposed approach both qualitatively and quantitatively, when possible, and found that it was able to detect all fiducials even those that could not be identified in the volumetric data and that its accuracy is equivalent to direct localization in the volumetric data.

ACKNOWLEDGMENTS

The work was funded by US Army grant W81XWH-04-1-0078. The content of this manuscript does not necessarily reflect the position or policy of the U.S. Government.

REFERENCES

1. D. Ritter, J. Orman, C. Schmidgunst, and R. Graumann, "3D soft tissue imaging with a mobile C-arm," *Comput. Med. Imaging Graph.* **31**(2), pp. 91–102, 2007.
2. M. J. Wallace *et al.*, "Three-dimensional C-arm cone-beam CT: Applications in the interventional suite," *Journal of Vascular and Interventional Radiology* **19**(6), pp. 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. Imag.* **19**(5), pp. 391–403, 2000.
4. D. Ritter, M. Mitschke, and R. Graumann, "Markerless navigation with the intra-operative imaging modality SIREMOBIL iso-C3D," *Electromedica* **70**(1), pp. 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**(24), pp. 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," *Medical Image Analysis* **10**(2), pp. 113–124, 2006.
7. Z. Yaniv, E. Wilson, D. Lindisch, and K. Cleary, "Electromagnetic tracking in the clinical environment," *Med. Phys. in press* .
8. B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Journal of the Optical Society of America A* **4**, pp. 629–642, April 1987.
9. D. Mazilu *et al.*, "Synthetic torso for training in and evaluation of urologic laparoscopic skills," *J Endourol. 2006 May;20(5):340-5.Click here to read* **20**(5), pp. 340–345, 2006.
10. Z. Yaniv and K. Cleary, "Fluoroscopy based accuracy assessment of electromagnetic tracking," in *SPIE Medical Imaging: Visualization, Image-Guided Procedures, and Display*, K. Cleary and R. Galloway, eds., pp. 168–174, SPIE, 2006.
11. J. K. Udupa *et al.*, "A framework for evaluating image segmentation algorithms," *Comput. Med. Imaging Graph.* **30**(2), pp. 75–87, 2006.
12. J. F. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Machine Intell.* **8**(6), pp. 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**(6), pp. 381–395, 1981.
14. R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, 2000.
15. W. Gander, G. H. Golub, and R. Strebel, "Least-squares fitting of circles and ellipses," *BIT* **34**(4), pp. 558–578, 1994.
16. R. Goldman, "Intersection of two lines in three-space," in *Graphics Gems*, p. 304, Academic Press, 1990.

APPENDIX A. N-D SPHERE

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

$$(\mathbf{p}-\mathbf{c})^T(\mathbf{p}-\mathbf{c}) = r^2 \tag{1}$$

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.

A.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.

A.2 Algebraic least squares

Following[15] 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:

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

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 $k \geq \mathbf{c}^T\mathbf{c}$. 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.

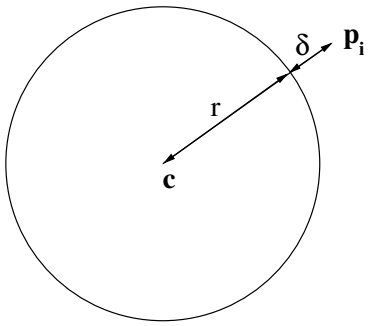

Figure 8. Orthogonal least squares, point-sphere distance.

A.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 (Figure 8):

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

The optimal sphere parameters are computed as:

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

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:

$$\begin{aligned} \frac{\partial \delta_i}{\partial c_i} &= \frac{c_i - p_i}{\sqrt{(\mathbf{p_i} - \mathbf{c})^T(\mathbf{p_i} - \mathbf{c})}} \\ \frac{\partial \delta_i}{\partial r} &= -1 \end{aligned}$$

APPENDIX B. RAY INTERSECTION

A ray is defined by the equation:

$$\mathbf{r}(t) = \mathbf{a} + t\mathbf{n} \quad \text{s.t.} \quad t \in [0, \infty) \tag{2}$$

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.

B.1 Estimation using 2 (minimal number) points

Following[16] we derive the intersection of two non-parallel rays. Given two such rays we have:

$$\begin{aligned} \mathbf{r_1}(t1) &= \mathbf{a_1} + t_1\mathbf{n_1} \\ \mathbf{r_2}(t2) &= \mathbf{a_2} + t_2\mathbf{n_2} \end{aligned}$$

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

$$\begin{aligned}
\mathbf{a_1} + t_1\mathbf{n_1} &= \mathbf{a_2} + t_2\mathbf{n_2} && \text{/subtracting } \mathbf{a_1} \\
t_1\mathbf{n_1} &= \mathbf{a_2} - \mathbf{a_1} + t_2\mathbf{n_2} && \text{/crossing with } \mathbf{n_2} \\
t_1(\mathbf{n_1} \times \mathbf{n_2}) &= (\mathbf{a_2} - \mathbf{a_1}) \times \mathbf{n_2} && \text{/dotting with } (\mathbf{n_1} \times \mathbf{n_2}) \\
t_1(\mathbf{n_1} \times \mathbf{n_2}) \cdot (\mathbf{n_1} \times \mathbf{n_2}) &= ((\mathbf{a_2} - \mathbf{a_1}) \times \mathbf{n_2}) \cdot (\mathbf{n_1} \times \mathbf{n_2}) && \text{/dividing by } (\|\mathbf{n_1} \times \mathbf{n_2}\|^2) \\
t_1 &= \tfrac{(\mathbf{a_2}-\mathbf{a_1})\times\mathbf{n_2})\cdot(\mathbf{n_1}\times\mathbf{n_2})}{\|\mathbf{n_1}\times\mathbf{n_2}\|^2}
\end{aligned}$$

Isolating t_2 is done in a similar manner, and the point of closest approach is given as $\frac{\mathbf{r_1}(t1)+\mathbf{r_2}(t2)}{2}$. Note that if the angular deviation between the rays is small we have a singular configuration as $(\mathbf{n_1} \times \mathbf{n_2}) \to 0$.

B.2 Geometric least squares

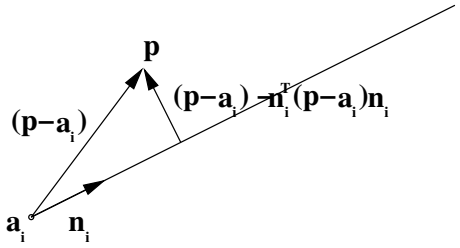

Figure 9. Distance between point and ray. Note that $\mathbf{n_i}, (\|\mathbf{n_i}\| = 1)$, is the ray direction and not the normal to the ray.

The distance between a point $\mathbf{p}$ and a ray $\mathbf{r}(t_i) = \mathbf{a_i} + t_i\mathbf{n_i}$ is (Figure 9):

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

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

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

and in explicit form:

$$\begin{aligned}
\Delta = \quad & 0.5 \textstyle\sum_{i=1}^m \|(\mathbf{p} - \mathbf{a_i}) - \mathbf{n_i}^T(\mathbf{p} - \mathbf{a_i})\mathbf{n_i}\|^2 = \\
& 0.5 \textstyle\sum_{i=1}^m [(\mathbf{p} - \mathbf{a_i})^T(\mathbf{p} - \mathbf{a_i}) - \mathbf{n_i}^T(\mathbf{p} - \mathbf{a_i})(\mathbf{p} - \mathbf{a_i})^T\mathbf{n_i}]
\end{aligned}$$

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

$$\begin{aligned}
\frac{\partial\Delta}{\partial\mathbf{p}} &= 0.5\sum_{i=1}^{m}\left[2(\mathbf{p} - \mathbf{a_i}) - 2\mathbf{n_i}^T(\mathbf{p} - \mathbf{a_i})\mathbf{n_i}\right] \\
&= m\mathbf{p} - \sum_{i=1}^{m}\mathbf{a_i} - \sum_{i=1}^{m}\mathbf{n_i}^T\mathbf{p}\mathbf{n_i} + \sum_{i=1}^{m}\mathbf{n_i}^T\mathbf{a_i}\mathbf{n_i}
\end{aligned}$$

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

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

and in matrix notation:

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

Note that when all rays are parallel we have a singular matrix:

$$A = \begin{bmatrix} m-mn_x^2 & -mn_x n_y & -mn_x n_z \\ -mn_x n_y & m-mn_y^2 & -mn_y n_z \\ -mn_x n_z & -mn_y n_z & m-mn_z^2 \end{bmatrix}$$

as $det(A) = m^3(1-n_x^2-n_y^2-n_z^2)$ and $\|n\| = 1$, which gives $det(A) = 0$.