

FTRAC—A robust fluoroscope tracking fiducial

Ameet Kumar Jain^{a)}

Department of Computer, Science, Johns Hopkins University, Baltimore, Maryland 21218

Tabish Mustafa

Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Yu Zhou

Department of Mechanical Engineering, Stony Brook University, Stony Brook, New York 11794

Clif Burdette

Acoustic MedSystems, Inc., 2110 Clearlake Blvd., Champaign, Illinois 61822

Gregory S. Chirikjian

Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218

Gabor Fichtinger^{a)}

Department of [Computer Science, Mechanical Engineering, Radiology], Johns Hopkins University, Baltimore, Maryland 21218

(Received 2 November 2004; revised 26 May 2005; accepted for publication 10 August 2005; published 29 September 2005)

C-arm fluoroscopy is ubiquitous in contemporary surgery, but it lacks the ability to accurately reconstruct three-dimensional (3D) information. A major obstacle in fluoroscopic reconstruction is discerning the pose of the x-ray image, in 3D space. Optical/magnetic trackers tend to be prohibitively expensive, intrusive and cumbersome in many applications. We present single-image-based fluoroscope tracking (FTRAC) with the use of an external radiographic fiducial consisting of a mathematically optimized set of ellipses, lines, and points. This is an improvement over contemporary fiducials, which use only points. The fiducial encodes six degrees of freedom in a single image by creating a unique view from any direction. A nonlinear optimizer can rapidly compute the pose of the fiducial using this image. The current embodiment has salient attributes: small dimensions ($3 \times 3 \times 5$ cm); need not be close to the anatomy of interest; and accurately segmentable. We tested the fiducial and the pose recovery method on synthetic data and also experimentally on a precisely machined mechanical phantom. Pose recovery in phantom experiments had an accuracy of 0.56 mm in translation and 0.33° in orientation. Object reconstruction had a mean error of 0.53 mm with 0.16 mm STD. The method offers accuracies similar to commercial tracking systems, and appears to be sufficiently robust for intraoperative quantitative C-arm fluoroscopy. Simulation experiments indicate that the size can be further reduced to $1 \times 1 \times 2$ cm, with only a marginal drop in accuracy. © 2005 American Association of Physicists in Medicine. [DOI: 10.1118/1.2047782]

Key words: C-arm, fluoroscopy, tracking, reconstruction, registration, prostate brachytherapy

I. INTRODUCTION

C-arm fluoroscopy is the most widely used intraoperative imaging modality in general surgery (some of the prominent works being Refs. 1–8), but it presently lacks the ability for robust and easy quantitative guidance.⁹ Quantitative fluoroscopy-guided surgery needs to solve four major problems: (1) C-arm image distortion; (2) the calibration of imaging parameters; (3) pose recovery or tracking; and (4) registration to imaging modalities. The first two are well-studied problems in the literature.^{1,3,9} On the other hand, pose recovery on unencoded C-arm machines is a major technical problem that presently does not have a clinically practical solution in many areas of application. In this paper, we propose a solution to accurately estimate the pose by using a radio-opaque fiducial. Moreover, it can also register the x-ray images to other imaging modalities.

The history of using radio-opaque fiducials to solve for the C-arm pose starts as early as 1987,¹⁰ where 16 spherical beads in a known configuration were introduced in the region of interest. Theorems from projection geometry were used to reconstruct the C-arm pose, calibration parameters and target details. This was later followed by other bead-based fiducials^{2,11–13} for solving both the problems of a C-arm pose and calibration (a small survey¹⁴). The number of beads varied from 20 to 613, with the patient sometimes fitting inside them. Though the accuracies were acceptable (0.5–1 mm), the fiducials tended to be (a) too large; (b) cumbersome to use in a clinical setting; (c) interfering with the anatomy in the image; and (d) nontrivial to segment. A smaller and well-encoded fiducial can solve these problems efficiently.

Off-the-shelf tracking devices, as they became available, were a natural choice to simplify the problem. In current commercial C-arm fluoroscopy surgical navigation systems,

the x-ray detector is localized in room coordinates by some auxiliary optical tracker^{15,16} or electromagnetic (EM) tracker.¹⁷ Unfortunately, auxiliary trackers sometimes become impractical for various reasons. They are expensive and add to the complexity of the operating room since they require an additional calibration step. Optical trackers require a line of sight, which becomes cumbersome in a clinical setting and requires an alteration in the standard workflow. The EM trackers can successfully overcome this issue, but become susceptible to field distortion from metal objects like surgical tools or the *C* arm itself, compromising on accuracy. This has led some recent researchers to prefer fiducial-based tracking. In a recent publication,¹⁸ the authors delineate the above problems and also say that using optical trackers reduces the useful imaging volume of the fluoroscope and potentially compromises the achievable accuracies. Despite using an optical tracker to track their surgical tools, they explicitly choose to not track the *C* arm using the tracker but instead use a radio-opaque fiducial. Their system has been fairly successful for various surgeries and has been in clinical use for the last four years. Similar preferences can be observed in other recent publications as well.^{3,19–22}

To make the fiducials feasible, recent publications have reported smaller fiducials by compact bead placement. It should be noted that this increase in clinical friendliness by decreasing the size was achieved at the cost of a decrease in accuracy. The typical number of beads were reported to be between 6–28,^{18–22} achieving translation accuracies ranging from 1–3 mm and orientation accuracies around 1°–2°. The variation in accuracies is governed by both bead configuration¹⁹ and implementation²⁰ choices. Increasing the number of beads greater than 6 provides little improvement in the accuracy.¹⁹ Moreover, nonlinear optimization fares far superior when compared to linear methods.^{20,23} Thus, from recent publications, it seems that a 1 mm error in translation and 1° error in rotation is probably the best that bead-based fiducials can achieve. The reason for this is that accurate segmentation of beads in the x-ray image is nontrivial. A 2 mm diameter radio-opaque bead would project as an ellipse with a 12 pixel long axis. The center of the ellipse does not necessarily correspond to the center of bead, contributing at least 1–2 pixel segmentation error. Moreover, a further decrease in fiducial size would greatly compromise the accuracy.²⁴

Methods to accommodate for and decrease the effect of such systematic biases²⁵ are studied by the computer vision community during their camera calibration procedures. Unfortunately, the improvement is not substantial. Moreover, point-based registration methods are known to be sensitive to segmentation errors.²⁴ Thus a totally new approach toward fiducial design will be required. Moreover, automatic segmentation and establishing correspondences between the fiducial beads and their projections in the image is itself a nontrivial task for arbitrary positions of the fiducial. In a synergistic problem, the end effector of a surgical robot was registered to a *C* arm using line fiducials³ mounted on the robot's tool holder. The idea was to register the x-ray image to a robot, in order to locate the target anatomy in the coordinate

frame of the surgical robot. Although their helical fiducial did not allow for sufficiently accurate and robust pose recovery, it was an important effort that inspired our work. Our present contribution to the state of the art is the development of a novel fiducial system and its mathematical framework. The fiducial uses ellipses and straight lines in addition to points. It offers a decrease in size, an enhancement in accuracy and robustness in *C*-arm pose recovery, allowing for precise object reconstruction. Some preliminary studies have been published earlier,^{21,26} while we present the complete mathematical framework, results and analysis in this paper.

II. MATERIALS AND METHODS

Our solution to image-based pose recovery is to mount a stationary fiducial in the field of the *C*-arm fluoroscope. The fiducial, by design, encodes six degrees of freedom (DOF) from a single image by creating a unique view from any direction. The accuracy and robustness of pose recovery would critically depend on the design and precise manufacturing of the fiducial. Software was created for segmentation of the fiducial and for numerical optimization that determines the six-DOF transformation between the coordinate frames of the *C* arm and the fiducial.

A. Fiducial design and manufacturing

For practical viability the fiducial should (a) be noninvasive with small dimensions; (b) not corrupt image quality; (c) not necessarily be close to the anatomy of interest; (d) be easily/accurately segmentable from x-ray images; and (e) have software that is robust, fast, and simple to use.

1. Choice of fiducial features

Geometric aspects of x-ray imaging are modeled using the perspective projection model (Sec. II B), under which we can limit the choice of features to well-behaved features like points, straight lines, ellipses, parabolas, helices, and other curves. Different features exhibit varying segmentation accuracies, each combination thus producing a fiducial with different characteristics. Segmentation is crucial, since it directly relates to achievable pose estimation accuracy. Image processing algorithms can more accurately segment constrained (parametric) curves when compared to unconstrained curves, usually to a subpixel precision.

Projected point features have a high segmentation error²⁵ (~1–2 pixel) but make the pose estimation problem very well constrained.²⁷ Straight lines are not as well constrained as points, but offer superior segmentation accuracies to points. Linear methods are available for both features,^{28–30} but it should be noted that due to an inability to incorporate all the rotation constraints, linear methods have been shown to not fare as well as nonlinear methods.^{20,23} A helix is a parametric curve that projects as a nonparametric curve with nondifferentiable segments, leading to inaccurate segmentation, which makes pose estimation prone to errors.³ Alternatively, ellipses, parabolas, and hyperbolas are conics that will

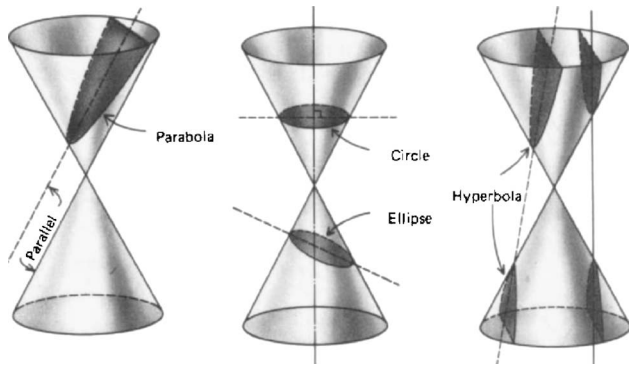


FIG. 1. The center of the cone is the x-ray source, while the base is the image plane. The intersection of this plane with the cone is an ellipse. Any 3D curve combined with the x-ray source defines a unique cone, which intersects the image plane as an ellipse. Any conic in general projects as another conic under perspective geometry.

always project as another member of the same family. The proof of this is self-explanatory from Fig. 1. Moreover, since the mean direction of the x-ray beam is near-orthogonal to the image plane, we can, without loss of generality, assume that the projection of all conics would be an ellipse on the image plane. Parabolas and hyperbolas would project as open ellipses, making segmentation prone to error. In comparison, a 3D ellipse will project as a closed ellipse in the image, making it an ideal candidate feature.^{31,32} A 3D circle is a special case of the ellipse.^{33–35} It should be noted here that even though ellipses are easy to segment, exact linear techniques for pose estimation from ellipses are not available.^{36,37} Some authors have attempted to linearize pose estimation using an ellipse,²⁹ but not with high accuracy.

Thus we choose a combination of ellipses, straight lines, and points as our primary features, balancing between accuracy and convergence. Ellipses make the algorithm more accurate, points make it more robust, and lines will achieve a bit of both.

2. Necessary and sufficient feature combinations

A combination of the above features can be used for pose estimation. The number and placement of the features will decide the robustness of the fiducial. It can be shown geometrically that the correspondence of at least three points is required for determining the pose of the fiducial uniquely. Using the point-line duality, we can also conclude that a minimum of three straight lines (not line segments but infinite lines) is necessary to find the pose uniquely. A single 3D circle limits the pose to four discrete symmetric choices, along with an indeterminable rotation along the normal axis. An ellipse is able to limit this to only 16 discrete positions. Thus, an ellipse/circle along with a point, can uniquely encode the pose of the fiducial. Open curves like the helix, parabola, and hyperbola can limit the pose to just two discrete rotations.

Note that these are purely theoretical considerations and using the bare minimum number of features will result in a loss of accuracy. Inaccuracies in the imaging model coupled




FIG. 2. Images of the FTRAC fiducial (a) cylindrical sheath design; (b) wire model; (c) photograph; (d) x-ray image.

with image quantization errors, would make slight variations from the actual pose produce nearly identical images. In other terms, the Jacobian for image formation loses rank and becomes singular. Singular matrices are undesirable because they cannot be inverted, and indicate that the system has become degenerate. Each feature has blind spots, near which the Jacobian would be singular. Near-singular matrices are theoretically invertible but numerically unstable, resulting in large errors during matrix inversion. For example, the ellipse projects as a straight line in its own plane, and hence the Jacobian would be ill conditioned (near-singular) when the actual pose is in the vicinity of this blind spot. Thus, the placement of all the features should be such that a singularity from one feature is cancelled by a good Jacobian from another. This suggests that a highly redundant structure is needed to accommodate a sufficient range of projection angles and compensate for inevitable calibration and image processing errors.

Our first motivating application is prostate brachytherapy (Sec. IV), where the fluoroscope has a limited view within a 25° cone around the AP axis. To tackle blind spot singularity, we choose to arrange two noncoplanar ellipses 60° from each other (Fig. 2). When one ellipse projects with high eccentricity, the other ellipse projects as a circle, forcing at least one ellipse to encode strongly in any direction. We added three line segments as features, further utilizing the endpoints of the segments as six point features. To make it nonsymmetric, three additional points were added at known distances along the lines. Thus, one segment has four points, another has three, while the third has just two. This fact is also used to distinguish between the ellipses, which otherwise have the same dimensions. These features constitute our fiducial, which henceforth we shall refer to as the fluoroscope track-

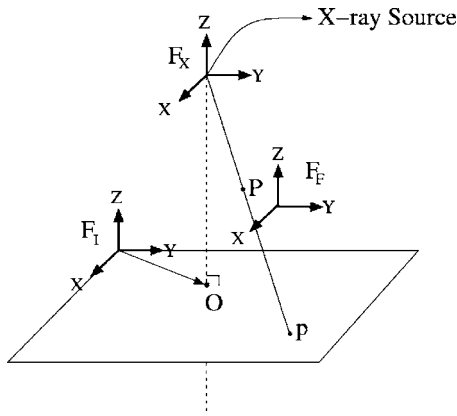


Fig. 3. Projective geometry and notations for fluoroscopic imaging.

ing (FTRAC) fiducial. The features are placed carefully so that the projection of radioactive seeds in the prostate does not overlap with that of the fiducial, even when they are beside each other. All the features of this fiducial are mounted on the surface of a hollow cylinder, which remains stationary. The diameter of the cylinder is 3 cm. Thus, all the features fit inside a $3 \times 3 \times 5$ cm volume.

3. Manufacturing

The design for the fiducial was refined by generating multiple rough prototypes from ABS (acrylonitrile butadiene styrene) using an FDM (fused deposition modeling) rapid prototype machine. The final fiducial design was then fabricated from a acetol rod using a four-axis CNC (computer numerical control) mill, with the grooves of the fiducial geometry machined into its surface to press fit 0.5 mm stainless steel wires and 2 mm stainless steel beads. The cylinder was press fit into a custom acetol mount that provided three mutually orthogonal sets of mounting holes for attaching to an accurate rotary table for validation.

B. Pose recovery mathematics

1. C-arm imaging

Geometric aspects of fluoroscopic imaging can be modeled as a perspective transformation with five parameters—focal length, image origin, and pixel size (Fig. 3). The transformation formula is given in Eqs. (1) and (2); where F_F, F_X, F_I are the coordinate frames of the fiducial, x-ray source, and the image, respectively; P is a 3D point; P_F are the homogenous coordinates of P in F_F ; p is the projection of P on the image plane; p_I are the homogenous coordinates (in pixels) of p in frame F_I ; ${}^X F_F$ is the 4×4 rigid transformation matrix that transforms a point in F_F to F_X ; ${}^I F_X$ is the 3×4 perspective projection matrix; f is the focal length; $O = (o_x, o_y)$ is the projection of the x-ray source on the image plane (later referred to as the origin); s_x and s_y are the pixel sizes along the X and Y axes of the image,

$$p_I = {}^I F_X {}^X F_F P_F, \quad (1)$$

$${}^I F_X = \begin{bmatrix} -f/s_x & 0 & o_x & 0 \\ 0 & -f/s_y & o_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad (2)$$

$${}^X F_F = \begin{bmatrix} r_{11} & r_{12} & r_{13} & T_1 \\ r_{21} & r_{22} & r_{23} & T_2 \\ r_{31} & r_{32} & r_{33} & T_3 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Assuming that the C arm is calibrated (${}^I F_X$ is known), the problem remains to estimate the pose of the fiducial (${}^X F_F$), given the image. ${}^X F_F$ consists of six independent parameters— T_1, T_2, T_3 are the translation parameters, while ϕ_1, ϕ_2, ϕ_3 are the three Euler rotation angles along the x, y , and z axes [which determine r_{ij} in Eq. (2) uniquely].

In this framework, the exact solution is rendered nonlinear and we use the iterative Newton's optimization method to solve it. We segment the fiducial from the image, solve the correspondence problem using the nonsymmetry of the design, and then feed this to the optimizer. The optimizer as a first step computes an expected image location of the fiducial. It also computes a Euclidean-distance-based multidimensional error function. It is assumed that at the correct pose, this error is zero. A closed form of the Jacobian is provided, to ensure quick and accurate convergence.

2. Error formulas for points, lines, and ellipse

The multidimensional error vector measures the Euclidian distance (in the image) between the segmentation of the observed projection of the fiducial, and that of the current projection. A segmentation algorithm for the current image is not needed since the position of any feature in the current image is computed directly using its 3D location on the fiducial. The distance from each feature in the current projection, to its corresponding counterpart in the observed image, constitutes one element in the vector. When all features match perfectly, each element in the vector is zero.

The Euclidian distance for corresponding point features is the difference in their image coordinates. Let (X_i^m, Y_i^m, Z_i^m) be the coordinate of P_i in F_F (model/FTRAC frame) and (X_i, Y_i, Z_i) be its coordinates in F_X . Let (x_i, y_i) be the current coordinates of p_i in F_I . p_i is a function of the six parameters $(\phi_1, \phi_2, \phi_3, T_1, T_2, T_3)$, which we shall refer to by a rotation matrix \mathcal{R} and a translation vector \mathcal{T} . Let (\bar{x}_i, \bar{y}_i) be the coordinates of the observed point p_i in the image. Since each iteration of the method has its own rotation and translation parameters, we shall refer to them by their iteration number k as, \mathcal{R}_k and \mathcal{T}_k , with \mathcal{R}_0 and \mathcal{T}_0 being the initial estimate. δ is the error function. Under this notation, the current estimate of the distance function for points can be written as

$$\delta x_i = x_i(\mathcal{R}_k, \mathcal{T}_k) - \bar{x}_i, \quad (3)$$

$$\delta y_i = y_i(\mathcal{R}_k, \mathcal{T}_k) - \bar{y}_i. \quad (4)$$

Note that the L_2 norm of these two values is the exact Euclidian distance between the two points, and that the optimizer minimizes the L_2 norm of the vector.

There is no notion of the Euclidian distance between two straight lines. Thus, we select a point on one straight line and measure the distance of this point from the other straight line. Any two points uniquely determine a line, and hence they can be sufficient to measure the distance between the straight lines. If the equation of a straight line in the image is $Ax+By+C=0$, then the distance of (x_i, y_i) from the line is

$$\delta(x_i, y_i) = \frac{Ax_i + By_i + C}{\sqrt{A^2 + B^2}}. \tag{5}$$

Any two points on the line segment can be used, but we use the endpoints for robustness. A similar technique can be used for an ellipse, but no simple closed-form formula exists for the Euclidian distance of a generic point from an ellipse. In fact, this distance is one of the solutions to a fourth degree polynomial,³⁸ from which the minimum distance solution has to be chosen. We derive an approximation that works well when the point is not too far away from the ellipse. The general equation of an ellipse (conic) is $P^T M P = 0$,

$$[x \ y \ 1] \begin{bmatrix} A & B/2 & D/2 \\ B/2 & C & E/2 \\ D/2 & E/2 & F \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0, \tag{6}$$

which can be rotated, translated, and scaled to give $P^T F^T M F P = 0 (P^T M' P = 0)$,

$$[x \ y \ 1] \begin{bmatrix} 1/a^2 & 0 & 0 \\ 0 & 1/b^2 & 0 \\ 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = 0, \tag{7}$$

where a, b are the length of the major and minor axes. $P^T M P$ is a good distance metric since it is rotation and translation invariant. Unfortunately, it is a nonlinear metric but can be scaled near the ellipse boundary to give an approximate Euclidian distance. Using Eq. (7), the scale factor after some first-order approximations turns out to be $K = ab(AC - B^2/4)/2\sqrt{a^2 + b^2} \det(M)$. The final distance function from any point to the ellipse is shown in Eq. (8). Theoretically, a minimum of 5 points are required to localize an ellipse exactly, though we use 12 for robustness,

$$\begin{aligned} \delta(x_i, y_i) &= K P^T M P \\ &= \frac{ab}{2\sqrt{a^2 + b^2}} \frac{(AC - B^2/4)}{\det(M)} [x_i \ y_i \ 1] \\ &\quad \times \begin{bmatrix} A & B/2 & D/2 \\ B/2 & C & E/2 \\ D/2 & E/2 & F \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}. \end{aligned} \tag{8}$$

3. Jacobian calculation

The error/distance metric is a 48-dimensional vector (18 from 9 points, 6 from 3 lines, and 24 from 2 ellipses), the




FIG. 4. The derivative of a point vector \mathbf{P} with respect to a direction \mathbf{n} is given by the cross-product of \mathbf{n} and \mathbf{P} .

zero value of which is computed using the Newton's method. To quickly find this global minimum is a practical challenge. Exact knowledge of the closed form Jacobian of this vector makes a significant improvement in speed and accuracy. The Jacobian is a matrix of size 48×6 , and of the form $\partial \delta_i / \partial \mathcal{P}_j$, where $i = 1 \dots 48, j = 1 \dots 6$; and \mathcal{P}_j representing the (six) pose parameters. We first compute the building blocks $\partial x_i / \partial \mathcal{P}_j$ and $\partial y_i / \partial \mathcal{P}_j$, for later use in the Jacobian computation. From Eqs. (1) and (2), we derive

$$\begin{aligned} \begin{bmatrix} x_i \\ y_i \end{bmatrix} &= \begin{bmatrix} -fX_i \\ s_x Z_i + o_x \\ -fY_i \\ s_x Z_i + o_y \end{bmatrix} \\ &= \begin{bmatrix} -\frac{f r_{11} X_i^m + r_{12} Y_i^m + r_{13} Z_i^m + T_1}{s_x r_{31} X_i^m + r_{32} Y_i^m + r_{33} Z_i^m + T_3} + o_x \\ -\frac{f r_{21} X_i^m + r_{22} Y_i^m + r_{23} Z_i^m + T_1}{s_y r_{31} X_i^m + r_{32} Y_i^m + r_{33} Z_i^m + T_3} + o_y \end{bmatrix}. \end{aligned} \tag{9}$$

From here, straightforward manipulations lead to the translational Jacobian,

$$\begin{aligned} \frac{\partial x_i}{\partial T_1} &= \frac{-f}{s_x Z_i}, & \frac{\partial x_i}{\partial T_2} &= 0, & \frac{\partial x_i}{\partial T_3} &= \frac{fX_i}{s_x Z_i^2}, \\ \frac{\partial y_i}{\partial T_1} &= 0, & \frac{\partial y_i}{\partial T_2} &= \frac{-f}{s_y Z_i}, & \frac{\partial y_i}{\partial T_3} &= \frac{fY_i}{s_y Z_i^2}. \end{aligned} \tag{10}$$

To obtain the rotational Jacobian, we first estimate $\partial \mathbf{P}_i^m / \partial \phi_j$, which we use to estimate $\partial \mathbf{P}_i / \partial \phi_j$, which will give us $\partial x_i / \partial \phi_j$ and $\partial y_i / \partial \phi_j$. Here \mathbf{P}_i is treated as a vector in space $\partial \mathbf{P}_i^m / \partial \phi_j$ can be obtained geometrically as a cross-product, as illustrated in Fig. 4,

$$\begin{aligned}\frac{\partial \mathbf{P}_i^m}{\partial \phi_1} &= [1, 0, 0]^T \times \mathbf{P}_i^m, \\ \frac{\partial \mathbf{P}_i^m}{\partial \phi_2} &= [0, 1, 0]^T \times \mathbf{P}_i^m, \\ \frac{\partial \mathbf{P}_i^m}{\partial \phi_3} &= [0, 0, 1]^T \times \mathbf{P}_i^m.\end{aligned}\quad (11)$$

$\partial \mathbf{P}_i^m / \partial \phi_j$ is defined in F_F . It needs to be rotated to change it to frame F_X ,

$$\frac{\partial \mathbf{P}_i}{\partial \phi_j} = \mathcal{R} \frac{\partial \mathbf{P}_i^m}{\partial \phi_j}. \quad (12)$$

From Eq. (9), $\partial x_i / \partial \phi_j$ and $\partial y_i / \partial \phi_j$ can be derived as

$$\frac{\partial x_i}{\partial \phi_j} = \frac{-f \left(Z_i \frac{\partial X_i}{\partial \phi_j} - X_i \frac{\partial Z_i}{\partial \phi_j} \right)}{s_x Z_i^2}, \quad \frac{\partial y_i}{\partial \phi_j} = \frac{-f \left(Z_i \frac{\partial Y_i}{\partial \phi_j} - Y_i \frac{\partial Z_i}{\partial \phi_j} \right)}{s_y Z_i^2}. \quad (13)$$

Using Eqs. (3)–(5), (8), (10), (12), and (13), we can derive the Jacobian for the point, line, and ellipse features as

$$\begin{aligned}\frac{\partial (\delta x_i)}{\partial \mathcal{P}_j} &= \frac{\partial x_i}{\partial \mathcal{P}_j}, \quad \frac{\partial (\delta y_i)}{\partial \mathcal{P}_j} = \frac{\partial y_i}{\partial \mathcal{P}_j}, \\ \frac{\partial \delta_j}{\partial \mathcal{P}_j} &= \frac{1}{\sqrt{A^2 + B^2}} \left(A \frac{\partial x_i}{\partial \mathcal{P}_j} + B \frac{\partial y_i}{\partial \mathcal{P}_j} \right), \\ \frac{\partial \delta_i}{\partial \mathcal{P}_j} &= K \left((2Ax_i + By_i + D) \frac{\partial x_i}{\partial \mathcal{P}_j} + (2Cy_i + Bx_i + E) \frac{\partial y_i}{\partial \mathcal{P}_j} \right).\end{aligned}\quad (14)$$

4. System of equations and update formula

Thus, for each feature distance, a first-order approximation can be written as

$$\sum_{j=1}^3 \left(\frac{\partial \delta_i}{\partial T_j} \Delta T_j + \frac{\partial \delta_i}{\partial \phi_j} \Delta \phi_j \right) = \delta_i, \quad (15)$$

which can be rewritten as

$$\mathbf{J} \Delta \mathcal{P} = \delta, \quad (16)$$

where \mathbf{J} is the Jacobian, \mathcal{P} are the six parameters, and δ is the multidimensional error function. The method proceeds by computing new estimates for \mathcal{R} and \mathcal{T} , and iterates until δ becomes acceptably small. The new estimates for \mathcal{P} are obtained by inverting Eq. (16). Given the current estimate \mathcal{R}_k and \mathcal{T}_k , and corrections $\Delta \mathcal{T}_k$ and $\Delta \phi_k$, the new estimates \mathcal{R}_{k+1} and \mathcal{T}_{k+1} can be computed as follows:

$$\begin{aligned}(1) \quad \begin{bmatrix} \mathcal{R}_{k+1} & \mathcal{T}_{k+1} \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} \mathcal{R}_k & \mathcal{T}_k \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \mathcal{R}_k & \Delta \mathcal{T}_k \\ 0 & 1 \end{bmatrix}; \\ (2) \quad \mathcal{T}_{k+1} &= \mathcal{T}_k + \Delta \mathcal{T}_k, \quad \phi_{k+1} = \phi_k + \Delta \phi_k;\end{aligned}$$

$$(3) \quad \mathcal{T}_{k+1} = \mathcal{T}_k + \Delta \mathcal{T}_k, \quad \mathcal{R}_{k+1} = \mathcal{R}_k \Delta \mathcal{R}_k.$$

The third update formula is most robust, while the first two are not stable, as they are extremely sensitive to numerical errors. It is also important that the order of $\Delta \phi_1$, $\Delta \phi_2$, and $\Delta \phi_3$ should not be changed while generating $\Delta \mathcal{R}_k$. Alternately, some other representation for the rotation matrix can also be used.

5. Least-square robustness

The fiducial rotation is measured in radian while translation is measured in mm, giving rise to numerical instabilities during the inversion of the Jacobian in Eq. (16). The inversion process can further worsen if the Jacobian is close to singularity. Besides, the assignment of different weights to different features may be desired, since within the region of interest some features will be statistically more robust than others. These problems are solved by using parameter and task variable scaling,^{39,40}

$$G(JH)(H^{-1} \Delta \mathcal{P}) = G \delta. \quad (17)$$

H scales the parameter space without changing the least-squares solution, since H can be chosen to make JH better conditioned than J . The H matrix can be constructed by using the parameter covariance matrix, for which we need an initial estimate of the solution and the standard deviation of each parameter. Column scaling⁴¹ is another popular method when no *a priori* information is available. Unfortunately, this can have a problem if the parameters are poorly identifiable⁴² in some poses. This is solved using the external scaling method, which does a column scaling on the most conservative estimates (the worst possible Jacobians). We ran 50 000 pose simulations to determine our H matrix using this method.

The G matrix is used for task variable (δ) scaling, which can be computed by the Gauss–Markov estimate by using a good estimate of the covariance matrix of the error function. In the absence of this matrix, another statistical approximation is to use a diagonal matrix, with each entry being the reciprocal of the standard deviation of that error element. This would make the L_2 norm a reasonable measure of the size of the least square error vector. We computed our G matrix using data from the same 50 000 pose simulations that we ran earlier.

6. Final optimization

The above framework was implemented using Newton's nonlinear optimization. We noticed that the optimization always converged in less than 25 iterations. The exit condition was defined as 50 iterations or a plateau in the error curve, whichever came later. We also incorporated some basic checks on the final answer to remove any mirror solutions. Other extraneous solutions can also act as a local minimum, in which case the initial estimate is perturbed a little to yield the correct convergence. A robust pose recovery algorithm was developed, by building on this optimization.

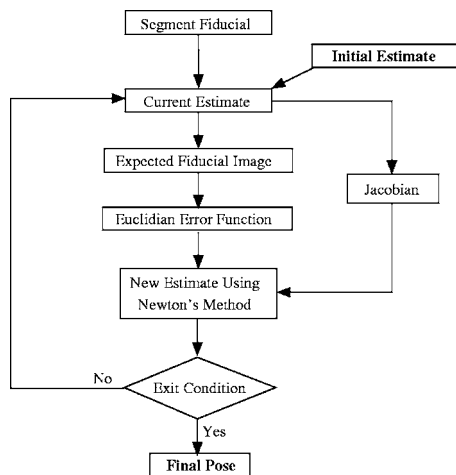


FIG. 5. The basic pose estimation algorithm.

7. Robust pose recovery algorithm

To make the pose recovery more robust, we have designed a two-phase pose recovery algorithm. The basic pose recovery algorithm is outlined in Fig. 5. Like all optimizations, this one is prone to fall into a local minima, from which it may not be able to recuperate. This is resolved by running the algorithm multiple times from various initial estimates for the pose. This approach is able to reduce the probability of failure exponentially. In a loose sense, if $p(\sim 0.3)$ is the probability of failure in any given iteration, then the probability of success after N independent runs is $1 - p^N$, which is sufficiently close to guaranteed success. Moreover, to make pose recovery more accurate, we use the observation that point features, though inaccurate, converge more reliably (see Table I). Thus, the basic algorithm runs 25 times using only the point features, from which we obtain a very good guess for the final pose. This is used as the initial estimate and now the basic algorithm runs another 25 times using all features. Though the number of iterations can be cut to one-fifth without compromising on the convergence, the extremely fast run time allows us to add a safety bracket by increasing the number of iterations. Figure 6 outlines the algorithm. This robust algorithm ran for tens of thousands of poses (spanning the complete workspace) without failure.

TABLE I. For the basic pose algorithm, different feature combinations give varying errors and percentage failures.

Points, lines and ellipse combinations	Translation (nm)	Rotation (°)	Failures (%)
6P+3L+2E	1.82	0.89	31.40
6P+2E	1.96	1.16	28.00
6P+3L	2.09	1.24	28.80
3L+2E	2.23	0.98	90.60
6P	2.48	1.53	27.00

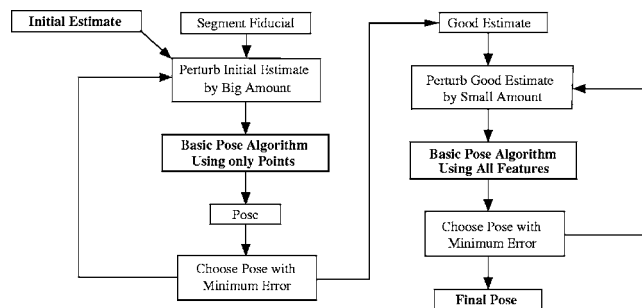


FIG. 6. The robust pose estimation algorithm.

8. Segmentation

Additional constraints on their coordinates make a segmentation of parametric curves (straight lines, ellipses, etc.) more accurate than points. Since automatic segmentation of parametric curves is not the research contribution of this work, we do not produce specific numerical results that would sidetrack the main theme of the paper. Instead we refer the readers to the extensive literature already available in the computer vision and image processing community, where ellipse segmentation is a field in itself.^{33,43-45} The parametric constraints can help segment the whole curve with high accuracy, even when most of the curve is occluded and only a small part is visible/segmented.⁴³ This makes pose estimation efficient. Moreover, the compact design of the fiducial decreases the likelihood of occlusion. We have developed and tested two methods of segmentation, both of which exhibit similar performances.

For many clinical applications, semiautomatic segmentation suffices. The operator clicks at a few points in the image that are close to the curves of interest. A minimum intensity centroid search in the neighborhood fine-tunes this location. Least square curve fitting is used to obtain the equation of the straight line and ellipse. Since the point features are known to be on the (extended) straight lines, their centroid positions are projected on the straight lines to get the final position. The semiautomatic MATLAB implementation takes about 10-15 s to process each image, which is acceptable for many applications.

Automatic segmentation of parametric curves is achieved by using the Hough transform,^{33,44} which searches in the parametric space. An edge image is created to search in the parametric space. Straight lines are segmented using the Hough transform, after which each line is fine-tuned by doing a search for minimum intensity points in a small neighborhood, followed by running a RANSAC-based least square fitting⁴³ on these points. The three strongest near-parallel straight lines are chosen. The ellipse is composed of five parameters, which renders the parametric search space huge. A decoupled approach is used where the Hough transform is used to estimate various possible axes of symmetry in the image. Selected intersecting orthogonal axes are possible candidates for an ellipse. A 2D Hough transform is used to find the best ellipse for each intersection pair, which is then fine-tuned. The two strongest ellipses are chosen. Our imple-

TABLE II. Pose recovery error (simulation) as a function of segmentation error.

Segmentation error (mm)	Translation (mm)			Rotation ($^{\circ}$)		
	X	Y	Z	X	Y	Z
0.00	0.18	0.19	0.75	0.02	0.02	0.01
0.50	0.19	0.18	0.79	0.05	0.06	0.04
1.00	0.19	0.19	0.79	0.09	0.11	0.07
2.00	0.22	0.22	0.98	0.18	0.22	0.14
3.00	0.26	0.26	1.24	0.26	0.32	0.21

mentation is a normalized and robust variation of Ref. 45. The points are segmented as end points of the lines. Alternatively, a curve following technique^{46,47} can also be used, though not implemented in our package.

III. RESULTS AND DISCUSSION

To measure sensitivity and robustness, the FTRAC fiducial was tested using simulated data and data that was obtained from real images of a precisely fabricated phantom. Finally, experiments measuring 3D fluoroscopic reconstruction of objects (*C*-arm tracked using the FTRAC fiducial) were conducted. The description and results for each experiment is followed immediately by the discussion.

A. Simulation studies

As a customary first step in the evaluation of fluoroscopic registration, simulations were conducted to evaluate the accuracy and robustness of the pose estimation software. Assuming an ideally calibrated fluoroscope, studies were conducted on synthetic images to analyse the effects of various governing parameters. MATLAB software was created to model x-ray imaging. Given *C*-arm parameters and fiducial pose, it generated a synthetic image and the exact location of each feature in the image. Appropriate noise was added to each imaging parameter, and its effect on pose recovery accuracy was observed. The error was modeled using a uniform and random probability density function, i.e., a 1 mm error means that a maximum error of *magnitude* 1 mm was added, with a uniform probability distribution. Based on our experience, the parameter values (for an average commercial *C* arm) were chosen to be 0.25 mm segmentation error, 1000 ± 2 mm focal length, 0.5 mm translation from the image center for the origin, and $0.5 \text{ mm} \pm 1 \mu\text{m}$ as the pixel size. The region of use was near the isocenter of the *C* arm, at about a distance of two-thirds the focal length from the x-ray source. With the above-mentioned error in the parameters, individual parameters were changed to understand the sensitivity of pose recovery. The simulations were carried out for randomly selected poses in the workspace. The initial guess was also chosen at random, far away from the actual pose.

1. Segmentation

Segmenting the projected fiducial from the x-ray images is one of the most important sources of error. The last step of our segmentation algorithm is to pick out multiple minimum

intensity points and fit a least square curve. Since in the synthetic data we know the precise position of each feature in the image, we added a segmentation error to the simulated point position. The erroneous points were then fit on a least square curve. This shifted the recovered position of the fiducial, which was compared to the known actual location.

Most of the translation error (Table II) was found to be along the imaging direction, which is intuitive since small changes in depth do not change the image significantly. The translation error parallel to the image plane is low, which is also intuitive since any small movement in that plane is magnified in the image. This can also be derived mathematically from Eq. (10), where the depth derivative has a quadratic term in the denominator, while the other derivatives are linear. Note that this could be an important source of error in depth sensitive surgeries, but in most cases is acceptable, especially when 3D information is projected into the image for guidance. The rotation error, on the other hand, is similar for all axes. The net translation and rotational error (based on 10 000 simulated poses) is displayed as a graph in Fig. 7(a). Most image processing algorithms achieve 1 pixel (0.4 mm) segmentation accuracy for ellipses, which indicates a 0.8 mm translation and 0.1° rotation error. This appears to be acceptable for most procedures and is similar to commercial image guidance systems that use optical/EM trackers. Moreover, the graph suggests that the pose recovery error does not grow exponentially with an increase in segmentation error. It can be concluded that the FTRAC fiducial is fairly robust to segmentation errors.

2. Focal length

The focal length, usually calibrated using an accurately machined phantom, could vary as much as a few mm when the *C* arm changes pose. Thus, an error in the focal length estimate is very likely, which could possibly lead to an inaccurate pose. Figure 7(b) (based on 20 000 simulated poses) shows the pose recovery error as a function of focal length calibration error. Orientation is always recovered with under 0.1° error, while translation error is nearly linear. Even a 5 mm error in focal length leads to under a 2 mm pose recovery error, with the largest component being in the imaging direction. We conclude that the orientation of the FTRAC fiducial is robust to focal length errors, while translation is a little sensitive.

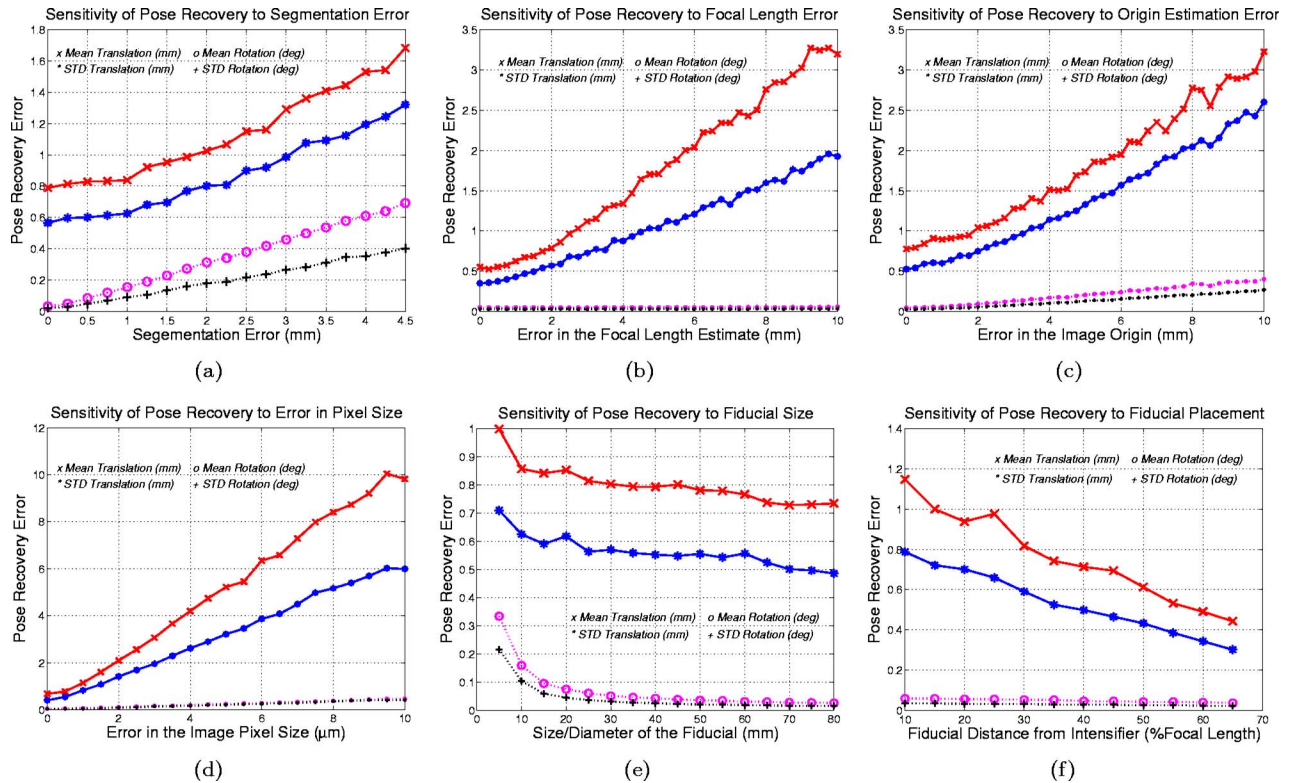


Fig. 7. Error in pose recovery as a function of (a) segmentation error; (b) focal length error; (c) origin error; (d) pixel size error; (e) fiducial size; (f) region of use.

3. Origin

The origin of the image, also calibrated using an accurately machined phantom, would change location with varying C -arm poses. Any error in the origin estimate could lead to a consistent shift in that particular direction. Figure 7(c) (based on 20 000 simulated poses) shows the rotational and translational sensitivity to the magnitude of error in origin estimation. Both translation and orientation errors are linear, though orientation is more robust. A 5 mm origin error leads to under a 2 mm translation error and 0.2° orientation error, with the largest component being parallel to the imaging direction. We conclude that the FTRAC fiducial is robust to origin estimation errors.

4. Pixel size

The pixel size along the x and y axes of the image does not change with fluoroscope movements and stays constant throughout the life of the fluoroscope. Hence a one-time calibration can provide very accurate estimates that can be used repeatedly, though any inaccuracies could have strong effects in fluoroscopic reconstruction. An error as small as $10\ \mu\text{m}$ could lead to a 5 mm shift near the boundary of the image. Figure 7(d) (based on 20 000 simulated poses) shows the sensitivity to this parameter. Both translation and orientation errors grow linearly, with most of the translation error being parallel to the image and orientation error less than 0.5° . We conclude that pose recovery orientation is robust to errors in pixel size, while translation is sensitive but linear.

5. Fiducial size

In the design of the FTRAC fiducial, the size of the features can influence the accuracy of the algorithm to some extent. A large fiducial tends to be more stable than a smaller one, and thus the size should be chosen depending on the work space constraints of the clinical application. We studied the pose recovery sensitivity by scaling the size (diameter and length), but without altering the relative configuration. Figure 7(e) (based on 8000 simulated poses) shows the translation and rotation errors as a function of size. It can be concluded that the curve practically plateaus at 30 mm, with a bigger fiducial exhibiting no improvement in accuracy. Moreover, a fiducial of $1 \times 1 \times 2$ cm size appears to be practical for most applications, with only a marginal drop in accuracy.

6. Region of interest

The region of use of the fiducial also has a strong effect on the accuracies obtained. A fiducial closer to the source has a larger magnification, thus scaling any small errors and making the algorithm converge better. This is evident also from Eq. (10), where the denominator in the Jacobian is the depth of the fiducial. Figure 7(f) (based on 6000 simulated poses) shows the pose recovery error as a function of the distance from the image plane, expressed as a percentage of the focal length. It can be concluded that orientation is independent of fiducial placement, while the translation error drops linearly in depth.

TABLE III. Pose recovery results (phantom) using the FTRAC fiducial phantom.

Trial number	Number of images	Translation (mm)			Rotation (°)		
		X	Y	Z	X	Y	Z
1	4	0.03	0.04	0.40	0.13	0.07	0.07
2	7	0.13	0.04	0.68	0.21	0.27	0.27
3	7	0.04	0.02	0.51	0.15	0.15	0.12
4	7	0.05	0.02	0.52	0.17	0.17	0.11
5	13	0.08	0.10	0.60	0.40	0.40	0.15
	Mean	0.07	0.04	0.55	0.21	0.21	0.15
	STD	0.05	0.03	0.32	0.12	0.14	0.09

7. Run time and convergence

The robust pose recovery algorithm always converged without failure. An unoptimized implementation using MATLAB 6.5.13, on a 2.4 GHz Intel P4 with 512 MB of RAM with a Windows 2000 OS, takes around 3.5 s to run (excluding image transfer and processing time). This indicates that an optimized C++ implementation would perform near real time.

B. Pose recovery studies on phantom

A highly precise phantom was manufactured and real x-ray images were taken using a fluoroscope (Philips BV 3000). The system-supplied parameters were read from the DICOM header, otherwise the fluoroscope was not explicitly calibrated. Moreover, the images were not distortion corrected (distortion ≤ 2 mm). The fiducial was mounted on a highly accurate 0.002° resolution rotational turntable (30000 Heavy Load Worm Gear Drive from Parker Automation, Irwin, Pennsylvania). The fluoroscope remains stationary while the fiducial moves in a known path, providing ground truth.

A fiducial mount was designed such that it produced zero translation and a known rotation when the turntable was rotated. The design supported the two independent rotation axes typical to C arms. Thus, given any starting pose, the relative motion between the current pose and the starting pose is known precisely from the turntable reading. The relative motion is also calculated using our algorithm (from computed current pose and the computed starting pose). The difference between the computed relative motion and known relative motion is the error. This setup was taken to the OR and five series of tests were carried out with different fiducial trajectories. The results are shown in Table III.

We recorded a mean accuracy of 0.56 mm for translation (STD 0.33 mm) and 0.33° for rotation (STD 0.21°), which appears to be acceptable for most surgeries. It can be argued that there may be a constant drift in the pose algorithm computation, which would not be captured by this experiment. All five trials have consistent results with the simulation experiments, indicating that a constant error drift is highly unlikely, which is further affirmed by the 3D reconstruction studies below.

C. Reconstruction studies on phantom

C-arm tracking is not an end in itself, but is a necessary link toward 3D object reconstruction using multiple x-ray images. The second set of phantom experiments involved taking x-ray images of a stationary object from different views, and reconstructing it in space. The FTRAC fiducial was kept stationary, while the C arm was moved and multiple images (20° around the AP axis) were taken. The 3D coordinates of the nine fiducial points were reconstructed using two to five images (the FTRAC fiducial tracked the C arm). The location of these points was known with respect to the fiducial by precise fabrication. The 3D reconstruction accuracy is the difference between the computed coordinates and the actual coordinates of the fiducial points. The points were reconstructed with a mean accuracy of 0.53 mm (STD 0.16 mm), which appears to be acceptable in most surgeries. A significant portion of the pose error is in the imaging direction, and hence in order to improve accuracy, the image separation should be increased, being orthogonal in the ideal case. Complete results are available in Table IV. Note that this reconstruction accuracy will also be a function of image warping and segmentation error. The FTRAC fiducial has also been successfully used to reconstruct phantom brachytherapy implants,⁴⁸ the performance of which is similar to that reported here.

D. Comparison to conventional fiducials

In Sec. I, we provided the reasons why there is a need for a small, yet accurate radio-opaque fiducial. In this section, we offer quantitative results indicating progress over previous work. When analyzing the basic pose recovery algorithm, simulation experiments from Sec. II show that though using only point features is most stable, it is also the least accurate (Table I). As reviewed in Sec. I, all previously proposed fiducials can be classified as bead-based fiducials (except one³ that we shall compare with in the end). Moreover, even among these fiducials, increasing the number of beads to more than 6 has little improvement in accuracy.¹⁹ For any given fiducial, various pose estimation algorithms can be designed, most of which fall into two categories—linear and nonlinear. Linear algorithms evaluate least-squares solutions to the pose problem, for which they typically linearize the

TABLE IV. Object reconstruction results using the FTRAC fiducial phantom.

	Number of images used for each reconstruction	MEAN	STD	MIN	MAX
TRIAL 1	2	0.55	0.17	0.20	1.18
	3	0.50	0.14	0.19	0.92
	4	0.48	0.14	0.19	0.75
	5	0.47	0.14	0.23	0.67
	2	0.66	0.25	0.27	1.70
TRIAL 2	3	0.55	0.16	0.29	1.13
	4	0.53	0.15	0.30	0.80
	5	0.52	0.14	0.33	0.73

rotation matrices. Though quick and easy to implement, the problem with this approach is that it neglects the rigid rotation constraints, leading to lower accuracies. Extensive results are available in the literature,^{20,23} where the accuracy of nonlinear methods is shown to be superior to linear methods. Thus, it should be sufficient to compare the performance of the FTRAC fiducial to a nonlinear pose estimation algorithm using six beads. To improve robustness of bead-based fiducials and to remove any biases that can arise due to a preferential configuration of the six beads, we compared the FTRAC fiducial to a nine-bead fiducial.

The predominant reason that the FTRAC fiducial offers better accuracies to other fiducials is that it relies on ellipses for evaluating the pose. Both for automatic and semiautomatic segmentation methods, achievable accuracies are dependent on multiple parameters like feature constraints, image resolution, x-ray wavelength, material used, choice of algorithm, etc. Thus, in the same image, different features provide different levels of accuracies. As reviewed in Sec. II, ellipses intrinsically offer better segmentation than beads. To quantitatively understand the improvement, we look at the residual errors from segmented datasets. The accuracy of each feature at the recovered pose is shown in Table V. It can be noticed that closed parametric curves (ellipses) segment with an overall accuracy of 0.13 mm, open parametric curves are a little more error prone with an accuracy of 0.19 mm, while points offer only 0.43 mm. The standard deviations

values are stable at under one-eighth of a pixel ($50 \mu\text{m}$). Thus, it seems that ellipse segmentation tends to be about three times more accurate than points.

We conducted simulations for over 7000 poses and compared the pose recovery results of a nine-bead fiducial to the FTRAC fiducial, keeping this disparity in mind. The results for translational and rotational pose recovery accuracies are available in Figs. 8(a) and 8(b), respectively. It can be noted that even with high amounts of segmentation error, the FTRAC fiducial exhibits only a marginal loss in accuracy. On the other hand, the errors for bead-based fiducials climbs steadily. This pattern is observed for both translation and rotation, where the error quickly climbs to the 1 mm and 1° mark for the bead-based fiducials, while it remains consistently low for the FTRAC fiducial. This indicates that the FTRAC fiducial is not only more accurate, but also more robust when compared to conventional fiducials. We do not explicitly compare the FTRAC fiducial to the helical fiducial,³ since the errors reported there are well over 1.5 mm in translation and 2° in rotation, while we obtain significantly superior results. Last, it should also be noted that the results obtained in Sec. III indicate a significant improvement to the previously published results (Sec. I).

IV. MOTIVATING APPLICATIONS

Though we have designed and tested the fiducial as a stand-alone piece of technology, we also intend to use it in a clinical application. The first use of this pose recovery technique will be in prostate brachytherapy, where C arms are ubiquitous, with over 60% of the practitioners using it in the operating room.⁴⁹ Unfortunately, it is used only for qualitative implant analysis and not for providing real-time intraoperative dosimetry. Thus, the ultimate goal will be to achieve registration between the seeds that are reconstructed from fluoroscopy and the soft tissue anatomy that is reconstructed from transrectal ultrasound (TRUS), which would allow us to make immediate provisions for any dose deviations from the intended plan. Several leading groups have published initial results favorably supporting C -arm fluoroscopy for intraoperative dosimetric analysis.⁵⁰⁻⁵² To achieve intraoperative dosimetry in prostate brachytherapy, we need

TABLE V. Residual distances for each feature from the five phantom trials. Parametric curves provide greater accuracies.

	Points	Lines	Ellipse
Residual error (mm)	0.44	0.14	0.15
	0.49	0.14	0.18
	0.43	0.28	0.11
	0.38	0.21	0.10
	0.42	0.19	0.11
MEAN	0.43	0.19	0.13
STD	0.04	0.05	0.03

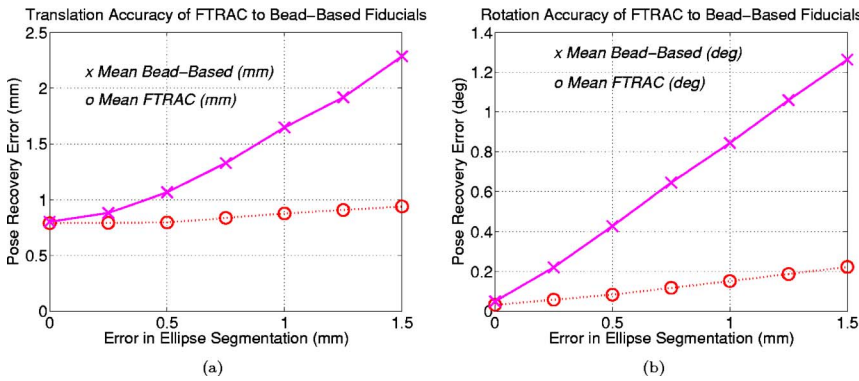


FIG. 8. A comparison of pose recovery accuracies for a nine-bead conventional fiducial compared to the FTRAC fiducial. Errors in (a) translation; (b) rotation. The FTRAC fiducial performs significantly superior.

to address the issues of (a) C-arm distortion correction and calibration; (b) C-arm pose tracking; (c) seed matching and reconstruction; and (d) registration of C-arm to TRUS images. The FTRAC fiducial is capable of addressing the issues of pose tracking and registration.

In a quick survey of registration, thin metal wire⁵³ inside a Foley catheter was used to visualize the prostatic urethra fluoroscopically in anterior-posterior and lateral projections. Recently, gold marker seeds⁵⁴ were implanted into the prostate and the relative positions of the needles and marker seeds were observed in fluoroscopy. Inserted implant needles⁵² were used as fiducials for the registration of the fluoroscopy and ultrasound spaces. Unfortunately, implanted markers of any sort are susceptible to motion, and thus are not reliable. Alternately, radio-opaque spherical beads have been rigidly attached to the TRUS probe,¹⁹ not only registering the TRUS to the x-ray images, but also addressing tracking. This process appears to require a permanent alteration to the probe, which for some practitioners might not be desirable.

Since in commercial brachytherapy systems the template is already registered to TRUS, the FTRAC fiducial will be attached to the implant template with a precision-machined mechanical connector in a known position. The spatial relationship between the FTRAC and TRUS is explained in Fig. 9. Let P_F be a point reconstructed in FTRAC space, P_U be the same point in ultrasound space, F_{FT} be the transformation between FTRAC and template, and F_{TU} be the transformation between template and ultrasound. Then the relationship between P_U and P_F can be expressed as $P_U = F_{TU} F_{FT} P_F$. Let us examine how accurate P_U is expected to be the accuracy of F_{FT} is defined by the accuracy of CNC machining and is in the neighborhood of 0.25 mm. The accuracy of P_F was found to be about 0.5 mm in our experiments. The accuracy of F_{TU} is specific to the FDA-approved brachytherapy system (CMS Interplant) that we will be using, and it is about 1.0 mm. Thus, the overall accuracy of the whole system is expected to be under 2 mm and can be further reduced by an improvement in the template-to-TRUS registration method. Experiments to evaluate the accuracy of the complete system are currently underway.

The fiducial is of general use, but has been optimized for prostate brachytherapy. The ellipses (and other features) are inclined at an angle of 30° to the vertical. Even when the

fiducial is kept close to the prostate and at varying heights, a 30° plane (to the vertical) divides the fiducial from the prostate (which contains the seeds). If the C-arm images are taken in any position inside the 30° cone around the AP axis, this design reduces the likelihood of overlap between the seed projections and the fiducial. As we have mentioned earlier, we will mount the FTRAC fiducial over the abdomen on the TRUS stepper, in a known calibrated pose relative to the template with the use of a precision-machined extender [Fig. 10(a)]. The distance between the FTRAC fiducial and the prostate will be about 15 cm, in which range the pose recovery and object reconstruction stabilities are adequate. The fallback plan,⁵⁵ as shown in Fig. 10(b), is to place the fiducial on a rectal sheath that supports the prostate, while the TRUS probe can move unhindered inside the sheath. The technical difficulty in this case was a robust coupling between the TRUS probe and the sheath, at acceptable loss in image quality. However, an increase in insertion diameter is of concern. A third alternative is to put the patient on a foam board incorporating one or several FTRAC fiducials in known relative offset with respect to the TRUS stepper.

Apart from brachytherapy, another promising application is fluoroguided robot-assisted surgery, as done in the literature.⁵ If mounted on the robot, the FTRAC fiducial can register the robot to the C-arm space, achieving virtual fluoroscopy. In general, surgical tool tracking under x-ray fluo-




FIG. 9. Frame transformations between the FTRAC and TRUS.




FIG. 10. The FTRAC fiducial mounted on (a) the template; (b) a hollow cylindrical sheath.

roscopy imaging can be achieved if the *C* arm is kept stationary, while the FTRAC fiducial is rigidly mounted on a surgical tool. Moreover, segmentation of the fiducial image and pose recovery have an inherent duality (i.e., knowing one directly relates to the other), which can further boost accuracy. The final pose can be used to improve segmentation, which, in turn, can enhance the pose. Iteratively doing this will not only improve accuracy, but also remove the need for an accurate segmentation. Alternately, an algorithm that does not optimize on the geometric error of the fiducial features, but optimizes directly on the intensity information, can be used. Both these techniques will eliminate the need for segmentation, while still achieving accurate pose recovery. The FTRAC fiducial can potentially also be used for both preoperative and intraoperative *C*-arm calibration. Additionally, the FTRAC fiducial can be instrumental in registering x-ray images to other imaging modalities. Thus, the FTRAC fiducial can facilitate a plurality of fluoroscopically guided procedures.

V. CONCLUSION AND OPEN QUESTIONS

A fiducial for robust and accurate *C*-arm tracking was designed, prototyped, computationally evaluated, and experimentally validated. The primary contribution of this work was the use of parametric curves (ellipses and straight lines), in addition to spherical beads that have been used conventionally. Point-based algorithms, though easy to implement, are sensitive to errors. On the other hand, parametric curves (a) segment accurately; (b) constrain pose recovery better; and (c) move pose recovery outside the framework of purely point-based methods. The proposed fiducial also addresses the issue of registration to other imaging modalities by precise placement.

Without distortion correction and extensive calibration, the FTRAC fiducial can track a *C* arm with a mean accuracy of 0.56 mm in translation (STD 0.33 mm) and 0.33° in rotation (STD 0.21°) and can reconstruct 3D coordinates to a mean accuracy of 0.53 mm (STD 0.16 mm). It not only promises to offer superior accuracies when compared to conventional fiducials, but also is significantly smaller in size ($3 \times 3 \times 5$ cm). The simulations suggest that the size of the FTRAC fiducial can be further reduced to $1 \times 1 \times 2$ cm, maintaining a similar level of accuracy.

One of the open questions in quantitative fluoroscopy is the importance of distortion correction and *C*-arm calibra-

tion. Generally speaking, distortion is widely prevalent and necessitates a correction procedure. This assures accurate tracking and reconstruction. Since properties vary from one machine to another, the distortion on some fluoroscopes might be small, while on others it might be large. Moreover, some applications might be more robust to distortion effects than others. Depending on the application and the amount of distortion, researchers might choose to (a) correct for a unique distortion at every pose, (b) use a constant correction for all poses, or (c) completely neglect distortion. This decision has to be made on a case by case basis. The results in this paper are on uncorrected images. The reason for robustness of our results to *C*-arm distortion can be attributed to the low amounts of distortion in the fluoroscope (average < 2 mm) and to the small physical size of the fiducial, which covers only a small area in the image. Since distortion correction is not the focus of this paper, we do not analyze its effects on reconstruction. The results and analysis are available elsewhere.⁴⁸

A similar issue arises for a calibration of the imaging parameters. Though it is believed that accurate calibration is a necessary prerequisite for quantitative fluoroscopy, our results indicate that it might not always be the case. Depending on the *C* arm used and the variation in the imaging parameters, inaccurate imaging parameters might still provide an accurate reconstruction of 3D information. These ideas can have desirable repercussions in a variety of clinical procedures, meriting further research. Though the FTRAC fiducial directly does not address these issues, it has provided the necessary intuition. We hope that these ideas, in collaboration with the FTRAC fiducial, will boost the clinical relevance and applicability of quantitative fluoroscopy on ordinary *C* arms.

ACKNOWLEDGMENTS

This work has been financially supported by NIH 1R43CA099374-01 and NSF EEC-9731478. We are grateful for the thoughtful advice of Professor Russell Taylor, and for Stephen Kubinak, John Sofranko, Jane Kwietkowski, and Scott Borzillary for help with the experiments.

^aElectronic mail: {jain.gabor}@cs.jhu.edu

¹R. Hofstetter, M. Slomczykowski, M. Sati, and L. Nolte, "Fluoroscopy as an imaging means for computer-assisted surgical navigation," *Comput. Aided Surg.* **4**, 65–76 (1999).

²L. Joskowicz, C. Milgrom, A. Simkin, L. Tockus, and Z. Yaniv, "Fracas: a system for computer-aided image-guided long bone fracture surgery," *Comput. Aided Surg.* **3**, 271–288 (1998).

³J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazanzides, and A. Guezic, "A *c*-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," *Comput. Aided Surg.* **5**, 373–390 (2000).

⁴J. C. Eyke, J. E. Ricciardi, W. Roesch, and T. S. Whitecloud, "Computer-assisted virtual fluoroscopy," *University of Pennsylvania Orthopaedic Journal*, 2002, pp. 53–59.

⁵K. Wiesent, K. Barth, N. Navab, P. Durlak, T. Brunner, O. Schuetz, and W. Seissler, "Enhanced 3-d-reconstruction algorithm for *c*-arm systems suitable for interventional procedures," *IEEE Trans. Med. Imaging* **19**, 412–422 (2000).

⁶S. A. M. Baert, W. J. Niessen, and M. A. Viergever, "Guide wire tracking during endovascular interventions," *IEEE Trans. Med. Imaging* **22**, 965–972 (2003).

- ⁷T. Yamazaki, T. Watanabe, Y. Nakajima, K. Sugamoto, T. Tomita, H. Yoshikawa, and S. Tamura, "Improvement of depth position in 2-d/3-d registration of knee implants using single-plane fluoroscopy," *IEEE Trans. Med. Imaging* **23**, 602–612 (2004).
- ⁸T. Rohlfing, D. Russako, M. Murphy, and C. Maurer, "An intensity-based registration algorithm for probabilistic images and its application for 2-d to 3-d image registration," in *SPIE Medical Imaging: Image Processing*, 2002, Vol. 4684, pp. 581–591.
- ⁹T. Tang, "Calibration and point-based registration of fluoroscope images," Master's thesis, Queen's University, January 1999.
- ¹⁰R. Siddon and N. Barth, "Stereotaxic localization of intracranial targets," *Int. J. Radiat. Oncol., Biol., Phys.* **13**, 1241–1246, August 1987.
- ¹¹R. Koppe, E. Klotz, J. O. d. Beek, and H. Aerts, "3d vessel reconstruction based on rotational angiography," in *CAR*, 1995, pp. 101–107.
- ¹²N. Navab, A. Bani-Hashemi, M. Mitschke, D. W. Holdsworth, R. Fahrig, A. J. Fox, and R. Graumann, "Dynamic geometrical calibration for 3-d cerebral angiography," in *SPIE Medical Imaging*, February 1996, Vol. 2708, pp. 361–370.
- ¹³C. Brack, M. Roth, A. Czopf, J. Moctezuma, F. Gtte, and A. Schweikard, "Towards accurate x-ray camera calibration in computer assisted robotic surgery," in *Computer-Assisted Radiology and Surgery (CARS)*, Berlin, 1996, pp. 721–728.
- ¹⁴M. Mitschke and N. Navab, "Recovering projection geometry: How a cheap camera can outperform an expensive stereo system," in *CVPR*, 2001, Vol. 1, pp. 193–200. Available online at citeseer.ist.psu.edu/456732.html
- ¹⁵VectorVision® Navigation System, "Brainlab, Inc., Heimstetten, Germany."
- ¹⁶StealtStation®, "Medtronic Surgical Navigation Technologies, Louisville, CO."
- ¹⁷OEC 9800 FluoroTrak™, "GE Healthcare, Waukesha, WI."
- ¹⁸R. Phillips, A. Mohsen, W. Viant, S. Malek, Q. Li, N. Shah, M. Bielby, and K. Sherman, "A phantom based approach to fluoroscopic navigation for orthopaedic surgery," in *MICCAI (2)*, 2004, pp. 621–628.
- ¹⁹M. Zhang, M. Zaider, M. Worman, and G. Cohen, "On the question of 3d seed reconstruction in prostate brachytherapy: the determination of x-ray source and film locations," *Phys. Med. Biol.* **49**, N335–345 (2004).
- ²⁰Z. Yaniv and L. Joscowicz, "Precise robot-assisted guide positioning for distal locking of intramedullary nails," *IEEE Trans. Med. Imaging* **24**, 624–635 (2005).
- ²¹A. K. Jain, T. Mustafa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (FTRAC) fiducial," in *SPIE Medical Imaging: Visualization, Image-Guided Procedures, and Display*, February 2005, Vol. 5744, pp. 798–809.
- ²²T. S. Y. Tang, N. J. MacIntyre, H. S. Gill, R. A. Fellows, N. A. Hill, D. R. Wilson, and R. E. Ellis, "Accurate assessment of patellar tracking using fiducial and intensity-based fluoroscopic techniques," *Med. Image Anal.* **8**, 343–351 (2004).
- ²³D. DeMenthon and L. Davis, "Model-based object pose in 25 lines of code," *Int. J. Comput. Vis.* **15**, 123–141 (1995).
- ²⁴J. Fitzpatrick and J. West, "The distribution of target registration error in rigid-body point-based registration," *IEEE Trans. Med. Imaging* **20**, 917–927 (2001).
- ²⁵J. Heikkila and O. Silven, "A four-step camera calibration procedure with implicit image correction," in *CVPR97*, 1997, pp. 1106–1112.
- ²⁶A. K. Jain, T. Mustafa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (FTRAC) fiducial," in *AAPM Annual Meeting*, Pittsburgh, July 2004.
- ²⁷O. D. Faugeras and B. Mourrain, "On the geometry and algebra of the point and line correspondences between n images," in *ICCV*, 1995, pp. 951–956. [Online]. Available: citeseer.ist.psu.edu/faugeras95geometry.html
- ²⁸A. Ansar and K. Daniilidis, "Linear pose estimation from points or lines," in *ECCV02*, 2002, Sec. IV, pp. 282ff.
- ²⁹Q. Ji, M. Costa, R. Haralick, and L. Shapiro, "An integrated linear technique for pose estimation from different features," *International Journal of Pattern Recognition and Artificial Intelligence*, June 1999. Available online at citeseer.ist.psu.edu/ji99integrated.html
- ³⁰C. Taylor and D. Kriegman, "Structure and motion from line segments in multiple images," *T-PAMI*, 1995, Vol. 17, pp. 1021–1032.
- ³¹R. Hu and Q. Ji, "Camera self-calibration from ellipse correspondences," in *ICRA*, 2001, pp. 2191–2196.
- ³²L. Wang, C. Grimm, and R. Pless, "A 3d pattern for post estimation for object capture," Available online at citeseer.ist.psu.edu/606684.html
- ³³E. Trucco and A. Verri, *Introductory Techniques for 3-D Computer Vision* (Prentice-Hall, Englewood Cliffs, NJ, 1998), pp. 283–294.
- ³⁴P. Yao, G. Evans, and A. Calway, "Face tracking and pose estimation using affine motion parameters," in *Proceedings of the 12th Scandinavian Conference on Image Analysis*, I. Austvoll, Ed. Norwegian Society for Image Processing and Pattern Recognition, 2001, pp. 531–536. [Online]. Available: citeseer.ist.psu.edu/yao01face.html
- ³⁵M. Dhome, J. La Preste, G. Rives, and M. Richetin, "Spatial localization of modelled objects of revolution in monocular perspective vision," in *ECCV90*, 1990, pp. 475–485.
- ³⁶S. Ma, "Conics-based stereo, motion estimation and pose determination," *IJCV*, February 1993, Vol. 10, No. 1, pp. 7–25.
- ³⁷M. Costa and L. Shapiro, "3d object recognition and pose with relational indexing," *CVIU*, September 2000, Vol. 79, No. 3, pp. 364–407.
- ³⁸D. Eberly, "Distance from a point to an ellipse in 2D," Oct 2002, www.magic-software.com, Magic Software, Inc. Available Online at www.magic-software.com
- ³⁹J. Hollerbach and C. Wampler, "The calibration index and taxonomy for robot kinematic calibration methods," *IJRR*, 1996, Vol. 15(6), pp. 573–591, 1996.
- ⁴⁰C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Prentice-Hall, Englewood Cliffs, NJ, 1974).
- ⁴¹R. Prager, R. Rohling, A. Gee, and L. Berman, "Automatic calibration for 3-d free-hand ultrasound," Cambridge University Department of Engineering, Tech. Rep. CUED/F-INFENG/TR 303, September 1997.
- ⁴²K. Schroer, *Theory of Kinematic Modeling and Numerical Procedures For Robot Calibration* (Chapman & Hall, London, 1993), pp. 157–196.
- ⁴³R. Halir and J. Flusser, "Numerically stable direct least squares fitting of ellipses," available online at citeseer.ist.psu.edu/350661.html
- ⁴⁴R. McLaughlin, "Randomized Hough transform: Improved ellipse detection with comparison," *Pattern Recogn. Lett.* **19**, 299–305 (1998). Available online at citeseer.ist.psu.edu/mclaughlin98randomized.html
- ⁴⁵Y. Lei and K. C. Wong, "Ellipse detection based on symmetry," *Pattern Recogn. Lett.* **21**, 41–47 (1999).
- ⁴⁶C.-L. T. Tsai, C. V. Stewart, H. L. Tanenbaum, and B. Roysam, "Model-based method for improving the accuracy and repeatability of estimating vascular bifurcations and crossovers from retinal fundus images," *IEEE Trans. Inf. Technol. Biomed.* **8**, 122–130 (2004).
- ⁴⁷A. Can, H. Shen, J. N. Turner, H. L. Tanenbaum, and B. Roysam, "Automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms," *IEEE Trans. Inf. Technol. Biomed.* **3**, 125–138 (1999).
- ⁴⁸A. K. Jain, Y. Zhou, T. Mustafa, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)," in *SPIE Medical Imaging: Visualization, Image-Guided Procedures, and Display*, February 2005, Vol. 5744, pp. 810–821.
- ⁴⁹B. Prestidge, J. Prete, T. Buchholz, J. Friedland, R. Stock, P. Grimm, and W. Bice, "A survey of current clinical practice of permanent prostate brachytherapy in the United States," *Int. J. Radiat. Oncol., Biol., Phys.* **15**, 461–465 (1998).
- ⁵⁰D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," *Phys. Med. Biol.* **47**, 2031–2048 (2002).
- ⁵¹D. Todor, M. Zaider, G. Cohen, M. Worman, and M. Zelefsky, "Intraoperative dynamic dosimetry for prostate implants," *Phys. Med. Biol.* **48**, 1153–1171 (2003).
- ⁵²L. Gong, P. Cho, B. Han, K. Wallner, S. Sutlief, S. Pathak, D. Haynor, and Y. Kim, "Ultrasonography and fluoroscopic fusion for prostate brachytherapy dosimetry," *Int. J. Radiat. Oncol., Biol., Phys.* **1**, 1322–1330 (2002).
- ⁵³K. Wallner, J. Roy, M. Zelefsky, Z. Fuks, and L. Harrison, "Fluoroscopic visualization of the prostatic urethra to guide transperineal prostate implantation," *Int. J. Radiat. Oncol., Biol., Phys.* **1**, 863–867 (1994).
- ⁵⁴M. Baird, R. Holt, and T. Selby, "Improvement of transperineal implant dosimetry by intraoperative cystoscopic confirmation of prostate anatomy," *J. Urol. (Paris)* **164**, 406–410 (2000).
- ⁵⁵J. Kemper, A. Burkholder, A. Jain, T. Mustafa, K. Wyrobek, C. Burdette, D. Song, A. Okamura, and G. Fichtinger, "Transrectal fiducial carrier for radiographic image registration in prostate brachytherapy," in *AAPM Annual Meeting*, Seattle, July 2005.