

Matching 3-D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines

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

Stéphane Lavallée
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 second octree spline to model the deformation function. The coarsest level of the deformation encodes the global (e.g., affine) transformation between the two surfaces, while finer levels encode smooth local displacements which bring the two surfaces into closer registration. We present experimental results on both synthetic and real 3-D surfaces.

1 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 [20, 9], 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 [28], octree splines [20], and hierarchical basis functions [30]. Our approach en-

ables 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 (manual or automatic) of features on the two surfaces, and which can work on arbitrarily shaped surfaces with highly complicated deformations.

The main intended 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 [18, 19]. 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. Such anatomical models can either be digitized from textbooks, or they can be built by elastically registering data sets from many patients.

A second application of our technique is in quantification of normal anatomical structures and deviations from normal (*morphometrics* [6]), 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.

Another application of our technique is the calibration of non-linear distortion in medical imaging devices such as MRI through elastic registration with CT images which do not have such distortions. Our technique can also be used in other applications of deformable model-based vision, e.g., in the tracking of deformable objects such as the beating heart [15] or animate motion [32]. Another novel application could be in performing *morphing* [3] between 3-D data sets for computer animation effects.

We begin the paper in section 2 with a review of related work in the fields of medical image registration, computer

vision, and computer graphics. In section 3, we present our formulation of the problem. In section 4, we present the Levenberg-Marquardt iterative nonlinear least-squares algorithm which we use to compute the deformation estimate. Section 5 presents the *octree spline* 3-D distance map we use to rapidly compute minimal distances between points on the two surfaces. Section 6 describes in more detail the hierarchical octree spline used to represent the multiresolution non-rigid deformation. Section 7 presents some experimental results, with both synthetic data and real data. Finally, the algorithm, extensions of the method, and future improvements are discussed in section 8.

2 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 or by interactively selecting pairs of matching points. An automated algorithm for matching 3-D surfaces with other 3-D surfaces (such as the head skin surface) has been developed by Pelizzari [25]. 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 [20, 9] (see also [4]). It also solved the more difficult problem of registering a 3-D surface with its 2-D projections [20, 7]. Methods based on the registration of 3-D curves (*ridges*) have been proposed in [14, 23].

The registration of 3-D medical images under non-rigid deformations has been studied by Bajcsy and Kovacic [1] using volumetric deformations based on physical properties, and more recently by Christensen et al. [10]. Another approach from Evans et al. [12] is based on approximating a 3D warping function by 3D thin-plate splines fitted to manually matched reference points. Jacq and Roux [17] estimate a global warping function by minimizing an L_1 -distance in 3D images using genetic algorithms. In computer graphics, non-rigid deformations are widely used for modeling and animation purposes [2, 35]. In computer vision, non-rigid deformations have been used for fitting flexible models to both image and range data [8, 33, 32], and for fitting models to 3D volumetric data [15, 22]. The approach we use in this paper is based on *free-form deformations* [28], which use volumetric, 3-D tensor-product splines to describe the warping or displacement of points embedded in the space. For increased efficiency, we represent the deformation spline using a multiresolution representation [13, 34], based on the concept of hierarchical basis functions [36, 30]. Deformable models have also been used for the segmentation of 3-D medical data. Leitner and Cinquin [21] use a tensor-product spline surface model to segment complicated structures such as vertebra.

3 Problem formulation

We formulate our deformable matching problem as follows. Given a *model object* in a coordinate system $\text{Ref}_{\text{model}}$ and *sensed data* in a coordinate system $\text{Ref}_{\text{sensor}}$, estimate the transformation \mathbf{T} parameterized by a parameter vector \mathbf{p} which relates $\text{Ref}_{\text{sensor}}$ to $\text{Ref}_{\text{model}}$. More specifically, we assume for now that both data sets are surfaces, with the sensed data represented as a collection of points $\{\mathbf{q}_i, i = 1 \dots N\}$, and the model surface S represented in some arbitrary way. The estimation 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

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

where $d(\mathbf{r}_i, S)$ is the minimum Euclidean distance from the point \mathbf{r}_i to the surface S ,

$$d(\mathbf{r}_i, S) = d(\mathbf{T}(\mathbf{q}_i; \mathbf{p}), S) = \min_{\mathbf{s} \in S} \|\mathbf{r}_i - \mathbf{s}\| \quad (2)$$

and σ_i^2 is the variance associated with point i [20, 4]. In addition to the distance-based cost function (2), which corresponds in a Bayesian interpretation to the log likelihood of the measurements \mathbf{q}_i , we add an additional regularizing or stabilizing cost function $\mathcal{P}(\mathbf{p})$ on the parameters \mathbf{p} , which corresponds to a prior distribution on the unknown parameters [29]. The exact form of this function depends on what parameters are used (see sections 3.1 and 3.3).

To solve the minimization problem, we require three components: an iterative minimization algorithm (section 4), an efficient method for computing $d(\mathbf{r}, S)$ along with its gradient (section 5), and a suitable representation for the geometric transformation $\mathbf{T}(\mathbf{q}_i; \mathbf{p})$, which we discuss next.

3.1 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 [20], 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 (e.g., on the operating table or in the scanner) is unknown.

A more general class of transformations are the affine transforms, which can be parameterized with 12 degrees of freedom

$$\mathbf{T}_g(\mathbf{q}_i; \mathbf{p}) = \begin{bmatrix} p_{00} & p_{01} & p_{02} & p_{03} \\ p_{10} & p_{11} & p_{12} & p_{13} \\ p_{20} & p_{21} & p_{22} & p_{23} \end{bmatrix} \begin{bmatrix} 1 \\ x_i \\ y_i \\ z_i \end{bmatrix} \quad (3)$$




Figure 1: Sample global deformations: rigid, affine, trilinear and quadratic.

where $\mathbf{q}_i = (x_i, y_i, z_i)$ are the coordinates of point i in $\text{Ref}_{\text{sensor}}$. Affine transforms have scaling along each dimension as well as shearing. The trilinear class of transformations adds another 12 degrees of freedom for a total of 24,

$$\mathbf{T}_g(\mathbf{q}_i; \mathbf{p}) = \begin{bmatrix} p_{00} & p_{01} & & p_{07} \\ p_{10} & p_{11} & \cdots & p_{17} \\ p_{20} & p_{21} & & p_{27} \end{bmatrix} \begin{bmatrix} 1 \\ x_i \\ \dots \\ x_i y_i z_i \end{bmatrix}. \quad (4)$$

The trilinear transformation can also be parameterized by the 3-D positions of the 8 corners of a bounding cube, but the equations are somewhat more complicated. Finally, the 30-parameter quadratic family of transformations can be similarly formed,

$$\mathbf{T}_g(\mathbf{q}_i; \mathbf{p}) = \begin{bmatrix} p_{00} & p_{01} & & p_{09} \\ p_{10} & p_{11} & \cdots & p_{19} \\ p_{20} & p_{21} & & p_{29} \end{bmatrix} \begin{bmatrix} 1 \\ x_i \\ \dots \\ z_i^2 \end{bmatrix}. \quad (5)$$

Figure 1 shows an illustration of the above four global transforms where a unit cube has been deformed by a representative sample from each transformation class.

Each of these global transformations is more general than the previous one, which means that it can model a wider range of deformations, but at the expense of having more parameters whose values may be difficult to estimate reliably. All of these transformations share the property that the position of the transformed point $\mathbf{r}_i = \mathbf{T}_g(\mathbf{q}_i; \mathbf{p})$ is a *linear* function of the parameters in \mathbf{p} , i.e., the transformation can be written as $\mathbf{r}_i = \mathbf{M}(\mathbf{q}_i)\mathbf{p} = \mathbf{M}_i\mathbf{p}$ [35]. For example, the affine deformation can be represented with

$$\mathbf{M}_i = \begin{bmatrix} 1 & 0 & 0 & x_i & 0 & 0 & y_i & 0 & 0 & z_i & 0 & 0 \\ 0 & 1 & 0 & 0 & x_i & 0 & 0 & y_i & 0 & 0 & z_i & 0 \\ 0 & 0 & 1 & 0 & 0 & x_i & 0 & 0 & y_i & 0 & 0 & z_i \end{bmatrix},$$

where we have ordered the parameter vector as $\mathbf{p} = [p_{00} \ p_{10} \ \dots \ p_{23}]^T$.

As more global parameters are added to the transformation model, the ability to recover them reliably starts to deteriorate. While in practice we have found that the rigid motion parameters are always recovered reliably, it is usually necessary to

add *prior constraints* on the parameters for affine or more general models. We do this using an additional cost function

$$\mathcal{P}(\mathbf{p}) = \sigma_p^{-2} \|\mathbf{p}\|^2, \quad (6)$$

for some suitably chosen σ_p .

3.2 Local spline deformations

To obtain a wider range of more flexible deformations, we could continue increasing the order of the polynomial. However, 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,

$$\mathbf{T}_l(\mathbf{q}_i; \mathbf{p}) = \mathbf{q}_i + \sum_{j,k,l} \mathbf{u}_{jkl} B_j(x_i) B_k(y_i) B_l(z_i), \quad (7)$$

where the \mathbf{u}_{jkl} are the spline deformation coefficients which comprise the parameter vector \mathbf{p} , and B_j , B_k , and B_l are B-spline basis functions [28]. The \mathbf{u}_{jkl} vectors are located on an adaptive 3-D grid, which we call the octree-spline (see section 6). Each basis function (a piecewise polynomial function) only has a limited range of support, i.e., it is non-zero only in the interval $x_{j-o} \leq x \leq x_{j+o}$, where the x_j , $j = 0 \dots M_x$, form a subdivision of each coordinate axis in $\text{Ref}_{\text{sensor}}$. This implies that only a small number of the \mathbf{u}_{jkl} will contribute to the value of \mathbf{r}_i , or, equivalently, that the matrices \mathbf{M}_i in this representation will be extremely sparse (only a few non-zero entries). In general, this means that 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. The \mathbf{u}_{jkl} can be thought of as local estimates of the displacements required to register the two surfaces, which are smoothed and interpolated through the action of the spline functions.

3.3 Regularization

To put prior constraints on the spline deformation parameters \mathbf{u}_{jkl} , we could simply take their norm as before. A more general solution, which we prefer, is to use *regularization* [1, 30]. A fairly general form for a regularizer is

$$\mathcal{P}(\mathbf{u}) = \frac{1}{2} \sum_{m=0}^p w_m \mathcal{R}_m(\mathbf{u}) \quad (8)$$

where \mathbf{u} is the d -dimensional displacement spline (the second term in (7)), w_m are weights, and

$$\mathcal{R}_m(\mathbf{u}) = \int \sum_{j_1 + \dots + j_d = m} \frac{m!}{j_1! \dots j_d!} \left| \frac{\partial^m \mathbf{u}(\mathbf{x})}{\partial x_1^{j_1} \dots \partial x_d^{j_d}} \right|^2 d\mathbf{x}$$

is the m -th order stabilizer [29]. The 0th order stabilizer is similar to the regular norm on the \mathbf{u}_{jkl} , while a first-order stabilizer penalizes linear variations in $\mathbf{u}(\mathbf{x})$. In practice, we have been using a linear combination of 0th and 1st order stabilizers.

To compute the discrete cost function defined on the \mathbf{u}_{jkl} corresponding to (8), we can take one of two approaches. The first is to use *finite element analysis*, which involves analytically evaluating (8) using (7). This has the advantage of yielding more accurate smoothness measures, but at the cost of more complicated discrete energy equations. The approach we take is to use *finite differences*, where we approximate the integrals with averaged squared discrete derivative estimates, e.g., for \mathcal{R}_0 we use $h^3 \sum_{jkl} |\mathbf{u}_{jkl}|^2$, whereas for \mathcal{R}_1 we use

$$\mathcal{R}_1(\{\mathbf{u}_{jkl}\}) \approx h \sum_{jkl} |\mathbf{u}_{j+1,k,l} - \mathbf{u}_{jkl}|^2 + \dots, \quad (9)$$

where h is the size of each cube in the spline domain.

4 Least squares minimization

To perform the nonlinear least squares minimization, we use the Levenberg-Marquardt algorithm because of its good convergence properties [26]. Least squares techniques work well when we have many uncorrelated noisy measurements with a normal (Gaussian) distribution.¹

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,

$$\frac{\partial d_i}{\partial \mathbf{p}} = \mathbf{g}_i \frac{\partial \mathbf{r}_i}{\partial \mathbf{p}} = \mathbf{g}_i \mathbf{M}_i \quad (10)$$

in the case when the transformation can be written as $\mathbf{r}_i = \mathbf{M}_i \mathbf{p}$.

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} ,

$$\mathbf{A} = \sum_{i=1}^N \frac{1}{\sigma_i^2} \frac{\partial d_i}{\partial \mathbf{p}} \left(\frac{\partial d_i}{\partial \mathbf{p}} \right)^T, \quad \mathbf{b} = - \sum_{i=1}^N \frac{1}{\sigma_i^2} d_i \frac{\partial d_i}{\partial \mathbf{p}} \quad (11)$$

¹Under these assumptions, the least squares criterion is equivalent to maximum likelihood estimation.

and then computes an increment $\delta \mathbf{p}$ towards the local minimum by solving

$$(\mathbf{A} + \lambda \text{diag}(\mathbf{A})) \Delta \mathbf{p}^{(k)} = \mathbf{b}, \quad (12)$$

where λ is a stabilizing factor which varies over time [26].² After setting $\mathbf{p}^{(k+1)} = \mathbf{p}^{(k)} + \Delta \mathbf{p}^{(k)}$, this process is repeated until $\mathcal{C}(\mathbf{p})$ 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.

The solution of the system of equations (12) is straightforward when the number of parameters being estimated is reasonably small (as in rigid and global deformations). We currently use the Gauss-Jordan elimination algorithm given in [26]. For systems with more parameters, such as local deformations, it becomes impractical to explicitly compute the Hessian \mathbf{A} , which takes $O(M^2)$ space, where M is the number of parameters (say $16^3 \dots 64^3$), or to directly solve (12), which takes $O(M^3)$ operations. Instead, we use a single step of preconditioned conjugate gradient descent [30]. Section 6 discusses how to use hierarchical basis preconditioning to make this process converge faster.

Once the Levenberg-Marquardt algorithm has converged, we compute a robust estimate of the parameter \mathbf{p} by throwing out the measurements where $d_i^2 \gg \sigma_i^2$ and performing some more iterations [20, 5] (this technique is called using *metrically Winsorised residuals* [16]). This process removes the influence of the *outliers* which are likely to occur when only part of the two data sets is in proper registration.³

When using a gradient descent technique such as Levenberg-Marquardt, there is a possibility that the minimization might fail because of local minima in the high-dimensional parameter space. We try to avoid this by first estimating the simplest possible transformation (a rigid matching), and then estimating successively more complex deformations (affine, then possibly trilinear or quadratic). For the local spline deformations, we begin with a low-resolution spline (typically a single cube), and then use the optimal spline parameters to initialize the estimates at the next finer level.

5 Fast distance computation using oct-tree 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*

²Note that the matrix \mathbf{A} is an approximation to the Hessian matrix, as the second derivative terms in d_i are left out. As mentioned in [26], inclusion of these terms can be destabilizing if the model fits badly or is contaminated by outlier points.

³The threshold for outlier rejection must be fixed according to application-specific knowledge or by experimentation. In our case, we chose to set this threshold to 3σ where σ is a mean *a priori* standard deviation of the noise.

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 [11].

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* [20, 7, 9]. 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 [27] 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 [20, 31] for details).

6 Hierarchical octree spline deformations

The benefits of using a hierarchical data structure for representing the distance map can be extended to the representation used for the tensor-product (local deformation) displacement spline. In this case, the coefficients associated with the spline are 3-D vectors representing the displacement between the model and sensor reference frames. As in the distance octree spline, we use the octree to determine at what resolution the spline coefficients are interpolated. To ensure continuity, we must enforce “crack filling” [27]. In this case, however, since the coefficients in the octree spline are unknown parameters being estimated, this is equivalent to controlling which coefficients are free variables and which are determined by their parents’ values.

As mentioned in section 4, the inner update loop of the Levenberg-Marquardt requires the solution of a large system of linear equations. Because direct methods are prohibitive, we use a single iteration of preconditioned conjugate gradient descent algorithm [26]. To accelerate its convergence, we use a *hierarchical basis* representation of the octree spline [36, 30]. In this representation (Figure 2), displacement values at finer levels are added to the displacements interpolated from the parents, making this a *relative* or *offset representation* [29, 34]. In a true hierarchical basis representation, only nodes not coincident with their parents in location have a non-zero offset (Figure 2), which keeps the number of variables in the hierarchical and regular (*nodal*) representations the same [36].

In our implementation, we keep both the hierarchical and nodal representations, and map between the two as required. For accumulating the distances and gradients required in (11), we compute the interpolated displacements and the derivatives with respect to the parameters in the nodal basis. We then use the hierarchical basis to *smooth* the residual vector \mathbf{b} before selecting a new conjugate direction and computing the optimal step size (for details, see Figure 3 and [30]). Using

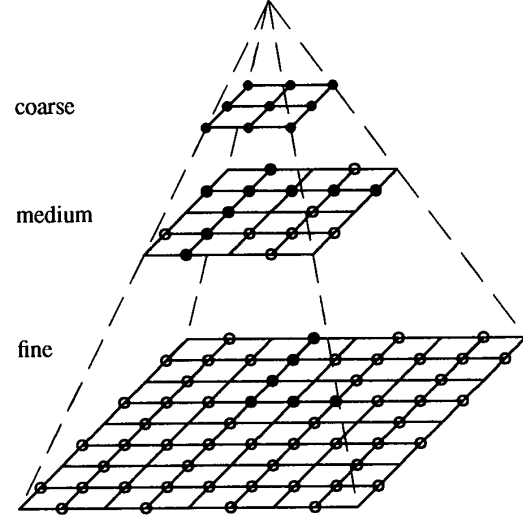


Figure 2: Multiresolution pyramid

The multiple resolution levels are a schematic representation of the octree spline (we represent it in 2D for simplicity). The circles indicate the nodes in the hierarchical basis. Filled circles (\bullet) are free variables in the restricted octree spline, while open circles (\circ) must be zero (see Figure 4).

this technique not only makes the convergence faster but also propagates local corrections over the whole domain, which tends to smooth the resulting deformation significantly.

The hierarchical basis also makes it easier to understand the influence that the octree has on which displacement parameters are free to change. In the hierarchical basis, a node is free to change (has non-zero value), if all of the cubes within its support region have been subdivided to at least its level. In other words, a basis function associated with a non-zero node cannot extend into a larger cube where its influence would not be accounted for (Figure 4).

To use the hierarchical octree spline for modeling defor-

1. $\mathbf{r}_n = \mathbf{b} - \mathbf{A}\mathbf{p}_n = -\nabla\mathcal{C}(\mathbf{p})$
2. † $\tilde{\mathbf{r}}_n = \mathbf{S}\mathbf{S}^T\mathbf{D}^{-1}\mathbf{r}_n$
3. $\beta_n = \tilde{\mathbf{r}}_n \cdot \mathbf{r}_n / \tilde{\mathbf{r}}_{n-1} \cdot \mathbf{r}_{n-1}$
4. † $\mathbf{d}_n = \tilde{\mathbf{r}}_n - \beta_n \mathbf{d}_{n-1}$
5. $\alpha_n = \mathbf{d}_n \cdot \mathbf{r}_n / \mathbf{d}_n^T \mathbf{A} \mathbf{d}_n$
6. $\mathbf{p}_{n+1} = \mathbf{p}_n + \alpha_n \mathbf{d}_n$
7. increment n , loop to 1.

- † $\mathbf{S} =$ mapping from hierarchical to nodal basis [30],
 $\mathbf{D} = \text{diag}(\mathbf{A})$
 ‡ at iteration $n = 0$, $\beta_n = 0$ and $\mathbf{d}_{n-1} = \mathbf{0}$.

Figure 3: Hierarchical basis preconditioned conjugate gradient descent algorithm (simplified from [30])




Figure 4: Quadtree associated with spline function. The nodes with filled circles (●) are free variables in the associated hierarchical basis, whereas the open circles (○) (and also the nodes not drawn) must be zero (in the nodal basis, these nodes are interpolated from their ancestors).

mations, we still have to devise a strategy for deciding how to adaptively subdivide the octree spline. The heuristic we use is to measure the distance to the surface at the center of each cube, and to subdivide all cubes which are within a given distance of the surface (this threshold is a linear function of the current octree resolution). We do this in a recursive top-down fashion, dividing cubes in the octree only after we have finished iterating over the displacement coefficients u_{jkl} at a given resolution. Figure 6d shows a slice through the deformation spline clearly showing how the octree cubes are divided more finely closer to the model surface.

The hierarchical octree spline favors smoother displacement solutions because it smooths out the local deformation corrections (the coarse-to-fine registration algorithm described in Section 4 also has a similar effect). However, to achieve sufficiently smooth displacements, we have to use regularization. To implement the regularization, we evaluate the average squared finite differences inside each cube of the octree. Thus for the 0th order stabilizer, we evaluate the average of the 8 squared corner values in each cube scaled by its volume. For the 1st order stabilizer, we evaluate the average of the squared differences along each of the 12 edges scaled by the cube size.

7 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 synthetic 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.⁴ We then initialize the Levenberg-Marquardt algorithm with some initial rigid pose

⁴For affine transforms, this inverse has a closed form solution. For other non-linear transforms, we use an iterative inversion technique.

Deform.	cube (affine)	cube (trilin.)	cube (quadr.)	face	vertebra
rigid	32.81	10.94	10.11	0.89	3.83
global	4.25	4.54	4.44	0.76	3.01
local				0.57	1.77

Table 1: Root mean square (RMS) error after each stage of algorithm for the five test data sets [31]. The RMS error is in octree units (the surface diameter is 32 units).

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 stabilize the motion estimates, we use a small prior penalty term of $\sigma_p = 1$ (compared to $\sigma_i = 0.5$). Table 1 shows the RMS errors between the two surfaces after the rigid and global non-rigid deformations [31].

To demonstrate the local non-rigid matching, we use two different sets of range data acquired with a Cyberware laser range scanner (Figures 5a and 5b). In their initial positions, the data sets overlap by about 50% and differ in orientation by about 10° (Figure 5c). Here, the octree spline distance map is computed on the larger of the two data sets (*george1* in Figure 5a), and the smaller of the two data sets is deformed (*heidi* in Figure 5b). After 8 iterations of rigid matching and 8 iterations of non-rigid affine matching, the registered data sets appear as in Figure 5d. We then perform 8 iterations at each level of the local displacement spline for 1 through 5 levels. The octree spline has a total of 5728 cubes for a total of about 17000 degrees of freedom. A full resolution displacement spline would have $3(2^5 + 1)^3 \approx 108000$ degrees of freedom. Even with our 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 5f, 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 *heidi* (Figure 5e) resembles that of *george1* (Figure 5a) more than its former (undeformed) self (Figure 5b).

As a final 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 6 shows the result of our matching. After rigid 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 (this may be an artifact of the asymmetrical nature of the cost function). By examining a slice through the octree deformation spline, we can see clearly both how the spline is subdivided more finely near the surface, and how adding reg-

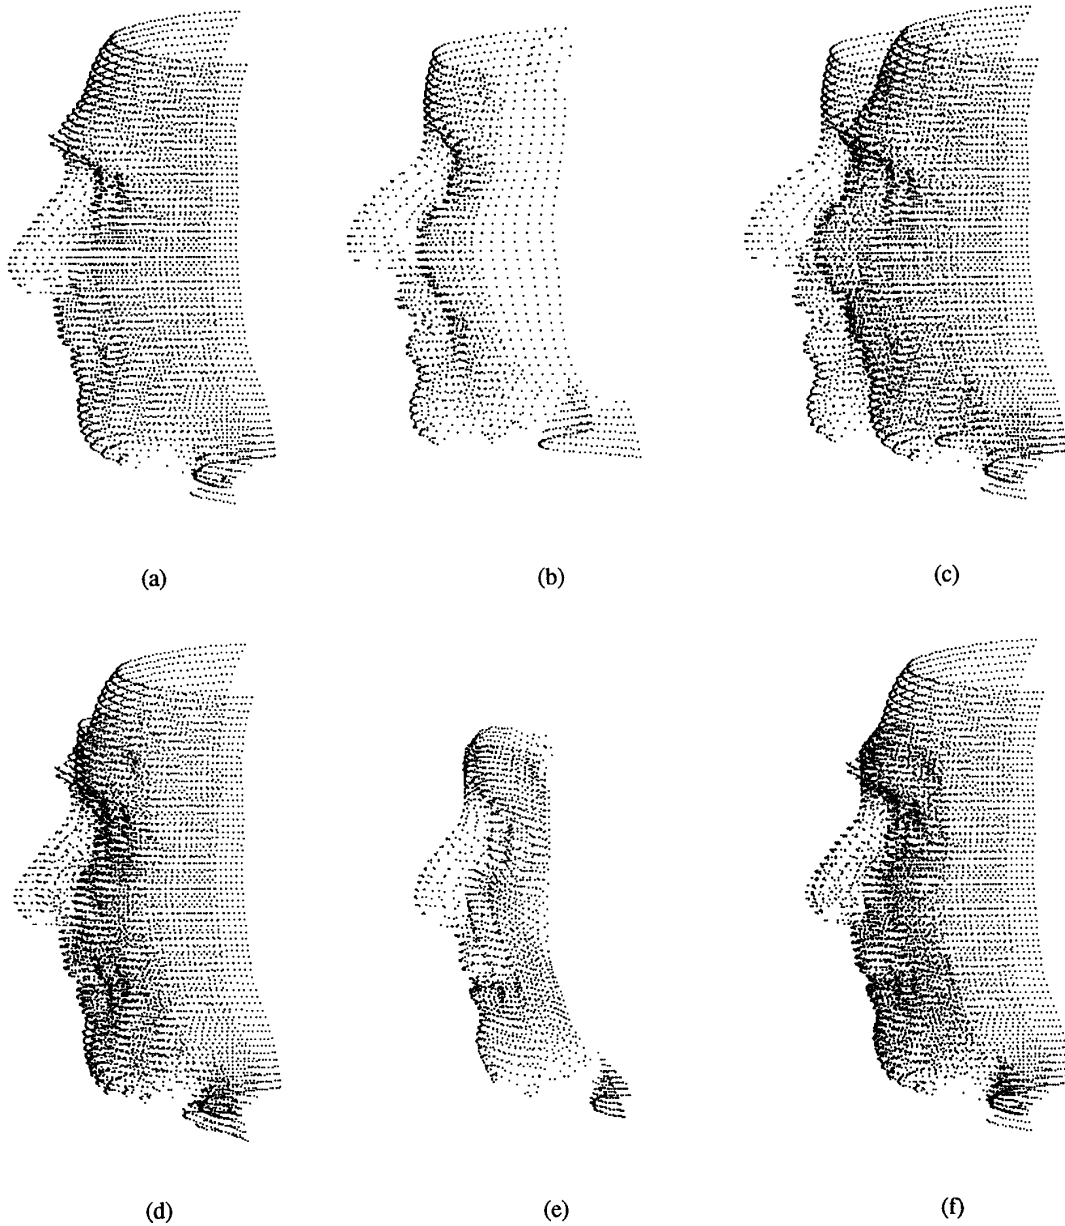


Figure 5: Registration of two face data sets:
 (a) model data set (george1) (b) sensor data set (heidi) (c) both data sets overlaid (initial position) (d) after affine registration (e) final deformed sensor data (f) final registered data sets

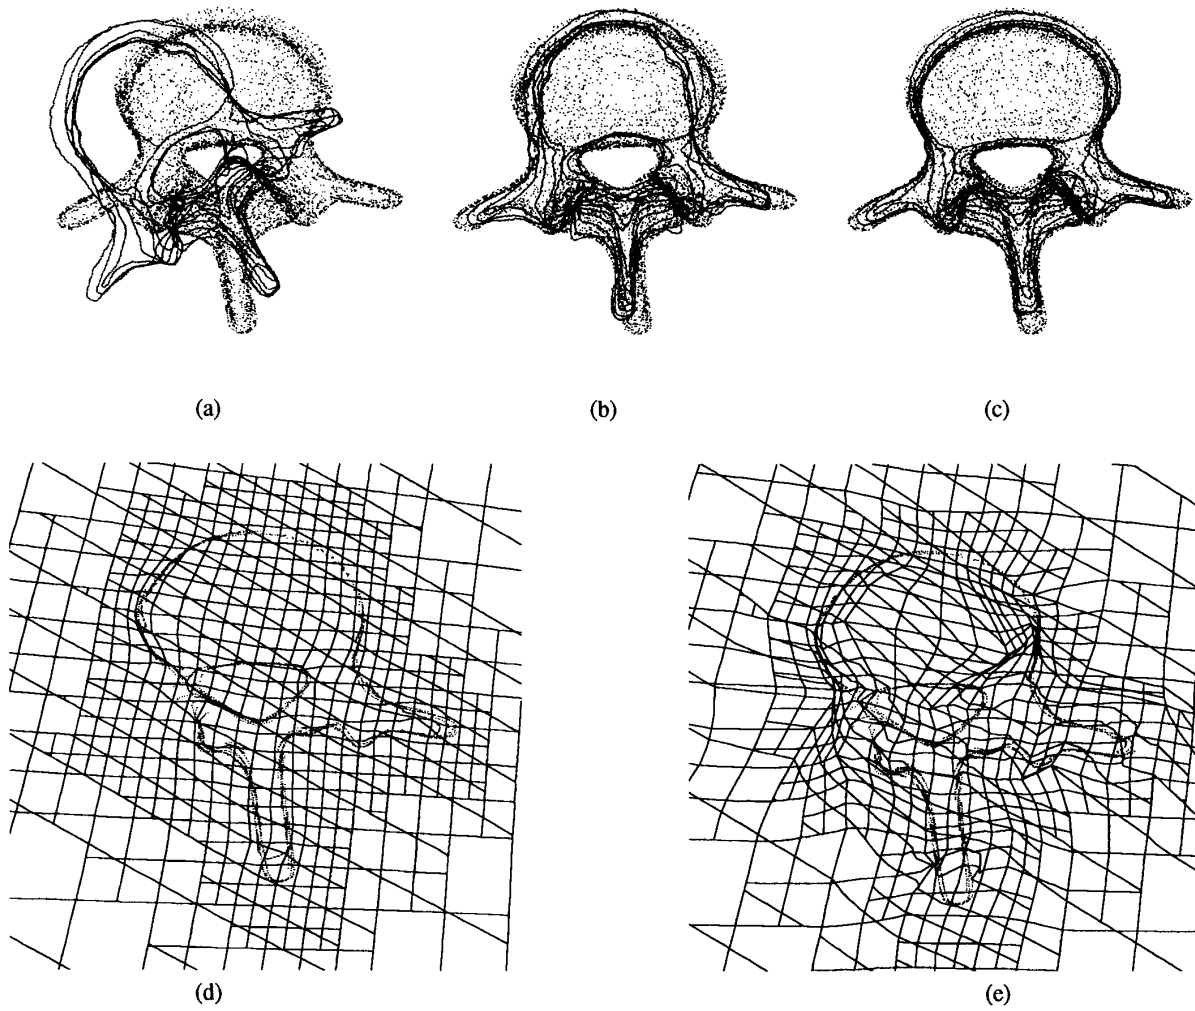


Figure 6: Registration of patient vertebra with plastic "phantom":
 (a) initial configuration (b) after rigid registration (c) after final local spline registration (d) slice through final displacement spline (e) slice through displacement spline without regularization (the local registration results are visually similar with/without regularization).

ularization dramatically smooths the displacement map. As before, the RMS error of the match is significantly smaller for the local deformation than for the global deformation (Table 1).

8 Discussion and Conclusions

In this paper, we have developed a technique for registering 3-D surfaces which are related by 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 [14, 23]) 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. The computational complexity argument assumed that if a small set of features was found and matched (say dozens of discrete features or hundreds of ridge points, as opposed to tens of thousands of surface points), the overall complexity of the algorithm would be drastically lower. However, by 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.

The second argument in favor of discrete features is related to the first, and examines the combinatorial complexity of certain discrete matching (correspondence) algorithms. 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, and manual intervention or guidance is often not a problem (e.g., in interactive analysis or surgical planning), 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 with elastic properties (as in [32, 15, 22]). While the latter approach may be more 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 smooth connected representation. Another advantage is that volumetric deformations prevent gross shape deformations and surface interpenetration.

Volumetric deformations also yield auxiliary information about the motion of nearby structures which did not participate in the matching, e.g., registrations performed on certain

easily identifiable brain structures can be used to estimate the registration for the whole brain. This suggests an incremental, model-driven segmentation process.

In theory, volumetric deformation models are of a higher complexity than elastic surface models, i.e., they represent the deformations over a dense region of 3-D, rather than over a 2-D manifold. However, the use of a multiresolution, coarse-to-fine strategy results in a very efficient matching algorithm. Furthermore, the use of an adaptive spatial subdivision can make a volumetric representation such as the octree competitive with a surface representation for sufficiently smooth surfaces [27].

The estimation of elastic deformations in general is an ill-posed problem for which it is crucial to define some notion of *minimal deformation* that relates directly to the notion of smoothness. In our method, we have introduced smoothness in the deformation in three ways: first, using regularization terms of order 0 and 1, second using a sweep up and down in the hierarchical basis, and third using a coarse to fine strategy in the minimization. The correct amount of smoothing is application dependent: for change detection and morphometrics, we wish to have conservative (smooth) deformations, while for model-based segmentation, we wish to fit the data exactly (less smoothness).

Our experiments to date have been performed on pre-segmented 3-D surfaces. We are planning to extend our algorithm to work directly with unsegmented 3-D medical images. Since we are matching a known anatomical model with patient data, one of the data sets will still be a segmented surface. Our choice is then to use gradients in the medical image to attract a deformable surface model, or to run an edge operator over the medical image [24] and use this as the (noisy) surface to be deformed. In the latter case, our algorithm must be sufficiently robust to reject large number of outliers [5]. Other areas of future investigation include an adaptive algorithm for deciding how to refine the local octree deformation spline, the choice of the order of the octree spline, and the choice of various smoothness constraints (regularizers).

To summarize, the technique we have developed allows us to rapidly register two segmented 3-D surfaces which are related by an unknown non-rigid deformation. Our use of octree splines makes this algorithm efficient, and our choice of weighted non-linear least squares makes it statistically optimal. Our algorithm has immediate applications to the registration of 3-D medical surfaces with anatomical models and to change detection in medical imagery. We are extending the algorithm to support the automatic model-based segmentation of 3-D medical images, and we expect this to further widen the applicability and utility of our approach.

References

- [1] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing*, 46:1–21, 1989.

- [2] Alan H. Barr. Global and local deformations of solid primitives. *Computer Graphics (SIGGRAPH'84)*, 18(3):21–30, July 1984.
- [3] T. Beier and S. Neely. Feature-based image metamorphosis. *Computer Graphics (SIGGRAPH'92)*, 26(2):35–42, July 1992.
- [4] 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, 1992.
- [5] E. Bittar, S. Lavallée, and R. Szeliski. A method for registering overlapping range images of arbitrarily shaped surfaces for 3-d object reconstruction. In *SPIE Vol. 2059, Sensor Fusion VI*, Boston, September 1993. Society of Photo-Optical Instrumentation Engineers.
- [6] 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.
- [7] L. Brunie, S. Lavallée, and R. Szeliski. Using forces fields derived from 3D distance maps for inferring the attitude of a 3D rigid object. In G. Sandini, editor, *Second European Conference on Computer Vision (ECCV'92)*, pages 670–675, Santa Margherita, Italy, May 1992. Springer Verlag.
- [8] D. J. Burr. A dynamic model for image registration. *Computer Graphics and Image Processing*, 15(2):102–112, February 1981.
- [9] G. Champeboux, S. Lavallée, R. Szeliski, and L. Brunie. From accurate range imaging sensor calibration to accurate model-based 3-D object localization. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'92)*, Champaign, Illinois, June 1992.
- [10] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. 3D brain mapping using a deformable neuronanatomy. *Physics in Medicine and Biology*, (accepted) October 1993.
- [11] P.-E. Danielson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.
- [12] 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.
- [13] D. R. Forsey and R. H. Bartels. Hierarchical B-spline refinement. *Computer Graphics (SIGGRAPH'88)*, 22(4):205–212, August 1988.
- [14] A. Guéziec 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.
- [15] 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.
- [16] P. J. Huber. *Robust Statistics*. John Wiley & Sons, New York, New York, 1981.
- [17] 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.
- [18] S. Lavallée. *Geste Medico-Chirurgicaux Assistés par Ordinateur : Application a la Neurochirurgie Stereotaxique*. PhD thesis, Grenoble University, France, December 1989.
- [19] 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.
- [20] 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.
- [21] S. Leitner, I. Marque, S. Lavallée, and P. Cinquin. Dynamic segmentation : finding the edge with spline snakes. In P.J. Laurent, editor, *International Conference on Curves and Surfaces*, Chamonix, 1991. Academic Press.
- [22] T. McInerney and D. Terzopoulos. A finite element model for 3D shape reconstruction and nonrigid motion tracking. In *Fourth International Conference on Computer Vision (ICCV'93)*, pages 518–523, Berlin, Germany, May 1993. IEEE Computer Society Press.
- [23] O. Monga, S. Benayoun, and N. Ayache. From partials derivatives of 3D density images to ridge lines. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'92)*, pages 354–359, Champaign, Illinois, June 1992.
- [24] O. Monga, R. Deriche, G. Malandrain, and J. P. Cocquerez. Recursive filtering and edge closing : Two primary tools for 3D edge detection. In *First European Conference on Computer Vision (ECCV'90)*, pages 56–65, Antibes, France, April 1990. Springer-Verlag.
- [25] C.A. Pelizzari, G.T.Y. Chen, D.R. Spelbring, R.R. Weichselbaum, and C-T. Chen. Accurate 3-D registration of CT, PET, and-or MR images of the brain. *J. Computer Assisted Tomography*, 13(1):20–26, 1989.
- [26] 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.
- [27] H. Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading, Massachusetts, 1989.
- [28] T.W. Sederberg and S.R. Parry. Free-form deformations of solid geometric models. *Computer Graphics (SIGGRAPH'86)*, 20(4):151–160, 1986.
- [29] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-Level Vision*. Kluwer Academic Publishers, Boston, Massachusetts, 1989.
- [30] R. Szeliski. Fast surface interpolation using hierarchical basis functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(6):513–528, June 1990.
- [31] 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.
- [32] D. Terzopoulos and D. Metaxas. Dynamic 3D models with local and global deformations: Deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):703–714, July 1991.
- [33] D. Terzopoulos, A. Witkin, and M. Kass. Constraints on deformable models: Recovering 3D shape and nonrigid motion. *Artificial Intelligence*, 36:91–123, 1988.
- [34] W. Welch and A. Witkin. Variational surface modeling. *Computer Graphics (SIGGRAPH'92)*, 26(2):157–166, July 1992.
- [35] A. Witkin and W. Welch. Fast animation and control of nonrigid structures. *Computer Graphics (SIGGRAPH'90)*, 24(4):243–252, August 1990.
- [36] H. Yserantant. On the multi-level splitting of finite elements spaces. *Numerische Mathematik*, 49:379–412, 1986.