Surface-Constrained Volumetric Brain Registration Using Harmonic Mappings

Anand A. Joshi, David W. Shattuck, Paul M. Thompson and Richard M. Leahy

Abstract—In order to compare anatomical and functional brain imaging data across subjects, the images must first be registered to a common coordinate system in which anatomical features are aligned. Intensity-based volume registration methods can align subcortical structures well, but the variability in sulcal folding patterns typically results in misalignment of the cortical surface. Conversely, surface-based registration using sulcal features can produce excellent cortical alignment but the mapping between brains is restricted to the cortical surface. Here we describe a method for volumetric registration that also produces an accurate one-to-one point correspondence between cortical surfaces. This is achieved by first parameterizing and aligning the cortical surfaces using sulcal landmarks. We then use a constrained harmonic mapping to extend this surface correspondence to the entire cortical volume. Finally, this mapping is refined using an intensity-based warp. We demonstrate the utility of the method by applying it to T1-weighted magnetic resonance images (MRI). We evaluate the performance of our proposed method relative to existing methods that use only intensity information; for this comparison we compute the inter-subject alignment of expert-labeled sub-cortical structures after registration.

Index Terms—Image Registration, deformable registration, brain mapping, harmonic mapping.

I. INTRODUCTION

Morphometric studies of anatomical changes over time or of differences between populations require that the data first be transformed to a common coordinate system in which anatomical structures are aligned. Similarly, inter-subject longitudinal studies or group analyses of functional data also require that the images first be anatomically aligned. Alignment is commonly performed either with respect to the entire volumetric space [1] or is restricted to the cortical surface [2]. Here we describe an approach to brain image registration based on harmonic maps that combines these two approaches producing a volumetric alignment in which there is also a one-to-one correspondence between points on the two cortical surfaces.

Talairach normalization based on a piecewise affine transformation [3] was the first commonly used volumetric alignment technique. Because it uses a restricted set of anatomical landmarks and is piecewise affine, it results in relatively poor alignment and has been largely replaced by automated intensity-based alignment methods that also allow non-rigid deformations [4], [5]. There are a vast array of such methods,

differing in how they measure the fit between the two images (e.g., squared error, correlation, mutual information), the parameterization of the transformation (e.g., polynomial, splines, discrete cosine transform or other eigenfunction bases), and the procedure used to regularize the transformation (e.g., elastic, biharmonic, or viscous fluid models) [6]. Polynomial warps and linear elastic deformations implicitly assume that deformations are small and do not guarantee preservation of topology for larger deformations [7]. The viscous fluid approach [8] and more recent approaches using large-deformation diffeomorphic metric mapping [9], [10] were developed to address the problem of ensuring diffeomorphic maps and are better able to register objects whose alignment requires large deformations while conserving their topology.

Since these intensity-based methods do not explicitly model the cortical surface, alignment can be rather poor. An illustration of this is shown in Fig. 1, where we have used the Automated Image Registration (AIR) software [5], [11] to align two brain volumes using a 5th order polynomial (168 parameters). While the regions of cortical grey matter exhibit reasonably good correspondence between the two images, the cortical surfaces themselves do not align well. Since cytoarchitectural and functional parcellation of the cortex is intimately related to the folding of the cortex, it is important when comparing cortical anatomy and function in two or more subjects that the surfaces are aligned. For this reason, there has been an increasing interest in analyzing the cerebral cortex based on alignment of surfaces rather than volumes.

Various surface-based techniques have been developed for inter-subject registration of two cortical models. One class of techniques involves flattening the two cortical surfaces to a plane [12] or to a sphere [13] using mechanical models or variational methods and then analyzing the data in the common flattened space [14]. Other surface based techniques work in the surface geometry itself rather than a plane or a sphere and choose to account for the surface metric in the inter-subject registration [15], [16]. The advantage of such techniques is that they produce registration results that are independent of the intermediate flat space (or, equivalently, the specific parameterization of the cortex) resulting in a more consistently accurate registration throughout the cortex. These approaches involve manually delineated sulcal landmark matching [16] in the intrinsic surface geometry. While some progress has been made recently towards automating the matching process using mutual information [17] or optical flows of mean-curvature images in the surface parameter space [18], [19], fully automatic alignment of high resolution cortical surfaces remains a challenging problem.

While the volume registration methods described above do not provide suitable cortical alignment, the cortical registration methods do not define any volumetric correspondence. One approach to this problem is to combine landmark points, curves and surfaces as additional constraints in an intensity-based warping method [20], [2], [21], [22], [23], [24], [25]. For example, landmarks, curves [25] and image matching [24] can be applied in a hierarchical manner in a large deformations framework ensuring generation of diffeomorphisms [26], [27]. Registration methods such as the Hierarchical Attribute Matching Mechanism for Image Registration (HAMMER) algorithm [28] incorporate surface as well as volume information for the alignment using geometric attributes of the images. Alignment of brain images often involves relatively large displacements which need to be obtained incrementally using large deformation or fluid models [29], [30] and hence are computationally expensive. Accurate alignment of the cortical surface as part of a volumetric registration procedure remains a challenging task mainly due to the complex folding pattern variability of the cortex.

In this paper, we (i) propose a novel landmark based surface matching technique based on elastic energy minimization in the intrinsic geometry of the cortex, (ii) propose a new method based on harmonic mappings for extending the surface matching to the entire cortical volume, and (iii) present a modified intensity alignment based on [31] to compute the final map. The resulting method, comprising the three steps outlined above, gives an inverse consistent map which is capable of aligning both subcortical and sulcal features.

II. PROBLEM STATEMENT AND FORMULATION

Here we address the following problem: produce a one-to-one mapping between two brain volumes such that subcortical structures and sulcal landmarks are aligned and that there is also a one-to-one correspondence between the cortical surfaces of the two volumes. Equivalently, given 3D manifolds M and N representing the two brain volumes, with boundaries ∂M and ∂N representing their respective cortical surfaces, we want to find a map from M to N such that ∂M, the surface of M, maps to ∂N, the surface of N, and the intensities of the images in the interior of M and N are matched. In addition the map must satisfy a sulcal matching constraint so that labelled sulci on the surface ∂M map onto homologous sulci on ∂N. The boundaries, ∂M and ∂N, are assumed to have a spherical topology.

We solve the mapping problem in three steps:

1) Surface matching, which computes a map between ∂M and ∂N, the cortical surfaces of the two brains. The mapping is based on minimization of an elastic strain energy subject to the constraint that a set of interactively labelled sulci are aligned, as described in Section III.

2) Extrapolation of the surface map to the entire cortical volume such that the cortical surfaces remain aligned. This is done by computing a harmonic map between M and N subject to a surface matching constraint. As we describe in Section IV, an intermediate spherical representation is used to facilitate enforcement of


Fig. 1. Cortical surface alignment after using AIR software for intensity based volumetric alignment with a 168 parameter 5^{th} order polynomial. Note that although the overall morphology is similar between the brains, the two cortical surfaces do not align well.

this constraint. We note also that while the sulci are constrained to remain in correspondence, the cortical surfaces can flow with respect to each other when computing the volume harmonic map provided we retain the one-to-one mapping between ∂M and ∂N.

3) Refinement of the harmonic map on the interiors of M and N to improve intensity alignment of subcortical structures. For this step we use an inverse consistent linear elastic registration method as described in Section V.

The mapping in Step 2 requires large scale deformation to ensure that ∂M and ∂N are aligned. Linear elastic or thin-plate spline registration based on landmarks cannot be used for this purpose [32]. Harmonic maps on the other hand are suitable since they are bijective provided that the boundary (the cortical surface in this case) is mapped bijectively. Conversely, the final step requires a more local refinement of the mapping to align subcortical structures so that use of linear elastic methods is appropriate.

III. SURFACE REGISTRATION

We describe a method that sets up one-to-one correspondence between the surface ∂M of brain M and the surface ∂N of the brain N, using labeled sulcal curves as constraints. Our earlier methods for surface registration were computed in two stages: (i) for each subject, parameterize the surface of each cortical hemisphere to either a unit square or disk, and (ii) find a vector field with respect to this parameterization that aligns sulcal landmarks between subjects. Registration can use linear elastic [33] or thin-plate bending energy [16] for regularizing the displacement field and covariant derivatives to make the alignment independent of parameterization. However, in order to solve the resulting variational minimization problem, numerical derivatives must be computed by resampling the brain on a uniform grid with respect to the parameterization. In addition to the computational cost of resampling and interpolation, this step results in a loss of resolution since the regular or semi-regular grid in flat space is not necessarily optimal for representing the brain in 3D space. In our new surface registration approach, we incorporate sulcal landmark alignment directly in our parameterization method and thus avoid the resampling and reparameterization step completely. This approach also has the advantage that the computation cost is relatively small and that the resulting alignment is inverse

consistent [30] as will become clear from the symmetry of the cost function defined below.

In order to generate such a parameterization with prealigned landmarks, we model the cortical surface as an elastic sheet and solve the associated linear elastic equilibrium equation using the Finite Element Method (FEM). We choose the more general elastic model over a surface based harmonic mapping method [34], [35], [36] because we found that the surface based harmonic mappings do not remain bijective when multiple sulcal landmark constraints are imposed on the interior of the flat parameter space. However, for the elastic model, we have found improved bijective behavior with appropriate choices of model parameters λ and μ. The reason for this situation, intuitively, is that relative to the power of the Laplacian alone, the Cauchy-Navier elasticity operator provides additional control over the gradient of the divergence of the surface vector field, and this indirectly controls the Jacobian of the mapping, constraining it from taking on extreme values and thereby violating the smoothness assumption. If the smoothness were a problem, the flow could be discretized in time and integrated to prevent singularities, as described by Christensen et al. [8].

A. Mathematical Formulation

We map the surfaces of each cortical hemisphere of a pair of brains M and N to the unit square such that in the flat map the manually delineated sulcal landmarks align in the square parameter space. The resulting parameterization then defines a correspondence between the two cortices. We describe the process for identifying the surfaces and delineating the sulci in Section IIIC. For now it is sufficient to assume two surfaces with spherical topology on which are traced two homologous sets of continuous, non-intersecting sulcal curves. The corpus callosum maps to the boundary of the square with the boundary condition assuring continuity between the two hemispheres of each brain when subsequently mapping these flat maps to the surface of a sphere.

Let $\phi = [\phi_1, \phi_2]^T$ be the two coordinates assigned to every point on a given cortical surface such that the coordinates ϕ satisfy the linear elastic equilibrium equation with Dirichlet boundary conditions on the boundary of each cortical hemisphere, represented by the corpus callosum. We constrain the corpus callosum to lie on the boundary of the unit square mapped as a uniform speed curve.

We solve the linear elastic equilibrium equation in the geometry of the cortical surface using the form:

$$\mu \Delta \phi + (\mu + \lambda) \nabla (\nabla \cdot \phi) = 0, \qquad (1)$$

where μ and λ are Lamé's constants which model the elastic material response to linear strain and shear respectively [37]. The operators Δ and ∇ represent the Laplace-Beltrami and covariant gradient operators, respectively, with respect to the surface geometry. The solution of this equation can be obtained variationally by minimizing the integral on the cortical surface [38]:

$$E(\phi) = \int_S \frac{\mu}{2} \operatorname{Tr} \left(((D\phi)^T + D\phi)^2 \right) + \lambda \operatorname{Tr} (D\phi)^2 dS, \qquad (2)$$

where $D\phi$ is the covariant derivative of the coordinate vector field ϕ. The integral $E(\phi)$ is the total *strain energy*. Though the

elastic equilibrium equation models only small deformations, we have found that in routine practice, our method is able to produce a flat map of the cortex by using the parameters $\mu = 100$ and $\lambda = 1$. Intuitively, they control the 'stretching' and 'bending' properties of the desired map.

Let ϕ_M and ϕ_N denote the 2D coordinates to be assigned to corresponding hemispheres of M and N respectively. Then we define the cost function $C(\phi_M, \phi_N)$ as

$$C(\phi_M, \phi_N) = E(\phi_M) + E(\phi_N) + \rho \sum_{k=1}^{K} (\phi_M(x_k) - \phi_N(y_k))^2, \qquad (3)$$

where $\phi_M(x_k)$ and $\phi_N(y_k)$ denote the coordinates assigned to the set of K sulcal landmarks $x_k \in M$, $y_k \in N$ and ρ is a penalty parameter. Here sulcal matching is represented by a set of point constraints obtained by sampling each curve at a fixed number of points taken along its length. Note that we do not constrain the locations of the sulci in the flat map but simply constrain homologous landmarks in the two maps to lie at the same coordinates. The method is easily modified for subject to template matching by fixing the parameters of the sulcal constraints in the template.

B. Finite Element Formulation

To minimize (3) on a tessellated surface we use an FEM to discretize the strain energy $E(\phi)$. Since the strain energy at a point is independent of coordinate system, it is justifiable to compute it locally at each vertex point by assigning a local coordinate system (x, y) to its neighborhood.

For each triangle the covariant derivative $D\phi$ in the local coordinates x, y becomes the Jacobian matrix:

$$D\phi = \begin{pmatrix} \frac{\partial \phi_1}{\partial x} & \frac{\partial \phi_1}{\partial y} \\ \frac{\partial \phi_2}{\partial x} & \frac{\partial \phi_2}{\partial y} \end{pmatrix}. \qquad (4)$$

We compute the strain energy $E_i(\phi)$ for the i^{th} triangle Δ_i using (2) as:

$$E_i(\phi) = \int_{\Delta_i} (2\mu + \lambda) \left((\frac{\partial \phi_1}{\partial x})^2 + (\frac{\partial \phi_2}{\partial y})^2 \right) \qquad (5)$$

$$+ 2(\mu + \lambda) \left(\frac{\partial \phi_1}{\partial y} \right) \left(\frac{\partial \phi_2}{\partial x} \right) + \mu \left((\frac{\partial \phi_1}{\partial y})^2 + (\frac{\partial \phi_2}{\partial x})^2 \right) dS.$$

We now describe the FEM discretization of the partial derivatives with respect to the local coordinates. Let α be any piecewise linear real-valued scalar function defined over the surface, and let x, y denote local coordinates for triangle i. Also denote the local coordinates of the three vertices as $(x_1, y_1), (x_2, y_2)$ and (x_3, y_3) respectively. Since α is linear on the i^{th} triangle, we can write,

$$\alpha(x, y) = a_0^i + a_1^i x + a_2^i y. \qquad (6)$$

Writing this expression at three vertices of the triangle i in matrix form,

$$\underbrace{\begin{pmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ 1 & x_3 & y_3 \end{pmatrix}}_{E^i} \begin{pmatrix} a_0^i \\ a_1^i \\ a_2^i \end{pmatrix} = \begin{pmatrix} \alpha(x_1, y_1) \\ \alpha(x_2, y_2) \\ \alpha(x_3, y_3) \end{pmatrix}. \qquad (7)$$

The coefficients a_0^i, a_1^i and a_2^i can be obtained by inverting the matrix E^i. From (6) and by inverting the matrix in (7), we obtain

$$\begin{pmatrix} \frac{\partial \alpha}{\partial x} \\ \frac{\partial \alpha}{\partial y} \end{pmatrix} = \begin{pmatrix} a_1^i \\ a_2^i \end{pmatrix} \qquad (8)$$

$$= \frac{1}{|E^i|} \begin{pmatrix} y_2 - y_1 & y_3 - y_1 & y_1 - y_2 \\ x_1 - x_2 & x_1 - x_3 & x_2 - x_1 \end{pmatrix} \begin{pmatrix} \alpha(x_1, y_1) \\ \alpha(x_2, y_2) \\ \alpha(x_3, y_3) \end{pmatrix}. \qquad (9)$$
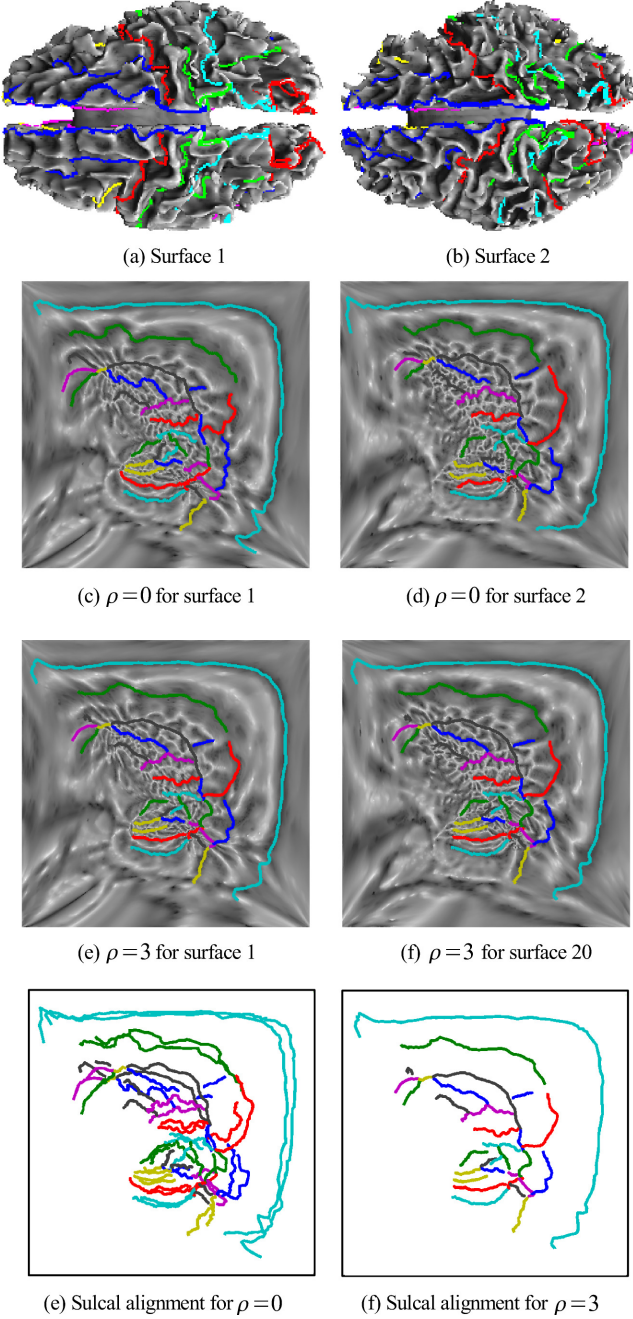
Denote the discretization of $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ at triangle i by D_x^i and D_y^i respectively. Also note that $|E^i| = 2A_i$ where A_i is area of the i^{th} triangle. Then we have:

$$D_x^i = \frac{1}{2A_i} \begin{pmatrix} y_2 - y_1 & y_3 - y_1 & y_1 - y_2 \end{pmatrix} \quad (10)$$

$$D_y^i = \frac{1}{2A_i} \begin{pmatrix} x_1 - x_2 & x_1 - x_3 & x_2 - x_1 \end{pmatrix}. \quad (11)$$

Substituting these in (5) and (3), we have

$$E(\phi) = \sum_i \frac{1}{4A_i} \begin{pmatrix} \phi_1^i \phi_2^i \end{pmatrix} M \begin{pmatrix} \phi_1^i \\ \phi_2^i \end{pmatrix} \quad (12)$$

$$= \sum_i \left\| \frac{1}{2\sqrt{A_i}} \begin{pmatrix} \sqrt{\lambda}D_x^i & \sqrt{\lambda}D_y^i \\ \sqrt{\mu}D_y^i & \sqrt{\mu}D_x^i \\ \sqrt{2\mu}D_x^i & 0 \\ 0 & \sqrt{2\mu}D_y^i \end{pmatrix} \phi^i \right\|^2, \quad (13)$$

where M is given by

$$M = \begin{pmatrix} (\lambda + 2\mu)D_x^{it} D_x^i + \mu D_y^{it} D_y & \lambda D_x^{it} D_y + \mu D_y^{it} D_x^i \\ \lambda D_y^{it} D_x^i + \mu D_x^{it} D_y^i & (\lambda + 2\mu)D_y^{it} D_y^i + \mu D_x^{it} D_x^i \end{pmatrix}. \quad (14)$$

This method is used to discretize both $E(\phi_M)$ and $E(\phi_N)$. The resulting cost function (3) is then minimized using a conjugate gradient method as described in the next section. We note from (13) and (3) that the cost function is quadratic. We minimize (3) with respect to both ϕ_M and ϕ_N, with the corpus callosum fixed at the boundary of the unit square, to compute the sulcally coregistered flat maps for both brains simultaneously.

C. Implementation

We assume as input two T1-weighted MR volumes. Cortical surfaces were extracted from volume images using the BrainSuite software [39]. BrainSuite includes a six stage cortical modeling sequence. First the brain is extracted from the surrounding skull and scalp tissues using a combination of edge detection and mathematical morphology. Next the intensities of the MRI are corrected for shading artifacts. Each voxel in the corrected image is labeled according to tissue type using a statistical classifier. A standard atlas with associated structure labels is aligned to the subject volume, providing a label for cerebellum, cerebrum, brainstem, and subcortical regions. These labels are combined with the tissue classification to automatically identify the cerebral white matter, to fill the ventricular spaces, and to remove the brainstem and cerebellum. This produces a volume whose boundary surface represents the outer white-matter surface of the cerebral cortex. It is likely that the tessellation of this volume will produce surfaces with topological handles. Prior to tessellation, these handles are identified and removed from the binary volume automatically using a graph based approach [40]. A tessellated isosurface of the resulting mask is then extracted to produce a genus zero surface based on the registered atlas labels that is subsequently split into two cortical hemispheres.

We then use BrainSuite to interactively label 23 major sulci on each cortical hemisphere according to a sulcal labeling protocol with established intra- and inter-rater reliability [41]. This protocol specifies that sulci do not intersect and that individual sulci are continuous curves that are not interrupted. If interruptions are present, the human raters specify the path across any interrupting gyri. In cases where a full set of sulci cannot be defined, a subset can be used without requiring any changes in the algorithm described here.

We note that the procedure implemented in BrainSuite to find the cortical surface uses the inner grey/white boundary of cortex as the surface. Consequently the images shown here do not include cortical grey matter but are restricted to white matter, ventricles and subcortical grey matter. However, the method can be used with any of the approaches for cortical segmentation that produce a genus-zero representation of the cortical surface, e.g., [42], [13].

(a) Surface 1 (b) Surface 2

(c) $\rho = 0$ for surface 1 (d) $\rho = 0$ for surface 2

(e) $\rho = 3$ for surface 1 (f) $\rho = 3$ for surface 20

(e) Sulcal alignment for $\rho = 0$ (f) Sulcal alignment for $\rho = 3$

Fig. 2. (a),(b) Two cortical surfaces with labeled sulci as colored curves; (c),(d) flat maps of a single hemisphere for each brains without the sulcal alignment constraint; (e),(f) flat maps with sulcal alignment; (g),(h) overlay of sulcal curves on the flat maps, without and with sulcal alignment.


Fig. 3. Result of mapping of sulcal landmarks from 5 subjects to a single brain using the linear elastic mapping described here (left) without and (right) with the sulcal alignment constraint.

Automated extraction of a topologically spherical surface from MRI using BrainSuite takes 4-5 min. and interactive sulcal labeling takes approximately 1-2 hours per brain.

As described above, $E(\phi_M)$ and $E(\phi_N)$ are discretized using (14) and (13). The resulting quadratic cost function (3) is minimized using a preconditioned conjugate gradient method with Jacobi preconditioner. In practice the minimization algorithm converges in approximately 500 iterations, requiring 2-3 mins on a desktop computer for tessellations with on the order of 10^5 vertices.

We show an example of this flattening and alignment procedure in Fig. 2. Shown in (a) and (b) are the sulci traced on the white matter surface for two brains. In (c) and (d) we show the flat maps of one hemisphere in each brain computed with the penalty parameter $\rho = 0$ in (3), which corresponds to the case where the two hemispheres are flattened independently. Similarly, (e) and (f) show the flat maps for the case $\rho = 3$, which forces approximate alignment of the sulci in the flat maps. Note that we only enforce the matching approximately via the penalty term in (3). Maps with and without sulcal alignment appear quite similar, so we show in (g) the two sets of curves in the flat space when they are not constrained to align ($\rho = 0$), and in (h) when they are constrained ($\rho = 3$). These figures illustrate the procedure and demonstrate alignment of the curves in the flat space. In Fig. 3 we show the results of applying the mapping procedure separately from five subjects to a single template brain. Each pair of flattenings produces a different parameterization of each subject and of the template. In each case, there is a one-to-one map between the surfaces so that the sulcal curves of each subject can be mapped back onto the template surface. Shown in the figure are the set of sulcal curves mapped onto the template for cases where the alignment constraint is ($\rho = 0$) and is not ($\rho = 3$) applied.

IV. HARMONIC MAPPING

The surface registration procedure described in Section III sets up a point to point correspondence between the two cortical surfaces, which represent the boundary of the two cerebral volumes. Extrapolating this correspondence from the boundary surface to the entire cerebral volume in a one-to-one manner is challenging due to the convoluted nature of the cortex. In fact, most of the linear models such as linear elastic or thin-plate splines become non-bijective under relatively mild landmark matching constraints [32]. 3D harmonic maps are attractive for this purpose due to their tendency to be bijective if the boundary (cortical surface) is mapped bijectively, which is the case here. In this section we describe a framework for computing a harmonic mapping between two 3D volumes as well as the computational approach used for implementation. Details of harmonic maps and their properties can be found in [43].

Let $u : M \rightarrow N$ be a C^∞ map from a 3 dimensional Riemannian manifold (M, g) to a 3 dimensional Riemannian manifold (N, h) where g and h are Riemannian metrics for M and N respectively. A Riemannian metric defines an inner product at every point in the manifold and thus helps in defining the notion of distance on the manifold [43]. Let $\{g_{ij}; i, j \in \{1, 2, 3\}\}$ denote components of the Riemannian metric tensor g and $\{h_{\alpha\beta}; \alpha, \beta \in \{1, 2, 3\}\}$ denote the components of the Riemannian metric tensor h. The inverse of the metric $g = \{g_{ij}\}$ is denoted by $\{g^{ij}\}$. Let (x^1, x^2, x^3) and (u^1, u^2, u^3) be local coordinates for x and $u(x)$ respectively. Let Du denote the derivative (generalized Jacobian) of the map. The *energy density function* $e(u)$ of map u is defined to be the norm of Du [44] and is given by

$$e(u)(x) = \frac{1}{2}|Du|^2 \tag{15}$$

$$= \frac{1}{2}\sum_{i,j=1}^{3}\sum_{\alpha,\beta=1}^{3} g^{ij}(x)h_{\alpha\beta}(u(x))\frac{\partial u^\alpha(x)}{\partial x^i}\frac{\partial u^\beta(x)}{\partial x^j}, \tag{16}$$

which can be thought of as the rate of expansion of the map u in orthogonal directions, at point $x \in M$ [44],. The *mapping energy* is

defined as

$$E(u) = \int_M e(u)(x)d\mu_g. \tag{17}$$

The mapping energy in coordinate form [44], is given by

$$E(u) = \frac{1}{2}\int_M \sum_{i,j=1}^{3}\sum_{\alpha,\beta=1}^{3} g^{ij}(x)h_{\alpha\beta}(u(x))\frac{\partial u^\alpha(x)}{\partial x^i}\frac{\partial u^\beta(x)}{\partial x^j}d\mu_g, \tag{18}$$

where the integration is over the manifold M with respect to the intrinsic measure $d\mu_g$ induced by its Riemannian metric g.

A *harmonic map* from (M, g) to (N, h) is defined to be a critical point of the mapping energy $E(u)$. In this sense harmonic maps are the least expanding maps in $C^\infty(M, N)$, the space of all smooth maps from M to N. Therefore, among all possible smooth maps between two manifolds, the harmonic maps have the tendency to avoid overlaps and folds in the map, resulting in a bijective map.

A number of existence, uniqueness, and regularity results have been proven for harmonic maps [45]. Eells and Sampson [46] proved the existence of a harmonic map from any compact Riemannian manifold to a compact Riemannian manifold of non-positive sectional curvature. Hamilton [47] generalized this result to manifolds with boundaries. In medical imaging, harmonic mappings and p-harmonic mappings, their generalized counterparts [48], have been used for various applications such as surface parameterization and registration [34], [49], [35] and image smoothing [36]. Wang et al. [50] describe a method for volumetric mapping of the brain to the unit ball $B(0, 1)$. Here we use harmonic maps to align two brain volumes so that both the brain volumes and cortical surfaces are aligned.

When computing the harmonic maps we could fix the correspondence between the two surfaces using the method from Section III and map only the interior of the two volumes. This would result in a suboptimal mapping with respect to the 3D mapping energy. To overcome this limitation, we instead allow the surface M to flow within the surface of N when computing the map. The only constraints placed on the surfaces are that the maps are aligned at the set of user defined sulcal landmarks and that the boundary ∂M maps onto ∂N. This less restrictive surface mapping constraint cannot be formulated directly in the ambient Euclidean 3D space since there is no analytical expression for the surfaces. It could be accomplished without parameterizing the surface using a level set approach [36], [51]. Here we use an intermediate representation for the manifolds which allows us to enforce the boundary matching constraint while allowing one boundary to flow within the other. We achieve this by first mapping to the unit ball as described below. This mapping to the unit ball results in a non-Euclidean representation of N thus requiring the use of the Riemannian metric in computing the harmonic map.

A. Mathematical Formulation

We find the map v of the 3D brain manifold N to the 3D unit ball $B(0, 1)$ [50] using the method described in Sec. IV-C. Let $v = (v^1, v^2, v^3)$ denote the three coordinates of the map v. This map is bijective and therefore we can treat the unit ball $B(0, 1)$ as an alternative representation (N, h) of the manifold N, with associated metric h, that has the advantage over the Euclidean space (N, I) that the cortical surface lies on the surface of the sphere (here I represents the identity metric for the Euclidean space); h is the metric induced by the map v. With this alternative representation of N, the components of its metric $h_{\alpha\beta}$ at point $x = (x^1, x^2, x^3)$ are given by

$$h_{\alpha\beta} = \sum_{i=1}^{3}\frac{\partial x^i}{\partial v^\alpha}\frac{\partial x^i}{\partial v^\beta}. \tag{19}$$

Now instead of needing to directly compute the harmonic map $u : (M, I) \rightarrow (N, I)$, we instead find the harmonic map $\tilde{u} : (M, I) \rightarrow (N, h) \approx B(0, 1)$ subject to the constraint that the cortical surface ∂M maps to the spherical boundary of the unit ball, as illustrated in Fig. 4.

Fig. 4. Illustration of our general framework for surface-constrained volume registration. We first compute the map v from brain manifold (N, I) to the unit ball to form manifold (N, h). We then compute a map $\tilde{u}$ from brain (M, I) to (N, h). The final harmonic map from (M, I) to (N, I) is then given by $u = v^{-1} \circ \tilde{u}$.


Fig. 6. Illustration of the deformation induced with respect to the Euclidean coordinates by mapping to the unit ball. Shown are iso-surfaces corresponding to the Euclidean coordinates for different radii in the unit ball. Distortions become increasingly pronounced towards the outer edge of the sphere where the entire convoluted cortical surface is mapped to the surface of the ball.

Since M remains in the Euclidean space, its metric is I, so $g^{ij}(x)$ is the identity operator and the harmonic mapping problem (18) becomes:

$$\tilde{u} = \arg\min_{\gamma} \int_M \sum_{i=1}^{3} \sum_{\alpha,\beta=1}^{3} h_{\alpha\beta}(\gamma(x)) \left(\frac{\partial \gamma^{\alpha}(x)}{\partial x^i} \right) \left(\frac{\partial \gamma^{\beta}(x)}{\partial x^i} \right) d\mu_g, \tag{20}$$

subject to $\|\tilde{u}(x)\|^2 = 1$ for $x \in \partial M$, the surface of M. Note that this constraint allows the surface map to flow within the spherical boundary. We also want to constrain the maps so that predefined sulcal landmarks are aligned. To achieve this we impose the additional constraints that $\tilde{u}(c) = u_c$ for $c \in M_c$ where M_c are the set of sulcal landmark points in M and u_c are the locations of the homologous landmarks in (N, h). Having obtained $\tilde{u}$ by minimizing the integral in (20), the final harmonic mapping from $u : (M, I) \to (N, I)$ can then be computed as $u = v^{-1} \circ \tilde{u}$ as illustrated in Fig. 4.

B. Initialization Procedure

Because the minimization problem (20) is nonquadratic, it is important to have a good initial estimate of the map $\tilde{u}$ in order to achieve convergence in reasonable time. We therefore generate an initial estimate $\tilde{u}_0$ of $\tilde{u}$ by computing a map of the second manifold (M, I) to the unit ball, just as we do for the first manifold (N, I) (Fig. 4). Thus our initialization generates a bijective initial map, which is not necessarily harmonic. The procedure is illustrated in Fig. 5.

The initialization consists of the following steps. We first compute flat maps to the unit square for each hemisphere of the two brains with aligned sulci as described in Section III. A stereographic projection then maps the two hemispheres of each brain to the unit sphere so that the corpus callosum that forms the boundary of the unit squares maps to the equator. Using these surface maps as constraints, we then map N and M to the unit ball to provide, respectively, the unit ball manifold (N, h) and an initial estimate $\tilde{u}_0$ of the desired map $\tilde{u}$ from (M, I) to (N, h). The initial map obtained in this manner is smooth and bijective. With this initialization, the 3D harmonic map is computed by minimizing (20) to obtain the final harmonic mapping from M to N.

C. Mapping to the Unit Ball $B(0, 1)$

In the special case when (M, g) and (N, h) are 3D Euclidean manifolds, then $h_{\alpha\beta} = \delta_\alpha^\beta$, $g_{ij} = g^{ij} = \delta_i^j$, the Kronecker delta, or identity tensor, for $\alpha, \beta, i, j \in 1, 2, 3$, and the mapping energy simplifies to

$$E(u) = \int_M |\nabla u|^2 dV, \tag{21}$$

where ∇ is the usual gradient operator in $3D$ Euclidean space and dV is the volume integral [50]. In order to map the given cortical brain volume M to the unit ball, this energy is minimized subject to the constraint that the surface of M maps to the surface of the unit ball using the point-to-point correspondence defined by the flat mapping obtained as described in Section III. This is computed by numerical integration over the voxel lattice using finite differences to approximate the gradients in (21). The resulting function is minimized using a preconditioned conjugate gradient method. The process of mapping to the unit ball is illustrated in Fig. 6 where we show iso-surfaces in brain coordinates corresponding to different radii, r, within the unit ball. At $r = 1$ we are at the outer surface of the brain and see the full cortical surface. As r is reduced we see successively less distortion since the harmonic map is driven entirely by the surface constraint.

D. Harmonic Mapping Between the Two Brains

The mapping to the unit ball is applied to both brain volumes M and N. The mapping of the Euclidean coordinates in M to the unit ball provides the initial estimate $\tilde{u}_0$ of the harmonic map $\tilde{u}$. We then refine this map by minimizing the harmonic energy in (20) from (M, I) to (N, h), the unit ball representation of N. Again, the problem is solved using numerical integration and finite difference operators, in this case accounting for the metric h according to (20) when computing these derivatives. In this mapping, the locations of the sulci in M are constrained using their initial mappings $\tilde{u}_0$ computed when flattening and matching the cortical surfaces. Other points on the surface are allowed to move freely to minimize the harmonic energy, subject to the constraint that all points on the surface map to $\|\tilde{u}\|^2 = 1$, which is achieved by adding a penalty function $\varrho(\|\tilde{u}\|^2 = 1)^2$ to the discretized form of (20). The penalty parameter ϱ is set to 1000 for the purpose of numerical implementation.

E. Implementation

We first describe a numerical method for computation of the metric $h_{ij}(x)$ and then outline the harmonic mapping method.

1) Computation of Metric: The metric $h_{ij}(x), x \in N$ is associated with the unit ball coordinates $B(0, 1)$ given to N by the map $v = (v^1, v^2, v^3)$ (Fig. 4). It is given by $h_{\alpha\beta}(p) = \sum_{i=1}^{3} \frac{\partial x^i}{\partial v^\alpha} \frac{\partial x^i}{\partial v^\beta}$ with $\alpha, \beta \in \{1, 2, 3\}$ at $x = (x^1, x^2, x^3)$. Note that although $x \in N$ is in the regular grid, $v(x) \in B(0, 1)$ is not necessarily so, and hence computation of partial derivatives with respect to v directly
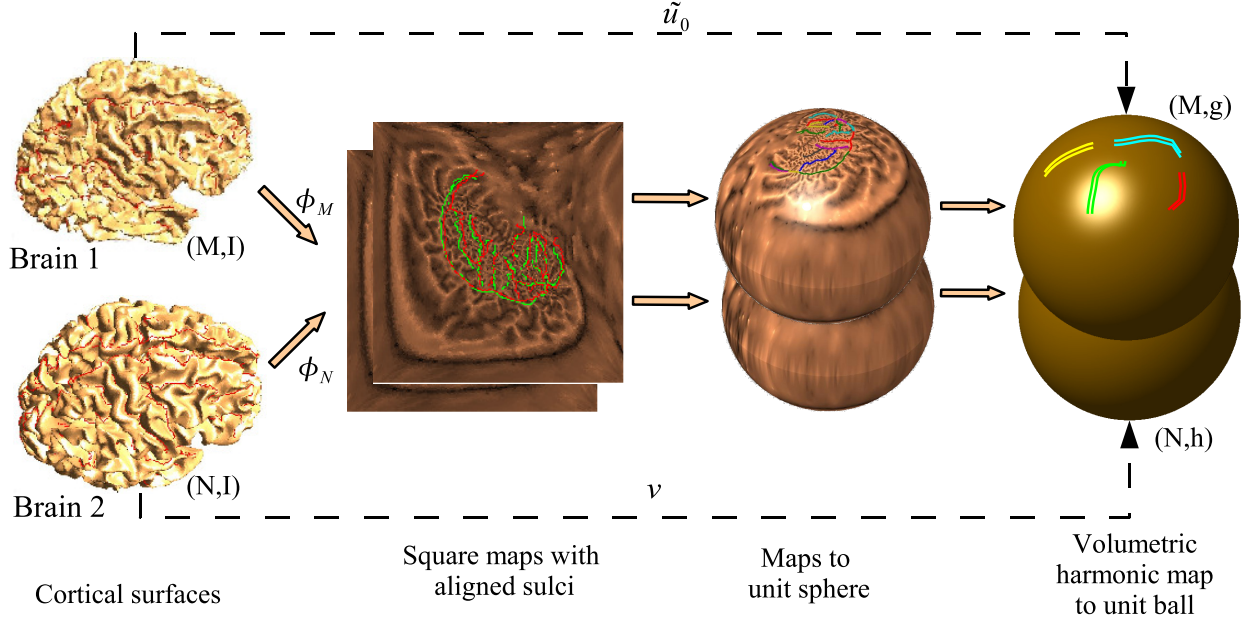
Fig. 5. Initialization for harmonic mapping from M to N. First we generate flat square maps of the two brains, one for each hemisphere, with pre-aligned sulci. The squares corresponding to each hemispheres are mapped to a disk and the disks are projected onto the unit sphere. We then generate a volumetric maps from each of the brains to the unit ball. Since all these maps are bijective, the resulting map results in a bijective point correspondence between the two brains. However, this correspondence is not optimal with respect to the harmonic energy of maps from the first brain to the second, but is used as an initialization for minimization of (20).

is difficult. In order to compute $\frac{\partial v^\alpha}{\partial x^i}$, first compute $\frac{\partial v^\gamma}{\partial x^j}$ using finite differences and then use the chain rule identity

$$\sum_{\gamma=1}^{3} \frac{\partial x^i}{\partial v^\gamma} \frac{\partial v^\gamma}{\partial x^j} = \frac{\partial x^i}{\partial x^j} = \delta_j^i \quad (22)$$

to solve for $\frac{\partial v^\gamma}{\partial x^j}$. The metric h_{ij} is computed by substituting these partial derivatives in the above equation.

2) Harmonic Mapping: The harmonic mapping procedure can now be summarized as follows:

1) Align the surfaces of both the brains M and N using the procedure described in Sec. III.
2) Map the unit squares to unit disks by the transformation $(x,y) \rightarrow (\frac{x}{\sqrt{x^2+y^2}}, \frac{y}{\sqrt{x^2+y^2}})$ and then project them onto two hemispheres using $(x,y) \rightarrow (x,y,\pm\sqrt{x^2+y^2})$.
3) Using this mapping of the cortical surface to the unit sphere as the boundary condition, generate volumetric harmonic maps of M and N to the unit ball $B(0,1)$ as described in Sec. IV-C.
4) Compute the metric h associated with the unit ball $B(0,1)$ coordinates of N as described above.
5) Minimize (20) holding the matched sulci fixed, and letting the cortical surface ∂M slide along boundary of the unit ball. This is done by minimizing (20) with the constraint that $\|\tilde{u}(x)\|^2 = 1$ for $x \in \partial M$ and $\tilde{u}(c) = \tilde{u}_0(c)$ for $c \in M_c$ where $M_c \subset M$ denotes the set of sulcal points on M. The partial derivatives in (20) are discretized by finite differences and the minimization is done by gradient descent.
6) Compute the displacement vector field $u(x) - x$ where $u = v^{-1} \circ \tilde{u}$ and apply this to map brain volume M to N. Trilinear interpolation is used for this deformation.

V. VOLUMETRIC INTENSITY REGISTRATION

The surface constrained harmonic mapping procedure described above produces a bijective mapping between the two brain volumes.

However, it uses only surface shape and sulcal labels and does not use the MRI intensity values to compute the map. The result is a large scale deformation that aligns surface features but will benefit from an intensity-based refinement aimed at aligning subcortical features. In order to do this refinement and also make the final map inverse consistent, we use linear elastic inverse consistent registration based on Christensen's approach [31] with the modifications described below to ensure that the entire mapping process, rather than just this last step, is inverse consistent. Alternatively, inverse consistency can be achieved by symmetric intensity based registration method suggested by Tagare et al. [52]

A. Formulation

The surface constrained volumetric harmonic mapping procedure described above can be used to generate two maps $u^M : M \rightarrow N$ and $u^N : N \rightarrow M$, each harmonic, but not necessarily inverses of each other. The corresponding displacement fields for these maps can be expressed as $d_u^M(x) = u^M(x) - x, x \in M$ and $d_u^N(x) = u^N(x) - x, x \in N$. Note that both of these displacement fields accurately align the two surfaces and corresponding sulci, and are one-to-one. These deformations are used to initialize the volumetric inverse consistent intensity registration procedure that we now describe.

Let $f_M(x), x \in M$ denote intensity at point $x \in M$ and $f_N(x), x \in N$ denote intensity at point $x \in N$. The situation can be summarized as follows and is illustrated in Fig. 7: We have maps $u^M : M \rightarrow N$ $u^N : N \rightarrow M$ that change the shapes of domains M and N to match their respective targets N and M. In order to align the intensities, we seek refinement maps $w^M : M \rightarrow M$ and $w^N : N \rightarrow N$ such that the mapped intensity value $f_M \circ w^M \circ u^N$ matches f_N (or equivalently $f_M \circ w^M$ matches $f_N \circ (u^N)^{-1}$), and $f_N \circ w^N \circ u^M$ matches f_M (or $f_N \circ w^N$ matches $f_M \circ (u^M)^{-1}$). For inverse consistency, we need $w^N \approx (u^M \circ w^M \circ u^N)^{-1}$ and $w^M \approx (u^N \circ w^N \circ u^M)^{-1}$. Let d_w^M, d_w^N denote the displacement fields corresponding to w^M, w^N and let $\tilde{d}_w^M, \tilde{d}_w^N$ denote the displacement fields for $(u^N \circ w^N \circ u^M)^{-1}, (u^M \circ w^M \circ u^N)^{-1}$.

Fig. 7. Schematic of the intensity alignment procedure. Once harmonic maps u^M and u^N are computed, we refine these with intensity driven warps w^M and w^N while imposing constraints so that the final deformations are inverse consistent.


Fig. 8. Illustration of the effects of the two stages of volumetric matching is shown by applying the deformations to a regular mesh representing one slice. Since the deformation is in 3D, the third in-paper value is represented by color. (a) Regular mesh representing one slice in the subject; (b) the regular mesh warped by the harmonic mapping which matches the subject cortical surface to the template cortical surface. Note that deformation is largest near the surface since the harmonic map is constrained only by the cortical surface; (c) Regular mesh representing one slice in the harmonically warped subject; (d) the intensity-based refinement now refines the deformation of the template to improve the match between subcortical structures. In this case the deformation is constrained to zero at the boundary and are confined to the interior of the volume.

The inverse consistency similarity cost function $C(d_w^M, d_w^N)$, can now be defined as the sum of three terms:

$$C(d_w^M, d_w^N) = C_{REG}(d_w^M, d_w^N) + \alpha C_{SIM}(d_w^M, d_w^N)$$
$$+ \beta C_{ICC}(d_w^M, d_w^N) \text{ subject to } d_w^N(u^M(x)) = 0, x \in \partial M \text{ and}$$
$$d_w^M(u^N(x)) = 0, x \in \partial N, \quad (23)$$

where the boundary constraints ensure that the cortices remain aligned after registration and the three constituent terms are defined as follows:

$$C_{REG}(d_w^M, d_w^N) = \|L_M d_w^M\|^2 + \|L_N d_w^N\|^2$$
$$C_{SIM}(d_w^M, d_w^N) = \|f_M(x + d_w^M(x)) - f_N(u^{N^{-1}}(x))\|^2 +$$
$$\|f_N(x + d_w^N(x)) - f_M(u^{M^{-1}}(x))\|^2$$
$$\approx \|f_M(x) + \nabla_M f_M(x) \cdot d_w^M(x) - f_N(u^{N^{-1}}(x))\|^2 +$$
$$\|f_N(x) + \nabla_N f_N(x) \cdot d_w^N(x) - f_M(u^{M^{-1}}(x))\|^2 +$$
$$C_{ICC}(d_w^M, d_w^N) = \|d_w^M(x) - \tilde{d}_w^M(x)\|^2 + \|d_w^N(x) - \tilde{d}_w^N(x)\|^2$$
$$(24)$$

The first term is the regularizer where $L_M = \alpha\nabla_M^2 + \beta\nabla_M(\nabla_M\cdot) + \gamma$ and $L_N = \alpha\nabla_N^2 + \beta\nabla_N(\nabla_N\cdot) + \gamma$ denote the Cauchy Navier elasticity operators in M and N respectively. The parameter values of $\alpha = 10, \beta = 2$ and $\gamma = 0.1$ are used for the purpose of numerical implementation. The second term measures the intensity match between the transformations in both directions and the third term is a measure of deviation from the inverse consistent condition. The similarity cost C_{SIM} is linearized by using Tayor series so that the approximation is quadratic. Note that the inverse consistency cost C_{ICC} and regularizing cost C_{REG} are quadratic. The resulting quadratic cost function C can be efficiently minimized by the conjugate gradient method. We use a preconditioned conjugate gradient method with Jacobi preconditioner for this purpose.

B. Implementation

1. First, the harmonic maps $u^M : M \to N$ and $u^N : N \to M$ are computed using the procedure described in Section IV-D.

2. The inverses of the map $u_M^{-1} : N \to M$ is computed. This is done by interpolating the correspondence $u_M^{-1} : u_M(x) \mapsto x$ from points to the regular voxel grid of N using Matlab's *griddata3* function with linear interpolation. This function implements the method based on Delauney triangulation as described in [53] although it can also be computed using the method described in [31]. $u_N^{-1} : M \to N$ is computed similarly.

3. Set $d_w^M = 0$ and $d_w^N = 0$.

4. Compute the maps $w^N(y) = y + d_w^N(y), y \in N, \tilde{w}^M = (u^N \circ w^N \circ u^M)^{-1}$ and $\tilde{d}_w^M(x) = \tilde{w}^M(x) - x$.

5. Compute the difference term $f_N(x) - f_M(u^{N^{-1}}(x))$.

6. Compute an updated estimate of the displacement field $\hat{d}_w^M$ from (23) using a preconditioned conjugate gradient method.

7. Repeat steps 4-6 with M and N interchanged.

8. Test inverse consistency error C_{ICC} for convergence, otherwise go to Step 4.

This final refinement completes the surface-constrained registration procedure. While there are several steps required to complete the registration, each step can be reduced to either a surface or a volume mapping cast as an energy minimization problem, with constraints, and can be effectively computed using a preconditioned conjugate gradient method. The different effects of the harmonic mapping, producing large scale deformations, and the linear elastic intensity-driven refinement, producing small scale deformations, are illustrated in Fig. 8

VI. RESULTS AND VALIDATION

In order to illustrate the application of our surface constrained registration procedure to T1-weighted MR brain images and validate

its performance, we obtained labeled brain data from the Internet Brain Segmentation Repository (IBSR) dataset at the Center for Morphometric Analysis at Massachusetts General Hospital. This consists of T1-weighted MR images with $1.5mm$ slice thickness as well as expert segmentations of 43 individual structures. The cortical masks were obtained and their topology corrected using the BrainSuite software as described in Sec. III-C. The cortical surfaces were then interactively labelled with 23 sulcal curves on each hemisphere using a standard labeling protocol [41]. Our registration algorithm was applied by performing surface matching, harmonic mapping and volumetric intensity registration as described above. Shown in Fig. 9 are three orthogonal views of a subject before and after alignment to the template image. Note that before alignment the surfaces of the subject and template are clearly different, while after the harmonic mapping the deformed subject surface almost exactly matches the morphology of that of the template. However, since at this point we do not take the image intensities into account, the interior structures do not align well. Following the final intensity-based alignment procedure the subcortical structures of the warped subject show improved agreement with those in the the template. Also shown in Fig. 9b are the labels provided by the IBSR data set before and after mapping.

Our method for evaluating the quality of our registration results is based on the following two desirable features:

1) Alignment of the cortical surface and sulcal landmarks. We expect the sulcal landmarks to be accurately aligned after registration and for the two surfaces to coincide.

2) Alignment of subcortical structures. We also expect the boundaries of subcortical structures (e.g., thalamus, lateral ventricles, corpus callosum) to be well aligned after registration.

To evaluate performance with respect to 1, we used a set of 6 MR volumes on which we labeled 23 sulci in each hemisphere. For comparison we use a 5th order polynomial intensity-driven warp computed using the AIR software [5], [11]. We also compare performance with the HAMMER [28], [54] algorithm. HAMMER is an automated method for volume registration which is able to achieve improved alignment of geometric features by basing the alignment on an attribute vector that includes a set of geometric moment invariants rather than simply the voxel intensities. We note that since our approach uses explicitly labelled sulci we can expect better performance than either AIR or HAMMER in terms of the alignment of these features. However, AIR and HAMMER provide a basis for comparison since they are among the most widely used and best performing algorithms for volumetric registration.

We measured the mean squared distance between pairs of homologous landmarks corresponding to uniform samples along each of the 23 labeled sulci. We repeated this procedure for each of the 30 possible pairwise registrations of two from six brains and computed the average mean squared distance over all registrations. We found that the mean squared misalignment between sulcal landmarks was $11.5mm$ for HAMMER, $11mm$ for AIR and $2.4mm$ for our cortically constrained method. Results for individual sulci are included in Fig. 10. The significantly lower error for our approach is unsurprising since matching of sulci is imposed as a constraint. The reason that the error is not zero is that the constraint is imposed using a penalty function rather than strictly using Lagrange multipliers.

To evaluate performance in terms of subcortical structures we used the manually labeled regions in the IBSR data set. To evaluate accuracy, we computed the Dice coefficients between the template and warped subject for each subcortical structure, where the structure names and boundaries were taken from the IBSR database. The Dice coefficient measures overlap between two sets representing regions S_1 and S_2, and is defined as $\frac{2|S_1 \cap S_2|}{|S_1|+|S_2|}$ where $|\cdot|$ denotes size of the set [55]. Values range from zero for disjoint sets to unity for identical sets. A comparison of the Dice coefficients for some major subcortical organs is shown in Table I, where we show Dice coefficients for our method before and after application of the intensity-based alignment step. This comparison shows similar results for all three methods, with each producing superior results
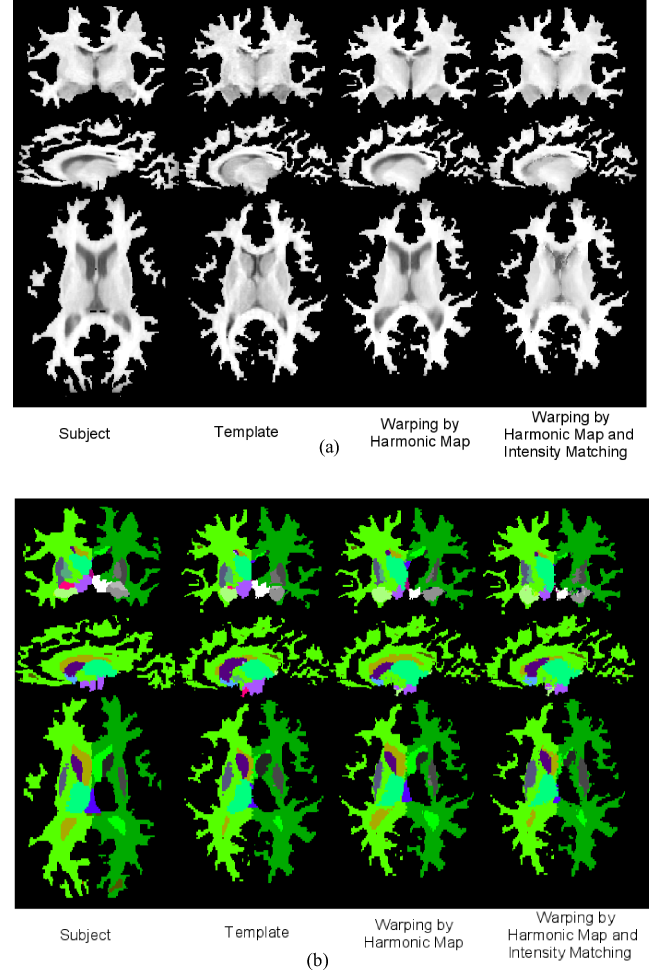

Fig. 9. Examples of surface constrained volumetric registration. (a) Original subject volume, original template, registration of subject to template using surface constrained harmonic mapping, intensity-based refinement of the harmonic map of subject to template is shown. Note that the surface of the warped subject matches to the surface of the template. (b) Anatomical labels of the subject and the template followed by labels of the subject warped by surface constrained harmonic mapping and intensity-based refinement of the harmonic map are shown.


Fig. 10. RMS errors in alignment of different sulci using AIR (5th order), HAMMER and our surface constrained mapping.

TABLE I
COMPARISON OF DICE COEFFICIENTS

Subcortical Structure	AIR	Harmonic	HAMMER	Harmonic with intensity
Left Thalamus	0.6588	0.5294	0.7303	0.5856
Left Caudate	0.4426	0.4336	0.5688	0.5716
Left Putamen	0.4079	0.3497	0.4905	0.5092
Left Hippocampus	0.4676	0.3069	0.3916	0.3930
Right Thalamus	0.6326	0.5018	0.7495	0.6230
Right Caudate	0.3671	0.3572	0.5098	0.5116
Right Putamen	0.3096	0.2358	0.4111	0.4679
Right Hippocampus	0.5391	0.3455	0.1989	0.4342
Avg. Dice coeff. for all structures	0.3021	0.3821	0.3621	0.4019
Std. Dev. of Dice coeff.	0.1937	0.2547	0.2390	0.2671

in some subcortical structures. For example, HAMMER produced superior results in thalamus, while our proposed method produced superior results in hippocampus. Thus the geometric invariants in HAMMER seem to improve performance relative to our intensity based alignment of deeper subcortical structures, while our use of a cortical constraint leads to superior performance with respect to sulcal alignment and structures that are more superficial with respect to the cerebral cortex, such as the hippocampus. This is a preliminary validation and larger scale validation is needed on a larger population with a larger range of brain structures.

VII. CONCLUSION

We have presented a framework for coregistration of brain volume data using harmonic maps. Through the use of an intermediate spherical map, we are able to constrain the surfaces of the two brain volumes to align while enforcing point matching only at a set of hand labeled sulcal curves. Using harmonic maps we are able to compute large scale deformations between brain volumes.

We have also described, as an initialization procedure, a new method for cortical surface parameterization and sulcal alignment in which the two problems are solved in a single step using a finite element method. This method has the properties that it is inverse consistent between the two brains and can be computed directly from a tessellated representation of the surface, rather than requiring resampling using a regular grid with respect to the induced parameterization.

The examples shown here demonstrate the cortical matching properties and the ability to also align subcortical structures. One of the limitations of this evaluation was that cortical grey matter was not included in the registration since the cortical surfaces were generated by BrainSuite [56], which selects the inner grey/white boundary as the cortical surface. However, this is a limitation of the preprocessing step rather than the method itself, and the process can be applied to the full cerebral volume provided that a genus-zero brain volume and sulcal labels are supplied. A second limitation is that the cerebellum and brainstem are not included in the analysis since the volume of interest that is mapped is restricted to the cerebrum, bounded by the outer cortical sheet. We can address this issue in practice by modifying the final intensity-based matching step by first adding the brainstem and cerebellum back to the cerebrum. This would also require extrapolation of the displacement field from the harmonic map outwards to these structures as an initialization of the intensity based warp. Alternatively, the cerebellum could also be explicitly modelled using a surface based approach (see, e.g., Hurdal et al. [12]), and its surface and enclosed volume could be treated in a similar fashion to the cerebrum.

ACKNOWLEDGMENT

REFERENCES

[1] G. E. Christensen, S. C. Joshi, and M. I. Miller, "Volumetric transformation of brain anatomy," *IEEE TMI*, vol. 16, no. 6, December 1997.

[2] P. M. Thompson and A. W. Toga, "A surface-based technique for warping 3-dimensional brain," *IEEE Transactions on Medical Imaging*, vol. 15, no. 4, pp. 1–16, 1996.

[3] J. Talairach and P. Tournoux, *Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging*. NY: Thieme Medical Publishers, New York, 1988.

[4] J. Ashburner and K. Friston, "Spatial normalization," in *Brain Warping*, A. Toga, Ed. Academic Press, 1999, pp. 27–44.

[5] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta, "Automated image registration: I. General methods and intrasubject, intramodality validation," *Journal of Computer Assisted Tomography*, vol. 22, pp. 139–152, 1998.

[6] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration," *Phys. Med. Biol.*, vol. 46, no. 4, pp. R1–R45, March 2001.

[7] G. E. Christensen, R. D. Rabbitt, M. I. Miller, S. C. Joshi, U. Grenander, T. A. Coogan, and D. C. V. Essen, "Topological properties of smooth anatomic maps," in *In 14 Conference on Information Processing in Medical Imaging, France*. Kluwer Academic Publishers, 1995, pp. 101–112.

[8] G. E. Christensen, R. D. Rabbit, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Transactions on Image Processing*, vol. 5, no. 10, pp. 1435–1447, 1996.

[9] J. Glaunés, M. Vaillant, and M. I. Miller, "Landmark matching via large deformation diffeomorphisms on the sphere," *J. Math. Imaging Vis.*, vol. 20, no. 1-2, pp. 179–200, 2004.

[10] B. B. Avants and J. C. Gee, "Shape averaging with diffeomorphic flows for atlas creation," in *ISBI*, 2004.

[11] R. P. Woods, S. T. Grafton, J. D. G. Watson, N. L. Sicotte, and J. C. Mazziotta, "Automated image registration: II. Intersubject validation of linear and nonlinear models," *Journal of Computer Assisted Tomography*, vol. 22, pp. 153–165, 1998.

[12] M. K. Hurdal, K. Stephenson, P. L. Bowers, D. W. L. Sumners, and D. A. Rottenberg, "Coordinate system for conformal cerebellar flat maps," *NeuroImage*, vol. 11, p. S467, 2000.

[13] B. Fischl, M. I. Sereno, R. B. H. Tootell, and A. M. Dale, "High-resolution inter-subject averaging and a coordinate system for the cortical surface," *Human Brain Mapping*, vol. 8, pp. 272–284, 1998.

[14] M. Bakircioglu, U. Grenander, N. Khaneja, and M. I. Miller, "Curve matching on brain surfaces using frenet distances," *Human Brain Mapping*, vol. 6, pp. 329–333, 1998.

[15] P. M. Thompson, R. P. Wood, M. S. Mega, and A. W. Toga, "Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain (invited paper)," *Human Brain Mapping*, vol. 9, no. 2, pp. 81–92, Feb. 2000.

[16] A. A. Joshi, D. W. Shattuck, P. M. Thompson, and R. M. Leahy, "A framework for registration, statistical characterization and classification of cortically constrained functional imaging data," in *Lecture Notes in Computer Science*, vol. 3565, July 2005, pp. 186–196.

[17] Y. Wang, M. C. Chiang, and P. M. Thompson, "Automated surface matching using mutual information applied to Riemann surface structures," in *MICCAI 2005, LNCS 3750*, J. Duncan and G. Gerig, Eds. Springer-Verlag Berlin Heidelberg, 2005, pp. 666–674.

[18] D. Tosun and J. L. Prince, "Cortical surface alignment using geometry driven multispectral optical flow," in *Information Processing in Medical Imaging*, ser. LNCS, vol. 3565, 2005, pp. 480–492.

[19] D. Tosun, M. E. Rettmann, and J. L. Prince, "Mapping techniques for aligning sulci across multiple brains," *Medical Image Analysis*, vol. 8, no. 3, pp. 295–309, 2005.

[20] C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen, "Accurate three-dimensional registration of CT, PET and/or MR images of the brain," *J. Comput. Assist. Tomogr.*, vol. 13, no. 1, pp. 22–26, 1989.

[21] N. Krahnstover and C. Lorenz, "Development of point-based shape representation of arbitrary three-dimensional medical objects suitable for statistical shape modeling," in *Proc. SPIE-Medical Imaging 1999:Image Processing*, vol. 3661, 1999, pp. 620–631.

[22] J. H. Downs, J. L. Lancaster, and P. T. Fox, "Surface based spatial normalization using convex hulls," in *Brain Warping*. San Diego, CA: Academic, 1999.

[23] T. Hartkens, D. Hill, A. D. Castellano-Smith, D. Hawkes, C. Maurer, A. Martin, W. Hall, and C. T. H. Liu, "Using points and surfaces to improve voxel-based non-rigid registration," in *MICCAI*, 2002, pp. 565–572. [Online]. Available: citeseer.ist.psu.edu/571972.html

[24] C. Davatzikos, J. Prince, and R. Bryan, "Image registration based on boundary mapping," *IEEE Transactions on Medical Imaging*, 1996.

[25] C. Davatzikos and J. Prince, "Brain image registration based on curve mapping," in *IEEE Workshop Biomedical Image Anal.*, 1994, pp. 245–254.

[26] S. C. Joshi and M. I. Miller, "Landmark matching via large deformation diffeomorphisms," *IEEE Transactions on Image Processing*, vol. 9, no. 8, August 2000.

[27] G. Gerig, S. Joshi, T. Fletcher, K. Gorczowski, S. Xu, S. M. Pizer, and M. Styner, "Statistics of population of images and its embedded objects: Driving applications in neuroimaging," in *ISBI*, April 2006, pp. 1120–1123.

[28] T. Liu, D. Shen, and C. Davatzikos, "Deformable registration of cortical structures via hybrid volumetric and surface warping," *NeuroImage*, vol. 22, no. 4, pp. 1790–1801, 2004.

[29] G. E. Christensen, P. Yin, M. W. Vannier, K. S. C. Chao, J. L. Dempsey, and J. F. Williamson, "Large-deformation image registration using fluid landmarks," in *Image Analysis and Interpretation, 2000. Proceedings. 4th IEEE Southwest Symposium*, 2000, pp. 269–273.

[30] H. J. Johnson and G. E. Christensen, "Consistent landmark and intensity-based image registration," *IEEE Transactions on Medical Imaging*, vol. 21, no. 5, pp. 450–461, 2002.

[31] G. E. Christensen, "Consistent linear-elastic transformations for image matching," *Lecture Notes in Computer Science*, vol. 1613, pp. 224–237, 1999.

[32] A. P. Eriksson and K. Åström, "On the bijectivity of thin plate transforms," in *Swedeish Symposium on Image Analysis*, 2005, pp. 53–56.

[33] P. M. Thompson, D. MacDonald, M. S. Mega, C. J. Holmes, A. C. Evans, and A. W. Toga, "Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces," *Journal of Computer Assisted Tomography*, vol. 21, no. 4, pp. 567–581, Jul.-Aug. 1997.

[34] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, "Laplace-Beltrami operator and brain surface flattening," *IEEE Transactions on Medical Imaging*, vol. 18, pp. 700–711, 1999.

[35] A. A. Joshi, R. M. Leahy, P. M. Thompson, and D. W. Shattuck, "Cortical surface parameterization by *p*-harmonic energy minimization." in *ISBI*, 2004, pp. 428–431.

[36] B. Tang, G. Sapiro, and V. Caselles, "Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case." *International Journal of Computer Vision*, vol. 36, no. 2, pp. 149–161, 2000.

[37] L. A. Segel, *Mathematics Applied to Continuum Mechanics*. Dover Publications, 1987.

[38] C. Chefd'Hotel, G. Hermosillo, and O. Faugeras, "A variational approach to multi-modal image matching," in *IEEE Workshop on Variational and Level Set Methods in Computer Vision*, 2001, pp. 21–28.

[39] D. W. Shattuck and R. M. Leahy, "Brainsuite: An automated cortical surface identification tool," *Medical Image Analysis*, vol. 8, no. 2, pp. 129–142, 2002.

[40] ——, "Graph based analysis and correction of cortical volume topology," *IEEE Transactions on Medical Imaging*, vol. 20, no. 11, pp. 1167–1177, 2001.

[41] P. M. Thompson, K. M. Hayashi, G. de Zubicaray, A. L. Janke, S. E. Rose, J. Semple, D. M. Doddrell, T. D. Cannon, and A. W. Toga, "Detecting dynamic and genetic effects on brain structure using high dimensional cortical pattern matching," in *Proceedings of ISBI*, 2002.

[42] X. Han, C. Xu, and J. Prince, "A topology preserving geometric deformable model and its application in brain cortical surface reconstruction," in *Geometric Level Set Methods in Imaging, Vision, and Graphics*, S. Osher and N. Paragios, Eds. Springer Verlag, 2003.

[43] J. Jost, *Riemannian geometry and geometric analysis*. Springer Verlag, 2002.

[44] S. Nishikawa, *Variational Problems in Geometry*, ser. Translations of Mathematical Monographs. AMS, 2001, vol. 205.

[45] Y. Xin, *Geometry of harmonic maps*. Birkhäuser, 1996.

[46] J. Eells and J. H. Sampson, "Harmonic mappings of Riemannian manifolds," *Ann. J. Math.*, pp. 109–160, 1964.

[47] R. Hamilton, "Harmonic maps of manifolds with boundary," in *Lecture Notes in Mathematics*, ser. 471. Springer, 1975.

[48] A. Fardoun and R. Regbaoui, "Heat flow for p-harmonic maps between compact Riemannian manifolds," *Indiana Univ. Math. J.*, vol. 51, pp. 1305–1320, 2002.

[49] T. Kanai, H. Suzuki, and F. Kimura, "Three-dimensional geometric metamorphosis based on harmonic maps," *The Visual Computer*, vol. 14, no. 4, pp. 166–176, 1998.

[50] Y. Wang, X. Gu, and S. T. Yau, "Volumetric harmonic map," *Communications in Information and Systems*, vol. 3, no. 3, pp. 191–202, 2004.

[51] F. Mémoli, G. Sapiro, and S. Osher, "Solving variational problems and partial differential equations mapping into general target manifolds," *J. Comput. Phys.*, vol. 195, no. 1, pp. 263–292, 2004.

[52] H. D. Tagare, D. Groisser, and O. Skrinjar, "A geometric theory of symmetric registration," *CVPRW*, vol. 0, p. 73, 2006.

[53] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa, "The quickhull algorithm for convex hulls," *ACM Transactions on Mathematical Software*, 1996.

[54] D. Shen and C. Davatzikos, "Hammer: Hierarchical attribute matching mechanism for elastic registration," *IEEE Transactions on Medical Imaging*, vol. 21, no. 11, 2002.

[55] A. P. Zijdenbos, B. M. Dawant, R. A. Margolin, and A. Palmer, "Morphometric analysis of white matter lesions in mr images," *IEEE Transactions on Medical Imaging*, vol. 13, pp. 716–724, Dec. 1994.

[56] D. W. Shattuck and R. M. Leahy, "BrainSuite: An automated cortical surface identification tool." in *MICCAI*, ser. Lecture Notes in Computer Science, S. L. Delp, A. M. DiGioia, and B. Jaramaz, Eds., vol. 1935. Springer, 2000, pp. 50–61.

[57] P. M. Thompson, K. M. Hayashi, E. R. Sowell, N. Gogtay, J. N. Giedd, J. L. Rapoport, G. I. de Zubicaray, A. L. Janke, S. E. Rose, J. Semple, D. M. Doddrell, Y. L. Wang, T. van Erp, T. D. Cannon, and A. W. Toga, "Mapping cortical change in Alzheimer's disease, brain development and schizophrenia," *NeuroImage*, vol. 23, no. 1, pp. S2–S18, Sept. 2004.

[58] A. Leow, P. M. Thompson, H. Protas, and S.-C. Huang, "Brain warping with implicit representations." in *ISBI*. IEEE, 2004, pp. 603–606.

[59] F. Mémoli, G. Sapiro, and P. Thompson, "Implicit brain imaging," *NeuroImage*, vol. 23, no. 1, pp. S179–S188, 2004.

[60] Y. Wang, X. Gu, K. Hayashi, T. Chan, P. Thompson, and S. Yau, "Brain surface parameterization using Riemann surface structure," in *MICCAI 2005*, 2005, pp. 657–665.

[61] U. Grenander and M. I. Miller, "Computational anatomy: an emerging discipline," *Q. Appl. Math.*, vol. LVI, no. 4, pp. 617–694, 1998.

[62] G. E. Christensen, R. Rabbitt, and M. I. Miller, "3D brain mapping using a deformable neuroanatomy," *Physics in Medicine and Biology*, vol. 39, pp. 609–618, March 1994.

[63] V. Camion and L. Younes, "Geodesic interpolating splines," *Lecture Notes in Computer Science*, pp. 513–527, 2001.

[64] Y. Ge, J. M. Fitzpatrick, R. M. Kessler, M. Jeske-Janicka, and R. A. Margolin, "Intersubject brain image registration using both cortical and subcortical landmarks," in *Proc. SPIE Vol. 2434, p. 81-95, Medical Imaging 1995: Image Processing, Murray H. Loew; Ed.*, May 1995, pp. 81–95.

[65] P. M. Thompson, C. Vidal, J. N. Giedd, P. Gochman, J. Blumenthal, R. Nicolson, A. W. Toga, and J. L. Rapoport, "Mapping adolescent brain change reveals dynamic wave of accelerated gray matter loss in very early-onset schizophrenia," *PNAS*, vol. 98, no. 20, pp. 11 650–11 655, 2001.

[66] E. R. Sowell, B. S. Peterson, P. M. Thompson, S. E. Welcome, A. L. Henkenius, and A. W. Toga, "Mapping cortical change across the human life span," *Nature Neuroscience*, vol. 6, pp. 309–315, 2003.

[67] P. M. Thompson, K. M. Hayashi, G. de Zubicaray, A. L. Janke, S. E. Rose, J. Semple, D. Herman, M. S. Hong, S. S. Dittmer, D. M. Doddrell, and A. W. Toga, "Dynamics of gray matter loss in Alzheimer's disease," *The Journal of Neuroscience*, vol. 23, no. 3, pp. 994–1005, 2003.

[68] D. W. Shattuck, A. MacKenzie-Graham, and A. W. Toga, "Duff: software tools for visualization and processing of neuroimaging data," *IEEE International Symposium on Biomedical Imaging: Macro to Nano, 2004*, vol. 1, pp. 644–647, April 2004.

[69] N. Ray and B. Levy, "Hierarchical least squares conformal map," in *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*. Washington, DC, USA: IEEE Computer Society, 2003, p. 263.

[70] J. S. George, C. J. Aine, J. C. Mosher, D. M. Schmidt, D. M. Ranken, H. A. Schlitt, C. C. Wood, J. D. Lewine, J. A. Sanders, and J. W. Belliveau, "Mapping function in the human brain with magnetoencephalography, anatomical magnetic resonance imaging, and functional

magnetic resonance imaging," *J Clin Neurophysiol*, vol. 12, no. 5, pp. 406–431, Sep. 1995.

[71] F. L. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 567–585, June 1989.

[72] I. Kreyzig, *Differential Geometry*.   Dover, 1999.

[73] M. Do Carmo, *Differential Geometry of Curves and Surfaces*.   Prentice-Hall, 1976.

[74] P. M. Thompson, M. S. Mega, C. Vidal, J. Rapoport, and A. W. Toga, "Detecting disease-specific patterns of brain structure using cortical pattern matching and a population-based probabilistic brain atlas," in *Proc. 17th IPMI2001, Davis, CA, USA*, 2001, pp. 488–501.