

Spectral Demons – Image Registration via Global Spectral Correspondence

Herve Lombaert^{1,2}, Leo Grady³, Xavier Pennec², Nicholas Ayache²,
and Farida Cheriet¹

¹ Ecole Polytechnique de Montreal, Canada

² INRIA Sophia Antipolis, France

³ Siemens Corporate Research, Princeton, NJ

Abstract. Image registration is a building block for many applications in computer vision and medical imaging. However the current methods are limited when large and highly non-local deformations are present. In this paper, we introduce a new direct feature matching technique for non-parametric image registration where efficient nearest-neighbor searches find global correspondences between intensity, spatial and geometric information. We exploit graph spectral representations that are invariant to isometry under complex deformations. Our direct feature matching technique is used within the established Demons framework for diffeomorphic image registration. Our method, called *Spectral Demons*, can capture very large, complex and highly non-local deformations between images. We evaluate the improvements of our method on 2D and 3D images and demonstrate substantial improvement over the conventional Demons algorithm for large deformations.

1 Introduction

The comparison of images or the need to find correspondences between them is key to many computer vision applications and crucial to establishing accurate diagnostics in medical imaging. Current methods register images with Euler-Lagrangian approaches that typically model elastic and viscous (non-rigid) deformations (recent state-of-the-art is surveyed in [1]). Parametric models ([2,3]) take advantage of restraining the optimization to a few parameters, while non-parametric models use the entire space of displacement fields. Recently, optimization is performed in the diffeomorphism group [4,5,6,7,8,9] (differentiability and reversibility of the transformation) that prevents an invalid folding of space and guarantees a smooth one-to-one mapping between points. The update schemes underlying these methods rely, however, on forces derived from the image gradients and are therefore fundamentally limited by their *local scope* (e.g., gradients are null in textureless areas and optimization is undermined by local minima). The typical response is to use a multilevel scheme to capture larger deformations in a higher resolution but fundamentally does not solve the local scope of their gradient-based update schemes.

In order to capture complex large deformations, we introduce a new approach for image registration based on a *direct feature matching* technique with a *global scope* that uses simple nearest-neighbor searches in a multi-dimensional space comprising information on *image texture* (e.g., pixel intensities), *space* (e.g., Euclidean

coordinates of pixels) and on *global image geometry*. For the latter, graph spectral representations [10], which are invariant to isometry (preserving geodesic distances), can capture large and complex deformations. Spectral methods [10,11,12,13,14] are popular for general graph partitioning and are applied in computer vision for mesh matching [15,16,17,18] and shape retrieval [19]. Pioneered in the late 80s, [12,13,14], spectral correspondence methods match points in shapes by comparing eigenvectors of a proximity matrix derived from the mesh structure. Recent work (surveyed in [20,21]) uses different types of proximity matrices, graph structures and spectrum deformations [15,16,18]. However, spectral representations of images (strongly linked with Normalized Cuts [22,23]) have never been used for image registration.

Our new direct feature matching technique provides a geometric component with a *global scope* that can be used for image registration, i.e., the nearest-neighbor search finds correspondences with the best compromise in image similarity, spatial regularity and geometric similarity. Unfortunately, diffeomorphism is not guaranteed since a simple nearest-neighbor search could leave unassigned correspondences. Therefore, our approach for image registration is to perform our direct feature matching technique within a diffeomorphic framework such as the efficient symmetric *Log-Demons* algorithm [24,25]. The new method, called **Spectral Demons**, enables a symmetric and diffeomorphic registration of images undergoing *large* and *complex* deformations. We describe in the next section our new direct feature matching and briefly review the *Log-Demons* algorithm. The results evaluate the capability and robustness of the *Spectral Demons* to register images with highly non-local deformations.

2 Method

Our registration approach aims at exploiting the *global scope* and the speed of nearest-neighbor search methods in order to capture very large deformations between two images. We begin our methodology with our simple and direct feature matching technique followed by how spectral representations are built for images. We finish by explaining how our new direct feature matching technique can be used within a diffeomorphic framework for image registration with very large deformations.

2.1 Direct Feature Matching

Image registration warps a moving image M toward a fixed image F through a transformation ϕ that maps points from F to M (i.e., features $\mathbf{F}(\cdot)$, such as point coordinates $\mathbf{x}(\cdot) = (x, y)$ or image intensity $I(\cdot)$, match those in the transformed features $\mathbf{M}(\phi(\cdot))$, or simply $\mathbf{F} \mapsto \mathbf{M} \circ \phi$). A *direct* approach for feature matching would find the point correspondence with a nearest-neighbor search in the feature space (e.g., with a Voronoi tessellation or a k -d tree) such that $\phi(i) = \operatorname{argmin}_{j \in M} \|\mathbf{F}(i) - \mathbf{M}(j)\|^2$ (e.g., if point $j \in M$ has the closest intensity from the one of point $i \in F$, then $\phi(i) = j$). Obviously, matching 1D features such as image intensity would result in a highly unsmooth mapping lacking any spatial regularity (points with the closest intensities might be far apart in the images). Spatial regularity can be introduced by incorporating Euclidean coordinates in the feature space where points now have the extended coordinates $\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F)$ and $\mathbf{M} = (\alpha_i I_M, \alpha_s \mathbf{x}_M)$ (with weights α_i and α_s). A nearest-neighbor search in such extended space effectively provides similarity

in pixel intensity and closeness in space between corresponding pixels and minimizes the similarity criterion:

$$\text{Sim}(F, M, \phi) = (I_F - I_{M \circ \phi})^2 + \frac{\alpha_s^2}{\alpha_i^2} \|\mathbf{x}_F - \mathbf{x}_{M \circ \phi}\|^2 \quad (1)$$

However, such method lacks information on the image geometry (how to naturally match objects in different poses?) and does not produce a diffeomorphic mapping (a one-to-one mapping is not guaranteed). Each issue is addressed below.

2.2 Spectral Correspondence

The spectral representation of shapes [10,11,12,13,14,15,16,18] has the strong property of being invariant to *isometry*, i.e., corresponding points between shapes in different poses would share the same *spectral coordinates* (or signature) even if they are far away in space (e.g., a point on a nose tip has a geometric description that is unique even if moved in space). By adding these spectral coordinates in our feature space, we enforce an intrinsic *geometric consistency* in our matching technique.

Spectral Graph Theory. From the pixels of I_Ω (the portion of an image I bounded by a contour Ω), the connected undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is constructed with the vertices $\mathcal{V}$ representing pixels and the edges $\mathcal{E}$ defined by the neighborhood structure of these vertices. Such graph can be represented with its adjacency matrix W in terms of affinity weights [11] where high weights are given to edges within a region of uniform intensity and low weights are given to edges crossing region boundaries, e.g., $W_{i,j} = \exp(-(I(i) - I(j))^2 / 2\sigma^2) / \|\mathbf{x}(i) - \mathbf{x}(j)\|^2$ if $(i, j) \in \mathcal{E}$ and 0 otherwise (parameter σ depends on the image noise and is set, without being limited, to $\sigma = \text{mean}\{|I(i) - I(j)|\}_{(i,j) \in \mathcal{E}}$). The (diagonal) degree matrix D provides the total weighting of all edges connected to each vertex ($D_{ii} = \sum_j W_{i,j}$) and the Laplacian matrix is defined by $L = D - W$. Here, we consider the general Laplacian operator on a graph $\mathcal{L} = G^{-1}(D - W)$ [11], i.e., a $|\mathcal{V}| \times |\mathcal{V}|$ sparse matrix where G is the (diagonal) node weighting matrix, e.g., $G = D$.

Spectral Coordinates. The graph spectrum [10] computed from the decomposition of the Laplacian $\mathcal{L} = \mathbf{X}^T \mathbf{A} \mathbf{X}$ comprises the eigenvalues (in increasing order) $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{|\mathcal{V}|})$ and their associated eigenvectors $\mathbf{X} = (\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(|\mathcal{V}|)})$ ($\mathbf{X}$ is a $|\mathcal{V}| \times |\mathcal{V}|$ sparse matrix where each column $\mathbf{x}^{(\cdot)}$ is an eigenvector). The first eigenvector $\mathbf{x}^{(0)}$ is the stationary distribution (related to the expected return time of a random walker). The following eigenvectors associated with the non-zero eigenvalues are the fundamental modes of vibrations of the shape (with free ends) depicted by I_Ω (in a Riemannian sense). We thus prefer the term eigenmode since they are effectively functions over I_Ω (visualized as images, see Fig. 1). The eigenmodes of lower modal frequencies are harmonics depicting coarse geometric properties of I_Ω while those associated with higher eigenvalues depict finer geometric details in I_Ω . Moreover, the oscillations at a modal frequency λ occur around nodal sets (where the eigenmode values equal 0) that reside on prominent demarcations of the shape geometry. For instance, graph-based segmentation methods (e.g., [22,23]) rely on the nodal set of $\mathbf{x}^{(1)}$, called the Fiedler vector [10], to find a binary partition of an image. The number




Fig. 1. The eigenmodes of the graph Laplacian are used as geometric descriptors. They remain stable under complex deformations (i.e., invariant to isometry). The lower eigenmodes describe coarse geometric properties, while higher eigenmodes describe finer details in the images (the image contours are overlaid on the eigenmodes).

of half waves in these oscillations (or the number of extrema in the eigenmodal values) is also given by the algebraic multiplicity of their eigenvalue, n_λ .

In our approach (summarized in Alg. (1)), we consider the first k eigenmodes of lower modal frequencies $\mathcal{x}^{(1..k)}$ (i.e., the strongest intrinsic geometric descriptors). Their components represent the k -dimensional *spectral coordinates* $\mathcal{x}$ in a spectral domain¹ where each point i has the coordinates $\mathcal{x}(i) = (\mathcal{x}^{(1)}(i), \mathcal{x}^{(2)}(i), \dots, \mathcal{x}^{(k)}(i))$ (a truncated line of matrix $\mathbf{X}$). The spectral representation $\mathcal{x}$ has the strong property of being invariant to isometry, i.e., if F and M are images of the same object in different poses, equivalent points would share similar coordinates $\mathcal{x}_F$ and $\mathcal{x}_M$. We use this property to improve our direct feature matching by extending Eq. (2) with these spectral coordinates (weighted with α_g). A nearest-neighbor search between $\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F, \alpha_g \mathcal{x}_F)$ and $\mathbf{M} = (\alpha_i I_M, \alpha_s \mathbf{x}_M, \alpha_g \mathcal{x}_M)$ provides similarity in intensity, space and in intrinsic geometric characteristics, effectively minimizing:

$$\text{Sim}(F, M, \phi) = (I_F - I_{M \circ \phi})^2 + \frac{\alpha_s^2}{\alpha_i^2} \|\mathbf{x}_F - \mathbf{x}_{M \circ \phi}\|^2 + \frac{\alpha_g^2}{\alpha_i^2} \|\mathcal{x}_F - \mathcal{x}_{M \circ \phi}\|^2, \quad (2)$$

where $\mathcal{x}_{M \circ \phi}$ are the spectral coordinates of the corresponding points in the transformed image $M \circ \phi$.

The choice of the number of spectral components $k = n_{\lambda_1}$ is motivated by the Colin de Verdière’s number [26] which is in this case the multiplicity of the Fiedler vector n_{λ_1} and is also related [26] to the maximal dimension of a space in which the graph $\mathcal{G}$ can be mapped (i.e., the eigenspace of the Fiedler eigenvalue reveals the principal symmetries in I_Ω and, $n_{\lambda_1} \leq 2$ in 2D, $n_{\lambda_1} \leq 3$ in 3D). More complex symmetries in the cyclic or dihedral group could be considered with a higher k , but is not required in our method.

¹ In our notation $\mathbf{x}$ is the Euclidean coordinates (e.g., x, y, z in 3D) and superscripted $\mathcal{x}^{(u)}$ is the u^{th} component of the spectral coordinates $\mathcal{x}$

Algorithm 1. Spectral Correspondence**Input:** Images F , M .**Output:** Correspondence c mapping F to M

- Compute general Laplacians $\mathcal{L}_F$, $\mathcal{L}_M$.
 $\mathcal{L} = D^{-1}(D - W)$, where
 $W_{ij} = \exp(-(I(i) - I(j))^2 / 2\sigma^2) / \|\mathbf{x}(i) - \mathbf{x}(j)\|^2$
 $D_{ii} = \sum_j W_{ij}$,
- Compute first k eigenmodes of Laplacians
- Reorder $\mathcal{X}_M$ with respect to $\mathcal{X}_F$ (Eq. (3))
- Build embeddings:
 $\mathbf{F} = (I_F, \mathbf{x}_F, \mathcal{X}_F)$; $\mathbf{M} = (I_M, \mathbf{x}_M, \mathcal{X}_M)$
- Find c mapping nearest points $\mathbf{F} \mapsto \mathbf{M}$

Rearrangement of the Spectra. Unfortunately, the spectral coordinates $\mathcal{X}_F$ and $\mathcal{X}_M$ of points in F and M may not be directly comparable as a result of two phenomena. *Firstly*, there is a sign and scaling ambiguity between corresponding eigenmodes (if $\mathcal{X}^{(\cdot)}$ is an eigenmode of $\mathcal{L}$, so is $-\alpha \mathcal{X}^{(\cdot)}$) which requires a sign check and a scaling correction between $\mathcal{X}_F$ and $\mathcal{X}_M$. *Secondly*, the order of the eigenmodes is undefined within an eigenspace (if two eigenmodes $\mathcal{X}^{(u)}$, $\mathcal{X}^{(v)}$ share the same eigenvalue, their order (u, v) may differ between two images). The order is additionally perturbed with imperfections in isometry (near-symmetry creates close but not equal eigenvalues and may change order between images). We rearrange the spectral coordinates using two new simple heuristics.

The *first* issue is addressed by scaling the values of each eigenmodes in order to fit the range $[-1; +1]$. The nodal set (where $\mathcal{X}^{(\cdot)} = 0$) is thought to remain on a prominent geometric feature (an axis of symmetry in a Riemannian sense) and should not be changed. We scale thus the positive values (where $\mathcal{X}^{(\cdot)} > 0$) with $\mathcal{X}^{(\cdot)+} \leftarrow \mathcal{X}^{(\cdot)+} / \max\{\mathcal{X}^{(\cdot)+}\}$ and the negative values (where $\mathcal{X}^{(\cdot)} < 0$) with $\mathcal{X}^{(\cdot)-} \leftarrow \mathcal{X}^{(\cdot)-} / \min\{\mathcal{X}^{(\cdot)-}\}$. The *second* issue is addressed by finding the optimal permutation π such that $\mathcal{X}_F^{(\cdot)}$ and $\mathcal{X}_M^{\pi \circ (\cdot)}$ correspond with each other. The Hungarian algorithm (also used in [16,18]) minimizes the following 2D dissimilarity matrix:

$$C(u, v) = \sqrt{\frac{1}{|I_\Omega|} \sum_{i \in I_\Omega} \left(\mathcal{X}_F^{(u)}(i) - \mathcal{X}_M^{(v)}(i) \right)^2} + \sqrt{\sum_{i,j} \left(h_F^{\mathcal{X}_F^{(u)}}(i, j) - h_M^{\mathcal{X}_M^{(v)}}(i, j) \right)^2} \quad (3)$$

The first term is the difference in eigenmodal values between the images and, the second term measures the dissimilarities between the joint histograms $h(i, j)$ (a 2D matrix where the element (i, j) is the joint probability of having at the same time a pixel with intensity i and eigenmodal value $\mathcal{X}^{(\cdot)} = j$). The sign ambiguity can be removed by using, instead, the dissimilarity matrix $Q(u, v) = \min\{C(u, v), C(u, -v)\}$. To keep the notation simple, in the next sections, we assume that the spectral coordinates have been appropriately signed, scaled and reordered.

The ordered spectral coordinates provides our geometric component in our new direct feature matching. We now briefly review how diffeomorphism can be achieved for image registration.

2.3 Diffeomorphic Registration

The minimization of Eq. (2) does not guarantee a one-to-one mapping between points (only closest points are assigned and undefined correspondences are possible). Such

Algorithm 2. Exponential $\phi = \exp(v)$ **Input:** Velocity field v .**Output:** Diffeomorphic map $\phi = \exp(v)$.

- Choose N such that $2^{-N}v$ is close to 0
e.g., such that $\max \|2^{-N}v\| \leq 0.5$ pixels
- Scale velocity field $\phi \leftarrow 2^{-N}v$.
- for** N times **do**
- Square $\phi \leftarrow \phi \circ \phi$.
- end for**

Algorithm 3. The Log-Demons Framework**Input:** Images F , M and initial velocity field v **Output:** Transformation $\phi = \exp(v)$ from F to M **repeat**

- Find updates $u_{F \rightarrow M}$ mapping F to $M \circ \exp(v)$
- Find updates $u_{M \rightarrow F}$ mapping M to $F \circ \exp(-v)$
- Average updates: $u \leftarrow \frac{1}{2}(u_{F \rightarrow M} - u_{M \rightarrow F})$.
- Smooth updates: $u \leftarrow K_{\text{fluid}} \star u$
- Update velocity field: $v \leftarrow \log(\exp(v) \circ \exp(u))$
(approximated with $v \leftarrow v + u$)
- Smooth velocity field: $v \leftarrow K_{\text{diff}} \star v$.

until convergence

property however exists in classical methods for diffeomorphic registration such as the *Log-Demons* algorithm [24]. Arising from the theory of Lie groups, a diffeomorphic transformation ϕ (on a Lie group structure) is related to the exponential map of a velocity field v (a Lie algebra), i.e., $\phi = \exp(v)$. In the case of stationary velocity fields, a practical and fast approximation is possible with the scaling-and-squaring method [24] (Alg. (2)). As an aside, the inverse of the transformation is simply $\phi^{-1} = \exp(-v)$.

The *Log-Demons* framework alternates, similarly to the Maxwell’s demon, between the optimization of a similarity term, e.g., $\text{Sim}(F, M \circ \exp(v)) = (I_F - I_{M \circ \exp(v)})^2$, and a regularization term, e.g., $\text{Reg}(v) = \|\nabla v\|^2$, through the introduction of a hidden variable (the correspondences c) which allows a small error between alternations, e.g., $\text{dist}(c, \phi) = \|c - v\|$. Moreover, invariance to the order of the input images is possible with the symmetric extension of the algorithm [25,9]. The energy of the symmetric *Log-Demons* can be written:

$$E(F, M, \exp(c), \exp(v)) = \frac{1}{2}\alpha_i^2 (\text{Sim}(F, M \circ \exp(c)) + \text{Sim}(F \circ \exp(-c), M)) \\ + \alpha_x^2 \text{dist}(c, v)^2 + \alpha_T^2 \text{Reg}(v), \quad (4)$$

where the Euler-Lagrangian updates are computed directly on the stationary velocity field v and consist of the average of the forward and backward updates $u_{F \rightarrow M}$, $u_{M \rightarrow F}$ mapping F to $M \circ \exp(c)$ and, M to $F \circ \exp(-c)$ such as (see [24] for more details):

$$u_{F \rightarrow M} = - \frac{I_F - I_{M \circ \phi}}{\|\nabla I_{M \circ \phi}\|^2 + \alpha_x^2 |I_F - I_{M \circ \phi}|^2} \nabla I_{M \circ \phi}. \quad (5)$$

In the first step, the transformation ϕ is fixed and the updates are computed (optionally smoothed with a kernel K_{fluid} , e.g., Gaussian with σ_{fluid}). In the second step, the velocity field is updated $v \leftarrow v + u$ (optionally smoothed with the a kernel $K_{\text{diffusion}}$, e.g., Gaussian with $\sigma_{\text{diffusion}}$, see [27]). The general symmetric diffeomorphic Demons framework is summarized in Alg. (3).

2.4 Spectral Demons

At this stage, we have described two methods. The first (the direct feature matching technique, Eq. (2)) can capture large deformations between images but does not guarantee diffeomorphism and symmetry. The second (the Demons algorithm, Eq. (4)) offers these properties but suffers from the local scope of the update forces derived



Fig. 2. Pairs of images used in the synthetic experiments (Lena, heart, baseball player). Each left image is a fixed image, each right side is a moving image generated with a random deformation of at most 25 pixels (difficulties in red). These transformations provide our ground truth.

from the image gradient (Eq. (5)). Both methods can benefit from the incorporation, with very little modifications, of our spectral correspondence approach in the *Log-Demons* framework.

The so-called *Spectral Demons* algorithm takes advantage of the efficient diffeomorphic framework offered by Alg. (3) and finds the correspondences between images F and M using our *global* spectral approach in place of using *local* gradient based updates. To be more precise, the first two steps of Alg. (3) (originally computing updates with Eq. (5)) now perform spectral correspondence between images F and $M \circ \exp(v)$ using Alg. (1) (and respectively between M and $F \circ \exp(-v)$). This modification enables large jumps in each iteration where points are moving toward their isometric equivalent even if they are far away in space. This virtually enables the capture of very large deformations (with invariance to isometry) between images as well as a faster convergence of the algorithm. The underlying energy being minimized has the form of Eq. (4) where its similarity term is expressed with Eq. (2).

Multilevel Scheme. A multilevel scheme is also possible with the *Spectral Demons* where large deformations are assumed to be related with coarse geometric information. Spectral updates can thus be safely used in a lower resolution level, while finer details and local deformations remain computed in the higher levels of resolutions using the classical update forces based on the image gradient. This multilevel scheme keeps the computation of the eigenmodes tractable. On the same note, the computation of the eigenmodes can be used with the efficient Lanczos method (used by Matlab) which has a running time of $O(n\sqrt{n}) + O(n^2)$ [22], where n is the number of pixels in I_Ω , while spectral matching can be performed with a k -d tree which is built in $O(n \log^2 n)$ and queried in $O(\log n)$.

3 Results

We evaluate here the performance of the *Log-Demons* (our benchmark) and *Spectral Demons* (our method) by registering images with large and highly non-local deformations. In our controlled experiments, the full power of the *Spectral Demons* can be appreciated with drastic deformations of the images. Accurate measurements, with known ground truth, are used to evaluate the improvements in registration accuracy with respect to the *Log-Demons*. We additionally provide a real application where two human brains are registered.




Fig. 3. Comparison of the iterations in the *Log-Demons* and the *Spectral Demons* (our method) within the same resolution level. The update forces (indicated with the arrows and scaled for visualization) have a local scope with the *Log-Demons* and a global scope with the *Spectral Demons*. This global scope allows a faster convergence while the *Log-Demons* remains in a local minimum.

3.1 Controlled Experiments

In the first controlled experiment, we evaluate the fundamental difference between the update schemes of the *Log-Demons* and *Spectral Demons*. To do so, we analyze the convergence rate of both algorithms and, since we are not interested here in their final performance, we compare them within the same level of resolution. The algorithms use the same parameter set ($\sigma_{\text{fluid,diffuse}} = 1, \alpha_x = 1, k = 2, \alpha_g = 0.05, \alpha_s = 0.15, \alpha_i = 0.8$). We register the images on Fig. 2 (Lena has size 128^2 , the heart is 75^2 and the baseball player is 110×75). Each moving image is generated with random diffeomorphic deformations ϕ_{truth} with displacements of at most 25 pixels, i.e., we take the exponential map of a random velocity field generated with 15 random displacements (control points randomly located) diffused across the image (Gaussian smoothing $\sigma = 10$ pixels). Notably, Lena’s hat, her neck and, the player’s arm are the highest registration challenges, while the cardiac image (a 2D slice of an MRI) shows a papillary muscle (red circle on Fig. 2) severely deformed and almost fully collapsed (the muscle forms, however, a dent in the image and provides a signature that *Spectral Demons* can understand). The iterations of the *Log-Demons* and *Spectral Demons* are compared in Fig. 3. It shows that within the same level of resolution, the update forces computed with spectral correspondence are coherent spatially and geometrically, i.e., points move toward their geometric equivalent, however the update forces derived from the image gradient lack any global information on the shape geometry and put the *Log-Demons* into an erroneous local minimum.

The use of a multilevel scheme allows the *Log-Demons* to capture larger deformations (Fig. 4) but does not change the inherent local scope of its update forces. For instance, *Log-Demons* even with 4 levels of resolution ultimately fails in recovering the extreme deformations on the anterior side of the heart, while *Spectral Demons* without a multilevel scheme can successfully register the whole myocardium with a 71% improvement in performance (mean square differences (or MSE) of intensities with ground truth from 19.9×10^{-3} to 5.7×10^{-3}). The performance is further improved when 4 levels are used (down to 2.4×10^{-3} MSE, or a 88% improvement in the heart registration). Similar results are observable with the other images. Lena’s hat, her neck and, the player’s arm are successfully registered using the *Spectral Demons*




Fig. 4. Final registrations for (*top row*) Multilevel *Log-Demons* with 4 levels, (*middle row*) *Spectral Demons* with 1 level, (*bottom row*) Multilevel *Spectral Demons* with 4 levels. The fixed image is in blue and misalignments of the registered image are in red. The mean square differences of intensities are reported along their standard deviations. *Log-Demons* is limited in areas of high deformations, while *Spectral Demons* can capture these large deformations.

with respectively 73% improvement over *Log-Demons* in Lena's image and 63% in the baseball image.

Additionally, the quality of the computed registration maps ϕ is evaluated in terms of difference of displacements (in pixels) with the ground truth $\|\phi - \phi_{\text{truth}}\|$. The *Log-Demons* results in registration maps (Fig. 5) with larger errors in high deformation areas (e.g., Lena's hat or neck), whereas the *Spectral Demons* results in a smoother registration map with significantly less errors (62% less) in these same areas.

The cost of the global scope offered by *Spectral Demons* is increased computation time. For instance, on Lena's image, 50 iterations requires 55.82 seconds with *Log-Demons* and 178.32 seconds with *Spectral Demons*; on the heart image, 21.01 seconds with *Log-Demons* and 41.03 seconds with *Spectral Demons*; and on the baseball player's image, 42.06 seconds with *Log-Demons* and 102.17 seconds with *Spectral Demons*. We used unoptimized Matlab code on a Core 2 Duo, 2.53GHz.

Robustness to Deformation. The *Spectral Demons*' robustness to deformation is evaluated on Lena's image by exaggerating the previous synthetic transformation from $\phi_0 = \exp(0v)$ (zero displacements) to $\phi_2 = \exp(2v)$ (creating a maximal displacement of 40 pixels, see samples on Fig. 6). The performance is evaluated with the transformation differences (in pixels) from the ground truth. Both *Log-Demons* and *Spectral Demons* perform with sub-pixel accuracy on deformations below 20




Fig. 5. Differences (in pixels) between computed transformations and ground truth for (top row) Multilevel *Log-Demons*, (middle row) *Spectral Demons* and, (bottom row) Multilevel *Spectral Demons*. The fixed images are overlaid on the error maps. The Multilevel *Spectral Demons* decreases the error by 60% from the Multilevel *Log-Demons*.

pixels, however they differ with larger deformations where our method shows a greater robustness, e.g., with a deformation of 40 pixels (more than 30% of the image size) the average transformation error is 5.9 pixels with *Log-Demons* and 1.6 pixels with *Spectral Demons* (a 73% decrease).

Robustness to Noise. The analysis on the robustness to noise reveals the current limitations of both algorithms. An increasing Gaussian noise is added to Lena’s image, from $\sigma = 0$ to 0.25 (samples on Fig. 6). *Spectral Demons* performs better with noise $\sigma < 0.075$, however the comparison with control images (the unregistered noisy images) reveals that *Log-Demons* stops transforming the images when noise is increased (due to trapping in a local minimum). For instance, Fig. 6 shows that when using *Log-Demons* with noise $\sigma = 0.13$, the registered image is similar to its initial state, while *Spectral Demons*, even though with a larger average error (6.8 pixels versus 3.7 pixels), continues to recover large deformations (see the hat area on Fig. 6). With noise $\sigma > 0.2$, the corrupted images become problematic for *Spectral Demons* (graph edge weights are almost null and may need a different heuristic weighting), whereas *Log-Demons* is almost immediately trapped in a local minimum.

3.2 Registration of Medical Images

The performance of the *Spectral Demons* is evaluated in a medical application with the registration of brain MR images from the Internet Brain Segmentation Repository




Fig. 6. Robustness to deformation and noise. *a)* Deformations of Fig. 2 are amplified to a max of 40 pixels (image size is 128^2 , deformed samples on top). The transformation differences with ground truth (*y*-axis in pixels) are smaller with *Spectral Demons* (blue) than with *Log-Demons* (red). *b)* Gaussian noise is used (samples around the hat area). *Spectral Demons* loses advantage after $\sigma > 0.075$, however even with a lower error, *Log-Demons* stops moving when increasing noise (error is similar with untransformed images (controls, in green)), whereas *Spectral Demons* continues to recover large deformations.

(IBSR, <http://www.cma.mgh.harvard.edu/ibsr>, our images are 64^3 volumes). The brain presents a wide variety of shapes and internal structures across individuals. While the cerebral cortex is particularly convoluted and is the focus of many specific surface matching techniques ([28,29,17,18]), the registration of internal components in the brain, such as the white or gray matter, requires a volumetric approach. We chose two individuals that have lateral ventricles with different sizes (Fig. 7, the moving image shows a longer ventricle). We evaluate the registration accuracy with the overlap of the provided manual segmentations of the white and gray matter (measured with the Dice metric defined as the ratio $2(A \cap B)/(|A| + |B|)$ with 1 being an optimal overlap) as well as with the mean squared differences of pixel intensities between the fixed and registered images. The original (unregistered) setting has a Dice metric of 0.65 in the white matter and 0.69 in the gray matter. Both algorithms (using 4 levels) increase the overlap of the white and gray matter to respectively above 0.71 and 0.75, with a slight advantage to the *Spectral Demons*, however the comparison of the mean squared differences of intensities reveals a 4% improvement in accuracy and precision when using *Spectral Demons* (from 29.0×10^{-3} error in the original setting, decreasing to 20.8×10^{-3} with *Log-Demons* and to 19.9×10^{-3} with *Spectral Demons*). This experiment showed that *Spectral Demons* offers an improved performance in a real application when registering medical images.

The computation and the current implementation shows again that there is room for improvements with our method. With downsampled images of size 32^3 , 50 iterations requires 108.49 seconds with *Log-Demons* and 201.43 seconds with *Spectral Demons*. Notably, memory becomes problematic with our unoptimized Matlab code as volumes beyond 32^3 require the decomposition of Laplacian matrices larger than $32^3 \times 32^3$ (although extremely sparse, our current code is not optimized for such large matrices).




Fig. 7. Registration in 3D between two brains from healthy subjects using *Log-Demons* and *Spectral Demons* (both with 4 levels). While observations show apparently similar results (ventricles are circled), the mean square differences of intensities (MSE) between the fixed and registered images reveal an increase of 4% in accuracy and precision when using *Spectral Demons*. The Dice metrics of the white and gray matter (measuring segmentation overlaps) also increase with *Spectral Demons*.

3.3 Note on Image Segmentation

The Laplacian eigenmodes have been demonstrated to have important properties in the field of spectral graph theory [10,22,30,11] by providing a probabilistic foundation [23,31] for graph-based segmentation methods. In particular, the normalized cut problem [22] finds a segmentation x by minimizing $\frac{x^T L x}{x^T D x}$ (revealed by the Fiedler vector of the normalized Laplacian $D^{-\frac{1}{2}} L D^{-\frac{1}{2}}$ [23]). *Spectral Demons* considers the more general Laplacian operator $\mathcal{L}$ and effectively exploits for registration the same global image description used by normalized cuts for segmentation, i.e., eigenmodes of the (general) Laplacian operator are used for image registration. Since the Fiedler vector is an inherent part of our algorithm, binary segmentations of images come at no extra cost by taking either the positive or negative values of the Fiedler vector, $x^{(1)+}$ or $x^{(1)-}$. For example, the implicit segmentation of Lena’s image, shown on the *right* (with overlaid image contours), was obtained with the nodal sets of the first and second eigenmodes (positive values of the Fiedler vector gave warmer colors, negative values gave cooler colors, see also 1st and 2nd eigenmodes on Fig. 1)). Nodal sets of higher frequency eigenmodes may additionally reveal important geometric features for meaningful segmentation, however, an exhaustive experimental study on the segmentation aspect of our registration method goes beyond the scope of this paper.



Implicit Segmentation from nodal sets

4 Conclusion

The methods in the state-of-the-art of image registration are currently limited by the local scope of their underlying update schemes (based on image gradients). The typical response tries to capture larger deformations with a multilevel scheme where large jumps are possible with images of lower resolutions. Although sufficient for most applications, a multilevel scheme does not change the inherent local scope of update forces derived from the image gradient and finds its limitation with very large deformations. We proposed a fundamentally new update scheme based on a direct feature matching technique that uses image, spatial and geometric information. The *global scope* of our new update scheme is captured through spectral representations of images that are invariant to isometry. Our approach consists of finding point correspondences with a nearest-neighbor search in a multi-dimensional space that embeds pixel intensity, Euclidean coordinates and spectral coordinates. Such direct feature matching technique can be used within classical frameworks for image registration. Among them, the *Log-Demons* algorithm [9] provides symmetry and diffeomorphism of the registration. Our new registration method, called *Spectral Demons*, successfully captures large, complex and highly non-local deformations. We evaluated its performance with displacements of more than 30% of the image size and observed an improvement of 73% over the *Log-Demons*. Here, we assumed that registered images have a similar topology (same global shape) with no occlusions, holes, or missing parts. This is generally true in most applications as we are often interested in comparing the same object or similar ones (e.g., the same organ in medical imaging). Future work will aim at improving our main limitation, the computational cost of the spectral decomposition, for instance by pre-computing the eigenmodes before registration and warping them during registration in order to avoid costly spectral decompositions in each iteration. We will also investigate the use of the Nyström approximation [32] (by sampling the Laplacian matrix). Additionally, the strong links with image segmentation [23,31] will be studied more carefully. Image registration may also further benefit from employing global optimization framework. Nevertheless, we believe that our fundamentally new approach, relying on the *global* image structure to drive correspondences via spectral representations, is a significant contribution in the state-of-the-art of image registration, i.e., our registration respects the *global* Riemannian structure of the image (isomorphism, brought by *spectral coordinates*) on top of the differentiable structure (diffeomorphism, from state-of-the-art methods). The *Spectral Demons* algorithm can improve applications that employ objects or study organs undergoing very large, complex and highly non-local deformations.

References

1. Crum, W.R., Hartkens, T., Hill, D.L.: Non-rigid image registration: theory and practice. Br. J. of Radiology 77, 140–153 (2004)
2. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast MR images. IEEE TMI 18, 712–721 (1999)
3. Chui, H., Rangarajan, A.: A new algorithm for non-rigid point matching. In: CVPR, pp. 44–51 (2000)

4. Miller, M.I., Trouné, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. *An. Review of Biomed. Eng.* 4, 375–405 (2002)
5. Beg, M.F., Miller, M.I., Trouné, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* 61, 139–157 (2005)
6. Bossa, M., Hernandez, M., Olmos, S.: Contributions to 3D Diffeomorphic Atlas Estimation: Application to Brain Images. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) *MICCAI 2007, Part I. LNCS*, vol. 4791, pp. 667–674. Springer, Heidelberg (2007)
7. Allasonnière, S., Amit, Y., Trouné, A.: Towards a coherent statistical framework for dense deformable template estimation. *Royal Stat. Soc.* 69, 3–29 (2007)
8. Durrleman, S., Fillard, P., Pennec, X., Trouné, A., Ayache, N.: Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage* 55, 1073–1090 (2011)
9. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* 45, 61–72 (2009)
10. Chung, F.R.K.: *Spectral Graph Theory*. AMS (1997)
11. Grady, L.J., Polimeni, J.R.: *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer (2010)
12. Umeyama, S.: An eigendecomposition approach to weighted graph matching problems. *IEEE TPAMI* 10, 695–703 (1988)
13. Scott, G.L., Longuet-Higgins, H.C.: An algorithm for associating the features of two images. *Royal Soc. Bio. Sciences* 244, 21–26 (1991)
14. Shapiro, L.S., Brady, J.M.: Feature-based correspondence: an eigenvector approach. *Image and Vision Computing* 10, 283–288 (1992)
15. Jain, V., Zhang, H.: Robust 3D shape correspondence in the spectral domain. In: *Int. Conf. on Shape Modeling and App.* (2006)
16. Mateus, D., Horaud, R., Knossow, D., Cuzzolin, F., Boyer, E.: Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In: *CVPR*, pp. 1–8 (2008)
17. Reuter, M.: Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *IJCV* 89, 287–308 (2009)
18. Lombaert, H., Grady, L., Polimeni, J.R., Cheriet, F.: Spectral correspondence for brain matching. In: *IPMI*, pp. 660–670 (2011)
19. Egozi, A., Keller, Y., Guterman, H.: Improving shape retrieval by spectral matching and meta similarity. *IEEE Trans. Image Processing* 19, 1319–1327 (2010)
20. Zhang, H., Van Kaick, O., Dyer, R.: Spectral mesh processing. *Eurographics* 29 (2010)
21. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. *Eurographics* 30, 1681–1707 (2011)
22. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE TPAMI* 22, 888–905 (2000)
23. Meila, M., Shi, J.: Learning segmentation by random walks. In: *NIPS* (2000)
24. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Non-parametric Diffeomorphic Image Registration with the Demons Algorithm. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) *MICCAI 2007, Part II. LNCS*, vol. 4792, pp. 319–326. Springer, Heidelberg (2007)
25. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric Log-Domain Diffeomorphic Registration: A Demons-Based Approach. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) *MICCAI 2008, Part I. LNCS*, vol. 5241, pp. 754–761. Springer, Heidelberg (2008)
26. Tlustý, T.: A relation between the multiplicity of the second eigenvalue of a graph Laplacian, Courant’s nodal line theorem and the substantial dimension of tight polyhedral surfaces. *Electronic Journal of Linear Algebra* 16, 315–324 (2010)

27. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: the PASHA algorithm. *CVIU* 89, 272–298 (2003)
28. Fischl, B., Sereno, M., Tootell, R., Dale, A.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *H. Brain Map.* 8, 272–284 (1999)
29. Yeo, B.T., Sabuncu, M.R., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: Fast diffeomorphic Landmark-Free surface registration. *IEEE TMI* 29 (2010)
30. Luxburg, U.: A tutorial on spectral clustering. *Stat. and Comp.* 17, 395–416 (2007)
31. Robles-Kelly, A.: Segmentation via Graph-Spectral Methods and Riemannian Geometry. In: Gagalowicz, A., Philips, W. (eds.) *CAIP 2005. LNCS*, vol. 3691, pp. 661–668. Springer, Heidelberg (2005)
32. Drineas, P., Mahoney, M.W.: On the Nyström method for approximating a Gram matrix for improved Kernel-Based learning. *J. Mach. Learn. Res.* 6, 2153–2175 (2005)