

FOCUSR: Feature Oriented Correspondence using Spectral Regularization – A Method for Precise Surface Matching

Herve Lombaert, Leo Grady, Jonathan R. Polimeni, Farida Cheriet

Abstract—Existing methods for surface matching are limited by the trade-off between precision and computational efficiency. Here we present an improved algorithm for dense vertex-to-vertex correspondence that uses direct matching of features defined on a surface and improves it by using spectral correspondence as a *regularization*. This algorithm has the speed of both feature matching and spectral matching while exhibiting greatly improved precision (distance errors of 1.4%). The method, FOCUSR, incorporates implicitly such additional features to calculate the correspondence and relies on the smoothness of the lowest-frequency harmonics of a graph Laplacian to spatially regularize the features. In its simplest form, FOCUSR is an improved spectral correspondence method that nonrigidly deforms spectral embeddings. We provide here a full realization of spectral correspondence where virtually *any* feature can be used as additional information using weights on graph edges, but also on graph nodes and as extra embedded coordinates. As an example, the full power of FOCUSR is demonstrated in a real case scenario with the challenging task of brain surface matching across several individuals. Our results show that combining features and regularizing them in a spectral embedding greatly improves the matching precision (to a sub-millimeter level) while performing at much greater speed than existing methods.

Index Terms—Registration, Surface fitting, Spectral methods, Graph Theory



1 INTRODUCTION

MESH correspondence is a key step in many applications of computer vision and whose precision and speed are crucial. It is at the core of studies on shape variability and on object and motion analysis. In the medical field, precision is essential and a fast method enables investigations on large studies between organs or individuals. The challenge of shape matching is to find the dense correspondences mapping all points on one surface to their equivalent points on a second surface. This task becomes particularly arduous when the matching involves highly convoluted surfaces or two surfaces representing different poses of an articulated object. Early solutions [8] to this problem, aligning surface models, were limited to rigid transformations [47], or relied on fiducial markers placed on the surfaces [5], [48]. Methods based on deformable surfaces [53], [30], [72], [67] could find nonrigid transformations. Such iterative deformation of surfaces, using minimal distortion [10], is also an adequate strategy for matching partial data [14], even with a change in topology [67], [63]. However, to keep these approaches tractable, prior knowledge on the underlying deformation between surfaces [35], [9], or the use of control or feature points [1], [61], [34], [52], [43], [64], is often required to restrain the search domain. Rather than optimizing for a deformation, other approaches would directly solve for the correspondence map [2], avoiding iterative deformations of the surfaces. Moreover, these surfaces may be correlated with measurable *features* other than their explicit mesh geometry. For example, the method used in FreeSurfer [25], a leading tool for brain surface reconstruction and

matching, uses geometric features such as cortical curvature and sulcal depth (the depth in the cortical folding pattern) [38] to drive the warping of one brain surface into another surface. However, despite its precision, FreeSurfer suffers from a substantial computational burden, taking hours to compute a correspondence map between typical brain surface models consisting of hundreds of thousands of vertices. Needless to say, the incorporation of additional features is a convenient aspect for a matching algorithm.

A direct method of matching two surfaces based on features (e.g., the geometry of the cerebral cortex in brain matching, or texture intensities for articulated object matching) is to treat the available features as characteristic *signatures* which can be used to identify each vertex within the surface mesh. With these signatures, a vertex on one surface could be mapped to the vertex on a second surface which most closely resembles the same characteristic features (e.g., by computing a Euclidean distance between the feature vectors). This feature matching technique would have the merits of being fast (e.g., computable within Voronoi cells) and flexible enough to allow any set of features to drive the matching. Unfortunately, this feature matching technique would completely ignore the spatial organization of the surface vertices and result in a highly non-smooth mapping between the surfaces. Our approach to the matching problem seeks to preserve the speed and flexibility of direct feature matching and address the problem of smooth mapping by using an improved spectral correspondence as a *regularization*.

Spectral correspondence [19] utilizes a graph (mesh) *spectrum*, which is the set of Laplacian eigenvalues and eigenvectors (illustrated on Fig. 1), to produce a vertex correspondence between two graphs (meshes). The key utility of spectral correspondence in our context is to provide a spatial regularization on the correspondence map. This regularization is enabled by the fact that the low-frequency eigenvectors (those corresponding to small

H. Lombaert is with McGill University, Canada.

L. Grady is with HeartFlow, Redwood, CA.

J.R. Polimeni is with the Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Harvard Medical School, Massachusetts General Hospital, Charlestown, MA.

F. Cheriet is with École Polytechnique de Montréal, Canada.

eigenvalues) are spatially smooth, as they represent low-frequency harmonics [28]. Put differently, all neighboring nodes will have a small change in the values of these harmonics, meaning that a correspondence driven by these harmonics will map neighboring nodes to neighboring locations in the range space. The value of the harmonics at each vertex are known as the *spectral coordinates* of the vertex. At its core, our technique for spectral regularization is to supplement the direct feature matching technique described above by extending the vertex *signature* to additionally include the *spectral coordinates* of each vertex. Matching vertices are revealed with similarities in such extended signatures. Fortunately, including the spectral coordinates in our matching maintains the speed and flexibility of the simple technique. We call our method FOCUSR for Feature Oriented Correspondence Using Spectral Regularization.

Spectral methods have been used in many fields, such as in the segmentation and registration of shapes in images [45], in the indexing of structures [57], or in the clustering of data [54], [7], [6]. Their use in shape matching is often limited to hierarchical matching (e.g., matchings of limbs in body models, or of large surface areas). Few medical applications of spectral methods exist and are targeted to brain studies in order to study the geometrical patterns of the anatomical structures of the brain such as the cortical folds [42], [46], [55], [56] and with the smoothing of cortical surfaces [3]. Spectral coordinates have also been used directly for graph partitioning [17]. Umeyama [65] and later Scott and Longuet-Higgins [50] pioneered the use of spectral methods for the correspondence problem. Shapiro and Brady [51] compared ordered eigenvectors of a proximity matrix to find correspondences. Their work served as a basis for future spectral correspondence methods. Variants include the use of different proximity matrices using different kernels, the use of the adjacency matrix, shock graphs [44], [58], different normalized Laplacian matrices, or the use of Multi-Dimensional Scaling [49], [23], [68], [10], [11]. Closely related are also methods based on the Heat Kernel [62], [43] which use multiscale descriptions of intrinsic shape properties, or on other conformal maps, for instance, the Ricci flow [71] (deforming the Riemannian metric underlying the original shape toward another conformal metric) or the Wave Kernel Signature [4] (using mechanical properties as an intrinsic metric). The blending of various conformal maps has also been shown to improve matching accuracy [34]. Learning of local shape descriptors may also be an alternative strategy [16]. Recent surveys covering the use of spectral methods in the past fifteen years are available in [66], [73]. Mateus *et al.* [40] proposed an original unsupervised spectral method with an alternative to eigenvalue ordering based on eigenvector histograms and refining the eigenvectors alignment with a probabilistic point matching within the framework of the EM algorithm [15]. Jain and Zhang [33] approach the eigenvector alignment problem with a nonrigid deformation based on Thin Plate Splines.

Spectral correspondence has presented several difficulties that act as a barrier to its widespread adoption. Specifically, when computing the eigenvectors for two surfaces, the signs of the eigenvectors need to be aligned (the eigenvectors are ambiguous to sign), the eigenvectors

sometimes require reordering (due to near algebraic multiplicity of the eigenvalues causing ordering changes of the spectral coordinates). Additionally, spectral matching methods typically start with a rigid alignment of the eigenvectors to account for translation and scaling of the spectral coordinates. Small variations however exist in the spectral coordinates (due to non perfect shape isometry, e.g., local expansion and compression within meshes). There is therefore a need for robust nonrigid point correspondence between spectral coordinates. Furthermore, the use of vertex features has not been fully realized in previous work on spectral correspondence, which have incorporated these features only to produce edge weights (measuring changes between neighboring features) rather than as node weights (using the features themselves). We address and improve all these aspects of spectral correspondence while additionally using the spectral coordinates to provide a smooth regularization of the simple feature matching technique. This work makes several contributions to dense surface matching:

- Extending simple surface feature matching with spectral regularization. The integration of spectral components in extended signatures alongside feature characteristics provides a natural means of regularization (i.e., matching extended signatures reveal matching vertices).
- Nonrigid alignment of the multidimensional embeddings (i.e., of the extended signatures, rather than only the spectral coordinates).
- The weighting of *nodes* in the graph Laplacian, which controls the influence of each node during correspondence.
- A global approach to handle automatically the sign ambiguity and the rearrangement of the graph Laplacian eigenvectors.

After detailing FOCUSR in the next section, we show in controlled experiments that it outperforms both direct feature matching and conventional spectral correspondence. Firstly, we demonstrate that nonrigid alignment of the spectral coordinates improves the direct matching method. We also show that FOCUSR can be used as a general method for mesh matching by benchmarking our method on generic meshes of animals or humans in various poses, faces having different expressions, and on the widely used TOSCA dataset. The computed correspondence maps have in fact a negligible error from ground truth. Secondly, we show the full power of FOCUSR with the use of additional features and assess its precision with the challenging task of brain surface matching. Indeed, while the sulcal and gyral folding pattern of the human cerebral cortex are somewhat stable across individuals, some geometric variability does exist [32], making the direct use of the folding geometry unsuitable for surface matching. This application to the problem of brain matching provides a platform for FOCUSR where the use of additional features available in the brain data—such as cortical Gaussian curvature, sulcal depth, and cortical thickness—can improve the matching precision. We show that FOCUSR produces results in a fraction of the time required by FreeSurfer while maintaining the same level of precision. We believe that this large gain in processing speed would make possible new brain studies

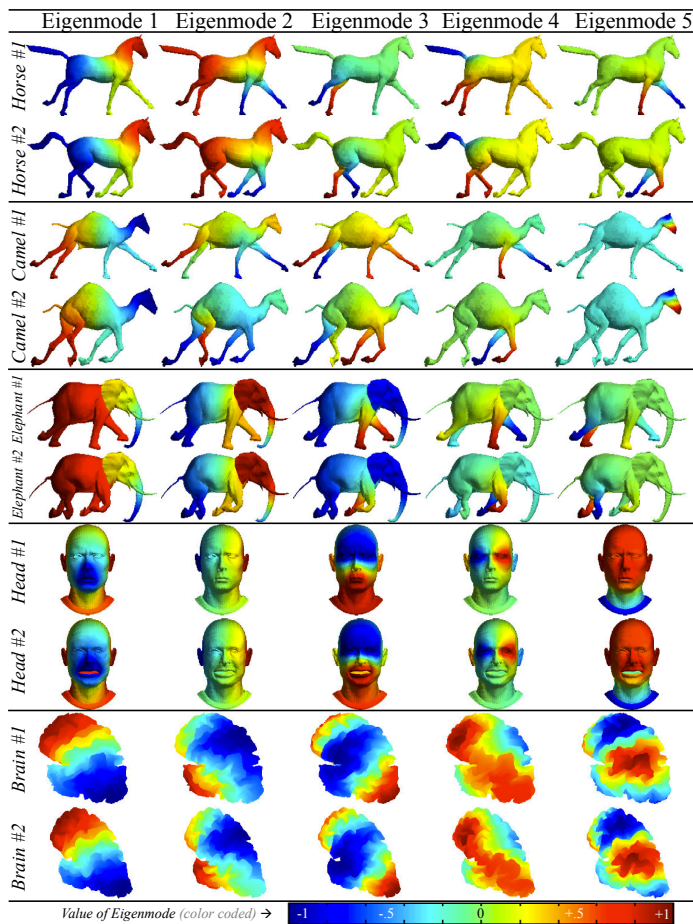


Fig. 1. Example of eigenmodes for pairs of animals and human brain surfaces. Each row shows the first five spectral components of a model (eigenmodes of the associated graph Laplacian, reordered and sign adjusted, so paired sets match). The color scale indicates the value of the spectral coordinate over the surface.

that were previously limited by computational burden, or, more generally, studies on meshes that wish to use non-standard features for driving the correspondence.

2 METHODS

We begin our exposition of FOCUSR by detailing a simple technique for feature matching that does not preserve smoothness of the mapping between surfaces (Fig. 2 a). We then describe how spectral coordinates can be used to regularize feature matching. We re-examine and improve each step in the spectral correspondence process to overcome previous limitations with spectral correspondence. The algorithm is summarized in Fig. 3. Code implementation in Matlab is available at <http://step.polymtl.ca/~rv101/focusr>¹.

2.1 Direct feature matching to provide vertex correspondence

Assume that we have two graphs, $\mathcal{G}_1 = \{\mathcal{V}_1, \mathcal{E}_1\}$ and $\mathcal{G}_2 = \{\mathcal{V}_2, \mathcal{E}_2\}$ (with vertices and edges) such that a correspondence $\phi: v_i \in \mathcal{V}_1 \rightarrow v_j \in \mathcal{V}_2$ is desired. Note that we

do not require that $|\mathcal{V}_1| = |\mathcal{V}_2|$ or $|\mathcal{E}_1| = |\mathcal{E}_2|$ (i.e., meshes can have different sizes and structures). Consequently, there is no guarantee that the mapping is one-to-one and may not be invertible. We will use the terms *node*, *vertex* and *point* interchangeably to describe a member of \mathcal{V}_1 or \mathcal{V}_2 . Given a set of K features \mathbf{X}_i at every node $v_i \in \mathcal{V}_1$, and a set of K features \mathbf{Y}_j at every node $v_j \in \mathcal{V}_2$, our goal is to use these features to produce a correspondence ϕ .

A direct feature matching approach to producing this correspondence would be to set

$$\phi(v_i) = \min_{v_j \in \mathcal{V}_2} \|\mathbf{X}_i - \mathbf{Y}_j\|, \quad (1)$$

which could be computed quickly by precomputing a Voronoi tessellation of the range space. Unfortunately, this simple technique has several inadequacies. Specifically, the technique based on the Voronoi tessellation does not properly account for global changes in the feature space (e.g., due to a global scaling or translation), nor does it utilize the neighborhood structure provided by the edge sets in any way (i.e., there is no spatial regularity to the mapping in the sense that neighbors in the domain are unlikely to remain neighbors in the range).

Global changes in the feature space can be accounted for by using a more sophisticated point correspondence than what is described in Eq. (1). Robust Point Matching [18] with a Thin Plate Spline-based transformation is often used for 2D or 3D registration. However, with this approach the final registration depends on the number and choice of the control points. A more recent approach to the point correspondence problem is the Coherent Point Drift (CPD) method [41] which is fast and demonstrates excellent performance. To summarize this method, the registration is treated as a Maximum Likelihood problem where Gaussian Mixture Model centroids are fit into a point set. There is no assumption on the global transformation between point sets. Instead, the evolution of the transformation is constrained with a motion coherence [41]. The CPD algorithm offers the possibility to perform matching on a subset of the points (for increased speed) while computing the transformation in the continuous domain (i.e., the continuous transformation, found with only a subsample of \mathcal{V}_1 and \mathcal{V}_2 , can be applied on all points of \mathcal{V}_2 and thus find a dense matching between \mathcal{V}_1 and \mathcal{V}_2). Furthermore, each feature (i.e., each coordinate of \mathbf{X}_i or \mathbf{Y}_j) can be weighted in order to accentuate or reduce its influence.

Although CPD provides a method to account for global transformation in the feature space between the two graphs, it is still necessary to incorporate spatial regularity into the mapping such that neighboring points in \mathcal{V}_1 map to neighboring points in \mathcal{V}_2 . Note that a strict neighbor-to-neighbor mapping is only possible when the two graphs are isomorphic. Since we target a more general scenario, we want to account for neighborhood relationships by promoting a correspondence that maps nearby nodes in \mathcal{V}_1 (based on \mathcal{E}_1) to nearby nodes in \mathcal{V}_2 (based on \mathcal{E}_2). Our strategy for promoting spatial regularization is to supplement \mathbf{X}_i and \mathbf{Y}_j with the *spectral coordinates* at nodes v_i and v_j before applying the CPD point correspondence. The values of the spectral coordinates over a few sample surfaces are illustrated on Fig. 1. The fundamental

1. This private link is going public after publication

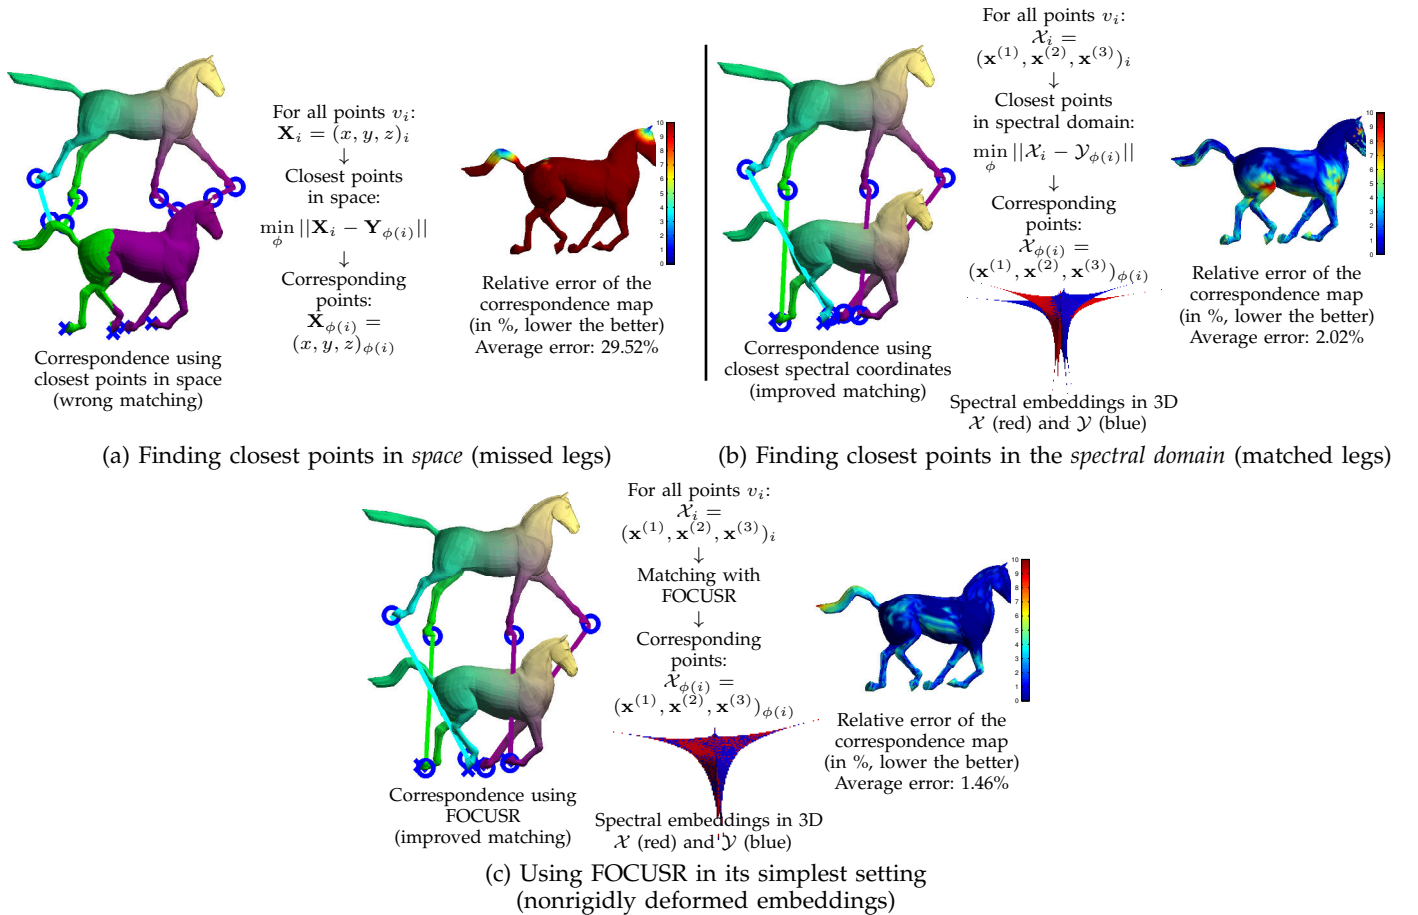


Fig. 2. *Direct matching* (coloring indicates correspondence, and links and circles indicate matching of leg extremities, crosses indicate ground truth) : (a) Finding closest points in *space*: this naive correspondence map is computed by finding for each point of model X its closest point in *space* of model Y (match X with Y). As illustrated, this strategy is dependent on rigid and nonrigid deformations and generates an inconsistent correspondence map. (b) Finding closest points in the *spectral domain*: the correspondence map is computed by finding for each point of model X its closest *spectral* equivalent in model Y (match \mathcal{X} with \mathcal{Y} instead of X with Y). As illustrated, even though the meshes are not aligned in space (they are translated), their spectral embeddings (red is \mathcal{X} , blue is \mathcal{Y} , both use three eigenmodes for 3D visualization) are almost perfectly superimposed. Spectral embeddings are much less dependent on rigid and nonrigid deformations, and finding closest points in the spectral domain generates a much better correspondence map (2.02% average error). *FOCUSR in its simplest setting* : (c) Our method performs matching in the *spectral domain* (with lower error over the surface) and improves the alignment of the spectral embeddings. Note that no additional features are used here in FOCUSR.

difference between the use of \mathbf{X} and \mathbf{Y} as general feature vectors (illustrated with 3D coordinates (x, y, z)) and as spectral coordinates is demonstrated in Fig. 2 b. The low-frequency spectral coordinates are dependent on the geometry of the surface, and these coordinates are effectively more stable across articulated shapes or highly deformable shapes, i.e., normalizing these shapes in a same referential. Additionally, they are known to be spatially smooth (see below) in accordance with the low-frequency harmonics of an elastic surface [28]. In the next section we will review spectral coordinates, and demonstrate improvements to traditional methods for solving some of the difficulties associated with comparing spectral coordinates from two graphs.

2.2 Spectral Coordinates

We may define the $|\mathcal{V}| \times |\mathcal{V}|$ adjacency matrix W of a graph in terms of *affinity weights* (see [28]), which are derived from a given distance metric $\text{dist}(i, j)$ between two neighboring vertices (v_i, v_j) . The elements of the weighted adjacency matrix are given by

$$W_{ij} = \begin{cases} 1/\text{dist}(i, j) & \text{if } \exists e_{ij} \in \mathcal{E}, \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The matrix W provides a weighting on the graph edges derived from the given distance metric. The distance may be derived from the geometry via the vertex coordinates $\mathbf{x} = (x, y, z)^T$ embedded in space (e.g., $\text{dist}(i, j) = \|\mathbf{x}_i - \mathbf{x}_j\|$, the distance between nodes v_i and v_j), from feature vectors (e.g., $\text{dist}(i, j) = \|\mathbf{F}_i - \mathbf{F}_j\|$, where $\mathbf{F} = (\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(K)})^T$ for K features), or both. The more general edge weighting

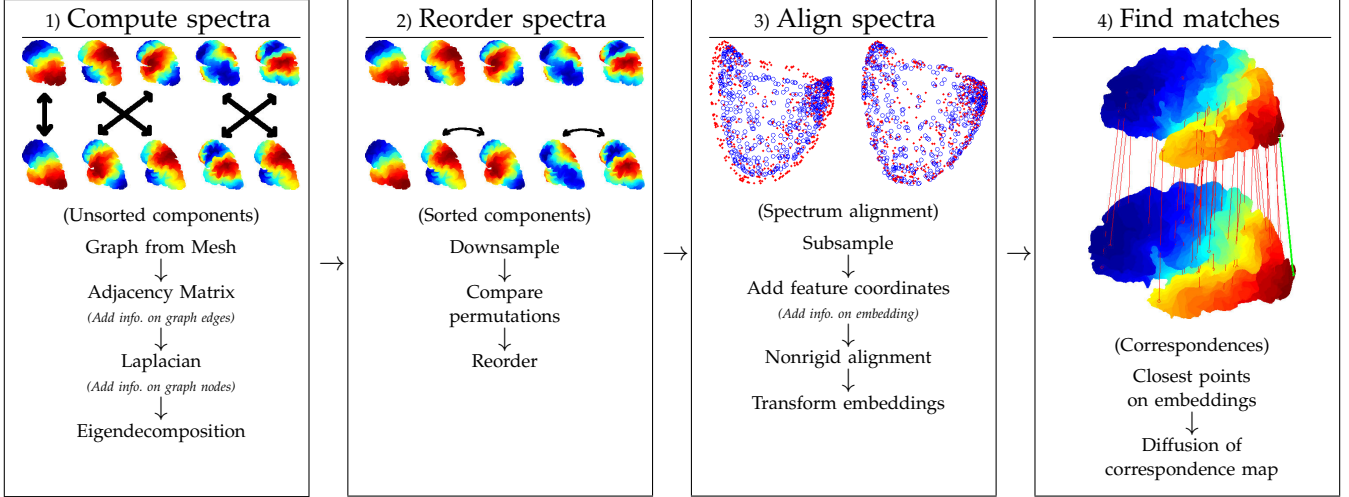


Fig. 3. FOCUSR overview for matching a pair of surfaces. *First*, we build a graph out of each surface mesh and set the graph edge weights and graph node weights to construct the Laplacian matrix. The eigendecomposition of each graph’s Laplacian matrix reveals its spectral components. *Second*, we reorder the spectral components by finding the optimal permutation of components between the pair of meshes. *Third*, regularization is performed by matching the spectral embeddings. *Finally*, corresponding points are found with closest points in both spectral embeddings, and the final correspondence map is diffused.

between vertices v_i and v_j uses the ℓ_2 norm between extended vectors:

$$w_{i,j} = \text{dist}(i, j) = \|(\mathbf{x}_i, \gamma \mathbf{F}_i) - (\mathbf{x}_j, \gamma \mathbf{F}_j)\|, \quad (3)$$

where $(\mathbf{x}, \gamma \mathbf{F})$ is the concatenation of the 3D coordinate values $\mathbf{x} = (x, y, z)^\top$ with the K feature values $\mathbf{F} = (\mathbf{f}^{(1)}, \dots, \mathbf{f}^{(K)})^\top$. The $K \times K$ diagonal matrix γ contains the K weights controlling the influence of each feature. To compensate for the different scalings of the feature values, each feature vector $\mathbf{f}^{(k)}$ is normalized with respect to the range of the 3D coordinate values \mathbf{x} (i.e., feature values are normalized such that $\min(\mathbf{f}^{(k)}) = \min(\mathbf{x})$ and $\max(\mathbf{f}^{(k)}) = \max(\mathbf{x})$).

The general Laplacian operator on a graph was formulated in [28] as a $|\mathcal{V}| \times |\mathcal{V}|$ matrix with the form:

$$\mathcal{L} = G^{-1}(D - W), \quad (4)$$

where D , the degree matrix, is a diagonal matrix defined as $D_{ii} = \sum_j W_{ij}$, and G is the diagonal matrix of node weights. Typically in spectral correspondence G is set to identity $G = I$, or to $G = D$. However, we propose here to replace the default assignment $G = D$ with *any* meaningful node weighting. In particular, we propose to use a function of feature *magnitudes* to establish the (positive-valued) node weighting based on the assumption that nodes with significant features are of more interest to match precisely (i.e., nodes with large weight have a greater influence on the spectral correspondence than low-weight nodes). For example, if half of the nodes in a graph had a large weight and the other half had a small weight, the Laplacian eigenvectors would closely resemble the eigenvectors of the large-weight subgraph. The diagonal of matrix G contains the general node weights for each vertex v_i :

$$w_i = G_{ii} = d_i \sum_{k=1}^K \gamma_i \rho \left(f_i^{(k)} \right), \quad (5)$$

where d_i is the node degree (i.e., D_{ii}), γ is the previously mentioned feature weights, and $\rho(\cdot)$ is a function that enforces positive values (e.g., $\rho(f) = f^2$ or $\rho(f) = \exp(f)$). The denominator in Eq. (5) contains the sum of the influences of each feature on vertex v_i . We used $\rho(f) = \exp(f)$ to promote correspondence between nodes having the largest feature components (which we assume indicate greatest significance).

The right eigenvectors of the Laplacian matrix comprise the graph spectrum $\mathcal{X} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)})^\top$, where $n = |\mathcal{V}|$ is the number of nodes. The values over surfaces for the five lowest frequency eigenvectors are shown on Fig. 1, and illustrates the stability of these eigenvectors between articulated or highly deformable shapes. Each eigenvector² $\mathbf{x}^{(u)}$ is a column matrix with n values, and represents a different (weighted) harmonic on a mesh surface that corresponds to an inherent property of the mesh geometry. This is in comparison with extrinsic properties such as the spatial location of points (i.e., point coordinates vary when the model takes a different pose). The n values $(x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(n)})$ give the *spectral coordinates* of node v_i (i.e., a coordinate in a *spectral domain*). The first eigenvector $\mathbf{x}^{(1)}$ is the trivial (uniform) eigenvector, and the eigenvectors associated with the lower non-zero eigenvalues (e.g., $\mathbf{x}^{(2)}, \mathbf{x}^{(3)}$) represent coarse (i.e., low-frequency) intrinsic geometric properties of the shape. The first of them $\mathbf{x}^{(2)}$ is called the *Fiedler vector* [19], while eigenvectors associated with higher eigenvalues (e.g., $\mathbf{x}^{(n-1)}, \mathbf{x}^{(n)}$) represent fine

2. In our notation \mathbf{x} represents the 3D coordinate in space (i.e., x, y, z), and the superscripted $\mathbf{x}^{(i)}$ represents the i^{th} spectral coordinate (i.e., the i^{th} eigenvector of the graph Laplacian).

(high-frequency) geometric properties. For example, in Fig. 1, the values of $\mathbf{x}^{(2)}$ increase along a virtual centerline depicting the global shape of the models (a coarse intrinsic property), while the values of $\mathbf{x}^{(5)}$ depict finer details of the models.

To illustrate why spectral coordinates corresponding to small eigenvalues transition smoothly and slowly across neighboring nodes, consider the Rayleigh quotient

$$\lambda = \frac{\mathbf{x}^\top L \mathbf{x}}{\mathbf{x}^\top G^{-1} \mathbf{x}} = \frac{\sum_{e_{ij} \in \mathcal{E}} w_{ij} (x_i - x_j)^2}{\sum_{v_i \in \mathcal{V}} \frac{1}{w_i} x_i^2}. \quad (6)$$

The minimum value of λ is the smallest eigenvalue for L . If the minimization of λ over x is conducted in the space orthogonal to the eigenvector corresponding to the smallest eigenvalue, then the minimum λ is the second smallest eigenvalue (the Fiedler value, in our case). Put differently, all of the eigenvectors corresponding to the smallest eigenvalues have small values of the Rayleigh quotient in Eq. (6). Examining the numerator of Eq. (6), we see that neighboring nodes must have a small change in the spectral coordinate (eigenvector) x in order for the corresponding eigenvalue to be small. However, a small edge weight indicates that the change in x across that edge may be large while still maintaining a small numerator (and therefore a small eigenvalue). Consequently, the edge weights act to enforce a smoother change between similar neighbors, but the spatial regularization is more relaxed for mapping neighboring points which are dissimilar. Examining the denominator of Eq. (6), we see that large node weights have the effect of reducing the influence of the node in the denominator, effectively pushing the node to take a value that minimizes the numerator (i.e., the average of its neighbors). By pushing the node to minimize the numerator, the large node weight effectively promotes maximal smoothness in the spectral coordinates at that node.

Consequently, we use the node features to enforce more spatial regularity between similar neighboring nodes (large edge weight) and to enforce more spatial regularity at unremarkable nodes (nodes with small feature magnitude and small node weight). Ultimately, this use of the node features to promote variable spatial regularization is designed to enforce a stronger correspondence between key nodes (nodes with large feature magnitude) and to enforce stronger spatial regularity between key nodes. In this manner, the key nodes (which are similar in the two meshes) are matched strongly, while the remaining nodes are matched to promote maximal spatial regularity.

2.3 Spectrum Ordering

Each node is represented with $M \ll |\mathcal{V}|$ spectral coordinates associated with the M smallest (non-trivial) eigenvalues, i.e., the embedded representations for meshes X and Y are $\mathcal{X}^M = (\mathbf{x}^{(2)}, \dots, \mathbf{x}^{(M+1)})^\top$ and $\mathcal{Y}^M = (\mathbf{y}^{(2)}, \dots, \mathbf{y}^{(M+1)})^\top$. Unfortunately, the spectral coordinates of the two meshes may not be directly comparable as a result of two phenomena. First, there exists a *sign ambiguity* when computing eigenvectors, i.e., if $Ax = \lambda x$ (the spectral decomposition of A) then $A(-x) = \lambda(-x)$, which requires checking that each corresponding eigenvector in the two meshes has the same sign. Additionally, as a result

of greater algebraic multiplicity of an eigenvalue, it may be possible that the *ordering* of the lowest eigenvectors will change, e.g., if two eigenvectors correspond to the same eigenvalue in both meshes, then the solver may compute these eigenvectors in one order for the first mesh and in the opposite order for the second mesh. A graph with an eigenvalue having algebraic multiplicity greater than one indicates symmetry in the mesh. For large meshes, symmetries (and near symmetries) is a common problem and the eigenvectors must be reordered.

Our approach to the eigenvector reordering is to favor pairs of eigenvectors that are most likely to match based on the similarity of their eigenvalues, histograms, and spatial distributions of their spectral coordinate values. The costs of pairing the u^{th} eigenvectors, $\mathbf{x}^{(u)}$, of mesh X with the v^{th} eigenvectors, $\mathbf{y}^{(v)}$, of mesh Y are gathered in a $M \times M$ dissimilarity matrix C . It consists of three terms:

$$C(u, v) = c_\lambda(u, v) \cdot c_{\text{hist}}(u, v) \cdot c_{\text{spatial}}.$$

The first term penalizes pairs of eigenvectors whose eigenvalues are far distant, $c_\lambda(u, v) = \exp(-(\lambda^{(u)} - \lambda^{(v)})^2 / 2\sigma^2)$ with a kernel width σ that depends on the average eigen-gap $\frac{1}{K-1} \sum_{k=1}^{K-1} \lambda^{(k+1)} - \lambda^{(k)}$. The second term $c_{\text{hist}}(u, v)$ favors pairing of eigenvectors that have similar histograms of spectral coordinate values. To ease comparison, spectral coordinates are normalized to the range $[0, 1]$. Pairing based on histograms is also used in [40], however, we found that using a logarithmic scale produces the best results (it minimizes the effect of over represented spectral coordinates values, such as those close to zero), such as in $c_{\text{hist}}(u, v) = |\log(\text{hist}(u)) - \log(\text{hist}(v))|$. The third term verifies the spatial coherence of the spectral coordinate values between two meshes. To speed up the reordering, all eigenvectors are subsampled by randomly selecting a subset of $N < |\mathcal{V}|$ nodes (we used 500 nodes or about 0.4% of the vertices in our experiments). The pairs of closest points within these subsampled points determine the correspondence map μ (i.e., vertex $v_i \in \mathcal{V}_1$, on the first mesh, is closest to point $v_{j=\mu(i)} \in \mathcal{V}_2$, on the second mesh). Then, we simply compute for all corresponding points ($v_i \in \mathcal{V}_1 \rightarrow v_{j=\mu(i)} \in \mathcal{V}_2$) the squared difference between the coordinate values $x_i^{(u)}$ and $y_{\mu(i)}^{(v)}$ and $c_{\text{spatial}} = \sum_{i=1}^N (x_i^{(u)} - y_{\mu(i)}^{(v)})^2$.

The Hungarian algorithm may be used to find an optimal permutation of eigenvectors $\mathbf{y}^{(v)}$ that minimizes dissimilarity. In the same step we can remove the sign ambiguity by calculating the minimal dissimilarity between all $\mathbf{x}^{(u)}$ and $\mathbf{y}^{(v)}$, as well as between all $\mathbf{x}^{(u)}$ and $-\mathbf{y}^{(v)}$. The cost matrix used in the Hungarian algorithm is thus $Q(u, v) = \min\{C(u, v), C(u, -v)\}$. After permutation π , any eigenvector $\mathbf{x}^{(u)}$ corresponds with $\mathbf{y}^{(\pi(u))}$, and its permutation cost $Q^{(u)}$ is stored for use in the spectral alignment.

To keep the notation simple, in the next sections we assume that the spectral coordinates have been appropriately reordered and signed (i.e., \mathcal{X}^M and $\mathcal{Y}^{\pi \circ M}$ will simply be denoted as \mathcal{X}^M and \mathcal{Y}^M such that $\mathbf{x}^{(u)}$, on the first mesh, corresponds with $\mathbf{y}^{(u)}$, on the second mesh).




Fig. 4. Nonrigid alignment of the two spectra corresponding to two brain surfaces. For visualization purposes, the first three eigenvectors ($\mathbf{x}^{(2)}$, $\mathbf{x}^{(3)}$, and $\mathbf{x}^{(4)}$) are used as 3D coordinates (x, y, z). Red and blue are the control points used to align both spectra. Initial spectra (a) before and (b) after final alignment.

2.4 Nonrigid Spectral Alignment

Once the reordering and sign adjustment of the eigenvectors have taken place, finding the closest points in the *spectral* domain between embeddings \mathcal{X} and \mathcal{Y} generates a smooth correspondence map (Fig. 2). However, these embedded representations contain slight differences, mostly due to perturbations of the shape isometries such as small changes in distances where the surface undergoes local expansion or compression between meshes. As illustrated on Fig. 4, nonrigid differences in the spectral embeddings become even more severe in highly convoluted surfaces such as brain cortices. Spectral representations need to be nonrigidly aligned.

Closest points in these nonrigidly aligned embedded representations would reveal corresponding points in both shapes (i.e., in the M -dimensional space (the *spectral* domain), if the point $v_i \in \mathcal{Y}_1$ with coordinates \mathcal{X}_i^M , is the closest point to $v_j \in \mathcal{Y}_2$ with coordinates \mathcal{Y}_j^M , then v_i corresponds to v_j). It is at this point where Eq. (1) is extended by combining the spectral coordinates, \mathcal{X}^M and \mathcal{Y}^M , with the feature vectors, $\mathbf{F}_x = (\mathbf{f}_x^{(1)}, \dots, \mathbf{f}_x^{(K)})^\top$ for nodes in model X , and $\mathbf{F}_y = (\mathbf{f}_y^{(1)}, \dots, \mathbf{f}_y^{(K)})^\top$ for nodes in model Y , to enable spatial regularization in the correspondence map. The extended vectors of Eq. (1) becomes:

$$\mathbf{X} = (c_x \mathcal{X}^M, \beta \mathbf{F}_x), \quad (7)$$

$$\mathbf{Y} = (c_y \mathcal{Y}^M, \beta \mathbf{F}_y), \quad (8)$$

where c_x and c_y are $M \times M$ diagonal matrices that contain weights influencing each spectral coordinate, and β is a $K \times K$ diagonal matrix containing the weights for each feature (to emphasize or reflect confidence). Each feature is initially scaled, as in Eq. (3), to fit the values of the Fiedler vector, $\mathbf{x}^{(2)}$ (i.e., $\min(\mathbf{f}^{(k)}) = \min(\mathbf{x}^{(2)})$ and $\max(\mathbf{f}^{(k)}) = \max(\mathbf{x}^{(2)})$). The weights c of the spectral coordinates takes into account the smoothness of an eigenvector (measured by its eigenvalue $\lambda^{(u)}$) and the confidence in the reordering (measured by the permutation cost $Q^{(u)}$). Specifically, the weight, $c^{(u)}$, of the u^{th} spectral coordinate is:

$$c^{(u)} = \exp(- (Q^{(u)} \lambda^{(u)})^2 / 2\sigma^2), \quad (9)$$




Fig. 5. Diffusion of the correspondence map: a) The closest corresponding points on the spectral embedding might not necessarily be coherent spatially, consequently, b) their spatial coordinates are smoothed using a mean filter within the neighborhood structure, this moves the corresponding points to regularized positions ($H(y_{\phi(i)})$), and, c) the final correspondence map is reset by finding the closest nodes to these regularized positions.

where σ is a normalization factor set to

$$\sigma = \text{mean} \left\{ Q^{(u)} \lambda^{(u)} \right\}_{u=1 \dots M}. \quad (10)$$

The alignment of these embeddings can be viewed as a nonrigid registration, $\mathbf{X} = \phi(\mathbf{Y})$. Fig. 4 shows the alignment challenge where the first three spectral components ($\mathbf{x}^{(2)}$, $\mathbf{x}^{(3)}$, $\mathbf{x}^{(4)}$) are used as 3D (x, y, z) coordinates for visualization purposes. The Robust Point Matching [18] with a Thin Plate Spline-based transformation is often used for 2D or 3D registration. However, with this approach, the final registration depends on the number and choice of the control points. We apply the recent Coherent Point Drift method [41] which is scalable to N dimensions, fast, and demonstrates excellent performance in this application.

To increase speed in FOCUSR, we take advantage of the property of the Coherent Point Drift method that a continuous transformation derived from a subset of the points can be applied to all nodes of the dense embeddings. In our case, we subsample \mathbf{X} and \mathbf{Y} by taking randomly a few points (in our experiments we chose 1% of the total number of vertices, roughly 1000 points).

2.5 Final Diffusion

After alignment, both embedded representations can be directly compared ($\mathbf{X} = \phi(\mathbf{Y})$), i.e. two points which are closest in the embedded representations could be treated as corresponding points in both meshes. However, the mapping is not guaranteed to be smooth, even after the CPD alignment. The spectral regularization promotes smoothness of the correspondence map, but it is possible to have irregularities in the smoothness when the features differ significantly between the two meshes. The resulting embeddings warped with the CPD, in the $K + M$ multi-dimensional space, can contain local spatial incoherence in the correspondence map (as illustrated on Fig. 5 a). Consequently, we include a postprocessing step to enforce additional smoothness of the correspondence map.

The correspondences obtained after CPD are used to map the second mesh vertices (target point y in Fig. 5 a) to the first mesh vertices (fixed points x in Fig. 5 a). The 3D coordinates of these mapped points on the second mesh are now treated as independent scalars and diffused on

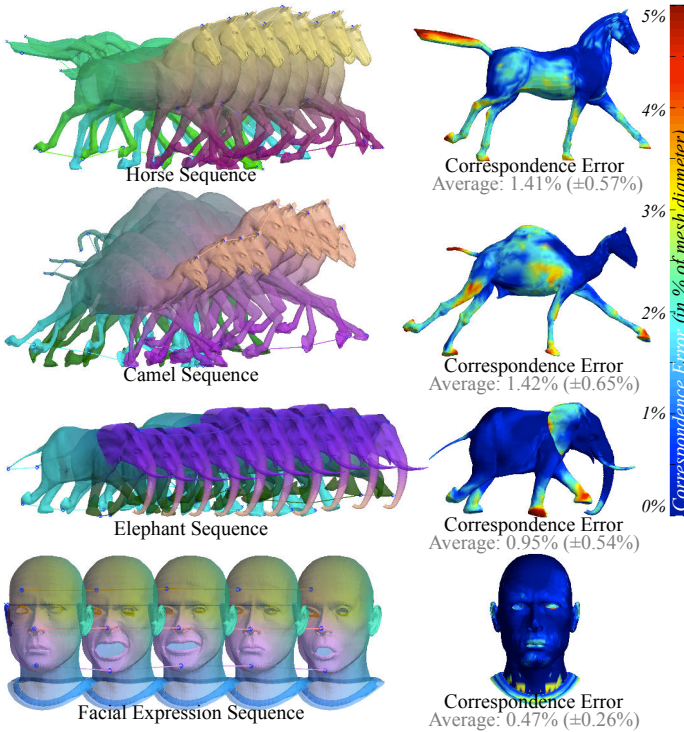


Fig. 6. Correspondences across animated sequences for 50 horses (average error of 1.41% ($\pm 0.57\%$)), 50 camels (1.42% ($\pm 0.65\%$)), 50 elephants (0.95% ($\pm 0.54\%$)), and 10 facial expression changes (0.47% ($\pm 0.26\%$)), for clarity, only a subset of each sequence is shown. Corresponding points have a unique color ((r,g,b) components given by the first three eigenmodes) and colored lines show five points tracked along the sequence for visualization (blue circles show corresponding points found with FOCUSR, blue crosses show ground truth).

the surface of the first mesh (i.e., this moves the points of the second mesh to positions obeying the (smooth) neighborhood system of the first mesh as illustrated with points $H(y)$ in Fig. 5 b). We used the smoothing method in [20] which is similar to the Laplacian smoothing, while other methods could also be used for this step. At this stage, the points x on the first mesh can be associated with either the smoothed coordinates $H(y)$ on the second mesh (i.e., vertices of the first mesh could be matched to coordinates in between the vertices of the second mesh), or with actual points on the second mesh. In our applications, we matched nodes to nodes, so the latter strategy is chosen. The correspondence map linking the first mesh to the second mesh is therefore updated by linking each point in the first mesh with the point in the second mesh which has the minimum Euclidean distance to the *diffused* geometric coordinates (shown with the new map in Fig. 5 c). In our experiments, 40 iterations were sufficient to diffuse the point coordinates. The fourth step in Fig. 3 shows a few corresponding points between two brain surfaces.

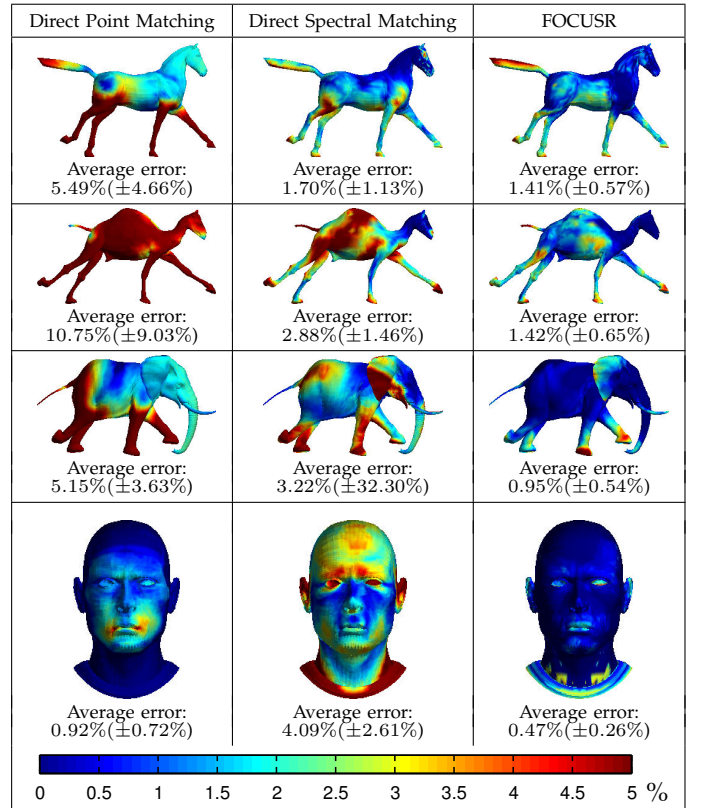


Fig. 7. Average relative error distance from the ground truth in each animated sequence (in percentage between 0% and 5%, e.g., if a mesh is 100 mm, 5% means an average error of 5 mm, the lower percentage the better). Three settings are used: a) Direct matching of closest points on surfaces, b) direct matching of closest points on rigidly aligned spectral embeddings, and c) matching using nonrigidly aligned spectral embeddings in FOCUSR. While matching points in the spectral domain clearly improves correspondence, FOCUSR gets additional precision by aligning the spectral embeddings. Note that no additional features are used here in FOCUSR.

3 RESULTS

To demonstrate the effectiveness of FOCUSR, we first match a variety of generic meshes in a controlled experiment (3 sequences of 50 and 10 frames of models in various poses, totaling the computation of 160 matchings, all with a known ground truth) and show that the use of a nonrigid alignment of spectral coordinates improves precision over a simple direct spectral matching method. In a second experiment, we evaluate FOCUSR on an established benchmark dataset and show that our method tracks the accuracy of the state-of-the-art with a simpler approach. We chose the high resolution TOSCA shapes [13] (totaling 80 objects with up to 50,000 vertices) since it has been tested on a variety of mesh matching methods [12], [34]. In a last experiment, we show the benefits of using FOCUSR in a real medical application, namely to brain surface matching where precision is crucial and where additional features are known to be meaningful to surface alignment. We do so by processing and analyzing the correspondence of 264 pairs of brain surfaces (each with

up to 135,000 vertices) using 15 different combinations of features (totaling the computation of 3,960 correspondence maps). This neuroimaging application reveals the full power of FOCUSR where the introduction of additional features significantly improves shape matching.

3.1 Benefits of Spectral Alignment

We first begin our validation by showing that FOCUSR can find efficiently and precisely a dense correspondence between generic meshes. We use the data from [61] (available publicly³) where animal models have been deformed in various poses. These meshes were created in [61] by transferring the deformation of a sequence of source meshes to target reference meshes. We use in the dataset the sequence of a galloping animal for a horse (8,431 vertices, 50 frames), an elephant (21,887 vertices, 50 frames), and a camel (42,321 vertices, 50 frames), all illustrated on Fig. 6. We want to recover the deformations and assess the precision of the correspondences between all models in a sequence and the reference model. For each gallop animation, the same mesh is deformed, and all vertices across the sequence maintain a direct one-to-one correspondence with the reference mesh (i.e., node i of any mesh in the animation corresponds with node i (the same index value i) in the reference mesh). This gives a ground truth for the correspondence maps in all animations (i.e., $\phi(i) = i$) on which we can compare our method.

We quantify precision by measuring the average distance between the locations of corresponding points found with FOCUSR and with the ground truth. That is, for all points $v_i \in \mathcal{V}_1$ in the first mesh matching the points $v_{\phi(i)} \in \mathcal{V}_2$ in the second mesh, the mean distance error is the average of the distances, $\frac{1}{N} \sum_i \|\mathbf{x}_i - \mathbf{x}_{\phi(i)}\|$, between the real locations of the corresponding points, \mathbf{x}_i , and their recovered locations on the second mesh, $\mathbf{x}_{\phi(i)}$. For each gallop animation, we computed the correspondence maps of the meshes of all frames with the reference mesh. Fig. 7 shows the average relative distance error for all sequences when finding the closest points in space, in the spectral domain, and when using FOCUSR in its simplest setting (i.e., $K = 0$ in Eq. (8)). Mismatches due to nonrigid deformations (e.g., articulated limbs of the galloping animals) are the most severe when matching in the spatial domain, while these errors are attenuated when matching occurs in the spectral domain (about a 60% increase in precision). FOCUSR improves precision over the simple spectral matching by about 50%.

The relative average distance error in FOCUSR with its standard deviation (expressed in percentage of the size of a mesh) is for the whole horse gallop animation: 1.41% ($\pm 0.57\%$) with an average computation time of 44 seconds, for the camel gallop: 1.42% ($\pm 0.65\%$) in 79 seconds, and for the elephant gallop: 0.95% ($\pm 0.54\%$) in 98 seconds (timing were performed on a 2.8 GHz Intel Pentium 4 using unoptimized Matlab code). We additionally ran the same experiment on an animation of changing facial expressions (15,941 vertices, 10 expressions) and found a relative average error of 0.47% ($\pm 0.26\%$) with




Fig. 8. Matching using FOCUSR on two models of the Michael dataset. Corresponding points have a unique color $((r, g, b)$ components given by the first three eigenmodes). Red lines indicates 50 random correspondences.

on average 40 seconds of computation. All these errors remain relatively small with corresponding points found at more or less 1% of the size of the mesh from their true locations (e.g., for a mesh of 100 mm, an error of 1% is a mismatch of 1 mm). Additionally, five points of interest were tracked along each animation (between the ears, the tail tip, right rear and front paw, and on the sternum of the animals; and the right ear, left upper eyelid, nose tip, lower lip, and chin of the head).

By applying a nonrigid alignment of spectral coordinates, FOCUSR exhibits an improved level of precision (of about 1.4% error) even in the absence of using additional features. Higher errors often occurs in areas of high non-rigid deformation, such as skin stretching (e.g., the side of the horse undergoing expansion and compression while galloping). One might also argue that displaced areas are not necessarily errors (e.g., the skin could move freely over a body by a few centimeters when galloping).

3.2 Benchmarking on Nonrigid Meshes

We now pursue our evaluation on a benchmark dataset that presents a broader variety of non rigid deformations. The high resolution TOSCA dataset [13] consists of 3 humanoids in various poses (Michael in 20 poses, each with 52,565 vertices; David, 7 poses with 52,565 vertices; Victoria, 12 poses, 45,659 vertices), a centaur (6 poses, 15,768 vertices), a cat (11 poses, 27,894 vertices), a dog (9 poses, 25,290 vertices), a horse (8 poses, 19,248 vertices) and a wolf (3 poses, 4,344 vertices). Meshes within the same class have again the same triangulation with vertices numbered in a compatible way, thus establishing a ground truth for correspondence maps (i.e., $\phi(i) = i$). We quantify precision in a similar fashion to the previous experiment, that is by measuring the displacement of correspondences from their ground truth positions.

We first ran our experiment by matching all models against their respective reference using the direct spectral matching approach used earlier (i.e., finding pairs

3. Meshes available at <http://people.csail.mit.edu/sumner/research/deftransfer>

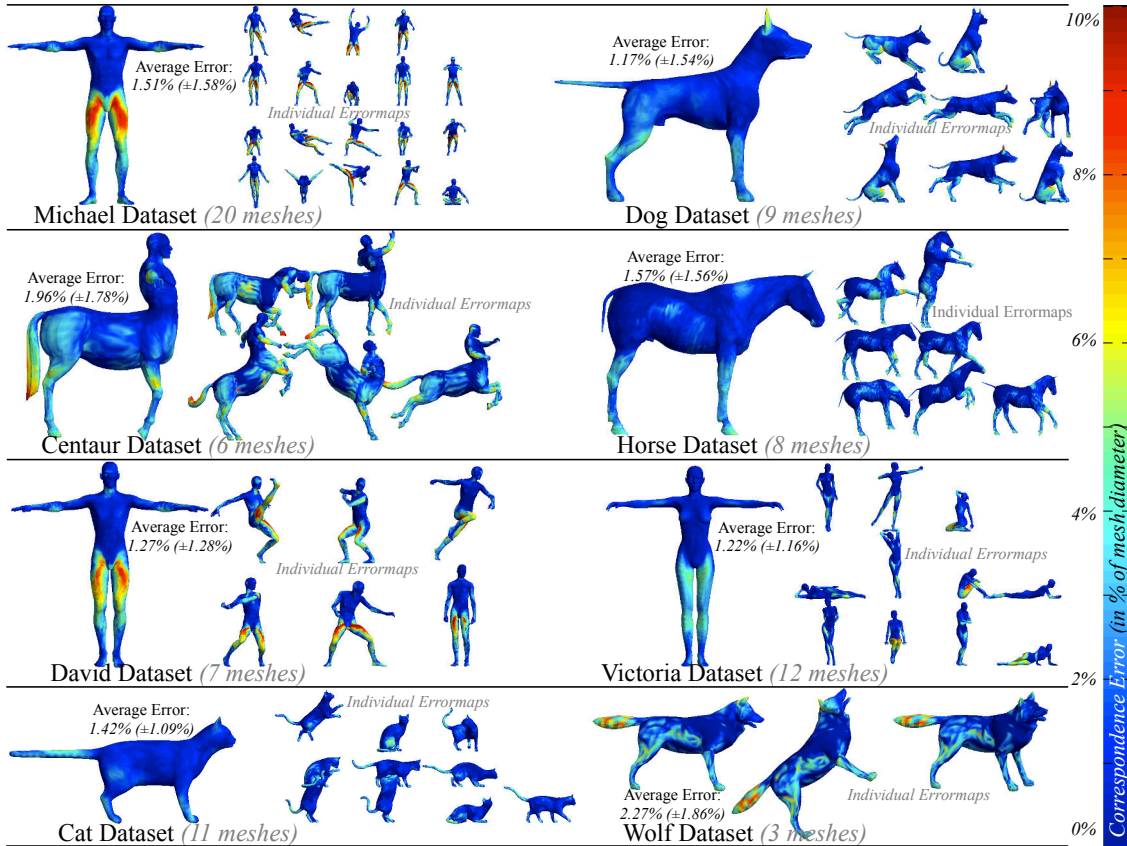


Fig. 9. Correspondence error when using FOCUSR on the TOSCA benchmark dataset (distance error color coded in percentage of the mesh size). The average error map within each class is shown on the reference mesh (larger models on the left) while errors in matching each individual model to the reference are shown on each right. FOCUSR demonstrates an average correspondence error of 1.46% on the TOSCA dataset.

of closest points in the spectral domain). The relative distance error from ground truth (expressed in percentage of the mesh size, with standard deviation) is on average $3.27\%(\pm 2.59\%)$, or within each dataset, for Michael: $4.06\%(\pm 2.87\%)$, for the centaur: $3.91\%(\pm 3.07\%)$, for David: $3.57\%(\pm 2.70\%)$, for the cat: $2.53\%(\pm 2.26\%)$, for the dog: $3.74\%(\pm 3.00\%)$, for the horse: $2.71\%(\pm 2.39\%)$, for Victoria: $2.26\%(\pm 2.05\%)$, and for the wolf: $2.21\%(\pm 2.21\%)$.

We then matched all models using FOCUSR in its simplest setting (i.e., $K = 0$ in Eq. (8)) in order to benchmark its general performance (without exploiting additional information) with respect to the direct matching of spectral embeddings as well as with the state-of-the-art. One example of matching is illustrated on Fig. 8 and the results for all pairs of matching are reported in Fig. 9 where the average correspondence error within each class is shown on each respective reference model (larger models on the left). The relative distance error from ground truth (expressed in percentage of the mesh size, with standard deviation) is on average $1.46\%(\pm 1.43\%)$, or within each dataset, for Michael: $1.51\%(\pm 1.58\%)$, for the centaur: $1.96\%(\pm 1.78\%)$, for David: $1.27\%(\pm 1.28\%)$, for the cat: $1.42\%(\pm 1.09\%)$, for the dog: $1.17\%(\pm 1.54\%)$, for the horse: $1.57\%(\pm 1.56\%)$, for Victoria: $1.22\%(\pm 1.16\%)$, and for the wolf: $2.27\%(\pm 1.86\%)$. Most errors look to appear on thighs, where again, ambiguity remains when there is a slight change in isometry (e.g., how to handle the

relative motion of the skin over the body). The increase in accuracy between the performances of the direct spectral matching and FOCUSR (a 55% improvement) illustrates the benefits of using a nonrigid alignment of spectral embeddings in spectral matching approaches.

For comparison purposes, we computed the geodesic distances from all corresponding points to their ground truth using the code provided by [34] and compared the results of FOCUSR with those reported in [34] for several methods, namely the Möbius Voting [36], the Generalized Multidimensional Scaling (GMDS) [10], and the Heat Kernel Maps method (HKM) [43]. FOCUSR produces an average geodesic distance error of 0.0470 which is significantly lower than the error produced by the Möbius Voting (0.0985), GMDS (0.3085) and HKM (0.2287). Details for each dataset are summarized in Table 1 and the cumulative distributions of the geodesic distance error are shown in Fig. 10. It is interesting to note that the Blended Intrinsic Maps method [34] produces the best results, however, it relies on multiple conformal maps that are blended together in order to establish correspondences. In this experiment, we evaluate the accuracy attained by FOCUSR with simple spectral coordinates, however, its accuracy may possibly increase with the use of various blended conformal maps.

Furthermore, from the results reported in [12], FOCUSR appears to perform at higher accuracy than other conven-

Dataset	FOCUSR	Blended	Möbius	GMDS	HKM
Michael:	0.0539	0.0329	0.1392	0.2820	0.2593
Centaur:	0.0476	0.0602	0.0407	0.1456	0.1205
David:	0.0486	0.0290	0.2118	0.2463	0.1799
Cat:	0.0369	0.0430	0.1000	0.1490	0.1815
Dog:	0.0357	0.0370	0.0499	0.5403	0.2291
Horse:	0.0713	0.0235	0.0342	0.3350	0.1568
Victoria:	0.0350	0.0306	0.1129	0.3918	0.3602
Wolf:	0.0581	0.0084	0.0145	0.3685	0.2444
Total:	0.0470	0.0342	0.0985	0.3085	0.2287

TABLE 1

Average geodesic error on TOSCA shapes using FOCUSR (our method), Blended Intrinsic Maps, Möbius Voting, GMDS, Heat Kernel Maps with 2 correspondences.




Fig. 10. Performance of various methods on TOSCA shapes. The x -axis depicts the geodesic distance from the computed corresponding points to the ground truth, and the y -axis indicates the percentage of correspondences having accuracy below the prescribed geodesic distance. FOCUSR (our method) performs better on TOSCA shapes than the Möbius Voting, GMDS, and Heat Kernel Maps. It is possible that the Blended Intrinsic Maps method, which uses a collection of conformal maps, may benefit from using maps generated with FOCUSR.

tional spectral methods such as [21], [22] (whose geodesic distance error is reported at 8.77 in the isometry test). It is, however, hard to assess the differences with other spectral methods, in particular with [40], yet our first controlled experiment showed that the nonrigid alignment of spectral components does improve the matching accuracy, whereas current state-of-the-art spectral methods, such as [40], currently rely on a simplistic rigid alignment. FOCUSR may therefore demonstrate further accuracy than conventional spectral methods due to its nonrigid spectral alignment.

3.3 Brain Surface Matching

We now show an application where FOCUSR demonstrates a significant improvement over typically used methods. Brain surface matching is an important topic for neuroimaging studies that requires the alignment of anatomical landmarks or functional activation across a population. Specifically, cognitive function can only be studied across individuals when correspondence is made from one individual to another between activation areas of the brains. The brain morphology offers the particularity that every individual has a unique folding pattern in

Coord.	Nodes	Edges	Error
S T C	S T C	S T C	
✓ ✓	✓ ✓	✓ ✓	0.14 mm
✓	✓	✓	0.13 mm
✓	✓	✓	0.29 mm
	✓	✓	0.50 mm
✓ ✓	✓ ✓	✓	0.07 mm
(with only spectral components)			0.38 mm
(simple feature matching)			53.02 mm

TABLE 2

Settings used in FOCUSR for the recovery of a synthetic deformation. Checkmarks indicate whether *sulcal depth* (S), *cortical thickness* (T), or *cortical Gaussian curvature* (C) is used as features, on graph node weights, or on graph edge weights. The synthetic deformation process did not distort sulcal depth or cortical thickness, but did distort Gaussian curvature. This experiment demonstrates that FOCUSR can profit from incorporating the meaningful (undistorted) features. The reported error gives the average error distance between the matched point and the ground truth across all available hemispheres. The use of FOCUSR without any features (with only spectral components) and the simple feature matching method are provided for comparison.

their cerebral cortical hemispheres while, at the same time, many large-scale similarities exist and allow correspondence between brain surfaces. Moreover, matching brain surfaces allows us to test the ability of FOCUSR to use extra features, such as the sulcal depth, the cortical Gaussian curvature, and the cortical thickness, that can potentially improve the precision of the correspondence beyond conventional spectral correspondence. We utilize the two features used by the FreeSurfer algorithm to drive alignment, which are the *sulcal depth* [25] at each point $\{s_1, s_2, \dots, s_n\}$ (as calculated by FreeSurfer), and the surface curvature at each point, $\{\kappa_1, \kappa_2, \dots, \kappa_n\}$. FreeSurfer outputs the mean curvature of a mesh, but in practice our method generated slightly better results when using the Gaussian curvature estimated with the method described in [60]. We thus chose to test the Gaussian curvature in our feature combinations in order to avoid exploding the number of feature combinations in our experiments. In addition, FreeSurfer also supplies gray matter cortical thickness (calculated from anatomical MRI image data [24]) at each point, $\{t_1, t_2, \dots, t_n\}$, which we can additionally test as a feature to drive the alignment with FOCUSR.

To demonstrate the flexibility of FOCUSR to handle different features, different combinations of these three additional features were used in our experiment. Additionally, we independently examine the effects of using the features to define only edge weights (in Eq. (3)), only node weights (in Eq. (5)), or only as coordinates for matching (in Eq. (8)).

3.3.1 Synthetic deformations

We begin with a synthetic experiment which is designed to demonstrate that FOCUSR profits from meaningful fea-

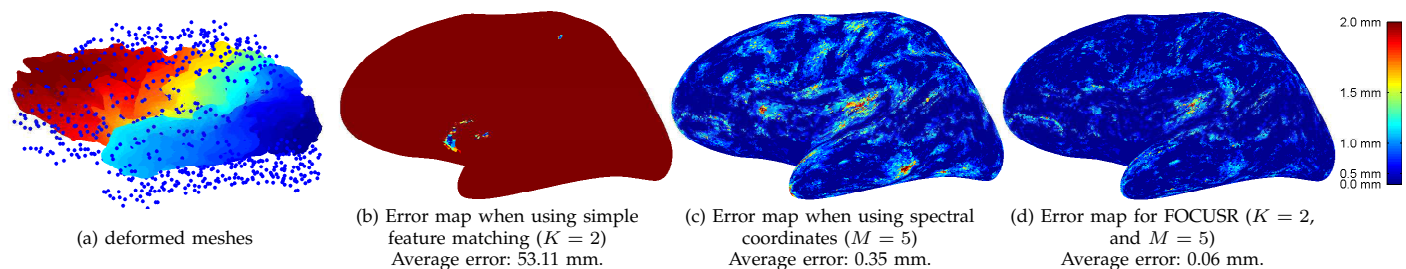


Fig. 11. Comparison with a synthetic ground truth using one brain hemisphere. (a): The deformed mesh (colored with its Fielder vector) overlaid with the original vertex positions illustrated by the blue dots. (b): When using simple feature matching ($K = 2, M = 0$), the mean distance error with the ground truth is as expected very high, 53.11 mm. (c): When using FOCUSR with only spectral components and no additional feature ($K = 0, M = 5$), the mean distance error is 0.35 mm. (d): When using FOCUSR with both spectral components and additional meaningful features ($K = 2, M = 5$), the mean distance error using this brain hemisphere is 0.06 mm. When iterating this experiment on all hemispheres, the mean distance error using FOCUSR is 0.07 mm (Note that surfaces are smoothed in (b,c,d) after correspondence to visualize the errors within the cortical foldings).

tures to produce a precise alignment. In this experiment, we synthetically deform a brain surface such that two of the features are preserved and one feature is distorted. FOCUSR is shown to perform better when the meaningful (preserved) features are included and worse when the noise feature is included. For our experiment, we match one brain hemisphere with a deformed version of itself. The vertex indexing remains the same in the deformed version. Similarly as the last experiment, the true matching is thus known (i.e., $\phi(i) = i$). We severely deform one cortical surface model, where for each point (x, y, z) , we apply the transformation $z' = (1 + \alpha)z$, i.e., a compression in the z -axis controlled by parameter α (we used $\alpha = 0.3$), and the transformation $x' = x + \beta r^2 / \max(r^2)$ with $r^2 = x^2 + y^2$, i.e., a radial distortion controlled by parameter β (we used $\beta = 15$). This simulates a deformation due to a drastic change in the head shape. The deformation however preserves the same mesh topology as it does not introduce any discontinuities or intersecting faces. Fig. 11 illustrates the position of the original hemisphere with the blue dots and the deformed hemisphere with the colored mesh. The sulcal depth and the cortical thickness are the same in both cortical meshes. The Gaussian curvature has been recomputed in the deformed mesh with the method described in [60]. Therefore two of the features (sulcal depth and cortical thickness) are meaningful under this distortion and one feature (Gaussian curvature) is a distracting noise feature. The goal of this experiment is to verify if the use of additional meaningful features helps the matching precision and to measure its improvement.

If we use the simple feature-only correspondence, the error is on average across all hemispheres 53.02 mm due to the fact that the correspondence map has virtually no mechanism to promote smoothness. When FOCUSR is used with only spectral components with no features (e.g., $K = 0$ and $M = 5$), we find for all hemispheres an average error distance of 0.38 mm as shown in the first error map of Fig. 11. Most errors appear to be located on the sulci extrema. By using FOCUSR to drive feature correspondence with spectral regularization, the error drops to 0.07 mm.

In FOCUSR, the surface features affect the correspon-

dence by using the features as coordinates in the point matching, and/or, by using the features to set edge weights, and/or by using the features to set node weights. Now we demonstrate that the greatest precision for FOCUSR is obtained by using the features in these three ways instead of just one or two of these ways. Specifically, we iterate through all of the 512 possible combinations ($2^{3 \times 3}$). Table 2 summarizes a few combinations. We tested FOCUSR using both sulcal depth and Gaussian curvature as additional features. The average error distance across all hemispheres is in this case 0.14 mm. Adding the sulcal depth as the only additional feature yields an error of 0.13 mm; adding only the cortex Gaussian curvature yields an error of 0.50 mm; and the cortical thickness yields an error of 0.29 mm. The best combination of features for FOCUSR was obtained when using sulcal depth and cortical thickness as additional coordinates and on graph nodes, and using cortical thickness on graph edges, yielding an error of 0.07 mm. It is expected that FOCUSR should perform best with these features, since they were not changed by the synthetic deformation, but the Gaussian curvature was. The error map on a single hemisphere is shown on Fig. 11. The best-performing combination of features demonstrates an almost perfect matching for FOCUSR.

This experiment shows that by incorporating meaningful features FOCUSR can indeed improve the matching precision. The weighting functions used here also differs slightly from the one used in [39] which used the exponentials of the additional features. This experiment confirms that using stable features between two cortices (i.e., the same sulcal depth and cortical thickness) improves the cortex matching precision.

3.3.2 Performance Evaluation on Real Data

Cortical surface matching is a challenging problem due to the wide variability in gyral morphology and topology between individuals. There is no ground truth available for perfect brain surface matching across individuals. However, FreeSurfer [25] has been demonstrated to provide highly precise cortical matchings that closely align cortical areas across subjects [32] and therefore provides a reliable surrogate for our comparison. The delineations of 81 sulcal

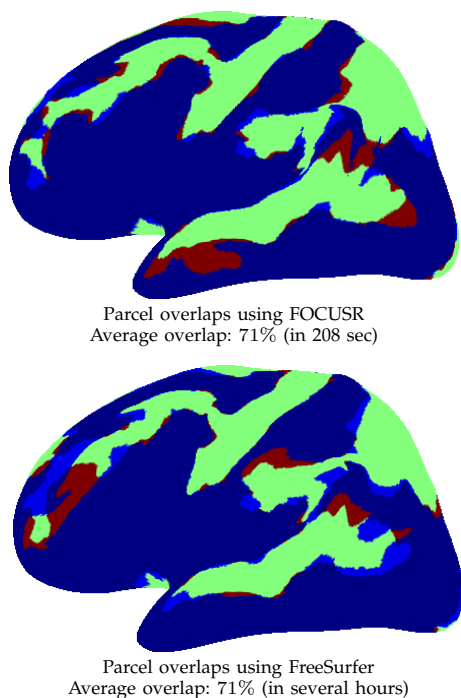


Fig. 12. In *green*, good overlap of projected sulcal regions, in *red*, wrong projection outside the sulcal regions, and in *light blue*, missed sulcal regions. (*First brain*) Correspondences computed in 208 seconds on average using FOCUSR, while (*second brain*) FreeSurfer required several hours (Note that in this visualization surfaces are smoothed to visualize the correspondence within the cortical foldings).

regions are available for 24 hemispheres (12 subjects). These sulcal regions were obtained using an automatic parcellation of the cortex [26] and are considered as our gold standard. Although parcellations of the cortex into named sulci and gyri are not expected to align between subjects in all cases (except for the primary folds), they do provide means to compare the two methods. We use correspondence maps generated by FreeSurfer and FOCUSR to project the parcellation areas onto different brain hemispheres and measure their overlap (illustrated on Fig. 12). To process a mesh of 135,000 vertices, FreeSurfer has a varying processing time which is currently on the order of several hours, while the time required by FOCUSR is on the order of 3–4 minutes. To process all our 264 possible pairs of left and right brain hemispheres, FOCUSR required on average 208 seconds (on a 2.8 GHz Intel Pentium 4 using unoptimized Matlab code). With reduced meshes of 20,000 vertices, FOCUSR performed in 19 seconds. The primary computational burden of the algorithm is the final diffusion of the correspondence map. This final step requires the smoothing of the matched mesh, which currently takes 84 seconds on average in Matlab. The total time to perform all our 264 correspondences using FOCUSR was 14 hours on a single computer, a substantial advantage compared to the several weeks required by FreeSurfer. Each overlap ratio is defined by the ratio of set intersection to set union. Fig. 13 shows the

Coord.	Nodes	Edges	Overlap	Overlap
S T C	S T C	S T C	Left	Right
			55.11%	55.18%
✓ ✓ ✓			70.79%	70.57%
✓ ✓ ✓	✓ ✓ ✓		70.77%	70.60%
✓ ✓ ✓		✓ ✓ ✓	70.77%	70.42%
✓ ✓ ✓	✓ ✓	✓ ✓	70.73%	70.66%
✓ ✓ ✓	✓ ✓	✓ ✓	69.65%	70.12%
✓ ✓ ✓	✓	✓	70.67%	70.65%
✓ ✓ ✓	✓	✓	65.51%	66.84%
✓ ✓ ✓		✓	70.67%	70.58%
✓ ✓ ✓	✓ ✓ ✓	✓ ✓ ✓	70.74%	70.41%
✓	✓	✓	71.10%	71.16%
✓	✓	✓	55.25%	56.77%
	✓	✓	55.28%	56.67%
✓	✓	✓ ✓	71.18%	71.11%
✓ ✓	✓ ✓	✓ ✓	69.64%	70.15%

TABLE 3

Different combinations of features with spectral correspondence (three modes: as matching coordinates, to define edge weights or to define node weights) demonstrate that using all three combination modes (as we advocate in FOCUSR) gives the best performance. Additionally, different combinations of three features, *sulcal depth* (S), *cortical thickness* (T), or *Gaussian curvature* (C), demonstrate that using sulcal depth and cortical curvature in FOCUSR provides performance most similar to FreeSurfer (as expected since it relies on similar features). In comparison, FreeSurfer's overlap ratios are 72.03% in the left hemispheres, and 70.95% in the right hemispheres. The experiment was performed over all 264 pairs (from 12 brains), averaged across the twelve largest parcels

overlap ratios for the twelve largest sulcal parcellations¹ using FOCUSR and FreeSurfer. The results of FOCUSR are correlated to FreeSurfer's overlaps with a correlation coefficient of $\rho = 0.897$.

From Fig. 13, we can see that FOCUSR closely matches the performance of FreeSurfer when using a similar feature set (sulcal depth and cortical curvature) to drive the correspondence (71.16% overlap for FOCUSR versus 70.95% overlap for FreeSurfer). In contrast, the pure feature matching or the use of FOCUSR with only spectral components produces results with a much lower precision (effectively null at 0.48% overlap). We now demonstrate that using features purely for edge or node weights (or purely as feature coordinates) also produces suboptimal results.

1. Sulcal regions: 9 (*G frontal middle*), 10 (*G frontal middle*), 18 (*G occipit temp med Lingual part*), 23 (*G parietal inferior Supramarginal part*), 24 (*G parietal superior*), 26 (*G precentral*), 41 (*Medial wall*), 42 (*Pole occipital*), 45 (*S central*), 47 (*S cingulate Main part and Intracingulate*), 59 (*S intraparietal and Parietal transverse*), 80 (*S temporal superior*).




Fig. 13. Average overlap ratios of the twelve largest sulcal regions on the right hemisphere over 264 matchings. (Dark blue) FOCUSR with features only, sulcal depth and cortical Gaussian curvature (i.e., $K = 2$, $M = 0$): $0.48\% \pm 0.28\%$ overlap. (Cyan) FOCUSR with spectral components only (i.e., $K = 0$, $M = 5$): $55.18\% \pm 9.09\%$ overlap. (Yellow) FOCUSR with features (sulcal depth and cortical Gaussian curvature) and spectral components (i.e., $K = 2$, $M = 5$): $71.11\% \pm 5.98\%$ overlap. (Red) FreeSurfer’s overlap ratios (requiring weeks of computations): $70.95\% \pm 7.27\%$ overlap. FOCUSR only required 14 hours to perform all 264 matchings and is strongly correlated with FreeSurfer (correlation coefficient of $\rho = 0.897$). The error bars show the standard deviation of each overlap ratio.

3.3.3 Testing with multiple configurations

We first analyze the matching performance using different configurations of the same features used by FreeSurfer, namely sulcal depth and cortical curvature. In a second step, to demonstrate the flexibility of FOCUSR, we introduce a different feature not used by FreeSurfer and tested several combinations of features to see whether any of these combinations performs better than FreeSurfer. Additional features were incorporated in FOCUSR using Eq. (8), Eq. (3), and Eq. (5), with $\gamma = 1.2$ and $\beta = 0.2$ (the description of the behavior of these parameters are described in [39]). Overall, fifteen different combinations of additional features were used. For each combination, we ran FOCUSR on the 132 pairs ($n(n-1)$ with $n = 12$ brains) of left brain hemispheres and on the 132 pairs of right brain hemisphere (totaling 3,960 matchings, 264×15). The results are summarized in Table 3. In comparison, FreeSurfer performs with an average overlap ratio for the largest parcels of $72.03\% (\pm 8.52\%)$ in the left hemispheres (the variation is the standard deviation of all overlap ratios), and $70.95\% (\pm 7.27\%)$ in the right hemispheres. Fig. 13 shows three relevant combinations.

- The *first* combination shown on Fig. 13 demonstrates the poor performance of the direct feature matching method where FOCUSR uses no spectral coordinates as described in Section 2.1 (i.e., $M = 0$, matching is a simple feature comparison using sulcal depth and cortical Gaussian curvature ($K = 2$) as used in Eq. (1)). The average overlap ratio on the largest parcels is effectively null at $0.38\% (\pm 0.19\%)$ in the left hemispheres ($0.48\% (\pm 0.28\%)$ in the right hemispheres).

- The *second* combination shows FOCUSR using no features and only spectral components ($K = 0$ and $M = 5$). The average overlap ratio on the largest parcels is only $55.11\% (\pm 10.73\%)$ in the left hemispheres ($55.18\% (\pm 9.09\%)$ in the right hemispheres).
- The *third* combination shows the full power of FOCUSR where it uses spectral components alongside sulcal depth and cortical curvature features, which are the same features driving the correspondence in FreeSurfer. The overlap ratio is as high as $71.18\% (\pm 7.63\%)$ in the left hemispheres ($71.11\% (\pm 5.98\%)$ in the right hemispheres). This is almost a perfect match with the overlap ratios in FreeSurfer (72.03% in the left side, and 70.95% in the right side).

FOCUSR is, in the left and right cortices, equivalent with FreeSurfer’s overlap ratios (71.18% vs. 72.03% in the left side, and 71.16% vs. 70.95% in the right side). It is important to note that there is no perfect combination of features to drive the correspondence. Our experiment shows that certain combinations perform better on particular parcels than on others. The best combination of extra features thus depends on which sulcal region of the brain should be matched. This finding concurs with a similar conclusion in [69].

3.3.4 Dependence on the number of spectral coordinates

In the previous section we demonstrated that it is optimal to use features to derive edge weights, node weights and as explicit feature coordinates. We now examine the dependence of the performance on the number of eigenvectors used as spectral coordinates by running the




Fig. 14. Robustness of FOCUSR to the number of eigenvectors used as spectral coordinates. When no eigenvector is used ($M = 0$, i.e., direct feature matching), the correspondence is weak, whereas the performance becomes stable after just three eigenvectors are included. The performance is measured with the overlap ratios of sulcal regions as used previously.

previous experiment with a varying number of eigenvectors. When no spectral regularization is used (i.e., direct feature matching with $M = 0$ eigenvectors), the algorithm relies solely on feature coordinates. As expected, the performance is weak. The plot in Fig. 14 shows, for the twelve largest parcels, indeed a low overlap ratio of 0.38% in the left hemispheres (0.48% in the right side) when using pure feature matching, sulcal depth and cortical Gaussian curvature, with no spectral coordinates (i.e., $M = 0$). The performance improves quickly when eigenvectors are used (i.e., $M > 0$) to spatially regularize the feature matching. These spectral coordinates provide additional means of discrimination during the optimization of the correspondence map. FOCUSR gains no further significant improvement in quality after $M > 3$. Essentially, this result demonstrates that the primary purpose of the spectral coordinates is to provide a spatial regularization, which is achieved by using only the lowest-frequency eigenvectors.

4 DISCUSSION AND CONCLUSIONS

This paper presents a novel method, based on spectral correspondence, for the challenging task of precise surface matching. Current methods, most of which are iterative and control local deformations of surfaces, are dependent on the extrinsic mesh geometry. They find their limitations when matching articulated or highly deformable shapes. Isometric deformations may be handled by using intrinsic metrics, however, this strategy often requires more complex methods, such as GMDS [10], which finds the embedding with the least distortion from one surface onto another (computationally expensive on meshes with $> 4,000$ vertices), approaches based on Heat Kernels [62], [43], which use multiscale geometric descriptors, or even other conformal maps [71]. We show in this paper that FOCUSR greatly generalizes and improves spectral correspondence, making it suitable for efficient and precise surface matching.

Additional information (e.g., texture, anatomical information, or landmark positions) can help in establishing a better correspondence. For example, this is the strategy that FreeSurfer [25] relies on to match brain surfaces. Spatial regularization becomes crucial in these methods. It

is clear from our experiments that direct feature matching, with no spatial regularization (i.e., $M = 0$) exhibits very poor performance. We decided to improve this strategy by using a spectral regularization of the feature matching method and to improve spectral matching by using a nonrigid alignment. The space of regularization (i.e., the spectral domain) is dependent on inherent properties of the mesh geometry. This modified strategy would free our method from the limitations of matching articulated or highly deformable shapes. Spectral methods provide a natural means of regularizing solutions at speeds of several orders of magnitude faster than current methods and are independent of the mesh extrinsic geometry. Our method can implicitly incorporate additional features to drive precise correspondence and it exploits the smoothness of the lowest-frequency harmonics of a graph Laplacian to regularize the correspondence map. Present day spectral correspondence methods are not fully realized and provide matchings that are not yet reaching a clinical level of precision. Currently, only intrinsic geometry can be embedded on graph edges and no additional information can be used. We provide a full realization of spectral correspondence where virtually any feature can be used as additional information as weights in graph edges, but also on graph nodes and as extra embedded coordinates with little or no computational expense. Furthermore, rigid transformations, or older point matching methods based on Thin Plate Splines are used [33] and are difficult to extend beyond 3D [59] to multidimensional feature space. Our approach is also not limited to genus zero surfaces [29] and can be applied to surfaces with arbitrary topology.

In its simplest form, FOCUSR is an improved spectral correspondence method that utilizes nonrigid point registration. We showed in our first experiment that a nonrigid alignment of the spectral coordinates improves significantly (by about 50%) the matching precision over a direct spectral matching. It has been demonstrated with a variety of generic models (animal and human models in various poses, varying facial expressions) that the error from a known ground truth is minimal (with 1.4% relative distance error for our matched models). FOCUSR also showed a greater accuracy on the TOSCA benchmark dataset than various state-of-the-art surface matching methods (Möbius Voting [36], the GMDS [10], and the Heat Kernel Maps method (HKM) [43]). The full power of FOCUSR is presented in a real-world application with the challenging task of brain surface matching across several individuals. We use FOCUSR with different combinations of additional features, such as sulcal depth, cortical Gaussian curvature, and cortical thickness, to improve the matching precision. The fast speed of our method allowed us to compute and analyze 3,960 correspondence maps (which is prohibitively expensive for FreeSurfer). When no regularization is used (e.g., $K = 2$ features and $M = 0$ spectral components), the correspondence generates a poor overlap ratio of 0.48% in the largest sulcal regions. When FOCUSR is used in its simplest form with no additional feature (e.g., $K = 0$ features and $M = 5$ spectral components), the overlap ratio is 55%. The performance of FOCUSR is improved by using additional information (e.g., $K = 2$ features and $M = 5$ spectral com-

ponents), and the overlap ratio increases to a level above 71% (versus 55% when using only spectral components). Our method is effectively equivalent to FreeSurfer’s level of precision (which is also around 71%) when aligning sulcal regions. However, the vast increase in speed (with a total processing time of 208 seconds on average for meshes of 135,000 vertices) and the added flexibility when using additional information gives new perspectives to previously computationally prohibitive experiments. New features (e.g., anatomical or functional features extracted from various data sources) can be quickly tested and evaluated to see if they improve cortex matching. Quick parameter sweeps can be performed to isolate the best parameter value sets, or alternatively, statistical learning can be used [16]. These computationally intensive experiments can help us to understand what features are consistently correlated with brain areas across individuals and what their role are during the development of the cortical folding pattern. FOCUSR may further improve accuracy by using other conformal maps as suggested by [34]. Additionally, the computational time could be significantly improved with a reimplement in C++ and with parallel programming for critical sections such as the eigendecomposition (e.g., LAPACK implementations on CUDA-enabled GPUs). Approximation methods for matrix eigendecomposition such as the Nyström approximation [27], the Gaussian projection [31], or the differentiable QR decomposition [6] could be used for additional speed up in processing time.

Spectral regularization promotes the smoothness of the correspondence map, but does not guarantee it. Better relaxation schemes, such as the Relaxation Labeling used in [74], might improve the matching precision. It is also important to consider which weighting function to use, for instance the cotangent weight has been shown to uniquely determine the discrete Riemannian metric [70], and to see how generalizable the parameter values are with a larger sample set. The use of different surface metrics [37] can be a promising area to investigate. The algorithm, as with other spectral methods, is also not symmetric (i.e., $\phi_{i \rightarrow j} \neq \phi_{j \rightarrow i}^{-1}$). The CPD alignment does not guarantee symmetry of the resulting transformation (i.e., the computed correspondence map matching nodes from mesh X to mesh Y might not be the same as the inverse correspondence map matching nodes from mesh Y to mesh X). Our method is also not tailored for matching partial data. However, FOCUSR may be used in any iterative method that drives a surface deformation with or without occlusion. Further improvement of the method will be toward achieving a better regularization and guarantee symmetry of the correspondence map. Nevertheless, FOCUSR already presents several clear advantages over present day methods for mesh correspondence and, in particular, conventional spectral matching. It provides a fast and precise solution for general mesh correspondence that can handle articulated or highly deformable surfaces, and creates a method that can implicitly use any set of additional features to drive improved precision.

ACKNOWLEDGMENTS

The authors would like to specially thank Bruce Fischl and Martin Reuter for their helpful comments, the anonymous

reviewers for their valuable comments and suggestions, Gareth Funka-Lea for supporting this project, as well as the financial support from the Alexander Graham Bell Canada Graduate Scholarships of the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. In *SIGGRAPH*, pages 587–594, 2003.¹
- [2] D. Anguelov, P. Srinivasan, H.-C. Pang, and D. Koller. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in Neural Information Processing Systems (NIPS)*, pages 33–40, 2004.¹
- [3] Q. Anqi, D. Bitouk, and M.I. Miller. Smooth functional and structural maps on the neocortex via orthonormal bases of the Laplace-Beltrami operator. *IEEE Transactions on Medical Imaging*, 25(10):1296–1306, 2006.²
- [4] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE International Conference on Computer Vision Workshops*, pages 1626–1633, 2011.²
- [5] M. A. Audette, F. P. Ferrie, and T. M. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis*, 4(3):201–217, 2000.¹
- [6] F. Bach and M. I. Jordan. Learning spectral clustering. In *Advances in Neural Information Processing Systems*, 2004.^{2, 16}
- [7] Y. Bengio, J.-F. Paiement, P. Vincent, O. Delalleau, N. Le Roux, and M. Ouimet. Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. In *Advances in Neural Information Processing Systems*. 2004.²
- [8] P. J. Besl. Geometric modeling and computer vision. *Proceedings of the IEEE*, 76(8):936–958, 1988.¹
- [9] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *SIGGRAPH*, pages 187–194, 1999.¹
- [10] A. M. Bronstein, M. I. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences of the United States of America*, 103(5):1168–1172, 2006.^{1, 2, 10, 15}
- [11] A. M. Bronstein, M. I. Bronstein, and R. Kimmel. Calculus of nonrigid surfaces for geometry and texture manipulation. *IEEE Transactions on Visualization and Computer Graphics*, 13:902–913, 2007.²
- [12] A. M. Bronstein, M. M. Bronstein, U. Castellani, A. Dubrovina, L. Guibas, R. Horaud, R. Kimmel, D. Knossow, E. von Lavante, D. Mateus, M. Ovsjanikov, and A. Sharma. SHREC: Robust correspondence benchmark. In *Eurographics Workshop 3DOR*, 2010.^{8, 10}
- [13] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes (Monographs in Computer Science)*. Springer, 1 edition, 2008.^{8, 9}
- [14] C. Cagniard, E. Boyer, and S. Ilic. Probabilistic deformable surface tracking from multiple videos. In *European Conference on Computer vision*, pages 326–339, 2010.¹
- [15] M. Carcassoni and E. Hancock. Spectral correspondence for point pattern matching. *Pattern Recognition*, 36(1):193–204, 2003.²
- [16] U. Castellani, M. Cristani, and V. Murino. Statistical 3D shape analysis by local generative descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2555–2560, 2011.^{2, 16}
- [17] T. Chan, J. Gilbert, and S.-H. Teng. Geometric spectral partitioning. Technical Report PARC CSL-94-15, Xerox, 1995.²
- [18] H. Chui. A new point matching algorithm for non-rigid registration. In *Computer Vision and Image Understanding*, volume 89, pages 114–141, 2003.^{3, 7}
- [19] F. Chung. *Spectral Graph Theory*. AMS, 1997.^{1, 5}
- [20] M. Desbrun, M. Meyer, P. Schröder, and A.H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH*, pages 317–324, 1999.⁸
- [21] A. Dubrovina and R. Kimmel. Matching shapes by eigendecomposition of the Laplace-Beltrami operator. In *3DPVT*, 2010.¹¹
- [22] A. Dubrovina and R. Kimmel. Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. *Advances in Adaptive Data Analysis*, pages 203–228, 2011.¹¹
- [23] A. Elad and R. Kimmel. On bending invariant signatures for surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1285–1295, 2003.²
- [24] B. Fischl and A. M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *Proceedings of the National Academy of Sciences of the United States of America*, 97(20):11050–5, 2000.¹¹

- [25] B. Fischl, M. I. Sereno, R. B. Tootell, and A. M. Dale. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8(4):272–284, 1999.^{1, 11, 12, 15}
- [26] B. Fischl, A. van der Kouwe, C. Destrieux, E. Halgren, F. Segonne, D. H. Salat, E. Busa, L. J. Seidman, J. Goldstein, D. Kennedy, V. Caviness, N. Makris, B. Rosen, and A. M. Dale. Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14(1):11–22, 2004.¹³
- [27] C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the nystrom method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):214–225, 2004.¹⁶
- [28] L. Grady and J. R. Polimeni. *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer, 2010.^{2, 4, 5}
- [29] X. Gu, Y. Wang, T. F. Chan, P. M. Thompson, and S.-T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23(8):949–958, 2004.¹⁵
- [30] D. Hahnel, S. Thrun, and W. Burgard. An extension of the ICP algorithm for modeling nonrigid objects with mobile robots. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 915–920, 2003.¹
- [31] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM*, 53(2):217–288, 2011.¹⁶
- [32] O. Hinds, N. Rajendran, J. R. Polimeni, J. C. Augustinack, G. Wiggins, L. L. Wald, D. H. Rosas, A. Potthast, E. L. Schwartz, and B. Fischl. Accurate prediction of V1 location from cortical folds in a surface coordinate system. *Neuroimage*, 39(4):1585–99, 2008.^{2, 12}
- [33] V. Jain and H. Zhang. Robust 3D shape correspondence in the spectral domain. In *IEEE International Conference on Shape Modeling and Applications*, page 19, 2006.^{2, 15}
- [34] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. In *SIGGRAPH*, 2011.^{1, 2, 8, 10, 16}
- [35] M. H. Lin. Tracking articulated objects in real-time range image sequences. In *IEEE International Conference on Computer Vision (ICCV)*, volume 1, pages 648–653 vol.1, 1999.¹
- [36] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. In *SIGGRAPH*, 2009.^{10, 15}
- [37] R. F. Liu, H. Zhang, A. Shamir, and D. Cohen-Or. A part-aware surface metric for shape analysis. *Eurographics*, 28(2), 2009.¹⁶
- [38] G. Lohmann, D. Y. von Cramon, and A. C. Colchester. Deep sulcal landmarks provide an organizing framework for human cortical folding. *Cerebral Cortex*, 18(6):1415–20, 2008.¹
- [39] H. Lombaert, L. Grady, J. R. Polimeni, and F. Chriet. Fast brain matching with spectral correspondence. In *Information Processing in Medical Imaging (IPMI)*, volume 22, pages 660–673, 2011.^{12, 14}
- [40] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.^{2, 6, 11}
- [41] A. Myronenko and X. Song. Point-set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2009.^{3, 7}
- [42] M. Niethammer, M. Reuter, F.E. Wolter, S. Bouix, N. Peinecke, M.S. Koo, and M. Shenton. Global medical shape analysis using the Laplace-Beltrami spectrum. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, volume 4791, pages 850–857, 2007.²
- [43] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.^{1, 2, 10, 15}
- [44] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1105–1120, 1999.²
- [45] M. Reuter. Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *International Journal of Computer Vision*, 89(2):287–308, 2009.²
- [46] M. Reuter, F. E. Wolter, M. Shenton, and M. Niethammer. Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer Aided Design*, 41(10):739–755, 2009.²
- [47] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. In *Proceedings on International Conference on 3-D Digital Imaging and Modeling*, pages 145–152, 2001.¹
- [48] J. Schreiner, A. Asirvatham, E. Praun, and H. Hoppe. Inter-surface mapping. In *SIGGRAPH*, pages 870–877, 2004.¹
- [49] E. L. Schwartz, A. Shaw, and E. Wolfson. A numerical solution to the generalized mapmaker’s problem: flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(9):1005–1008, 1989.²
- [50] G. L. Scott and H. C. Longuet-Higgins. An algorithm for associating the features of two images. *Proceedings of the Royal Society Biological Sciences*, 244(1309):21–26, 1991.²
- [51] L.S. Shapiro and J.M. Brady. Feature-based correspondence: an eigenvector approach. *Image and Vision Computing*, 10(5):283–288, 1992.²
- [52] A. Sharma, R. Horaud, J. Cech, and E. Boyer. Topologically-robust 3D shape matching based on diffusion geometry and seed growing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2481–2488, 2011.¹
- [53] C. R. Shelton. Morphable surface models. In *IEEE International Journal of Computer Vision*, volume 38, pages 75–91, 2000.¹
- [54] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.²
- [55] Y. Shi, I. Dinov, and A. W. Toga. Cortical shape analysis in the Laplace-Beltrami feature space. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 208–215, 2009.²
- [56] Y. Shi, B. Sun, R. Lai, I. Dinov, and A. W. Toga. Automated sulci identification via intrinsic modeling of cortical anatomy. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 49–56, 2010.²
- [57] A. Shokoufandeh, D. Macrini, S. Dickinson, K. Siddiqi, and S. W. Zucker. Indexing hierarchical structures using graph spectra. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(7):1125–1140, 2005.²
- [58] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999.²
- [59] R. Sprengel, K. Rohr, and H. S. Stiehl. Thin-plate spline approximation for image registration. In *Engineering in Medicine and Biology Society*, volume 3, pages 1190–1191, 1996.¹⁵
- [60] D. C. Steiner and J. M. Morvan. Restricted Delaunay triangulations and normal cycle. In *Proceedings of the Symposium on Computational Geometry*, pages 312–321, 2003.^{11, 12}
- [61] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. In *SIGGRAPH*, pages 399–405, 2004.^{1, 9}
- [62] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Symposium on Geometry Processing*, pages 1383–1392, 2009.^{2, 15}
- [63] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H. P. Seidel. Isometric registration of ambiguous and partial data. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1185–1192, 2009.¹
- [64] T. Tung and T. Matsuyama. Dynamic surface matching by geodesic mapping for 3D animation transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1402–1409, 2010.¹
- [65] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):695–703, 1988.²
- [66] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Eurographics*, 30(6):1681–1707, 2011.²
- [67] K. Varanasi, A. Zaharescu, E. Boyer, and R. Horaud. Temporal surface tracking using mesh evolution. In *European Conference on Computer Vision*, pages 30–43, 2008.¹
- [68] S. Wuhrer, Chang Shu, and P. Bose. Posture invariant correspondence of triangular meshes in shape space. In *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1574–1581, 2009.²
- [69] B. T. T. Yeo, M. R. Sabuncu, T. Vercauteren, D. J. Holt, K. Amunts, K. Zilles, P. Golland, and B. Fischl. Learning task-optimal registration cost functions for localizing cytoarchitecture and function in the cerebral cortex. *IEEE Transactions on Medical Imaging*, 29(7):1424–41, 2010.¹⁴
- [70] W. Zeng, R. Guo, F. Luo, and X. Gu. Discrete heat kernel determines discrete riemannian metric. *Graphical Models*, 74(4):121–129, 2012.¹⁶
- [71] W. Zeng, D. Samaras, and D. Gu. Ricci flow for 3D shape analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(4):662–677, 2010.^{2, 15}
- [72] H. Zhang, A. Sheffer, D. Cohen Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi. Deformation-driven shape correspondence. In *Proceedings of the Symposium on Geometry Processing*, pages 1431–1439, 2008.¹
- [73] H. Zhang, O. Van Kaick, and R. Dyer. Spectral mesh processing. *Eurographics*, 29(6):1865–1894, 2010.²
- [74] Y. Zheng and D. Doermann. Robust point matching for nonrigid shapes by preserving local neighborhood structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):643–649, 2006.¹⁶



Herve Lombaert earned the PhD degree in Computer Engineering from the École Polytechnique de Montréal in 2012, in collaboration with Siemens Corporate Research (Princeton) and INRIA Sophia Antipolis - Méditerranée (France), and the engineering degree from the École Polytechnique de Montréal in 2003. He was a Research Associate at Siemens Corporate Research between 2004 and 2005, and is currently a Postdoctoral Fellow in the Centre for Intelligent Machine at McGill University. He is interested in finding structures in images, understanding correspondences between images, and extracting 3D and 4D (3D+t) information from images, with applications in brain and cardiac imaging. He is a member of the IEEE.

structures in images, understanding correspondences between images, and extracting 3D and 4D (3D+t) information from images, with applications in brain and cardiac imaging. He is a member of the IEEE.



Farida Cheriet received the B.Sc. degree in computer science from the University of Science and Technology Houari Boumediene, Algiers, Algeria, in 1984, the D.E.A. degree in the field of languages, algorithms, and programming from the University of Paris VI, Paris, France, in 1986, and the Ph.D. degree in computer science from the University of Montreal, Montreal, QC, Canada, in 1996. Since 1999, she has been with the Department of Computer Engineering, École Polytechnique de Montreal, Montreal, where she is currently a Full Professor. She is also with Sainte-Justine Hospital Research Center, Montreal. Her research interests include 3-D reconstruction of bone structures from X-rays, calibration of X-ray imaging systems, non-invasive 3-D modeling of scoliosis deformities, 3-D navigation systems for minimally invasive surgery, 3-D reconstruction of vascular structures from angiographic images, and 3-D motion estimation from spatiotemporal sequences. She is a member of the IEEE.

currently a Full Professor. She is also with Sainte-Justine Hospital Research Center, Montreal. Her research interests include 3-D reconstruction of bone structures from X-rays, calibration of X-ray imaging systems, non-invasive 3-D modeling of scoliosis deformities, 3-D navigation systems for minimally invasive surgery, 3-D reconstruction of vascular structures from angiographic images, and 3-D motion estimation from spatiotemporal sequences. She is a member of the IEEE.



Leo Grady earned a BSc in Electrical Engineering at the University of Vermont in 1999 and a PhD in the Cognitive and Neural Systems department at Boston University in 2003. From 2003-2012 he served as a Principal Research Scientist at Siemens Corporate Research and since 2012 he has been the VP of R&D at HeartFlow, Inc. His research has included work on image segmentation, graph theory, discrete calculus, optimization and PDEs, including two books on applications of graph theory to computer vision. This research

has been applied primarily to problems in medical imaging, with a recent focus on cardiovascular applications. He holds 28 granted patents, including winning the Thomas Edison Patent Award for his contribution to image segmentation. He is a member of the IEEE and the IEEE Computer Society.



Jonathan R. Polimeni is an Instructor in Radiology at Harvard Medical School and an Assistant in Biomedical Engineering in the Department of Radiology of the Massachusetts General Hospital. Dr. Polimeni received his B.S. in Electrical and Computer Engineering from the Johns Hopkins University where he worked in the Sensory Communication and Microsystems Laboratory under Prof. Andreas Andreou. He later received his Ph.D. in Electrical and Computer Engineering from Boston University where he worked in the Computational Neuroscience and Computer Vision Laboratory under Prof. Eric Schwartz. His Ph.D. work was in the measurement and modeling of visuotopic maps in macaque and human visual cortex. Dr. Polimeni's postdoctoral training was under the supervision of Prof. Lawrence L. Wald at the Athinoula A. Martinos Center for Biomedical Imaging at the Massachusetts General Hospital and focused on technological development for increasing the spatial resolution and accuracy of the functional measurements. He joined the faculty of the Martinos Center in 2010. His research is currently focused on investigating the functional architecture of the human cerebral cortex using high-resolution functional MRI using 7 Tesla field strength scanners, and on characterizing and understanding the biological limits on spatial specificity of the fMRI signals. He is a member of the IEEE.

computational Neuroscience and Computer Vision Laboratory under Prof. Eric Schwartz. His Ph.D. work was in the measurement and modeling of visuotopic maps in macaque and human visual cortex. Dr. Polimeni's postdoctoral training was under the supervision of Prof. Lawrence L. Wald at the Athinoula A. Martinos Center for Biomedical Imaging at the Massachusetts General Hospital and focused on technological development for increasing the spatial resolution and accuracy of the functional measurements. He joined the faculty of the Martinos Center in 2010. His research is currently focused on investigating the functional architecture of the human cerebral cortex using high-resolution functional MRI using 7 Tesla field strength scanners, and on characterizing and understanding the biological limits on spatial specificity of the fMRI signals. He is a member of the IEEE.