

Consistent Image Registration

G. E. Christensen* and H. J. Johnson

Abstract—This paper presents a new method for image registration based on jointly estimating the forward and reverse transformations between two images while constraining these transforms to be inverses of one another. This approach produces a consistent set of transformations that have less pairwise registration error, i.e., better correspondence, than traditional methods that estimate the forward and reverse transformations independently. The transformations are estimated iteratively and are restricted to preserve topology by constraining them to obey the laws of continuum mechanics. The transformations are parameterized by a Fourier series to diagonalize the covariance structure imposed by the continuum mechanics constraints and to provide a computationally efficient numerical implementation. Results using a linear elastic material constraint are presented using both magnetic resonance and X-ray computed tomography image data. The results show that the joint estimation of a consistent set of forward and reverse transformations constrained by linear-elasticity give better registration results than using either constraint alone or none at all.

Index Terms—Correspondence, deformable templates, image registration, inverse transformation.

I. INTRODUCTION

IMAGE registration has many uses in medicine such as multimodality fusion, image segmentation, deformable atlas registration, functional brain mapping, image-guided surgery, and characterization of normal versus abnormal anatomical shape and variation. The fundamental assumption in each of these applications is that image registration can be used to define a meaningful correspondence mapping between anatomical images collected from imaging devices such as computed tomography (CT), magnetic resonance imaging (MRI), cryosectioning, etc. It is often assumed that this correspondence mapping or transformation is one-to-one, i.e., each point in image T is mapped to only one point in image S and *vice versa*. A fundamental problem with a large class of image registration techniques is that the estimated transformation from image T to S does not equal the inverse of the estimated transform from S to T . This inconsistency is a result of the matching criteria's inability to uniquely describe the correspondences between two images. This paper seeks to overcome this limitation by jointly estimating the transformation from T to

S and from S to T while enforcing the consistency constraint that these transforms are inverses of one another.

The forward transformation h from image T to S and the reverse transformation g from S to T are shown in Fig. 1. Ideally, the transformations h and g should be uniquely determined and should be inverses of one another. Estimating h and g independently very rarely results in a consistent set of transformations due to a large number of local minima. To overcome this deficiency in current registration systems, we propose to jointly estimate h and g while constraining these transforms to be inverses of one another. Jointly estimating the forward and reverse transformations provides additional correspondence information and helps ensure that these transformations define a consistent correspondence between the images. Although uniqueness is very difficult to achieve in medical image registration, the joint estimation should lead to more consistent and biologically meaningful results.

Image registration algorithms use landmarks [1]–[4], contours [5]–[7], surfaces [8]–[11], volumes [12]–[18], [6], [19]–[21], or a combination of these features [22] to manually, semi-automatically or automatically define correspondences between two images. The need to impose the invertibility consistency constraint depends on the particular application and on the correspondence model used for registration. In general, registration techniques that do not uniquely determine the correspondence between image volumes should benefit from the consistency constraint. This is because such techniques often rely on minimize/maximize a similarity measure which has a large number of local minima/maxima due to the correspondence ambiguity. Examples include methods that minimizing/maximizing similarity measures between features in the source and target images such as image intensities, object boundaries/surfaces, etc. In theory, the higher the dimension of the transformation, the more local minima these similarity measures have. Methods that use specified correspondences for registration will benefit less or not at all from the invertibility consistency constraint. For example, landmark-based registration methods implicitly impose an invertibility constraint at the landmarks because the correspondence defined between landmarks is the same for estimating the forward and reverse transformations. However, the drawbacks of specifying correspondences include requiring user interaction to specify landmarks, unique correspondences cannot always be specified, and such methods usually only provide coarse registration due to the small number of correspondences specified.

In this paper, we will restrict our analysis to the class of applications that can be solved using diffeomorphic transformations. A diffeomorphic transformation is defined to be continuous, one-to-one, onto, and differentiable. The diffeomorphic restriction is valid for a large number of problems in which the

Manuscript received October 14, 1999; revised April 26, 2001. This work was supported in part by the National Institutes of Health (NIH) under Grant NS35368, Grant CA75371, and Grant DC03590, and in part by a grant from the Whitaker Foundation. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was M. W. Vannier. *As-terisk indicates corresponding author.*

*G. E. Christensen is with the Department of Electrical and Computer Engineering, University of Iowa, 205 CC, Iowa City, IA 52242 USA (e-mail: gary-christensen@uiowa.edu).

H. J. Johnson is with the Department of Electrical and Computer Engineering, University of Iowa, Iowa City, IA 52242 USA.

Publisher Item Identifier S 0278-0062(01)05523-9.

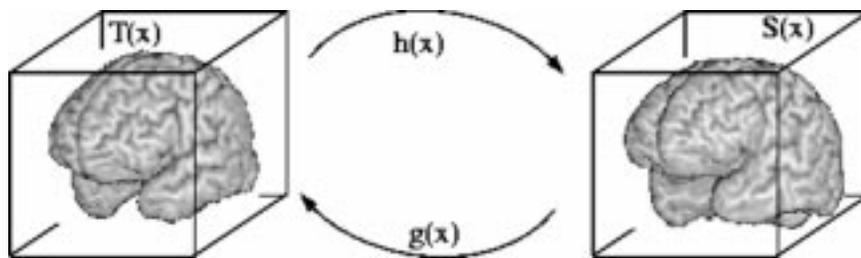


Fig. 1. Consistent image registration is based on the principle that the mappings h from T to S and g from S to T define a point by point correspondence between T and S that are consistent with each other. This consistency is enforced mathematically by jointly estimating h and g while constraining h and g to be inverse mappings of one another.

two images have the same structures and neighborhood relationships but have different shapes.

Diffeomorphic transformations maintain the topology and guarantee that connected subregions of an image remain connected, neighborhood relationships between structures are preserved, and surfaces are mapped to surfaces. Preserving topology is important for synthesizing individualized electronic atlases the knowledge base of the atlas maybe transferred to the target anatomy through the topology preserving transformation providing automatic labeling and segmentation. If total volume of a nucleus, ventricle, or cortical subregion are an important statistic it can be generated automatically. Topology preserving transformations that map the template to the target also can be used to study the physical properties of the target anatomy such as mean shape and variation. Likewise, preserving topology allows data from multiple individuals to be mapped to a standard atlas coordinate space [23]. Registration to an atlas removes individual anatomical variation and allows information from many experiments to be combined and associated with a single canonical anatomy.

II. REGISTRATION ALGORITHM

A. Problem Statement

Traditionally, the image registration problem has been stated as: Find the transformation h that maps the template image volume T into correspondence with the target image volume S . Alternatively, the problem can be stated as: Find the transformation g that transforms S into correspondence with T . For this paper, the previous two statements are combined into a single problem and restated as follows.

Problem Statement: Jointly estimate the transformations h and g such that h maps T to S and g maps S to T subject to the constraint that $h = g^{-1}$.

The image registration algorithm is formulated on the continuum and is discretized for implementation. Let T and S represent three-dimensional (3-D) image volumes of voxels dimension $N_1 \times N_2 \times N_3$ corresponding to medical imaging modalities such as MRI, functional MRI, CT, cryosection imagery, etc. Let $\Omega_d = \{(n_1, n_2, n_3) | 0 \leq n_1 < N_1; 0 \leq n_2 < N_2; 0 \leq n_3 < N_3; \text{and } n_1, n_2, n_3 \text{ are integers}\}$ correspond to the set voxel lattice coordinates of the discrete images T and S and let $\Omega_c = [0, 1]^3$. Continuous functions will be denoted with a subscript c and parentheses while discrete functions will be denoted with a subscript d and square brackets. The subscripts

are dropped when a statement is true for both continuous and discrete functions or where the context is clear due to the use of parentheses or square brackets. Let the continuous image $T_c(x)$ for $x \in \Omega_c$ be related to the discrete image $T_d[n]$ for $n \in \Omega_d$ in the normal multidimensional sampling sense $T_d[n] = T_c(n/N)$ where the notation n/N is defined as the 3×1 column vector $[n_1/N_1, n_2/N_2, n_3/N_3]^T$. Likewise, let $S_d[n] = S_c(n/N)$ for $n \in \Omega_d$. The discrete images are extended to the continuum using trilinear interpolation.

The transformations are vector-valued functions that map the image coordinate system Ω_c to itself, i.e., $h_c: \Omega_c \mapsto \Omega_c$ and $g_c: \Omega_c \mapsto \Omega_c$. Regularization constraints are placed on h and g so that they preserve topology. Throughout it is assumed that $h_c(x) = x + u_c(x)$, $h_c^{-1}(x) = x + \tilde{u}_c(x)$, $g_c(x) = x + w_c(x)$ and $g_c^{-1}(x) = x + \tilde{w}_c(x)$ where $h_c(h_c^{-1}(x)) = x$ and $g_c(g_c^{-1}(x)) = x$. The vector-valued functions u , w , \tilde{u} , and \tilde{w} are called displacement fields since they define the transformation in terms of a displacement from a location x . All of the functions h_c , g_c , h_c^{-1} , g_c^{-1} , u_c , \tilde{u}_c , w_c , and \tilde{w}_c are (3×1) vector-valued functions defined on the Ω_c . The continuous transformations and displacement fields are extended to the continuum from their corresponding discrete representations using trilinear interpolation. For example, $h_d[n] = h_c(n/N) = (n/N) + u_c(n/N) = (n/N) + u_d[n]$.

Registration is defined using a symmetric similarity cost function that describes the distance between the transformed template $T \circ h$ and target S , and the distance between the transformed target $S \circ g$ and template T . To ensure the desired properties, the transformations h and g are jointly estimated by minimizing the similarity cost function while satisfying regularization constraints and inverse transformation consistency constraints. The regularization constraints are enforced on the transformations by constraining the them to satisfy the laws of continuum mechanics [24].

B. Symmetric Similarity Cost Function

The problem with many image registration techniques is that the image similarity function does not uniquely determine the correspondence between two image volumes. In general, similarity cost functions have many local minima due to the complexity of the images being matched and the dimensionality of the transformation. It is these local minima (ambiguities) that cause the estimated transformation from image T to S to be different from the inverse of the estimated transformation from S to T . In general, this becomes more of a problem as the dimensionality of the transformation increases.

To overcome correspondence ambiguities, the transformations from image T to S and from S to T are jointly estimated. This is accomplished by defining a cost function to measure the shape differences between the deformed image $T \circ h$ and image S and the differences between the deformed image $S \circ g$ and image T . Ideally, the transformations h and g should be inverses of one another, i.e., $h = g^{-1}$. In this paper, the transformations h and g are estimated by minimizing a cost function

$$\begin{aligned} C_{\text{SIM}}(T \circ h, S) + C_{\text{SIM}}(S \circ g, T) \\ = \int_{\Omega_c} |T_c(h_c(x)) - S_c(x)|^2 dx \\ + \int_{\Omega_c} |S_c(g_c(x)) - T_c(x)|^2 dx \end{aligned} \quad (1)$$

where the intensities of T_c and S_c are assumed to be scaled between 0 and 1. To use this similarity function, the images T and S must correspond to the same imaging modality and they may require preprocessing to equalize the intensities of the image. In practice, MRI images require intensity equalization while CT images do not. A simple but effective method for intensity equalizing MRI data is to compute the histograms of the two images, scale the axis of one histogram so that the gray and white matter maximums match, and then apply the intensity scaling to the image.

In practice, the images S and T are discrete and the integrals in (1) are discretized and implemented as summations

$$\begin{aligned} C_{\text{SIM}}(T \circ h, S) + C_{\text{SIM}}(S \circ g, T) \\ = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} \left| T_c \left(h_c \left(\frac{n}{N} \right) \right) - S_c \left(\frac{n}{N} \right) \right|^2 \\ + \left| S_c \left(g_c \left(\frac{n}{N} \right) \right) - T_c \left(\frac{n}{N} \right) \right|^2 \\ = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} |T_c(h_d[n]) - S_d[n]|^2 \\ + |S_c(g_d[n]) - T_d[n]|^2 \\ = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} |T_d[Nh_d[n]] - S_d[n]|^2 \\ + |S_d[Ng_d[n]] - T_d[n]|^2 \end{aligned} \quad (2)$$

where the notation $Nh_d[n]$ is defined as the 3×1 column vector $[N_1 h_d^{(1)}[n], N_2 h_d^{(2)}[n], N_3 h_d^{(3)}[n]]^T$ and the terms $T_d[Nh_d[n]]$ and $S_d[Ng_d[n]]$ are evaluated using trilinear interpolation. Equation (2) is a discrete approximation to the integral cost defined in (1) which becomes exact as $N_1 \mapsto \infty$, $N_2 \mapsto \infty$, and $N_3 \mapsto \infty$. The discretized cost defined in (2) demonstrates a tradeoff between computational time and accuracy. This tradeoff can be exploited by using a low-resolution computational grid Ω_d during the early iterations of the algorithm when accuracy is less important and using a high-resolution grid Ω_d in the later iterations to get an accurate final result.

Note that this joint estimation approach applies to both linear and nonlinear transformations. In general, the squared-error similarity functions in (1) can be replaced by any suitable similarity function—mutual information [25], [26], demons [6], an intensity variance cost function [27], etc.—where the

choice is dependent on the particular registration application (see Section IV).

C. Inverse Consistency Constraint

Minimizing a symmetric cost function like (1) is not sufficient to guarantee that h and g are inverses of each other because the contributions of h and g to the cost function are independent. In order to couple the estimation of h with that of g , an inverse consistency constraint is imposed that is minimized when $h = g^{-1}$. The inverse consistency constraint is given by

$$\begin{aligned} C_{\text{ICC}}(u, \tilde{w}) + C_{\text{ICC}}(w, \tilde{u}) \\ = \int_{\Omega_c} \|h_c(x) - g_c^{-1}(x)\|^2 dx + \int_{\Omega_c} \|g_c(x) - h_c^{-1}(x)\|^2 dx \\ = \int_{\Omega_c} \|u_c(x) - \tilde{w}_c(x)\|^2 dx + \int_{\Omega_c} \|w_c(x) - \tilde{u}_c(x)\|^2 dx \end{aligned} \quad (3)$$

where $h_c(x) = x + u_c(x)$, $h_c^{-1}(x) = x + \tilde{u}_c(x)$, $g_c(x) = x + w_c(x)$ and $g_c^{-1}(x) = x + \tilde{w}_c(x)$. Notice that the inverse consistency constraint is written in a symmetric form like the symmetric cost function for similar reasons. The discretized inverse consistency constraint is given by

$$\begin{aligned} C_{\text{ICC}}(u, \tilde{w}) + C_{\text{ICC}}(w, \tilde{u}) \\ = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} \|u_d[n] - \tilde{w}_d[n]\|^2 + \|w_d[n] - \tilde{u}_d[n]\|^2. \end{aligned} \quad (4)$$

The procedure used to compute the inverse transformation of a transformation with minimum Jacobian greater than zero is as follows. Assume that $h_c(x)$ is a continuously differentiable transformation that maps Ω_c onto Ω_c and has a positive Jacobian for all $x \in \Omega_c$. The fact that the Jacobian is positive at a point $x \in \Omega_c$ implies that it is locally one-to-one and, therefore, has a local inverse. It is, therefore, possible to select a point $y \in \Omega_c$ and iteratively search for a point $x \in \Omega_c$ such that $\|y - h(x)\|$ is less than some threshold provided that the initial guess of x is close to the final value of x .

The discretized inverse transformation is related to the continuous inverse transformation by $h_d[n] = h_c(n/N)$. This implies that the discrete inverse transformation only needs to be calculated at each sample point $n \in \Omega_d$. The following procedure is used to compute the discrete inverse h_d^{-1} of the transformation h_d .

```

For each  $n \in \Omega_d$  do {
  Set  $\delta = [1, 1, 1]^T$ ,  $x = n/N$ , iteration = 0.
  While ( $\|\delta\| > \text{threshold}$ ) do {
     $\delta = n/N - h_d[Nx]$ 
     $x = x + \delta/2$ 
    iteration = iteration + 1
    if (iteration > max_iteration) then
      Report algorithm failed to converge
      and exit.
  }
   $h_d^{-1}[n] = x$ 
}

```

As before

$$\frac{n}{N} = \left[\frac{n_1}{N_1}, \frac{n_2}{N_2}, \frac{n_3}{N_3} \right]^T$$

$$Nx = [N_1x_1, N_2x_2, N_3x_3]^T$$

and $h_d[Nx]$ is computed using trilinear interpolation. We typically set $\text{threshold} = 10^{-4}$ and $\text{max_iteration} = 1000$. In practice, the algorithm converges when the minimum Jacobian of h is greater than zero although we have not proved this mathematically. Reducing the value of the threshold gives a more accurate inverse but increases the iteration time. To reduce computation time, the above algorithm is modified by rastering through the elements of Ω_d and initializing $x = (n/N) + \varepsilon$ where ε is equal to the displacement of the previously estimated grid point.

D. Regularization Constraint

Minimizing the cost function in (3) does not ensure that the transformations h and g are diffeomorphic transformations except for when $C_{\text{ICC}} = 0$. Continuum mechanical models such as linear elasticity [28], [22] and viscous fluid [22], [15] can be used to regularize the transformations. In this paper, a linear-elastic constraint of the form

$$C_{\text{REG}}(u) + C_{\text{REG}}(w)$$

$$= \int_{\Omega} \|L_c u_c(x)\|^2 dx + \int_{\Omega} \|L_c w_c(x)\|^2 dx \quad (5)$$

is used to regularize the transformations where $u_c(x) = h_c(x) - x$ and $w_c(x) = g_c(x) - x$. The linear elasticity operator L_c has the form $L_c u_c(x) = -\alpha \nabla^2 u_c(x) - \beta \nabla(\nabla \cdot u_c(x)) + \gamma u_c(x)$ where $\nabla = [\partial/\partial x_1, \partial/\partial x_2, \partial/\partial x_3]$ and $\nabla^2 = \nabla \cdot \nabla = [\partial^2/\partial x_1^2 + \partial^2/\partial x_2^2 + \partial^2/\partial x_3^2]$. In general, L_c can be any non-singular linear differential operator [29]. The limitation of using linear differential operators is that they cannot prevent the transformation from folding onto itself, i.e., destroying the topology of the images under transformation [30]. This includes the linear elasticity and thin-plate spline models. The linear elasticity operator is used in this paper to help prevent the Jacobian of the transformation from going negative. At each iteration the Jacobian of the transformation is checked to make sure that it is positive for all points in Ω_d which implies that the transformation preserves topology when transforming images.

The purpose of the regularization constraint is to ensure that the transformations maintain the topology of the images T and S . Thus, the elasticity constraint can be replaced by or combined with other regularization constraints that maintain desirable properties of the template (source) and target when deformed. An example would be a constraint that prevented the Jacobian of both the forward and reverse transformations from going to zero or infinity. A constraint that penalizes small and large Jacobian values is given by

$$C_{\text{Jac}}(h) + C_{\text{Jac}}(g)$$

$$= \int_{\Omega} (J(h_c(x)))^2 + \left(\frac{1}{J(h_c(x))} \right)^2$$

$$+ (J(g_c(x)))^2 + \left(\frac{1}{J(g_c(x))} \right)^2 dx$$

where J denotes the Jacobian operator. Further examples of regularization constraints that penalize large and small Jacobians can be found in Ashburner *et al.* [21].

E. Transformation Parameterization

A 3-D Fourier series representation [17] is used to parameterize the forward and reverse transformations. This parameterization is simpler than the parameterizations used in our previous work [14], [28], [29] and each basis coefficient can be interpreted as the weight of a harmonic component in a single coordinate direction. Let $k = [k_1, k_2, k_3]$ and $n = [n_1, n_2, n_3]$. The displacement fields are defined to have the form

$$u_c(x) = \sum_{k \in \Omega_d} \mu[k] e^{j \langle x, N\theta[k] \rangle}$$

and

$$w_c(x) = \sum_{k \in \Omega_d} \eta[k] e^{j \langle x, N\theta[k] \rangle} \quad (6)$$

for $x \in \Omega_c$ where $\mu[k]$ and $\eta[k]$ are (3×1) complex-valued vectors, $\theta[k] = [2\pi k_1/N_1, 2\pi k_2/N_2, 2\pi k_3/N_3]^T$ and $N\theta[k] = [2\pi k_1, 2\pi k_2, 2\pi k_3]^T$. The notation $\langle \cdot, \cdot \rangle$ denotes the dot product of two vectors such that $\langle x, N\theta[k] \rangle = 2\pi k_1 x_1 + 2\pi k_2 x_2 + 2\pi k_3 x_3$. The discretized displacement fields are defined as

$$u_d[n] = u_c\left(\frac{n}{N}\right) = \sum_{k \in \Omega_d} \mu[k] e^{j \langle n, \theta[k] \rangle}$$

and

$$w_d[n] = w_c\left(\frac{n}{N}\right) = \sum_{k \in \Omega_d} \eta[k] e^{j \langle n, \theta[k] \rangle} \quad (7)$$

for $n \in \Omega_d$. The basis coefficients are defined as

$$\mu[k] = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} u_d[n] e^{-j \langle n, \theta[k] \rangle}$$

and

$$\eta[k] = \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} w_d[n] e^{-j \langle n, \theta[k] \rangle}$$

for $k \in \Omega_d$. The displacement fields associated with the inverse of the forward and reverse transformations are given by replacing u , w , μ , and η in (7) with \tilde{u} , \tilde{w} , $\tilde{\mu}$, and $\tilde{\eta}$, respectively.

The Fourier series parameterization is periodic in x and, therefore, has cyclic boundary conditions for x on the boundary of Ω . Further more, the following proposition shows that the displacement fields are real assuming that the coefficients $\mu[k]$ and $\eta[k]$ have complex conjugate symmetry. It is shown in Section II-F that $\mu[k]$ and $\eta[k]$ have complex conjugate symmetry by construction.

Proposition 1: A displacement field of the form $u_d[n] = \sum_{k \in \Omega_d} \mu[k] e^{j \langle n, \theta[k] \rangle}$, $n \in \Omega_d$ is real and can be written as

$$u_d[n] = 2 \sum_{k_1=0}^{(N_1/2)-1} \sum_{k_2=0}^{N_2-1} \sum_{k_3=0}^{N_3-1}$$

$$\cdot \left(a[k] \text{Re} \left\{ e^{j \langle n, \theta[k] \rangle} \right\} - b[k] \text{Im} \left\{ e^{j \langle n, \theta[k] \rangle} \right\} \right)$$

$$n \in \Omega_d \quad (8)$$

if the (3×1) vector $\mu[k] = a[k] + jb[k]$ has complex conjugate symmetry.¹

Proof: The displacement field u_d can be written as

$$\begin{aligned} u_d[n] &= \sum_{k \in \Omega_d} (a[k] + jb[k])e^{j\langle n, \theta[k] \rangle} \\ &= \sum_{k_1=0}^{(N_1/2)-1} \sum_{k_2=0}^{N_2-1} \sum_{k_3=0}^{N_3-1} (a[k] + jb[k])e^{j\langle n, \theta[k] \rangle} \\ &\quad + (a[k] - jb[k])e^{-j\langle n, \theta[k] \rangle} \end{aligned}$$

because the $\mu[k]$ coefficients are complex conjugate symmetric, i.e., $\mu[k] = \mu^*[N - k]$. Simplifying the summand gives the result. Q.E.D.

The Fourier series parameterization in (6) is useful for simplifying the linear elasticity constraint given in (5). The operator L_c can be thought of as a (3×3) matrix differential operator [29] such that the linear elasticity operator as shown in the equation at the bottom of the page. Substituting (7) into (5) and discretizing² the continuous partial derivatives of L_c gives

$$\begin{aligned} C_{\text{REG}}(u) + C_{\text{REG}}(w) \\ = \sum_{k \in \Omega_d} \mu^\dagger[k] D^2[k] \mu[k] + \eta^\dagger[k] D^2[k] \eta[k] \end{aligned} \quad (9) \quad \text{and}$$

where \dagger is the complex conjugate transpose. $D[k]$ is a real-valued, (3×3) matrix with elements

$$[D[k]]_{rs} = \begin{cases} 2\alpha \{ N_1^2(1 - \cos(\theta_1[k])) \\ \quad + N_2^2(1 - \cos(\theta_2[k])) \\ \quad + N_3^2(1 - \cos(\theta_3[k])) \} \\ \quad + 2\beta N_r^2(1 - \cos(\theta_r[k])) + \gamma, & r = s \\ \beta N_r N_s \sin(\theta_r[k]) \sin(\theta_s[k]), & r \neq s. \end{cases}$$

¹This proposition assumes that N_1 is even. A similar statement can be made when N_1 is odd.

²Symmetric difference equations were used to discretize the continuous derivatives. For example, let $\Delta_1 = 1/N_1$ and $\Delta_2 = 1/N_2$ and then $\partial^2 f / \partial x_1^2$ was replaced by $[f(x_1 + \Delta_1, x_2, x_3) - f(x_1 - \Delta_1, x_2, x_3)] / \Delta_1^2$ and $\partial^2 f / \partial x_1 \partial x_2$ was replaced by $[f(x_1 + \Delta_1, x_2 + \Delta_2, x_3) - f(x_1 - \Delta_1, x_2 + \Delta_2, x_3) + f(x_1 - \Delta_1, x_2 - \Delta_2, x_3) - f(x_1 + \Delta_1, x_2 - \Delta_2, x_3)] / \Delta_1 \Delta_2$.

Likewise, the inverse consistency constraint (3) can be simplified in a similar manner. Substituting (7) into (4) and discretizing gives

$$\begin{aligned} C_{\text{ICC}}(u, \tilde{w}) + C_{\text{ICC}}(w, \tilde{u}) \\ = \sum_{k \in \Omega_d} (\mu[k] - \tilde{\eta}[k])^\dagger (\mu[k] - \tilde{\eta}[k]) \\ \quad + (\eta[k] - \tilde{\mu}[k])^\dagger (\eta[k] - \tilde{\mu}[k]). \end{aligned} \quad (10)$$

F. Estimation Procedure

The Fourier series parameterization in (6) is a multiresolution decomposition of the displacement fields. Let $\Omega_d[r] = \Omega_d \setminus G[r]$ represent a family of subsets of Ω_d where $G[r] = \{n \in \Omega_d | r_1 < n_1 < N_1 - r_1; r_2 < n_2 < N_2 - r_2; r_3 < n_3 < N_3 - r_3\}$ and the set subtraction notation $A \setminus B$ is defined as all elements of A not in B . Fig. 2 illustrates the definition of $\Omega_d[r]$ using a two-dimensional example. In practice, the low-frequency basis coefficients are estimated before the higher ones allowing the global image features to be registered before the local features. This is accomplished by replacing (7) by

$$\begin{aligned} u_d[n, r] &= \sum_{k \in \Omega_d[r]} \mu[k] e^{j\langle n, \theta[k] \rangle} \\ w_d[n, r] &= \sum_{k \in \Omega_d[r]} \eta[k] e^{j\langle n, \theta[k] \rangle} \end{aligned} \quad (11)$$

where $r \in \Omega_d$ determines the number of harmonics used to represent the displacement fields. Let the set of multiresolution transformations be defined as $h_d[n, r] = (n/N) + u_d[n, r]$ and $g_d[n, r] = (n/N) + w_d[n, r]$. The components of r are initially set small and are periodically increased throughout the iterative minimization procedure. The set $\Omega_d[r]$ can be replaced by Ω_d when all of the components of r are greater than or equal to $(N - 1)/2$ since the set $G[r]$ is empty. The constants r_1, r_2 , and r_3 represent the largest x_1, x_2 , and x_3 harmonic components of the displacement fields. Each displacement field in (11) is efficiently computed using three $N_1 \times N_2 \times N_3$ fast Fourier transforms (FFTs), i.e., each component of the 3×1 vectors u_d and w_d are computed with a FFT after zeroing out the coefficients not present in the summations.

The image registration problem can be stated mathematically by combining (2)–(4), and (9) and estimating the set of parameters $\{\hat{\mu}[k], \hat{\eta}[k]\}$ that minimize the combined cost function as shown in (12) at the bottom of the next page. The constants σ ,

$$L_c = -\alpha \nabla^2 - \beta \nabla \cdot \nabla + \gamma = \begin{bmatrix} -\alpha \nabla^2 - \beta \frac{\partial^2}{\partial x_1^2} + \gamma & -\beta \frac{\partial^2}{\partial x_1 \partial x_2} & -\beta \frac{\partial^2}{\partial x_1 \partial x_3} \\ -\beta \frac{\partial^2}{\partial x_1 \partial x_2} & -\alpha \nabla^2 - \beta \frac{\partial^2}{\partial x_2^2} + \gamma & -\beta \frac{\partial^2}{\partial x_2 \partial x_3} \\ -\beta \frac{\partial^2}{\partial x_1 \partial x_3} & -\beta \frac{\partial^2}{\partial x_2 \partial x_3} & -\alpha \nabla^2 - \beta \frac{\partial^2}{\partial x_3^2} + \gamma \end{bmatrix}$$

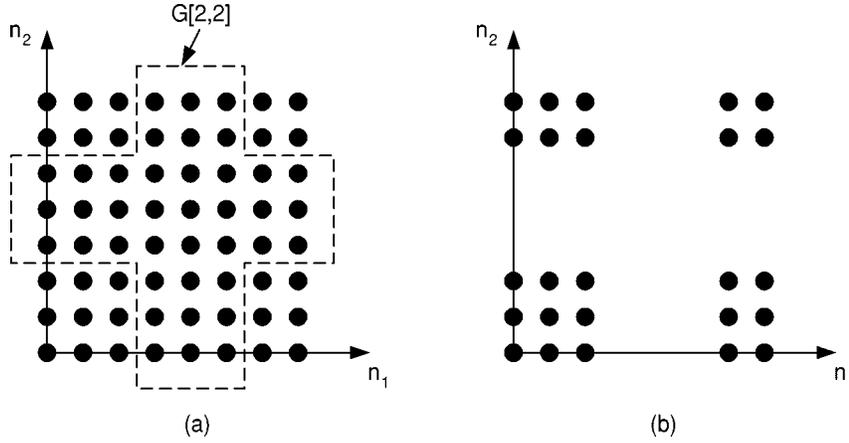


Fig. 2. Example of elements contained in an 8×8 lattice. (a) Ω_d contains all of the points while $G[2, 2]$ contains the points within the dotted lines. (b) Elements of $\Omega_d[2, 2] = \Omega_d \setminus G[2, 2]$.

χ and ρ are used to enforce/balance the constraints. The effects of varying these parameters are studied later in the paper.

The basis coefficients $\{\mu[k], \eta[k]\} = \{a[k] + jb[k], c[k] + jd[k]\}$ are determined by gradient descent using

$$\begin{aligned} a^{(n+1)}[k] &= a^{(n)}[k] + \varepsilon \frac{\partial C(\mu^{(n)}, \eta^{(n)})}{\partial a[k]} \\ b^{(n+1)}[k] &= b^{(n)}[k] + \varepsilon \frac{\partial C(\mu^{(n)}, \eta^{(n)})}{\partial b[k]} \end{aligned} \quad (13)$$

$$\begin{aligned} c^{(n+1)}[k] &= c^{(n)}[k] + \varepsilon \frac{\partial C(\mu^{(n)}, \eta^{(n)})}{\partial c[k]} \\ d^{(n+1)}[k] &= d^{(n)}[k] + \varepsilon \frac{\partial C(\mu^{(n)}, \eta^{(n)})}{\partial d[k]} \end{aligned} \quad (14)$$

where ε is the step size. The superscript (n) corresponds to the value of μ and η at the n th iteration. Each equation above represents a 3×1 vector equation,³ i.e., one equation for each component of the vectors a , b , c , and d . The partial derivatives⁴ of $C(\mu, \eta)$ used in (13) and (14) are

$$\begin{aligned} \frac{\partial C(\mu, \eta)}{\partial a[k]} &= \text{Re}\{I_0[k] + I_2[k]\} \\ \frac{\partial C(\mu, \eta)}{\partial b[k]} &= \text{Im}\{-I_0[k] + I_2[k]\} \end{aligned}$$

³The notation $\partial C / \partial a$ where C is a scalar-valued function and a is a 3×1 vector is defined as the 3×1 vector $[\partial C / \partial a_1, \partial C / \partial a_2, \partial C / \partial a_3]^T$.

⁴The coefficients $\tilde{\mu}[k]$ and $\tilde{\eta}[k]$ are assumed to be constant when computing the partial derivatives.

$$\frac{\partial C(\mu, \eta)}{\partial c[k]} = \text{Re}\{I_1[k] + I_3[k]\}$$

$$\frac{\partial C(\mu, \eta)}{\partial d[k]} = \text{Im}\{-I_1[k] + I_3[k]\}$$

where

$$\begin{aligned} I_0[k] &= \frac{4}{N_1 N_2 N_3} \sum_{n \in \Omega_d} \\ &\cdot \left(\sigma(T_d[Nh_d[n, r]] - S_d[n]) \nabla T_d \Big|_{Nh_d[n, r]} \right. \\ &\quad \left. + \chi(u_d[n, r] - \tilde{u}_d[n, r]) \right) e^{j(x, \theta[k])} \end{aligned}$$

$$\begin{aligned} I_1[k] &= 4N_1 N_2 N_3 \sum_{n \in \Omega_d} \\ &\cdot \left(\sigma(S_d[Ng_d[n, r]] - T_d[n]) \nabla S_d \Big|_{Ng_d[n, r]} \right. \\ &\quad \left. + \chi(w_d[n, r] - \tilde{w}_d[n, r]) \right) e^{j(x, \theta[k])} \end{aligned}$$

$$I_2[k] = 4\rho D^2[k] \mu[k]$$

$$I_3[k] = 4\rho D^2[k] \eta[k].$$

The terms I_0 and I_1 are 3×1 vectors and are efficiently computed for all k using 3-D FFTs; each component of I_0 and I_1 requires a separate 3-D FFT for a total of six FFTs. The notation $\nabla T_d|_{Nh_d[n, r]}$ represents the gradient of T_d evaluated at the point $Nh_d[n, r]$. The gradient is computed using symmetric dif-

$$\begin{aligned} C(\mu, \eta, r) &= \\ &\sigma \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} |T_d[Nh_d[n, r]] - S_d[n]|^2 + |S_d[Ng_d[n, r]] - T_d[n]|^2 + \chi \frac{1}{N_1 N_2 N_3} \sum_{n \in \Omega_d} \\ &\cdot \|u_d[n, r] - \tilde{u}_d[n, r]\|^2 + \|w_d[n, r] - \tilde{w}_d[n, r]\|^2 + \rho \sum_{k \in \Omega_d[r]} \mu^\dagger[k] D^2[k] \mu[k] + \eta^\dagger[k] D^2[k] \eta[k]. \end{aligned} \quad (12)$$

ference equations⁵ and trilinear interpolation is used to evaluate the gradient at the point $Nh_d[n, r]$. The notation $\nabla S_d|_{Ng_d[n, r]}$ is defined similarly. The gradient descent insures that μ and η have complex conjugate symmetry. Notice that I_0 and I_1 are the Fourier transforms of a real signal implying that I_0 and I_1 have complex conjugate symmetry. This in turn implies that the gradients in (13) and (14) have conjugate symmetry and, therefore, the basis coefficients μ and η have complex conjugate symmetry.

The steps involved in estimating the basis coefficients μ and η are summarized in the following algorithm.

Algorithm:

1. Set $\mu[k] = \eta[k] = 0$ for $k \in \Omega_d$ and set $r = [1, 1, 1]^T$.
2. Compute $g_d^{-1}[n, r]$ using the procedure described in Section II-C and set $\tilde{w}_d[n, r] = g_d^{-1}[n, r] - (n/N)$.
3. Compute new forward basis coefficients $\mu[k] = a[k] + jb[k]$ using (13).
4. Compute the forward displacement field $u_d[n, r]$ using (11).
5. Compute $h_d^{-1}[n, r]$ using the procedure described in Section II-C and set $\tilde{u}_d[n, r] = h_d^{-1}[n, r] - (n/N)$.
6. Compute new reverse basis coefficients $\eta[k] = c[k] + jd[k]$ using (14).
7. Compute the reverse displacement field $w_d[n, r]$ using (11).
8. If the criteria is met to increase the number of basis functions then set $r = r + 1$, and set the new coefficients in (11) to zero.
9. If the algorithm has not converged or reached the maximum number of iterations goto step 2.
10. Use forward displacement field $u_d[n]$ to transform the template image $T_d[n]$ and reverse displacement field $w_d[n]$ to transform the target image $S_d[n]$.

III. RESULTS

Two sets of experiments were performed to demonstrate the linear-elastic consistent image registration algorithm. The first set of experiments were designed to show the effect of varying the transformation parameters and the second set of experiments demonstrate typical results that can be expected using a multiresolution approach to transform full resolution MRI and CT data.

A. Parameter Evaluation

Two MRI and two CT image volumes were used to investigate the effect of varying the parameters used in the consistent image registration algorithm. The data sets were collected from

⁵The gradient of T_d evaluated at n is computed as $[(N_1/2)(T_d[n_1 + 1, n_2, n_3] - T_d[n_1 - 1, n_2, n_3]), (N_2/2)(T_d[n_1, n_2 + 1, n_3] - T_d[n_1, n_2 - 1, n_3]), (N_3/2)(T_d[n_1, n_2, n_3 + 1] - T_d[n_1, n_2, n_3 - 1])]^T$

different individuals using the same MR and CT machines and the same scan parameters. The MRI data sets correspond to two normal adults and the CT data sets correspond to two 3-mo.-old infants, one normal and one abnormal (bilateral coronal synostosis). The MRI and CT data sets were chosen to test registration algorithm when matching anatomies with similar and dissimilar shapes, respectively.

The MRI data were preprocessed by normalizing the image intensities, correcting for translation and rotation, and segmenting the brain from the head using Analyze (Mayo Clinic, Rochester, MN). The translation aligned the anterior commissure points, and the rotation aligned the corresponding axial and sagittal planes containing the anterior and posterior commissure points, respectively. The MRI data sets were down-sampled and zero padded to form a $64 \times 64 \times 80$ voxel lattice. The CT data sets were corrected for translation and rotation and down-sampled to form a $64 \times 64 \times 48$ voxel lattice. The translation aligned the basion skull landmarks, and the rotation aligned the corresponding Frankfort horizontal and midsagittal planes, respectively.

Tables I and II show the results of 32 experiments for to MRI-to-MRI and CT-to-CT registration, respectively, as the weighting values ρ and χ were varied. The weight for the similarity cost σ was set to one for all of the experiments. The values of ρ and χ ranged from 0.0 to 0.0125 and 0 to 5000 for the MRI-to-MRI experiments, respectively. The values of ρ and χ ranged from 0.0 to 0.00125 and 0 to 1275 for the CT-to-CT experiments, respectively. The gradient descent step size was set to 0.00004 for the MRI experiments and 0.0001 for the CT experiments. The difference in the parameters is due to the different intensity characteristics of each modality. These ranges can be used as a guide for determining parameter settings for registration of other modalities. We have found that there is no need to adjust the parameters for additional data sets of the same modality.

The data sets were registered initially with zero and first-order harmonics. Each experiment was run for 1000 iterations unless the algorithm failed to converge. After every 100th iteration, the maximum harmonic was increased by one. Each experiment that ran for 1000 iterations took approximately 1.5 h to run on an AlphaPC clone using a single 667-MHz, alpha 21 264 processor. It is expected that this time can be significantly decreased by optimizing the code and using a better optimization technique than gradient descent. In some of the experiments the Jacobian of the transformation went negative due to insufficient regularization or due to a bad choice of parameters. In these cases, the experiments were stopped before the Jacobian went negative to report the results. The numbers reported for the Similarity cost C_{SIM} , the linear elasticity cost C_{REG} , and the inverse consistency cost C_{ICC} , were scaled by 10 000 for presentation.

Experiments MRI01 and CT01 correspond to unconstrained estimation in which the forward and reverse transformations were estimated independently. These experiments produced the worst registration results as evident by the largest values of C_{SIM} , C_{REG} , and C_{ICC} in the respective tables. These experiments were stopped before the 1000th iteration because the Jacobian went negative during the gradient descent. This was expected since the regularization terms help prevent the Jacobian

TABLE I
MRI-TO-MRI CONSISTENT LINEAR-ELASTIC REGISTRATION EXPERIMENTAL PARAMETERS AND FINAL VALUES

Experiment	REG	ICC	Iterations	ρ	χ	C_{SIM}		C_{REG}		C_{ICC}		Jacobian(h)		Jacobian(g)	
						For.	Rev.	For.	Rev.	For.	Rev.	min	1/max	min	1/max
MRI01	No	No	800	0.0	0	217	185	342000	243000	2.76	2.41	0.0268	0.408	0.237	0.273
MRI02	No	Yes	1000	0.0	1275	278	207	127000	119000	0.00962	0.00804	0.340	0.651	0.585	0.387
MRI03	No	Yes	1000	0.0	2500	328	245	87100	73100	0.00423	0.00349	0.344	0.682	0.665	0.392
MRI04	No	Yes	1000	0.0	5000	399	301	59600	46000	0.00190	0.00158	0.365	0.707	0.700	0.414
MRI05	Yes	No	800	0.000125	0	183	159	302000	217000	2.45	2.17	0.0458	0.485	0.417	0.283
MRI06	Yes	Yes	1000	0.000125	1275	289	216	96400	87400	0.00722	0.00587	0.374	0.694	0.667	0.411
MRI07	Yes	Yes	1000	0.000125	2500	335	251	72500	59700	0.00342	0.00281	0.369	0.702	0.692	0.410
MRI08	Yes	Yes	1000	0.000125	5000	403	305	53300	40900	0.00165	0.00138	0.383	0.719	0.717	0.425
MRI09	Yes	No	1000	0.00125	0	276	228	72900	48400	1.01	0.922	0.381	0.686	0.658	0.486
MRI10	Yes	Yes	1000	0.00125	1275	339	254	45700	37700	0.00288	0.00215	0.489	0.737	0.738	0.502
MRI11	Yes	Yes	1000	0.00125	2500	372	279	39300	31200	0.00153	0.00123	0.479	0.757	0.756	0.498
MRI12	Yes	Yes	1000	0.00125	5000	431	327	32900	25100	0.000838	0.000722	0.477	0.772	0.778	0.502
MRI13	Yes	No	800	0.0125	0	449	338	20300	16300	0.413	0.392	0.501	0.779	0.754	0.563
MRI14	Yes	Yes	700	0.0125	1275	478	368	17700	13600	0.0014	0.00123	0.555	0.790	0.788	0.566
MRI15	Yes	Yes	700	0.0125	2500	502	390	16400	12700	0.000686	0.000642	0.569	0.799	0.798	0.568
MRI16	Yes	Yes	500	0.0125	5000	602	479	14500	11000	0.00104	0.000952	0.601	0.786	0.790	0.602

REG: Linear-elastic regularization; ICC: Inverse consistency constraint; ρ : REG weight; χ : ICC weight; $\sigma = 1$: Similarity cost weight. The values of C_{SIM} , C_{REG} , and C_{ICC} have been scaled by 10000

TABLE II
CT-TO-CT CONSISTENT LINEAR-ELASTIC REGISTRATION EXPERIMENTAL PARAMETERS AND FINAL VALUES

Experiment	REG	ICC	Iterations	ρ	χ	C_{SIM}		C_{REG}		C_{ICC}		Jacobian(h)		Jacobian(g)	
						For.	Rev.	For.	Rev.	For.	Rev.	min	1/max	min	1/max
CT01	No	No	600	0.0	0	170	104	98000	113000	7.97	5.95	0.0272	0.487	0.336	0.194
CT02	No	Yes	1000	0.0	300	82.8	58.7	204000	144000	0.0173	0.0144	0.170	0.444	0.347	0.209
CT03	No	Yes	1000	0.0	600	99.1	73.8	130000	89500	0.00799	0.00652	0.198	0.533	0.449	0.248
CT04	No	Yes	1000	0.0	1275	135	106	72600	50200	0.00345	0.00290	0.238	0.578	0.577	0.300
CT05	Yes	No	600	0.0000125	0	130	75.1	169000	181000	10.2	7.38	0.0158	0.325	0.0980	0.132
CT06	Yes	Yes	1000	0.0000125	300	88.8	60.0	187000	132000	0.0159	0.0132	0.173	0.463	0.374	0.213
CT07	Yes	Yes	1000	0.0000125	600	99.7	75.0	123000	84400	0.00758	0.0062	0.201	0.542	0.463	0.251
CT08	Yes	Yes	1000	0.0000125	1275	135	107	70000	48400	0.00340	0.00281	0.241	0.583	0.584	0.303
CT09	Yes	No	700	0.000125	0	109	65.9	155000	148000	9.15	6.87	0.0211	0.419	0.165	0.128
CT10	Yes	Yes	1000	0.000125	300	90.2	66.1	118000	83700	0.0102	0.00853	0.209	0.561	0.516	0.249
CT11	Yes	Yes	1000	0.000125	600	105	79.7	87200	60800	0.00538	0.00449	0.229	0.590	0.567	0.278
CT12	Yes	Yes	1000	0.000125	1275	139	110	55500	38800	0.00262	0.00225	0.269	0.618	0.620	0.327
CT13	Yes	No	1000	0.00125	0	108	76.0	36600	34400	4.70	4.25	0.348	0.639	0.636	0.373
CT14	Yes	Yes	1000	0.00125	300	114	90.1	35500	26200	0.00534	0.00451	0.361	0.704	0.701	0.388
CT15	Yes	Yes	1000	0.00125	600	129	103	30600	22600	0.00233	0.00208	0.377	0.719	0.717	0.397
CT16	Yes	Yes	1000	0.00125	1275	169	136	22400	17000	0.00102	0.000977	0.469	0.768	0.771	0.477

REG: Linear-elastic regularization; ICC: Inverse consistency constraint; ρ : REG weight; χ : ICC weight; $\sigma = 1$: Similarity cost weight. The values of C_{SIM} , C_{REG} , and C_{ICC} have been scaled by 10000

from going negative. The similarity cost is the lowest for these experiments since the algorithm finds the best match between the images without any constraint preventing the Jacobian from going negative.

Experiments MRI05, MRI09, MRI13, CT05, CT09, and CT13 demonstrate the effect of estimating the forward and reverse transformations independently while varying ρ the weight of the linear elastic cost. As before, the large difference between the forward and reverse displacement fields as reported by C_{SIM} confirms that linear elasticity alone is not sufficient to guarantee that the forward and reverse transformations are inverses of one another. However, the linear elasticity constraint did improve the transformation over the unconstrained case because the minimum Jacobian and the inverse of the maximum Jacobian is far from being singular.

Experiments MRI02, MRI03, MRI04, CT02, CT03, and CT04 demonstrate the effect of jointly estimating the forward

and reverse transformations without enforcing the linear elasticity constraint. The C_{ICC} values for these experiments are much lower than the previous cases since they are being minimized. The forward and reverse transformations are inverses of each other when C_{ICC} are zero so that the smaller the costs C_{ICC} , the closer the transformations are to being inverses of each other.

The remaining experiments show the effect of jointly estimating the forward and reverse transformations while varying the weights on both the linear elasticity constraint and the inverse consistency constraint. These experiments show that it is possible to find a set of parameters that produce better results using both constraints than only using one or none. Notice that increasing the constraint weights causes the similarity cost to increase indicating a worse intensity match between the images. At the same time, the worst case values of the Jacobian increase as the constraint weights increase indicating less spatial distor-

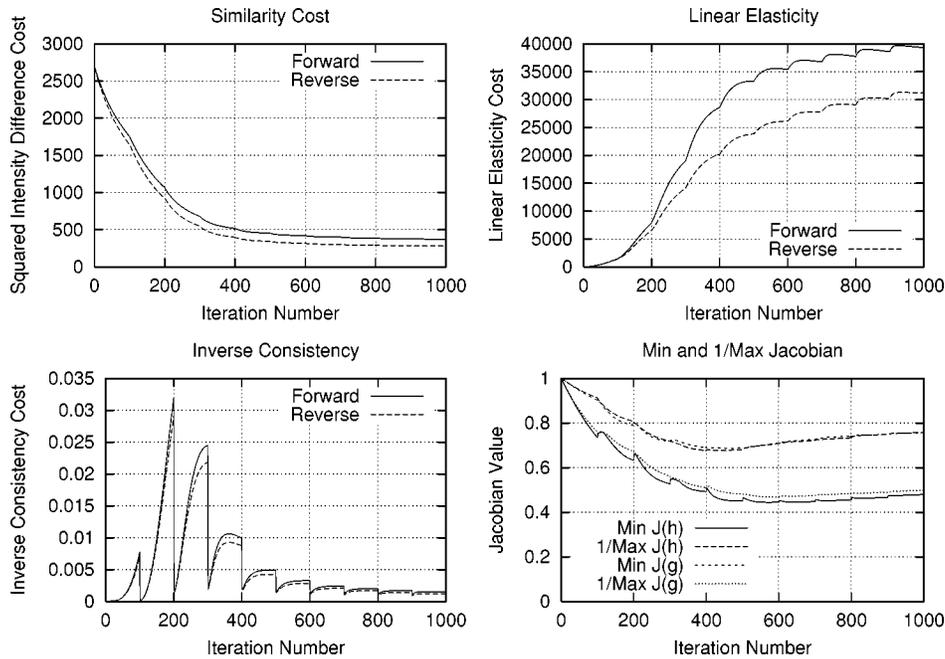


Fig. 3. Statistics associated with the MRI11 experiment.

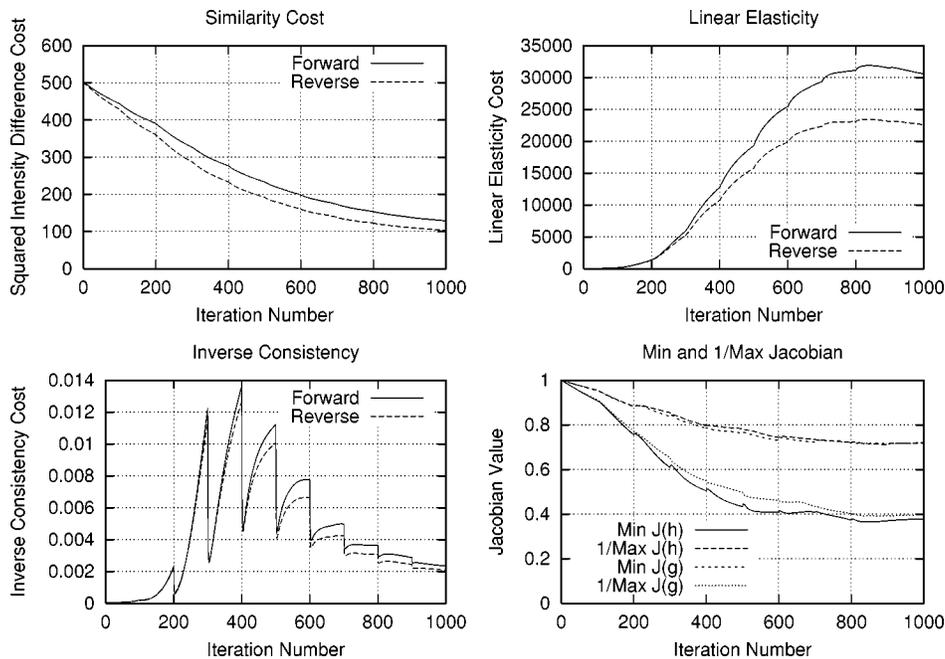


Fig. 4. Statistics associated with the CT15 experiment.

tion. The optimal set of parameters should be chosen to provide a good intensity match while producing the least amount of spatial distortion as measured by the Jacobian and an acceptable level of inverse consistency error.

The time series statistics for experiments MRI11 and CT15 are shown in Figs. 3 and 4, respectively. These figures show that the gradient descent algorithm converged for each set of transformation harmonics. In both cases, the similarity cost C_{SIM} decreased at each iteration while the prior terms increased before decreasing. Notice that the inverse consistency constraint increased as the images deformed for each particular harmonic resolution. Then when the number of harmonics were increased,

the inverse constraint decreased before increasing again. This is due to the fact that a low-dimensional Fourier series does not have the degrees of freedom (DOFs) to faithfully represent the inverse of a low-dimensional Fourier series. This is seen by looking at the high dimensionality of a Taylor series representation of the inverse transformation. Finally, notice that the inverse consistency constraint caused the extremal Jacobian values of the forward and reverse transformations to track together. The extremal Jacobian values correspond to the worst case distortions produced by the transformations.

Fig. 5 shows the effect of varying χ and ρ on the inverse consistency cost C_{ICC} as a function of iteration. Fig. 5(a) shows

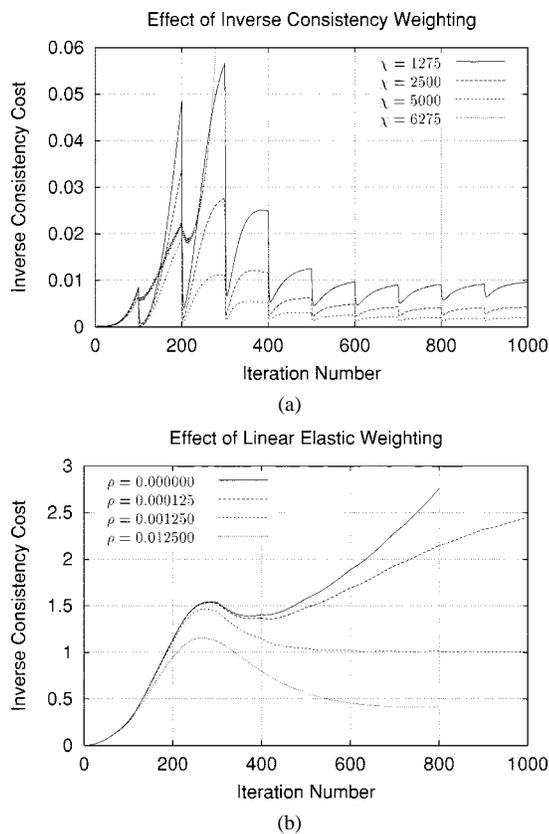


Fig. 5. Plots demonstrating the effect of (a) modifying the inverse consistency constraint weight χ while $\rho = 0$ and (b) modifying the linear-elasticity regularization weight ρ while $\chi = 0$.

that C_{ICC} increases with iteration and then drops every 100 iterations when additional parameters (DOFs) are added to the transformation. The curves decrease in amplitude as χ is increased until χ becomes too large and the algorithm fails to converge. Fig. 5(b) shows that C_{ICC} increases as the linear elasticity weight ρ is increased. This makes sense because the two regularization terms fight one another. The inverse consistency cost increases as the linear elasticity cost is penalized more.

B. Multiresolution Registration

A spatial and frequency multiresolution procedure was used to estimate the full resolution registration of the data sets used in the previous section. In this approach, global structures are matched before local to reduce the likelihood of incorrect local registration errors and to increase convergence. Table III shows the number of iterations, the harmonic increment iteration, the initial number of harmonics, and the final number of harmonics at each resolution. This schedule proceeds from low resolution to high resolution starting at one eighth the spatial resolution and increases to full resolution.

A set of parameters were chosen from Tables I and II that gave a good tradeoff between the image intensity match, the inverse consistency of the transformations, and low spatial distortion. The parameters used were time step = 0.00004, $\sigma = 1.0$, $\rho = 0.00125$, and $\chi = 2500$ for the MRI-to-MRI registration and time step = 0.0001, $\sigma = 1.0$, $\rho = 0.00125$, and $\chi = 600$ for the CT-to-CT registration. These parameters were used at all resolutions.

TABLE III
MULTIRESOLUTION ITERATION SCHEDULE USED TO GENERATE THE FULL RESOLUTION MRI-TO-MRI AND CT-TO-CT REGISTRATION RESULTS

Resolution	Iterations	Harmonic Increment Iteration	Initial Number of Harmonics	Final Number of Harmonics
1/8	500	100	1	5
1/4	500	100	5	9
1/2	100	50	9	10
1	10	10	10	10

Fig. 6 shows three transverse sections from the 3-D result from the template T , deformed target $S \circ g$, target S , deformed template $T \circ h$ MRI data sets. The first two columns and the last two columns should look like for a good registration. These pairs of columns look similar with respect to the global structures but have small local differences as seen by the difference images shown in the first two columns of Fig. 7. Notice that the outer contour of the deformed images match their respective target data sets and that there is good correspondence in the region of the ventricles. The local mismatch is mostly due to differences in the topology of the gray matter folds and due to the low-frequency Fourier series parameterization of the transformations.

The last two columns of Fig. 7 shows the normed difference between the forward and reverse transformations for the MRI-to-MRI experiment. These figures show the spatial locations of where the forward and reverse transformations have the largest inverse consistency errors. The range on the difference for the entire 3-D volume is from 0 to 0.002234. This maximum difference corresponds to a registration error between 0.571 and 0.749 voxel units.⁶ Notice that most of the error is internal to the brain and that most of the error appears in the cortex regions. The similarities between the absolute difference intensity images to the normed transformation difference images suggest that most of the inverse consistency error occurs where the transformed images are still mismatched. A further description and additional figures showing the effect of using or not using the inverse consistency cost as it relates to the spatial inverse consistency error can be found in [31].

Three transverse slices from the 3-D full resolution CT-to-CT experiment are shown in Fig. 8. Notice the good global registration of the corresponding CT data sets. The first two columns of Fig. 9 shows the absolute intensity difference between these slices. As before, the errors show up along the boundaries of the objects. The last two columns of this figure show the normed difference between the forward and reverse transformations. Again we see similarities between the intensity differences and the transformation differences. The maximum inverse consistency error for this experiment is between 0.871 and 1.16 voxel units.⁷

The MRI-to-MRI registration following the schedule in Table III took approximately 4, 40, 60, and 55 min to compute at the $32 \times 32 \times 40$, $64 \times 64 \times 80$, $128 \times 128 \times 160$, and $256 \times 256 \times 320$ voxel resolutions, respectively. The CT-to-CT registration took approximately 2, 21, 33, and 30 min to

⁶The minimum was computed as 256×0.002234 and the maximum as 320×0.002234 .

⁷The minimum was computed as 192×0.004538 and the maximum as 256×0.004538 .

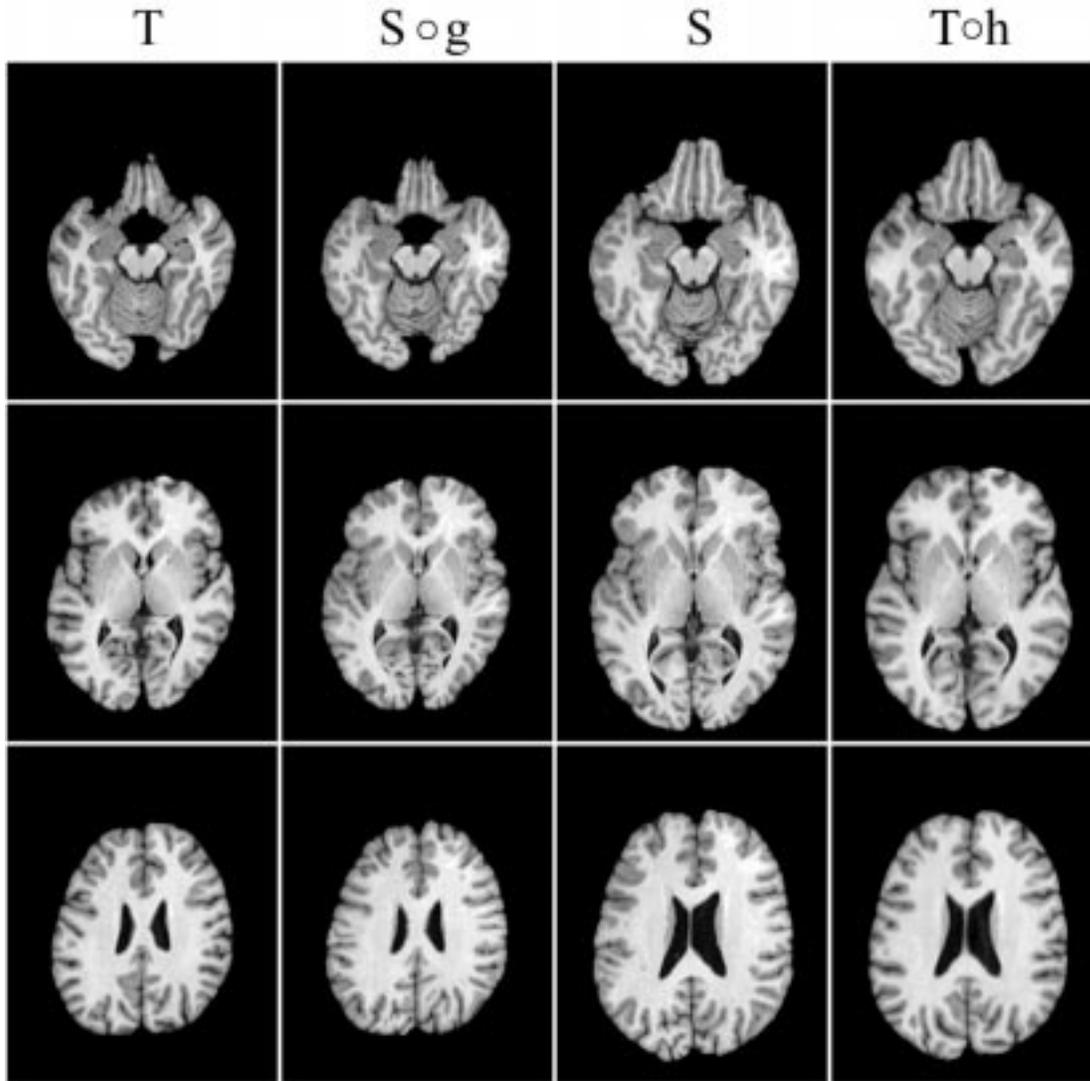


Fig. 6. Transverse slices 109, 135, and 165 (rows top to bottom) from the full resolution MRI-to-MRI registration experiment. The columns from left to right correspond to the template T , the deformed target $S \circ g$, the target S and the deformed template $T \circ h$. The intensities are on a range from 0 to 1.0.

compute at the $32 \times 32 \times 24$, $64 \times 64 \times 48$, $128 \times 128 \times 96$, and $256 \times 256 \times 192$ voxel resolutions, respectively. All times are for an AlphaPC clone using a single 667-MHz, alpha 21 264 processor.

IV. DISCUSSION

The experiments presented in this paper were designed to test the validity of the new inverse transformation consistency constraint as applied to a linear-elastic transformation algorithm. As such, there was no effort made to optimize the rate of convergence of the algorithm. The convergence rate of the algorithm can be greatly improved by using a more efficient optimization technique than gradient descent such as conjugate gradient at each parameterization resolution. In addition, a convergence criteria can be used to determine when to increment the number of parameters in the model. The CT data used in the experiments was selected to stress the registration algorithm. The convergence of the algorithm would have been much faster if the data sets were adjusted for global scale initially.

A. Measurement of Transformation Distortion

It is important to track both the minimum and maximum values of the Jacobian during the estimation procedure. The Jacobian measures the differential volume change of a point being mapped through the transformation. At the start of the estimation, the transformation is the identity mapping and, therefore, has a Jacobian of one. If the minimum Jacobian goes negative, the transformation is no longer a one-to-one mapping and as a result folds the domain inside out [30]. Conversely, the reciprocal of the maximum value of the Jacobian corresponds to the minimum value of the Jacobian of the inverse mapping. Thus, as the maximum value of the Jacobian goes to infinity, the minimum value of the Jacobian of the inverse mapping goes to zero. In the present approach, the inverse transformation consistency constraint was used to penalize transformations that deviated from their inverse transformation. A limitation of this approach is that cost function in (3) is an average metric and cannot enforce the pointwise constraints that $\min_x \{J(h)\} = 1 / \max_x \{J(g)\}$ and

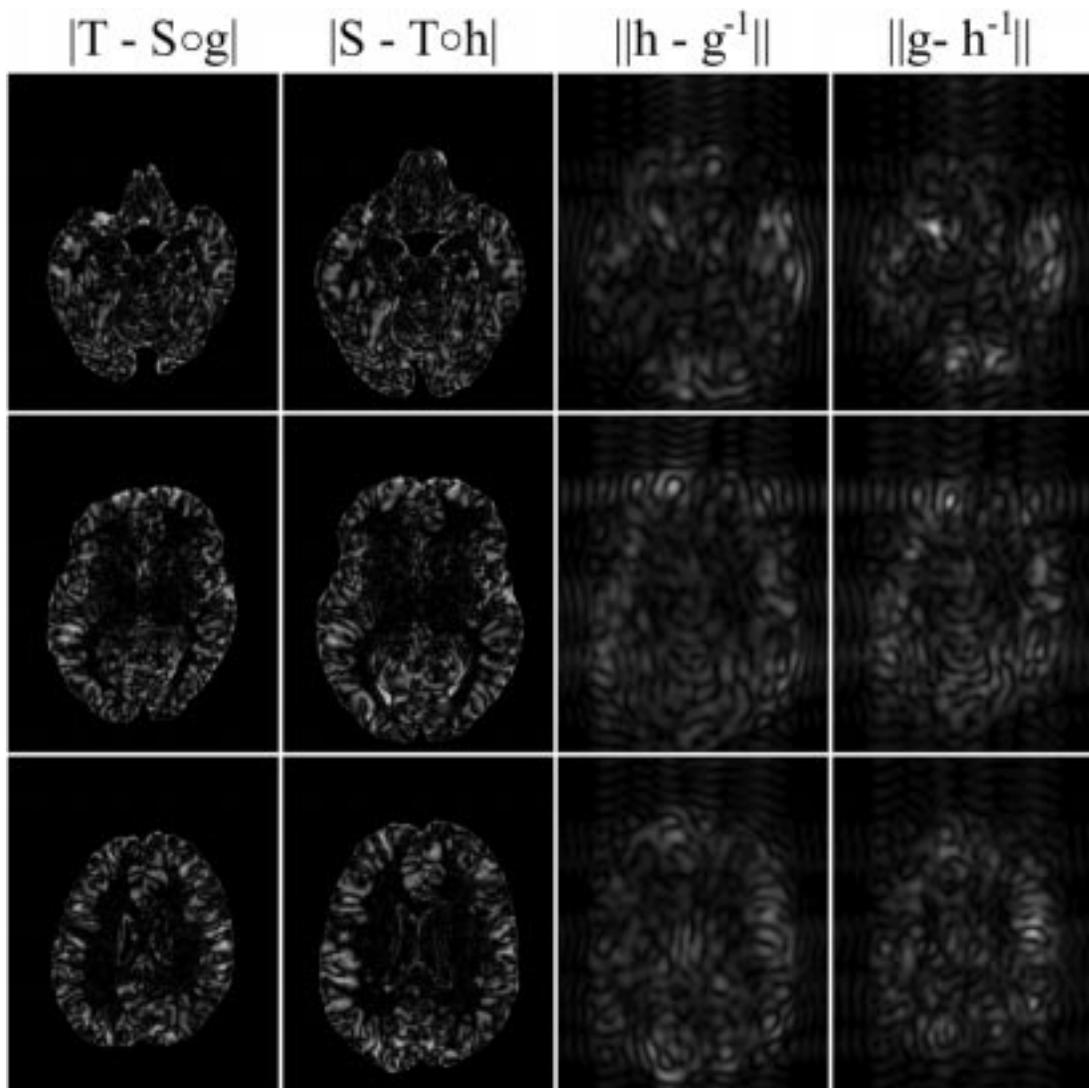


Fig. 7. The first two columns correspond to the absolute intensity difference between the template and the deformed target $|T - S \circ g|$ (column one) and between the target and the deformed template $|S - T \circ h|$ (column two) in Fig. 6. The intensity range for the absolute difference is on a range from 0 to 0.892. The last two columns correspond to the normed difference between the forward and inverse of the reverse transformation $\|h - g^{-1}\|$ (column three), and between the reverse transformation and the inverse of the forward transformation $\|g - h^{-1}\|$ (column four). The intensity range for the normed transformation differences is from 0 to 0.002234.

$\min_x \{J(g)\} = 1 / \max_x \{J(h)\}$. This point is illustrated by Tables I and II by the fact that the minimum values of $J(h)$ and $J(g)$ differ from the reciprocal of the maximum values of $J(g)$ and $J(h)$, respectively. However, these extremal Jacobian values do give an upper bound on the worst case distortions produced by the transformations demonstrating the consistency between the forward and reverse transformations.

B. Spatial Multiresolution

The minimization problem is discretized so it can be implemented on a digital computer. The higher the sampling rate the more accurate the discrete approximation is to the continuous case. An advantage of discretizing a continuous formulation is that the problem can be solved at different spatial sampling rates. The approach that is taken is to solve the minimization at a coarse resolution initially to approximate the solution. The

advantage of solving the problem on a coarse grid is that the algorithm requires fewer computations/iteration than a finer grid. This results in reduced computation time at low resolution. Each time the resolution of the grid is increased by a factor of two in each dimension, the computation time increases by a factor of eight. The drawback of solving the problem at low resolution is that there can be significant registration errors due to the loss of detail in the down sampling procedure. The tradeoff between quicker execution times at low resolution and more accurate registration at higher resolution can be exploited by solving the registration problem at low spatial resolution during the initial iterations to approximate the result and then increasing the spatial resolution to get a more accurate result at the later iterations.

The spatial multiresolution approach works well with the frequency multiresolution approach provided by increasing the number of harmonics used to represent the displacement fields. The number of harmonics used to represent the displacement fields is initially set small and then increased as the number

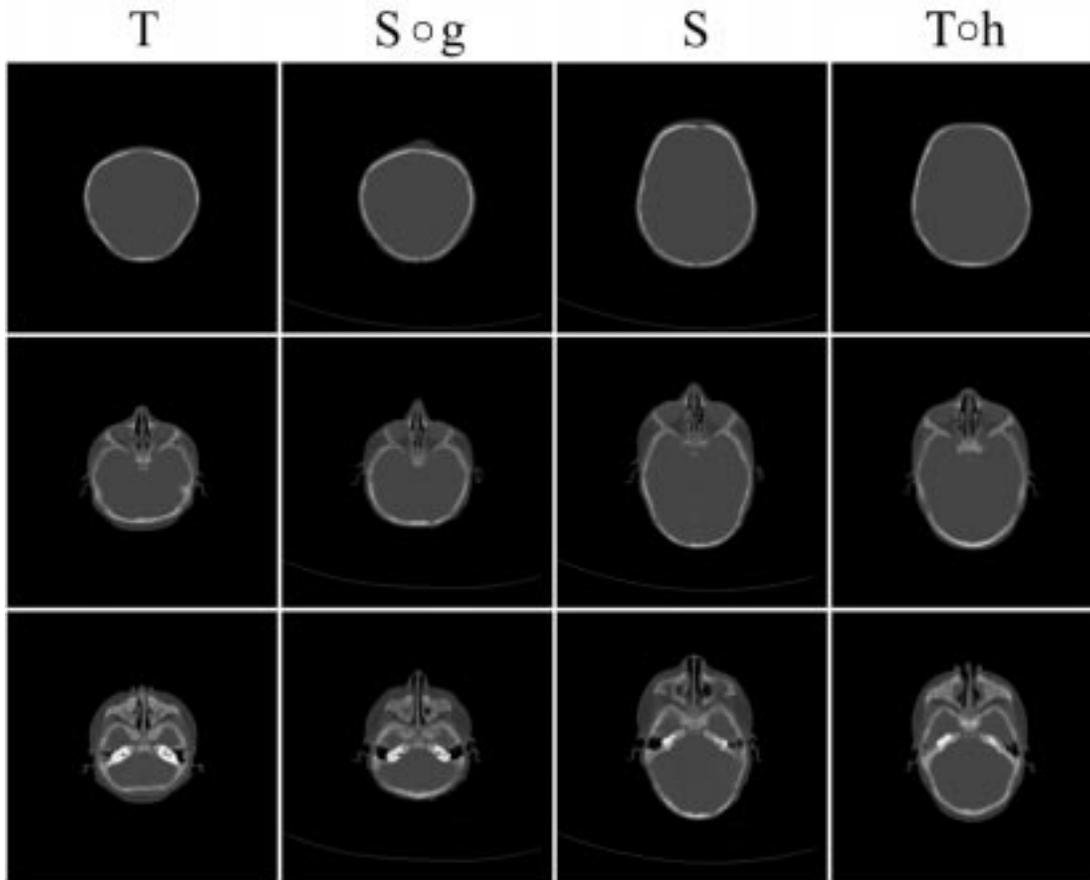


Fig. 8. Transverse slices 068, 116, and 124 (rows top to bottom) from the full resolution CT-to-CT registration experiment. The columns from left to right correspond to the template T , the deformed target $S \circ g$, the target S and the deformed template $T \circ h$. The intensities are on a range from 0 to 1.0.

of iterations are increased. A low-frequency registration result is an approximation of the desired high-frequency registration result. Computing the gradient descent for a low-frequency basis coefficient at low spatial resolution gives approximately the same answer as using high spatial resolution but the computational burden is much less.

C. Comparison to Other Methods

Other investigators have proposed methods for enforcing pairwise consistent transformations. For example, Woods *et al.* [27] computes all pairwise registrations of a population of image volumes using a linear transformation model, i.e., a 3×3 matrix transformation. They then average the transformation from T to S with all the transformations from T to X to S . The original transformation from T to S is replaced with average transformation. The procedure is repeated for all the image pairs until convergence. This technique is limited by the fact that it cannot be applied to two data sets. Also, there is no guarantee that the generated set of consistent transformations are valid. For example, a poorly registered pair of images can adversely effect all of the pairwise transformations.

The method proposed in this paper is most similar to the heuristic approach described by Thirion [6]. Thirion's idea was to iteratively estimate the forward h , reverse g , and residual $r = h \circ g$ transformations in order to register the images T and S . At each iteration, half of the residual r is added to h and half

of the residual r is mapped through h and added to g . After performing this operation, $h \circ g$ is close to the identity transformation. The advantage of Thirion's method is that it enforces the inverse consistency constraint without having to explicitly compute the inverse transformations as in (3). The residual method is an approximation to the inverse consistency method in that the residual method approximates the correspondences between the forward and reverse transformations while the inverse consistency method computes those correspondences. Thus, the residual approach only works under a small deformation assumption since the residual is computed between points that do not correspond to one another. This drawback limits the residual approach to small deformations and it, therefore, cannot be extended to nonlinear transformation models. On the other hand, the approach presented in this paper can be extended to the nonlinear case by modifying the procedure used to calculate the inverse transformation to include nonlinear transformations.

D. Limitations of Diffeomorphic Transformations

Diffeomorphic transformations are valid for registering images collected from the same individual imaged by two different modalities such as MRI and CT, but it is not necessarily valid when registering images before and after surgery. Likewise, a diffeomorphic mapping assumption may be valid for registering MRI data from two different normal individuals if the goal

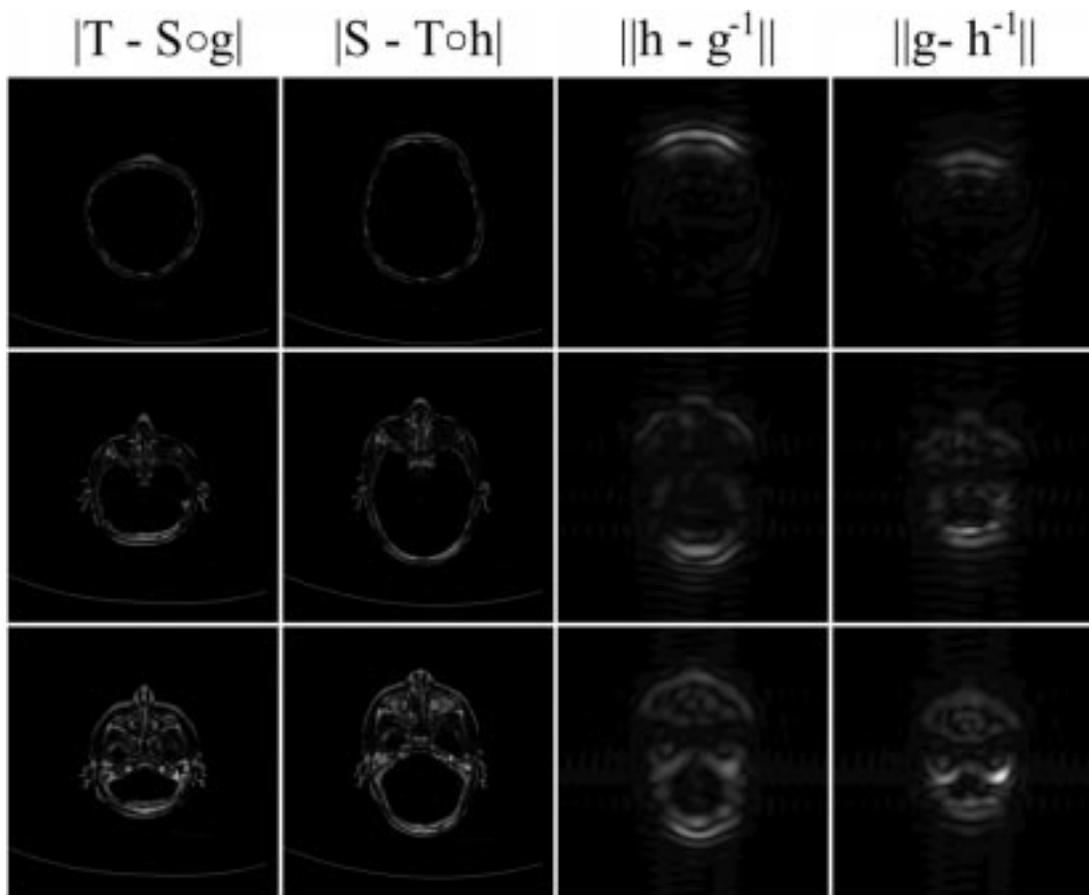


Fig. 9. The first two columns correspond to the absolute intensity difference between the template and the deformed target $|T - S \circ g|$ (column one) and between the target and the deformed template $|S - T \circ h|$ (column two) in Fig. 6. The intensity range for the absolute difference is on a range from 0 to 0.776. The last two columns correspond to the normed difference between the forward and inverse of the reverse transformation $\|h - g^{-1}\|$ (column three), and between the reverse transformation and the inverse of the forward transformation $\|g - h^{-1}\|$ (column four). The intensity range for the normed transformation differences is from 0 to 0.004538.

is to match the deep nuclei of the brain, but it may not be valid for the same data sets if the goal is to match the sulcal patterns.

Alternatively, diffeomorphic transformations may be used to identify areas where two image volumes differ topologically by analyzing the properties of the resulting transformation. For example, consider the problem of matching an MRI image with a tumor to one without a tumor. A possibly valid diffeomorphic transformation would be one that registers all of the corresponding brain structures by shrinking the tumor to a small point. Such a transformation would have an unusually small Jacobian which could be used to detect or identify the location of the tumor. Conversely, consider the inverse problem of matching the image without the tumor to the one with the tumor. A valid registration in this case is to register all of the corresponding brain structures by allowing the transformation to “tear” (i.e., not be diffeomorphic) at the site of the tumor [29]. Just as valid could be a diffeomorphic transformation that registers all of the corresponding brain structures by allowing the transformation to stretch at the site of the tumor.

As in the previous examples, we assume that a valid transformation is diffeomorphic everywhere except possibly in regions where the source and target images differ topologically, e.g., in the neighborhood of the tumor. These ideas can be extended to

nondiffeomorphic mappings by including the proper boundary conditions around regions that differ topologically.

V. SUMMARY AND CONCLUSION

This paper presented a new algorithm for jointly estimating a consistent set of transformations that map one image to another and *vice versa*. A new parameterization based on the Fourier series was presented and was used to simplify the discretized linear-elasticity constraint. The Fourier series parameterization is simpler than our previous parameterizations and each basis coefficient can be interpreted as the weight of a harmonic component in a single coordinate direction. The algorithm was tested on both MRI and CT data. It was found that the unconstrained estimation leads to singular or near singular transformations. It was also shown that the linear-elastic constraint alone is not sufficient to guarantee that the forward and reverse transformations are inverses of one another. Results were presented that suggest that even though the inverse consistency constraint is not guaranteed to generate nonsingular transformations, in practice it may be possible to use the inverse consistency as the only constraint. Finally, it was shown that the most consistent transformations were generated using both the inverse consistency and the linear-elastic constraints.

ACKNOWLEDGMENT

The authors would like to thank J. Haller and M. W. Vannier of the Department of Radiology, The University of Iowa for providing the MRI data. They would also like to thank J. L. Marsh of the Department of Surgery, Washington University School of Medicine for providing the CT data. Finally, they would like to thank M. I. Miller and S. C. Joshi for their helpful insights over the last eight plus years.

REFERENCES

- [1] F. L. Bookstein, *The Measurement of Biological Shape and Shape Change*. New York: Springer-Verlag, 1978, vol. 24, Lecture Notes in Biomathematics.
- [2] —, *Morphometric Tools for Landmark Data*. New York: Cambridge Univ. Press, 1991.
- [3] A. C. Evans, C. Beil, S. Marret, C. J. Thompson, and A. Hakim, "Anatomical-functional correlation using an adjustable MRI-based region of interest atlas with positron emission tomography," *J. Cereb. Blood Flow Metab.*, vol. 8, pp. 513–530, 1988.
- [4] K. Rohr, M. Fornefett, and H. S. Stiehl, "Approximating thin-plate splines for elastic registration: Integration of landmark errors and orientation attributes," in *Information Processing in Medical Imaging*, A. Kuba and M. Samal, Eds. Berlin, Germany: Springer-Verlag, June 1999, LCNS 1613, pp. 252–265.
- [5] C. A. Davatzikos, J. L. Prince, and R. N. Bryan, "Image registration based on boundary mapping," *IEEE Trans. Med. Imag.*, vol. 15, pp. 112–115, Feb. 1996.
- [6] J. P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.*, vol. 2, pp. 243–260, 1998.
- [7] G. Subsol, "Crest lines for curve-based warping," in *Brain Warping*. San Diego, CA: Academic, 1999, pp. 241–262.
- [8] C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen, "Accurate three-dimensional registration of CT, PET, and/or MR images of the brain," *J. Comput. Assist. Tomogr.*, vol. 13, no. 1, pp. 20–26, 1989.
- [9] P. M. Thompson and A. W. Toga, "A surface-based technique for warping three-dimensional images of the brain," *IEEE Trans. Med. Imag.*, vol. 15, pp. 1–16, Aug. 1996.
- [10] N. Krahnstover and C. Lorenz, "Development of a point-based shape representation of arbitrary three-dimensional medical objects suitable for statistical shape modeling," in *Proc. SPIE—Medical Imaging 1999: Image Processing*, vol. 3661, K. M. Hanson, Ed., Feb. 1999, pp. 620–631.
- [11] J. H. Downs, J. L. Lancaster, and P. T. Fox, "Surface-based spatial normalization using convex hulls," in *Brain Warping*. San Diego, CA: Academic, 1999, pp. 263–282.
- [12] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Comput. Vis. Graph. Image Processing*, vol. 46, pp. 1–21, 1989.
- [13] R. P. Woods, J. C. Mazziotta, and S. R. Cherry, "MRI–PET registration with automated algorithm," *J. Comput. Assist. Tomogr.*, vol. 17, no. 4, pp. 536–546, July/August 1993.
- [14] M. I. Miller, G. E. Christensen, Y. Amit, and U. Grenander, "Mathematical textbook of deformable neuroanatomies," *Proc. Nat. Acad. Sci.*, vol. 90, no. 24, pp. 11 944–11 948, Dec. 1993.
- [15] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Processing*, vol. 5, pp. 1435–1447, Oct. 1996.
- [16] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, "Automatic 3-D intersubject registration on MR volumetric data in standardized talairach space," *J. Comput. Assist. Tomogr.*, vol. 18, no. 2, pp. 192–205, March/April 1994.
- [17] Y. Amit, "A nonlinear variational problem for image matching," *SIAM J. Scientific Comput.*, vol. 15, no. 1, Jan. 1994.
- [18] J. G. Gee, L. Le Briquer, and C. Barillot, "Probabilistic matching of brain images," in *Information Processing in Medical Imaging*, Y. Bizais, C. Braillet, and R. Di Paola, Eds. Boston, MA: Kluwer Academic, June 1995, pp. 113–125.
- [19] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta, "Automated image registration: I. General methods and intrasubject, intramodality validation," *J. Comput. Assist. Tomogr.*, vol. 22, no. 1, pp. 139–152, 1998.
- [20] J. Ashburner and K. J. Friston, "Spatial normalization," in *Brain Warping*. San Diego, CA: Academic, 1999, pp. 27–44.
- [21] J. Ashburner, J. L. R. Andersson, and K. J. Friston, "High-dimensional image registration using symmetric priors," *NeuroImage*, vol. 9, pp. 619–628, 1999.
- [22] G. E. Christensen, S. C. Joshi, and M. I. Miller, "Volumetric transformation of brain anatomy," *IEEE Trans. Med. Imag.*, vol. 16, pp. 864–877, Dec. 1997.
- [23] J. Talairach and P. Tournoux, *Co-Planar Stereotactic Atlas of the Human Brain*. Stuttgart, Germany: Beorg Thieme Verlag, 1988.
- [24] L. A. Segel, *Mathematics Applied to Continuum Mechanics*. New York: Dover, 1987.
- [25] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, pp. 187–198, Apr. 1997.
- [26] W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Med. Image Anal.*, vol. 1, no. 1, pp. 35–51, 1996.
- [27] R. P. Woods, S. T. Grafton, J. D. Watson, N. L. Sicotte, and J. C. Mazziotta, "Automated image registration: II. Intersubject validation of linear and nonlinear models," *J. Comput. Assist. Tomogr.*, vol. 22, no. 1, pp. 153–165, 1998.
- [28] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "3-D brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.*, vol. 39, pp. 609–618, 1994.
- [29] M. I. Miller, A. Banerjee, G. E. Christensen, S. C. Joshi, N. Khaneja, U. Grenander, and L. Matejic, "Statistical methods in computational anatomy," *Stat. Meth. Med. Res.*, vol. 6, pp. 267–299, 1997.
- [30] G. E. Christensen, R. D. Rabbitt, M. I. Miller, S. C. Joshi, U. Grenander, T. A. Coogan, and D. C. Van Essen, "Topological properties of smooth anatomic maps," in *Information Processing in Medical Imaging*, Y. Bizais, C. Braillet, and R. Di Paola, Eds. Boston, MA: Kluwer Academic, June 1995, vol. 3, pp. 101–112.
- [31] H. J. Johnson, "Method for consistent linear-elastic medical image registration," master's thesis, Dept. Elect. Comput. Eng., Univ. Iowa, Iowa City, May 2000.
- [32] A. Toga, Ed., *Brain Warping*. San Diego, CA: Academic, 1999.