

Fiducial Registration from a Single X-Ray Image: A New Technique for Fluoroscopic Guidance and Radiotherapy

T.S.Y. Tang

R.E. Ellis

G. Fichtinger

Computing and Information Science, Queen's University at Kingston, Canada
CISST/ERC Johns Hopkins University, Baltimore, Maryland, USA
contact: ellis@cs.queensu.ca

Abstract. Fiducial registration is useful both in applications where other registration techniques have poor performance and for validation of new registration techniques. Registration of 3D CT or MR images to 2D X-ray images is particularly difficult, in part because automated contour extraction from 2D images is not yet well solved and in part because of the considerable computational expense in matching the contours to the 3D images.

This work addresses the problem of fiducial registration from a single X-ray image. We have developed an algorithm for fast, efficient registration of 3D fiducial locations to the lines cast from the X-ray source to the 2D projective image that is 60 times faster than the popular iterated closest-point algorithm. The algorithm has been tested on fluoroscopic images from portable C-arms and on portal images from a radiotherapy device. On these images, six or seven fiducials can be registered within seconds to an absolute accuracy of about one millimeter and two degrees.

1 Introduction

Image-guided minimally invasive therapy and conformal radiotherapy can benefit from registration of a patient or tool to a preoperative image. In the case of image-guided therapy, registration of a rigid (or near-rigid) tissue to a preoperative image permits intraoperative guidance; for example, in many fluoro-guided percutaneous procedures the soft target volume is determined in preoperative CT or MR scans but the target is not visible intraoperatively. In the case of conformal radiotherapy and fractionated radiosurgery, it is essential that the patient is registered on the couch of the linear accelerator in the position and orientation that was planned in pre-operative CT and/or MRI images.

These functions can be performed by registering multiple projected X-ray images, e.g., from several positions or with several imaging systems. However, multi-image registration not only increases the ionizing radiation dose to the patient but also requires additional time (with the risk of patient motion occurring between the imaging instants) or additional resources in the form of additional imaging systems (which increases the cost). The ideal registration technique would use an automatic registration of contours or intensities from the intraoperative image. However, automatic procedures are not yet sufficiently fast and robust for reliable therapeutic use so an alternative technique is worthy of examination.

This work addresses the problem of how to estimate the pose of a known object with a single 2D image using fiducial markers. The location of the fiducial markers in the anatomy is known from a preoperative 3D image, and the specific problem is to estimate the pose of the markers in the coordinates of the 2D imaging system. Our pose estimation process has two steps: the calibration of the image and imaging device, and the reconstruction of the pose from the image and the known spatial distribution of the fiducial markers. We will describe an accurate algorithm for estimating the pose by minimizing the distance between the set of 3D fiducial points and the lines that project from the X-ray source to the 2D image sensor. By taking advantage of the geometric constraints inherent in the problem we have developed an algorithm that is more than 60 times faster than the iterative-closest-point algorithm and that can achieve submillimeter registration accuracy from a single fluoroscopic image.

2 Background

Single-image registration requires a calibrated image and a registration algorithm. Here, we briefly review the main issues and previous work on these issues.

2.1 Image Calibration

We are interested in using a single-X-ray image for registration, concentrating on C-arm fluoroscopy and portal imaging. For portal images only minimal calibration is needed, as the images are practically free of distortion and the machine's internal parameters (such as the source-to-screen distance and pixel size) are readily determined from routine procedures. But it is critical that a C-arm fluoroscope be calibrated properly, as there is a substantial amount of distortion in fluoroscopic images and substantial mechanical deformation of the C-arm beam. Various methods for determining global [10, 12] and local [3, 9] spatial transformations for "unwarping" image distortion have been devised, typically using a grid of lines or points with known spatial coordinates. The deformation for mechanical beams, which is dependent on the C-arm orientation, can be determined beforehand [13]. For finding the internal parameters of the imaging device, X-ray calibration methods [6, 11] and camera calibration methods [5, 14] can be applied.

2.2 Point-Based Image Registration

Point-based 3D/3D registration of a pair of corresponding point sets, also known as the absolute orientation problem, has been well addressed. Arun *et al.* [1] provided a least-squares solution based on singular value decomposition, and Horn [7, 8] gave closed-form solutions using quaternions and orthonormal matrices respectively.

If the correspondence between points is not known, or if one of the sets is not a point set, the iterative closest point algorithm (ICP) of Besl and McKay [2] can be used. ICP is a general algorithm for registering a point set to a geometric model, which may be a set of points, lines, or surfaces. Typically the model is constructed preoperatively and the data are gathered intraoperatively. ICP can also be accelerated [2] to speed the iteration in regions where the error gradient is smooth.




Fig. 1. Imaging geometry.

By reversing the temporal sequence, i.e., by gathering the points preoperatively and constructing a line set model from the intraoperative image, one can use ICP to solve the registration problem. From the X-ray image of the object, which is a 2D point set, a set of *back-projected lines* can be constructed from the 2D image points to the 3D X-ray source location (which is known from the device calibration). ICP can then be applied between this set of 3D lines and the 3D point set of the object.

However, as we will show below, accelerated ICP needs thousands of iterations to converge and often yields incorrect answers as it converges to a local non-global minimum. A main reason for this is that ICP may start with a wrong correspondence as a point could correspond to two or more lines. Even if ICP is modified to force a unique correspondence, results are not satisfactory because of long computation time and high failure rate when the number of points are small. (ICP takes 172 CPU seconds on a SUN Ultra 60 to register just five points, with a failure rate of more than 16%.)

3 Computational Methods

The main weakness of ICP in this application is the incorrect correspondence that is found from the various initial pose estimates. The initial estimate is very important in this application, because a good initial estimate should lead to a correct correspondence and thus the correct pose estimate. The algorithm described here is in two parts. First the algorithm quickly finds a set of initial estimates of the pose, each of which should be very close to the correct estimate; then, for each initial estimate, the algorithm refines the estimate by using iterative gradient descent to minimize the total least-squares distances from the points to the corresponding model lines.

Formally, the registration problem is defined here as

Given:

- An object O that is a set of 3D points $\mathcal{P}^B = \{p_1^B, \dots, p_n^B\}$ in frame B ;
- A set of n coplanar image points $\mathcal{J}^C = \{j_1^C, \dots, j_n^C\}$ in frame C ;
- An X-ray source location s^C ;
- A distance metric $\rho(\cdot, \cdot)$ that measures the distance between a point z and the closest point on line $L(t) = a + tb$; and

– An objective function $F(\cdot)$ of a set of values;

Construct:

the set of back-projected lines from the image points to the X-ray source $\mathcal{L}(t)^C$, where $\mathbf{L}_i(t) = \mathbf{j}_i^C + t(\mathbf{s}^C - \mathbf{j}_i^C)$, then

Find:

The rotation R_B^C and translation t_B^C that minimizes $F(\rho)$, where

$$\rho_i = \rho(R_B^C \mathbf{p}_i^B + \mathbf{t}_B^C, \mathbf{j}_i^C + t(\mathbf{s}^C - \mathbf{j}_i^C))$$

is minimized.

The distance metric used here is the squared Euclidean distance between a point \mathbf{z} and a line $\mathbf{L}(t) = \mathbf{a} + t\mathbf{b}$, and $F(\rho)$, so finding the registration is a least-squares minimization problem.

3.1 Finding Initial Estimates

We have previously established that the algorithm for finding initial estimates that Besl and McKay originally proposed [2], which was based on symmetries in $SO(3)$ orientations, is inadequate [13].

To find good initial estimates, consider the problem of registering three points to three lines. Suppose there is no symmetry in \mathcal{P} , i.e., that by accident or design there is no rotation and/or reflection of \mathcal{P} similar to the original. Pick three points in \mathcal{J}^C that are not collinear and use these three back-projected lines to try to register with all non-collinear 3-permutations of the n points in \mathcal{P}^B . The idea is that, for each permutation, we generate the (two) initial estimates by finding the orientation and position of the triangle formed by the three points, say \mathbf{p}_i^B for $i = 1 \dots 3$. For brevity, denote $\mathbf{L}_{j_i^C}(t)$ be the line formed by \mathbf{j}_i^C and \mathbf{s}^C , and let $\mathbf{a}_i = \mathbf{L}_{j_i^C}(t_i)$ and $\mathbf{a}'_i = \mathbf{L}_{j_i^C}(t'_i)$ for some t_i and t'_i .

Let the triangle with vertices \mathbf{a}_i be denoted as \mathcal{A} . There are two possible registrations, so there are two such triangles \mathcal{A} . The problem is to find some $\mathbf{a}_i, i = 1 \dots 3$, such that the triangle \mathcal{A} is congruent to the triangle \mathcal{P} formed by \mathbf{p}_i . To find the orientation of the triangle, observe that if a point \mathbf{a}'_1 is arbitrarily chosen then there is a triangle \mathcal{A}' that contains \mathbf{a}'_1 and that is a scaled version of the registration \mathcal{A} . The triangle \mathcal{A}' is therefore similar to the triangle \mathcal{P} , so

$$\frac{\|\mathbf{p}_1 - \mathbf{p}_2\|}{\|\mathbf{a}'_1 - \mathbf{a}'_2\|} = \frac{\|\mathbf{p}_1 - \mathbf{p}_3\|}{\|\mathbf{a}'_1 - \mathbf{a}'_3\|} = \frac{\|\mathbf{p}_2 - \mathbf{p}_3\|}{\|\mathbf{a}'_2 - \mathbf{a}'_3\|} \tag{1}$$

Two constants, r and s can be defined as

$$r = \frac{\|\mathbf{p}_1 - \mathbf{p}_2\|}{\|\mathbf{p}_1 - \mathbf{p}_3\|} = \frac{\|\mathbf{a}'_1 - \mathbf{a}'_2\|}{\|\mathbf{a}'_1 - \mathbf{a}'_3\|} \quad s = \frac{\|\mathbf{p}_1 - \mathbf{p}_2\|}{\|\mathbf{p}_2 - \mathbf{p}_3\|} = \frac{\|\mathbf{a}'_1 - \mathbf{a}'_2\|}{\|\mathbf{a}'_2 - \mathbf{a}'_3\|} \tag{2}$$

Rewriting Equations (1) with the terms from Equation (2) gives two equations of constraint for the points \mathbf{a}'_2 and \mathbf{a}'_3 :

$$\begin{aligned} f(t'_2, t'_3) &= r^2 \|\mathbf{a}'_1 - \mathbf{a}'_3\|^2 - \|\mathbf{a}'_1 - \mathbf{a}'_2\|^2 = 0 \\ g(t'_2, t'_3) &= s^2 \|\mathbf{a}'_2 - \mathbf{a}'_3\|^2 - \|\mathbf{a}'_1 - \mathbf{a}'_2\|^2 = 0 \end{aligned} \tag{3}$$

Partial derivatives of these two non-linear functions can be found, so standard optimization methods for solving two equations with two unknowns can be applied. In this work the Levenberg-Marquardt algorithm was used. Because $L_{j_i^C}(t)$ is defined such that $t = 0$ is at the source, and $t = 1$ is on the image plane, all t_i are between 0 and 1. If t'_1 is fixed at 0.5, by using all combinations of 0 and 1 as initial values for t'_2 and t'_3 two solutions were always found.

The registration \mathcal{A} can be determined by scaling the solution \mathcal{A}' so that \mathcal{A} is congruent to \mathcal{P} .

Two points, say \mathbf{a}'_u and \mathbf{a}'_v , and the source \mathbf{s} form a triangle on the plane containing the two lines $L_{j_u^C}$ and $L_{j_v^C}$. We want to find the triangle $\triangle s\mathbf{a}_u\mathbf{a}_v$ that is similar to $\triangle s\mathbf{a}'_u\mathbf{a}'_v$ and $\|\mathbf{a}_u - \mathbf{a}_v\| = \|\mathbf{p}_u - \mathbf{p}_v\|$. The two similar triangles yield the constraint

$$\frac{\|\mathbf{s} - \mathbf{a}_u\|}{\|\mathbf{s} - \mathbf{a}'_u\|} = \frac{\|\mathbf{a}_u - \mathbf{a}_v\|}{\|\mathbf{a}'_u - \mathbf{a}'_v\|} \tag{4}$$

From the line $L_{j_u^C}(t)$, it can be determined that

$$\|\mathbf{s} - \mathbf{a}_u\| = \|\mathbf{s} - L_{j_u^C}(t_u)\| = \frac{\|\mathbf{a}_u - \mathbf{a}_v\| \cdot \|\mathbf{s} - \mathbf{a}'_u\|}{\|\mathbf{a}'_u - \mathbf{a}'_v\|} \tag{5}$$

Expanding this equation and solving for t_u we find

$$t_u = \frac{\|\mathbf{a}_u - \mathbf{a}_v\| \cdot \|\mathbf{s} - \mathbf{a}'_u\|}{\|\mathbf{a}'_u - \mathbf{a}'_v\| \cdot \|\mathbf{s} - \mathbf{j}_u\|} \tag{6}$$

Using this algorithm, all t_i can be found and thus the simplified problem is solved. We call this algorithm the *3-point-line* algorithm (3PL).

3.2 The 3PLFLS Registration Algorithm

For a 3-permutation \mathcal{M}^B of the point set \mathcal{P}^B there are two sets of approximate solutions, $\hat{\mathcal{R}}_1^C$ and $\hat{\mathcal{R}}_2^C$, that are found from 3PL. Each solution consists of three points that lie on the back-projected lines. The transformations $\hat{T}_{i_B^C}$ from \mathcal{M}^B to $\hat{\mathcal{R}}_i^C$ can be determined by any solution of the absolute-orientation problem [1, 7, 8]. Applying transformation $\hat{T}_{i_B^C}$ to all points in \mathcal{P}^B transforms the points to $\hat{\mathcal{Q}}^C$. For each transformation, the correspondence is chosen by matching a point to the closest line.

Unlike in ICP, this correspondence is now *fixed* for all subsequent iterations.¹ We call this the *Fixed Least Squares* algorithm, abbreviated as FLS. Using \hat{T}_{iB}^C as the initial estimate, the least-squares error between the points and their corresponding lines is minimized by a Levenberg-Marquardt optimization. The final solution is the correspondence that gives the least RMS error from FLS.

We call the registration algorithm that uses three point-line correspondences as initial estimates, and then fixes the correspondences for subsequent least-squares error minimization, the 3PLFLS algorithm. Figure 2 summarizes this algorithm in pseudo-code.

This algorithm is efficient because it is polynomial in the number of markers used. For n points, there are $2n(n-1)(n-2)$ initial 3PL estimates to be computed. Because it would be uncommon to use a large number of fiducial markers, the number of initial estimates is always small (even for 10 markers there are only 1440 initial estimates). Because the optimization algorithm uses fixed correspondences, the computationally expensive nearest-neighbor calculation that tends to dominate ICP calculations is avoided. If the structure of the fiducial points is known beforehand it may be possible to limit the number of Levenberg-Marquardt iterations, so the computation time can indeed be seen to grow cubically with the number of fiducial markers.

4 Experimental Methods

The 3PLFLS algorithm was tested on fluoroscopic and portal images. The fluoroscopic tests were conducted at Kingston General Hospital (Canada) and the portal tests were conducted at Johns Hopkins University (USA). The fluoroscopic tests were performed in conjunction with highly accurate independent 3D detection methods that provided both a relative and an absolute measurement of registration accuracy. The portal images were evaluated with a relative accuracy measurement only.

4.1 Fluoroscopic Methods

The fluoroscopic tests were conducted on two single-plane C-arms (models BV25 and BV26, Philips). Two objects, plastic models of a femur and tibia, were instrumented with six 0.8 mm tantalum fiducial markers each. The location of the fiducial markers were determined by Roentgen stereophotogrammetry and cross-validated with CT, as we have previously reported [4]; the maximum diameter of the point sets was 72 mm. Arrays of infrared light-emitting diodes (IRED's) were attached to the C-arm and the test object; the markers were contacted with an optically tracked probe (Optotrak 3020, Northern Digital) and Horn's algorithm was used to register the location of the markers to the C-arm coordinate system via the local IRED array.

The C-arms were calibrated using our previously reported methods [13]; the image-to-source distance was found to be approximately 900 mm for each C-arm. Images of each test object were taken and the video signals from the C-arm were digitized directly

¹ The original ICP algorithm reselects the correspondence in each iteration.

Algorithm 3PLFLS

```

 $\mathcal{L}^C \leftarrow \{\mathbf{L}_{j_i^C} | 1 \leq i \leq n\}$ 
 $\mathcal{L}_3^C \leftarrow \{\mathbf{L}_{j_i^C} | i = i_1, i_2, i_3 \text{ where } j_i^C \text{ are not collinear}\}$ 
for  $\mathcal{M}^B = \text{permutation of 3 points in } \mathcal{P}^B$  do
  if  $\mathcal{M}^B$  is collinear then
    continue
  end if
   $(\hat{\mathcal{R}}_1^C, \hat{\mathcal{R}}_2^C) \leftarrow \text{3PL}(\mathcal{M}^B, \mathcal{L}_3^C)$ 
  foreach  $\hat{\mathcal{R}}_i^C$  do
     $\hat{T}_B^C \leftarrow \text{AO}(\mathcal{M}^B, \hat{\mathcal{R}}_i^C)$ 
     $\hat{Q}^C \leftarrow \text{Transform}(\hat{T}_B^C, \mathcal{P}^B)$ 
     $\text{Cor}(\cdot) \leftarrow \text{correspondence between } \hat{Q}^C \text{ and } \mathcal{L}^C$ 
     $(\hat{T}_B^C, \epsilon_{rms}) \leftarrow \text{FLS}(\text{Cor}(\mathcal{P}^B), \mathcal{L}^C, \hat{T}_B^C)$ 
    if  $\hat{T}_B^C$  has the least RMS error then
       $T_B^C \leftarrow \hat{T}_B^C$ 
       $\text{BestCor}(\cdot) \leftarrow \text{Cor}(\cdot)$ 
    end if
  end for
end for
 $(T_B^C, \epsilon_{rms}) \leftarrow \text{FLS}(\text{BestCor}(\mathcal{P}^B), \mathcal{L}^C, T_B^C)$ 
 $Q^B \leftarrow \text{Transform}(T_B^C, \mathcal{P}^B)$ 
end algorithm

```

Fig. 2. The 3PLFLS algorithm for finding correspondence and registration. $T_{F_1}^{F_2} = \text{AO}(\mathcal{P}^{F_1}, \mathcal{P}^{F_2})$ calculates the transformation using an absolute orientation solution such that $T_{F_1}^{F_2}$ transforms a point set \mathcal{P}^{F_1} in frame F_1 to a point set \mathcal{P}^{F_2} in frame F_2 . The function 3PL() calculates the two solutions with the three-point-line back-projection algorithm. The pose of each object was estimated with the accelerated ICP algorithm and with the 3PLFLS algorithm; the ICP algorithm was initialized, as recommended by Besl and McKay [2], from a set of 312 distinct poses. Registration accuracy was tested by using the ground-truth data provided by the optical tracking system.

4.2 Portal-Imaging Methods

Portal-imaging presents significant challenges due to its relatively poor spatial resolution and high X-ray energy. Pixels are large (0.9 mm) and the “hard” beam easily penetrates small steel or tantalum objects, leaving no visible trace in the image.

The portal-imaging tests were conducted on a conventional radiotherapy linear accelerator (Varian Systems). A plastic vertebral phantom was instrumented with sharp metal screws of 9-10 mm length, gradually widening, with 3 mm diameter at the end. We selected screws, rather than spherical objects, to better simulate actual clinical circumstances. The phantom was cased in a rigid transparent plastic box that had a thin cross-hair on the top. This arrangement allowed us to set the phantom box first in the center of the CT gantry, then in the isocenter of the linear accelerator, using the laser set-up mechanisms of the therapy device.

The phantom was placed on the patient table and imaged from a nominal AP view. The gantry was then rotated 45° and imaged again. The phantom was slightly rotated and translated, then the imaging protocol was repeated. The four images were processed by manually selecting the center locations of the markers on the portal-imaging console.

5 Experimental Results

The fluoroscopic results were statistically indistinguishable for the two C-arms tested, so the results were pooled. The ICP algorithm was deemed to have failed if it converged to a point-line correspondence that was known *a priori* to be incorrect. The performance of results calculated by ICP and 3PLFLS are shown in Table 1. Relative and absolute errors of registration using 3PLFLS are shown in Table 2.

Table 1. Performance of ICP and 3PLFLS on fluoroscopic images. n is the number of fiducial markers involved; time is in CPU seconds on a SUN Ultra 60. A registration is considered as *failed* if the correspondence of the fiducial markers is incorrect.

Algorithm	$n = 4$		$n = 6$	
	CPU Time	Failure Rate	CPU Time	Failure Rate
ICP	88.52	36.7%	154.58	0
3PLFLS	0.53	0	2.19	0

Table 2. Accuracy of registration for fluoroscopic images. $\delta\theta$ is the rotation error in degrees, δd is the translation error in millimeters. SD is the standard deviation of the errors. MAX is the maximum error.

Object	RMS Error	Absolute Errors	
		$\delta\theta$ (SD; MAX)	δd (SD; MAX)
Femur	0.10 (0.02/0.12)	0.86 (0.40; 1.75)	1.57 (0.77; 3.29)
Tibia	0.06 (0.02/0.09)	0.76 (0.42; 1.77)	1.40 (0.69; 3.40)

Table 3. Root-mean-square (RMS) errors, in millimeters, and angular errors, in degrees, of registration for portal images of an instrumented lumbar vertebral phantom. For each pose of the phantom a portal image was taken at 0° and at 45° orientation of the accelerator beam. The relative angular error is the estimated pose difference from 45° .

Phantom	RMS error, 0° image	RMS error, 45° image	Angular error
Pose 1	0.47	0.43	$+1.44^\circ$
Pose 2	0.46	0.34	-0.66°

6 Discussion and Conclusions

We have developed an efficient and effective algorithm for registering a set of 3D points to a single 2D X-ray image. The algorithm is faster than the well known ICP algorithm by more than a factor of 60. The algorithm makes use of geometric constraints that are natural in the problem domain and exploits the combinatorics of point-line matching.

Efficacy of the algorithm has been established on two fluoroscopic C-arms and one portal imaging system.

The errors in estimating the pose are extraordinarily small, given the difficulty of the problem. The point sets were of only 70 mm in diameter, yet could be registered from a single fluoroscopic image to within 2 mm absolute accuracy for an image-to-source distance of over 900 mm. This result includes all sources of error, including unwarping of the fluoroscopic imaging physics, estimation of the 3D X-ray source location with respect to the image, and estimation of the projected fiducial centroid in the image. In particular, we found that we could estimate the pose within ± 0.5 mm in the plane parallel to the image, and that the translational component normal to the image was the primary contributor to the registration error.

This problem is much harder than the problems encountered in virtual fluoroscopic navigation, where an instrument is tracked optoelectronically and then virtually superimposed on the image by performing forward calculations of perspective projection. In virtual navigation the error in locating the X-ray source location are effectively cancelled by perspective, whereas here such cancellation is not inherent in the problem. Cancellation does, however, occur if *two* point sets are tracked. In such a case the relative pose of the objects can be very accurately estimated, particularly in planes parallel to the image.

The algorithm scales as a moderate polynomial in the number of markers. If there are n markers, then the number of triplets of markers to be exhaustively searched is $O(n^3)$ complexity. If a large number of markers is required by the surgical application, alternative registration algorithms (such as geometric hashing [15]) could be considered. However, the purpose of the proposed algorithm is to accurately solve the registration problem for surgical guidance in which there are not many markers. In particular, the algorithm works very well when the number of markers is minimal (four) or nearly minimal.

There are many potential clinical applications of single-image registration. The technique is directly applicable to conformal radiotherapy, where the technique can localize target anatomy that has been previously instrumented. For percutaneous RF ablation of liver, radio-opaque markers can be placed around a tumor and subsequently be used in fluoroscopic guidance of an ablative instrument (the liver parenchyma is relatively stiff and locally moves as a rigid body). For cancer in the pelvic region, the pelvic bone can be instrumented and nearby soft tissues can potentially be treated percutaneously. In these and other applications, the speed and convenience of single-image registration may be able to improve the therapeutic outcome and improve the quality of life of the patient.

Acknowledgments

This research was supported in part by Communications and Information Technology Ontario, the Institute for Robotics and Intelligent Systems, and the Natural Sciences and Engineering Research Council. Fluoroscopic and CT images were acquired with the

assistance of the Department of Surgery and the Department of Diagnostic Radiology at Kingston General Hospital, Canada. Portal and CT images were acquired with the assistance of the Department of Radiology and the Department of Radiation Oncology at Johns Hopkins university, USA.

References

- [1] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(5):698–700, 1987.
- [2] Paul J. Besl and Neil D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.
- [3] John M. Boone, J. Anthony Seibert, William A. Barrett, and Eric A. Blood. Analysis and correction of imperfections in the image intensifier-TV-digitizer imaging chain. *Medical Physics*, 18(2):236–242, 1991.
- [4] R. E. Ellis, S. Toksvig-Larsen, M. Marcacci, D. Caramella, and M. Fadda. Use of a biocompatible fiducial marker in evaluating the accuracy of CT image registration. *Investigative Radiology*, 31(10):658–667, 1996.
- [5] Keith D. Gremban, Charles E. Thorpe, and Takeo Kanade. Geometric camera calibration using systems of linear equation. In *Proceedings IEEE International Conference on Robotics and Automation*, pages 562–567, Philadelphia, PA, 1988.
- [6] Kenneth R. Hoffmann, Jacqueline Esthappan, Shidong Li, and Charles A. Pelizzari. A simple technique for calibrating imaging geometries. In *Medical Imaging: Physics of Medical Imaging*, pages 371–376. SPIE, 1996.
- [7] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642, April 1987.
- [8] Berthold K. P. Horn, Hugh M. Hilden, and Shahriar Negahdaripour. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7):1127–1135, July 1988.
- [9] Laurent Launay, Catherine Picard, Eric Maurincombe, René Anxionnat, Pierre Bouchet, and Luc Picard. Quantitative evaluation of an algorithm for correcting geometric distortions in DSA images: applications to stereotaxy. In *Medical Imaging: Image Processing*, volume 2434, pages 520–529. SPIE, 1995.
- [10] Ewa Pietka and H. K. Huang. Correction of aberration in image-intensifier systems. *Computerized Medical Imaging and Graphics*, 16(4):253–258, 1992.
- [11] Anne Rougée, Catherine Picard, Cyril Ponchut, and Yves Troussel. Geometrical calibration of X-ray imaging chains for three-dimensional reconstruction. *Computerized Medical Imaging and Graphics*, 17(4/5):295–300, 1993.
- [12] Stephen Rudin, Daniel R. Bednarek, and Roland Wong. Accurate characterization of image intensifier distortion. *Medical Physics*, 18(6):1145–1151, 1991.
- [13] Thomas S Y Tang. Calibration and point-based registration of fluoroscopic images. Master's thesis, Queen's University, Kingston, Ontario, Canada, 1999.
- [14] Roger Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, August 1987.
- [15] H. J. Wolfson and I. Rigoutsos. Geometric hashing: An overview. *IEEE Computational Science and Engineering*, 4(4):10–21, 1997.