Brain Transfer: Spectral Analysis of Cortical Surfaces and Functional Maps

Herve Lombaert, Michael Arcaro, Nicholas Ayache

► To cite this version:

Herve Lombaert, Michael Arcaro, Nicholas Ayache. Brain Transfer: Spectral Analysis of Cortical Surfaces and Functional Maps. Information Processing in Medical Imaging (IPMI 2015), Jul 2015, Scotland, United Kingdom. pp.474-487, 10.1007/978-3-319-19992-4_37 . hal-01203570

HAL Id: hal-01203570

<https://hal.inria.fr/hal-01203570>

Submitted on 23 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Brain Transfer: Spectral Analysis of Cortical Surfaces and Functional Maps

Herve Lombaert¹, Michael Arcaro², and Nicholas Ayache¹

¹ INRIA Sophia-Antipolis, Asclepios Team, France

² Princeton Neuroscience Institute, Princeton University

Abstract. The study of brain functions using fMRI often requires an accurate matching of cortical surface data for comparing brain activation across a population. In this context, several tasks are critical, such as surface inflation for cortical visualizations and measurements, surface matching and alignment of functional data for group-level analyses. Present methods typically treat each task separately and can be computationally expensive. It takes for example several hours to smooth and match a single pair of cortical surfaces. Furthermore, conventional methods rely on anatomical features to drive the alignment of functional data across individuals, whereas their relation to function can vary across a population. To address these issues, we propose *Brain Transfer*, a spectral framework that unifies cortical smoothing, point matching with confidence regions, and transfer of functional maps, all within minutes of computation. Spectral methods have the advantage of decomposing shapes into intrinsic geometrical harmonics, but suffer from the inherent instability of these harmonics. This limits their direct comparison in surface matching, and prevents the *spectral transfer* of functions. Our contributions consist of, first, the optimization of a spectral transformation matrix, which combines both, point correspondence and change of eigenbasis, and second, a localized spectral decomposition of functional data, via *focused harmonics*. *Brain Transfer* enables the transfer of surface functions across interchangeable cortical spaces, accounts for localized confidence, and gives a new way to perform statistics on surfaces. We illustrate the benefits of spectral transfers by exploring the shape and functional variability of retinotopy, which remains challenging with conventional methods. We find a higher degree of accuracy in the alignment of retinotopy, exceeding those of conventional methods.

1 Introduction

Major brain activities occur on the cerebral cortex. The analysis of this thin, highly convoluted surface is therefore of particular interest to neuroscience, for instance, in studying vision and perception. The geometry of the cortex can be extracted from anatomical MRI, while neural activity can be measured via changes of blood flow in functional MRI. In order to establish shape and functional relationships across a population, fast and accurate algorithms are often sought for matching surfaces and functions on surfaces. Early approaches based on volumetric warping [1] ignore the complex geometry of the cortical folds, and consequently, produce misaligned cortical areas [2]. Recent surface-based approaches, either slowly deform surfaces until matching [3, 4], or inflate cortical surfaces to a spherical template [5–7]. Methods that flow surfaces into one another, such as LDDMM [8] and Currents [9, 10], provide an elegant mathematical framework, with guarantees on diffeomorphism, but remain

Shape Spectrum – Surface Functions can be written as composition of shape harmonics

Fig. 1. Spectral Decomposition on Shapes – Any function $f(x)$ on surface 1 (*top row*), or $g(x)$ on surface 2 (*bottom row*) can be written as a sum of weighted harmonics (coloring is harmonic amplitude). Spectral weights, $\mathbf{w}^{(f)}$ and $\mathbf{w}^{(g)}$, are not directly transferable between surfaces since harmonics may change sign (bases 8), and not correspond (differences in black isolines, incompatible bases 80).

computationally expensive, and typically take several hours to match meshes with a few thousand vertices. On the other hand, spherical methods, such as FreeSurfer [5] and Spherical Demons [7], find correspondences on simplified spherical models of the cortex. These methods subsequently exploit extrinsic shape features, such as sulcal depth and mean curvature, to drive the alignment of the spherical models. Unfortunately, these approaches rely on an expensive inflation of cortical surfaces into spheres [11], which also requires a few hours to process high-resolution meshes. Moreover, current methods typically rely on anatomical features to drive the alignment of functional data across individuals, whereas their relation to function can vary across a population [12].

Spectral graph theory [13] recently offers a fast alternative for matching high-resolution cortical surfaces within minutes on a conventional laptop computer [14, 15]. Spectral methods facilitate the correspondence problem by matching shapes in a spectral domain, rather than in the Euclidean space. Complex geometries, such as cortical surfaces, become reduced embeddings that are isometry-invariant, i.e., two shapes having the same intrinsic geometry, with identical geodesic distances between points, yield identical spectral embeddings, even if they have different extrinsic geometries. This is equivalent to comparing intrinsic vibration properties of shapes rather than their external configurations. However, a perturbation in shape isometry, such as expansion and compression of surfaces, changes the Laplacian eigenvectors in the spectral embeddings. This limited the use of spectral methods to coarse alignment [16] or global analysis [17, 18]. Attempts were made to correct the perturbed embeddings with rigid [19] and nonrigid transformations [20, 14], but all assume that Laplacian eigenvectors are *compatible* between shapes. These methods are in fact inherently flawed by the instability of the eigendecomposition. Eigenvectors can indeed change sign, orientation and shape, due to possible multiplicities of eigenvalues, to ambiguities in shape symmetry, and to numerical instabilities. For instance, the Laplacian eigenvectors of a sphere have ambiguous orientations. This makes spectral embeddings hardly comparable for near spherical shapes. The underlying graph metric may also be adapted by an expensive conformal metric correction [21]. Laplacian eigenvectors provide nonetheless a set of basis functions on complex geometries [22–24], and enable, therefore, any surface function to be represented as a linear com-

Fig. 2. Spectral Transfer between Shapes – a) Reconstruction of a function (here the point coordinates of mesh 1, $f(x) = (x, y, z)$), is correct when using the right eigenbasis of mesh 1. b) Direct transfer of spectral weights, $\mathbf{w}^{(f)}$, onto mesh 2 yields a wrong reconstruction. c) Transfer must account for an optimal spectral transformation R . Color represents sulcal depth on respective meshes. Transfer of $\mathbf{w}^{(f)}$ onto mesh 2 enables a direct comparison between depth maps (a,c).

bination of harmonics. This is for instance exploited in Laplacian smoothing [25, 26], where meshes are reconstructed with the coarser harmonics. The weights associated with each harmonic, capture intrinsic geometrical properties of the represented shape or function. These weights remain, however, incompatible between surfaces.

This paper proposes to ameliorate spectral approaches by enabling a *Spectral Transfer* of such weights across interchangeable harmonic bases. Surface functions, such as point coordinates or activation maps, become transferable from one surface onto another via a *spectral transformation* of the harmonic weights. We indeed assume that if any function can be expressed using a surface basis, there exist a transformation allowing the harmonic weights on one surface to be transferred on another surface [23, 24]. In other words, a set of harmonic weights from one surface must be translated to the same language as on another surface. This requires the optimization of a spectral transformation matrix that combines both, a surface correspondence and a change of harmonic basis. Such spectral transfer provides a more robust formulation for spectral methods and handles naturally the sign changes as well as differences across Laplacian eigenvectors, including the *mixing* of eigenvectors in higher frequencies. In addition, this paper proposes *Focused Harmonics* in order to better capture geometrical properties within a region of interest. This is achieved by building a confidence map with a graph node weighting, which guides the spectral decomposition within regions of higher confidence. Localized functional maps, such as the visual area in the occipital lobe, can therefore be expressed with dedicated harmonics.

Our framework, *Brain Transfer*, enables the transfer of harmonic weights representing cortical shape and functional data across individuals. These harmonic weights capture intrinsic geometrical properties at multiscales, and can be linearly composed in order to reconstruct shapes and functions between subjects. We explore the parameter space of these harmonic weights via a principal component analysis over a dataset on retinotopy. Our results shows that our framework achieves similar accuracy in a fraction of the time as compared with conventional methods when matching cortical surfaces, and outperforms the state-of-the-art when matching functional data within the visual cortex.

Fig. 3. Retinotopy – fMRI activation maps of visual inputs over the visual field: (Left) Polar angle map, varying with $[-\pi/2, +\pi/2]$ (lower/upper field), (Right) Eccentricity map, varying with $[0, \pi]$ (center/peripheral). Brain shows occipital lobe.

2 Spectral Decomposition on Shapes

We begin by reviewing the fundamentals of spectral decomposition on shapes.

Surface Function Decomposition – Let us consider a compact manifold $\mathcal{M}$, and $f(x)$ a smooth function on $\mathcal{M}$. The Laplace-Beltrami operator on f is defined as the divergence of the function gradient, $\Delta f = -\text{div} \nabla f$, and admits an eigendecomposition $\Delta \phi = \lambda \phi$. This can be interpreted as finding the natural vibration amplitudes ϕ , at harmonic frequencies λ , of a membrane with shape $\mathcal{M}$. Since ϕ forms a basis on $\mathcal{M}$, any smooth function f can be represented as a linear combination $f = \sum_{i=0}^{\infty} w_i \phi_i$, where w_i 's are harmonic weights. On a discrete representation, e.g., a triangulated mesh of $\mathcal{M}$, the general Graph Laplacian is often used. Let us build the graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ from the set of vertices, with positions $\mathbf{x} = (x, y, z)$, and the set of edges. We define the $|\mathcal{V}| \times |\mathcal{V}|$ weighted adjacency matrix W in terms of node affinities, e.g., $W_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|^{-2}$, and the diagonal degree matrix D as the sum $D_i = \sum_j W_{ij}$. The General Graph Laplacian is formulated as $\mathcal{L} = G^{-1}(D - W)$ [27], where G is a general node weighting matrix, typically diagonal with $G = D$. The decomposition, $\mathcal{L} = \mathbf{U} \Lambda \mathbf{U}^{-1}$, provides the eigenvalues $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{|\mathcal{V}|})$ and its associated eigenfunctions $\mathbf{U} = (U_0, U_1, \dots, U_{|\mathcal{V}|})$, that correspond respectively to the shape harmonic frequencies and bases on $\mathcal{M}$. Fig. 1 shows how harmonics describe increasingly more complex geometrical properties as frequency augments. A surface function f is represented with:

$$f(x) = \sum_{i=0}^{|\mathcal{V}|} w_i U_i(x), \quad \text{or simply, } \mathbf{f} = \mathbf{U} \mathbf{w}, \quad (1)$$

where $\mathbf{w} = (w_0, w_1, \dots, w_{|\mathcal{V}|})^T = \mathbf{U}^{-1} \mathbf{f}$ is a column vector of harmonic weights. When $\mathbf{U}$ is truncated, we compute the first k coefficients with: $\mathbf{w} = (\mathbf{U}^T \mathbf{U})^{-1} \mathbf{U}^T \mathbf{f}$.

Focused Decomposition – We propose a new focused spectral decomposition in a region of interest by exploiting a graph node weighting as a confidence map. This is of interest to fMRI, where signal is often localized in specific neural areas, such as in the visual cortex for retinotopy (Fig. 3). Node weighting is typically $G = D$, which makes the Graph Laplacian a stochastic matrix. Each node has equal chances of occurring, i.e., row sum of $\mathcal{L}$ is 1. When a confidence map $g(x)$ is used in place, the total node probability is changed with $G = \text{diag}(g(x))$ in $G^{-1}(D - W)$. Nodes are given importance with $g(x)$, and stronger nodes become dominant in the Laplacian

Fig. 4. Focused Spectral Decomposition in Region of Interest – A confidence map indicates regions of confidence of a surface function. (*Left*) Higher confidence is given to the visual cortex, (*Right*) 1st to 10th harmonics. Stronger graph node weights produce focused harmonics, with more description in ROI.

matrix. The spectral decomposition consequently becomes focused on these stronger graph nodes, yielding what we call *Focused Harmonics* (Fig. 4). In our experiment on functional data, the map is set as the signal difference in a common area between subjects, i.e., high signal difference between subjects indicates low confidence areas.

Spectral Smoothing – Functions on surfaces can be represented with general, or focused harmonics, $\mathbf{f} = \mathbf{U}\mathbf{w}$. One application is smoothing of meshes or surface functions. For instance, surface point coordinates $f(x) = (x, y, z)$ can be reconstructed using the first k low-frequency harmonics: $(x, y, z) \leftarrow \mathbf{U}((\mathbf{U}^T\mathbf{U})^{-1}\mathbf{U}^T(x, y, z))$ where the basis is truncated with $\mathbf{U} = (U_0, \dots, U_k)$, and has size $|\mathcal{V}| \times k$. Similarly, surface functions, such as retinotopy, can be smoothed on the surface by reconstructing the original maps with k harmonics. Fig. 5 shows such reconstructions and illustrates how geometrical details appear with more harmonics.

3 Spectral Transfer

Let us now consider the functions $\mathbf{f} = \mathbf{U}^{(f)}\mathbf{w}^{(f)}$ on a mesh $\mathcal{M}^{(f)}$, and $\mathbf{g} = \mathbf{U}^{(g)}\mathbf{w}^{(g)}$ on a mesh $\mathcal{M}^{(g)}$, and see how their intrinsic geometrical properties are transferred across meshes.

Brain Matching – Conventional spectral matching methods rely on the principle that harmonic bases are *compatible* between meshes. A dense point matching is established via fast nearest-neighbor searches between spectral representations, typically $\mathbf{U}\mathbf{A}^{-\frac{1}{2}}$. However, this approach inherently suffers from the instability of eigenvectors, due to a sign or multiplicity ambiguity in a spectral decomposition. Eigenvectors may even be *mixed* in higher frequencies, which yield incompatible bases between meshes. Spectral methods typically compensate for these issues with a rigid or non-rigid transformation of spectral representations, but fail in addressing the fundamental incompatibility of harmonics between meshes. Recently, a double-layered spectral graph decomposition was proposed in [15], where the graph Laplacian is augmented by incorporating in the same Laplacian matrix, the two original meshes with a correspondence map between them. The decomposition of such joint graph Laplacian produces one unique set of harmonics, common to both surfaces. Correspondence links between such joint Harmonics may be interpreted as shortest random-walk distances between surfaces [15].

Brain Transfer – We propose a new approach for brain matching, where intrinsic geometrical properties are transferred across brains instead of being matched.

Fig. 5. Spectral Smoothing – (*Top*) Reconstruction of a surface using 20 to 600 eigenbases. Note how the brain convolves with more eigenbases. (*Bottom*) Reconstruction of retinotopy (polar angle map) using 5 to 200 eigenbases. Few focused harmonics are sufficient to reconstruct the polar map.

This transfer addresses the inherent issues on eigenvector incompatibility in spectral methods. It was shown earlier that intrinsic geometrical properties of shapes or functions are captured in harmonic weights, and that shapes and functions can be reconstructed using such weights. However, harmonic weights are not directly interchangeable between meshes. Fig. 2-b shows an invalid direct transfer with $\mathbf{U}^{(g)}\mathbf{w}^{(f)}$, which does not produce the expected shape of $\mathcal{M}^{(f)}$. Central to our method is the ability to transfer intrinsic geometrical information between brains, via an optimal spectral transformation of harmonic weights. We assume there exists a spectral transformation $R^{(f \rightarrow g)}$ between $\mathbf{U}^{(f)}$ and $\mathbf{U}^{(g)}$ [22–24]. This corresponds to a change of basis $R = (\mathbf{U}^{(g)})^{-1}\mathbf{U}^{(f \circ c)}$, where c is an unknown correspondence map that matches rows of $\mathbf{U}^{(f \circ c)}$ with equivalent rows of $\mathbf{U}^{(g)}$. Such transformation permits the spectral transfer of $\mathbf{w}^{(f)}$ between meshes, such that $\mathbf{U}^{(f)}\mathbf{w}^{(f)}$ becomes equivalent to $\mathbf{U}^{(g)}R\mathbf{w}^{(f)}$, which uses a different basis. In practice, only $k < 20$ coefficients $\mathbf{w} = (w_0, \dots, w_k)^T$ are required in brain matching. The $k \times k$ spectral transformation matrix from $\mathcal{M}^f$ to $\mathcal{M}^g$ is computed with the truncated basis $\mathbf{U}^{(f)}$ and $\mathbf{U}^{(g)}$:

$$R^{(f \rightarrow g)} = \left((\mathbf{U}^{(g)})^T \mathbf{U}^{(g)} \right)^{-1} \left((\mathbf{U}^{(g)})^T \mathbf{U}^{(f \circ c)} \right), \quad (2)$$

such that $\mathbf{U}^{(f \circ c)}\mathbf{w}^{(f)}$ becomes equivalent to $\mathbf{U}^{(g)}R^{(f \rightarrow g)}\mathbf{w}^{(f)}$.

Optimization – The optimal spectral transformation requires solving for the point mapping c from $\mathcal{M}^{(f)}$ to $\mathcal{M}^{(g)}$. In effect, this correspondence minimizes several criteria, such as transferred functions across shapes: $E_u^{(f \rightarrow g)} = \|\mathbf{U}^{(f \circ c)}\mathbf{w}^{(f)} - \mathbf{U}^{(g)}R^{(f \rightarrow g)}\mathbf{w}^{(f)}\|^2$, or surface data, such as a sulcal depth map $\mathbf{s}$ on the surface, and/or its gradient: $E_{\text{data}}^{(f \rightarrow g)} = \|\mathbf{s}^{(f \circ c)} - \mathbf{s}^{(g)}\|^2 + \|\nabla \mathbf{s}^{(f \circ c)} - \nabla \mathbf{s}^{(g)}\|^2$. Symmetry is enforced in adding the inverse mapping, c^{-1} from $\mathcal{M}^{(g)}$ to $\mathcal{M}^{(f)}$, in the energy:

$$E^{(\text{sym})}(f, g, c) = E^{(f \rightarrow g)}(f \circ c, g) + E^{(g \rightarrow f)}(g, f \circ c^{-1}), \text{ where} \quad (3)$$

$$E^{(f \rightarrow g)}(f \circ c, g) = \|\mathbf{U}^{(f \circ c)}\mathbf{w}^{(f)} - \mathbf{U}^{(g)}R^{(f \rightarrow g)}\mathbf{w}^{(f)}\|^2 \quad (4)$$

$$+ \alpha_s \|\mathbf{s}^{(f \circ c)} - \mathbf{s}^{(g)}\|^2 + \alpha_{\nabla s} \|\nabla \mathbf{s}^{(f \circ c)} - \nabla \mathbf{s}^{(g)}\|^2,$$

Fig. 6. Spectral Interpolation – Reconstructions of surfaces using interpolation of harmonic weights, from $p = 0$ (left, subject 1), $p = 0.5$ (middle, synthetic average) to $p = 1$ (right, subject 2): Using a common eigenbasis with $\mathbf{w} = (2 - p)\mathbf{w}^{(1)} + pR\mathbf{w}^{(2)}$. Smoothing occurs by removing harmonic coefficients (600 vs. 150 eigenbases).

where α_s ensures that corresponding points between meshes should have similar sulcal depth maps, while $\alpha_{\nabla s}$ ensures that corresponding sides of sulci should be matched. The harmonic weights are set to $\mathbf{w} = \Lambda^{-\frac{1}{2}}$, which favors eigenfunctions in low, rather than high frequency. Global geometrical characteristics are favored over fine details.

The energies in Eq. 3 can be formulated as the L_2 norm of the difference between high-dimensional vectors: $E^{(f \rightarrow g)} = \|X^{(f \circ c)} - X^{(f \rightarrow g)}\|^2$ and $E^{(g \rightarrow f)} = \|X^{(g \circ c^{-1})} - X^{(g \rightarrow f)}\|^2$, where: $X^{(f)} = (\mathbf{U}^{(f)} \Lambda^{(f)}{}^{-\frac{1}{2}}, \alpha_s \mathbf{s}^{(f)}, \alpha_{\nabla s} \nabla \mathbf{s}^{(f)})$, $X^{(g)} = (\mathbf{U}^{(g)} \Lambda^{(g)}{}^{-\frac{1}{2}}, \alpha_s \mathbf{s}^{(g)}, \alpha_{\nabla s} \nabla \mathbf{s}^{(g)})$, $X^{(f \rightarrow g)} = (\mathbf{U}^{(g)} R^{(f \rightarrow g)} \Lambda^{(f)}{}^{-\frac{1}{2}}, \alpha_s \mathbf{s}^{(g)}, \alpha_{\nabla s} \nabla \mathbf{s}^{(g)})$, and $X^{(g \rightarrow f)} = (\mathbf{U}^{(f)} R^{(g \rightarrow f)} \Lambda^{(g)}{}^{-\frac{1}{2}}, \alpha_s \mathbf{s}^{(f)}, \alpha_{\nabla s} \nabla \mathbf{s}^{(f)})$. Each vector is a combination of intrinsic geometrical information and a data term based on sulcal depth. The data term may also consist of functional data, such as polar and eccentricity maps. The mapping c is established with simple nearest-neighbor searches between vectors $X^{(f)}$ and $X^{(f \rightarrow g)}$, and c^{-1} is established similarly between $X^{(g)}$ and $X^{(g \rightarrow f)}$: $c = \arg \min_c \|X^{(f \circ c)} - X^{(f \rightarrow g)}\|^2$ and $c^{-1} = \arg \min_{c^{-1}} \|X^{(g \circ c^{-1})} - X^{(g \rightarrow f)}\|^2$.

The correspondences may be further refined by decomposing the joint graph Laplacian formed with c and c^{-1} between $\mathcal{M}^{(f)}$ and $\mathcal{M}^{(g)}$, as used in [15]. This may be interpreted as a random walker wandering between surfaces, with probabilities given by the links formed by c and c^{-1} .

The optimization iteratively alternates between 1) a correspondence map phase, which finds the mapping c and c^{-1} and 2) an update phase, which refines the spectral transformation matrices $R^{(f \rightarrow g)}$ and $R^{(g \rightarrow f)}$. The iterations are initialized by using a 3×3 identity matrix for R and a truncated $\mathbf{U}$ with 3 eigenfunctions. The number of eigenfunctions is increased at each iteration, up to 20 in our experiments.

- Initialize $R^{(f \rightarrow g)}$ and $R^{(g \rightarrow f)}$ with identity
- Iterate by increasing $k = 3$ to $k = K$
 - 1) Build vectors $X^{(f)}$, $X^{(f \rightarrow g)}$, and $X^{(g)}$, $X^{(g \rightarrow f)}$ with k eigenfunctions
 - 2) Find mapping c via nearest-neighbor search between $X^{(f)}$ and $X^{(f \rightarrow g)}$
 - 3) Find mapping c^{-1} via nearest-neighbor search between $X^{(g)}$ and $X^{(g \rightarrow f)}$
 - 4) Refine c and c^{-1} with a Joint Graph Laplacian between $\mathcal{M}^{(f)}$ and $\mathcal{M}^{(g)}$
 - 5) Update $R^{(f \rightarrow g)}$ and $R^{(g \rightarrow f)}$ with new c and c^{-1} (Eq. 2)

Spectral Statistics – The harmonic weights $\mathbf{w}$ form a vector space in k dimensions that is spanned by an eigenbasis $\mathbf{U}$. In this context, one may wonder what this space

Fig. 7. Accuracy of Shape Alignment – *Top*: Average Dice Index across 240 alignments of cortical parcellations, using FreeSurfer (dark blue, Dice: 0.84), Spherical Demons (light blue, Dice: 0.85), Spectral Matching (light green, Dice: 0.82), Joint Spectral Matching (orange, Dice: 0.83), and our Brain Transfer (dark red, Dice: 0.84) – *Bottom*: One example showing the alignment of 12 major parcellations (light blue for one cortex, light red for aligned second cortex) – Our method yields similar accuracy as compared with the state-of-the-art, with significant speedup.

looks like. Let us consider two set of transferred harmonic weights for two surface functions: $R^{(f \rightarrow \text{ref})} \mathbf{w}^{(f)}$ and $R^{(g \rightarrow \text{ref})} \mathbf{w}^{(g)}$. We define the linear interpolation between them as $\mathbf{w}^{(p)} = (1-p)R^{(f \rightarrow \text{ref})} \mathbf{w}^{(f)} + pR^{(g \rightarrow \text{ref})} \mathbf{w}^{(g)}$, where $\mathbf{U}^{(\text{ref})}$ may be either $\mathbf{U}^{(f)}$, $\mathbf{U}^{(g)}$, or else. The average weight would be at mid-point $p = 0.5$. Fig. 6 shows the reconstruction of shapes from such interpolated harmonic weights.

The average weights for N surface functions become $\mathbf{w}^{(\mu)} = \frac{1}{N} \sum_{i=1}^N R^{(i \rightarrow \text{ref})} \mathbf{w}^{(i)}$. The covariance matrix of N transferred harmonic weights is $\Sigma = WW^T$, where $W = (\text{vec}(\mathbf{w}^{(1)}), \text{vec}(\mathbf{w}^{(2)}), \dots, \text{vec}(\mathbf{w}^{(N)}))^T - \mathbf{w}^{(\mu)T}$, with the vector operator $\text{vec}(\mathbf{w})$, linearizing $\mathbf{w}$ into a column vector. At last, the decomposition of $\Sigma = \mathbf{Q}\Lambda_{\Sigma}\mathbf{Q}^T$ provides the principal modes of variations [28, 29] of the harmonic weights.

4 Results

We evaluate the performance of Brain Transfer in two folds, first, in aligning cortical surfaces, and second, in aligning retinotopy.

4.1 Alignment using Shape Features

Shape alignment is evaluated with a dataset of 16 cortical surfaces with 109k to 174k vertices, each with a manual labeling of cortical parcellations. We align all 240 possible pairs of cortices using FreeSurfer (FS), Spherical Demons (SD), Spectral Matching (SM), Joint Spectral Matching (JSM), and Brain Transfer (BT). The average Dice overlap ($2|A \cap B|/(|A| + |B|)$) for 12 major parcellations, among all pairs, is 0.84 (±0.08) for FS (timing is 3 hours for one matching), 0.85 (±0.07) for SD (2 hours), 0.82 (±0.08) for SM (110 seconds), 0.83 (±0.08) for JSM (140 seconds), and 0.84 (±0.06) for Brain Transfer (220 seconds). Running times are measured on a 2.6GHz Core i7 laptop with 16GB of RAM. The matching accuracy is arguably

Fig. 8. Alignment of Retinotopy with Brain Transfer – (Top) Polar and (Bottom) Eccentricity map, before and after use of functional features. Surface maps become non-rigidly warped on surface. Transferred harmonic weights enable map comparisons. Matching accuracy increases by 17%, computation time is 220sec. Brain shows occipital lobe.

similar among all methods with an average Dice score of 0.84. Spectral Methods have, however, a notable speed advantage over conventional methods, from hours to minutes of computation.

4.2 Alignment using Functional Features

The advantage of Brain Transfer is better assessed when using focused harmonics on functional data, such as retinotopy on the visual cortex. Our dataset is collected from 9 subjects, each with a polar angle map (θ between $\pm\pi/2$), an eccentricity map (ρ between 0 and π), as well as a manual labeling of 16 visual areas used for evaluation. Data were projected on cortical surfaces of 117k to 164k vertices. We set the graph node weighting with the difference of surface data in the overlapping regions between two shape-aligned cortices: $g(x) = \exp(-\beta((\theta^{(1)} - \theta^{(2)})^2 + (\rho^{(1)} - \rho^{(2)})^2))$, where $\beta = 1/2$.

Since Brain Transfer showed an equivalent accuracy as compared with conventional methods in cortical alignment, we use it in order to compare matching with sulcal features only, and with functional data. This reflects the benefit of using functional data via focused harmonics within the same framework. Fig. 9 shows the average Dice overlap for all 16 visual areas, using Brain Transfer with shape features only: 0.42, and functional features: 0.49. These alignments reveal the challenges of matching visual cortices, whereas large cortical parcellations, based on anatomy, were used in the previous experiment. The use of functional data, via focused harmonics, shows an increase in matching accuracy, up to 54%, notably in areas that are away from the calcarine fissure [30, 12].

4.3 Variability Study of Cortical Shape and Functional Data

Brain Transfer captures intrinsic properties of shapes and functions in harmonic weights. We now provide a Principal Component Analysis (PCA) in order to explore the parameter space spanned by the harmonic weights.

Areas	BT w/Shape	BT w/Function	%Improv.
V1 ventral	0.65	0.70	+8%
V1 dorsal	0.65	0.71	+8%
V2 ventral	0.58	0.67	+18%
V2 dorsal	0.50	0.61	+22%
V3 ventral	0.43	0.57	+31%
V3 dorsal	0.33	0.45	+34%
V3A	0.55	0.62	+12%
V3B	0.38	0.41	+10%
V7	0.45	0.53	+18%
LO1	0.40	0.41	+3%
LO2	0.33	0.33	0%
hV4	0.40	0.54	+37%
VO1	0.23	0.35	+54%
VO2	0.32	0.39	+21%
PHC1	0.22	0.23	+6%
PHC2	0.30	0.27	-8%
Total	0.42	0.49	+17%

Fig. 9. Accuracy of Retinotopy Alignment – Average Dice Index across 72 alignments of 16 visual areas, using Brain Transfer (BT) with pure Shape Features (sulcal depth), with average Dice: 0.42, and Functional Features (polar and eccentricity maps, with focused harmonics), with Dice: 0.49. An example showing the initial and aligned visual areas using Brain Transfer with functional features (green contour for one subject, red for second subject). Brains show occipital lobe.

Shape Variability – We reuse the harmonic weights computed in the former experiment on cortical alignment. PCA is applied on all 16 transferred harmonic weights, each representing surface point coordinates. The target space is chosen arbitrarily as the harmonic basis of the first subject. Our aim is to explore the variability of weights in a common space, regardless of finding an optimal unbiased space. Fig. 10-a shows the reconstruction of cortical surfaces using the average of transferred harmonic weights, and the first 3 principal modes of variations within 2 standard deviations. The first mode appears to capture a global rotation of the cortices, which interestingly, were not initially aligned in space. This indicates that Brain Transfer can correctly handle datasets that are not initially aligned. Further modes capture variabilities in the cortical folds, such as, shape differences in the occipital and temporal lobes.

Functional Variability – We now reuse the harmonic weights computed in the latter experiment on retinotopy alignment. PCA is applied on all 9 transferred harmonic weights, each representing polar and eccentricity maps, expressed in the focused harmonic basis of the first subject. Fig. 10-b shows the reconstruction of polar maps using the average of transferred weights, and the first 3 principal modes of variations. The modes capture specific variabilities in the distinct striped pattern of the polar angle map. The variability study is directly performed on the focused harmonic weights, which captures functional information that is independent from the cortical shape. This contrasts with conventional approaches for studying variabilities where shape and function may be mixed by coupling surface values to point positions. Statistics are therefore performed directly on intrinsic geometrical properties, rather than on pure surface point values.

5 Conclusion

Our method, *Brain Transfer*, contributes to the challenging problem of matching shapes and functional data on cortical surfaces. Spectral methods have demonstrated

Fig. 10. Variability of Harmonic Weights – PCA on harmonic weights, (*Left*) using shape-based weights (function is point coordinates $f(x) = (x, y, z)$) and, (*Right*) using the functional-based weights (function is retinotopy $f(x) = (\theta, \rho)$), with focused harmonics. Showing the first 3 principal modes of variations within $\pm 2\sigma$ from the average.

a tremendous speed advantage over conventional cortical matching methods, but remain inherently limited by the instability of the Laplacian eigenvectors. Incompatible surface harmonics typically require to compensate for the sign flip, ordering, and differences of eigenvectors. Our approach ameliorates spectral methods by enabling a natural transfer of intrinsic geometrical properties across surfaces, where harmonic weights are optimally expressed using different surface harmonics. In other words, intrinsic geometrical information travels across surfaces by translating it into the language of a different surface. Functional data [31] is often localized on the cortical surface, such as retinotopy in the visual area. *Focused harmonics* are, therefore, proposed to better capture surface data in regions of interests. Our results show a 17% improvement in matching visual areas when functional features are used instead of only shape features. *Brain Transfer* also offers a new way to perform statistics across surfaces, by directly analyzing intrinsic geometrical properties of shapes or of surface functions. This analysis on harmonic weights contrasts with conventional approaches where surface values are typically coupled with surface points. Possible extensions to images [32, 33] may be also provide new applications. To conclude, our method, *Brain Transfer*, offers a better formulation for spectral methods by enabling the spectral transfer of intrinsic geometrical properties across surfaces. It also addresses the instability issue of Laplacian eigenvectors. Finally, our improved alignment of functional data may potentially translate into a lower number of required acquisitions in fMRI studies, and still preserve an equivalent statistical power.

Acknowledgment – This research is partially funded by the ERC Advanced Grant MedYMA, and the Research Council of Canada (NSERC).

References

1. Talairach, J., Szikla, G., Tournoux, P., Prosalenti, A., Bordas-Ferrier, M., Covello, L., Iacob, 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 (1996)

4. Thompson, P., Toga, A.W.: A surface-based technique for warping three-dimensional images of the brain. *TMI* (1996)
5. 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* (1999)
6. 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* (2007)
7. Yeo, T., Sabuncu, M., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: fast diffeomorphic landmark-free surface registration. *TMI* (2010)
8. Beg, F., Miller, M., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* (2005)
9. Vaillant, M., Glaunas, J.: Surface matching via currents. In: *IPMI*. (2005)
10. Durrleman, S., Pennec, X., Trounev, A., Ayache, N.: Statistical models of sets of curves and surfaces based on currents. *MedIA* (2009)
11. Segonne, F., Pacheco, J., Fischl, B.: Geometrically accurate Topology-Correction of cortical surfaces using nonseparating loops. *TMI* (2007)
12. Haxby, J.V., Guntupalli, J.S., Connolly, A.C., Halchenko, Y.O., Conroy, B.R., Gobbini, M.I., Hanke, M., Ramadge, P.J.: A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* (2011)
13. Chung, F.: *Spectral Graph Theory*. AMS (1997)
14. Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F.: FOCUSR: Feature Oriented Correspondence using Spectral Regularization, A Method for Accurate Surface Matching. *PAMI* (2012)
15. Lombaert, H., Sporring, J., Siddiqi, K.: Diffeomorphic Spectral Matching of Cortical Surfaces. In: *IPMI*. (2013)
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.: Global Shape Analysis Using the Laplace-Beltrami Spectrum. In: *MICCAI*. (2007)
18. Shi, Y., Lai, R., Kern, K., Sicotte, N., Dinov, I., Toga, A.W.: Harmonic surface mapping with Laplace-Beltrami eigenmaps. In: *MICCAI*. (2008)
19. Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using Laplacian eigenfunctions and unsupervised registration. In: *CVPR*. (2008)
20. Jain, V., Zhang, H.: Robust 3D shape correspondence in the spectral domain. In: *CSMA*. (2006)
21. Shi, Y., Lai, R., Wang, D.J.J., Pelletier, D., Mohr, D., Sicotte, N., Toga, A.W.: Metric optimization for surface analysis in the Laplace-Beltrami embedding space. *TMI* (2014)
22. Vallet, B., Lévy, B.: Spectral Geometry Processing with Manifold Harmonics. *CG Forum* (2008)
23. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps. *ACM Transactions on Graphics* (2012)
24. Kovnatsky, a., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. *CG Forum* (2013)
25. Chung, M.K., Robbins, S.M., Dalton, K.M., Davidson, R.J., Alexander, A.L., Evans, A.C.: Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage* (2005)
26. Anqi, Q., Bitouk, D., Miller, M.I.: Smooth functional and structural maps on the neocortex via orthonormal bases of the Laplace-Beltrami operator. *TMI* (2006)
27. Grady, L., Polimeni, J.R.: *Discrete Calculus: Analysis on Graphs*. Springer (2010)
28. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J.J., Shenton, M.E., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *Insight* (2006)
29. Gao, Y., Riklin-Raviv, T., Bouix, S.: Shape analysis, a field in need of careful validation. *Human Brain Mapping* (2014)
30. Wang, L., Mruczek, R.E.B., Arcaro, M.J., Kastner, S.: Probabilistic maps of visual topography in human cortex. *Cerebral Cortex* (2014)
31. Sabuncu, M.R., Singer, B.D., Conroy, B., Bryan, R.E., Ramadge, P.J., Haxby, J.V.: Function-based intersubject alignment of human cortical anatomy. *Cereb Cortex* (2010)
32. Lombaert, H., Grady, L., Pennec, X., Ayache, N., Cheriet, F.: Spectral Demons - Image Registration via Global Spectral Correspondence. In: *ECCV*. (2012)
33. Lombaert, H., Grady, L., Pennec, X., Ayache, N., Cheriet, F.: Spectral Log-Demons: Diffeomorphic Image Registration with Very Large Deformations. *IJCV* (2013)