

Matching 3-D Anatomical Surfaces with Non-Rigid Volumetric Deformations

Richard Szeliski

Digital Equipment Corporation,
Cambridge Research Lab,
One Kendall Square, Bldg. 700,
Cambridge, MA 02139

Stéphane Lavallée

TIMB - TIMC - IMAG,
Faculté de Médecine de Grenoble,
38 700 La Tronche, France

Abstract

This paper presents a new method for determining the minimal non-rigid deformation between two 3-D surfaces, such as those which describe anatomical structures in 3-D medical images. Although we match surfaces, we represent the deformation as a volumetric transformation. Our method performs a least squares minimization of the distance between the two surfaces of interest. To quickly and accurately compute distances between points on the two surfaces, we use a precomputed distance map represented using an *octree spline* whose resolution increases near the surface. To quickly and robustly compute the deformation, we use a volumetric spline to model the deformation function. We present experimental results on both synthetic and real 3-D surfaces.

Introduction

The matching of 3-D anatomical surfaces between anatomical atlases and patient data, or between different patient data sets, is an important element of 3-D medical image analysis and quantification. Matching between atlases and patient data enables more accurate and reliable segmentation and the functional labeling of medical images, as well as multimodality data registration and integration. In computer vision, this problem corresponds to finding the non-rigid deformation between two surfaces, with applications to model-based object recognition and deformable object tracking.

In previous work [1], we developed a fast and accurate technique for determining the rigid transformation between two surfaces, and also between a 3-D surface and its 2-D projections. In this paper, we extend our technique to recover smooth non-rigid deformations between 3-D surfaces. Our approach is based on describing the deformation as a warping of the space containing one of the surfaces. In particular, we use a multiresolution warp or displacement field based on concepts from free-form deformations [2]. Our approach enables us to locally adapt the resolution of the deformation field to bring the two surfaces into registration, while maintaining smoothness and avoiding unnecessary computation. The result is a rapid and efficient registration algorithm which does not require the extraction of features from the

two surfaces (see [3] for more details on our algorithm and results).

The main application of our technique is model-based segmentation of 3-D medical images. Segmenting medical structures is important both in medical diagnosis and as a component of computer assisted medical interventions [4]. While unsupervised segmentation based on purely local operators can be very difficult, the *a priori* knowledge contained in anatomical models can make the segmentation process easier and more robust.

A second application of our technique is in quantification of normal anatomical structures and deviations from normal (*morphometrics* [5]), e.g., studies of brain asymmetry, and in deviation from self over time, e.g., changes in liver tumors. In some cases, a volume registration between a patient data set (e.g., a 3D MRI data set) and a model (e.g., a brain atlas) is made possible by the existence of some common reference surfaces in both data sets (e.g., the surface of brain ventricles). This registration can be used to *infer* the location of specific features (e.g., thalamus brain nuclei) in the patient data set. Such an inference would not be possible if the elastic registration was a surface deformation instead of a volume deformation.

Previous work

The registration of 3-D and 2-D images has been widely studied in both medical image processing and computer vision. In medical imaging, the problem of image registration is usually solved using external fiducial markers placed on the body of the patient [6] or by interactively selecting pairs of matching points [7]. Our previous algorithm solved this problem by minimizing the sum of squared distances between the transformed points on one surface and a stationary description of the other surface [1]. It also solved the more difficult problem of registering a 3-D surface with its 2-D projections [1].

The registration of 3-D medical images under non-rigid deformations has been studied by Bajcsy and Kovacic [8] using volumetric deformations based on physical properties. Another approach from Evans et al. [9] is based on approximating a 3D warping function by 3D thin-plate splines fitted to manually matched reference points. Jacq and Roux [10] estimate a global warping function by minimizing an

distance in 3D images using genetic algorithms. In computer graphics, non-rigid deformations are widely used for modeling and animation purposes [11]. In computer vision, non-rigid deformations have been used for fitting flexible models to both image [12] and volumetric data [13]. The approach we use in this paper is based on *free-form deformations* [2], which use volumetric, 3-D tensor-product splines to describe the warping or displacement of points embedded in space.

Problem formulation

We formulate our deformable matching problem as follows. Given a *model object* and *sensed data*, estimate the transformation \mathbf{T} parameterized by a parameter vector \mathbf{p} which registers the two coordinate systems. More specifically, we assume that both data sets are surfaces, with the sensed data represented as a collection of points $\{\mathbf{q}_i, i = 1 \dots N\}$, and model surface S represented in some arbitrary way. The registration task is then to find a geometric transformation \mathbf{T} such that the transformed coordinates $\mathbf{r}_i = \mathbf{T}(\mathbf{q}_i; \mathbf{p})$ all lie on the surface S .

In practice, due to noise and the inability to perfectly register two surfaces, this condition will never be satisfied. Instead, we pose the problem as a minimization of the cost function

$$C(\mathbf{p}) = \sum_{i=1}^N \frac{1}{\sigma_i^2} [d(\mathbf{r}_i, S)]^2, \quad (1)$$

where $d(\mathbf{r}_i, S) = \min_{\mathbf{s} \in S} \|\mathbf{r}_i - \mathbf{s}\|$ is the minimum Euclidean distance from the point \mathbf{r}_i to S and σ_i^2 is the variance associated with point i [1].

To solve the minimization problem, we require three components: a suitable representation for the geometric transformation $\mathbf{T}(\mathbf{q}_i; \mathbf{p})$, an iterative minimization algorithm, and an efficient method for computing $d(\mathbf{r}, S)$ along with its gradient.

Global polynomial deformations

The simplest representation for $\mathbf{T}(\mathbf{q}_i; \mathbf{p})$ is a rigid body transformation which can be parameterized by 6 degrees of freedom (3 for the translation, and 3 for rotation.) In our previous work [1], we used the Euler angle representation for the rotation matrix \mathbf{R} . The rigid body representation is appropriate when working with rigid anatomical structures where only the *pose* of the patient is unknown.

A more general class of transformations are the affine transforms, which can be parameterized with 12 degrees of freedom. The trilinear class of transformations adds another 12 degrees of freedom for a total of 24. Finally, the quadratic class of transformations has 30 free parameters (see [3] for details). For all of these transformations, the position of the transformed point $\mathbf{r}_i = \mathbf{T}(\mathbf{q}_i; \mathbf{p})$ is a *linear* function of the parameters in \mathbf{p} , i.e., the transformation can be written as $\mathbf{r}_i = \mathbf{M}_i \mathbf{p}$ [11].

Local spline deformations

To obtain a wider range of more flexible deformations, we need to continue increasing the order of the polynomial. How-

ever, this suffers from the same problems as high degree polynomial fitting, e.g., instability and the presence of ringing. Instead, we model the deformation using a family of volumetric tensor product splines,

$$\mathbf{T}(\mathbf{q}_i; \mathbf{p}) = \sum_{j,k,l} \mathbf{d}_{jkl} B_j(x_i) B_k(y_i) B_l(z_i), \quad (2)$$

where the \mathbf{d}_{jkl} are the spline coefficients which comprise the parameter vector \mathbf{p} , and B_j , B_k , and B_l are B-spline basis functions [2]. With splines, the deformations required to bring a local area of two surfaces into registration will not affect the registration at far-away portions of the surface.

For our initial experiments, we first find a set of global transformations, starting with a rigid transformation and then fitting an affine transformation. We then estimate a local trilinear spline deformation. To better decouple local *elastic* deformations from global effects such as scale change, we use both transformations in series.

Least squares minimization

To perform the nonlinear least squares minimization, we use the Levenberg-Marquardt algorithm because of its good convergence properties [14]. Least squares techniques work well when we have many uncorrelated noisy measurements with a normal (Gaussian) distribution. In order to update the current estimate of the parameters $\mathbf{p}^{(k)}$, Levenberg-Marquardt requires the evaluation of the distance function $d(\mathbf{r}_i, S)$ along with its derivative with respect to all of the unknown parameters. Efficient techniques for computing the distance function d , as well as its spatial gradient $\mathbf{g} = \nabla_{\mathbf{r}} d$, are presented in the next section. The evaluation of the derivative involves a straightforward application of the chain rule.

Once the distance samples d_i and their derivatives $\partial d_i / \partial \mathbf{p}$ have been computed, the Levenberg-Marquardt algorithm forms the approximate Hessian matrix \mathbf{A} and the weighted error gradient vector \mathbf{b} [14], and then computes an increment $\delta \mathbf{p}$ towards the local minimum by solving

$$(\mathbf{A} + \lambda \mathbf{I}) \delta \mathbf{p}^{(k)} = \mathbf{b}, \quad (3)$$

where λ is a stabilizing factor which varies over time [14]. After setting $\mathbf{p}^{(k+1)} = \mathbf{p}^{(k)} + \Delta \mathbf{p}^{(k)}$, this process is repeated until $C(\mathbf{p})$ is below a fixed threshold, the difference between parameters $|\mathbf{p}^{(k)} - \mathbf{p}^{(k-1)}|$ at two successive iterations is below a fixed threshold, or a maximum number of iterations is reached.

When the number of parameters being estimated is reasonably small (as in global deformations), we use the Gauss-Jordan elimination algorithm [14] to solve (3). For systems with more parameters, such as local deformations, we use conjugate gradient descent [14].

Fast distances using octree splines

The method described in the previous section relies on the fast computation of the distance $d(\mathbf{r}, S)$ and its gradient. To speed up this computation, we precompute a 3-D *distance map*, which is a function that gives the minimum distance to S from any point \mathbf{r} inside a bounding volume V that encloses S [15].




Figure 1: Registration of two face data sets:

(a) model data set (george1) (b) sensor data set (heidi) (c) final deformed sensor data (d) final registered data sets

In looking for an improved trade-off between memory space, accuracy, speed of computation, and speed of construction, we developed a new kind of distance map which we call the *octree spline* [1]. The intuitive idea behind this geometrical representation is to have more detailed information (i.e., more accuracy) near the surface than far away from it. We start with the classical octree representation associated with the surface S [16] and then extend it to represent a continuous 3-D function that approximates the Euclidean distance to the surface. This representation combines advantages of adaptive spline functions and hierarchical data structures (see [1, 3] for details).

Experimental results

To determine the reliability of our global and local deformation estimates, we first performed a series of experiments on both real and synthetic surfaces under simulated (known) motion. For a given surface, we first compute the octree distance map. Next, we select a subset of the surface points (typically 5%) and transform these through the inverse of the deformation we are simulating. We then initialize the Levenberg-Marquardt algorithm with some initial rigid pose estimate and set the non-rigid parameter estimates to 0. Finally, we run the Levenberg-Marquardt algorithm in stages, first computing the best rigid match, then the best global non-rigid match, and finally the best local displacement warp.

To demonstrate the local non-rigid matching, we use two sets of range data acquired with a Cyberware laser range scanner (Figures 1a and 1b). The octree spline distance map is computed on the larger of the two data sets (Figure 1a), and the smaller of the two data sets is deformed (Figure 1b). In their initial positions, the data sets overlap by about 50% and differ in orientation by about 10° . We first perform 8 iterations of rigid matching and 8 iterations of non-rigid affine matching. We then perform 8 iterations at each of 4 levels of the local displacement spline. The finest octree spline level has $(2^4 + 1)^3 \approx 5000$ nodes for a total of about 15000 degrees of freedom. Even with this large number of parameters, the algorithm converges very quickly, because it is always in the vicinity of a good solution (a typical iteration at the finest level takes about 2 seconds). From Figure 1d, we

see that the two data sets are registered well, except for the eyebrows, which would require a more detailed deformation. We also note that the deformed face of Figure 1c resembles that of Figure 1a more than its former (undeformed) self (Figure 1b).

As a second example of our algorithm, we matched the surface of a real patient vertebra to the surface of a plastic "phantom" vertebra (both 3-D images sets were acquired with a CT scanner). Figure 2 shows the result of our matching. After affine registration, a fair amount of discrepancy remains. After the local spline registration, most of the patient vertebra (contour lines) matches the phantom model (cloud of dots), except for the tips of the vertebra which have not been pulled into registration.

Discussion and Conclusions

We have developed a technique for registering 3-D surfaces with non-rigid deformations. Our approach represents the deformation using a combination of global polynomial deformations and local displacement splines (free-form deformations). We use an algorithm which directly minimizes the squared distances between the two surfaces, rather than identifying sparse features (e.g., ridge lines or feature points [17]) and then trying to match them.

We believe that the direct matching of surfaces has better accuracy and removes the need for a feature detection stage, which may not always operate reliably. Two arguments which favored feature-based approaches in the past were computational complexity and global correspondence search. Using the octree spline distance map, the complexity of each iterative adjustment step in our algorithm is linear in the number of sensed surface points. In our approach, we use a modified gradient descent which avoids combinatorial search but only finds locally optimal matches. In practice, we have found that false local minima are usually far away from the true solution. In medical applications, *a priori* knowledge about position and shape is usually available, so that only small displacements and local deformations have to be estimated.

Our approach embeds one of the surfaces being matched into a deformable space, rather than equipping it directly




Figure 2: Registration of patient vertebra with plastic "phantom":
 (a) original image, (b) after affine registration, (c) after final local spline registration (c) selected lines in deformation spline

elastic properties. While the latter approach may be physically realistic if an anatomically and biomechanically correct model is constructed, our approach enables us to deal with arbitrary surfaces which may not even have a good connected representation. Volumetric deformations yield auxiliary information about the motion of nearby structures which did not participate in the matching, e.g., distortions performed on certain easily identifiable brain structures can be used to estimate the registration for the whole brain.

Our experiments to date have been performed on pre-rendered 3-D surfaces. We are planning to extend our algorithm to work directly with unsegmented 3-D medical images. Other areas of future investigation include an adaptive algorithm for deciding how to refine the local deformation spline, the choice of the order of the deformation spline, the choice of various smoothness constraints (regularization).

References

S. Lavallée, R. Szeliski, and L. Brunie. Matching 3-D smooth surfaces with their 2-D projections using 3-D distance maps. In *SPIE Vol. 1570 Geometric Methods in Computer Vision*, pages 322–336, San Diego, CA, July 1991.

T.W. Sederberg and S.R. Parry. Free-form deformations of solid geometric models. *Computer Graphics (SIGGRAPH'86)*, 20(4):151–160, 1986.

R. Szeliski and S. Lavallée. Matching 3-D anatomical surfaces with non-rigid deformations. In *SPIE Vol. 2031 Geometric Methods in Computer Vision II*, pages 306–315, San Diego, July 1993. Society of Photo-Optical Instrumentation Engineers.

S. Lavallée, L. Brunie, B. Mazier, and P. Cinquin. Matching of medical images for computer and robot assisted surgery. In *IEEE EMBS Conference*, pages 39–40, Orlando, Florida, November 1991.

F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.

B.A. Kall, P.J. Kelly, and S.J. Goerss. Comprehensive computer-assisted data collection treatment planning and in-

teractive surgery. In *SPIE, Medical Imaging, Vol. 767*, pages 27–35, 1987.

- [7] C. Schiers, U. Tiede, and K.H. Hohne. Interactive 3D registration of image volumes from different sources. In H.U. Lemke, editor, *Computer Assisted Radiology, CAR 89*, pages 667–669, Berlin, June 1989. Springer-Verlag.
- [8] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.
- [9] A. C. Evans, W. Dai, L. Collins, P. Neelin, and S. Marett. Warping of a computerized 3-d atlas to match brain image volumes for quantitative neuroanatomical and functional analysis. In *SPIE Vol. 1445, Medical Imaging V*, pages 236–247, 1991.
- [10] J. J. Jacq and C. Roux. Automatic registration of 3D images using a simple genetic algorithm with a stochastic performance function. In *IEEE Engineering Medicine Biology Society (EMBS)*, pages 126–127, San Diego, October 1993.
- [11] A. Witkin and W. Welch. Fast animation and control of nonrigid structures. *Computer Graphics (SIGGRAPH'90)*, 24(4):243–252, August 1990.
- [12] D. Terzopoulos, A. Witkin, and M. Kass. Constraints on deformable models: Recovering 3D shape and nonrigid motion. *Artificial Intelligence*, 36:91–123, 1988.
- [13] B. Horowitz and A. Pentland. Recovery of non-rigid motion and structure. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'91)*, pages 325–330, Maui, Hawaii, June 1991. IEEE Computer Society Press.
- [14] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, second edition, 1992.
- [15] P.-E. Danielson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.
- [16] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading, Massachusetts, 1989.
- [17] A. Guézic and N. Ayache. Smoothing and matching of 3-d space curves. In G. Sandini, editor, *Second European Conference on Computer Vision (ECCV'92)*, pages 620–629, Santa Margherita, Italy, May 1992. Springer Verlag, LNCS Series Vol. 588.