Medical Image Analysis 14 (2010) 633–642

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

GRAM: A framework for geodesic registration on anatomical manifolds

Jihun Hamm [a,*], Dong Hye Ye [b], Ragini Verma [a,b], Christos Davatzikos [a,b]

[a] Department of Radiology, University of Pennsylvania, 3600 Market Street, Suite 380, Philadelphia, PA 19104, USA
[b] Department of Bioengineering, 210 S. 33rd Street, Room 240 Skirkanich Hall, Philadelphia, PA 19104, USA

ARTICLE INFO

Article history:
Available online 8 June 2010

Keywords:
Geodesic registration
Large deformation
Diffeomorphism
Manifold learning

ABSTRACT

Medical image registration is a challenging problem, especially when there is large anatomical variation in the anatomies. Geodesic registration methods have been proposed to solve the large deformation registration problem. However, analytically defined geodesic paths may not coincide with biologically plausible paths of registration, since the manifold of diffeomorphisms is immensely broader than the manifold spanned by diffeomorphisms between real anatomies. In this paper, we propose a novel framework for large deformation registration using the learned manifold of anatomical variation in the data. In this framework, a large deformation between two images is decomposed into a series of small deformations along the shortest path on an empirical manifold that represents anatomical variation. Using a manifold learning technique, the major variation of the data can be visualized by a low-dimensional embedding, and the optimal group template is chosen as the geodesic mean on the manifold. We demonstrate the advantages of the proposed framework over direct registration with both simulated and real databases of brain images.

1. Introduction

1.1. Problem description

Medical image registration plays an indispensable role in the analysis of functional and structural variation of human anatomy. Due to the inevitable differences in the human anatomy in the population under study, an accurate and reliable method is required to transform the images into a common reference frame to perform statistical tests. A large volume of work in the registration methods has been proposed since the 80s. The early developments in image registration method focused on the elastically-constrained deformations (Bajcsy et al., 1983; Bookstein, 1991). In the basic setting, the problem of registering two images boils down to minimizing the weighted sum of dissimilarity and smoothness of the deformation field. However, the problem becomes particularly challenging in the presence of a large shape difference. Despite the fact that smoothness relieves the ill-posedness of this high-dimensional optimization problem, a smoothness term such as Laplacian of the field alone does not guarantee the preservation of topology, which may results in abrupt compressions, expansions, and foldings in the warped images and the loss of one-to-one correspondence. Furthermore, the dissimilarity term such as Mean-Squared Error (MSE) or Mutual Information (MI) is a highly nonlinear function of the deformation field, and therefore the optimization process is likely to be trapped in a local minimum. Simultaneous minimization of dissimilarity and preservation of topology is hard to achieve with a single regularization function of the deformation field. To preserve topology, one can add explicit constraints such as bounds on the determinant of the Jacobian of the fields (Karaçali and Davatzikos, 2004; Haber and Modersitzki, 2007), or restriction of the displacements (Rueckert et al., 2006). However, the difficulty of registering two dissimilar images remains unabated. In this paper, we take an alternative approach to minimize the dissimilarity and preserve the topology; we aim to find a sequence of deformation fields that gradually warps an image to another, as illustrated in Fig. 1.

1.2. Geodesic registration on the manifold of diffeomorphisms

Large deformation registration methods have been proposed (Christensen et al., 1996; Dupuis et al., 1998; Grenander and Miller, 1998) to cope with the shortcomings of the previous approaches. In particular, the geodesic registration methods try to preserve the topology of the deformation by considering the time-varying velocity field $f_t(x)$ the whole path of registration instead of only the final deformation $f(x)$. In its general form, the final deformation is the end point of the flow of a time-dependent velocity vector field $v_t : \Omega \rightarrow \mathbb{R}^D, t \in [0, 1]$, defined by the differential equation: $\frac{df}{dt} = v_t(f_t)$, where $f_{t=0} = Id$ is an identity map. By taking the variational approach, we can find the optimal velocity vector field v_t

* Corresponding author.
E-mail addresses: jihun.hamm@uphs.upenn.edu (J. Hamm), dong.ye@uphs.upenn.edu (D.H. Ye), ragini.verma@uphs.upenn.edu.com (R. Verma), christos.davatzikos@uphs.upenn.edu (C. Davatzikos).

1361-8415/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2010.06.001

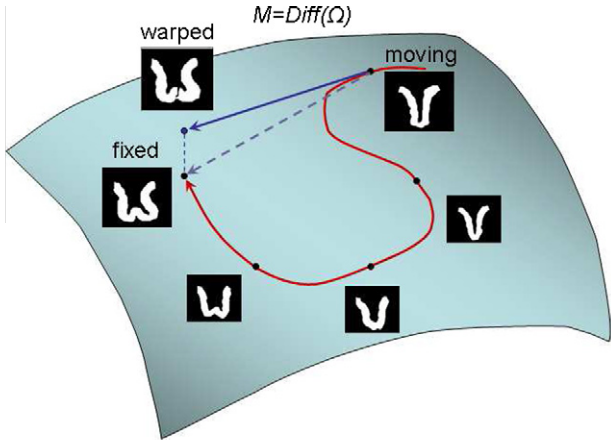

Fig. 1. The figure illustrates the difficulty of registration between two dissimilar shapes. The V-shaped image on the right is the moving image and the W-shaped image on the left is the fixed image. When the smoothing factor is too small, the direct path of minimization (blue curve) using a demons-like algorithm deviates from the space of diffeomorphisms and incurs unnatural distortion such as the tear in the warped image. Increasing the smoothness term, on the other hand, results in a under-registration where the residual MSE remains large. We aim to find the registration path that is both diffeomorphic and biologically relevant (red curve). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in the space of smooth vector fields as the solution to the following cost (Beg et al., 2005):

$$\mathscr{C}(v) = \int_{\Omega} |I_1(f_1^{-1}(x)) - I_2(x)|^2 dx + w \int_0^1 \|v_t\|_V^2, \tag{1}$$

where I_1 and I_2 are two real-valued images or volumes, w is the weight, $\|v_t\|_V$ is a smoothness term such as $\|v_t\|_V = \|Lv\|_2$, and L is a differential operator in the space of velocity vector fields. When the differences between images are only the diffeomorphic change of shape, that is, if they are contained in the single orbit of diffeomorphisms, the similarity term vanishes and the minimum of the cost $\mathscr{C}$ endows the images with a true metric structure: $d^2(I_1, I_2) = \min_v \mathscr{C}(v)$. Recent developments in the geodesic registration method built around this common framework includes the symmetric formulation of the cost function (Avants et al., 2008), and the unbiased estimation of the mean template for groupwise registration (Joshi et al., 2004) to name a few.

1.3. Geodesic registration on the space of learned manifold

Geodesic registration approach has provided a mathematically elegant solution to the large deformation problem. However, the numerical computation of the velocity fields is quite time-consuming which can outweigh its benefits. More importantly, geodesic registration calculates the geodesics on the manifold of diffeomorphisms (Grenander, 1993), which is still a largely unconstrained space. A geodesic path on this space can extend outside the space of "true anatomical variation", which the large deformation kinematics does not prevent from happening. Ideally we want to calculate geodesics on the manifold of transformations that represent only the biologically relevant variation. However, such manifold cannot be represented analytically.

In this paper, we propose a registration framework that achieves this goal. The key idea of the paper is that we approximate analytical geodesic paths with finite sequences of small deformations observed in the actual anatomies in data. In another point-of-view, we are constructing empirical manifolds from data, a technique known as manifold learning (Hamm et al., 2004), instead of dealing with an analytical manifold of diffeomorphisms. In particular we borrow an idea from Isomap algorithm (Tenenbaum et al., 2000), which replaces the geodesic path of the analytical manifold by the shortest path on a k-nearest-neighbor (kNN) graph that approximates the metric structure of the empirical manifold. We refer to our approach as a framework for **G**eodesic **R**egistration on **A**natomical **M**anifolds (**GRAM**).

The GRAM has the following beneficial properties:

1. *Learning of anatomical manifolds*: GRAM computes the geodesics path from the observed anatomical variation of the actual data, which is the key difference to the previous approaches to the large deformation problem. The anatomical manifold learned from a database is reusable: to register test images from a new database to a template image in the old database, we can compute the new registration paths by utilizing the learned deformations.
2. *Efficiency*: Since in this framework the deformations are computed between two close images, we can use simpler and faster registration algorithms such as Diffeomorphic Demons algorithm (Vercauteren et al., 2007), rather than more elaborate algorithms such as Large Deformation Diffeomorphic Metric Mapping (LDDMM) (Beg et al., 2005). In GRAM framework, a registration algorithm is *an interchangeable component*, and therefore different kinds of registration algorithms may be used in the framework (more will be discussed in Section 4). The only requirement is that the component registration algorithm results in diffeomorphic deformation fields for two similar images.
3. *Visualization and automatic template selection*: From the analysis of the shortest paths, GRAM computes a Euclidean embedding of the data which allows us to preview the overall structure of the data such as existence of multiple clusters or the major mode of variation. It also finds an optimal template among the samples for groupwise registration.

1.4. Related work

This paper builds on our previous work (Hamm et al., 2009) and has been extended by new experiments and in-depth analysis of the algorithm. In this framework we adopt the Isomap algorithm (Tenenbaum et al., 2000) to compute and visualize the Euclidean embedding of the metric structure of the data after pairwise registration. Several authors have proposed related algorithms to analyze metric structure of the data and visualize them. Blezek and Miller (2006) proposed Atlas Stratification, which finds multiple modes of the images by mean-shift and visualizes the distribution of the data by Multidimensional Scaling. Images are affinely registered, using Mutual Information as a metric between two images, although it is not a metric, strictly speaking. Sabuncu et al. (2008) proposed an algorithm that also finds the multiple modes of the images by Generalized Expectation-Maximization-based clustering. Images are registered by B-spline, and the membership probability of an image belonging to multiple templates are calculated iteratively. The use of geodesic distances to discover the manifold structure of data, has been proposed by Rohde et al. (2008) and Gerber et al. (2009). These two papers commonly use LDDMM and Multidimensional Scaling to visualize the manifold structure of data, and the latter further uses a kernel regression to reconstruct unseen images from the manifold. However, the two methods directly register all image pairs, which can be difficult and slow for image pairs that are very dissimilar. Our framework distinguishes itself from the aforementioned methods by the following facts: we not only compute the low-dimensional embeddings to visualize the data, but we also compute actual large deformation from each image to a common template for groupwise registration. Furthermore, these large deformations are computed efficiently via sequences of small deformations on the anatomical manifold learned from data.

The remainder of the paper is organized as follows. Section 2 describes the proposed algorithm in detail. Section 3 demonstrates the proposed framework with several simulated and real databases, including simulated 2D images, 3D cortical surfaces from OASIS database, and 3D Fractional Anisotropy map of mouse brains. Section 4 discusses the limitations and extensions of the proposed method, and Section 5 concludes the paper with discussion on the future work.

2. Methods

In this section we provide the algorithmic details of the GRAM framework. The overall training procedure consists of three stages. First, we analyze the data structure by coarse registrations between all image pairs. From this we find a kNN graph structure and a low-dimensional embedding of the data. In the second stage, we choose a template automatically from the graph structure, and identify geodesic paths[1] from the template to other images on anatomical manifolds. In the third stage, we compute the large deformation from the template to each image by composing small deformations between adjacent images along the paths. In addition to the training procedures, we also describe how to use the trained manifold to register a new set of images by updating the previously found geodesic paths. Each stage is described in more detail in the following sections.

Throughout the paper, let us assume the dataset $\mathscr{I}$ consists of n images $I_1,\ldots,I_n$, and each image is a nonnegative real function on a 2D or a 3D domain Ω.

2.1. Construction of empirical manifolds

In the first stage we construct the empirical manifold of data by investigating its metric structure. For this purpose we represent the data as a graph whose vertices correspond to the image samples. Below is the summary of the required steps.

1. Perform coarse registrations between all pair of images. The edge e_{ij} is assigned a weight equal to the distance d_{ij} between two images after registration. The definition of distance d_{ij} is dependent on the specific algorithm used for registration, and we use a weighted sum of a similarity term and a smoothness term.
2. Construct a connected kNN or ϵ-NN graph based on the edge lengths.
3. Find the geodesics (=shortest paths on the graph) between all pairs of vertices, e.g., by Dijkstra's or Floyd-Warshall algorithm. The length g_{ij} of a geodesic is the sum of its edge lengths d_{kl} along the path.
4. (Optional) Visualize the Euclidean embedding of the data by solving eigenvalue problems (refer to Tenenbaum et al. (2000) for details).

The distance d_{ij} is asymmetric in general, that is $d_{ij} \neq d_{ji}$. To make it symmetric we can use the average $0.5(d_{ij} + d_{ji})$, or we can compute d_{ij} for $i < j, i,j = 1,2,\ldots,n$ and assign $d_{ji} = d_{ij}$ to reduce the computation to a half. The latter is possible since d_{ij} and d_{ji} are usually highly correlated. By enforcing symmetry the shortest path length g_{ij} becomes a valid metric, since triangle inequality is fulfilled by the definition of shortest paths.

The size of the neighborhood k in kNN is a parameter the user should select. For a small value of k, the graph is not connected and has multiple disjoint subgraphs. For a too large value of k, the graph becomes completely connected and the shortest path is the same as the direct path. A convenient heuristic is to choose the smallest value that makes the kNN graph connected. More will be discussed in Section 4. An alternative to kNN selection is ϵ-neighbor selection, in which two images I_i and I_j are considered neighbors of each other if $d_{ij} < \epsilon$ for some $\epsilon > 0$. The advantage of this method is that we can strictly set an upperbound to the distance of the edges that will be used for registration. However, finding the smallest ϵ that makes the whole graph connected still requires searching through all values of ϵ.

The most time-consuming part in practice is the pairwise registration between all images which requires $O(n^2)$ number of registrations. To reduce the overhead we can perform the registration on coarse-resolution images of the original data and also use fewer number of iterations than the final registration in the later stage. Although such approximation is not ideal, it may be necessary to keep the computation time practical for databases with a large number of images. To further speed up the pairwise registration, we can distribute the registration tasks over multiple CPUs, since the registration of one pair is independent of the other pairs.

2.2. Automatic template selection

An unbiased template of the given data can be defined as the geodesic mean of the data (Joshi et al., 2004; Avants and Gee, 2004). From the graph derived in the previous section, we can choose a template from the population that is closest to the geodesic mean:

$$I_T = \arg\min_i \sum_j g^2(I_i, I_j),$$

where g is the shortest path length. Since the shortest path length is only an approximation, the chosen template is different from those of Joshi et al. (2004) and Avants and Gee (2004). However, the advantage of this approach is that, the template is chosen with little additional computation. Since we have already computed the geodesic lengths g_{ij}, the template can be chosen by looking up the values.

Two other variants to the mean are the center

$$I_T = \arg\min_i \max_j g(I_i, I_j),$$

and the median

$$I_T = \arg\min_i \sum_j g(I_i, I_j).$$

The three templates look similar in our experiments, but we choose the median as the template due to its resilience to outlying samples in the data.

2.3. Computation of large deformations

We compute the large deformation from the template I_T to any node I_j by a recursive composition of the small deformations from its edges along the geodesic path. Let $f_{i,j}:\Omega_i \to \Omega_j$ denote the deformation field computed from the registration of I_j to I_i. Given the two fields $f_{i,j}$ and $f_{j,k}$, we can easily compute the composition field $f_{i,k} = f_{j,k} \cdot f_{i,j}:\Omega_i \to \Omega_k$ by resampling and interpolating the two fields. The final deformation $\hat{f}_{T,j}$ is the refinement on the composed field $f_{T,j}$ by a few additional iterations of registration. Below is the summary of the procedure.

1. Identify n geodesic paths from I_T to the rest I_j, $\forall j \in 1,\ldots,n$.

[1] From now on, a "geodesic path" refers to the shortest path on the anatomical manifold which will be clear from the context.

2. Enumerate all edges $\mathscr{E}$ used in any of the shortest paths. Perform accurate registration between $(I_i, I_j),\ \forall e_{ij} \in \mathscr{E}$.
3. For each $j \in 1,\ldots,n$,
   (a) Let $s = (s_1 = T,\ldots,s_m = j)$ be the geodesic path from I_T to I_j.
   (b) If $f_{T,j}$ is already computed then exit.
   (c) Otherwise, recursively compute $f_{s_1,s_m} = f_{s_{m-1},s_m} \cdot f_{s_1,s_{m-1}}$.
   (d) Fine-tune f_{s_1,s_m} by additional iterations of registration.

Note that we needed only coarse registration results in the previous stages, and this stage is where we actually perform accurate registrations. Step 2 may seem to be a huge computational burden at first since the number of all the edges in a graph can be as large as n^2. In fact, we only need to update the registration for $n-1$ edges, that is, no more than the number of direct registration for a conventional approach. This is due to property of the graph that the shortest paths from the template to the rest forms a spanning tree. Furthermore, the registration converges faster since the two adjacent images are similar by construction. The condition that each deformation field of the edge being diffeomorphic is sufficient for the composed field to be diffeomorphic as well.

The fine-tuning is a crucial part of the procedure. It is required since the transitivity (Christensen and Johnson, 2003) is not guaranteed for registration algorithms in general, that is, the composed field $f_{j,k} \cdot f_{i,j}$ of the two registration results is not the same as the field $f_{i,k}$ computed directly from the registration between I_i and I_k. In summary, the composed field serves as the initial field to start the registration which helps to avoid the local minimum of direct registration path, and the fine-tuning serves as the minimization of the transitive error.

2.4. Registration of new data

The learned manifold of training images can be used to facilitate the registration of new test images not included in the training database. When the new images are introduced, the manifold can be reused without recomputing the geodesic paths from the beginning. The geodesic deformation for the test image can be computed by registering the new image to the closest image in the training data and then composing the field with the known learned deformation field of the closest image to the template. Below is the summary of the procedure.

1. Register the new test image to the training images to compute the distances $\tilde{d}_i$ and the deformation $\tilde{f}_i$, where $i = 1,\ldots,n$ is the index of the training images.
2. Update the distance from the template to the test image by adding $\tilde{d}_i$ and $d_{T,i}$, where $d_{T,i}$ is the known distance from the template to I_i.
3. Choose the shortest path from above.
4. Compose the fields $\tilde{f}_i$ and $f_{T,i}$, where $f_{T,i}$ is the known field from the template to I_i.
5. Fine-tune the field by additional iterations of registration.

For this approach to work, the new dataset must not be too heterogeneous to the training dataset. Otherwise, the new data will be equally distant from all training images and gain no benefit from the learned deformations of the training data.

3. Experiments

In this section we test the proposed framework with several simulated and real databases, including simulated 2D images, 3D cortical surfaces from OASIS database, and 3D Fractional Anisotropy volume of mouse brains. To demonstrate its advantages, we compare the proposed method with the direct registration method which does not use the geodesic path. Since we do not have ground truth for the 'best' registration for these databases, we measure the quality of the registration results in terms of MSE, Harmonic Energy (HE)[2] and Maximum Jacobian Determinant (MJD) where maximum is computed over all voxels.[3]

3.1. Validation with simulated data

We first test the proposed framework on a dataset of simulated 2D cortical patches. The aim of this section is to demonstrate the properties of the proposed method and to check the validity of the algorithm under varying parameters. The data consist of 60 binary 2D images of size 140×140 which simulate a patch of a cortex varying in the thickness and the number of folds. We use an ITK (Ibanez et al., 2005) version of the Diffeomorphic Demons by Vercauteren et al. (2007) for registration due to its fast speed. The images are registered with three levels of resolution for coarse pairwise registration, and with the original resolution for fine-tuning, with a smoothness parameter of $\sigma = 1.5$. The whole procedure takes about an hour on our cluster server (Sun Grid Engine), which has 22 multi-core nodes and 4–8 GB. Since the server is a shared resource, the exact time can vary. The computing time of the first stage, which is the dominant stage, can be estimated more accurately from the equation $0.5n(n-1)T$, where n is the number of images and the T is the average time to register a single image to a template under a given computing resource.

From the coarse pairwise registration we define the distance in Section 2.1 as the weighted sum of: (1) MSE between the fixed and the warped images and (2) HE of the deformation field:

$$d_{ij} = w\,\mathrm{MSE}(I_i, I_j(f_{i,j})) + (1-w)\,\mathrm{HE}(f_{i,j}). \tag{2}$$

The smoothness parameter affects the registration results significantly. A too small value of σ reduces the final MSE, but also increases HE and MJD significantly (over-registration). A too large value of σ can make the final MSE many times larger than it is with a small σ although it reduces HE and MJD (under-registration). Since the parameter selection is the choice associated with the component algorithm and not with our framework, we do not perform repeated experiments for a full range of σ. Instead, we have chosen an appropriate σ by checking that the deformation field has no negative Jacobian, that is, the field is diffeomorphic. However, the parameter w and k remains to be decided.

We first show the results with a fixed value of $w = 0.75$ and $k = 16$.[4] Fig. 2 shows the two-dimensional Euclidean embedding of the simulated data. The embedding conveniently summarizes the major shape variation of the population which have three prototypical shapes (which resemble the letter U, V, and W).

To measure improvements in registration due to the geodesic approach, we calculate the relative change of MSE

$$\delta\mathrm{MSE} = 100 \times \frac{\mathrm{MSE}_{\mathrm{geodesic}} - \mathrm{MSE}_{\mathrm{direct}}}{\mathrm{MSE}_{\mathrm{direct}}},$$

and similarly for HE and MJD.

We perform the groupwise registration using the automatically chosen template with the proposed method and the direct method. The registration results are shown in Fig. 3 which shows the paths and the final warped images of the five samples which has the largest decrease in MSE. The image warped by the proposed method is

[2] Harmonic Energy is the mean Frobenius norm of the Jacobian of the deformation field.

[3] We report the 99 percentile of the Jacobian Determinant instead of the maximum since the maximum is prone to noise.

[4] The w here is not an absolute value but a relative weight between the similarity term and the smoothness term. We normalize the two term to have a unit l_2 norm summed over all images.

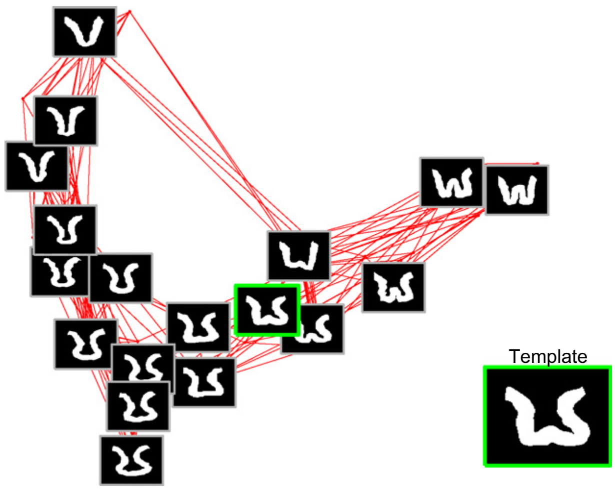

Fig. 2. Two-dimensional embedding of the manifold of simulated shapes. Only a subset of the samples is shown to avoid clutter. The template determined to be the median of the graph is marked by a green box, and the red lines denote the nearest-neighbor relationship. The embedding reveals that there are three major variants (which resemble the letters U, V, and W) and the rest of the images lie in-between the three prototypical shapes. The template marked by the green box is chosen from the median of the geodesic distances. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

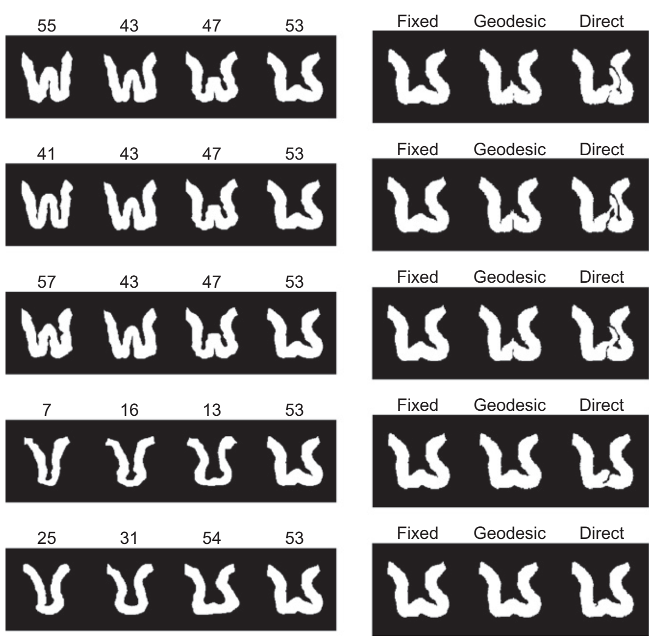

Fig. 3. Left: geodesic paths of simulated shapes. The images are sample paths from the leftmost image (moving) to the rightmost image (fixed). The number on top of each image is the sample index. Note the gradual change of shape along each path. Right: comparison of the final warped images from the geodesic versus the direct registration using the same registration method and parameters. Warping the W-shaped images (55, 41, 57) to the fixed image (53) requires a large deformation near the middle fold in the image. The proposed method finds such path that gradually flattens the middle fold, whereas the direct registration aggressive fit the image by squeezing the middle fold towards the right side of the image, resulting in artificial fissures in the image.

noticeably better than those of the direct method which present severe distortions. These samples have 43.4%, 40.1%, 38.1%, 37.6%, and 37.2% decrease in MSE, respectively. The average decrease of MSE over all samples is 11%, and the number of samples that shows decrease is 79%. Table 1 summarizes the relative changes in MSE, HE and MJD. Although the decrease in MSE is of our main concern, HE also decreases significantly. This shows us that our framework can achieve more accurate and smoother registration simultaneously on average. Also note that in the worst case, geodesic method results in higher errors than direct method. The explanation for this is as follows. The registration errors from geodesic method have two opposing factors. One is the desirable decrement due to the avoidance of the local minimum of direct registration, and the other is the undesirable increment from the transitive error of the composition of deformation fields. Our experiment shows that the summed effect of the two factors is beneficial on average, but it can be negative for a fraction of the whole samples. In practice, we can always register images using both direct and geodesic methods and choose the better of the two methods for each sample, since direct registration using Demons is computationally inexpensive.

To check the robustness of the framework to the change of parameters, we repeat the experiments with three values of w (0.25, 0.5, 0.75). As we mentioned in Section 2.1, a heuristic of choosing k is to find the smallest value c that makes the kNN graph connected. We also repeat the experiments with $k = c$, $c + 2$, $c + 4$. The w and k change the topology of kNN graph and subsequently the paths and the template. Fig. 4 shows two-dimensional Euclidean embeddings with these parameters. The overall shape of the embedding and the chosen templates seems to be affected by the parameters. However, the groupwise registration results of Table 2 shows that the improvements in MSE, HE and MJD vary within a small range. Note that MSE and HE decrease consistently whereas the average MJD increases sometimes, which may be due to the fact that HE and MJD measure different aspects of 'smoothness'. From these experiments we conclude tentatively that a small difference in the parameters does not adversely affect the final outcome.

3.2. Registration of new data

We demonstrate the capability of our method described in Section 2.4: the learned manifold of the training samples can be used to facilitate the registration of new test images not included in the training database. For this purpose we generate additional simulated images that are similar to but different from the training images.

To visualize the test images along with the training images, we need to compute the coordinates of the new test points in the embedding of the training images. To do this, we first register the test images to the training images and compute the distances from (2). Using these new distances, and the known embedding and pairwise distances of the training images, we compute the coordinates from the algorithm described in (de Silva and Tenenbaum, 2002). Fig. 5 shows the two-dimensional Euclidean embedding of four test images superimposed on the embedding of the trained simulated data. The embedding provides information on the homogeneity (or heterogeneity) of the test data to the training data. One of the test images is slightly apart from the training population due to its relatively distinctive shape, whereas the remaining test images blend well into the population.

We register the test images to the template determined from the training data, using the method in Section 2.4. The registration results are shown in Fig. 6 which shows the paths and the final warped images of the test images. Compared with the MSE obtained by registering the test images to the template directly, the MSE obtained from the proposed method has a decrease of 3.4%, 34.6%, 3.2%, and 38.8%, for the four test images, respectively.

3.3. 3D cortical surfaces of human brains

We test the algorithm on a database of real brain images. The Open Access Series of Imaging Studies (OASIS) databases is a

Table 1
Summary of registration results in simulated data. The decrease in Mean Square Error (MSE), Harmonic Energy (HE), and Maximum Jacobian Determinant (MJD) are shown in percentage. A higher value means a larger improvement.

Measure	MSE			HE			MJD		
	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)
Decrease	10.7 ± 16.1	43.4	−16.3	8.9 ± 8.4	32.4	−13.5	0.7 ± 10.8	30.7	−32.4

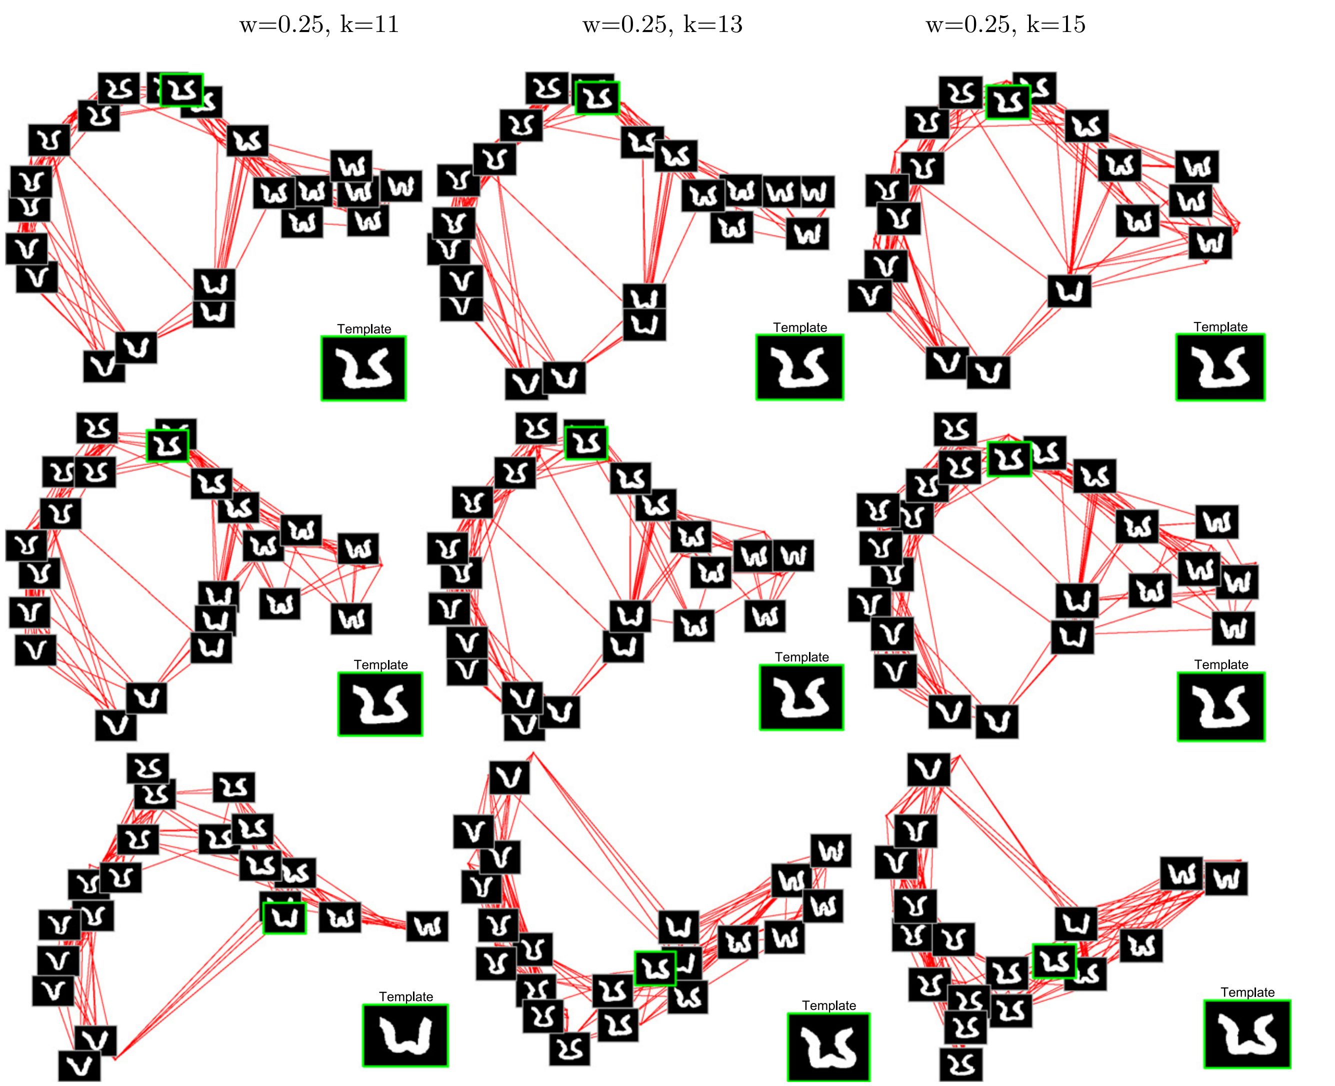

Fig. 4. Embeddings and templates under varying parameters. These nine figures looks different but they all show the gradual variation of the shapes between the three prototypical shapes (U, V, and W).

Table 2
Summary of registration results in simulated data with different parameters. The maximum and the average decrease in Mean Square Error (MSE), Harmonic Energy (HE), and Maximum Jacobian Determinant (MJD) are shown in percentage. A higher value means a larger improvement.

Measure	MSE		HE		MJD	
	Mean (%)	Max (%)	Mean (%)	Max (%)	Mean (%)	Max (%)
$w = 0.25$, $k = 11$	8.9	43.1	3.0	40.1	1.3	40.2
$w = 0.25$, $k = 13$	8.2	44.4	3.6	40.3	−2.3	32.8
$w = 0.25$, $k = 15$	7.8	43.9	4.2	40.3	−2.0	32.8
$w = 0.50$, $k = 11$	8.1	47.5	4.3	40.3	−1.4	41.1
$w = 0.50$, $k = 13$	8.4	45.8	4.2	40.3	−1.8	41.1
$w = 0.50$, $k = 15$	9.9	43.9	3.4	40.7	−2.5	41.1
$w = 0.75$, $k = 12$	5.5	39.2	3.7	33.8	1.1	19.0
$w = 0.75$, $k = 14$	10.9	43.4	8.5	32.4	0.7	30.7
$w = 0.75$, $k = 16$	10.7	43.4	8.9	32.4	0.7	30.7

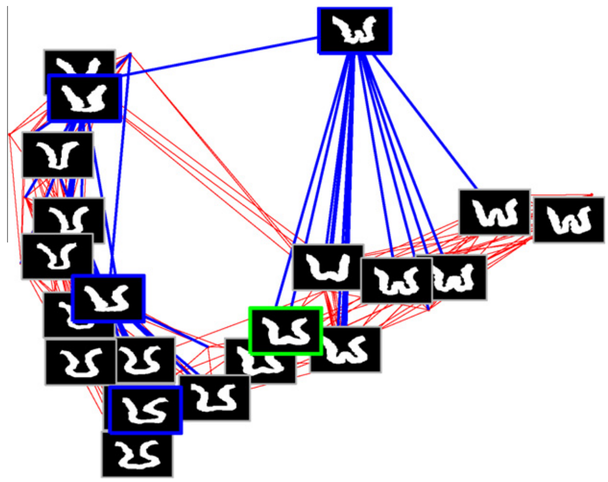

Fig. 5. Two-dimensional embedding of four test images (marked by blue boxes) superimposed on the embedding of the training data. The template is marked by a green box. The nearest-neighbors of the training data are connected by red lines, and the nearest-neighbors of the test data within the training data are connected by thick blue lines. Note that the similarity of the test images and their neighboring images in the training data, except for one test image on the upper right that is relatively distinctive from the other images. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

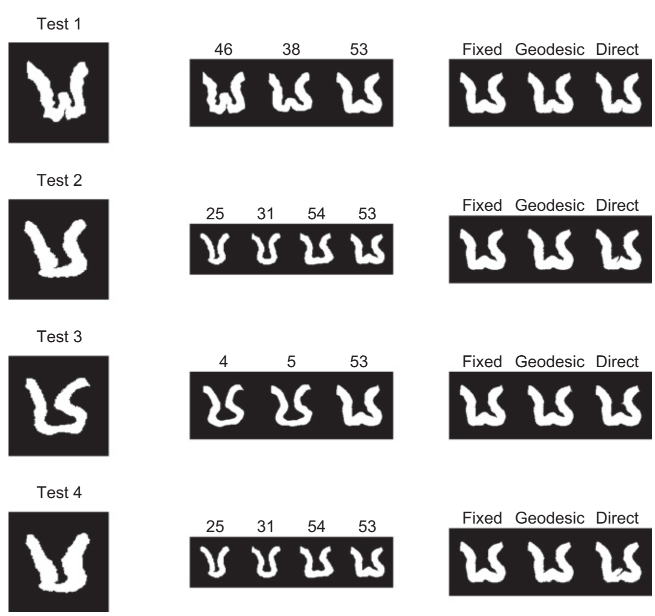

Fig. 6. Left: four randomly generated test images. Middle: geodesic paths on the training data corresponding to the test images. The leftmost image is the closest training image to the test image. Right: comparison of the final warped images from the geodesic versus the direct registration using the same registration method and parameters.

publicly available collection of MRIs (Marcus et al., 2007). This data set consists of a cross-sectional collection of 416 subjects covering the adult life span aged 18–96 including individuals with neurodegeneration. The subjects are all right-handed and include both men and women. One hundred of the included subjects over the age of 60 have been clinically diagnosed with very mild to moderate Alzheimer's disease. In this study we focus on the variation of cortical patterns in a small volume of interest (VOI). The VOI is cropped in the region that contains right superiorfrontal cortex. We use the segmentation provided with the data to extract surfaces between the gray matter and the cerebrospinal fluid. The size of each volume is resized to $68 \times 56 \times 72$ and affinely aligned.

Out of 416 images we discard 23 outlier images that are not connected to the rest of the data with 24-NN graph. The images are registered with three levels of resolution for the coarse pairwise registration, and with two levels of resolution for fine-tuning, with a smoothness parameter of $\sigma = 1.0$. The whole procedure takes about 24 h on our cluster server.

Fig. 7 shows the two-dimensional embedding of the OASIS data. The cortical surfaces of the VOI are rendered to aid visualization of the results, using the curvature information computed from the smoothed surface. At a glance, the OASIS data contain complex variation of cortical patterns in contrast to the simulated data. Note that in the first axis (from left to right) the embedding shows change in the depth of sulcus/gyrus which may be ascribed to the atrophy of the subjects with age and the Alzheimer's disease. The biological plausibility of the geodesic paths are demonstrated in Fig. 8 with the four samples that have the largest decrease in MSE. Since the samples 54, 52, 188, and 221 are quite different from the fixed image 108, the registration is still not perfect. However the circled areas in the figure shows that proposed method can avoid unnatural collapsing of the gyri in the direct registration method and produces more realistic patterns. The advantage is also evidenced by the improvement of MSE in the four samples (17%, 16%, 14% and 14%, respectively).

We now look at the overall statistics of the data. The distribution of length of the paths (=the number of vertices along the path) is as follows. The numbers of the paths of length 2, 3, 4, 5, and 6 are: 24 (6%), 196 (50%), 139 (35%), 30 (7%), and 3 (1%), respectively. The average decrease in MSE for these paths are 0%, 2.8%, 2.3%, 1.9% and -0.8%. The paths of length 2 have no change obviously since there is no intermediate sample in the path. The number of paths of length 6 are only three and shows no improvement in MSE.

Table 3 summarizes the improvements by geodesic paths. This also shows that we achieve improvements in both MSE and the smoothness measures HE and MJD, although the average amount of improvement is less significant than the simulated data. In this experiment we also use Demons algorithm with the segmented

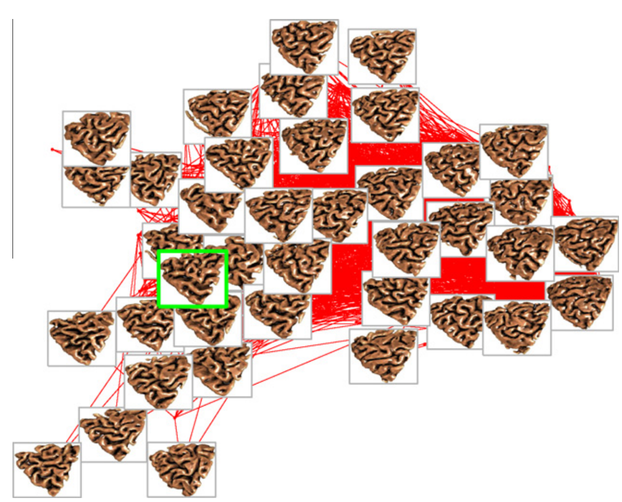

Fig. 7. Two-dimensional embedding of the manifold of OASIS. The cortical surfaces of the VOI are superimposed on the embedding. Only a subset of the samples is shown to avoid clutter. The template determined to be the median of the graph is marked by a green box, and the red lines denote the nearest-neighbor relationship. Although the variation of the sulcul and gyral patterns in the embedding is too complex to describe concisely, there is tendency of atrophy left to right (compare the leftmost surfaces with rightmost surfaces). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

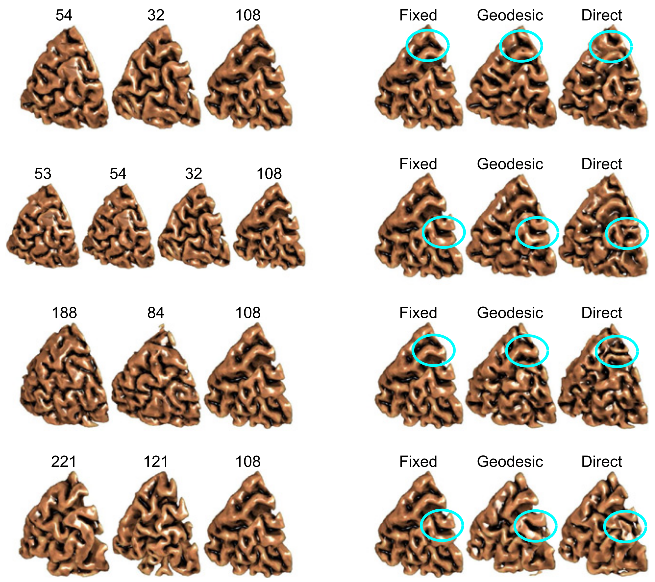

Fig. 8. Left: geodesic paths of OASIS data. The images are sample paths from the leftmost surface (moving) to the rightmost surface (template). Note the relatively gradual change of cortical patterns through the paths. Right: comparison of the final warped surface from the geodesic versus the direct registration using the same registration method and parameters. The results of former method present more smooth and realistic warping of the cortices when compared to the unnatural warping from the latter method. Representative regions in which the geodesic registration is markedly better are shown by cyan-colored circles. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

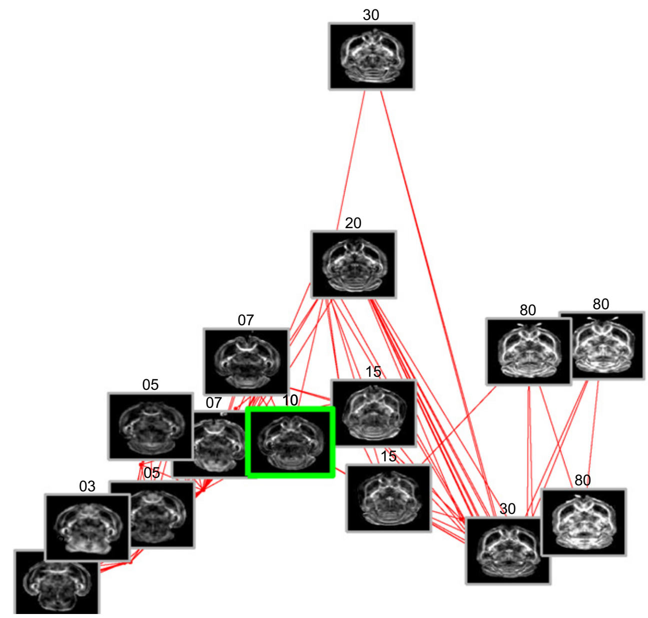

Fig. 9. Two-dimensional embedding of the manifold of mouse FA map. A mid-axial slice is shown for each mouse brain volume. Only a subset of the samples is shown to avoid clutter. The template determined to be the median of the graph is marked by a green box, and the red lines denote the nearest-neighbor relationship. The number on top of each image indicates the age of the brain. The embedding reveals that the major variation of the data is the age factor which increases gradually from left to right. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

volumes. To study the cortical patterns better, we plan to use a surface-based registration such as Spherical Demons (Yeo et al., 2009).

3.4. Fractional Anisotropy map of mouse brains

Finally, we show that the proposed method can be applied to image database that has large variation in both shape and appearance. Data of mouse brains are collected in our lab with the aim of creating a normative atlas of a developing mouse brain. The data consist of 69 Fractional Anisotropy maps of the brains sampled at 2, 3, 4, 7, 10, 15, 20, 30, 45, and 80 days of age. Each volume is resized to $150 \times 150 \times 100$ and affinely aligned. The images are registered with three levels of resolution for coarse pairwise registration, and with the original resolution for fine-tuning, with a smoothness parameter of $\sigma = 1.5$. The whole procedure takes about 6 h on our cluster server. The images in this dataset not only have a larger number of voxels than the other experiments but they are more challenging for registration due to their large shape and appearance variation from different ages and the degrees of maturation of tracts.

The two-dimensional embedding of the data in Fig. 9 provides a glimpse of its manifold structure. From the figure we can observe that the major variability of the data comes from age. The importance of the age factor is also observed in Fig. 10: a path that connects two brain images of different ages passes through brains of intermediate ages in a monotonic fashion. These findings are consistent with our prior knowledge of the data that the developmental stage is the major factor of the variation in the data. Fig. 10 also shows that the proposed method produces better registration results than those from the direct method. The decrease of MSE is 13.0%, 8.3%, 7.8%, 7.6%, and 6.8% for the five examples, respectively.

Table 4 summarizes the overall improvements by geodesic paths. MSE and HE decrease significantly (especially HE), and MJD remains unchanged. For mouse data we use histogram normalized intensity difference to compute MSE and geodesic distances. However, the large appearance variation in addition to the shape variation may require different model of the data manifold and revised definitions of the metric (Trouvé and Younes, 2005), which is out of the scope of this paper.

4. Discussion

In this section we discuss several aspects of the proposed framework and their practical implications for registration.

4.1. Number of samples

The proposed registration method is motivated by Isomap algorithm, which is based on the premise that the true geodesic on a convex set can be approximated well by the shortest path on the

Table 3
Summary of registration results in the OASIS data. The decrease in Mean Square Error (MSE), Harmonic Energy (HE), and Maximum Jacobian Determinant (MJD) are shown in percentage. A higher value means a larger improvement.

Measure	MSE			HE			MJD		
	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)
Decrease	2.6 ± 4.4	17.0	−7.7	2.0 ± 7.4	21.4	−25.1	9.7 ± 14.6	46.8	−63.0

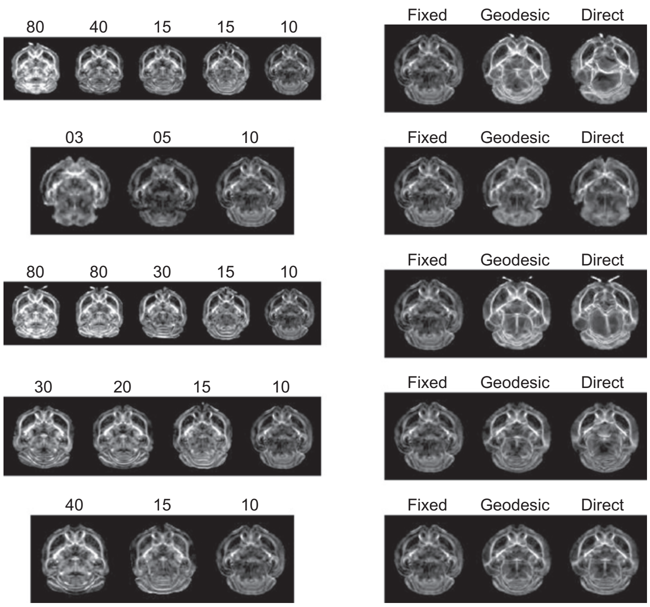

Fig. 10. Left: geodesic paths of mouse FA map. The images are sample paths from the leftmost image (moving) to the rightmost image (fixed). Each path reflects the changes in both the shapes and the appearances of developing brains. The numbers on top is the age of the brain which can be the same for two different images. Note that the age either increases or decreases monotonically through the paths. Right: comparison of the final warped images from the geodesic versus the direct registration using the same registration method and parameters. Brains in different developmental stages are quite different, and therefore it is hard to impose strict one-to-one correspondences. However, the warped images from the geodesic method are more similar to the fixed image than those of the direct method. Note the asymmetry of the latter images in the first and the third examples.

*k*NN graph connecting the data samples. Therefore the framework also inherits the limitations of Isomap. The number of samples necessary increases exponentially with the intrinsic dimensionality of the data, which cannot be determined a priori. For the shortest paths to be faithful approximations we need a large database whose size is proportional to the degree of freedom of variation in the data. However, the number of available images in a study is typically limited to a few hundreds at most. Considering that a brain image lies in a huge-dimension Euclidean space of $O(10^6)$ voxels, a few hundred is still a relatively small number. Nevertheless, the small amount of data in our experiments have been shown to provide improved registration results compared to direct registration, despite the approximate nature of our geodesic paths.

4.2. Multiple clusters

Our model assumes that the whole data lie on a single manifold of deformation variation. This assumption can be restrictive when the data are clustered around a few distant cluster centers rather than evenly distributed on a low-dimensional manifold. In that case the size k of the *k*NN to make the graph connected can grow very large (>100), which undermines the advantages of our method. However, even in the extreme case that every image is a neighbor of the other images, the method simply reverts to the coarse pairwise registration and no worse. Such cases may be handled by imposing connectivity of the graph via a minimum spanning tree, or by computing a few large deformation paths that connect the clusters via a numerical geodesic registration method. However, if the data are known to have disjoint clusters a priori, approaches based on the cluster assumption will be more appropriate to analyze the data (Blezek and Miller, 2006; Sabuncu et al., 2008).

4.3. Component registration algorithms

As we stated in the introduction, the component registration algorithm of the framework is interchangeable as long as the field it produces is diffeomorphic between two nearby images. There are many alternatives to Demons algorithm we used in this paper, including B-spline free-form deformation (Rueckert et al., 1999), elastically deformable model (Davatzikos, 1997), and feature-based algorithms such as HAMMER (Shen and Davatzikos, 2002) and DRAMMS (Ou and Davatzikos, 2009). Furthermore, the framework can be adopted for registering different representations of imagery, such as point set (of landmarks), curves, or surfaces. Depending on the component algorithm and the data types, the definition of distance between two images has to change accordingly. Note that such distance need not strictly be a true metric or a Riemannian distance since the shortest path on the graph impart the metric properties to the geodesic distance. The question of which algorithm and representation is optimal for the given data, is left to empirical studies.

5. Conclusion

In this paper, we propose a novel framework for Geodesic Registration on Anatomical Manifold (**GRAM**). The most distinguishing feature of the method is that it computes the geodesics on the manifold of the anatomical variation learned from the data, instead of computing the analytic geodesics of all diffeomorphisms. This warrants that any deformation field, as well as geodesic path, calculated in our framework represents real brain morphology, and is not merely a diffeomorphic transformation of a template, which can represent an unrealistically distorted morphologies. The learned manifold also provides a visualization of the data structure and allows us to choose an optimal template among the samples for groupwise registration. The experiments on simulated images, human cortical surfaces, and mouse FA maps show that the proposed method can achieve smaller MSEs with smoother deformation fields than those computed without using the geodesic paths. This attests to the hypothesized benefits of utilizing anatomical variation of the actual data. It is left as our future work to perform cross-database tests using the framework and to compare the results with numerical geodesic registration methods.

Finally, GRAM is intended to be a meta-registration framework to efficiently compute large deformations, which allows a large class of registration algorithms to be used as its component. The code for GRAM framework will be made available on the web to encourage evaluation from the community.

Table 4
Summary of registration results in the mouse data. The decrease in Mean Square Error (MSE), Harmonic Energy (HE), and Maximum Jacobian Determinant (MJD) are shown in percentage. A higher value means a larger improvement.

Measure	MSE			HE			MJD		
	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)	Mean ± std (%)	Max (%)	Min (%)
Decrease	3.1 ± 3.2	13.1	−4.1	10.6 ± 6.0	29.1	3.4	0.1 ± 1.1	2.5	−2.7

Acknowledgments

We thankfully acknowledge the support from NIH Grants R01-MH079938 and R01-AG014971. We also acknowledge the use of the OASIS database supported with Grants P50 AG05681, P01 AG03991, R01 AG021910, P20 MH071616, and U24 RR021382.

References

Avants, B., Gee, J.C., 2004. Geodesic estimation for large deformation anatomical shape averaging and interpolation. NeuroImage 23, S139–S150.
Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical Image Analysis 12, 26–41.
Bajcsy, R., Lieberson, R., Reivich, M., 1983. A computerized system for the elastic matching of deformed radiographic images to idealized atlas images. Journal of Computer Assisted Tomography 7 (4), 618.
Beg, M.F., Miller, M.I., Trouvé, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61 (2), 139–157.
Blezek, D.J., Miller, J.V., 2006. Atlas stratification. MICCAI, 712–719.
Bookstein, F., 1991. Thin-plate splines and the atlas problem for biomedical images. Information Processing in Medical Imaging 511, 326–342.
Christensen, G., Rabbitt, R., Miller, M., 1996. Deformable templates using large deformation kinematics. IEEE Transactions on Image Processing 5 (10), 1435–1447.
Christensen, G.E., Johnson, H.J., 2003. Invertibility and transitivity analysis for nonrigid image registration. SPIE Electronic Imaging 12 (1), 106–117.
Davatzikos, C., 1997. Spatial transformation and registration of brain images using elastically deformable models. Computer Vision of Image Understanding 66 (2), 207–222.
de Silva, V., Tenenbaum, J.B., 2002. Global versus local methods in nonlinear dimensionality reduction. Advances in Neural Information Processing Systems 15, 705–712.
Dupuis, P., Grenander, U., Miller, M., 1998. Variational problems on flows of diffeomorphisms for image matching. Quarterly of Applied Mathematics 56 (3), 587.
Gerber, S., Tasdizen, T., Joshi, S.C., Whitaker, R.T., 2009. On the manifold structure of the space of brain images. MICCAI, 305–312.
Grenander, U., 1993. General Pattern Theory: A Mathematical Study of Regular Structures. Oxford University Press, USA.
Grenander, U., Miller, M.I., 1998. Computational anatomy: an emerging discipline. Quarterly of Applied Mathematics LVI (4), 617–694.
Haber, E., Modersitzki, J., 2007. Image registration with guaranteed displacement regularity. International Journal of Computer Vision 71 (3), 361–372.
Hamm, J., Davatzikos, C., Verma, R., 2009. Efficient large deformation registration via geodesics on a learned manifold of images. MICCAI, 680–687.
Hamm, J., Lee, D.D., Mika, S., Schölkopf, B., 2004. A kernel view of the dimensionality reduction of manifolds. In: ICML '04: Proceedings of the Twenty-first International Conference on Machine Learning, p. 47+.
Ibanez, L., Schroeder, W., Ng, L., Cates, J., 2005. The ITK Software Guide. Kitware, Inc. <http://www.itk.org/ItkSoftwareGuide.pdf>.
Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23 (Suppl. 1), 151–160.
Karaçali, B., Davatzikos, C., 2004. Estimating topology preserving and smooth displacement fields. IEEE Transaction on Medical Imaging 23 (7), 868–880.
Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open access series of imaging studies (OASIS): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. Journal of Cognitive Neuroscience 19 (9), 1498–1507.
Ou, Y., Davatzikos, C., 2009. Dramms: deformable registration via attribute matching and mutual-saliency weighting. Information Processing in Medical Imaging, 50–62.
Rohde, G.K., Ribeiro, A.J.S., Dahl, K.N., Murphy, R.F., 2008. Deformation-based nuclear morphometry: capturing nuclear shape variation in hela cells. Cytometry Part A 73A (4), 341–350.
Rueckert, D., Aljabar, P., Heckemann, R.A., Hajnal, J.V., Hammers, A., 2006. Diffeomorphic registration using b-splines. MICCAI, 702–709.
Rueckert, D., Sonoda, L.I., Hayes, D.L.G.H.C., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transaction on Medical Imaging, 712–721.
Sabuncu, M.R., Balci, S.K., Golland, P., 2008. Discovering modes of an image population through mixture modeling. MICCAI 11 (Pt 2), 381–389.
Shen, D., Davatzikos, C., 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. IEEE Transactions on Medical Imaging 21 (11), 1421–1439.
Tenenbaum, J.B., de Silva, V., Langford, J.C., 2000. A global geometric framework for nonlinear dimensionality reduction. Science 290 (5500), 2319–2323.
Trouvé, A., Younes, L., 2005. Metamorphoses through lie group action. Foundations of Computational Mathematics 5 (2), 173–198.
Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2007. Non-parametric diffeomorphic image registration with the demons algorithm. MICCAI, 319–326.
Yeo, B.T.T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P., 2009. Spherical demons: fast diffeomorphic landmark-free surface registration. IEEE Transactions on Medical Imaging.