Frameless Patient Registration Using Ultrasonic Imaging

Thomas Ault, M. W. Siegel
The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

ABSTRACT

Locating the patient accurately in real-time is necessary for many problems in medical robotics. Current solutions are expensive or inconvenient. We show progress toward solving these problems by using ultrasound to track CT reference features in real-time. We compute the locations of target features in the ultrasonic images relative to the CT reference features.

Introduction

Problem Definition

An important limiting factor in the performance of autonomous or semi-autonomous robots in the surgical arena has been accurately and precisely registering the patient to the robot's reference frame. Presumably, the robot operates with a model of the patient indicating the location of the target structures as well as those structures to be strictly avoided, such as nerves or major blood vessels. Even though a model may represent the location of the patient with respect to the robot at the start of the procedure, should the patient move significantly during the procedure, the model would no longer reflect accurately the position of structures within the patient relative to the robot's instruments. Consequently, the robot might cut through a nerve that now lies in an area marked as "safe" in its model, or it might not remove all of a target tumor. Were the robot able to track the location of the features of the patient relevant to the operation, it could compensate in its internal model for motion of the patient and avoid such hazards.

Solution Space

Furthermore, being able to locate key features of the patient relative to the robot coordinate frame accurately and in real-time enables other technologies. One possibility involves overlaying data acquired pre-operatively on the view of a laprascope, allowing the surgeon to effectively see beneath the surface of the tissues visible through the laprascope, helping her to avoid accidentally cutting or damaging important tissues. Remote surgery could be similarly assisted by such a system. Not only could the remote surgeon have a view of the location of the patient's nerves, bones and blood vessels relative to her instruments, but the remote system could also define a safety zone around structures that must not be touched. The system could then implement an emergency shutdown should the surgeon accidentally attempt to move her instruments into these areas.

Framed vs. Frameless

The ideal patient-tracking system would be accurate, inexpensive, place no extra burden on the surgeon, and pose no unnecessary hazard to the patient. No existing system achieves all these goals. Several systems that track the patient's location during surgery have been designed and used. The strategies these systems use may be divided into two categories, framed and frameless. Framed strategies place some fiduciary marks (*e.g.* pins) in the rigid structures (*e.g.* bones) of the patient and track them. Framed strategies require the use of an additional operation to insert the pins. An example of a framed system is the ISS Robodoc/ Orthodoc system [1]. Frameless strategies either track easily-removed markers on deformable structures, such as the skin, or track natural features of rigid structures in the body, avoiding the use of fiduciary markers altogether. Frameless strategies that track markers on deformable structures suffer from errors introduced by the distortion of those structures. Examples of systems using frameless strategies include the retrospective patient image registration system developed by Tan, *et al* [2] and the surgical navigation system developed by Kato, *et al* [3].

Frameless Ultrasonic

Given the particular requirements of the patient-location problem (accurate, inexpensive, safe), ultrasound appears to provide an almost ideal frameless solution. Typical ultrasonic imaging equipment has a resolution of 0.5 mm axially and 1.2 mm laterally [6]. It is inexpensive compared with other imaging technologies. Furthermore, the side effects of long-term exposure to diagnostic levels of ultrasonic radiation are believed to pose minimal hazard to the patient and the surgeon, provided that the spatial-peak, temporal-average (SPTA) of the dosage is less than 100 mW/ cm^2 [7].

Implementation

Our system locates the patient by identifying and tracking "corners," places where the curvature[1] reaches a local extremum[2], on a designated bone in the patient. We track features on the bone because it is a rigid, and therefore does not deform when the patient moves, and because bone contours are easily identified in ultrasonic images. Corners are identified in a set of cross-sectional images of the bone. These corners are matched to a reference model of the bone built off-line from a CT scan. CT provides precise, high-density data, but it is expensive and not real time. Thus, we use it to create a reference off-line that ultrasonic images,

1. *Curvature* at a point on a curve is defined as the reciprocal of the radius of the circle that best matches the curve near that point.

2. An *extremum* is a minimum or a maximum.

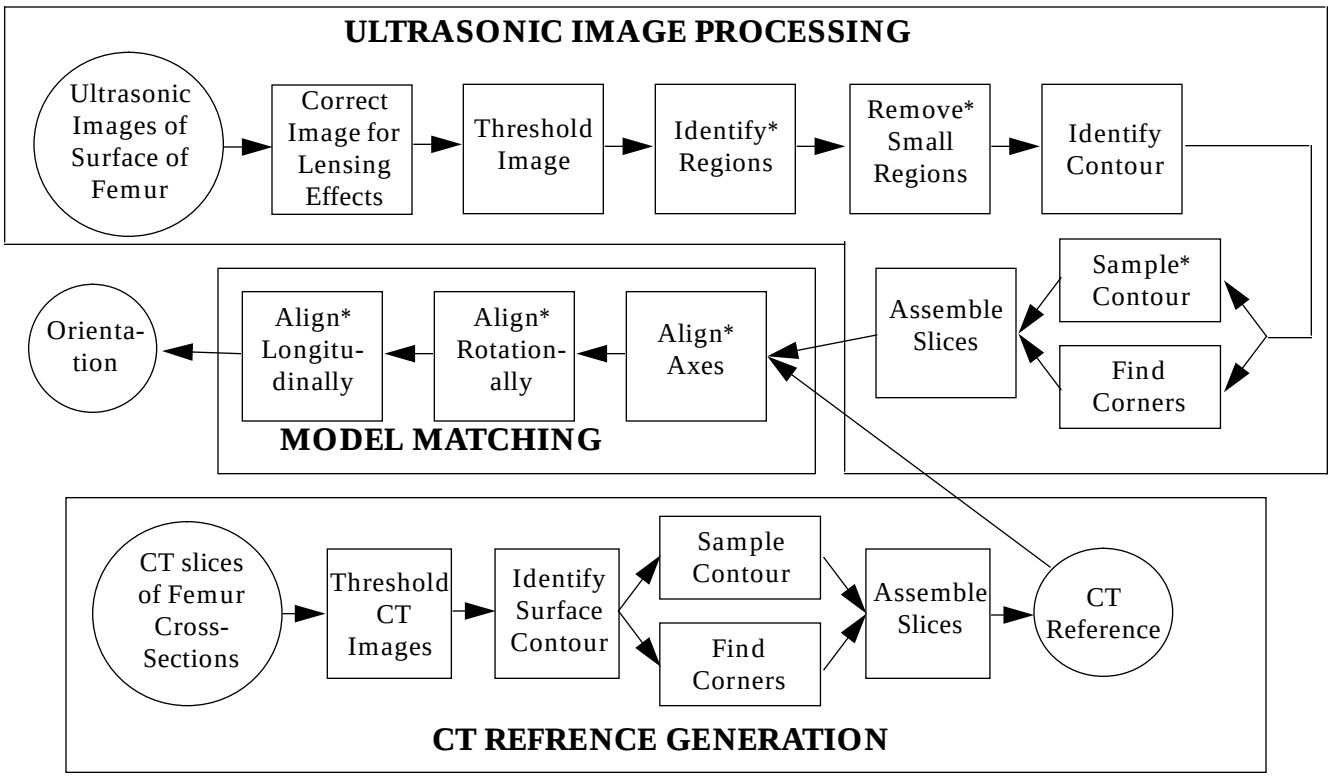

Figure 1. Block diagram of the current system. Blocks are operations performed on some data. Circles are input to or output from major modules. Many algorithms of the present system can be incorporated in a real-time system. Starred blocks will be new in an automated system. In a future system, the CT reference will be computed off-line before surgery. Ultrasonic image processing and model matching will take place in real-time.

processed in real-time, can be matched to. The matching process gives the orientation of and direction to the bone relative to the sensor path. The distance from the sensor path to the bone can be computed from the ultrasonic images themselves. Since the location and orientation of the sensor relative to the operating room is known, both the location and orientation of the bone relative to the operating room can thus be calculated.

We are presently working with a human femur. None of the techniques we use are specific to the femur, but the search procedure we describe is most naturally applicable to elongated bones. We plan to generalize our techniques to include other types of bones correspondingly naturally.

Our current system tracks off-line; we are using it to identify critical features and techniques for eventual real-time tracking. In particular, we are assessing the usefulness of "corners" in tracking; our system is unique among the ultrasonic patient registration systems we know of in tracking these corners.

A sequence of ultrasonic images taken over the length of the femur is processed to locate corners, then the corners from the individual ultrasonic images are assembled to form a surface picture of the femur. The corners in this picture are then matched to the corners in the reference model. The matching algorithm calculates the location and orientation of the femur. The accuracy of registration of ultrasonic data with a CT reference is currently being assessed. We expect to identify which regions of the femur contain the best corners to track.

A block diagram of the present (off-line tracking) system can be found in figure 1.

Related Ultrasonic Techniques

Our system is not the first to use ultrasound in surgical positioning. Other systems use non-imaging ultrasonic rangefinders to locate the position of a device relative to the patient [4], or use intra-operative ultrasonic imaging to verify the position of a target once registration has been done by other means [5,8]. To our knowledge, no system other than our own uses ultrasonic imaging to do the registration *per se*.

Method

Apparatus

We use a Niemark 9 Ultrasonic Imaging machine and a 5 MHz linear array probe 4 cm in width with a transducer density of 16 crystals/cm. The ultrasonic images[3] are saved to a video tape in VHS format. Off-line, we digitize the images with a Matrox video board in an IBM PC.

We built a femur phantom from a purchased femur and a 45 cm long piece of plastic pipe 15.25 cm (6 in) in diameter, sealed at one end. A flat 6 cm wide was machined on one side. The flat is ruled with lines perpendicular to the cylinder axis every 5 mm for positioning the sensor. The flat permits intimate contact between the rigid transducer and the rigid

3. We use the term "images" to refer to the highly-processed B-scans the Niemark 9 machine reports in 2D Real-Time imaging mode.

tube; intimate contact is impossible on the curved areas of the tube. The phantom is filled with water, and the femur is immersed in it with the distal end at the bottom and the proximal end above the waterline. The distal end of the femur pivots about a plastic rod centered on the cylinder bottom inserted into a hole drilled in the bone. Rotating the femur has the same effect as rotating the ultrasound sensor about the bone by the same angle in the opposite direction. Thus, data from all azimuths of the bone can be collected using a single flat.

Coordinate System and Position Measurement Conventions

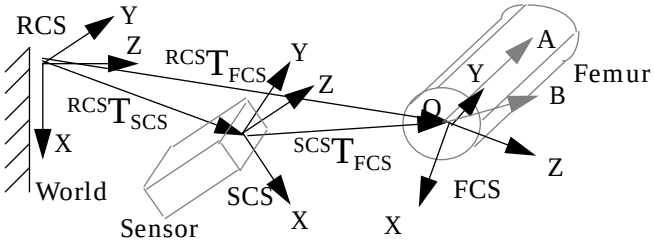

Figure 2. Coordinate systems we use. Features of the bone are located in the Femur Coordinate System (FCS). Ultrasonograms are obtained in the Sensor Coordinate System (SCS). The operating room defines the Room Coordinate System (RCS).

We define three coordinate systems for our system - the *Room Coordinate System (RCS)*, the *Femur Coordinate System (FCS)*, and the *Sensor Coordinate System (SCS)*. The FCS is attached to the femur and defines position of its features, *e.g.* corners. The SCS is attached to the ultrasonic transducer; the echograms of the bone obtained from ultrasound are in the SCS. The RCS is attached to the operating room. Our system attempts to measure the location and orientation of the FCS in the RCS. Since the transformation between the SCS and the RCS is assumed to be known, computing the position and orientation of the FCS relative to the SCS allows us to compute the position and orientation of the FCS in the RCS. The relationship between the coordinate systems are shown in figure 2.

The FCS is instantiated using two points, that we call A and B, marked on the proximal end of the femur, and by the center of the hole, that we call point O, into which the plastic rod is inserted. The locations of the three points A, B and O are measured in the RCS. O is defined to be the origin of the FCS. The Y-axis of the FCS is defined to lie along the vector from O to A, and the Z-axis along the cross-product of the vectors from O to A and O to B (see figure 3). Placing the Y-axis of the FCS close to the longitudinal axis of the bone is not only natural, but it also makes measurements related to calibration easier, since points A and B are above the water line in the phantom and their locations relative to the phantom can be easily measured.

The SCS is instantiated with its origin where the sensor begins its sweep, its Z-axis perpendicular and pointing downwards from the (flat) head of the ultrasonic sensor, its X-axis parallel to the head of the sensor and pointing left as one looks at the sensor face, and the Y-axis as the cross-product of Z and Y. The sensor travels along the Y-axis of the SCS. As with the FCS, the location of the origin and the orientations of the axes are measured in the RCS. The RCS is arbitrarily fixed in the operating room.

Experimental

Protocol and Data Format

Ultrasonic images of the femur were taken every 0.5 cm starting 10.0 cm from the bottom of the phantom and ending 4.5 cm from the top for a total of seventy-one images along Y_{FCS}. We call each group of images corresponding to a particular rotational orientation of the femur a run; a total of six runs were measured. A typical image can be seen in figure 3. The ultrasonic imaging apparatus delivers its data

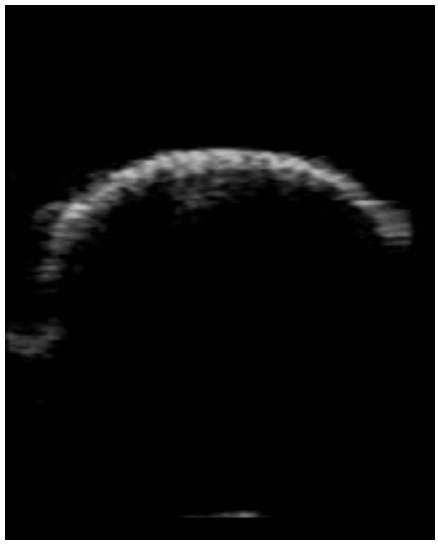

Figure 3. Typical cross-sectional image of surface of the femur obtained from ultrasound.

as an NTSC television signal in which horizontal position (pixel number) represents the width dimension across the transducer array, vertical position (line number) represents the depth dimension away from the transducer array, and black-to-white intensity represents the relative strength of the echo signal from each location.

Two of the runs are of the same nominal orientation. The replications are used to estimate measurement errors. Important sources of measurement error include: misplacement of the sensor while taking data, motion of the sensor while taking data, quantization errors in the digitization process, uncertainty in the measurements of the locations of points A and B, and rotation of the bone from its initial position during a run.

The ultrasonic array is provided with a “focal-zone” control that results in a converging beam (presumably achieved through phasing the transducer elements). The difference in curvature between the flat and the inner surface of the tube augments this effect. Consequently, there is a significant lensing effect, and the apparent widths of objects varies as a function of their distance from the sensor. We correct this effect by calibration of the sensor and phantom with an object of known width at several known depths. See figure 4 for the calibration curve. Figure 5 shows an image corrected for this effect. The corresponding correction *in vivo* would be straightforward, inasmuch as the CT data are

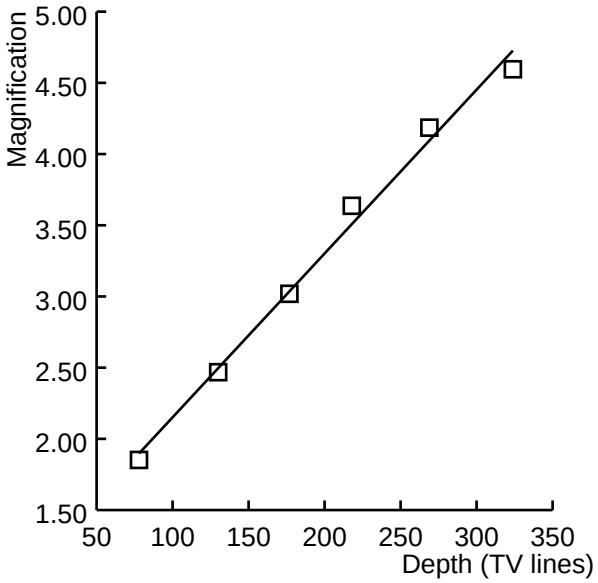

Figure 4. Calibration curve. The magnification for a given TV line is M(z)=f(z)/f(167) (167 lines = 4 cm). The width of an object at depth z in the corrected image is w/M(z), where w is its width in the uncorrected image

Figure 5. An ultrasonic image after correction for focussing effects of the transducer and lensing effects of the plastic tube. The correction is linear with depth. The widths are also scaled so that objects at TV line 167 (4cm) are the same width in the corrected image as they are in the original image.

well calibrated. Thus, features, such as the bone diameter, whose dimensions are well-known from the CT data and observable in the ultrasonic data could be used to rescale the surface contour data obtained from the ultrasonic images independent of the lensing effects occurring from different soft tissue properties, coupling effects between the transducer and skin, etc.

Corner Finding Algorithm

Corners are computed from a curve fit to the data in each image. The image is thresholded[4] to remove noise. Since the ratio of the acoustical impedance of bone to water (or soft-tissues) is large, the image of the bone/water interface is much brighter than its surroundings, and thus a threshold filter removes a large portion of the noise in the image. Comparing figure 6 with figure 5 shows the benefits of thresholding. After thresholding, the image is segmented into regions, where a region is an eight-connected[5] group of pixels of non-zero intensity. Groups that are smaller than one-tenth of the area of the largest group are removed. This eliminates the bright but localized spikes of noise that survive thresholding. Figures 7 and 8 show the groups found in figure 7 and the results of elimination. The remaining groups are assumed to be valid portions of the image of the bone. Their centroid[6] defines the observed contour of the bone. Points on this contour are obtained by searching each column of pixels (depth dimension) in the two-dimensional array that is the image from its top and bottom for the beginning and ending, respectively, of the group in that column. Then a mean depth value, weighted by intensity, is computed for the group in that column using:

$$y_{mean}(x_j) = \frac{\sum y_i I(x_j, y_i)}{\sum y_i}$$

where y_i is a discrete depth value in column j, and x_j is the corresponding horizontal position.

Figure 9 gives the intensity vs. depth profile for a typical column and indicates the point on the contour we computed. Columns that do not have any non-zero pixels (and therefore do not intersect any groups) obtain their values from their neighbors via linear interpolation. An example of a bone contour can be seen in figure 10.

The above process describes the bone contour as a discrete function $y_{mean}(x_j)$. Since $y_{mean}(x_j)$ is not very smooth, a smoothing operator, described by:

$$y(x_j) = \sum_{i=-2}^{2} w_i y_{mean}(x_{j+i}) \qquad w_i = \frac{3-j}{9}$$

4. Thresholding is a process by which all the pixels of an images are examined and those whose values fall below a designated level have their values set to zero, effectively erasing them.

5. An eight-connected group is a set of pixels such that one can trace a path from any pixel in the group to any other pixel in the group by moving horizontally, vertically or diagonally without crossing any empty pixels on a rectilinear grid. Pixels in an eight-connected group may thus touch other pixels "weakly," at corners, not just "strongly," at edges.

6. Strictly speaking, the centroid is defined as the center of mass. Naturally, we substitute intensity for mass in our calculations.

is applied to it. Then the coefficients of a continuous approximating polynomial are computed for the $y(x_j)$:

$$p(x) = \left[\sum_{i=1} a_i T_i(x)\right] - \frac{1}{2}a_0$$

$T_i(x)$ is the i-th degree Chebyshev polynomial and a_i is the weight assigned to it. (See [8] for the algorithm to compute the a_i's for a particular $y(x_j)$). The curvature of $p(x)$, called $c(x)$, and its first derivative, $c'(x)$,

$$c(x) = \frac{f''(x)}{(1 + [f'(x)]^2)^{3/2}}$$

$$c'(x) = \frac{f'''(x)\,[g(x)]^2 - 3f'(x)\,[f''(x)]\,\sqrt{g(x)}}{[g(x)]^3}$$

$$g(x) = 1 + [f'(x)]^2$$

are then calculated and used to locate the corner points, which occur at the extrema of $c(x)$. Sometimes extraneous corners will be found between legitimate ones. These are removed by scanning the corners in order from left to right. Corners that are approximately co-linear with and lie between their neighbors are removed. The results of the curve-fitting step can be seen in figure 11. Once all the corners have been found, they are positioned relative to the SCS, producing a three-dimensional image of the bone that we call its *reconstruction*. An example of a reconstruction of a run of ultrasonic data is shown in figure 12.

Analysis: Correlation with CT Reference

To compute the orientation of the femur, the system must first have a reference defining a canonical orientation. This reference is obtained from a CT scan of the femur with 5mm intervals between slices. Curve-fitting and corner-finding algorithms similar to those used for the ultrasonic images are applied to the CT scan slices. For the CT images, it is more appropriate to compute the contours in polar than in Cartesian coordinates because the cross-sections of the femur are more-or-less circular. It is also appropriate to plot the corners in the FCS instead of the SCS because the SCS is not defined for the CT reference. Figure 13 explains the algorithm used to find points on the contour and figure 14 provides an example of the results. Since curvature is a property of the shape of the femur itself and not the coordinate system used to describe it, the same corners will be found with a polar fit as with a Cartesian fit. Figure 15 displays the CT reference.

The CT reference and the US reconstruction have the same interval between slices, but in general the US cross-sections will not align exactly with the CT slices. Thus, the alignment algorithm must be able to interpolate between slices to assure a good match can be found.

Figure 6. Corrected image after thresholding. Much of the noise has been removed, but some still remains to be removed by region culling. (See figure 7.)

Figure 7. Regions in an ultrasonic image. Each region has a different color. Only the largest regions are used to establish a contour. Smaller ones are deleted (“culled”).

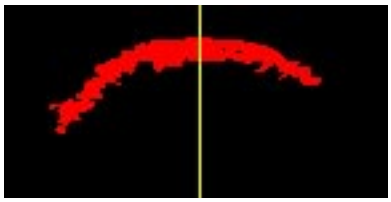

Figure 8. Largest regions in an ultrasonic image. Contours will follow the intensity-weighted mean Y values of the columns of these groups. The yellow line indicates the column in which an intensity vs. depth profile will be generated

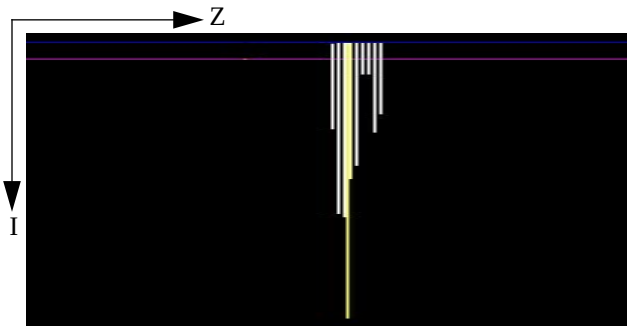

Figure 9. Intensity vs. depth graph for one column. We use an intensity-weighted mean depth of all the points in the group in this column as a point on the contour. This point is indicated by the yellow line in the graph.

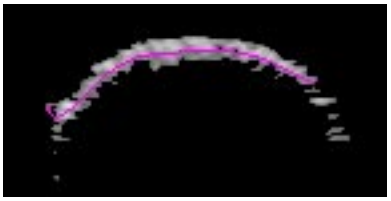

Figure 10. Discrete contours of bone in an ultrasonic image. Linear interpolation is used to obtain values between columns. These contours are sampled to obtain the triangular mesh of the ultrasonic reconstruction.

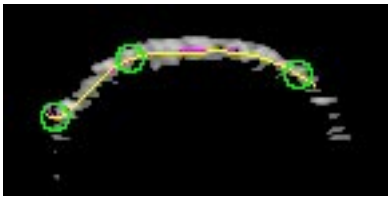

Figure 11. Contour with Chebyshev polynomial fit (yellow) to provide a smooth curve. Corners (green) are identified via the curvature of the polynomial. Curvature is inherent to the shape of the femur itself; thus, “corners” - places where the curvature achieves an extremum - are good features to track.

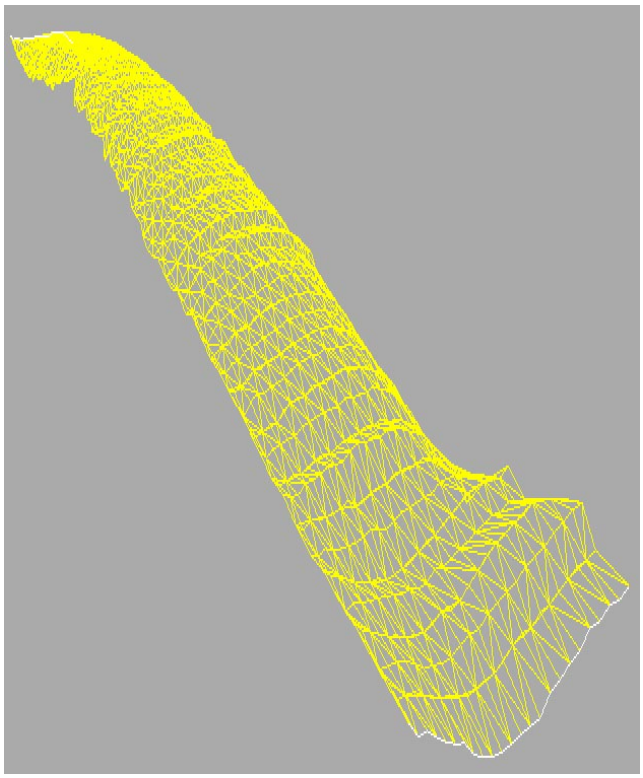

Figure 12. Ultrasonic reconstruction from contiguous ultrasonic images, rendered as a triangular mesh. The white contours are the beginning and ending of the reconstruction. For clarity, the corners are omitted from this figure.

Search Algorithm

The reconstruction of the femur is matched to the corner picture of the model using three successive searches. The first search starts from an initial guess of the rotation to align the long axes of the reconstruction and the reference; it focuses on the central portion of the bone. This portion is well approximated by a cylinder (it’s surface runs parallel to the axis through its center). The properties of this section enable us to compute an approximation to the long (Y) axes of the FCS and SCS without search (see figure 16). One of these properties is that any displacement between corresponding points (*i.e.* corners) in the ultrasonic cross-sections must be parallel to the axis running through the center of the nearly cylindrical portion. The rotation between this axis and the Y-axis of the FCS can be computed from the CT reference. Thus an approximation to the Y-axis of the reconstruction can be computed by fitting a line to several corresponding corners in the center of the run and applying the aforementioned rotation to align it with the Y-axis of the reconstruction. The properties of the nearly cylindrical section also allow us to separate the longitudinal and rotational components of the search to align the FCS with the SCS, resulting in two tractable one-dimensional searches instead of one intractable two-dimensional search. Figures 17 and 18 illustrate this process.

After the Y-axes of the reference and the reconstruction are aligned, the corners are rotated about the now co-incident Y-axes until a good match, defined as the orientation where the error is minimal, is obtained. For this search and the two that follow, the error in the position of the bone is defined as the sum of the absolute distances between corners in the reconstruction and their nearest neighbors in the CT reference. In obtaining this approximation to rotation about the Y-axis, we use the fact that the shape of the bone changes little as one moves along the long axis of the nearly cylindrical section, and thus the center section of the picture and the center section of the model can be displaced significantly with respect to one another along the Y-axis without seriously affecting the results. The outcome of the rotational search can be seen in figure 19.

Once the Y-axes have been aligned in direction, and all three axes have been aligned rotationally, we slide the bone longitudinally until a minimum in the error computed over all the corners of the bone, not just those in the central (nearly cylindrical) section, is found (see figure 20). We then search for a more accurate matching between the corners of the picture and the corners of the model by making small adjustments in the rotations about the axes and the longitudinal displacement with a hill-climbing algorithm[9] that minimizes error over all the corners guiding the selection of moves. Once a minimum is found, the final orientation of the bone has been computed.

Since this process necessarily identifies a particular corner in the reconstruction with a particular corner in the reference, the displacement of the femur with respect to the sensor can be calculated by adding the displacement from

the origin of the SCS to any one corner of the bone to the displacement of that corner from the origin of the FCS as seen from the sensor. In practice, we would probably calculate the displacement vectors and error estimates thereof for several corners and report the average weighted by the reciprocal of the errors:

$$\hat{\vec{v}} = \frac{\sum_i \frac{\hat{\vec{v}}_i}{s_i}}{\sum_i \frac{1}{s_i}}$$

$$s_i = \sqrt{\Delta\hat{\vec{v}}_i \bullet \Delta\hat{\vec{v}}_i}$$

Conclusion

Present Work

We have constructed the CT reference model from raw CT slices and located the corners within it. We have built the ultrasonic reconstruction for a run and located its corners as well. Pending implementation of automatic orientation-measuring algorithms, we are matching the reconstruction and the CT reference by hand. Repeated manual matching experiments conservatively suggest a reproducibility on the order of about +/- 5mm displacement in the plane tangent to the local surface.

Near Future

By taking advantage of the shape of the human femur and the characteristics of ultrasonic imaging, it appears that a frameless registration system using ultrasonic imaging will be feasible. By applying suitable algorithms to extract bone contours from ultrasonic images, matching prominent features within those images (*e.g.* corners) to a high-resolution CT reference computed off-line, and by taking advantage of the elongated shape and symmetries of the femur to reduce the search space, in light of the encouraging manual results, a fully automated procedure seems entirely feasible. We are working on building the individual algorithmic components and integrating them into a functional system.

Long Term

Our long-range goal is to produce an ultrasonic real-time positioning system for use in the operating room and applicable to many different bones. We conservatively expect that automation of the matching process will improve the reproducibility of our results to the order of +/- 3mm. The current system takes a step in that direction by demonstrating the basic premise that corner matching to a CT reference can be used to compute the orientation and location of the bone. Repeated manual matching experiments suggest accuracy adequate for surgical use. We expect that automated matching will be sufficiently reproducible that the method's accuracy will not be an issue in evaluating its ultimate utility. After matching is automated, it will be appropriate to investigate the minimum number of images our system needs to successfully identify the location and orientation of the femur. We will also explore issues such as the trade-offs between slices chosen over a small longitudinal extent and many rotations vs. slices chosen over a large longitudinal extent and few rotations.

Acknowledgment

We wish to thank Robert O'Toole of Shadyside Hospital for his valuable assistance in the laboratory, and in many useful conversations.

References

[1] Paul, *et al*, A surgical robot for total hip replacement surgery, in *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, Nice, France, May 1992

[2] Tan, *et al*, A frameless stereotaxic approach to neurosurgical planning based on retrospective patient-image registration, in *J. Neurosurg* **79**: 296-303 (1993)

[3]Kato, *et al* A frameless, armless navigational system for computer-assisted surgery, in *J. Neurosurg* **74**: 845-849 (1991)

[4]Reinhardt, Horst, and Gratzl, Sonic stereometry in microsurgical procedures for deep-seated brain tumors and vascular malformations, in *Neurosurgery* **32**(1): 51-57 (1993)

[5]Friets, Strohbehn, Hatch, and Roberts, A frameless stereotaxic operating microscope for neurosurgery, in *IEEE Transactions on Biomedical Engineering* **36**(6): 608-617 (1993)

[6] Advanced Technology Laboratories, UM9 Scanhead Image Acceptance Criteria, Bothell, WA (1994)

[7] Miller, Update on the safety of diagnostic ultrasonography, in *Journal of Clinical Ultrasound* **19**(9): 531-40 (1991)

[8] Press, Teukolsky, Vettering, Flannery, Numerical Recipes in C: The Art of Scientific Computing, 2nd Ed., Cambridge University Press (1992), pp. 190-194

[9] Rich and Knight, Artificial Intelligence, 2nd Ed., McGraw-Hill, Inc. (1991), pp. 66-70

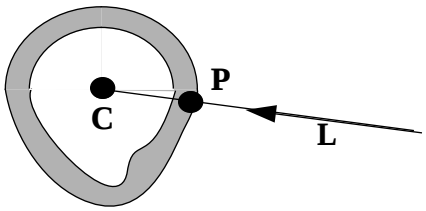

Figure 13. Illustration of contour construction algorithm for the CT reference. Line L extends a fixed distance from the center of opacity (C). L is searched inward from its farthest point. Its first intersection with the bone density map is P, a point on the outer contour. L is rotated about C at fixed angular intervals, sampling the outer contour of the density map. The samples are linearly interpolated to construct a continuous function r=f(θ) in polar coordinates with its origin at C.

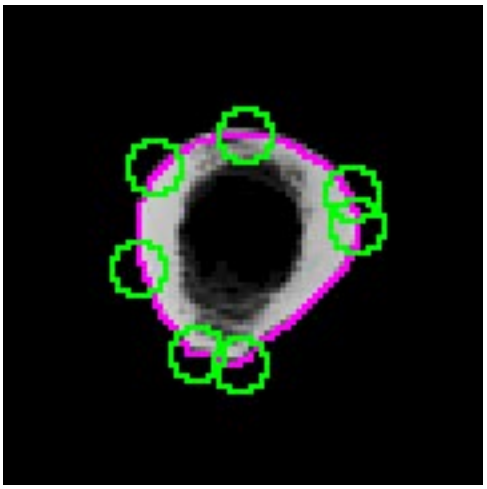

Figure 14. Results of curve-fitting and corner finding algorithms on a typical CT slice. We use algorithms similar to those used in processing the ultrasonograms as described in the text.

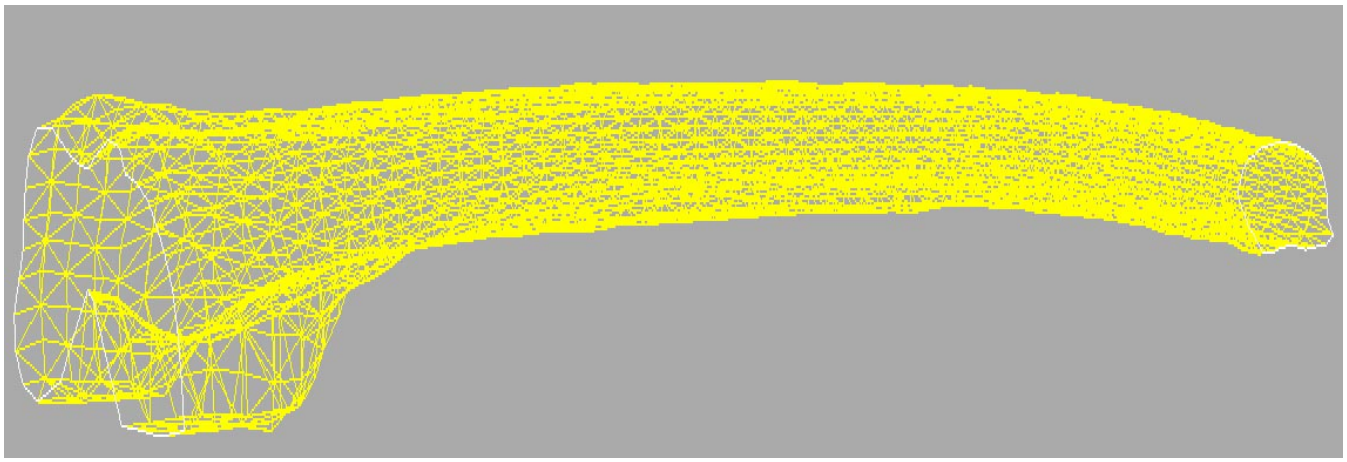

Figure 15. CT reference of femur obtained from 5mm slices of CT data. CT data are more accurate and less noisy than ultrasound, but CT is more expensive and cumbersome than ultrasound, and it is not generally a real-time technique. Thus, the CT reference is built off-line, and the real-time ultrasonograms are matched to it.

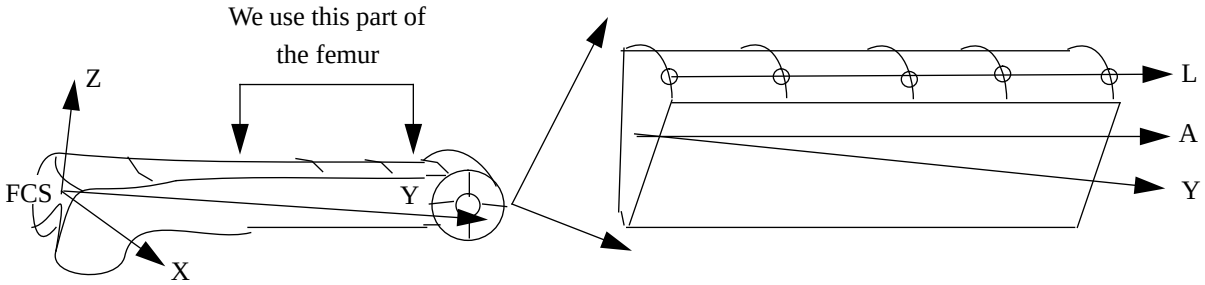

Figure 16. Computation of an approximation to Y_{FCS} that the searches uses as a starting point. We use the middle section of the femur, which is nearly cylindrical. L is a line through corresponding points (*e.g.* corners - shown as circles in the diagram) in the ultrasonic images. L is necessarily parallel to A, the axis of the nearly cylindrical section of the femur. Although A is not exactly parallel to Y_{FCS}, the angle between the two is a constant for a given femur and can be computed off-line from the CT reference. Applying this rotation to L thus gives Y_{FCS}.

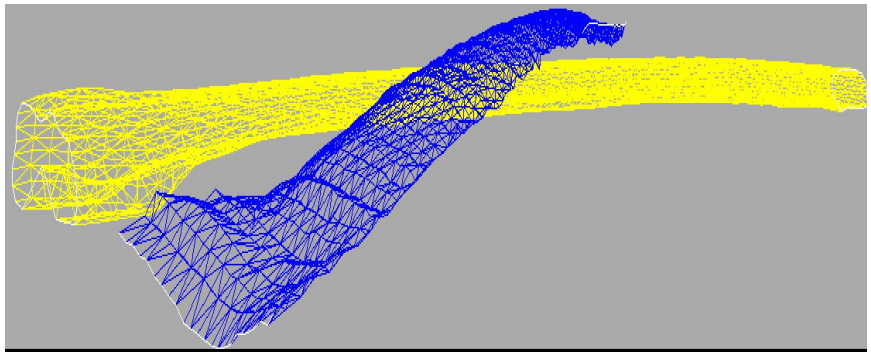

Figure 17. To compute the position and orientation of the femur with respect to the ultrasound transducer, the reconstruction (blue) must be matched to the CT reference (yellow). The normally intractable three-dimensional search can be reduced to a series of tractable one-dimensional searches. This is so because the central portion of the femur is very nearly a cylinder. An approximation to the long axes of reference and reconstruction can be found by calculating the axes of these portions.

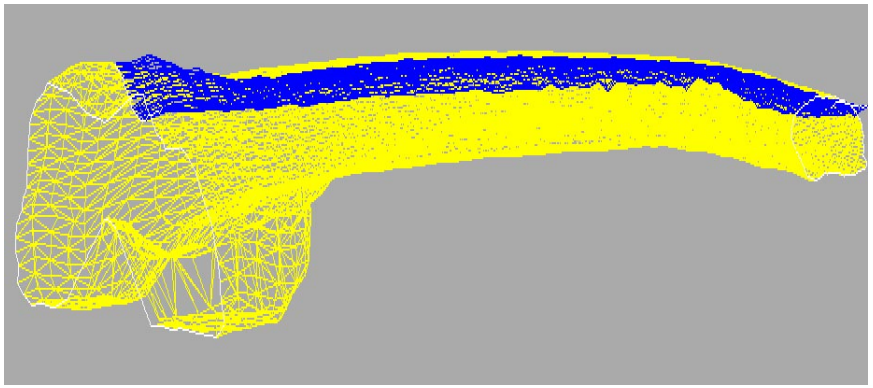

Figure 18. The major axes of the ultrasonic reconstruction and the CT reference have been aligned. This is done with the calculation outlined in figure 16. After the long axes are aligned, the reconstruction will be rotated about the CT reference until a good match is found for the middle section. Because the center section is very nearly a cylinder, any difference along the axes will not be significant; they can be resolved in the next search.

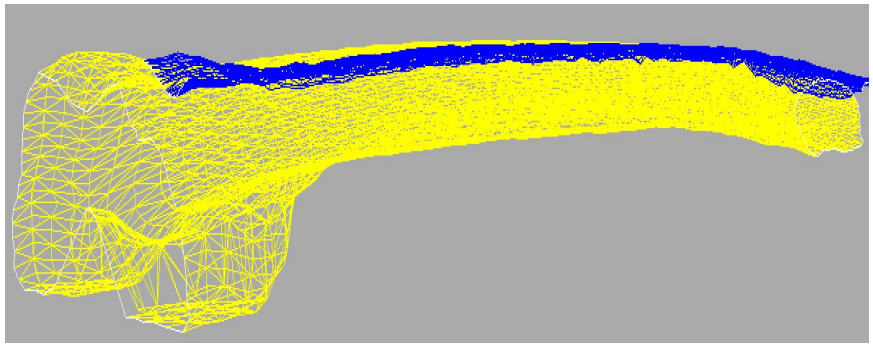

Figure 19. The ultrasonic reconstruction has been aligned axially and rotationally with the CT reference. All that remains to be done is to slide the reconstruction along the major axis until a close match is found in regions outside of the cylindrical center section.

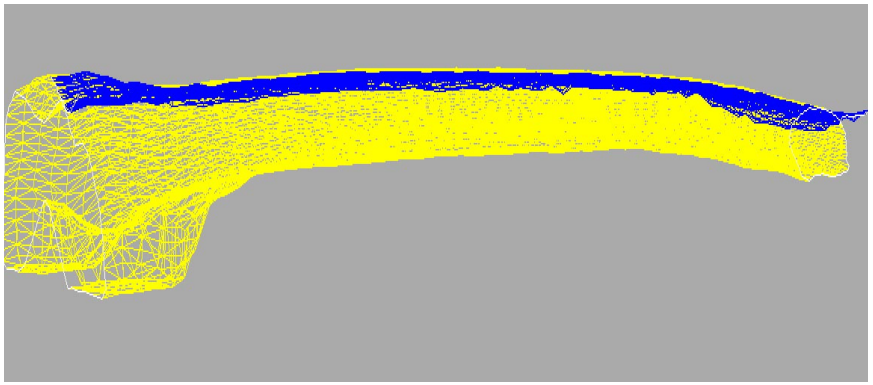

Figure 20. The ultrasonic reconstruction and the CT reference have been aligned axially, rotationally and laterally. Minor adjustments will be made to these components using a hill-climbing algorithm to further refine the match. The match gives the position and orientation of the femur with respect to the ultrasonic sensor.