The Adaptive Bases Algorithm for Intensity-Based Nonrigid Image Registration

Gustavo K. Rohde, *Student Member, IEEE*, Akram Aldroubi, and Benoit M. Dawant*, *Senior Member, IEEE*

***Abstract*—Nonrigid registration of medical images is important for a number of applications such as the creation of population averages, atlas-based segmentation, or geometric correction of functional magnetic resonance imaging (fMRI) images to name a few. In recent years, a number of methods have been proposed to solve this problem, one class of which involves maximizing a mutual information (MI)-based objective function over a regular grid of splines. This approach has produced good results but its computational complexity is proportional to the compliance of the transformation required to register the smallest structures in the image. Here, we propose a method that permits the spatial adaptation of the transformation's compliance. This spatial adaptation allows us to reduce the number of degrees of freedom in the overall transformation, thus speeding up the process and improving its convergence properties. To develop this method, we introduce several novelties: 1) we rely on radially symmetric basis functions rather than B-splines traditionally used to model the deformation field; 2) we propose a metric to identify regions that are poorly registered and over which the transformation needs to be improved; 3) we partition the global registration problem into several smaller ones; and 4) we introduce a new constraint scheme that allows us to produce transformations that are topologically correct. We compare the approach we propose to more traditional ones and show that our new algorithm compares favorably to those in current use.**

***Index Terms*—Adaptive bases algorithm, mutual information, nonrigid image registration.**

I. Introduction

NONRIGID medical image registration, also known in the literature as spatial normalization or warping, is often an essential step in automated medical image analysis. A number of methods have been proposed over the years to solve this problem. For instance (a more complete review on the subject can be found in [1]), Collins [2] proposes a technique in which the overall transformation is obtained as a set of local affine ones. Bajcsy *et al.* [3], [4] use an elastic model approach. Algorithms based on viscous fluids are put forth by Christensen [5] and Bro-Nielsen [6]. Thirion [7] uses a method called "demons" that is similar to an optical flow approach for small displacements. Recently and following their success for rigid body registration problems [8], [9], mutual information (MI)-based methods have also been used for nonrigid registration problems. Meyer [10] relies on a technique based on thin-plate splines in which an optimizer is used to adjust the position of homolgous control points. Rueckert [11] and Studholme [12], [13] use a similar approach but with B-splines. Although implementations vary, these intensity-based techniques can be viewed in an optimization framework in which the registration problem consists of deforming a source image $B(\mathbf{x})$ to "best" match a target image $A(\mathbf{x})$ under a chosen similarity measure. Mathematically, this can be expressed as

$$\arg\max_{\mathbf{x}'} F\left(B(\mathbf{x}'), A(\mathbf{x}), \mathbf{x}'\right) \tag{1}$$

in which

$$\mathbf{x}' = \mathbf{x} + \mathbf{v}(\mathbf{x}) \tag{2}$$

and F is an intensity-based similarity measure (the cost function), $\mathbf{x}$ a coordinate vector in $\Re^d$, with d being the dimensionality of the images, and

$$\mathbf{v}(\mathbf{x}) = \{v_x(\mathbf{x}), v_y(\mathbf{x}), v_z(\mathbf{x}), \ldots, v_d(\mathbf{x})\}$$

a deformation field that warps image $B(\mathbf{x})$; thus, $\mathbf{v}(\mathbf{x})$ is what is computed in the registration problem. As mentioned earlier, a number of authors have proposed to use linear combinations of B-splines placed on a regular grid to model the deformation field $\mathbf{v}(\mathbf{x})$ [11], [12], [14], [15]. Because the splines are placed on a regular grid, the characteristics of the warping transformation [e.g., the number of degrees of freedom (DOFs) it possesses] does not vary spatially and we refer to this model as being spatially invariant. The major disadvantage of this approach is that the computational complexity of the method is proportional to the compliance (i.e., the number of basis functions) of the transformation needed to register the smaller structures in the image. Many structures of interest in medical images, especially in the brain, are in the order of millimeters. Deforming such structures requires placing basis functions at approximately every couple of millimeters which can require the optimization of a few hundred thousand basis function coefficients. For example, registering two typical three–dimensional (3-D) MR image volumes ($256 \times 256 \times 128$ pixels large) using a $64 \times 64 \times 32$ regular grid of splines generates a 393 216-dimensional search space. Finding an optimum in such a search space is not only time consuming but difficult because of the possibility of convergence to local optima. The work presented here proposes a new approach to nonrigid image registration that reduces the computational complexity and improves the convergence properties of methods proposed so far. The new approach is based on the idea that much can be gained if the compliance of the transformation is adapted locally. We have derived

Manuscript received December 18, 2001; revised June 29, 2003. This work was supported in part by the National Science Foundation (NSF) DMS under Grant 0103104 and Grant 0139740. Portions of this work were presented at the SPIE Medical Imaging Conferences of 2001 and 2002. *Asterisk indicates corresponding author.*

G. K. Rohde is with the STBB/LIMB/NICHD, National Institutes of Health, Bethesda, MD 02872 USA, and also with the Applied Mathematics and Scientific Computation Program, University of Maryland, College Park, MD 02792 USA. (e-mail: rohdeg@helix.nih.gov).

A. Aldroubi is with the Department of Mathematics, Vanderbilt University, Nashville TN 37235 USA.

*B. M. Dawant is with the Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville TN 37235 USA (e-mail: benoit.dawant@vanderbilt.edu).

Digital Object Identifier 10.1109/TMI.2003.819299

an approach that automatically identifies regions where the images are misregistered, and we focus on these regions only, thus avoiding useless computations on regions that are already correctly registered. Moreover, to prevent the optimization process from producing transformations that are physically incorrect, we have developed a new, precise, and fast way of enforcing that the Jacobian matrix of the deformation field remains uniformly invertible throughout the domain of the images. Together, these ideas were used to derive a new nonrigid registration algorithm that we call the Adaptive Bases Registration Algorithm.

The remainder of this paper is organized as follows. Section II describes the new method in detail, including the method for region of misregistration identification, our local deformation field optimization scheme, as well as our constraint scheme. Section III presents results we have obtained with this algorithm and it includes a comparison with a regular grid approach. Section IV summarizes the main contributions of this paper and suggests possible future work in this area. The appendix provides details on the constraint scheme we propose.

II. Method

A. Problem Statement

The goal in nonrigid registration is to generate a mapping relating any point in the domain of the source image $B(\mathbf{x})$ to a point in the domain of the target image $A(\mathbf{x})$. Let $D := [0,1]^d \subset \Re^d$, represent the domain of the images $B(\mathbf{x})$ and $A(\mathbf{x})$, where d is the dimensionality of the data sets, i.e., $d = 2$ for two-dimensional (2-D) images, and $d = 3$ for volumetric images. Registering images $B(\mathbf{x})$ and $A(\mathbf{x})$ is equivalent to finding the deformation $\mathbf{v}: D \rightarrow D$ such that $I + \mathbf{v}: D \rightarrow D$, where I is the identity transformation, is a one to one onto continuous map with continuous inverse (homeomorphism) and for which some cost function $F(A(\mathbf{x}), B(\mathbf{x}+\mathbf{v}), \mathbf{v})$ is optimized. These constraints on the transformation preserve the natural topology of the image, impeding the transformation from producing artifacts known as "folding" and "tearing" of the image. Here, we have used compactly supported radial basis functions to model the deformation $\mathbf{v}$, while the cost function optimized is the normalized MI (NMI) [16] between images $B(\mathbf{x}+\mathbf{v})$ and $A(\mathbf{x})$, although the approach we propose is not limited to this particular similarity measure. In our implementation, the NMI is estimated using the joint histogram of the source and target images while the value of image $B(\mathbf{x})$ at an arbitrary point is evaluated using trilinear interpolation. The value of the NMI is always evaluated over all the voxels belonging to the overlapping domain of images $B(\mathbf{x}+\mathbf{v})$ and $A(\mathbf{x})$.

B. Local Deformation Fields and Radial Basis Functions

As previously stated, rather than modeling the deformation field with a linear combination of cubic B-splines placed on a regular grid as is usually done [11]–[15], we build our deformation field incrementally, region by region, focusing on regions that are misregistered. The method by which we identify the relevant regions is detailed below but the total deformation field $\mathbf{v}(\mathbf{x})$ is modeled as a linear combination of a set of basis functions irregularly spaced over the image domain, i.e.,

$$\mathbf{v}(\mathbf{x}) = \sum_{i=1}^{N} \mathbf{c}_i \Phi(\mathbf{x} - \mathbf{x}_i) \tag{3}$$

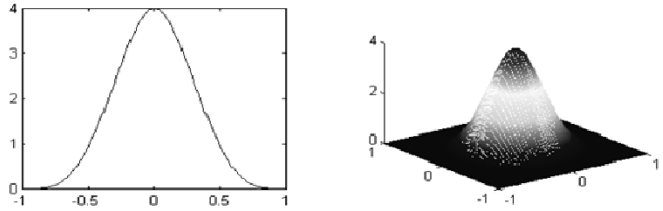

Fig. 1. Plot of the radial basis function whose equation is given in (6) with radius equal to one. Left: in one dimension; right: in two dimensions.

with coefficients $\mathbf{c}_1, \ldots, \mathbf{c}_N \in \Re^d$ and a function $\Phi: \Re^d \rightarrow \Re$ that is *positive definite* on $\Re^d$ in the following sense: for all sets $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_N\}$ of finitely many distinct points $\mathbf{x}_1, \ldots, \mathbf{x}_N$ in $\Re^d$, the matrix $M = (\phi(\mathbf{x}_k - \mathbf{x}_j))_{1 \le j,k \le N}$ is *positive definite* which guarantees the solvability of the system

$$\mathbf{y}_k = \sum_{i=1}^{N} \mathbf{c}_i \Phi(\mathbf{x}_k - \mathbf{x}_i), \quad 1 \le k \le N. \tag{4}$$

This property is important for registration problems for it guarantees that the model allows for the construction of any given deformation field solution prescribed by points placed at arbitrary locations. In the context of nonrigid registration, this means that any deformation field specified at an arbitrary set of points placed on an irregular grid can be modeled. For this to be true, the basis functions used to model the deformation field need to possess what is know as the universal interpolation property. While it is known that radial basis functions possess this property (see for instance [17]–[19]), it is not known whether or not B-splines possess it. This lead us to use one of Wu's compactly supported positive definite radial basis functions to model the deformation field

$$\Phi(\mathbf{x}) = \phi\left(\frac{||\mathbf{x}||_2}{s}\right), \quad \mathbf{x} \in \Re^d \tag{5}$$

with

$$\phi(r) = (1-r)_+^4 (3r^3 + 12r^2 + 16r + 4) \text{ for } r \ge 0 \tag{6}$$

where $(1-r)_+ = \max(1-r, 0)$, s is a predetermined scale for the basis function, and $||\cdot||_2$ is the usual Euclidean norm on $\Re^d$. Fig. 1 shows a plot of this function in one and two dimensions.

There are several advantages in using a compactly supported basis function such as (6) in registration problems. First, compact support means that for each value of x, the sum in (3) can be reduced to relatively few terms. This also means that under many circumstances optimization can be confined to a finite part of the domain D, improving the computational efficiency of the overall method. Moreover, (6) and, therefore, (5) have been shown to possess C^2 continuity. Smoothness properties are important in registration problems since the first and second derivatives of the deformation field are often used for the computation of the gradient, and sometimes Hessian, of the cost function with respect to the optimization parameters. These quantities are used in several optimization algorithms applicable to this type of registration problem, e.g., conjugate gradient descent, or Newton methods.

C. Multiscale and Multiresolution Approach

The algorithm proposed here approaches the final deformation field iteratively across scales and resolutions. Here, resolution means the spatial resolution of the image while the scale is related to the transformation itself. A standard image pyramid is

created to apply the algorithm at different resolutions. At each resolution, the scale of the transformation is adapted by modifying the region of support and the number of basis functions. The scale of the transformation is proportional to the bases' region of support (i.e., a large region of support leads to a transformation at a large scale). Typically, the algorithm is initialized on a low-resolution image with few basis functions having large support. As the algorithm progresses to finer resolutions and smaller scales, the region of support of the basis functions is reduced. Following this approach, the final deformation field is computed as

$$\mathbf{v}(\mathbf{x}) = \mathbf{v}_1(\mathbf{x}) + \cdots + \mathbf{v}_M(\mathbf{x}) \tag{7}$$

with M the total number of levels (in the remainder of this paper, a level refers to a particular combination of scale and resolution). It should be noted that the universal interpolation property discussed above holds only if all the basis functions have the same scale. Here we model the overall deformation field as a sum of deformation fields each computed at a different resolution and scale. Although the region of support for the basis functions changes from scale to scale and from resolution to resolution, at a particular scale and resolution the deformation field is computed with bases that have the same scale and region of support.

D. Regions of Misregistration Identification

One of the key features of the algorithm we propose is to adjust the transformation only where it needs to be adjusted. This requires identifying regions where the two images are not well registered at the current level and adjusting the deformation field over these regions. To achieve this, a local measure of misregistration needs to be developed. The approach we have used is as follows. When the algorithm moves from one level to another, we first place basis functions on a regular grid and we model the deformation field as

$$d(\mathbf{x}) = \mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x}) + \sum_{i=1}^{N} \mathbf{c}_i \Phi(\frac{\mathbf{x}}{s} - k_i), \text{with } k_i = \frac{\mathbf{x}_i}{s} \tag{8}$$

with $\mathbf{x}_i$ the position of the basis functions , s their scale, and $\sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x})$ the sum of the deformation fields obtained up to level $m-1$. This equation states that, when moving from one level to the other, an additional set of basis functions temporarily placed on a regular grid is used to model the deformation field; we call this grid Θ. Next, the gradient $\mathbf{G}$ of the cost function $NMI(A(\mathbf{x}), B(d(\mathbf{x}))$ with respect to the basis function's coefficients $\mathbf{c}_i$ is evaluated through finite differences. The value of $\mathbf{G}$ is then used to determine which regions in the images $A(\mathbf{x})$ and $B\left(\mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x})\right)$ are most likely to be misregistered at the current level. The idea behind using $\mathbf{G}$ to decide on regions of mismatch is as follows: if the magnitude of the gradient of the cost function with respect to the coefficient $\mathbf{c}_i$ is large, then the cost function is not at a minimum with respect to $\mathbf{c}_i$. If the cost function is not at a minimum at the location corresponding to $\mathbf{c}_i$ then it is likely that the region where the corresponding basis function is located is misregistered. Therefore, registration in this particular area could be improved at the current level. If, on the other hand, the magnitude of the gradient with respect to coefficient $\mathbf{c}_j$ is small, two situations are possible. Either the images are reasonably well registered over that region at the current level or the images could be significantly misregistered at that location but the cost function is at a local extremum. In either case, further gradient-based optimization in this region is unlikely to be fruitful and we assume that it can be neglected.

The algorithm we use for identifying regions over which to concentrate starts by evaluating $\mathbf{G}$ as described above. Once this is done, the individual components $\mathbf{G}_i$ of $\mathbf{G}$ are sorted in decreasing order according to their magnitude $||\mathbf{G}_i||$. The center of the regions of misregistration $\mathbf{x}_i$ are chosen as the location of the basis function for which $\mathbf{G}_i$ is above the selected threshold. Once a center is selected, the adjacent locations are eliminated from the list; this is done to force regions of interest (ROIs) to be disjoint (i.e., prevent overlap between these regions). This is discussed further in the next section.

E. Local Optimization

Once ROIs have been identified, the local deformation fields need to be computed. One possible approach would be to optimize all the coefficients $\mathbf{c}_i$ associated with the ROIs chosen in the previous step simultaneously. This would amount to optimizing the coefficients of basis functions scattered throughout the image domain which would be akin to the approach recently proposed by Schnabel *et al.* [20]. Here, however, we propose a solution that allows us to reduce the dimensionality of the optimization process by partitioning it. Given a location representing the center of a ROI $\mathbf{x}^j$ and the current resolution and scale s, we choose eight locations $\mathbf{x}^j_{[1,\ldots,8]}$ arranged in the form of a cube around $\mathbf{x}^j$ as centers for the basis functions that will be used for computing the deformation field associated with a particular region of the image. For 2-D registrations a square around the center location $\mathbf{x}^j$ is used, in 3-D we use a cube. This gives us the ability to build local deformations with eight DOFs in 2-D and 24 DOFs in 3-D around location $\mathbf{x}^j$. The support of the basis functions placed around location $\mathbf{x}^j$ is also s. Note that the value for s is obtained from the support of the basis functions used in the automatic ROI identification algorithm presented earlier. The local deformation field is thus adjusted at the current scale and resolution. A steepest gradient descent algorithm combined with the quadratic interpolation four-point bracketing update method of line minimization is then applied to the coefficients of the cube of basis functions under the following cost function:

$$F(A(\mathbf{x}), B(\mathbf{x}'), \mathbf{x}')$$

in which

$$\mathbf{x}' = \mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x}) + \mathbf{v}_m(\mathbf{x})$$

and

$$\mathbf{v}_m(\mathbf{x}) = \sum_{j=1}^{J} \sum_{i=1}^{8} \mathbf{c}_i^j \Phi(\mathbf{x} - \mathbf{x}_i^j). \tag{9}$$

In these equations, the values $\mathbf{x}^j_{[1,\ldots,8]}$ represent the aforementioned cube of center locations, and J is the number of regions of mismatch identified at this level. Therefore, for each local field, the set of coefficients $\mathbf{c}^{[1,\ldots,J]}_{[1,\ldots,8]}$ must be optimized. Because the ROIs are chosen in such a way that they do not overlap, the

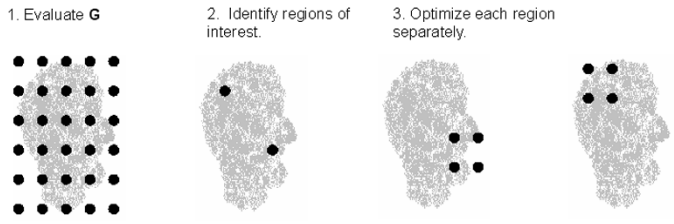

Fig. 2. Graphical illustration of the sequence of steps through which the algorithm goes at each level.

optimization of the set of coefficients $\mathbf{c}_{[1...8]}^{[j]}$ can be done independently of all other sets of coefficients $\mathbf{c}_{[1...8]}^{[l]}$, with $j \neq l$. Thus, we optimize (9) one region at a time. Fig. 2 illustrates the operation of the algorithm.

Conducting the optimization one region at a time reduces one large optimization problem to a series of small ones with at most 24 parameters each. This strategy has one main advantage: by decoupling the optimization of coefficients associated with disjoint regions into separate optimization problems, we eliminate the possibility that undesirable effects such as noise and local minima associated with optimization over one region do interfere with optimization over other regions, thus leading to more accurate results as will be illustrated further. Furthermore, this strategy permits the parallelization of the algorithm.

F. Optimization Constraint Scheme

Previous attempts at constraining spline-based deformation field models to produce consistent topological deformations can be found in Rueckert *et al.* [11] or Studholme *et al.* [13] where the basic idea is to optimize the similarity measure while regularizing the deformation field by minimizing its second derivative. While this technique can reduce folding artifacts, it does not guarantee the positive definiteness of the Jacobian of the transformation. Here we not only make explicit the relationship between a smoothness constraint and the Jacobian of the deformation field, but we derive precise bounds for the basis functions coefficients $\boldsymbol{c_i}$ in (3) that guarantee this positiveness. Consider the following deformation field:

$$T(\mathbf{x}) = \mathbf{x} + v_1(\mathbf{x}) + v_2(\mathbf{x}) + \cdots + \mathbf{v}_N(\mathbf{x}) \tag{10}$$

from $\Re^3 \rightarrow \Re^3$, where N is the number of levels utilized during registration. Let

$$J(T) = \mathbf{I} + \alpha_1(\mathbf{x}) + \alpha_2(\mathbf{x}) + \cdots + \alpha_N(\mathbf{x})$$

be the Jacobian of the transformation $T(\mathbf{x})$, where I is the identity matrix and $\alpha_i(\mathbf{x})$ is the Jacobian matrix of the displacement field $\mathbf{v}_i(\mathbf{x})$. In the appendix, we prove that if the constraint

$$||\alpha_{n+1}||_\infty < \frac{1}{3}\left(1 - 3\left\|\sum_{i=1}^{n}\alpha_i\right\|_\infty\right) \tag{11}$$

is satisfied, then the Jacobian remains positive. In other words, for any given deformation field (10), a displacement field $\mathbf{v}_{N+1}(\mathbf{x})$ can be added without violating the topology constraint as long as relationship (11) is satisfied for $n = 1, \ldots, N$. Relationship (11) can be used to design a number of possible constraint schemes for registration procedures that use splines, radial basis functions, or other types of bases. One possible scheme is to compute $\|\sum_{i=1}^{n}\alpha_i\|_\infty$ after each level. Then (11) can be used to compute bounds on the basis function coefficients $\mathbf{c}_i$ composing the $(n+1)$th displacement field. This, however, would be computationally expensive since in our registration algorithm the evaluation of $\|\sum_{i=1}^{n}\alpha_i(\mathbf{x})\|_\infty$ for all coordinates $\mathbf{x}$ would have to be done explicitly through finite differences. Another possible constraint scheme is to use (11) to compute bounds for the basis function coefficients by assuming worst-case estimates for $\|\sum_{i=1}^{n}\alpha_i(\mathbf{x})\|_\infty$. For example, if the number N of levels is fixed *a priori*, then (11) will be satisfied if $1 - 3\|\sum_{i=1}^{n}\alpha_i\|_\infty$ is positive for any value of $n = 1, \ldots, N$. Using the triangle inequality $\|\sum_{i=1}^{n}\alpha_i\|_\infty \leq \sum_{i=1}^{n}\|\alpha_i\|_\infty$, we derive the bound

$$\max_{i=1,\ldots,N}||\alpha_i||_\infty \leq \lambda < \frac{1}{3N}. \tag{12}$$

In practice, we use (12) to obtain the bounds for the basis function's coefficients at each level of the algorithm. At every iteration, the value of the coefficients is computed and checked against their upper limit. If this limit is exceeded the maximum allowable value is substituted which, in effect, forces the displacements to be compatible with our topological constraints.

G. Deformation Field Derivative Estimation

Since $\alpha_i = \text{Jacobian}(\mathbf{v}_i)$ the use of the constraint scheme described above requires a fast and accurate method for estimating the partial derivatives of the deformation field. Because at each level the displacement field $\mathbf{v}_i$ is a linear combination of basis functions, these partial derivatives can be directly computed from the coefficients of the linear combination. Thus, a constraint on the coefficients will translate into a constraint on the lefthand side of (12). For example, for a one-dimensional (1-D) deformation field

$$v(x) = \sum_{k=1}^{M} c_k \phi(x - x_k)$$

we have $\max_{x\in\Re}|v'(x)| \leq \|\phi'\|_\infty \sum_{k=1}^{M}|c_k|$. Therefore, to impose $\max_{x\in\Re}|v'(x)| \leq \lambda$ all we need is to impose

$$\sum_{k=1}^{M}|c_k| \leq \frac{\lambda}{\|\phi'\|_\infty}.$$

In practice, this can be achieved simply by verifying that the constraint is not violated after each step of the $\mathbf{c}_i$ optimization process. When the constraint is violated, the optimization process is interrupted. Note that the aforementioned estimate is crude (i.e., it is overconstraining) and sharper bounds can be obtained by taking into account the spatial arrangement of the basis functions (e.g., their regions of support and amount of overlap). For instance, if it is known that the basis functions have the same radius and are placed in such a way that the support of one does not extend beyond the center of the next one, one can show that the maximum value of the derivative of the deformation field obeys the following inequality:

$$\max_{x\in[x_i,x_{i+1}]}|v'(x)| \leq \frac{\|\Phi'\|_\infty |c_{i+1} - c_i|}{r} \tag{13}$$

in which v is a 1-D deformation field modeled by a linear combination of basis functions located at x_i and x_{i+1} with coefficients c_i and c_{i+1} and radius $r = r_i = r_{i+1}$.

```
Initialize A(x), B(x'), x' at the lowest resolution and scale.
    For l=1... Number of resolutions.
        For J=1.... Number of scales at current resolution.
            Create regular grid Θ at current resolution and scale and
              compute region of support for basis functions.
            Identify regions of misregistration.
            Optimize each region independently from each other
        End for
        Upsample A(x), B(x'), x'
    End for
Ouput  B(x') and x'.
```

Fig. 3. The adaptive bases registration algorithm.

H. Summary of the Adaptive Bases Algorithm

Fig. 3 summarizes the algorithm we propose. At first, input images $A(\mathbf{x})$ and $B(\mathbf{x})$ are downsampled to the lowest user-specified resolution and a bounding box is computed from the union of the foreground of both images. Initially, v is set to zero. The parameters needed by the algorithm are the number of resolutions and the scales at which the transformation needs to be computed at each resolution (the scales are specified by the number of basis functions to be used when creating the regular grid Θ; the lower the number of basis functions, the larger the scale). At each resolution and scale, the region of support for the basis functions is calculated as a constant times the distance between two adjacent grid points in Θ (in practice the constant is between 1.5 and 2).

III. Results

A. Algorithms Comparisons: Simulated Data

As explained above, the adaptive basis registration algorithm introduces three novel concepts to intensity-based image registration. The first is an adaptation step, where at each level only the regions that are mismatched at the current scale are optimized. Secondly, we propose optimizing disjoint regions one at a time, instead of jointly. Lastly, we also introduce a novel constraint scheme. The purpose of this section is to demonstrate the effect of the adaptation step separately from the effects of disjoint optimization. To that end, four 2-D registration algorithms were implemented. The first treats the problem with the conventional stationary approach described in [11], [12], [14], and [15]. The second method referred to as the adaptive method (AM) uses the region of mismatch identification algorithm presented earlier but the coefficients of all the basis functions are optimized concurrently whether or not these basis functions define overlapping regions. This approach is similar to the methods investigated in [20] and [21]. The third method implements our approach in which the algorithm operates on disjoint regions. As described before, in this implementation regions where the images are misregistered are first identified then four basis functions are placed around each identified location while making sure that regions over which the algorithm operates do not overlap. Here we optimize over all the basis functions at once and we refer to this implementation as the adaptive disjoint parallel method (ADPM). The last algorithm is similar to the third one except that we optimize over each disjoint region in sequence. We refer to this algorithm as the adaptive disjoint serial method (ADSM). The source and target images used here are shown in Fig. 4. For this example the target image

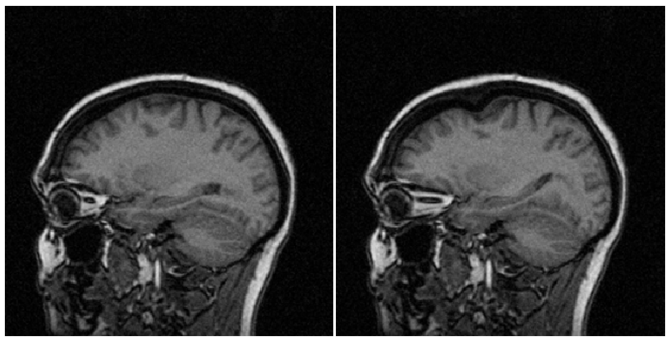

Fig. 4. From left to right, the source and target images used in the 2-D simulated experiments.

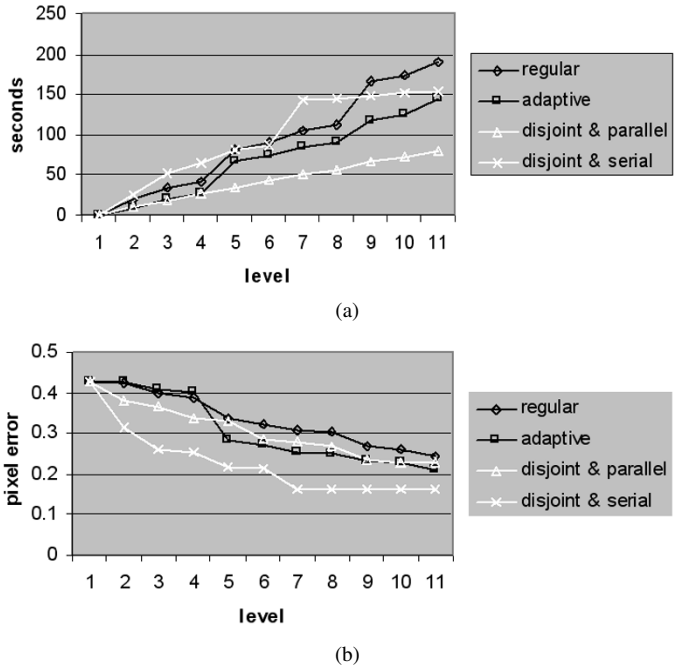

Fig. 5. Timing (a) and average (b) pixel error for various optimization schemes.

was generated by applying a known transformation to the source image. The deformation modified the image in three disjoint locations (around the eye, on the top of the cortex, and around the hippocampus). It was built using the radial basis function $\phi(r) = (1-r)^4 + (4r+1)$, for $r \geq 0$; it is, thus, different from the basis used to compute the deformation field that registers the two images. The radius of the basis function used in generating the simulated warping function was 20 pixels. Gaussian distributed noise was added to both images. When using the ADSM or ADPM algorithms to register the source image to the generated target image, we do not control the position of the basis functions. These are placed automatically over regions that have been identified as misregistered.

The four algorithms were set to span ten levels, starting with a grid of basis functions of 10×10 up to a grid of 20×20. At the end of each level, the system time as well as the error between the current displacement field solution and the true displacement field was computed. The plots of system time and deformation field error for each algorithm are shown in Fig. 5. Several conclusions can be drawn from these experiments. First, the traditional method based on a regular grid is the slowest. As could be

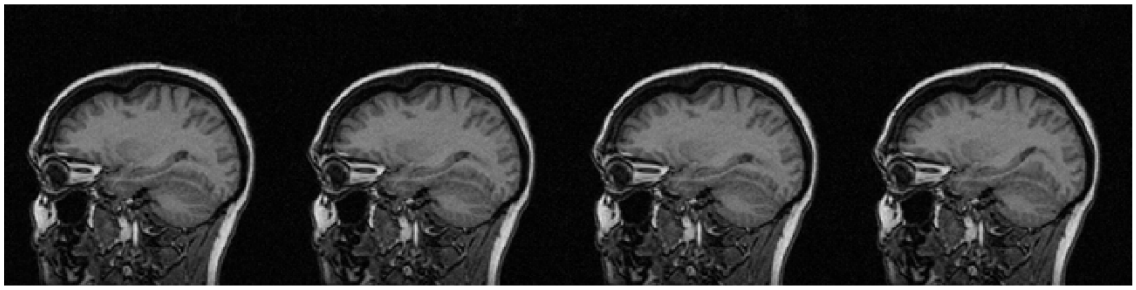

Fig. 6. From left to right, registration results obtained with regular grid optimization, the AM, the ADPM, and the ADSM.

expected because the dimension of the search space is reduced, the AM method speeds up the process. It also leads to an average pixel error inferior to the one obtained with the traditional approach. The disjoint region optimization approach we propose when performed simultaneously over all regions of mismatch is the fastest but lead to accuracy numbers that are similar to the numbers obtained with the AM. This suggests that the algorithm does converge rapidly toward a local minimum. The best accuracy numbers we have obtained are with the ADSM although it is slower than the ADPM. Note that the results presented in this figure are only indicative of what can be achieved. Relative timing and accuracy numbers between methods depend on both the images being registered and parameters being used (results presented in the next section on a series of 3-D volumes do show more than a threefold gain in speed between the traditional approach and the ADSM). In our experience, however, the traditional method is the least accurate and the slowest while the ADSM method is, in general, the most accurate. This suggests that optimizing the cost functions on small regions one at a time rather than together reduces the overall process' sensitivity to local minima. As discussed before it is also worth noting that the ADSM lends itself to parallelization which could improve performance substantially. Fig. 6 illustrates visually differences between the four algorithms. The results obtained with the traditional approach, the AM, and the ADPM appear very similar. The ADSM is the only one that has been able to deform correctly the ventricular region at the end of the hippocampus, indicating that this approach permits focusing on small regions over which the images are misregistered.

B. Algorithms Comparisons: Real Data

In this section we evaluate the performance of the Adaptive Bases Registration Algorithm we propose on a series of 3-D MR data sets. The program was written in the C++ programming language, and all experiments were run on an IBM compatible PC, with an Intel Pentium 4 processor (1.7 GHz) running Windows ME.

For the purpose of comparison, we have also implemented a method which is similar in nature to the methods used in [11], [12], [14], [15], except for the additional constraint scheme used in our method. Both programs were built by compiling identical codes, with the exception that the identification of misregistered regions was turned off for the second program leading to a gradient descent optimization of the cost function (5) using a spatially invariant model for the deformation field. The data used in this set of experiments were 11 MRI volumes. The MR brain images used here were obtained with high-resolution 3-D SPGR pulse sequences (FOV 24×24 cm, 256×256, 1.3-mm thickness, 0-mm gap, $\mathrm{TE} = 1.9$ ms, flip angle $= 20$ deg, $\mathrm{TI} = 450$ ms, $\mathrm{TR} = 11.9$ ms, 128 slices). We have used a set of 11 volumes and the task was to register ten volumes to the 11th one chosen as the target image for all registrations, by using both registration algorithms. The parameters used with both algorithms were identical. Six levels were used, with the radius of the basis functions varying from about 200 mm (with Θ of size $2 \times 2 \times 2$) to about 6 mm (with Θ of size $17 \times 17 \times 15$). For both algorithms, the optimization of any set of basis functions was halted when improvements of at least 0.0005 in the cost function could not be detected. In this set of experiments, we have used the NMI [16] computed from the joint histogram of the overlapping regions between the two image volumes as the cost function. The NMI is computed from the joint histogram of the entire overlapping area between the two image volumes even though we only alter small regions of the deformation field. We have used 32 bins to generate the joint histograms and a two-point finite difference formula to compute the gradient of the cost function. The number of bins was chosen experimentally as a good compromise between speed and accuracy on a few volumes and then kept constant for the entire study. Although we could have changed the number of bins in the histogram when moving from one resolution to the other we have chosen no to do it to avoid having to specify another set of parameters. Prior to applying our algorithm, each volume was registered to the target volume using a nine DOFs transformation computed with an independent implementation of an MI-based method similar to the one proposed by Maes [8].

The time required for each registration for both algorithms as well as the final value of the cost function are given in Tables I and II. These tables show that, on average, the adaptive bases algorithm is about 3.5 times faster than the more traditional approach for this set of experiments. For every registration, the

TABLE I
FINAL VALUE OF COST FUNCTION (NMI) AND TOTAL TIME TAKEN TO ACHIEVE REGISTRATION (SECONDS) BY THE ADAPTIVE ALGORITHM

	1	2	3	4	5	6	7	8	9	10
Time	9773	9781	10058	9869	10155	9928	9862	9802	9950	10165
NMI	1.25	1.23	1.242	1.24	1.242	1.23	1.25	1.25	1.23	1.228

TABLE II
FINAL VALUE OF COST FUNCTION (NMI) AND TOTAL TIME TAKEN TO ACHIEVE REGISTRATION (SECONDS) BY THE TRADITIONAL APPROACH

	1	2	3	4	5	6	7	8	9	10
Time	24405	28973	29679	41331	36413	46698	24207	29441	35473	48012
NMI	1.224	1.211	1.216	1.213	1.222	1.213	1.217	1.232	1.219	1.213

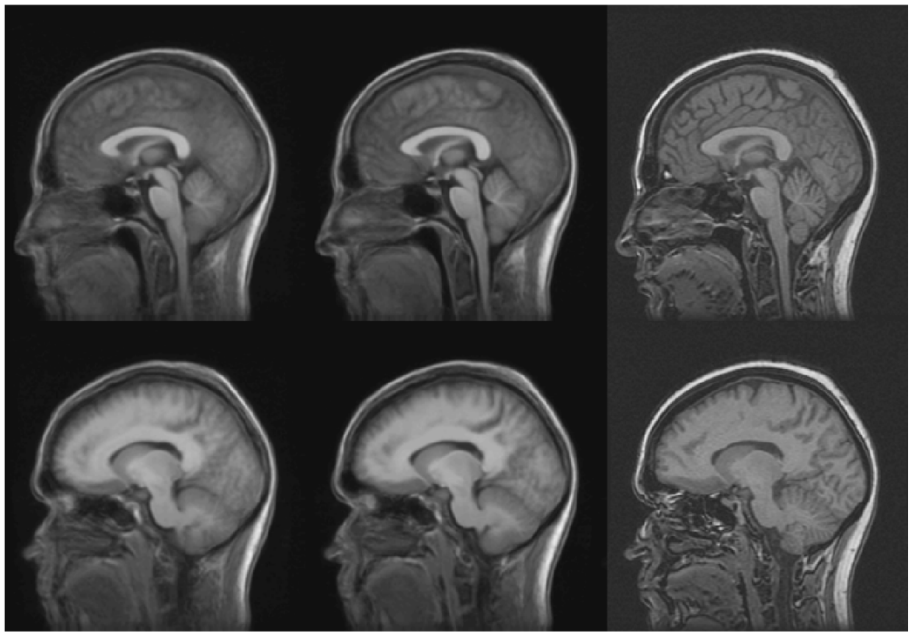

Fig. 7. Representative slices of atlases constructed through nonrigid registration with both adaptive and nonadaptive approaches. From left to right: average after nonrigid registration using nonadaptive approach, average constructed after registration using adaptive approach, and target image.

adaptive bases approach is at least twice as fast as its nonadaptive counterpart. Moreover, Tables I and II also show that for every registration, the adaptive bases registration algorithm is at least as accurate, or more, as measured by the final value of the cost function, than the traditional approach. The average final value for the cost function for the registrations which used the adaptive approach was 1.239, while it is 1.218 for the nonadaptive approach.

To further assess the quality of the registrations, for one cannot always use the cost function itself to assess the quality of the results, we also show the average brains computed through registration using both approaches. These are shown in Fig. 7. In this figure, the leftmost image is the average computed after nonrigid registration with the traditional approach. The middle image is the average computed after nonrigid registration with the approach we propose. The target image is also shown on the right. The average computed after nonrigid registration with the new approach is visibly sharper than the average computed after registration with the traditional approach, indicating that, overall, the new approach succeeded in matching the target image better than the traditional one.

Finally, note that the smallest radius of the basis functions used in the registrations described above was about 6 mm. Further improvements in accuracy can be achieved by adding more levels, which use radial basis functions with even smaller radii. Fig. 8 shows an example of a registration achieved using seven levels of the adaptive bases registration algorithm, with radial basis function ranging in radius from about 200 mm to 3 mm. The images clearly show that the adaptive bases registration algorithm is capable of producing high-quality matches, even for the smallest visible structures.

C. Atlas-Based Segmentation Results

The accuracy of the adaptive bases algorithm for nonrigid registration problems was assessed quantitatively by means of an atlas-based segmentation task. ROIs (whole brain, eyes and optic nerves, and spine) were manually delineated in the atlas used in Section III-B and binary masks were created. The regions were chosen because these present a range of difficulty with the whole brain contours being the easiest and the optic nerves the most difficult. The deformation fields between the atlas and the ten other volumes were computed and used to project the masks from the atlas onto each of the remaining volumes. Contours were manually drawn on a few slices chosen at random in each volume (four slices/volume/structure). Manual contours and contours obtained automatically were then compared using an accepted similarity index defined as two times the number of pixels in the intersection of the contours divided by the sum of the number of pixels within each contour [22]. This index varies between zero (complete disagreement) and one (complete agreement) and is sensitive to both displacement and differences in size and shape. Table III lists mean values for the similarity index for each structure. It is customarily accepted that a value of the similarity index above 0.80 indicates a very

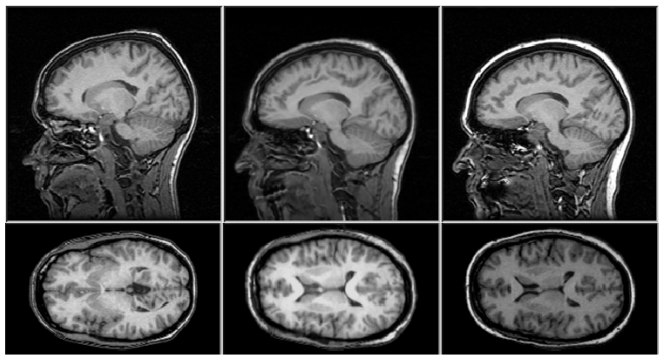

Fig. 8. Representative slices of an example registration using the adaptive grid registration algorithm. From left to right: The source image, the source image resampled to match the target image, and the target image.

TABLE III
AVERAGE VALUES OF S FOR A NUMBER OF STRUCTURES

Volume	1	2	3	4	5	6	7	8	9	10
Whole Head	0.97	0.96	0.96	0.96	0.97	0.97	0.97	0.97	0.96	0.96
Spine	0.93	0.92	0.93	0.92	0.93	0.92	0.91	0.95	0.92	0.92
Left Eye	0.91	0.9	0.9	0.88	0.9	0.86	0.86	0.92	0.91	0.82
Right Eye	0.89	0.86	0.81	0.88	0.9	0.82	0.86	0.89	0.92	0.82

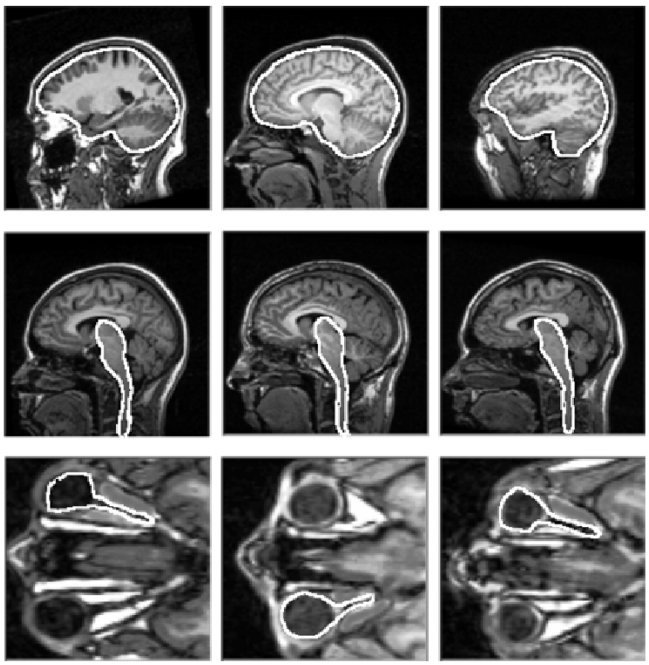

Fig. 9. Illustrative contours obtained automatically with our algorithms by deforming templates in the atlas.

good agreement between contours. Our results are well above this value except for the eyes. The major source of errors is the optic nerve which is a thin and elongated structure. Because of the size of this structure, an error of a few pixels reduces the value of the similarity index considerably. Fig. 9 shows a few representative examples of contours that have been obtained automatically.

D. Constraint Scheme Demonstration

We now demonstrate the effect of the constraint scheme presented earlier on a 3-D registration problem. Fig. 10 shows representative slices of the source image $B(\mathbf{x})$ and the target image $A(\mathbf{x})$ used in this demonstration. We have used our algorithm together with the constraint scheme described above to register image $B(\mathbf{x})$ to image $A(\mathbf{x})$. We have repeated the experiment four times, each with a different value for λ. The total number of levels used was 6, while the last level used a grid Θ with $17 \times 17 \times 15$ center locations. Fig. 11 shows the registration results $TB(\mathbf{x})$ for each value of λ in (12). The results shown in Fig. 11 were generated with λ equal to (from left to right): 0.3, 0.2, 0.1, and 0.05. Note that the results look identical in nearly all of the regions of the slices shown in Fig. 11. The same has been visually confirmed for all the other slices in these images. Differences can be noticed mainly around the mouth. The bottom row of Fig. 11 displays the mouth region for each output image.

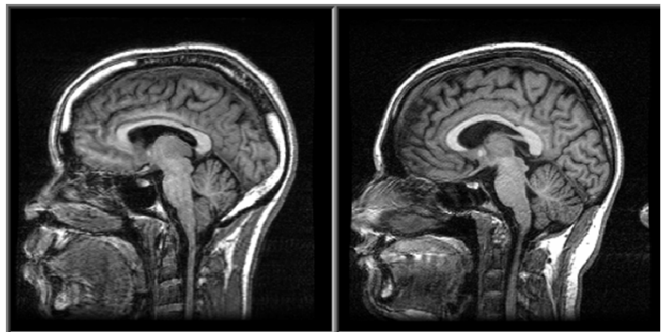

Fig. 10. Images used in constraint scheme demonstration. Left: source image $B(\mathbf{x})$. Right: target image $A(\mathbf{x})$.

As can be seen by focusing on the bottom left corner of these images, the results generated with $\lambda = 0.3$ does not satisfy the topological constraint for it produces folding. The folding effect becomes decreasingly noticeable as the threshold used is lowered. Using (12) our theory predicts a value of $\lambda = 0.0556$ or less for $N = 6$. From Fig. 11 one can verify that the folding effect completely disappears in the result image generated with $\lambda = 0.05$, confirming what has been theoretically predicted.

IV. DISCUSSION

Over the years, nonrigid registration using NMI and B-splines placed on a regular grid has been shown to be both accurate and robust. In this paper, we present an approach which, while inspired by existing work, presents several novel elements. First, we do not rely on basis functions placed on a regular grid. This allows us to adjust the deformation field only where it needs to be adjusted at the current scale and resolution. By doing so we reduce the dimension of the search space, thus speeding up the process. We simplify the process further by computing local deformations on disjoint regions. By doing so we transform one large optimization process into a series of smaller ones with at most 24 DOFs (in 3-D). While it is possible to show theoretically that the global optimum reached when optimizing over the entire region is the same as the one reached when optimizing over smaller, separate, regions when the cost function is separable (as would be the case if we were using the L2 norm), we have not been able to prove the same with a nonseparable cost function such as MI without making stationarity assumptions that may not always be verified in practice [23]. There is, therefore, no theoretical guarantee that by optimizing on separate regions we will be able to reach the global optimum. In practice, however, the various local minima in the MI cost function limit the ability of numerical algorithms commonly used to solve the nonrigid registration problem to reach this optimum. Our results show that by optimizing on small separate regions in series we reduce the effect of local minima and reach solutions that are better than solutions arrived at when trying to solve the problem globally.

The method we have used to identify ROIs is not the only one possible but in our experience it is robust and reliable. In [24], we did try to use the local correlation as an index of regional misregistration but we found the current strategy to be superior. The method we use to select misregistered regions involves a threshold whose value is set experimentally. Chosen too high, some regions at a particular resolution and scale could be

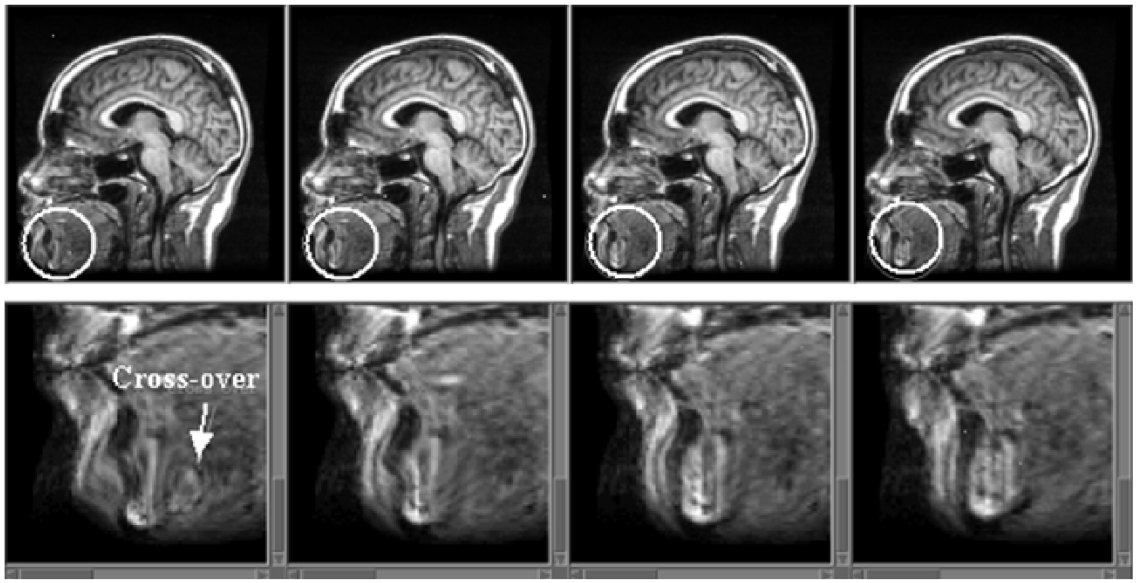

Fig. 11. Each row from left to right: registration results obtained with $\lambda = 0.3$, $\lambda = 0.2$, $\lambda = 0.1$, and $= 0.05$. Bottom row are zoomed in versions of the images on top row.

skipped and may not be recoverable at the next level. Chosen too low, many regions will be optimized at the current scale when they could be better registered at the next one. At the time of writing we have not studied the effect of this threshold on the registration accuracy. We also did not investigate the possibility of changing the number of bins in the joint histogram when moving from one resolution to the other to take into account the change in pixel numbers on the images when their spatial resolution is changed. It is possible that convergence properties of the algorithm could be improved by doing so.

Since we published earlier versions of this paper [24], [25] others have also proposed techniques by which the compliance of the transformation could be adapted within a B-spline framework [20], [21]. The approach these authors use is to fix the coefficients of the B-splines whose region of support have been labeled as passive. Criteria used to identify passive regions include local statistical measures such as the joint entropy or the gradient measure we introduced in [25]. The constraint scheme we propose to guarantee the topological correctness of the overall transformation is another novelty we introduce. Constraints proposed in the literature are somewhat ad hoc, attempting to limit the folding problem by adding a smoothness constraint to the cost function. The method we propose here explicitly enforces the correctness of the transformation and our experimental results have verified our theoretical predictions. The results we have shown demonstrate that when using basis functions with the same supports the method we propose is both more accurate and faster than a method relying on a regular grid. Although the accuracy of nonrigid registration algorithms is difficult to assess, we have shown that the value of the final similarity measure is larger with our method than with a more traditional approach. We have also shown qualitatively that the accuracy of the registration is better with our approach by comparing the sharpness of average volumes. Quantitatively, we have shown that our algorithm, when used for atlas-based segmentation tasks, produces results that are comparable to those obtained manually. As is the case for all intensity-based algorithms, the main limitation of our approach is the lack of *a priori* and anatomical information that can lead to erroneous deformation. This is particularly true in the cortical area where nothing prevents the algorithm to match sulci that do not correspond to each other. Addressing this problem will require adding constraints in a way similar to the approach proposed by Hellier *et al.* [26].

APPENDIX

Let $J(T) = \mathbf{I} + \alpha_1(\mathbf{x}) + \alpha_2(\mathbf{x}) + \cdots + \alpha_N(\mathbf{x})$ be the Jacobian of $T(\mathbf{x})$ in (10), $\mathbf{I}$ is the identity matrix, and $\alpha_i(\mathbf{x})$ the Jacobian matrix of the displacement field $\mathbf{v}_i(\mathbf{x})$. If $\|\alpha_1\|_{op_2} = \sup_{\|y\|_2=1} \|\alpha_1 y\|_2 < 1$ (here $\|y\|_2 = \sqrt{\sum_{i=1}^{3} y_i^2}$ is the usual L^2-norm), then it is well known that $(\mathbf{I}+\alpha_1)^{-1}$ exists and is given by the Neuman series

$$(\mathbf{I}+\alpha_1)^{-1} = \sum_{m=0}^{\infty} (-1)^m \alpha_1^m = \mathbf{I} - \alpha_1 + \alpha_1^2 - \cdots + \alpha_1^p - \cdots. \quad (14)$$

Moreover, we have

$$\frac{I}{I + \|\alpha_1\|_{op_2}} \leq \left\|(I+\alpha_1)^{-1}\right\|_{op_2} \leq \frac{I}{I - \|\alpha_1\|_{op_2}} \quad (15)$$

Since $(\mathbf{I}+\alpha_1)^{-1}$ exists for all $\|\alpha_1\|_{op_2} < 1$, we conclude that $\det(\mathbf{I}+\alpha_1) \neq 0$ for $\|\alpha_1\|_{op_2} < 1$. Moreover, since 1) the mapping $\det(\mathbf{I}+\alpha_1)$ that takes the vector space M_3 of (3×3) matrices into $\Re$ is continuous; 2) $\det \mathbf{I} > 0$; and 3) $\|\alpha_1\|_{op_2} < 1$ is a connected set, it follows that $\det(\mathbf{I}+\alpha_1) > 0$ (if $\det(\mathbf{I}+\alpha_1)$ switches sign, the continuity would imply that $\det(\mathbf{I}+\alpha_1)$ must become null for some matrix α_1 contradicting the fact that $\mathbf{I}+\alpha_1$ has an inverse).

Let $A_1 = \mathbf{I} + \alpha_1$, and $A_2 = A_1 + \alpha_2$. For $\|\alpha_1\|_{op_2} < 1$, $A_1^{-1} = (\mathbf{I}+\alpha_1)^{-1}$ exists, and we write A_2 as $A_2 = A_1 + \alpha_2 = A_1(\mathbf{I} + A_1^{-1}\alpha_2)$. Thus, A_2 is invertible if and only if $\mathbf{I} + A_1^{-1}\alpha_2$

is invertible and $A_2^{-1} = (\mathbf{I} + A_1^{-1}\alpha_2)^{-1}A_1^{-1}$. Using (14), it follows that A_2^{-1} exists if $||A_1^{-1}\alpha_2||_{op_2} < 1$ and as before

$$\det A_2 = \det A_1(\mathbf{I} + A_1^{-1}\alpha_2) = \det A_1 \det(\mathbf{I} + A_1^{-1}\alpha_2)$$

remains positive since it is the product of positive values. But using (15) together with a well known norm inequality we get

$$\left\|A_1^{-1}\alpha_2\right\|_{op_2} \leq \left\|A_1^{-1}\right\|_{op_2} ||\alpha_2||_{op_2} \leq \frac{||\alpha_2||_{op_2}}{1 - ||\alpha_1||_{op_2}}. \quad (16)$$

Thus assuming $||\alpha_1||_{op_2} < 1$, if $||\alpha_2||_{op_2}(1 - ||\alpha_1||_{op_2})^{-1} < 1$, then $\det A_2$ remains positive. Equivalently, $\det A_2$ does not change sign if $||\alpha_2||_{op_2} < 1 - ||\alpha_1||_{op_2}$. This last condition gives a bound for $||\alpha_2||_{op_2}$ in terms of $||\alpha_1||_{op_2}$. By induction, we obtain that

$$||\alpha_{n+1}||_{op_2} < 1 - \left\|\sum_{i=1}^{n}\alpha_i\right\|_{op_2} \quad \text{for } n = 1, \ldots, N-1 \quad (17)$$

is sufficient for $A_N = I + \alpha_1 + \alpha_2 + \cdots + \alpha_N$ to be invertible, and for $\det A_N$ to remain positive. Therefore, to satisfy the topological constraint on the displacement field $\mathbf{v}_i(\mathbf{x})$, it is sufficient to enforce (17) for all possible coordinates $\mathbf{x}$. This can be achieved by restricting the range of the coefficients of the radial bases functions used in modeling the displacement field $\mathbf{v}_i(\mathbf{x})$. However, computing the operator norm $||\alpha_i||_{op_2}$ of matrix α_i during optimization is cumbersome for it involves the computation of eigenvalues. To avoid computing operator norms $||\alpha_i||_{op_2}$, we can use the infinity norm $||\alpha||_\infty := \max_{m,n}|\alpha_{m.n}|$, which can be computed cheaply from the coefficients $\mathbf{c}_i$ in (3). Using the well known matrix relation $||\alpha||_{op_2} \leq ||\alpha||_2$ together with a simple calculation that yields $||\alpha||_2^2 \leq 3||\alpha||_\infty$, we obtain that if (11) is satisfied for $n = 1, \ldots, N-1$ then (17) is satisfied. Thus, for any given deformation field

$$T(\mathbf{x}) = I + \sum_{i=1}^{N}\mathbf{v}_i(\mathbf{x})$$

a displacement field $\mathbf{v}_{N+1}(\mathbf{x})$ can be added without violating the topology constraint.

ACKNOWLEDGMENT

The authors would like to thank R. Li and P.-F. D'Haese for their help in generating some of the results presented in this paper.

REFERENCES

[1] J. B. Antoine Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, no. 1, pp. 1–36, 1998.

[2] D. L. Collins, P. Neelin, T. M. Peters, and A. C. Evans, "Automatic 3D intersubject registration of MR volumetric data in standardized talairach space," *J. Comput. Assist. Tomogr.*, vol. 18, no. 2, pp. 192–205, 1994.

[3] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Comput. Vis., Graphics, Image Processing*, vol. 46, pp. 1–21, 1989.

[4] J. C. Gee, M. Reivich, and R. Bajcsy, "Elastically deforming 3D atlas to match anatomical brain images," *J. Comput. Assist. Tomogr.*, vol. 17, no. 2, pp. 225–236, 1993.

[5] G. E. Christensen, M. I. Miller, and M. Vannier, "3D brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.*, vol. 39, pp. 609–618, 1994.

[6] M. Bro-Nielsen and C. Gramkow, "Fast fluid registration of medical images," in *Proc. Visualization in Biomedical Computing*, K. H. Hohne and R. Kikinis, Eds., Hamburg, Germany, 1995, pp. 267–276.

[7] J.-P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.*, vol. 2, no. 3, pp. 243–260, 1998.

[8] 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. 187–98, Apr. 1997.

[9] 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.

[10] C. R. Meyer, J. L. Boes, B. Kim, P. Bland, K. R. Zasadny, P. V. Kison, K. Koral, K. A. Frey, and R. L. Wahl, "Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin plate spline warped geometric deformations," *Med. Image Anal.*, vol. 3, pp. 195–206, 1997.

[11] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, pp. 712–721, Aug. 1999.

[12] C. Studholme, R. T. Constable, and J. Duncan, "Accurate alignment of functional EPI data toanatomical MRI using a physics-based distortion model," *IEEE Trans. Med. Imag.*, vol. 19, pp. 1115–1127, Nov. 2000.

[13] C. Studholme, V. A. Cardenas, and M. W. Weiner, "Multiscale image and multiscale deformation of brain anatomy for building average brain atlases," *Proc. SPIE (Medical Imaging 2001: Image Processing)*, vol. 4322, pp. 557–568.

[14] J. Kybic, P. Thévenaz, A. Nirkko, and M. Unser, "Unwarping of unidirectionally distorted EPI images," *IEEE Trans. Med. Imag.*, vol. 19, pp. 80–93, Feb. 2000.

[15] T. Rohlfing, C. R. Maurer Jr, W. G. O'Dell, M. C. Schell, and J. Zhong, "Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images," *Proc. SPIE (Medical Imaging 2001: Image Processing)*, vol. 4319, pp. 337–348.

[16] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recogn.*, vol. 32, no. 1, pp. 71–86, 1999.

[17] R. Schaback, "Creating surfaces from scattered data using radial basis functions," in *Mathematical Models for Curves and Surfaces*. Nashville, TN: Vanderbilt Univ. Press, 1995, pp. 447–496.

[18] M. J. D. Powel, "The theory of radial basis functions in 1990," in *Advances in Numerical Analsysis II: Wavelets, Subdivision, and Radial Basis Functions*, W. Light, Ed. Oxford, U.K.: Oxford Univ. Press, 1992, pp. 105–210.

[19] C. A. Micchelli, "Interpolation of scattered data: Distance matrices and conditionally positive definite functions," *Constr. Approx.*, vol. 2, pp. 11–22, 1986.

[20] J. Schnabel *et al.*, "A generic framework for nonrigid registration based on nonuniform multi-level free form deformation," in *Lecture Notes in Computer Science*, W. J. Niessen and M. A. Viergerver, Eds. Berlin, Germany: Springer-Verlag, 2001, vol. 2208, Proceedings of MICCAI 2001, pp. 573–581.

[21] T. Rohlfing and C. Maurer, "Intensity-based nonrigid registration using adaptive multilevel free-form deformation with an incompressibility constraint," in *Lecture Notes in Computer Science*, W. J. Niessen and M. A. Viergerver, Eds. Berlin, Germany: Springer-Verlag, 2001, vol. 2208, Proceedings of MICCAI 2001, pp. 111–119.

[22] A. P. Zijdenbos, B. M. Dawant, and R. Margolin, "Morphometric analysis of white matter lesions in MR images: Method and validation," *IEEE Trans. Med. Imag.*, vol. 13, pp. 716–724, Aug. 1994.

[23] G. K. Rohde, "The Adaptive grid registration algorithm: A new spline modeling approach for automatic intensity-based nonrigid registration," masters thesis, Vanderbilt Univ, Nashville, TN, 2001.

[24] T. Dinsenbacher, G. K. Rohde, G. Hardin, A. Aldroubi, and B. M. Dawant, "Multiscale nonrigid data registration using adaptive basis functions," in *SPIE-Wavelet Applications in Signal and Image Processing VIII*, vol. 4119, 2000, pp. 1076–1083.

[25] G. K. Rohde, A. Aldroubi, and B. M. Dawant, "Adaptive free-form deformation for inter-patient medical image registration," *Proc. SPIE (Medical Imaging 2001: Image Processing)*, vol. 4322, pp. 1578–1587.

[26] P. Hellier, C. Barillot, E. Memin, and P. Perez, "Hierarchical estimation of a dense deformation field for 3-D robust registration," *IEEE Trans. Med. Imag.*, vol. 20, pp. 388–402, May 2000.