Proceedings of the 3rd IEEE Workshop on Applications of Computer Vision, Dec. 2-4, 1996, Sarasota, Florida, pp. 181-186.

Pose Estimation of Artificial Knee Implants in Fluoroscopy Images Using A Template Matching Technique[1]

William A. Hoff (whoff@mines.edu), Richard D. Komistek (rkomiste@mines.edu), Douglas A. Dennis, Scott Walker (sawalker@mines.edu), Eric Northcut (enorthcu@mines.edu), Keith Spargo (kspargo@mines.edu)

Division of Engineering
Colorado School of Mines
Golden, CO 80401

Rose Musculoskeletal Research Center
2425 S. Colorado Blvd., Suite 200
Denver, CO 90220

Abstract

This paper describes an algorithm to estimate the position and orientation (pose) of artificial knee implants from X-ray fluoroscopy images using computer vision. The resulting information is used to determine the kinematics of bone motion in implanted knees. This determination can be used to support the development of improved prosthetic knee implants, which currently have a limited life span due to premature wear of the polyethylene material at the joint surface. Our algorithm determines the full 6 degree of freedom translation and rotation of knee components. This is necessary for artificial knees which have shown significant rotation out of the sagittal plane, in particular internal/external rotations. By creating a library of images of components at known orientation and performing a template matching technique, the 3-D pose of the femoral and tibial components are determined. The entire process, when used at certain knee angles, gives a representation of the positions in contact during normal knee motion.

1. Introduction

More than 100 types of arthritis now afflict millions of Americans, often resulting in progressive joint destruction in which the articular cartilage (joint cushion) is worn away causing friction between the aburnated (uncovered) bone ends. This painful and crippling condition frequently requires total joint replacement using implants with polyethylene inserts (Figure 1).

Although artificial knee joints are expected to last 10 to 15 years, research indicates that most implants last an average of just 5.6 years [1]. With population longevity rising, many patients will require additional surgery to replace dysfunctional prosthetic joints to achieve two decades or more of use. Currently, more than 400,000 Americans receive lower extremity implants per year, accounting for over $1.5 billion annually in health care costs.

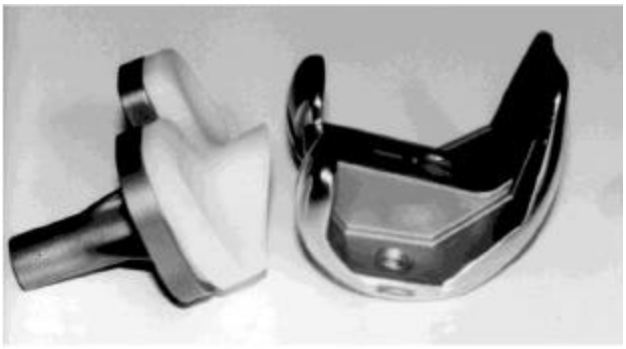

Figure 1 Components of knee implant, with polyethylene insert (white material).

A key problem with knee (and hip) implants is the premature wearing of the polyethylene inserts. Repeated frictional contact generates loose particles which lead to inflammation of the joint, osteolysis, and eventual implant loosening. It is hypothesized that the kinematics of artificial knees are different than in normal knees, and may involve excessive sliding and rotational motions which lead to high shear stresses. If the kinematics of artificial knee joints can be measured "in vivo", this information could be used to help design implants that better replicate normal knee kinematics.

X-ray fluoroscopy is a useful tool to analyze knees "in vivo". X-rays are emitted from a tube, pass through the knee and strike a fluorescent screen where the images are intensified and recorded via video tape [2]. The result is a perspective projections of the knee, recorded as a continuous series of images (Figure 2(a)).

[1] The support of the Rose Foundation and Johnson & Johnson Professional Incorporated is gratefully acknowledged.

Our problem was to analyze the fluoroscopy images to determine the relative pose (position and orientation) of the two implant components (tibial and femoral) with respect to each other. This paper describes a technique we developed to do this, based on template matching. CAD models of the components were developed and a library of silhouette, or template images was created. These silhouette images were matched against the extracted silhouettes from the fluoroscopy images. The pose of each component, as well as their relative pose, was then derived. Figure 2(b) shows an image of the CAD models overlaid on top of the original X-ray fluoroscopy image. Our technique has been used to analyze the kinematics of several different knee types and to quantify the effect of several kinematic phenomena, including sliding and edge lift-off [3-6].

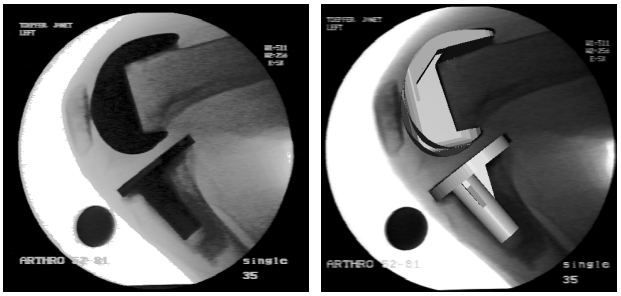

Figure 2. (a) Fluoroscopy image of artificial knee. (b) Image with CAD model overlay.

Section 2 describes previous related work in this area. Section 3 describes our developed technique in detail. Section 4 presents the results of tests on synthetic and real data to measure the accuracy of the method, and Section 5 provides conclusions.

2. Previous work

Previous work with fluoroscopy images has primarily concentrated on measuring the rotation and translation of components within the plane of the image. Stiehl and Komistek *et. al.* [7] analyzed still photographs of the fluoroscopic video to measure rotations and translations of the knee joint members. Their work was limited to the in-plane rotations and translations of the implant components. However, the actual motion of the components also includes rotations and translations out of the plane of the image.

In computer vision, the problem of estimating the pose of a known rigid object from 2-D images has been widely studied. In our application, the primary feature that can be extracted from the image is the extremal contour, or silhouette, of the object. There have been a number of algorithms described in the literature that make use of the silhouette.

Lowe [8] describes an iterative least-squares algorithm for aligning the projected extreme contours of the model with edges found in the image. However, this technique assumes a polyhedral model; and would be difficult to apply to the knee implants, which have smooth, complex surfaces. Kriegman and Ponce [9] used rational surface patches, implicit algebraic equations, and elimination theory to obtain analytic expressions for the projected contours. However, this method is restricted to objects with only a few patches, and would be difficult to apply to knee components, which have highly complex surfaces. Lavallee, *et. al.* [10] describe an algorithm which minimizes the 3-D distances between the rays (corresponding to the points on the contour) and the closest point on the surface of the object. A 3-D distance map is pre-computed that stores the distance from any point in the neighborhood of the object to the closest point on the surface. Lavallee developed an octree-spline technique to speed up the construction of the distance map, which otherwise would be prohibitively slow.

Another, perhaps simpler, approach for pose estimation is to use a template matching technique. If the complete silhouette of the object is visible, an algorithm can match the entire silhouette of the object with a pre-computed template. For two-dimensional applications such as character recognition, the object and the template differ only by a translation and a rotation within the plane of the image. For three-dimensional applications such as automatic target recognition, the silhouette of the object changes shape as it rotates out of the image plane. A solution to this problem is to pre-calculate and store a complete set of templates of the object, representing the object over a range of possible orientations.

Banks and Hodge [11, 12] used this approach to measure the full six degree of freedom motion of knee prostheses by matching the projected silhouette contour of the prosthesis against a library of shapes representing the contour of the object over a range of possible orientations. They measured the accuracy of the technique by comparing it to known translations and rotations of prostheses "in vitro". They report 0.74 degrees of rotational accuracy (in all three axes), 0.2 mm of translational accuracy parallel to the image plane, and 5.0 mm of translational accuracy perpendicular to the image plane. Our process is similar to this method. The difference is that we use a direct template matching technique instead of Fourier descriptors. We also utilize larger image libraries to increase the resolution of the angular measurements.

3. Process description

The central idea in our process is that we can determine the pose of the object (*i.e.*, a knee implant component) from a single perspective image by measuring the size and shape of its projection in the image. A simple technique for doing this is to match the silhouette of the object against a library of synthetic images of the object, each rendered at a known position and orientation. The image with the best match would directly yield the position and orientation of the object in the input image. However, a library which encompasses all possible rotations and translations would be prohibitively large. As an example, suppose we want to determine an object's pose within 5 mm and 2 degrees, and the allowable range is 50 mm translation along and 20 degrees rotation about each of the axes. Dividing each range by the resolution results in a 11^6 or 1,771,561 entry library.

To reduce the size of the required library, we use a simplified perspective model. This model assumes that the shape of an object's silhouette remains unchanged as the object is moved towards or away from the imaging sensor. This is not strictly true because the fluoroscope is a true perspective projection. However, it is a reasonable approximation if this translational motion is small. In the case of a subject flexing their knee in the sagittal plane, the translation out of the sagittal plane is in fact typically small.

With the simplified perspective assumption, the shape of the object is independent of its distance from the imaging sensor (although its size is dependent on the distance). Therefore, we generate a library of images of the object, all rendered at a constant (nominal) distance from the sensor. When we process an input image, we correct for any difference in distance by scaling the size of the unknown object silhouette so that its area is equal to the area of the library silhouettes.

The library of images consists of views of the object rendered at different out-of-plane rotations. The object is rotated at 1° increments about the x axis, and at 1° increments about the y axis. The object is always centered in the middle of the image. The object is always rotated within the plane of the image so that its principal axis is aligned with the horizontal (x) axis. Thus, the library is only two dimensional rather than 6 dimensional. The range of rotation is ±15 degrees about the x axis, and ±20 degrees about the y axis. Figure 3 shows a portion of the library of images for a femoral component. There are a total of 41 x 31 = 1271 images in each library.

3.1 Creating the library

The software modeling program AutoCAD™ was used to create and render (in perspective projection) 3-D models of the implant components. The settings for the perspective projection were determined by calibrating the fluoroscope using a pinhole camera model. We then converted the rendered images into "canonical" form. This was achieved by scaling the silhouette to a specified area (15000 pixels) and rotating the object so that its principal axis was aligned with the x axis. The amount of scaling and rotation was recorded, for later use by the pose estimation algorithm. Finally, the library images were converted to binary form (one bit per pixel), and stacked to form a single multi-frame image. The size of a library for a single component was about 7 Mbytes of data.

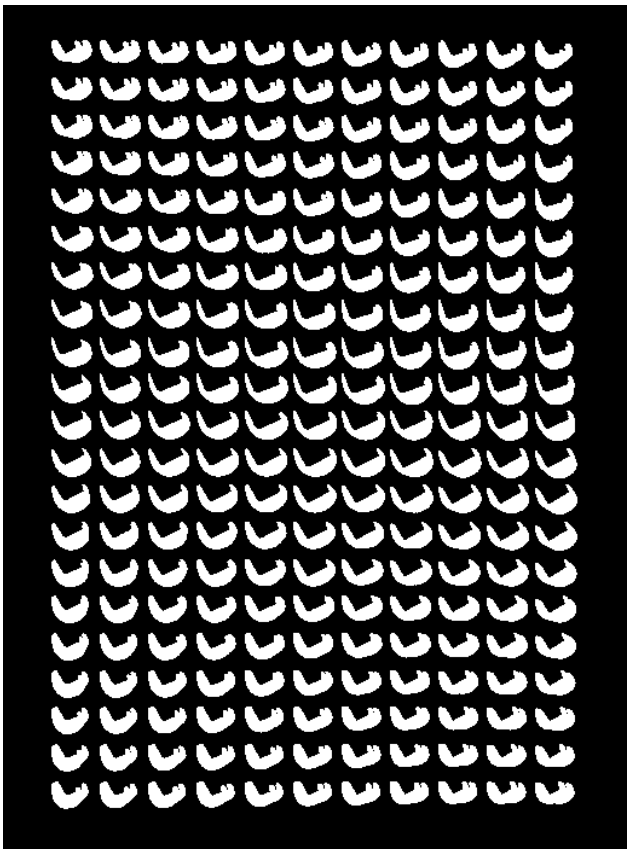

Figure 3. Portion of library for femoral implant component.

3.2 Pose estimation

Analysis of a fluoroscopy video begins with digitizing selected images from the sequence on a Silicon Graphics workstation. These images are then stretched to equalize the horizontal and vertical scale. The images are then input to the software that extracts the silhouette of the component and estimates its pose.

The image processing algorithms described in this paper were implemented using a public domain image

processing software package called Khoros, from Khoral Research Inc.. Khoros is extensible and many new programs were developed and integrated during the course of this work. Figure 4 shows the visual programming workspace that performs the silhouette extraction and pose estimation. Each box, or "glyph," within the workspace performs an operation on an image (or images), and the lines between the glyphs show the transfer of data.

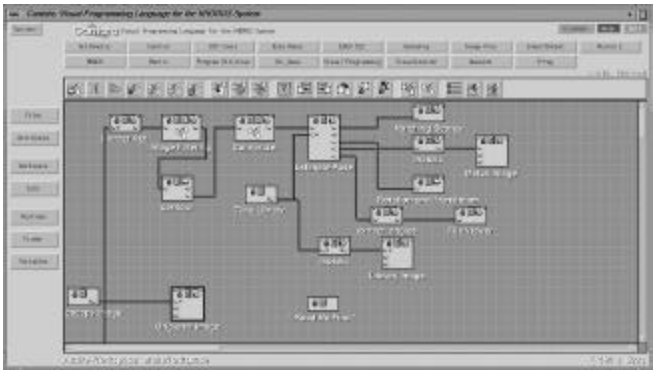

Figure 4 Data flow in Khoros workspace.

The process begins by manually extracting a rough region of interest from the input fluoroscopy image, that contains the component of interest. This is done to reduce the size of the image to be processed and speed up the computation time. The reduced region of interest image is passed through a median filter to reduce the effect of noise.

Next, the contour (or silhouette) of the implant component is extracted. Currently, this is done manually, by having the user designate points around the boundary of the implant with the mouse. We have found it difficult to reliably extract the contour automatically, due to the presence of nearby objects such as bone cement that are nearly the same intensity as the implant. The resulting binary image is passed to a process called "Canonize", which automatically converts the silhouette image to a canonical form. As described earlier, the canonization process centers the silhouette, scales it to a constant area, and rotates it so that its principal axis is aligned with the horizontal (X) axis. The resulting canonical image can be directly matched with the library.

The next step finds the best match of the input canonical image with a library image. This is done by systematically subtracting the input canonical image with each of the library images and generating a "score" for each, which is the number of unmatched pixels. Figure 5 shows the matching results for a particular image. The black areas indicate the unmatched pixels. The library image with the best match determines the two out-of-plane rotation angles of the object (θ_X, θ_Y).

We then find the remaining degrees of freedom of the object. The in-plane rotation angle θ_X is determined by taking the difference between the input image's in-plane rotation angle and the library image's in-plane rotation angle:

$$q_Z = q_Z^{input} - q_Z^{library}$$

The Z position of the object is determined by dividing scale of the fluoroscopy image by the scale of the library and multiplying that by the initial Z distance that the library image was rendered:

$$Z = Z_{library} \cdot \left(s_{input} / s_{library} \right)$$

Figure 5. Matching results.

To determine the x, y position of the object, we compute the 2D image vector from the image centroid of the object to its projected origin. We then calculate the (x,y) location of the object (in inches) relative to the sensor's coordinate frame. Finally, we correct the θ_X and θ_Y rotation angles to take into account the effect of x,y translation on the apparent rotation. The result is the full 6 DOF pose of the object (X, Y, Z, θ_X, θ_Y, θ_Z) relative to the sensor frame.

As a check, the CAD model of the implant is projected onto the original X-ray image, using the derived pose data and the known model of the imaging system. The model should fit the actual image silhouette closely. Figure 2(b) shows an example of the CAD models for the femoral and tibial components overlaid on the original image.

4. Experiments

To check the accuracy of the pose estimation process, we performed two experiments. For the first experiment, we created a set of synthetic images of the implant components, rendered at pre-determined poses. Three different component models were used: one tibial implant and two femoral implants. Each component was rendered in 8 different poses (Table 1) for a total of 24 test images.

Table 1 Poses used for synthetic images.

Pose	θ_X	θ_Y	θ_Z	X	Y	Z
1	0	0	-60°	0	0	0
2	8°	0	0	0	0	0
3	0	8°	0	0	0	0
4	8°	5°	-60°	0	0	0
5	0	0	0	1 in	0	0
6	0	0	0	0	1 in	0
7	0	0	0	0	0	1 in
8	0	0	0	1 in	1 in	1 in

The root mean squared (RMS) pose errors of the components are shown in Table 2, and agree approximately with the results reported by Banks and Hodge. Note that errors are largest for out-of-plane rotations and out-of-plane translations.

Table 2 RMS pose error, synthetic images.

	Femoral1	Femoral2	Tibial
In-Plane, Rotational	0.14°	0.22°	0.28°
Out-of-Plane, Rotational	1.50°	1.24°	1.31°
In-Plane, Translational	0.013 in	0.018 in	0.013 in
Out-of-Plane, Translational	0.088 in	0.054 in	0.036 in

In the next experiment, we mounted the femoral and tibial components on rotational platforms and took fluoroscopy images of them at known orientations. Figure 6 shows the experimental setup. Although the rotational platforms can be rotated very precisely, we did not know the precise absolute pose of the components relative to the sensor, due to uncertainties in mounting and initial alignment. We therefore measured only the relative pose between the two components.

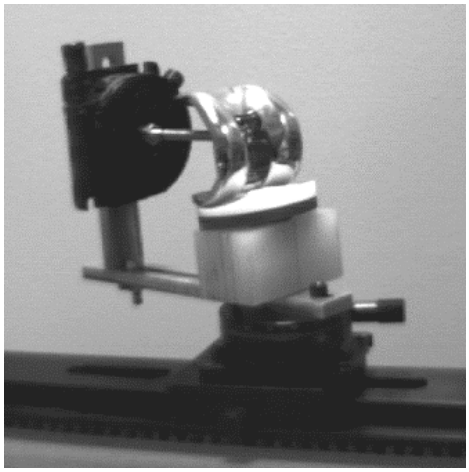

Figure 6 Apparatus for pose tests.

Our apparatus allowed us to rotate the femoral component relative to the tibial component about the horizontal axis. This angle of rotation is called the "lift-off" angle and is of great interest in the orthopedics community. A non-zero lift-off angle indicates that the weight of the femur is being borne by only one of the two condyles (the two curved runners on each side of the femoral implant component). If such a situation occurred, it could lead to increased wear on that side and eventual implant failure. Therefore, a key goal of our tests was to determine how accurately we could measure lift-off angle.

Our apparatus also allowed us to rotate the tibial component about the horizontal (X) and vertical (Y) axes. The femoral component was fixed with respect to the tibia (except for the single degree of freedom corresponding to the lift-off angle) and thus was rotated as well. The complete set of test poses are shown in Table 3 (all angles are expressed in degrees).

Table 3 Test poses, real images.

Case	θ_x	θ_y	Lift-off angle
1	8	-5	-5
2	8	-5	-15
3	8	-10	-3
4	8	-10	-7
5	8	-15	-6
6	12	3	-2

Case	θ_x	θ_y	Lift-off angle
7	12	3	-6
8	12	3	-10
9	12	8	4
10	12	13	1
11	12	13	9

We processed the fluoroscopy images and computed the relative orientation of the femoral component relative to the tibial component for each of the 11 cases. The angle of rotation about the horizontal (X) axis was computed and compared to the known ground truth lift-off angle for each case, shown in Table 3. In addition, we also measured the angle of rotation about the other two axes (Y

and Z). These angles *should* have been identically zero (for the Y axis) and identically 90° (for the Z axis), since the femoral component was rigidly mounted with respect to the tibial component, except for the single degree of freedom rotation about the X axis. The results are shown in Table 4.

Table 4 Pose errors, real images.

	Angular error (absolute value) in degrees		
Case	*X axis (lift-off)*	*Y axis*	*Z axis*
1	0.684	0.346	1.02
2	11.540	0.240	0.92
3	0.590	0.266	0.1
4	2.930	0.027	0.36
5	0.092	0.613	1.98
6	16.215	0.306	0.13
7	0.884	1.772	1.48
8	0.171	0.147	1.72
9	18.396	0.706	1.55
10	0.439	0.360	0.49
11	2.881	0.653	0.56
AVG	4.984	0.494	0.937
RMS	8.275	0.670	1.131

In Table 4, note the unusually large errors for cases 2, 6, and 9. The reason for this is that the pose estimation algorithm chose the incorrect library match for several of the tibial images. This is very easy to do because the tibial component is very symmetrical about the Y axis. Thus, the silhouette for a negative rotation about the Y axis is very similar to the silhouette for a positive rotation of the same magnitude.

To correct this problem, which also arises occasionally in processing "in vivo" (patient) images, we always verify the correct match by visually inspecting the overlay of the model onto the original image. The operator can usually detect mismatches by noticing small errors in the overlay. Another method we use to verify correct matches is to check whether the relative poses of the tibial and femoral components are consistent. For example, in the case of an incorrect match, the femoral component usually intersects the tibial component (a physically impossible situation).

If a matching error is detected, the operator can go back and force the system to choose a different match. The operator can also make minor adjustments on the rotation angles to achieve a better overlay with the image. We applied this procedure to our test images and re-computed the pose errors. The revised pose error results are shown in Table 5. As can be seen, manually adjusting the match resulted in much smaller angular errors. In our experiments, the person performing these operations did not have knowledge of the "correct" poses and thus was not biased in choosing a different match or adjusting the angles.

Table 5 Revised pose errors, real images.

	Angular error in degrees		
	X axis (lift-off)	*Y axis*	*Z axis*
AVG	0.256	0.145	0.123
RMS	0.298	0.195	0.240

5. Conclusions

This paper has described a simple yet effective technique for estimating the pose of artificial knee implants from X-ray fluoroscopy images. The pose estimation process has been a useful aid for determining "in vivo" knee kinematics in implanted knees. "In vitro" experiments show that we can measure lift-off angle with an RMS error of about 0.3°.

One of the limitations is that the entire process requires a significant amount of human interaction for development of libraries, and external knowledge of the implants for contour extraction. The contour extraction is the area where most of the error can be attributed. This is due to human variability in picking the vertices for the contour.

6. References

[1] M. M. Landy and P. S. Walker, "Wear of ultra-high molecular weight polyethylene components of 90 retrieved knee prostheses," *Journal of Arthoplasty*, Vol. 3, No. pp. s73-s85, 1988.

[2] R. L. Perry, "Principles of conventional radiography and fluoroscopy," *Veterinary Clinics of North America: Small Animal Practice*, Vol. 23, No. 2, pp. 235-252, 1983.

[3] J. B. Stiehl, R. D. Komistek, D. A. Dennis, R. D. Paxson, and W. A. Hoff, "Fluoroscopic analysis of kinematics after posterior-cruciate retaining knee arthroplasty," *Journal of Bone and Joint Surgery-British*, Vol. 77-B, No. 6, pp. 884-889, 1995.

[4] R. D. Komistek, D. A. Dennis, and W. A. Hoff, "A Kinematic Comparison of Prosthetically Implanted and Nonimplanted Knees Using Dynamic Fluoroscopy," *Proc. of 19th Annual Meeting of the American Society of Biomechanics*, Stanford, California, August, 1995.

[5] D. A. Dennis, R. D. Komistek, J. B. Stiehl, and E. Cheal, "An In Vivo Determination of Condylar Lift-Off Using an Inverse Perspective Technique that Utilizes Fluoroscopy," *Orthopaedic Transactions*, Vol. 1996.

[6] S. A. Walker, W. A. Hoff, D. A. Dennis, and R. D. Komistek, ""In Vivo" Pose Estimation of Artificial Knee

Implants Using Computer Vision," *Proc. of Annual Rocky Mountain Bioengineering Symposium*, ISA, United States Air Force Academy, Colorado Springs, April, pp. 143-150, 1996.

[7] J. B. Stiehl, R. D. Komistek, D. A. Dennis, R. D. Paxson, and W. A. Hoff, "Biomechanical Aspects of Posterior Cruciate Retaining Total Knee Arthroplasty," *Journal of Contemporary Orthopaedics*, Vol. 1995.

[8] D. G. Lowe, "Fitting Parameterized Three-Dimensional Models to Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 13, No. 5, pp. 441-450, 1991.

[9] D. J. Kriegman and J. Ponce, "On Recognizing and Positioning Curved 3D Objects From Image Contours," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. PAMI-12, No. 12, pp. 1127-1137, 1990.

[10] S. Lavallee, R. Szeliski, and L. Brunie, "Anatomy-Based Registration of Three-Dimensional Medical Images, Range Images, X-Ray Projections, and Three-Dimensional Models Using Octree-Splines," in *Computer Integrated Surgery: Technology and Clinical Applications*, R. Taylor, Ed., Cambridge, Massachusetts, MIT Press, pp. 115-143, 1996.

[11] S. A. Banks and W. A. Hodge, "Direct measurement of 3D knee prosthesis kinematics using single plane fluoroscopy," *Proc. of Orthopedic Research Society*, pp. 428, 1993.

[12] S. A. Banks and W. A. Hodge, "Comparison of dynamic TKR kinematics and tracking patterns measured "in vivo"," *Proc. of Orthopedic Research Society*, pp. 665, 1994.