

Spotlights: A Robust Method for Surface-Based Registration in Orthopedic Surgery

B. Ma, R. E. Ellis, and D. J. Fleet

Department of Computing and Information Science
Queen's University at Kingston, Ontario, Canada
ellis@cs.queensu.ca

Abstract. Fast, simple and effective registration methods are needed in a wide variety of computer-assisted surgical procedures in which readily locatable anatomical landmarks are not available. Surface-based least-squares registration methods can be used, but are susceptible to poor initial pose estimates and to error contamination during intraoperative data collection.

We have developed a fast, statistically robust method for surface-based registration during orthopedic surgery. The method, based on the iterative closest point (ICP) algorithm, fits a set of sparsely measured data points to a planar facet model. A first registration estimate is obtained by having the user contact the anatomy in a set of general anatomical regions (rather than contacting distinctive features). A small number of additional data points are acquired to refine the registration. Starting from the refined estimate, a robust scored perturbation method is used to find a better registration. This is followed by an M-estimate registration that is taken as the final registration. Simulation results show that this method is robust for data sets containing up to 25% gross outliers. The method has been tested *in vitro* on plastic bone models, where it outperformed the least-squares estimate and maintained the required $1\text{mm}/2^\circ$ accuracy. The *in vivo* use of spotlights in computer-enhanced osteotomies of the knee have confirmed the usefulness of the method.

1 Introduction

Registration of a patient to a medical image is often performed by finding a rigid transformation that minimizes the squared residual error between the surgical points and points on a model derived from a 3D medical image. One widely cited surface-based registration method is the iterative-closest-point (ICP) method of Besl and McKay [2]. Two widely acknowledged problems with ICP-like surface registration methods are the need for a good initial estimate and that minimizing the sum-of-squared residual error is optimal only when the measurement errors have Gaussian distributions. If measurements are accidentally taken far from the target anatomy, a least-squares error measure can produce poor results.

Here we present a fast, robust method for surface-based registration that provides reliability and accuracy in the presence of spurious data. The user interface presents the surgeon with a visualization of the target anatomy in which

restricted regions, termed *Spotlights*, are individually illuminated; Figure 1 shows an example surface mesh of the proximal tibia and four spotlight regions. Initial contact points, one from each region, are used to compute a registration. Additional points, collected from the exposed anatomy, are then used in a robustly perturbed M-estimation process to refine the final registration estimate. The principal contribution of this work is our consideration of the effects that



Fig. 1. CT-based surface meshes (not to scale) of a phantom tibia and vertebra, and spotlight regions for registration. The spheres mark the spotlight centers.

local-search methods have on the estimate of the registration. ICP, and previously investigated robust estimators, are susceptible to convergence on a local, non-global minimum. The method presented here is mathematically robust and was designed with the intent of avoiding false local convergence.

2 Robust Surface-Based Registration

Rigid-body surface-based registration is the process of finding a transformation from a set of measured points on the target anatomy to the model surface derived from medical images. Let $P = \{\mathbf{p}_i\}$ be a set of n surface-data points measured from the target anatomy by the surgeon, let $X = \{\mathbf{x}_i\}$ be the set of all points on the surface model, and let $T(\boldsymbol{\tau})\mathbf{z} = R(\hat{\mathbf{q}})\mathbf{z} + \mathbf{t}$ be a rigid transformation of a point \mathbf{z} with pose parameters $\boldsymbol{\tau}$. The registration goal is to find both the rigid-body transformation T and some n -element subset Y of model surface locations X to which the target anatomy locations P project under T . The anatomical points P will not in general project exactly onto Y . A least-squares solution to the surface-based registration can be stated as the minimum, over τ and $Y \subseteq X$, of

$$F_2(\tau, Y) = \sum_{i=1}^n \|\mathbf{y}_i - T(\tau)\mathbf{p}_i\|^2 \quad (1)$$

where $\mathbf{y}_i \in Y$. In the general case this is a non-convex minimization problem with multiple local minima.

Statisticians have long been aware of the need for robust methods of parameter estimation [9,14]. Robust methods have been applied widely in the

computer vision community for many estimation problems, including pose estimation [6,10,11] which is mathematically similar to the registration problem. Grimson *et al.* [5] performed point-to-surface matching by progressively refining the registration using a series of objective functions. Although it was not stated explicitly, their final objective function is the Huber [9] estimate.

Many robust estimation techniques use *M-estimation*, in which the L_2 norm in (1) is replaced with a robust norm to yield an objective function of the form

$$F_M(\tau, Y) = \sum_{i=1}^n \rho(\mathbf{y}_i - T(\tau)\mathbf{p}_i; \sigma) \quad (2)$$

where $\rho(r; \sigma)$ is the robust norm applied to the residual r , and σ is a scale parameter that depends on the form of the expected error distribution. One robust estimator that has reportedly provided good performance on 3D range data [13] is the Tukey biweight:

$$\rho(z; \sigma) = \begin{cases} \frac{\sigma^2}{2} (1 - (1 - \frac{z^2}{\sigma^2}))^3 & \text{if } z \leq |\sigma| \\ \sigma^2/2 & \text{otherwise} \end{cases} \quad (3)$$

We used ranked scores of perturbations of the least-squares estimator to improve the least-squares surface registration. We then used a robust M-estimator version of ICP to refine the registration further. These estimators form the basis of a fast and accurate method for surface-based registration.

The main stages of our surface-based registration method are:

Spotlight: Surface data are gathered intraoperatively. The surgeon contacts points on the exposed anatomical regions that correspond to the spotlights shown on the model.

Centroid: The initial contact points are first matched to the spotlight *centroids* on the model using a simple least-squares minimization method [8].

Patch: The initial contact points are then matched to the spotlight *surface regions* on the model, using a least-squares ICP method.

Refinement: The surgeon then contacts another set of points on the exposed anatomical region. These locations should be chosen to cover the anatomy that will be involved in the image-guided surgery, and should provide sufficient translational and rotational constraints on the registration.

Perturbation: The initial registration is repeatedly perturbed, and residuals are calculated. The perturbation with the largest number of residuals that are all less than a user-supplied threshold is taken as the best registration estimate.

Final: The perturbation registration estimate is refined by use of an ICP algorithm that incorporates the robust Tukey-biweight M-estimator.

Each iteration of the ICP algorithm actually involves two estimation steps: given a registration estimate, one needs to find the set of closest points on the surface to the transformed data points. It is important that this search for the closest points be fast because it is one of the most computationally demanding steps of the algorithm. From these closest points on the surface, one then needs to update the registration estimate.

2.1 Finding Nearest Neighbors on a Surface Model

Given a transformation, at each iteration of the algorithm the ICP method requires solution of the nearest-neighbor problem: For each point \mathbf{p}_i and a model X , the point \mathbf{x}_i in X that is nearest to \mathbf{p}_i under the transformation must be found. If X is a triangulated surface mesh then the facet containing \mathbf{x}_i must be found so that \mathbf{x}_i can be calculated. Exhaustive search over all facets is impractical for large models.

Although heuristics have been proposed for finding the facet containing the nearest point [1,12], we can guarantee that a nearest neighbor is always found. For each model facet, we precompute the centroid and the largest centroid-to-vertex distance and record the largest centroid-to-vertex distance found over all facets, R , which is used to build a k-D tree from the facet centroids. To find a nearest neighbor, we first find the nearest centroid and compute its distance r . We then use the region search algorithm of [1] to find all facet centroids within the sphere of radius $R + r$ centered on the data point. Finally, we exhaustively search all returned facets to find the true nearest neighbor.

This algorithm can return a large number of facets, especially when the datum \mathbf{p}_i is very far away from the model or if the model contains some unusually large facets. The requirement can be relaxed by limiting the number of returned facets to a fixed number. Compared to exhaustive search, reductions in computation time of more than two orders of magnitude were observed for models that contained tens of thousands of facets.

2.2 Robust Registration Estimation

A robust version of ICP was produced by modifying the process of updating the registration. This requires a solution to the absolute orientation problem, for which Horn's method provides a common least-squares solution.

To obtain an M-estimate of absolute orientation, we use an iteratively reweighted least-squares modification [7,6] of Horn's method [8]. The scale parameter σ in Equation (3), is estimated, following Rousseeuw [14], as a function of the parameters τ by using the median of absolute deviations of the residuals: $\mathbf{r}_i(\tau) = \mathbf{y}_i - T(\tau)\mathbf{p}_i$:

$$\sigma = 1.4826 \operatorname{median}_{i=1..n} \left(\|\mathbf{r}_i(\tau)\| - \operatorname{median}_{j=1..n} \|\mathbf{r}_j(\tau)\| \right) \quad (4)$$

2.3 Refinement of Registration Using Perturbation

Even when started from a reasonable spotlight estimate, traditional ICP and simplistic robust variants suffer from "trapping" by converging to a local non-global minimum of the registration parameters. To alleviate trapping we locally explore the parameter space to find a better estimate, as suggested by Grimson *et al.* [5]. Our implementation explored the effect of rotational parameters on the estimate by assessing sixty-four uniformly distributed axes of rotation. The surgical data were rotated, about their mutual centroid, around each of these

axes by ± 3 degrees and residual errors were calculated. For each of the 128 rotations, if half of the transformed surgical data had residuals that were less than a provided threshold value (1 mm) then the rotation and maximum residual were noted.

The perturbation that produced the minimum maximum residual for at least half the surgical data was provided as the initial estimate for the final robust iterative calculation of the registration transformation.

3 *In Vitro* Experiments

One application of computer-enhanced orthopedic surgery is to the high tibial osteotomy, for which the surgical exposure is limited to the anterolateral aspect of the proximal tibia. The only distinctive landmarks are the tibial tubercle (which is concealed by the patellar tendon) and the fibular head (which is mobilized from the tibia by osteoclasis). Spotlight registration was examined as an alternative to fiducial registration, which is very accurate but invasive.

As a standard comparison, we also considered the procedure of pedicle-screw insertion into a lumbar vertebra, for which the posterior aspect of the ends of the transverse and superior articular processes provide prominent landmarks.

3.1 Materials and Preparation

A plastic tibia and L1-vertebra (Sawbones, Vashon, WA) were instrumented with three titanium-alloy anchor screws of 1.9 mm diameter (Wright Medical Devices) that acted as fiducial markers. The phantoms were imaged by computed tomography, and decimated isosurface models were produced. The tibial mesh contained 34,537 vertices and 68,564 triangular faces, and the vertebral mesh contained 27,096 vertices and 54,904 faces. The fiducial locations in CT coordinates were found using a previously validated center-of-mass calculation [3] and the locations were verified with Roentgen stereogrammetric analysis.

The phantoms were fixed in frames and the fiducial markers were contacted using a six-degree-of-freedom mechanical pointer (FARO Technologies, Lake Mary, FL) to obtain a registration that bore a known error to ground truth [4]. For the tibial phantom, twelve 10 mm \times 10 mm squares were drawn on the surface in the area of typical surgical exposure and 100 data points were collected for each square, attempting to keep the datum spacing as uniform as possible. For the vertebral phantom, eight 8 mm \times 8 mm squares were drawn on the surface and data were collected as for the tibial phantom. Four spotlights of radius 10 mm were sampled with 100 data points each from the tibial phantom, and similarly for the vertebral phantom (with spotlights of radius 7.5 mm). Figure 1 shows the 3D tibial model and the spotlight locations.

3.2 Methods

A data set consisted of one point from each spotlight and one refinement point from each square on the surface, yielding a total of sixteen points for the tibial

phantom and twelve points for the vertebral phantom. One thousand sets were randomly selected and assessed by six different methods. For all but Method 1, all data in the set were matched to the entire isosurface:

1. Paired-point least-squares registration of initial data to spotlight centers.
2. Scored perturbation registration, starting from the estimate of Method 1.
3. Robust Tukey-estimator registration, starting from the estimate of Method 2.
4. ICP least-squares registration, starting from the estimate of Method 1.
5. Robust Tukey-estimator registration, starting from the estimate of Method 1.
6. Robust Tukey-estimator registration, starting from the estimate of Method 4.

The purpose of Methods 1, 2, and 3 are to provide an overall estimate of robust registration. Method 4 is the traditional ICP registration, to which robust estimates can be compared (it also acts as an initial estimate for a robust estimates). Method 5 is a robust M-estimator started from a naive initial estimate, and Method 6 is a robust M-estimator started from an ICP estimate.

Traditionally the results of a registration with parameters τ are reported in terms of the root-mean-square of the residual errors between the data P and the nearest points $Y \subseteq X$ derived from the model points X . We have previously shown [4] that this fails to describe the errors arising from incorrect estimates of the rotational parameters, so we used an axis-angle decomposition for analysis.

Suppose that the ideal transformation is $T_I(\tau)$. One can form the *residual* transformation between a given $T(\tau)$ and the ideal $T_I(\tau)$ as

$$D(\tau) = T(\tau)T_I(\tau)^{-1}$$

The matrix R_D of the transformation $D(\tau)$ is a rotation about an axis k by an angle θ . This angle is the angular error of the given $T(\tau)$ and is important because the angular error produces an increasingly large positional error of a transformed point as the point is increasingly far from the region from which the registration was derived. By comparison the translational error, which is the translational component of $D(\tau)$, is constant for all points.

To compare the results of surface registrations to the fiducial registration, the surface registration was applied to the measured location of fiducial marker nearest to the spotlight centroid. The distance between the transformed marker and the CT coordinate of the fiducial location was then calculated. The ideal registration transformation is unknown, so $T_I(\tau)$ was taken to be a registration derived from the fiducial markers (which were adjacent to the spotlight regions).

3.3 Results

The experiments produced an ensemble of 1,000 registrations for analysis, which represent a sampling of how the spotlight registration to an anatomical region might perform in practice. For each registration in this ensemble we calculated the angular error as the rotational difference between the sample registration and the fiducial registration. Histograms of the registration results of each method were produced; results for the tibial phantom are shown in Figure 2.

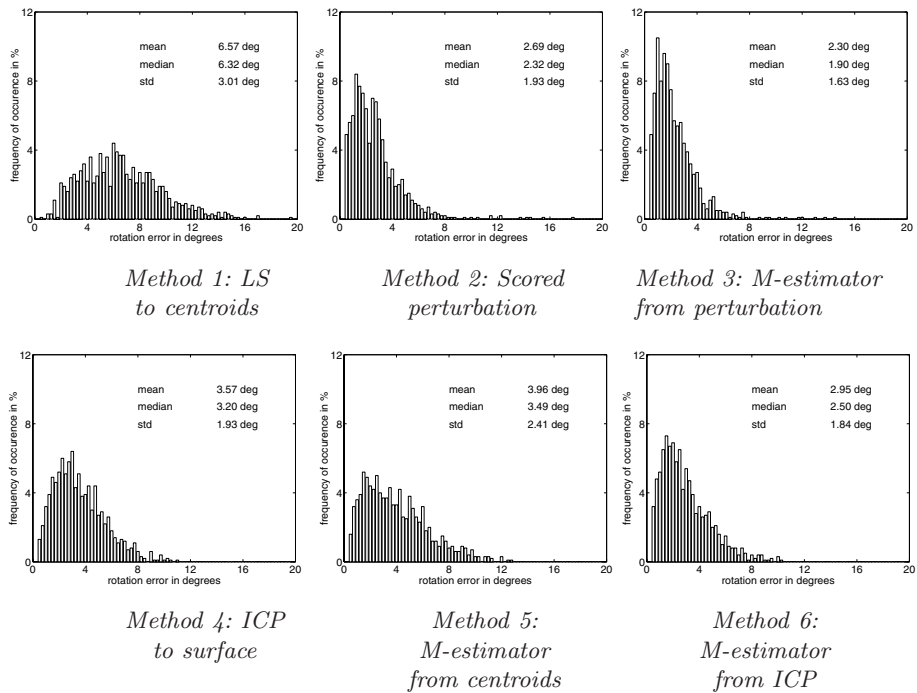


Fig. 2. Tibial rotational errors from 1000 sets of physical surface data, 16 points per set. Rotational error was calculated as the maximum expected feasible deviation of nearby fiducial points.

4 An *In Vivo* Pilot Clinical Study

Spotlight registration has been conducted on six patients in Kingston General Hospital. Each patient presented with osteoarthritis confined to the medial tibiofemoral compartment and was deemed appropriate for high tibial osteotomy. Five of the six patients were instrumented with the type of fiducial markers used in the *in vitro* study. In each case the process of drilling 4mm Kirschner guide wires for a modified Coventry procedure was performed with the spotlight registration. Registration was validated visually by contacting bony surfaces both within and outside the spotlight regions, and by contacting the fiducial markers when they were present and unmoved by dissection.

The ultimate use of registration is in providing an appropriate treatment, so a standard outcome measure was used. Postoperative A/P radiographs were measured to determine the radiographic angle between the tibial plateau and the tibial shaft. From this angle was subtracted the intended correction angle. The resulting correction errors are tabulated in Table 1.

Table 1. Osteotomy correction errors arising from computer-enhanced surgery with spotlight registration.

	Patient Number					
Valgus correction error (degrees)	1	2	3	4	5	6
	+1.5°	-1.5°	-1.5°	-1.5°	+0.5°	-1.0°

5 Discussion

The *in vitro* results for the tibial phantom demonstrate the utility of robust registration, as well as the sensitivity of robust estimators to the choice of initial estimate. The final M-estimates of the registration had a median rotational error that was about 60% of the conventional ICP estimate. However, if the M-estimator was started from the ICP registrations then the error was a little less than 80% of the ICP estimate, which is significantly different from our method (#3) and from the ICP method (#4). Naive use of the M-estimator, starting it from a closed-form registration to the spotlight centroids, produced registrations for which the median rotational error was almost 10% larger than for ICP.

The *in vitro* results for the vertebral phantom showed that ICP produced accurate registrations, with a median rotational error of slightly more than 2°. With our robust method, starting the M-estimator from ICP produced a stable result for which the median error was about 10% less than that of ICP. This suggests that, for surfaces with numerous distinctive features, traditional least-squares surface registration may be adequate.

The *in vivo* results measure the accuracy of the entire process of computer-enhanced surgery: CT, isosurface extraction, decimation, computer-based planning, registration, the physical processes of resection and reduction, and angular measurement from plain radiographs. The maximum error of 2° in the pilot study compares very favorably with the results by traditional methods.

The main contributions of this work are the development of an intraoperative data-collection scheme that is easy to use, and the implementation of a pair of robust statistical methods for estimating 3D surface registration. The methods have been tested extensively in the laboratory and have been used in early clinical trials. The registration codes run in a few seconds on common UNIX workstations.

Robust statistical methods are important for registration because they provide a sound mathematical basis for attenuating the influence of spurious data. For intraoperative use we suggest that they are superior to manual editing of the data and to *ad hoc* methods of attenuation. However, the implementation of robust methods requires care: robust methods, like nonlinear least-squares methods, can converge to local non-global minima.

Robust surface registration is potentially useful in computer-enhanced surgery. The local nature of the search still leaves the method subject to “trap-

ping”, and we recommend that such methods continue to undergo visual verification by the surgeon until validated global registration methods are devised.

Acknowledgments

This research was supported in part by Communications and Information Technology Ontario, the Institute for Robotics and Intelligent Systems, the Natural Sciences and Engineering Research Council, and by an Alfred P. Sloan Research Fellowship to DJF. Interventional procedures were performed in Kingston General Hospital by Dr. Paul Fenton, Dr. Mark Harrison, and Dr. John Rudan.

References

1. J. L. Bentley. “Multidimensional binary search trees used for associative searching”. *Communications of the ACM*, 18(9):509–517, September 1975. 939
2. P. J. Besl and N. D. McKay. “A method for registration of 3-d shapes”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992. 936
3. R. E. Ellis, S. Toksvig-Larsen, M. Marcacci, D. Caramella, and M. Fadda. “Use of a biocompatible fiducial marker in evaluating the accuracy of ct image registration”. *Investigative Radiology*, 31(10):658–667, October 1996. 940
4. R. E. Ellis, D. J. Fleet, J. T. Bryant, J. Rudan, and P. Fenton. “A method for evaluating ct-based surgical registration”. In J. Troccaz, E. Grimson, and R. Mösges, editors, *CVRMed-MRCAS’97*, pages 141–150. Springer-Verlag, March 1997. 940, 941
5. W. E. L. Grimson, G. J. Ettinger, S. J. White, P. L. Gleason, T. Lozano-Pérez, W. M. W. III, and R. Kikinis. “Evaluating and validating an automated registration system for enhanced reality visualization in surgery”. In N. Ayache, editor, *Computer Vision, Virtual Reality and Robotics in Medicine*, pages 3–12. Springer-Verlag, April 1995. 938, 939
6. R. M. Haralick, H. Joo, C. Lee, X. Zhuang, V. G. Vaidya, and M. B. Kim. “Pose estimation from corresponding point data”. In H. Freeman, editor, *Machine Vision for Inspection and Measurement*. Academic Press, Inc., 1989, pages 1–84. 938, 939
7. D. C. Hoaglin, F. Mosteller, and J. W. Tukey, editors. *Understanding Robust and Exploratory Data Analysis*. John Wiley & Sons, 1983. 939
8. B. K. P. Horn. “Closed-form solution of absolute orientation using unit quaternions”. *Journal of the Optical Society of America A*, 4(4):629–642, April 1987. 938, 939
9. P. J. Huber. *Robust Statistics*. John Wiley & Sons, 1981. 937, 938
10. R. Kumar and A. R. Hanson. “Analysis of different robust methods for pose refinement”. In *International Workshop on Robust Computer Vision*, 1990. 938
11. T. Masuda and N. Yokoya. “A robust method for registration and segmentation of multiple range images”. *Computer Vision and Image Understanding*, 61(3):295–307, May 1995. 938
12. C. R. Maurer, J. M. Fitzpatrick, R. J. Maciunas, and G. S. Allen. “Registration of 3-d images using weighted geometrical features”. *IEEE Transactions on Medical Imaging*, 15(6):836–849, December 1996. 939

13. M. J. Mirza and K. L. Boyer. "Performance evaluation of a class of m-estimators for surface parameter estimation in noisy range data". *IEEE Transactions on Robotics and Automation*, 9(1):75–85, February 1993. 938
14. P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, 1987. 937, 939