

# Interactive Pre-operative Selection of Cutting Constraints, and Interactive Force Controlled Knee Surgery by a Surgical Robot

S.J.Harris\*, M. Jakopec\*, R.D. Hibberd\*, J. Cobb†, B.L. Davies\*

\* Mechatronics in Medicine Laboratory,  
Department of Mechanical Engineering  
Imperial College of Science, Technology and Medicine,  
Exhibition Road, London SW7 2BX, England

† Department of Orthopaedic Surgery,  
The Middlesex Hospital,  
Mortimer Street, London W1N 8AA, England

email: [s.j.harris@ic.ac.uk](mailto:s.j.harris@ic.ac.uk)

**Abstract.** This paper describes a low-cost computer system that takes CT images of the knee, and with three-dimensional models of knee prostheses allows a surgeon to position the prosthesis correctly pre-operatively in an interactive manner. Once in position the computer can process bone and prosthesis geometry to derive a set of constraint boundaries that constitute a safe cutting area for a force controlled robot (i.e. that avoids soft tissue such as ligaments), and provides the correct cutting planes for good prosthesis/bone alignment. This boundary information is used to program the robot, allowing a surgeon to move the robot within predefined regions to machine away bone accurately whilst preventing damage to soft tissue.

**Keywords:** robotics, orthopaedics, force control, imaging

## 1. Introduction

Total knee replacement is a common surgical procedure used particularly with the elderly to replace worn and painful arthritic joints. To achieve good, pain-free mobility with a reliable joint, a degree of precision is necessary both with regard to the positioning of the prosthesis components relative to the bone axes and to each other, and with regard to the mating surfaces of the bone and prosthesis. The prosthesis components are complex shapes. To achieve the correct cuts in the bone, a series of jigs and fixtures are applied sequentially which ideally should produce a set of accurate cuts. However, the sequential nature of the application can result in small errors becoming compounded, as the reference frames are changed, compromising the requirement for accuracy.

To reduce the cumulative errors, a robotic aid is being developed which measures its position for each cut from a single reference frame. This removes the need for a series of jigs and fixtures as the surfaces of the bone for cutting are held within a computer's memory. Emphasis should be placed on the fact that the robot is an aid to the surgeon rather than a replacement. Automated surgical procedures have been tried by this

group [1][2] and others [3], however, it has been found after discussions with surgeons that there has been some resistance to them, due to the lack of continuous control and tactile sensation by the surgeon during the operation. The robotic aid under development can be used by the surgeon to locate a bone-cutting rotary milling-bit in the correct position for each cut and prevent it from being moved into dangerous regions and from removing too much bone. The surgeon has a 'hands-on' approach, enabling him to feel the quality of the bone being cut, and adjust his cutting forces and speed appropriately. Thus, the surgeon uses his skills and senses to the full, while the robot provides the geometrical accuracy required for a good fit and the physical constraints to avoid soft tissue damage.

## **2. Image Acquisition**

Before the bones can be cut to fit the prosthesis, images of the joint and surrounding area are necessary to allow the surgeon to position the prosthesis optimally on the bone. For manual surgery assessing the optimal component sizes is done with templates positioned on x-ray images of the joint while the optimal position is found per-operatively.

For the robotic system it was decided that CT images of the leg would be used. These have a number of advantages over normal x-rays, particularly of use in a robotic system: (1) there is no variable, depth related image magnification, thus accurate measurements can be made; (2) full three dimensional data is available, so x-ray images and bone cuts can be simulated from any angle and three dimensional reconstructions made; (3) the positions of fiducial markers can be found in three dimensions allowing registration of the CT data and physical bone.

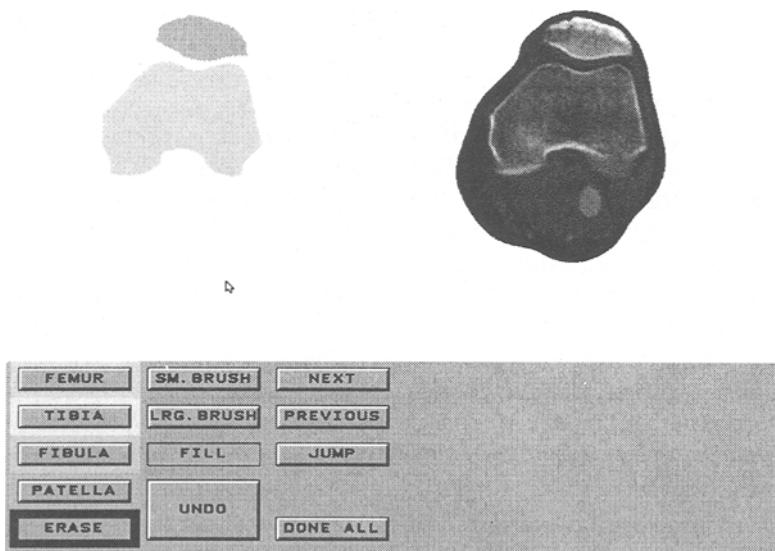
Because of the accuracy required for positioning the prosthesis, CT images are needed with a fine slice spacing (1 mm or better). The use of CT images gives the surgeon the ability to perform all prosthesis sizing and positioning pre-operatively, reducing the amount of time required during surgery and allowing him to plan the optimal cutting sequence. The imaging, modelling and planning phases have been implemented on a low-cost PC which is robust enough for clinical use, but nevertheless is adequate for the modelling, planning and implementation procedures.

## **3. CT pre-processing**

Before the images can be used, they must be processed to distinguish between soft tissue, femur, tibia and fiducial markers. To distinguish between bone and soft tissue, a bone threshold level is set. A flood fill procedure is then run on each image, starting from outside of any bone, set to replace any pixels with black up to a border of intensity greater than the threshold. This preserves the medullary canal within the bone. Often this has an intensity level similar to that of soft tissue and would hence be removed by normal thresholding methods. This procedure preserves it because the surrounding high intensity bone image provides a border that prevents the flood-fill from reaching it. Once soft tissue has been removed, the bone is thresholded to black and white. This is processed interactively to distinguish between femur, tibia and other leg bones. Starting from the top of the leg, the bone image is taken and each individual area of bone is coloured separately. The surgeon can then re-colour the areas in one of

four colours depending on bone type (figure 1). He also has basic painting and deleting tools to remove artefacts, separate touching bones, and replace bone in the image that was below the threshold level (this can happen if the x-ray beam just skims the surface of bone). Having completed the first slice, the next is presented. However, on this one the computer checks each bone area in the image, comparing it with the image above. Regions in the current slice that contact regions that the surgeon has marked in the previous slice are coloured to match those already marked. Because of the need for small slice spacing, and hence many images, this semi-automatic process speeds up the bone identification procedure, as the surgeon only needs to take action when the bone types change, or when the image requires correction for artefacts.

To give the surgeon more complete landmark information, as each bone image is displayed, the unprocessed CT slice is shown alongside, allowing him to determine whether the computer has mistakenly joined two bones or deleted bone material.

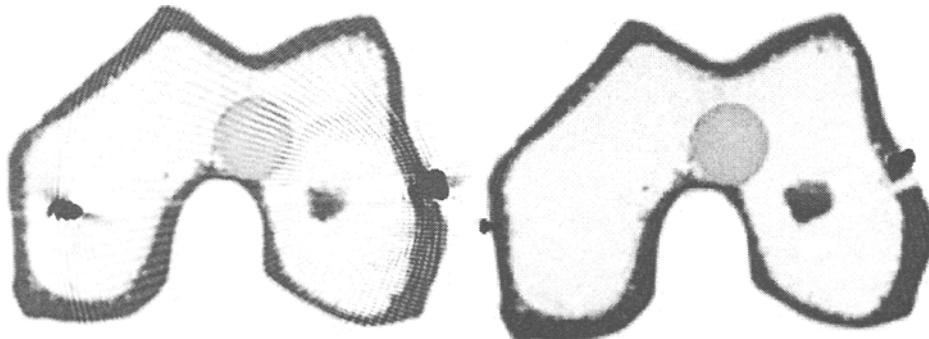


**Figure 1** - Annotation of CT images to discriminate between bone. Note the simplified image (left) and the original CT (right) to provide the surgeon with landmark features

#### 4. Detection of fiducial markers

For initial trials, it is intended that metal screws are used as fiducial markers. These indicate reference points that can be used to align the coordinate frames of the robot with the CT images. Trials with different screw types have revealed that stainless steel screws have a much greater intensity than bone, and hence should be easy to find. However, they are too bright, and can cause the following problems with the image: *blooming* - the region around the screw becomes very bright, resulting in an uncertainty in the screw position; *image distortion* - bone intensities in the image near

the screw are elevated, so bone/soft tissue discrimination may give errors in such regions; *image interference* - reflections off the screws results in patterns of dark and bright banding across the image that can make image processing difficult. These problems can be seen in figure 2. Note that the CT image has been inverted to show bone in darker intensities as this appears more clearly in printed images.



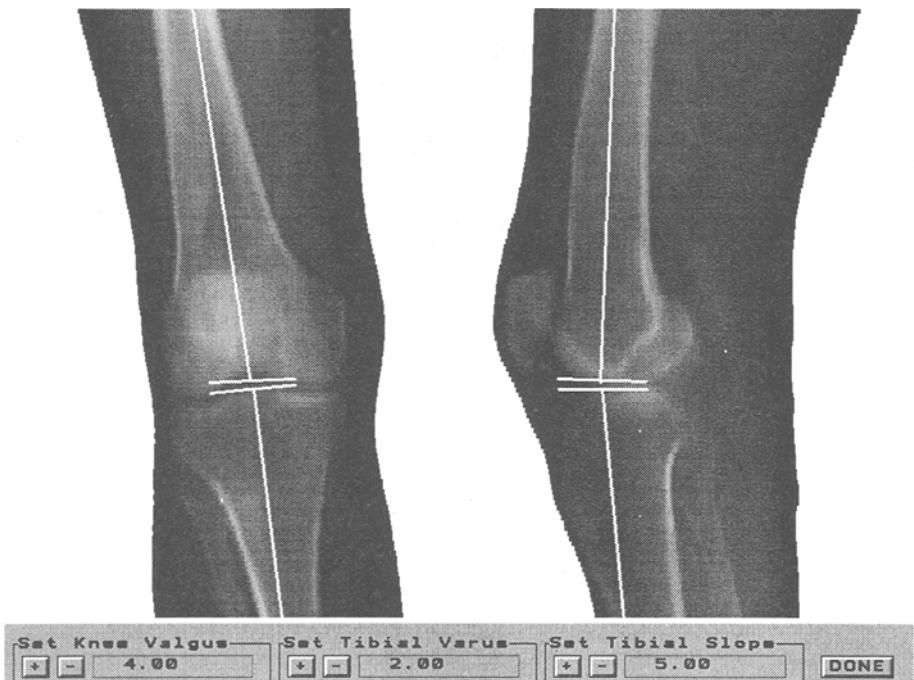
**Figure 2** - Comparison of the effects of steel (left) and titanium (right) fiducial pins on the image quality. Note the blooming of steel pins and the pattern superimposed over the image.

As an alternative, titanium screws have been tried. These have an intensity in the CT image not much greater than bone (although still distinguishable by thresholding) and do not cause so much distortion or interference. Because of the small intensity difference, artefacts such as noise in the image can cause false detection of pixels above the screw detection threshold. To discriminate between this noise and real screws, pixels above the threshold are checked with the images above and below in the CT record. If either of these images also contains pixels at the same coordinates which are above the screw threshold, then the pixels under test are assumed to be screws. As the CT scan is processed, a record is made of the coordinate of every remaining pixel found to be above the screw threshold.

Clusters are created by grouping pixels that are close together. Each cluster is then processed in turn to compute its position and major axis. The outermost end point of the cluster along the major axis is chosen as the contact point for mechanical registration of the robot during surgery. Very small clusters are ignored as they are assumed to be noise or artefacts.

## 5. Initial positioning of the prosthesis

After pre-processing the CT to determine bone types and fiducial marker positions, the surgeon provides an initial estimate of the prosthesis position. For normal posture, there should be approximately 5 to 8° valgus angle between the long axes of the tibia and femur, while for the prostheses used there should be a varus tilt of approximately 0 to 2°, and a posterior slope of 5° [4]. For initial positioning of the prosthesis, the surgeon can mark on X-ray images computed from the CTs where the long bone axes are and specify the appropriate angles using a simple graphical interface. An example of long bone axis selection can be seen in figure 3.



**Figure 3** - Interactive alignment of prosthesis with long bone axes. Axes are selected by dragging marker lines on the simulated X-rays. Angles are selected using the control panel

Having obtained an initial estimate of the prosthesis position the surgeon can fine tune it by moving the prosthesis parts on the CT images, and selecting appropriately sized components to find a good match with the bone geometry. During this procedure the bone can be viewed either as X-ray images (as shown in figure 4) or CT slices to give an overall view or fine detail at particular depths, and in any of three orientations (slice by slice, medial-lateral or anterior-posterior). The prosthesis can be viewed as a wire-frame or as a solid model (figure 4). This process mimics quite closely the procedure currently used by surgeons manually to size the prosthesis components, presenting them with a task with which they are already familiar. It does, however, allow them to assess the prosthesis size and position from views that would otherwise be impossible (e.g. the CT slice views allow the tibial component to be matched to the shape and orientation of the tibial plateau). In addition the interference between the prosthesis and the bone in individual CT slices can be shown. This allows the surgeon to detect whether, for example, the top of the femoral component is notching the femur, and creating a weak point in the bone. To visualise the completed procedure, the surgeon can view a three-dimensional rendering of the bone with the appropriate cuts made and the prosthesis in position.

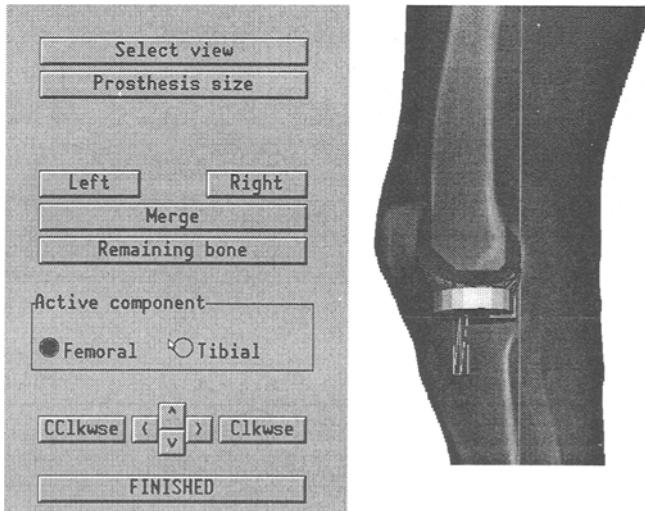


Figure 4 shows a simulated X-ray with a model of the prosthesis components superimposed

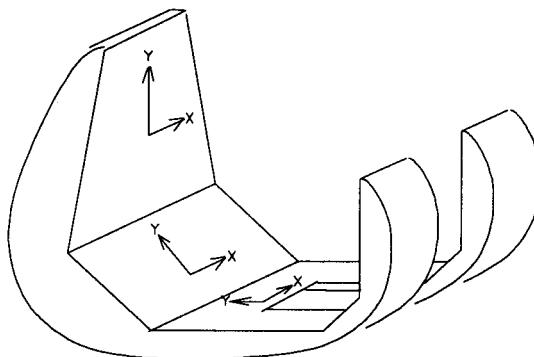
## 6. Computing the cutting boundary

When cutting leg bones to accept a prosthesis, it is usual to cut to the edge of the bone rather than to the edge of the prosthesis. Since prostheses come in a limited number of sizes, this prevents the generation of a brittle bony lip around the prosthesis that could both interfere with prosthesis component mating and be susceptible to breaking and damaging nearby tissue. Once the prosthesis position has been finalised, it is therefore necessary to compute the cutting planes required to generate the five surfaces on the femur and the one surface on the tibia, and to find the outline of bone required to be cut on these surfaces.

The current requirement (neglecting fixation and stabilisation shafts and vanes) for each mating surface is that they are flat planes. Therefore, each surface has a coordinate (its centre point) and a pair of unit vectors associated with it (figure 5 shows a typical component). These unit vectors describe the direction taken in the space of the CT image to move one unit on a cutting plane in the X and Y direction, assuming the plane is being viewed along its perpendicular. The cutting boundary will then be defined by points in this plane and the orientation and position of the plane. Initially the centre point and unit vectors are set up with the prosthesis not rotated and referenced to the CT origin. They are subjected to the same transformations as the prosthesis components as they are moved into position on the CT. To generate the cutting outline a two-dimensional image array is set up and voxels in the CT bone data record mapped onto the plane using the centre point and unit vectors. If the centre point for a plane is  $(C_x, C_y, C_z)$  and the two unit vectors are  $(X_x, X_y, X_z)$  and  $(Y_x, Y_y, Y_z)$ , then a point  $(a, b)$  on the cutting plane is taken from the CT bone data

record at point  $(C_x + aX_x + bY_y, C_y + aX_y + bY_y, C_z + aX_z + bY_z)$ . Any bone other than the one expected to be cut for the current plane is ignored.

Once the cutting plane image has been created a Roberts edge extraction filter is run over it. Although the Roberts filter does not fare well with low contrast images, the ones dealt with here are already thresholded to a high contrast so this simple, fast edge extractor is sufficient for the task. This gives an outline for the bone area to be cut, which is chain-encoded to follow its outline and produce an ordered cutting region.



**Figure 5-** Directions of unit vectors for three of the cutting planes required for a typical femoral component, relative to the component itself. These vectors are rotated to match the rotation of the prosthesis as it is moved into position

An additional task performed is to run through the mapping of CT data onto the cutting plane, but this time to scan outwards along the perpendicular to the cutting plane until the CT data registers that the edge of the bone has been reached. A count of the number of voxels passed through from the cutting plane to the bone edge is recorded for each pixel in the cutting plane. This procedure generates a height map, indicating the shape of the region of bone that will be removed.

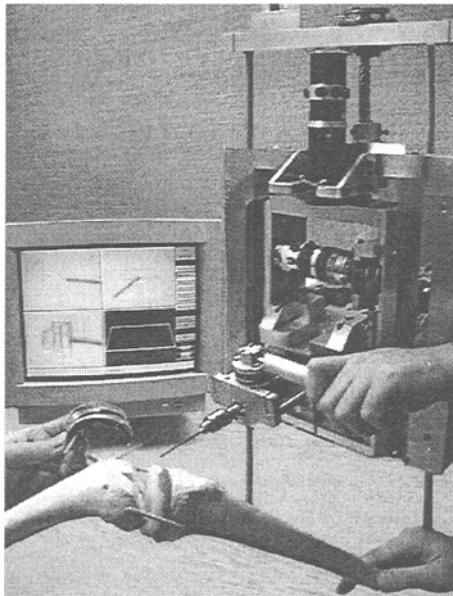
These processes are repeated for each of the six cutting planes and the resulting data stored before it is passed onto the robot.

## 7. The Acrobot active constraint robot

The robot aid is a four axis system (vertical motion, in-out motion, yaw and pitch) as shown in the photograph of figure 6 with a rotating milling bit as the end-effector. The mechanical and electrical arrangement has been described in detail elsewhere [5][6]. Backdriveable gearboxes on the motors enable it to be moved around by hand as well as under computer control, or using servo assist as some mixture of the two. A 6-axis force sensor with joystick type handle allows the robot to sense the surgeon's movement forces and move accordingly. This provides the robot with the surgeon's force requests in all three dimensions, and torque measurements around three axes.

Under normal operation the robot is controlled by the surgeon, who can direct its movements using the joystick handle, but is constrained to resist movements outside of

a safe region, as defined by the pre-operative processing described above. This is achieved using a variable force control strategy. It should be emphasised that the low and equal impedance which gives sensitivity of force control, together with good stiffness when required is a compromise which is quite different from the requirements for, say, registering the tip of a conventional industrial robot under force control.



**Figure 6** - Photograph of the Acrobot knee surgery robot in bench trials showing the robot (right) and a typical interactive display (left)

For yaw and pitch motions, the gear ratios were chosen to give approximately equal backdriveable force, so the surgeon has no sensation of being 'guided' in a specific direction unless at a boundary. The in-out axis requires a greater backdriving force, but preferential motions in other axes are compensated for by servo assistance on this axis. The vertical motion axis is not driven under surgeon control. The position on this axis is selected by the computer to allow optimum access to each cutting surface.

Movement of the cutting robot is limited to a small operating environment. Thus to position and align it correctly with respect to the knee, a passive positioning system is used. This carries the robot and allows it to be moved under manual control in three linear axes until the surgeon is satisfied that the cutter is correctly placed to obtain the desired cutting planes. A fourth axis allows the robot to be rotated by  $\pm 45^\circ$  with respect to the vertical (orientation shown in figure 6). This allows it easy access to all bone surfaces. The milling cutter and control strategy allow the robot to mill away bone with the cutter perpendicular, at an angle, or parallel to the required cutting plane. The latter method is fastest and is used preferentially where the geometry of the cuts and bone access allows.

## 8. Control of the robot

As with the pre-operative processing functions, the real-time robot control system uses graphics interactively with the surgeon, allowing him to select cutting planes using a simple set of on-screen control buttons and graphics. These provide views of the position of the robot's cutter in relation to the plane currently being cut, and the depth of bone currently cut. The robot controller comprises two computer systems. A PC compatible system controls user interaction, providing the graphical interface and storage for the outline database as generated by the pre-operative process, and overall control and monitoring. A DSP controlled system, hosted by the PC, provides control of the robot's motors and signal processing from the force sensor system.

The robot is force controlled on the yaw, pitch and in-out control motors (as shown in figure 6) by adjusting the torque of the motors depending on the force applied by the surgeon, and the current position of the robot in relation to the cutting boundaries. The vertical motion motor is under position control, and is used to move the cutter to the optimum height for each cutting plane. To allow the surgeon to feel how close to the boundary he is, the control strategy varies the admittance (i.e. the relationship between the force applied by the surgeon and the torque of the motors) in varying regions, by providing a two level control system. At the lower level is a closed-loop PD system whose gains can be adjusted to control stiffness. This control loop includes friction and gravity compensation to remove preferential motion that would otherwise result from gravity or differences in friction in each motor linkage. An outer control loop monitors the position of the robot relative to boundaries, and also the force and direction of the surgeon's requested movements. The controller adjusts the inner loop parameters to allow a strategy of servo assistance or neutral control, increasing stiffness or position control depending on whether the robot is in a safe region far from boundary regions, getting increasingly close to the boundary, or whether the surgeon has forced the robot over the boundary. In the first case, the surgeon is free to move anywhere within this region without constraint, and servo assist is used to balance the admittances on the various movement axes. In the second case, as the surgeon moves towards the boundary, the stiffness is increased, making it difficult to cross into a dangerous region, and providing force feedback to the surgeon that he is nearing the boundary. In the last case if the surgeon has overstepped the boundary (a situation that can't be completely eliminated as this would otherwise require very large, powerful motors) the robot provides position control normal to the constraint boundary to bring the cutter back to a safe region, while giving low resistance force control tangential to the boundary. Thus it is possible to move around the constraint boundary while the computer resists movement outside it.

Because of the nature of the prosthesis used, the control regime is set up to allow each plane to be cut in sequence. This simplifies the control task since each constraint region can be reduced to a two-dimensional boundary of coordinates from the pre-operative planning system and a flat plane. These two cutting constraints can be controlled individually and result in a prismatic (or 'tunnel-like') constraint, terminated in a flat plane. Considering the plane part of the regime, the stiffness is adjusted in the direction of the normal to the plane. This increases as the distance to the plane decreases if the cutter is pushed towards the plane. For the prismatic boundary, the

nearest point on the constraint curve is found (i.e. when projected onto the cutting plane, how close would the cutter be to the edge of the safe region). The stiffness of the robot is then adjusted independently in normal and tangential directions to the boundary curve at this nearest point. These stiffnesses, each of which are computed along three orthogonal axes (normal to the appropriate constraint surface and in two directions along the surface), from the two control calculations are then combined before being passed to the inner loop.

## 9. Positional accuracy of the robot and control algorithm

Preliminary tests have been performed to acquire information about the robot's accuracy and repeatability in terms of positioning, and its ability to cut flat planes.

To test positional accuracy, the milling tip was replaced by a 4mm diameter ball-ended probe. This probe was then placed in 30 holes of 3mm diameter in known positions in a perspex panel. Being ball ended, the probe tip is automatically centred on the hole. This process was repeated 10 times for each hole. From this test it was found that each hole could be found to within 0.13mm of its expected position on average, while the repeatability was measured to be 0.16mm (1 standard deviation).

Tests were made of the robot's accuracy when cutting flat planes in three modes - with the cutter axis perpendicular to the plane, with the cutter axis at 45° to the plane, and with the axis parallel to the plane. During the cutting process, position data from the robot was obtained and the positions of the deepest cuts, along the length of the milling cutter, were logged on a 0.1mm resolution grid covering the cutting plane. It was found that the cutter could be used to cut with little deviation from the cutting plane. For each of the cutting regimes, results were as follows:

Cutting method	Average distance of cut from plane	Max distance of cut from plane	Standard deviation
Perpendicular	-0.19mm	-0.35mm	0.039mm
Parallel	-0.15mm	-0.38mm	0.065mm
45° to plane	-0.15mm	-0.28mm	0.034mm

Negative values indicate depths past the required cutting plane. These results show that cuts are on average less than 0.2mm deeper than required. Although deeper than the required cutting plane, they are still well within the tolerances demanded by surgeons. The small standard deviations indicate that there is very little 'waviness' over the surface of the plane.

## 10. Conclusions

The pre-operative system provides a user-friendly method of setting up the information required by the robot for its cutting outline. Where ever possible, the procedures performed on screen have been designed to mimic those that the surgeon would normally do manually with x-ray images. The use of CT instead of x-ray images enables the surgeon to do pre-operatively many of the functions that were previously done during surgery (e.g. final decisions as to the prosthesis size).

The accuracy of the robot is well within the required accuracy required by surgeons (who have indicated that a 1mm error is acceptable overall). It should be emphasised

however that this is only one link in the chain. For accurate imaging and modelling CT image sets are required with sub-millimetre slice spacings, which generates large amounts of data, and exposes the patient to increased radiation (although this accuracy is only needed in the immediate vicinity of the knee). The control strategies have been found to enable planes to be cut with a good degree of flatness.

Initial tests indicate that the robot can match the shape of the prosthesis allowing a good match between prosthesis component and bone. Further tests are required to determine the accuracy of the complete system, in addition to that of the robot.

The approach of co-operation between surgeon and machine enables both the accuracy of the robot and the senses and intelligence of the surgeon to be used together for maximum effect, and provides a more acceptable solution than that of an autonomous system.

## 11. Acknowledgements

We would like to acknowledge the financial assistance of the UK Department of Health and the 'LINK' Medical Implants Fund in support of this project.

## 12. References

1. Harris, S.J. *et al.*: Transurethral Resection of the Prostate: The evolution of a robotic system to meet users' needs, Proceedings of the First International Symposium on Andrology Bioengineering and Sexual Rehabilitation (1995), 73-80
2. Davies, B.L.: Safety of Medical Robots, Safety Critical Systems, Chapman Hall Press, Part 4, chapter 15 (1993) 193-201
3. Taylor, R.H, *et al.*: An Image-Directed Robotic System for Precise Orthopaedic Surgery, Computer Integrated Surgery, Chapter 28, The MIT Press (1996) 379-396
4. Howmedica: Kinemax modular total knee system - Instrument Surgical Technique (1995)
5. Harris, S.J., *et al.*: Experiences with Robotic Systems for Knee Surgery, Proceedings of CVRMed-MRCAS'97, Springer-Verlag (1997) 756-766
6. Davies, B.L. *et al.*: Active compliance in robotic surgery - the use of force control as a dynamic constraint, Proceedings of the Institution of Mechanical Engineers, Part H (1997) 285-292