

Post-operative Measurement of Acetabular Cup Position Using X-ray/CT Registration

David LaRose^{1,2}, Laura Cassenti², Branislav Jaramaz^{2,3}, James Moody³,
Takeo Kanade^{1,2}, and Anthony DiGioia^{2,3}

¹ Department of Electrical and Computer Engineering, Carnegie Mellon University,
Pittsburgh, PA, USA

² Center for Medical Robotics and Computer Assisted Surgery, Robotics Institute,
Carnegie Mellon University, Pittsburgh, PA, USA

³ Center for Medical Robotics and Computer Assisted Surgery, UPMC Shadyside
Hospital, Pittsburgh, PA, USA

Abstract. This paper describes a system for measuring acetabular implant orientation following total hip replacement surgery. After a manual initialization procedure, the position of the pelvis is established relative to a pair of nearly orthogonal radiographs by automatically registering to pre-operative pelvic CT data. The pose of the cup is then recovered by projecting a 3D surface model into the two images. A phantom study is presented in which this pose is expressed relative to well defined anatomical landmarks and compared to measurements obtained using an image-guided surgery system.

1 Introduction

Following total hip replacement surgery, acetabular implant orientation is typically measured using anterior-posterior (AP) X-rays[9]. Implant orientation has been shown to be predictive of post-operative outcome, and of complications such as dislocation[6][7]. Despite recent improvements in measurement technique[3], these measurements are still unreliable due to unknown pelvic flexion at the time of X-ray acquisition.


This paper describes a system which measures acetabular implant placement with respect to the pelvis independent of pelvic flexion. The pose of the implant is initially recovered in the coordinate system of the CT volume by 2D/3D registration between planar X-ray images, the pre-operative CT volume, and a 3D surface model of the implant. This pose is then expressed relative to stable anatomic landmarks which are defined in the CT volume, providing an accurate basis for measurements of cup orientation and post-operative evaluation.

Initial results are presented in which pose estimates are generated using images of a high-density pelvis phantom. These estimates are compared with ground truth results obtained using the HipNav image-guided system for total hip replacement surgery [1].


2 Problem Description

During post-operative evaluation, the pose of the acetabular implant must be determined with respect to the pelvis. AP and lateral X-ray images are acquired with known source-to-film distances. The approximate projection centers of each X-ray image are known, but the pose of the patient is not. In particular, the images are not known to be true AP or lateral views, and each image may be acquired with the patient either lying or standing. In addition to the X-ray images, the pre-operative pelvic CT volume and a triangulated surface model of the implanted cup are available.

Although the X-ray images are not acquired simultaneously, it is assumed that the pose of the cup with respect to the pelvis does not change between acquisitions. A schematic drawing of this scenario is shown in Fig. 1(a). There are three unknown coordinate transformations in this figure: the transformation from the coordinate system of the acetabular cup to the CT coordinate system, which is labeled ${}^{ct}T_{cup}$; the transformation from the CT coordinate system to the coordinate system of the X-ray imager at the time of AP image acquisition, labeled ${}^{i0}T_{ct}$; and the transformation from the CT coordinate system to the coordinate system of the X-ray imager at the time of lateral image acquisition, labeled ${}^{i1}T_{ct}$. Each of these unknown coordinate transformations has six degrees of freedom, for a total of 18 unknown parameters.



(a) Schematic of problem.



(b) User specified seed points.

Fig. 1. Schematic of problem and input image with marked points.

3 Approach

The problem is broken down into 3 steps: 1) All parameters are initialized to reasonable starting values using input from the user; 2) The pose of the pelvis is estimated in each image using an iterative intensity-based 2D/3D registration

scheme; and 3) The pose of the acetabular implant is estimated with respect to the coordinate system of the CT volume by simultaneously matching the projection of the implant surface model to contours in each X-ray image. These three steps are discussed in sections 3.1, 3.2, and 3.3 respectively.

3.1 Initialization

The 18 unknown degrees of freedom are initialized to reasonable starting values by user input. To initialize the coordinate transformations ${}^{i0}T_{ct}$ and ${}^{i1}T_{ct}$, the user indicates the approximate positions of at least three anatomical landmarks in each radiograph and then enters the corresponding 3D coordinates from the CT volume. The initial pose is estimated using a point based 2D/3D registration.

The initial pose of the acetabular cup implant is known from the pre-operative plan. In addition, the user provides the image positions of several points on the border of the cup in each image as shown in Fig. 1(b). These points are used to further constrain the position of the cup as described in Sec. 3.3.

3.2 X-ray/CT Registration

After initialization, the pose of the pelvis with respect to the X-ray imaging apparatus is recovered for each image by iterative comparison between the input images and synthetic X-ray images. These synthetic images are known as Digitally Reconstructed Radiographs (DRRs).

DRR Generation Traditional DRR generation is computationally expensive, and involves numerically integrating through the CT volume along the path which connects each pixel to the X-ray source, as shown in Fig. 2. To avoid this time consuming integration during registration, DRRs are generated using an intermediate data structure called a Transgraph [4]. The Transgraph is a 4D database of precomputed integral values, and is based on data structures used in the computer graphics area called *view-based rendering* [2][5]. DRRs can be generated over a limited range of pose parameters by interpolating among the values in the database. This interpolation procedure is much faster than conventional DRR generation and allows generation of 621x512pixel images in roughly 2 seconds on a 550 MHz Pentium III computer, without requiring clipping of the CT volume. Work is currently underway to implement the interpolation using off the shelf texture mapping hardware as described in [2]. We expect this change to yield a significant additional speedup.

Image Comparison. Since the pose of the pelvis is estimated by iterative comparison of DRRs with the input radiograph, registration accuracy depends on an appropriate choice of image comparison metric. The images are compared using a local *semi-normalized* correlation measure which works as follows:




Fig. 2. Path of a single ray from radiation source to film surface. The box indicates the volume in space which is represented by the CT dataset. p_1 and p_2 represent the points at which the ray enters and exits this volume.

1. For each pixel in the synthetic image, define a surrounding region of interest. For the rest of this paper, we assume that the region of interest is an 11 pixel by 11 pixel window.
2. Compute the variance of the synthetic image over each 11x11 window, excepting those windows which extend past the border of the image.
3. Compute the normalized correlation between each of these 11x11 windows and the corresponding region of the input image.
4. Scale each correlation value by the associated variance, and compute the weighted average.

More precisely, we compute the local semi-normalized correlation by:

$$\hat{C}(I_0, I_1) = \frac{\sum_{x=5}^{M-6} \sum_{y=5}^{N-6} C(I_0, I_1, x, y)}{\sum_{x=5}^{M-6} \sum_{y=5}^{N-6} (S(I_1, I_1, x, y) / K^2)} \tag{1}$$

$$C(I_0, I_1, x, y) = \frac{S(I_0, I_1, x, y) \sqrt{S(I_1, I_1, x, y)}}{K^2 \sqrt{S(I_0, I_0, x, y)}} \tag{2}$$

$$S(I_p, I_q, x, y) = K \sum_{i,j \in W(x,y)} I_0(x+i, y+j) I_1(x+i, y+j) \tag{3}$$

$$- \sum_{i,j \in W(x,y)} I_0(x+i, y+j) \sum_{i,j \in W(x,y)} I_1(x+i, y+j) , \tag{4}$$

where \hat{C} is the weighted average described in step 4 above, C is the scaled correlation value computed over a single region of interest, I_0 is the input X-ray image, and I_1 is the DRR. The two images are the same shape, and the constants M and N represent their width and height in pixels. $W(i, j)$ is the 11x11 region of interest surrounding pixel (i, j) and the constant K is equal to $11^2 = 121$. Note that the pixel values of I_1 depend on the estimated pose of the pelvis, although this dependence is not made explicit in the notation. In Sec. 3.2 we introduce the pose parameter vector θ which specifies this estimated pose. In order to evaluate these expressions efficiently, all of the summations within S are computed using recursive filters.

The local semi-normalized correlation function has three key advantages: first, the intensity response function of the film/digitizer need not be known, as long as it is approximately linear over small variations in irradiance; second, weighting by the variance of the synthetic image concentrates attention in those regions of the image which have high signal strength; and third, any unmodeled variations in the X-ray images (such as those caused by soft tissue deformation) are effectively filtered out, provided that they have low spatial frequency compared to the correlation window size.

Occasionally, one or more of the 11x11 regions in the input image has all pixels at the same intensity. When this happens, the normalized correlation coefficient for that window is undefined, and a value of zero is substituted.

Minimization. The local semi-normalized correlation function reaches a maximum of 1.0 when the input image and the synthetic image match each other perfectly. A non-negative function which has a theoretical minimum of zero is generated by subtracting the local semi-normalized correlation value from 1.0, and this function is minimized using the quasi-Newton algorithm of Broyden, Fletcher, Goldfarb and Shanno, as described in [8]. The minimization is run for each input image, to recover the coordinate transformations ${}^{i0}T_{ct}$ and ${}^{i1}T_{ct}$.

Gradient Computation. Quasi-Newton minimization requires that the first derivative of the objective function be computed. The free parameters in the optimization describe the pose of the pelvis with respect to the X-ray imaging apparatus, and we represent these parameters using the vector θ . Repeated application of the chain rule to equation (2) gives

$$\frac{\partial}{\partial \theta_i} \widehat{C}(I_0, I_1) = \frac{K^2 \sum_{x,y} \frac{\partial}{\partial \theta_i} C(I_0, I_1, x, y)}{\sum_{x,y} S(I_1, I_1, x, y)} - \frac{2K^2 \left(\sum_{x,y} C(I_0, I_1, x, y) \right) \left(\sum_{x,y} S(I_1, \frac{\partial}{\partial \theta_i} I_1, x, y) \right)}{\left(\sum_{x,y} S(I_1, I_1, x, y) \right)^2}$$

$$\frac{\partial}{\partial \theta_i} C(I_0, I_1, x, y) = \frac{S(I_0, \frac{\partial}{\partial \theta_i} I_1, x, y)S(I_1, I_1, x, y) + S(I_0, I_1, x, y)S(I_1, \frac{\partial}{\partial \theta_i} I_1, x, y)}{K^2 \sqrt{S(I_0, I_1, x, y)S(I_1, I_1, x, y)}} ,$$

where the notation $\sum_{x=5}^{M-6} \sum_{y=5}^{N-6}$ from equation (2) has been replaced with $\sum_{x,y} \frac{\partial}{\partial \theta_i} I_1$ is the the first derivative of DRR pixel intensity with respect to the parameter θ_i , and all other variables are defined as in equation (2).

In general, the derivative of DRR pixel intensity with respect to pelvis pose parameters is difficult to compute. In this case, however, the DRR pixel values are found by interpolating among the values stored in the Transgraph, and the first derivative is easily found.

3.3 Determination of Cup Position

Once the pose of the pelvis has been estimated in both images, the pose of the acetabular cup implant is estimated with respect to the coordinate system of the CT volume. This is done by simultaneously matching the projection of the implant surface model to contours in the two X-ray images.

For a given pose, the silhouette of the surface model is projected into each image, and an error measure is computed based on the image positions of the projected points. This error measure is then minimized over the parameter space of the rigid body transformation ${}^{ct}T_{cup}$. The first subsection below describes how the silhouette is computed, the second describes a rough point-based registration used to approximate the actual cup position, and the third describes a final minimization which increases the registration accuracy.

Silhouette Generation. To generate the silhouette of the cup, the vertices of the cup surface model are projected into each image. The projected vertices define a set of 2D triangles corresponding to the 3D triangles of the surface model. The silhouette is generated by culling those vertices which lie interior to any of the projected triangles. To speed this culling process, the projected triangles are organized into a quadtree data structure, and each vertex is compared against only those triangles which lie in or intersect its cell in the quadtree.

Approximate Solution. The pose of the cup is initially computed based on the image coordinates supplied by the user during manual initialization. We define the objective function

$$f(\gamma) = \frac{1}{|R|} \sum_{r \in R} \min_{t \in H_0(\gamma)} (\|r - t\|) + \frac{1}{|S|} \sum_{s \in S} \min_{u \in H_1(\gamma)} (\|s - u\|) , \quad (5)$$

where γ is the vector of parameters describing the rigid body transformation ${}^{ct}T_{cup}$, R is the set of user-supplied initialization points in the AP image, and $|R|$ is the number of points in this set. S is the set of user-supplied initialization points in the lateral image, and $|S|$ is the number of points in this set. $H_0(\gamma)$ is the set of points comprising the silhouette of the cup in the AP image at the pose specified by γ , and $H_1(\gamma)$ is the set of points comprising the silhouette in the lateral image. The notation $\|x\|$ denotes the magnitude of vector x .

The objective function in equation 5 reaches a minimum of zero when every user supplied point is exactly overlapped by one of the points on the silhouette of the projected model. In practice, this minimum is nearly met when the boundaries of the projected surface model lie close to the edges of the cup in the X-ray images. The objective function is minimized using the downhill simplex method of Nelder and Mead, as described in [8].

Refinement of Approximate Solution. The minimization above gives a good approximation to the pose of the acetabular cup. There are, however, small

inaccuracies. This is because the initialization points may not lie exactly on the boundaries of the implant in the X-ray images, and because these points may not match well with the points which make up the silhouette of the projected cup. The estimate is, however, good enough to initialize a more precise search.

The gradient of pixel intensity with respect to pixel coordinates x and y is computed for each input image, and an objective function is defined

$$g(\gamma) = \frac{1}{|H_0(\gamma)|} \sum_{t \in H_0(\gamma)} \|\nabla_{x,y} I_0(t[0], t[1])\| \\ + \frac{1}{|H_1(\gamma)|} \sum_{u \in H_1(\gamma)} \|\nabla_{x,y} I_1(u[0], u[1])\| ,$$

where γ is the vector of parameters describing ${}^{\text{ct}}T_{\text{cup}}$, $|H_0|$ is the number of points in the silhouette corresponding to the AP image, and $t[0]$ and $t[1]$ are the x and y coordinates of projected vertex t . Similarly, H_1 , $u[0]$ and $u[1]$ refer to the silhouette corresponding to the lateral image. $\nabla_{x,y} I_0()$ and $\nabla_{x,y} I_1()$ are the spatial gradients the two images.

4 Results

A preliminary study was conducted using phantom data to evaluate the performance of the registration algorithm. A 62mm diameter VerSys Acetabular cup (Zimmer, Inc.) was fitted to a high density pelvis phantom. The pelvis coordinate system was defined with respect to the left and right anterior iliac spines, and the left and right pubic symphysis points. These points were marked on the model with 1mm diameter fiducial markers. A CT dataset was acquired, having an intra-slice pixel spacing of approximately 0.74mm, a slice thickness of 1mm, and an inter-slice spacing of 1mm. The fiducials were identified in the CT and used to compute the coordinate transformation between the pelvis coordinate system and the CT coordinate system. In order to prevent the fiducial markers from biasing the X-ray/CT registration, an image manipulation program was used to paint out any visible markers in the subsequent X-ray images.

Four series of images were acquired. In each series, AP films were taken with an approximate source-to-film distance of 40 inches, while lateral films were taken with an approximate source-to-film distance of 72 inches. The first series involved 3 AP images and 3 lateral images. Each image was acquired from a slightly different angle, and in four of the six images, household objects were placed in the field of view in order to simulate occluding patient anatomy as shown in Fig. 3(a).

For the second series of images, a simulated torso, surrounding the pelvis, was constructed out of plastic film and filled with rice to simulate soft tissue. Small balloons were inserted into the rice to simulate the effects of bowel gas and soft tissue inhomogeneity. In these images, the difference in density between the rice and the pelvis phantom was very small, and there was not enough contrast for successful registration.

For the third and fourth series of images, the rice was replaced with oatmeal in order to improve the visibility of the pelvis. Although the contrast in these images is still quite low, there is sufficient detail for successful registration. These two series differ in the arrangement of the soft tissue and in the placement of the bowel gas. In addition, the acetabular cup was removed prior to acquisition of the fourth series, and reattached in a different orientation. As before, the AP images of each series span a range of roughly 20° in flexion, while the lateral images span a range of roughly 15° rotation around the superior-inferior axis. Typical images from these series are shown in Figs 3(b) and 3(c).

Ground truth measurements were performed using the HipNav system [1]. Repeated measurements of the implant orientation for the first three series had a mean abduction of 45.2° and a mean flexion of 10.4° , with standard deviations of 0.11° and 0.22° respectively. After repositioning for the fourth series, the cup was measured to have an orientation of 52.6° abduction and 48.9° flexion.

All of the films were digitized and resampled to resolution of 36 dpi, giving a final image size of 621x512 pixels. The center of projection for each image was assumed to lie at pixel coordinates (310, 255.5). No further attempt was made to calibrate the X-ray imaging system. The images from each series were grouped into pairs, each consisting of one AP image and one lateral image. The registration algorithm was run using each pair as input to recover the pose of the acetabular cup implant with respect to the coordinate system of the CT. This transformation was composed with the CT-to-pelvis coordinate transformation, and flexion/abduction measurements were calculated. The outline of the cup surface model was projected into each image so that the accuracy of registration could be visually assessed.


Abduction and flexion measurements are presented in 1, and visual registration results for cup alignment are shown in Fig. 3. The poor result in case 7 is due to incorrect convergence of the X-ray/CT registration, and is discussed further in Sec. 5.

Series 1	Abduction	Flexion	Series 2	Abduction	Flexion	Series 3	Abduction	Flexion
HipNav	45.2°	10.4°	HipNav	45.2°	10.4°	HipNav	52.6°	48.9°
Case 1	45.3°	13.0°	Case 4	46.2°	11.9°	Case 7	59.3°	-10.7°
Case 2	45.6°	13.1°	Case 5	45.4°	10.3°	Case 8	52.6°	47.0°
Case 3	45.2°	12.3°	Case 6	45.8°	10.8°	Case 9	50.8°	50.2°


Table 1. The results closely reflect the nominal pose of the acetabular cup.

5 Discussion and Conclusion


These initial results are encouraging. The recovered cup orientation matches the measured ground truth to within 2° abduction and 3° flexion in all except one



(a)



(b)



(c)

Fig. 3. (a) A pair of input images from the first series of radiographs. The inset shows recovered cup position, and a peanut butter jar is visible in each image. (b) A pair of input images from the third series of radiographs. In the lateral image, the superior boundary of the simulated torso runs almost parallel to the superior edges of the iliac crests. (c) A pair of input images from the fourth series of radiographs for which the X-ray/CT registration was inaccurate, making it impossible to simultaneously align the cup in both images. The bright line running superior-inferior in the lateral image is a lexan plate to which the pelvis is attached.

of the trial cases. RMS error among these “correctly” converged cases was 0.78° abduction and 1.78° flexion. The results are presented in this space since these are the measures most commonly used intra-operatively.

Although the X-ray/CT registration algorithm converges reliably for AP images, correct convergence for lateral images is somewhat sensitive to the initial conditions. This is illustrated by the failure to converge correctly in case 7 above. One possible reason for this sensitivity is that the bilateral symmetry of the pelvis leads to a pose ambiguity when the view direction is very nearly lateral. In these cases it becomes very difficult, for example, to distinguish the right iliac crest from the left iliac crest. This is especially true when the source-to-film distance is large, as is the case in this study, since a large source-to-film distance decreases foreshortening due to perspective projection.

Future work includes automating the pose initialization to eliminate the need for user input, resolving the pose ambiguity in direct lateral images, and applying the system to real patient data.

References

- [1] A. M. DiGioia, B. Jaramaz, M. Blackwell, D. A. Simon, F. Morgan, J. E. Moody, C. Nikou, B. D. Colgan, C. A. Aston, R. S. Labarca, E. Kischell, and T. Kanade. The otto aufranc award. image guided navigation system to intraoperatively measure acetabular implant alignment. *Clinical Orthopaedics and Related Research*, 355:8–22, October 1998.
- [2] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen. The lumigraph. In *Computer Graphics Proceedings, Annual Conference Series*, pages 43–54, 528, New Orleans, LA, USA, August 1996. ACM SIGGRAPH.
- [3] B. Jaramaz, C. Nikou, and T. J. Levison. Cupalign: Computer-assisted postoperative radiographic measurement of acetabular components following total hip arthroplasty. In *Proceedings of MICCAI '99*, pages 876–882, Cambridge, UK, September 1999.
- [4] D. LaRose, J. Bayouth, and T. Kanade. Transgraph: Interactive intensity-based 2d/3d registration of x-ray and ct data. In *Proceedings of SPIE - International Society for Optical Engineering, Medical Imaging 2000 (In press)*, San Diego, CA, USA, February 2000.
- [5] M. Levoy and P. Hanrahan. Light field rendering. In *Computer Graphics Proceedings, Annual Conference Series*, pages 31–42, New Orleans, LA, USA, August 1996. ACM SIGGRAPH.
- [6] G. E. Lewinnek, J. L. Lewis, and R. Tarr. Dislocations after total hip replacement arthroplasties. *J Bone Joint Surg*, 60A:217–220, 1978.
- [7] D. E. McCollum and W. J. Gray. Dislocation after total hip arthroplasty: Causes and prevention. *Clinical Orthopaedics and Related Research*, 261:159–170, 1990.
- [8] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 1988.
- [9] R. Sellers, D. Lyles, and L. Dorr. The effect of pelvic rotation on alpha and theta angles in total hip arthroplasty. *Contemporary Orthop.*, 17:67–69, 1988.