

Diffeomorphic Spectral Matching of Cortical Surfaces

Herve Lombaert¹, Jon Sporring^{1,2}, and Kaleem Siddiqi¹

¹ Centre for Intelligent Machines, McGill University, Montreal

² eScience Center, Computer Science, University of Copenhagen, Denmark

Abstract. Accurate matching of cortical surfaces is necessary in many neuroscience applications. In this context diffeomorphisms are often sought, because they facilitate further statistical analysis and atlas building. Present methods for computing diffeomorphisms are based on optimizing flows or on inflating surfaces to a common template, but they are often computationally expensive. It typically takes several hours on a conventional desktop computer to match a single pair of cortical surfaces having a few hundred thousand vertices. We propose a very fast alternative based on an application of spectral graph theory on a novel association graph. Our symmetric approach can generate a diffeomorphic correspondence map within a few minutes on high-resolution meshes while avoiding the sign and multiplicity ambiguities of conventional spectral matching methods. The eigenfunctions are shared between surfaces and provide a smooth parameterization of surfaces. These properties are exploited to compute differentials on highly folded cortical surfaces. Diffeomorphisms can thus be verified and invalid surface folding detected. Our method is demonstrated to attain a vertex accuracy that is at least as good as that of FreeSurfer and Spherical Demons but in only a fraction of their processing time. As a practical experiment, we construct an unbiased atlas of cortical surfaces with a speed several orders of magnitude faster than current methods.

1 Introduction

The cerebral cortex is the center of many important functional activities, including vision and perception, and these are often studied by establishing properties which hold across a large population. These studies thus require fast and accurate algorithms for cortical surface matching are often sought. Early approaches based on volumetric comparisons [1] ignore the complex geometry of cortical folds, and therefore, produce misaligned cortical areas [2]. Recent surface-based approaches either optimize flows on surfaces [3,4,5] or inflate cortical surfaces to a spherical template [6,7,8]. Methods that “flow” surfaces into one another, such as LDDMM [9] and Currents [10,11], provide an elegant mathematical framework that guarantees diffeomorphic deformations between surfaces, i.e., they provide smooth and invertible correspondence maps. However, these methods are computationally expensive and typically require several hours on a conventional desktop

computer to process meshes containing a few thousand vertices. On the other hand, spherical methods, such as FreeSurfer [6] and Spherical Demons [8], establish correspondences on simplified spherical models of the cortex. They handle the complexity of the cortical folds by exploiting metrics that are derived from the original surfaces, such as sulcal depth and mean curvature. Unfortunately, in these approaches, the surfaces need to be inflated to spheres via an expensive process [12]. They too require a few hours to process high-resolution meshes. Current methods may therefore be computationally prohibitive for neuro-studies that involve several thousands of individuals.

Spectral graph theory [13] offers a fast alternative where surfaces containing several hundred thousand vertices can be matched within minutes on a conventional computer [14,15]. Spectral methods facilitate the correspondence problem by matching shapes in the spectral domain. Their spectral representations are in fact invariant to isometry, i.e., two shapes with identical geodesic distances between points have identical spectral representations. However, perturbations in shape isometry, such as expansion and compression of surfaces, change these spectral representations, and thus, alter the matching accuracy. This has limited the use of spectral methods in matching coarse hierarchical structures [16] or in defining global metrics for shape analysis [17]. Previous work attempted to correct these spectral representations with rigid [18] and nonrigid transformations [19,14,15]. In addition, a vertex accuracy of up to 88% of FreeSurfer's performance is achieved in [14] by embedding additional information, such as sulcal depth, in extra dimensions to the spectral representations. However, these extended representations may no longer be smooth and, consequently, the correspondence maps are not guaranteed to be diffeomorphic.

This paper proposes a new accurate surface matching approach that retains the speed advantage of spectral matching methods while guaranteeing diffeomorphic correspondence maps between cortical models of several hundred thousand vertices. Our method exploits a novel association graph that is formed with two meshes and a preliminary correspondence map generated from conventional spectral matching. The spectral decomposition of this unique association graph creates a shared set of eigenvectors that enables a direct comparison between meshes. This contrasts with the conventional spectral methods that produce two separate sets of eigenvectors, and thus, need to handle ambiguities inherent to the sign and multiplicity of eigenvectors, as well as perturbations in isometry. Additionally, the eigenvectors computed with our method provide a smooth parameterization of surfaces. We exploit this property to compute differentials on highly curved cortical surfaces. More precisely, we define a novel Jacobian operator on surfaces to verify diffeomorphisms of correspondence maps and to detect invalid folding of surfaces. Our new diffeomorphic method is demonstrated to produce, in only 350 seconds on a conventional laptop computer (2.53GHz Intel Core 2 Duo), a vertex accuracy that is at least as good as that of FreeSurfer and Spherical Demons. We finally show that our method can be used to construct accurate and unbiased atlases with a significant speed advantage over current competing methods.




Fig. 1. Spectral matching – Correspondences are made using spectral representations of surfaces (built with Laplacian eigenmodes). To aid visualization, equivalent points have a unique color on the above surfaces ($\text{rgb} = u^{(1,2,3)}$).

2 Spectral Matching

We begin by reviewing the basic concepts for matching two shapes with spectral methods.

Graph Laplacian. Let us build the graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ from the set of vertices (with position x) and edges of a surface model S . We may define the $|V| \times |V|$ *weighted adjacency* matrix W in terms of node affinities, e.g., $w_{ij} = \|x_i - x_j\|^{-2}$ if $\exists e_{ij} \in \mathcal{E}$ (the inverse distance between neighboring points) or 0 otherwise. The diagonal *node degree* matrix D is the sum of all point affinities $d_i = \sum_j W_{ij}$. The *general Laplacian operator* on a graph was formulated in [20] as a $|V| \times |V|$ matrix with the form $\mathcal{L} = G^{-1} (D - W)$ where G is the diagonal *node weighting* matrix ($G = I$, $G = D$ or any meaningful node weighting). It was found [14,15] that setting higher weights on gyral points produces better cortical matchings, e.g., the node weighting, $g_i = \exp(-h_i)$, is the exponential of the sulcal depth h_i (computed with FreeSurfer) at point i .

Spectral Coordinates. The spectral decomposition of the graph Laplacian $\mathcal{L} = U\Lambda U^{-1}$ provides the eigenvalues $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{|V|})$ and the associated eigenvectors $U = (u^{(0)}, u^{(1)}, \dots, u^{(|V|)})$, where $u^{(\cdot)}$ is a column of U . The values of $u^{(\cdot)}$ depict in fact a vibration mode of the shape S , which is a surface function, [13], and thus, the term *eigenmode* is used. These eigenmodes must be additionally corrected for their sign ambiguity, multiplicity, and perturbation in isometry (see [14,15] for more details). We denote as the *spectral representation*, $\text{spec}(S)$, a k -dimensional embedding of the shape S where a point has the *spectral coordinates* defined as $u^{(1, \dots, k)}$, which is a row of the truncated matrix U^k .

Spectral Matching with Vertex Accuracy. The correspondence problem is to match a point x_i on S_1 , with $y_{c(i)}$ on S_2 . The map $c : x_i \mapsto y_{c(i)}$ (also denoted here as $S_1 \mapsto S_2 \circ c$) is established with pairs of closest points (u, v) between spectral representations $\text{spec}(S_1)$ and $\text{spec}(S_2)$, as illustrated in Fig. 1. However, in order to achieve vertex accuracy in cortical surface matching, [14,15] proposed to incorporate extra information, such as sulcal depth h , in extended

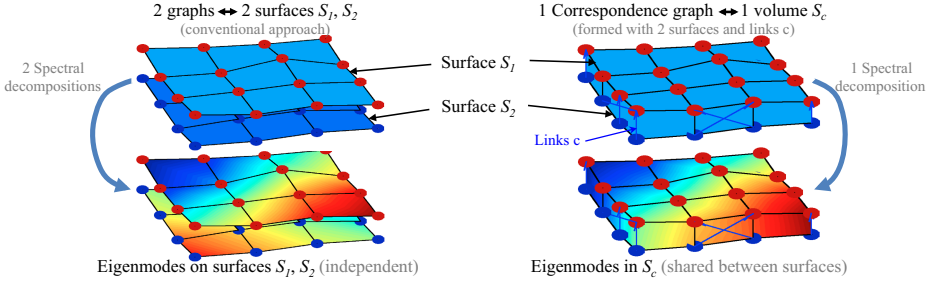


Fig. 2. Correspondence Graph – *Left:* Conventional spectral matching decomposes 2 independent graphs (built from 2 surfaces) – *Right:* Our approach decomposes 1 correspondence graph (built from one volumetric entity formed with the surfaces and a map c between them). Note that in our approach eigenmodes (u, v) do match without the explicit handling of their sign flips, reordering or changes in isometry. They are smooth even if c is irregular (shown here with many-to-one and crossing links c).

spectral representations where point coordinates are then (u, h) . The map $c(i) = \arg \min_j \|(u_i, h_i)_{S_1} - (u_j, h_j)_{S_2}\|^2$ is solved with a simple nearest-neighbor search between these extended representations. Unfortunately, this incorporation of extra information creates discontinuities in the correspondence map in the sense that neighbors in space may no longer be neighbors in the extended spectral representations.

3 Diffomorphic Spectral Matching

We now show that the spectral decomposition of a novel associative graph provides the ability to compute diffeomorphic maps between two surfaces.

3.1 Correspondence Graph

We define the *correspondence graph* $\mathcal{G}_c = \{\mathcal{V}_{1,2}, \mathcal{E}_{1,2,c}\}$ as the union of the set of vertices and edges of two surfaces $S_{1,2}$ with an initial set of correspondence links c between both surfaces. The initial map c , not necessarily dense, may be computed with a conventional spectral matching methods. We used [14] since it is optimized for cortical matching. This correspondence graph is a 2-split graph, where each surface is a separable set, and has its $|\mathcal{V}_1 \cup \mathcal{V}_2| \times |\mathcal{V}_1 \cup \mathcal{V}_2|$ weighted adjacency matrix in the form:

$$W_c = \begin{bmatrix} W_1 & W_{12} \\ W_{21} & W_2 \end{bmatrix}, \tag{1}$$

where W_1 and W_2 are the weighted adjacency matrices of both surfaces and W_{12} , W_{21} are the weighted adjacency matrices defined by the links c between surfaces, i.e., $w_{ij} = \|x_i - x_j\|^{-2}$ if nodes (i, j) or (j, i) are connected. The graph Laplacian operator \mathcal{L}_c of \mathcal{G}_c is defined as previously. Such a graph is illustrated in Fig. 2 where surfaces have been simplified to small patches to ease illustration.




Fig. 3. Spectral Decomposition of the Correspondence Graph – One graph is decomposed (rather than two as in conventional methods). This produces a smooth and shared parameterization (u, v) (color-coded between $[-1, +1]$) between surfaces. This prevents handling sign flips, reordering and corrections in isometry of (u, v) . The principal eigenmodes are shown on smoothed meshes to aid visualization.

Shared Parameterization. The graph \mathcal{G}_c embeds both surfaces within one single entity S_c , which is a double layered graph that represents two cortical surfaces interconnected with c . The spectral decomposition, $\mathcal{L}_c = U_c \Lambda_c U_c^{-1}$, therefore provides one orthonormal basis of the whole entity S_c . This contrasts with conventional spectral matching methods that produce two independent sets of eigenmodes. Moreover, each eigenmode $u_c^{(\cdot)}$, a $|\mathcal{V}_1 \cup \mathcal{V}_2|$ column vector of U_c , is separable back into two functions: $u^{(\cdot)}$, the first $|\mathcal{V}_1|$ values of $u_c^{(\cdot)}$, is a surface function on S_1 , and $v^{(\cdot)}$, which is the last $|\mathcal{V}_2|$ values of $u_c^{(\cdot)}$, is a surface function on S_2 (illustrated in Fig. 3). They share in fact the same eigenvalue and represent the same vibration mode. There is therefore no need to correct for a sign flip between $u^{(\cdot)}$ and $v^{(\cdot)}$, nor to reorder the set of eigenmodes or correct for perturbations in isometry since the surface functions $u^{(\cdot)}$ and $v^{(\cdot)}$ are derived from the same single entity and not from two independent surfaces.

Smooth Parameterization. Courant’s nodal line theorem [21,22,23] demonstrates that any n^{th} eigenmode of the graph Laplacian has the remarkable property of being *smooth* and *monotonous* between at most n poles of vibrations. This phenomenon is observed on Fig. 3 where $u^{(1)}$ and $u^{(2)}$ have two poles (red and blue spots), $u^{(3)}$ has three poles, $u^{(4)}$ has four poles, and so on.

Since S_c always remains one single double-layered entity, the function $u^{(\cdot)}$ varies smoothly within the global shape formed by S_c (as illustrated in Fig. 2 *right*, both $u^{(1)}$ and $v^{(1)}$ always increase smoothly in the same direction even if c has crossing links). The coordinate system formed by $u^{(1, \dots, k)}$ and $v^{(1, \dots, k)}$ provides, therefore, a smooth parameterization within the entity S_c and is shared between surfaces $S_{1,2}$.

Differentiable Space. We further consider each eigenmode $u^{(\cdot)}$ as continuous and differentiable between the discrete values at the mesh nodes. This is motivated by the fact that the graph Laplacian approximates the Laplace-Beltrami operator on Riemannian manifolds [24]. Since the mesh nodes associate a smooth spectral coordinate with a position in space, $u \mapsto x$ on S_1 and $v \mapsto y$ on S_2 , it is

possible to model the positions of points on S_1 and S_2 that have equal spectral coordinates, using any type of differentiable interpolation. We chose the Gaussian kernel for its regularity (differentiable in C^∞) and asymptotic behavior [25]. For instance, the node i on S_1 has its equivalent point on S_2 with position:

$$y'_i = \frac{\sum_{j \in \mathcal{N}_{\psi(i)}} \omega_{ij} y_j}{\sum_{j \in \mathcal{N}_{\psi(i)}} \omega_{ij}}, \tag{2}$$

where $\mathcal{N}_{\psi(i)}$ is the set of neighboring nodes of $\psi(i)$ on S_2 ; $\psi(i)$ is the node on S_2 with the closest spectral coordinate to u_i (the mapping ψ is found as previously with a nearest-neighbor search $\psi(i) = \arg \min_j \|u_i - v_j\|^2$); and ω_{ij} is the spectral similarity, e.g, $\omega_{ij} = \exp(-\|v_i - v_j\|^2 / 2\sigma^2)$. This simple modeling scheme creates a k -D differentiable space \mathcal{S}^k between S_1 and S_2 .

Diffeomorphic Mapping. The differentiable space \mathcal{S}^k allows us to establish symmetric correspondences between surfaces by locating points on S_1 and S_2 that have equal spectral coordinates, $\phi_{1 \rightarrow 2} : x_i \mapsto y'_i$ using Eq. 2 and conversely $\phi_{2 \rightarrow 1} : y_j \mapsto x'_j$. Such a mapping prevents an invalid folding of space (crossing of links is avoided since \mathcal{S}^k is monotonous between poles of vibration) as well as the collapse of space (many-to-one correspondences are prevented since equivalent points (x_i, y'_i) have unique spectral coordinates). The mapping ϕ between $S_{1,2}$ is consequently diffeomorphic (ϕ is smooth, bijective, and invertible $\phi_{1 \rightarrow 2}^{-1} = \phi_{2 \rightarrow 1}$).

Jacobian Operator on Surfaces. The continuous spectral parameterization (u, v) may be used to compute differentials on the curved space defined by $S_{1,2}$. We express the correspondence map ϕ in terms of (u, v) with $\phi_u : u_i \mapsto v'_i$ and define its Jacobian matrix as $J(\phi_u) = \left(\frac{\partial \phi_{u(1)}}{\partial u^{(1)}} \dots \frac{\partial \phi_{u(1)}}{\partial u^{(k)}}; \dots; \frac{\partial \phi_{u(k)}}{\partial u^{(1)}} \dots \frac{\partial \phi_{u(k)}}{\partial u^{(k)}} \right)$ where $\partial \phi_{u(\cdot)}$ and $\partial u_i^{(\cdot)}$ are estimated in the neighborhood \mathcal{N}_i using Taylor series approximation for central differences. Its determinant is simply denoted as the Jacobian $|J|$. To illustrate, two equivalent surface elements on $S_{1,2}$ may appear flipped in space, and thus, generate a negative Jacobian $|J(\phi)|$ when it is expressed in Cartesian coordinates, while in fact, these vectors may have been always pointing outwards from the surface. Expressing the variations in the spectral coordinates generates in this case a positive Jacobian $|J(\phi_u)|$. We, therefore, propose to study the variations of a surface map ϕ using the Jacobian matrix $J(\phi_u)$ expressed in terms of spectral coordinates (u, v) .

Objective Function. The initial map c is generated using [14], which minimizes $E_c = (u_{S_1} - u_{S_2} \circ c)^2 + (h_{S_1} - h_{S_2} \circ c)^2$. To summarize our method, c is reused to build the correspondence graph \mathcal{G}_c , whose spectral decomposition creates a differentiable parameterization $(u, v)_{S_c}$ in $\text{spec}(S_c)$. A symmetric and diffeomorphic map ϕ is then found by minimizing $E_\phi = (u_{S_c} - v_{S_c} \circ \phi)^2$ via simple nearest-neighbor searches.

3.2 Unbiased Atlas Construction

The computation of an average shape typically relies on an iterative evolution of an initial reference shape [26,27,27]. This process may, however, be biased by the choice of the initial reference [28]. We define the average cortical surface S_0 as the geometric mean of all surfaces in a dataset \mathcal{S} . The position of its vertices is defined with $\bar{x}_i = \frac{1}{|\mathcal{S}|} \sum_{t \in \mathcal{S}} x_i^{(t)}$, where $x_i^{(t)}$ is the interpolated position of point i on surface S_t , computed with Eq. 2. This requires the computation of mappings $\{\phi_{0 \rightarrow t}\}_{t \in \mathcal{S}}$, as well as an initial reference surface S_0^{init} .

Transitivity. Since ϕ is diffeomorphic, it is also transitive [29]. The composition of mappings $\phi_{s \rightarrow t} \circ \phi_{t \rightarrow u}$, from S_s to S_t to S_u , is therefore identical to the mapping $\phi_{s \rightarrow u}$, from S_s to S_u . This transitive relation implies that the mapping from the initial reference $\phi_{0 \rightarrow t}$ is equivalent to the composition with any other intermediate mapping, $\phi_{0 \rightarrow t} = \phi_{0 \rightarrow s} \circ \phi_{s \rightarrow t}$, and the average position \bar{x}_i is consequently unbiased to the choice of S_0^{init} since the mapping $\phi_{0 \rightarrow t}$ could be composed with any intermediate surface S_s in the dataset. Our diffeomorphic spectral method, therefore, has the advantage that it constructs an unbiased atlas with a direct, one-step, approach (without the need for an iterative evolution of S_0^{init} [28]).

4 Results

We begin our validation by verifying the key properties of our method and then assess its matching accuracy using synthetic and real cortical surfaces. Our dataset consists of 16 real cortical surfaces ranging from 109k to 174k vertices with an average resolution of 0.88mm, generated from MRI.

4.1 Verifying Properties of Our Method

We verify through a series of simple experiments that our method produces *a*) a shared parameterization between surfaces, *b*) diffeomorphic mappings from irregular (i.e., non-smooth) correspondences, and *c*) mappings that maintain transitivity.




Fig. 4. Shared Parameterization – Surfaces show eigenmodes $(u^{(1)}, v^{(1)})$ color-coded between $[-1, +1]$. The eigenmodes computed on separate surfaces may not overlap on them (*left*), whereas the spectral decomposition of 1 correspondence graph produces a common set of eigenmodes (*right*). 3 samples out of 1000 are shown. Overlap errors are measured with $SSD_{u,v}$.




Fig. 5. Diffeomorphic Mapping — An arbitrary and randomly perturbed correspondence map c (seen in the sliced views) folds over triangles on a sphere (blue indicates a face normal pointing inward), whereas the diffeomorphic map generated from c (by decomposing \mathcal{G}_c) recovers the original correspondences c and guarantees no folding of space by ϕ . The mapping ϕ is therefore diffeomorphic since (u, v) is smooth and bijective between surfaces.

Shared Parameterization. We study the stability of eigenmodes on circular shapes. Since these shapes have ambiguous axes of symmetry, $\text{spec}(S_1)$ and $\text{spec}(S_2)$ should differ by an arbitrary rotation on disks. However, $\text{spec}(S_c)$ should theoretically produce a shared set of eigenmodes (u, v) between surfaces. We use two uniform meshes of a disk, whose node positions have been perturbed with Gaussian noise (this perturbs isometry between disks). Their ground truth correspondence map is defined as their direct overlap, $c(i) = i$. This mapping c is then perturbed with Gaussian noise $c'(i) = \arg \min_j \|(y_i + \text{noise}) - y_j\|^2$ (as shown in Fig. 4). Firstly, we generate $u = \text{spec}(S_1)$ and $v = \text{spec}(S_2)$ for both disks. This is the conventional spectral matching approach. Secondly, we generate $(u, v) = \text{spec}(S_c)$ using the correspondence graph formed by c' . We repeat the experiment a thousand times and measure the overlap $SSD_{u,v} = \sum_{i \in S_1} \|u_i - v_i\|^2$. As expected, (u, v) rotates arbitrarily on the disks when they are computed separately. The worst overlap among the thousand samples gives $SSD_{u,v} = 0.95$. In contrast, (u, v) demonstrate a better overlap when decomposing the correspondence graph. The worst overlap, at $SSD_{u,v} = 0.05$, actually still produces aligned eigenmodes (Fig. 4). The decomposition of a correspondence graph is therefore demonstrated to generate a stable set (u, v) of eigenmodes that is in common between surfaces.

Diffeomorphic Mapping. We now verify the smoothness of correspondence maps by testing for folding of the triangulation between S_1 and $S_2 \circ \phi$. We use two uniform models $S_{1,2}$ of a sphere and rotate them arbitrarily. The face normals on both models are all initially pointing outward. We use the direct mapping $c(i) = i$ between $S_{1,2}$ as ground truth. We randomly perturb it as in the previous experiment (see slice view in Fig. 5). A discontinuity in c' is detected with a triangle fold-over (we check if a face normal points inward from the sphere). We generate a diffeomorphic map ϕ by decomposing the correspondence




Fig. 6. Transitive Mapping – The average distance error between 4,096 (16^3) composed surfaces $S_1 \circ \phi_{1 \rightarrow 2} \circ \phi_{2 \rightarrow 3}$ and $S_1 \circ \phi_{1 \rightarrow 3}$. FreeSurfer and the Spherical Demons have an error below the mesh resolution (0.88mm). The non-diffeomorphic spectral matching fails to achieve transitivity since it does not generate diffeomorphic maps. The new diffeomorphic spectral matching demonstrates transitivity with an error of 0.00292(± 0.00432) mm.

graph formed with c' and check if any face normal points inward. We repeat the experiment a thousand times with random perturbation of c and random rotations of spheres. Each time, ϕ is found to be identical to c (a one-to-one map) and no triangle fold-over is observed across the experiment (smooth map). The map ϕ produced by our method is, therefore, demonstrated to be diffeomorphic (smooth and bijective).

Transitive Mapping. We check for transitivity of our correspondence maps using real cortical surfaces. We first compute all 256 correspondence maps ϕ between all pairs of cortical surfaces with four methods: FreeSurfer (FS) [6], Spherical Demons (SD) [8], Spectral Matching (SM) [14], and our diffeomorphic spectral matching. We then generate pairs of surfaces $S_1 \circ \phi_{1 \rightarrow 2} \circ \phi_{2 \rightarrow 3}$ and $S_1 \circ \phi_{1 \rightarrow 3}$ using all 4,096 (16^3) possible compositions $\phi_{1 \rightarrow 2} \circ \phi_{2 \rightarrow 3}$ (using Eq. 2 and compositions of vertex indexing). We finally verify transitivity by measuring the distances between equivalent points (x, y) on composed surfaces, $\sqrt{\sum_{i \in S_1} \|x_{\phi_{1 \rightarrow 2} \circ \phi_{2 \rightarrow 3}(i)} - y_{\phi_{1 \rightarrow 2 \rightarrow 3}(i)}\|^2}$ (Fig. 6). The non-diffeomorphic SM did not satisfy transitivity with an error of 4.54 mm, i.e., $S_{1 \rightarrow 2 \rightarrow 3} \neq S_{1 \rightarrow 3}$. FS and SD demonstrated transitivity with errors of 0.43mm, 0.36mm (below the mesh resolution of 0.88mm) while our method produced a lower error of 0.003mm, i.e., $S_{1 \rightarrow 2 \rightarrow 3} \approx S_{1 \rightarrow 3}$.

4.2 Validating Correspondence Maps

We now assess the matching accuracy by testing for robustness to deformation and we compare our method with FreeSurfer and Spherical Demons.

Robustness to Deformation. The matching accuracy is evaluated in this synthetic experiment by monitoring the difference between the computed map ϕ and the ground truth $c(i) = i$, while increasing the level of deformation of the surface. A surface is deformed with the transformation $z' = (1 - \alpha)z$. This simulated head compression creates a controlled environment that does not produce any triangle fold-over, nor any intersecting face. We deform all 16 cortical surfaces

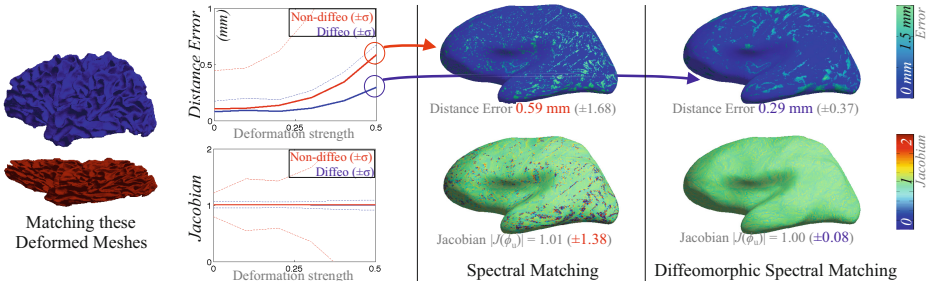


Fig. 7. Robustness to Deformation – We monitor the correspondence error between two meshes while increasing the deformation strength. The errors at maximal deformation (in *mm*) and the Jacobian of the correspondence map are shown. Our diffeomorphic spectral matching produces smaller errors and smoother correspondence maps.

by varying $\alpha \in [0; 0.5]$ and observe the differences between conventional spectral matching [14] and our diffeomorphic method. At maximal deformation $\alpha = 0.5$, SM gives an average error of 0.58mm (± 1.68) while our method gives an error of 0.29mm (± 0.08). The Jacobian $|J(\phi_u)|$ is strikingly different in both methods (Fig. 7): SM produces negative Jacobians (non-smooth map, breaking diffeomorphism) with an average $|J(\phi_u)| = 1.01(\pm 1.38)$, while our method gives strictly positive Jacobian (i.e., always diffeomorphic) with $|J(\phi_u)| = 1.00(\pm 0.08)$.

Benchmark with FreeSurfer and Spherical Demons. A ground truth mapping between real cortical surfaces is unfortunately unknown. However, the widely used FreeSurfer and Spherical Demons methods, can be used as a benchmark for evaluating the matching accuracy of our method. Our dataset has for all surfaces a labeling of cortical parcellations [30]. We verify how our method aligns these parcellations between individuals and compare these overlaps with the performance of FS, SD and SM. We measure the Dice overlap ($2|A \cap B|/(|A| + |B|)$) for 12 major parcellations between all 256 pairs of cortices (Fig. 8) and found an average Dice overlap of 0.84 (± 0.08) for FS (one matching took 3 hours), 0.85 (± 0.07) for SD (one matching took 2 hours), 0.82 (± 0.08) for SM (one matching took 250 seconds), and 0.83 (± 0.08) for our method (one matching took 350 seconds). Timing was measured on a 2.53GHz Core 2 Duo laptop with 8GB of RAM. The matching accuracy of spectral methods is arguably similar, achieving 99% accuracy of FS and SD, but they have a clear speed advantage. Clearly, conventional SM produces highly irregular, or non-smooth, maps with an average Jacobian of $|J(\phi_u)| = 1.61(\pm 1.00)$. This is observed as holes and island patches in the projected parcels in Fig. 8. Our method has the notable advantage of producing diffeomorphic maps with strictly positive and smooth $|J(\phi_u)| = 1.00(\pm 0.12)$.

4.3 Building Unbiased Atlases

We conclude by building atlases of cortical surfaces and verify that their constructions are unbiased to an initial reference. To do so, an arbitrary initial reference

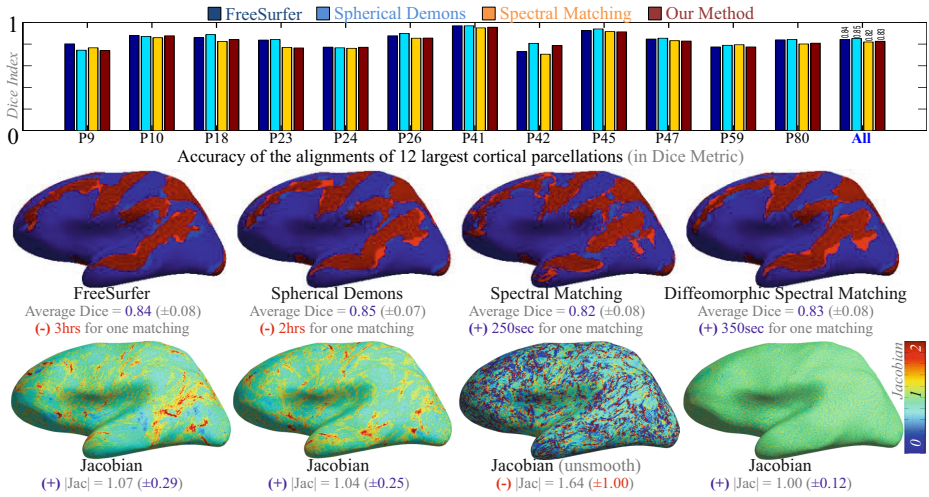


Fig. 8. Accuracy of Parcellation Alignments – Bar plot on *Top*: Average accuracy of 256 alignments of cortical parcellations (measured with Dice Index), using FreeSurfer (dark blue), Spherical Demons (light blue), Spectral Matching (yellow), and our Diffeomorphic Spectral Matching (red) – *Middle*: One example showing the alignment of 12 major parcellations (light blue for first cortex, light red for projected second cortex) – *Bottom*: Determinant of Jacobian (conventional spectral matching produces non-smooth maps). Our method yields similar accuracy than the state-of-the-art but in a fraction of its time.




Fig. 9. Unbiased Atlas Construction – Overlap of 16 average shapes computed from 16 different initial references using FS, SD, SM, and our method. To aid visualization, surface contours are of the same slice. Our method produces average shapes that are unbiased to the initial reference (all contours overlap with a variability of 0.001mm).

is first defined as one of the cortical surfaces in the dataset, e.g., $S_0^{\text{init}} = S_1$. All surfaces are then matched to this reference and the atlas is built using the average position \bar{x} of all mapped points as described in Sec. 3.2. We build 16 different atlases by iterating $S_0^{\text{init}} = S_i$. The bias to the initial reference is evaluated by measuring the average standard deviation of the distances between equivalent points across all 16 different atlases. This measures the variability of distances between atlas boundaries. We find a boundary variability of 0.29mm (± 0.07) for FS, 0.28mm (± 0.08) for SD, 1.96mm (± 0.81) for SM, and 0.0014mm (± 0.0009) for our method. Fig. 9 shows the contour overlays of all 16 atlases and shows that our

method produces stable atlases that are unbiased to the choice of initial reference, i.e., they all overlap.

5 Conclusion

Our method contributes to the challenging problem of cortical matching. First, we enhanced current spectral approaches to achieve diffeomorphism, and second, we provided a very fast algorithm that has a vertex accuracy comparable to FreeSurfer and Spherical Demons, in fact, 20 times faster with equivalent accuracy. Besides the clear speed advantage, we tackled the diffeomorphic matching problem with a direct, one-step approach that fundamentally contrasts with current iterative solutions. Our smooth and bijective correspondences are found via simple nearest-neighbor searches in the spectral domain rather than with an optimization of flows or with an inflation of surfaces to spheres. In fact, these current iterative approaches, including the LDDMM, could even reuse our diffeomorphic maps as initialization, and perhaps, gain speed and accuracy. Our method currently relies on an initial correspondence map c . We used the map generated from [14], which allowed us to obtain our high-accuracy maps. Our spectral approach can also be regarded as an action that builds a diffeomorphic map ϕ from an irregular map c . In this context, it would be interesting to see how our method performs with other maps. For instance, c could be built from fast graph matching methods [31,32] or even from manual pairing of surface landmarks. Moreover, our approach may also be related to methods based on the heat kernel [33], however, we do not exploit their multiscale properties. To conclude, our new diffeomorphic spectral matching method provides a fast and accurate alternative to current methods specialized for cortical matching.

Acknowledgments. The authors would like to thank Stanley Durrleman and BT Thomas Yeo for helpful comments, and Jonathan Polimeni for suggestions and for providing data. Funding is from the Fonds de Recherche du Québec (FRQNT) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

1. Talairach, J., Szikla, G., Tournoux, P., Prosalenti, A., Bordas-Ferrier, M., Covello, L., Jacob, M., Mempel, E.: Atlas stereotaxique du telencephale. Masson (1967)
2. Amunts, K., Malikovic, A., Mohlberg, H., Schormann, T., Zilles, K.: Brodmann's areas 17 and 18 brought into stereotaxic space-where and how variable? *NeuroImage* (2000)
3. Drury, H., Van Essen, D., Joshi, S., Miller, M.: Analysis and comparison of areal partitioning schemes using 2-D fluid deformations. *NeuroImage* 3 (1996)
4. Van Essen, D., Drury, H.: Structural and functional analyses of human cerebral cortex using a surface-based atlas. *Neuroscience* 17(18) (1997)
5. Thompson, P., Toga, A.W.: A surface-based technique for warping three-dimensional images of the brain. *TMI* 15(4) (1996)

6. Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M.: High-resolution intersubject averaging and a coordinate system for cortical surface. *Human Brain Mapping* 8 (1999)
7. Fischl, B., Rajendran, N., Busa, E., Augustinack, J., Hinds, O., Yeo, T., Mohlberg, H., Amunts, K., Zilles, K.: Cortical folding patterns and predicting cytoarchitecture. *Cereb. Cortex* 18(8) (2007)
8. Yeo, T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: fast diffeomorphic landmark-free surface registration. *TMI* 29(3) (2010)
9. Beg, F., Miller, M., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* 61 (2005)
10. Vaillant, M., Glaunès, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) *IPMI 2005*. LNCS, vol. 3565, pp. 381–392. Springer, Heidelberg (2005)
11. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. *MedIA* 13(5) (2009)
12. Segonne, F., Pacheco, J., Fischl, B.: Geometrically accurate Topology-Correction of cortical surfaces using nonseparating loops. *TMI* 26(4) (2007)
13. Chung, F.: *Spectral Graph Theory*. AMS (1997)
14. Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F.: Fast brain matching with spectral correspondence. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011*. LNCS, vol. 6801, pp. 660–673. Springer, Heidelberg (2011)
15. Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F.: FOCUSR: Feature Oriented Correspondence using Spectral Regularization - A Method for Accurate Surface Matching. *TPAMI* (2012)
16. Reuter, M.: Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *IJCV* (2009)
17. Niethammer, M., Reuter, M., Wolter, F.-E., Bouix, S., Peinecke, N., Koo, M.-S., Shenton, M.E.: Global Medical Shape Analysis Using the Laplace-Beltrami Spectrum. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) *MICCAI 2007, Part I*. LNCS, vol. 4791, pp. 850–857. Springer, Heidelberg (2007)
18. Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using Laplacian eigenfunctions and unsupervised registration. In: *CVPR* (2008)
19. Jain, V., Zhang, H.: Robust 3D shape correspondence in the spectral domain. In: *CSMA* (2006)
20. Grady, L., Polimeni, J.R.: *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer (2010)
21. Courant, R., Hilbert, D.: *Methods of Mathematical Physics*. Wiley (1989)
22. Tlustý, T.: A relation between the multiplicity of the second eigenvalue of a graph Laplacian, Courant’s nodal line theorem and the substantial dimension of tight polyhedral surfaces. *Linear Algebra* 16 (2010)
23. Colin de Verdière, Y.: Multiplicités des valeurs propres. *Laplaciens discrets et laplaciens continus*. *Rendiconti di Matematica* 13(7) (1993)
24. Belkin, M., Niyogi, P.: Convergence of Laplacian eigenmaps. In: *NIPS* (2006)
25. Nielsen, M., Andresen, P.R.: Feature displacement interpolation. In: *ICIP* (1998)
26. Studholme, C., Cardenas, V.: A template free approach to volumetric spatial normalization of brain anatomy. *Pattern Recogn. Lett.* 25(10) (2004)
27. Zöllei, L., Learned-Miller, E., Grimson, W., Wells, W.: Efficient population registration of 3D data. In: Liu, Y., Jiang, T.-Z., Zhang, C. (eds.) *CVBIA 2005*. LNCS, vol. 3765, pp. 291–301. Springer, Heidelberg (2005)

28. Guimond, A., Meunier, J., Thirion, J.P.: Average brain models: A convergence study. *CVIU* 77(2) (2000)
29. Christensen, G., Johnson, H.: Invertibility and transitivity analysis for nonrigid image registration. *J. Elec. Im.* 12(1) (2003)
30. Fischl, B., van der Kouwe, A., Destrieux, C., Halgren, E., Segonne, F., Salat, D.H., Busa, E., Seidman, L.J., Goldstein, J., Kennedy, D., Caviness, V., Makris, N., Rosen, B., Dale, A.M.: Automatically parcellating the human cerebral cortex. *Cereb. Cortex* 14(1) (2004)
31. Gold, S., Rangarajan, A.: A graduated assignment algorithm for graph matching. *TPAMI* 18(4) (1996)
32. Zheng, Y., Doermann, D.: Robust point matching for nonrigid shapes by preserving local neighborhood structures. *TPAMI* 28(4) (2006)
33. Ovsjanikov, M., M erigot, Q., M emoli, F., Guibas, L.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29(5) (2010)