



Spectral Log-Demons: Diffeomorphic Image Registration with Very Large Deformations

Herve Lombaert, Leo Grady, Xavier Pennec, Nicholas Ayache, Farida Cheriet

► To cite this version:

Herve Lombaert, Leo Grady, Xavier Pennec, Nicholas Ayache, Farida Cheriet. Spectral Log-Demons: Diffeomorphic Image Registration with Very Large Deformations. International Journal of Computer Vision, 2014, 107 (3), pp.254-271. 10.1007/s11263-013-0681-5 . hal-00979616

HAL Id: hal-00979616

<https://inria.hal.science/hal-00979616>

Submitted on 18 Nov 2014

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

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



Distributed under a Creative Commons Public Domain Mark 4.0 International License

Spectral Log-Demons – Diffeomorphic Image Registration with Very Large Deformations

Herve Lombaert · Leo Grady · Xavier Pennec · Nicholas Ayache · Farida Cheriet

Abstract This paper presents a new framework for capturing large and complex deformations in image registration and atlas construction. This challenging and recurrent problem in computer vision and medical imaging currently relies on iterative and local approaches, which are prone to local minima and, therefore, limit present methods to relatively small deformations. Our general framework introduces to this effect a new direct feature matching technique that finds global correspondences between images via simple nearest-neighbor searches. More specifically, very large image deformations are captured in *Spectral Forces*, which are derived from an improved graph spectral representation. We illustrate the benefits of our framework through a new enhanced version of the popular *Log-Demons* algorithm, named the *Spectral Log-Demons*, as well as through a groupwise extension, named the *Groupwise Spectral Log-Demons*, which is relevant for atlas construction. The evaluations of these extended versions demonstrate substantial improvements in accuracy and robustness to large deformations over the conventional Demons approaches.

Keywords Image Registration · Atlas Construction · Spectral Correspondence · Graph Laplacian

1 Introduction

The ability to capture complex image deformations and establish accurate pointwise correspondence is key to many computer vision applications that involve image

registration and atlas construction. Unfortunately, these properties become particularly challenging when the object depicted on the images undergoes a severe deformation or has in general a high shape variability. Current methods register images using Euler-Lagrangian approaches that slowly warp images until a satisfying overlap is attained. This is typically achieved by minimizing a deformation energy governed by laws of continuum mechanics that model elastic or viscous, or non-rigid, deformations. A state-of-the-art is surveyed in (Crum et al, 2004). Parametric models (Rueckert et al, 1999; Chui and Rangarajan, 2000) simplify the registration by restraining the minimization to a few parameters, while non-parametric models use the entire space of displacement fields. The minimization is now often restrained to the diffeomorphism group (Miller et al, 2002; Beg et al, 2005; Bossa et al, 2007; Allasonnière et al, 2007; Durrleman et al, 2011; Vercauteren et al, 2009), which consists of differentiable and reversible transformations. This prevents an invalid folding of the deformation field and guarantees a smooth one-to-one mapping between points. The computation of an average shape, or an average image (Studholme and Cardenas, 2004; Bhatia et al, 2004; Zollei et al, 2005), also often referred as atlas construction, is facilitated with the use of diffeomorphic transformations. They are key for the unbiased construction of atlases (Joshi et al, 2004; Avants and Gee, 2004; Marsland et al, 2003), where transformations are paths on a Riemannian manifold that represents the space of diffeomorphism. Geodesic paths in such space represent optimal transformations of images for constructing an atlas. They can be elegantly modeled using the Large Diffeomorphic Deformation Mapping Metric (LDDMM) framework (Beg et al, 2005; Beg and Khan, 2006; Bossa et al, 2007) or forward scheme approaches (Allasonnière et al, 2007; Durrleman et al, 2011). Guimond et al.

Herve Lombaert is at McGill University, Montreal, Canada and in the Asclepios Team at INRIA Sophia Antipolis, France
Leo Grady is at HeartFlow Inc.
Xavier Pennec and Nicholas Ayache are in the Asclepios Team at INRIA Sophia Antipolis, France
Farida Cheriet is at École Polytechnique de Montréal, Canada

(Guimond et al, 2000) proposed a fast and efficient algorithm for atlas construction that uses sequential image registrations. This approach is further refined in order to produce diffeomorphic transformations (Peyrat et al, 2007; Lombaert et al, 2011b, 2012c; Wu et al, 2011). The update schemes underlying these methods rely, however, on forces derived from the image gradients and are, therefore, fundamentally limited by their *local scope*. For example, gradients are null in textureless areas and optimization is undermined by local minima. A typical response to this limitation is to capture larger displacements on coarser versions of images. A multilevel scheme consequently propagates these larger displacements back to the original resolution. Preconditioning steps could also be used (Zikic et al, 2011). However, such approaches do not fundamentally solve the underlying problem: the scope of gradient-based updates remains local, even in coarser resolutions. This local scope typically limits current registration and atlas construction methods to small and local deformations between images. Large deformations were tackled in computer vision by matching region-based descriptors. For instance, SIFT Flow (Liu et al, 2008, 2011) finds a dense correspondence from local patch descriptors by approximating a discrete optimization. The Large Displacement Optical Flow (Brox et al, 2009; Brox and Malik, 2011) also uses region-based descriptors that produces a dense mapping whose non-convex optimization is facilitated with a hierarchical matching of super regions. Discrete optimization schemes (Glocker et al, 2011; Zikic et al, 2010; Shekhovtsov et al, 2007) have attractive optimality properties but remain, however, limited to finite sets of displacement vectors.

In order to capture very large and complex deformations, we introduce a new approach for image registration based on a *direct feature matching* technique that generates updates with a *global scope*. To do so, the pointwise correspondence between images is established with simple nearest-neighbor searches in a multidimensional space (Lombaert et al, 2011a) that comprises information on *image feature*, e.g., pixel intensities, *space*, e.g., Euclidean coordinates of pixels, and on *global image geometry*, e.g., geometric characteristics. Closest points in this multidimensional space highlight in fact the best compromise between these three properties. The simplicity of this technique allows the use of graph spectral representations (Chung, 1997) as matching features. Since they are invariant to isometry, when geodesic distances are preserved, objects with large and complex deformations would share similar spectral representations. Such strong property can be exploited to determine an accurate dense pointwise correspondence between images.

Spectral methods (Chung, 1997; Grady and Polimeni, 2010; Umeyama, 1988; Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992) are popular for general graph partitioning problems and are applied in computer vision for matching meshes (Jain and Zhang, 2006; Mateus et al, 2008; Reuter, 2009; Lombaert et al, 2011a, 2013a), point sets (Carcassoni and Hancock, 2000, 2003), and for representing shapes (Wilson et al, 2010; Konukoglu et al, 2012) and improving shape retrieval (Egozi et al, 2010). Pioneered in the late 80s, (Umeyama, 1988; Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992), spectral correspondence methods find shape similarities by comparing eigenvectors of a proximity matrix derived from the mesh structure. Recent work, surveyed in (Zhang et al, 2010; van Kaick et al, 2011), uses different types of proximity matrices, graph structures and spectrum deformations (Jain and Zhang, 2006; Mateus et al, 2008; Lombaert et al, 2011a, 2013a,b). However, the spectral representation of images, which is also strongly linked with Normalized Cuts (Shi and Malik, 2000; Meila and Shi, 2000; Weiss, 1999), have never been used for dense and accurate image registration. Diffeomorphism is, however, not guaranteed by spectral correspondence methods, nor by our direct feature matching technique, since simple nearest-neighbor searches could leave unassigned correspondences. Our approach consequently consists of performing our direct feature matching technique within a diffeomorphic framework for image registration or atlas construction, such as for instance the symmetric *Log-Demons* algorithm (Vercauteren et al, 2007, 2008). The resulting method, called **Spectral Log-Demons** (Lombaert et al, 2012a), enables a symmetric and diffeomorphic registration of images undergoing *large* and *complex* deformations. In fact, any algorithm that currently relies on *local* gradient-based updates could be adapted to use our new **Global Spectral Forces**. A preliminary version of *Spectral Log-Demons* has been published in (Lombaert et al, 2012a). Here, the method is fully explained and we provide extensive details as well as the general intuition behind the *Spectral Forces*. A second example is provided by extending the conventional symmetric Demons algorithm (Vercauteren et al, 2008) in order to perform groupwise registration (Lombaert et al, 2012b), i.e., the atlas is computed in parallel to the registration process rather than with a series of pairwise registrations. Similarly, this new method, named **Groupwise Log-Demons**, or *GL-Demons*, can be adapted to use Spectral Forces, yielding the **Groupwise Spectral Log-Demons** algorithm, or *GSL-Demons*.

The next section will describe our new direct feature matching technique, briefly review the *Log-Demons* al-

gorithm, describe its extension for groupwise registration in *GL-Demons*, and above all, explain how the new *Spectral Forces*, may be used for image registration and atlas construction, for instance in the *Spectral Log-Demons* and *GSL-Demons* algorithms. The subsequent evaluation will focus on illustrating and assessing the properties, advantages and main lines of the new algorithms. The experiments will effectively show that this fundamentally new approach can naturally capture very large and complex deformations and can additionally demonstrate substantial improvements over conventional approaches.

2 Methods

Our strategy consists of exploiting the *global scope* and the speed of nearest-neighbor search methods for the purpose of capturing very large deformations between images. We begin our methodology with our simple and direct feature matching technique followed by how spectral representations can be utilized as geometrical features. Finally, we explain how our new extended direct feature matching technique can be used within diffeomorphic frameworks for image registration and atlas construction of images 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 . In this paper, we interpret the registration problem as a *direct* feature matching problem, where feature vectors \mathbf{F} and \mathbf{M} , representing images F and M , are matched via an unknown correspondence map ϕ . For instance, features can include image intensity and spatial information: $\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F)$ and $\mathbf{M} = (\alpha_i I_M, \alpha_s \mathbf{x}_M)$, where $I(\cdot)$ is a pixel intensity and $\mathbf{x}(\cdot) = (x, y)$ is a point coordinate, both weighted with parameters $\alpha_{i,s}$.

Our general matching problem consists of finding an optimal correspondence map:

$$\phi(i) = \underset{j \in M}{\operatorname{argmin}} \|\mathbf{F}(i) - \mathbf{M}(j)\|^2, \quad (1)$$

where $\mathbf{F}(i)$, the multidimensional characteristic of point $i \in F$, should correspond to $\mathbf{M}(\phi(i))$, the most similar characteristic found in image M . This is also written as $\mathbf{F} \mapsto \mathbf{M} \circ \phi$ and can be *directly* solved with a nearest-neighbor search between \mathbf{F} and \mathbf{M} with, for instance, a Voronoi tessellation or a k -d tree. If point $j \in M$ is found to have the closest characteristic to that of point $i \in F$, then $\phi(i) = j$. This effectively provides similarity

in pixel intensity and closeness in space between corresponding pixels, as illustrated in Fig. 1. It minimizes the similarity criterion:

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

where $\|I_F - I_{M \circ \phi}\|^2 = \sum_{i \in F} (I_F(i) - I_M(\phi(i)))^2$ is the sum of intensity differences, and $\|\mathbf{x}_F - \mathbf{x}_{M \circ \phi}\|^2 = \sum_{i \in F} \|\mathbf{x}(i) - \mathbf{x}(\phi(i))\|_{L_2}^2$ integrates the L_2 norms of the displacements between corresponding points, which effectively acts as a form of spatial regularization.

However, such criterium lacks information on the intrinsic image geometry: objects in images can have different poses or be severely deformed, in which case, equivalent points on deformed objects will result in very different Euclidean coordinates. This is illustrated on Fig. 1 with the wrong correspondence link, in red. Additionally, a nearest-neighbor search does not guarantee a one-to-one mapping, therefore, the correspondence map is not diffeomorphic. Each issue is addressed below.

2.2 Spectral Correspondence

The spectral representation of shapes (Chung, 1997; Grady and Polimeni, 2010; Umeyama, 1988; Scott and Longuet-Higgins, 1991; Shapiro and Brady, 1992; Jain and Zhang, 2006; Mateus et al, 2008; Lombaert et al, 2011a) can be regarded as a unique shape description or signature, and has the strong property of being invariant to *isometry*, i.e., corresponding points between shapes in different poses would share the same shape coordinates, called here *spectral coordinates*, even if these points are far away in space. For instance, a point on a nose tip has a unique geometric description even if it moves in space. By adding these spectral coordinates in our feature space, we enforce an intrinsic *geometric consistency* in our matching technique.

2.2.1 Spectral Graph Theory

The connected undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is constructed with the vertices \mathcal{V} representing pixels of an image I_Ω , bounded by a contour Ω , 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 (Grady and Polimeni, 2010) where high weights are given to edges within a region of uniform intensity and low weights are given to edges crossing region boundaries. For instance, $W_{i,j} = \exp(-(I(i) - I(j))^2 / 2\sigma^2) / \|\mathbf{x}(i) - \mathbf{x}(j)\|^2$ if $(i, j) \in \mathcal{E}$ and 0 otherwise. The parameter σ depends on the image

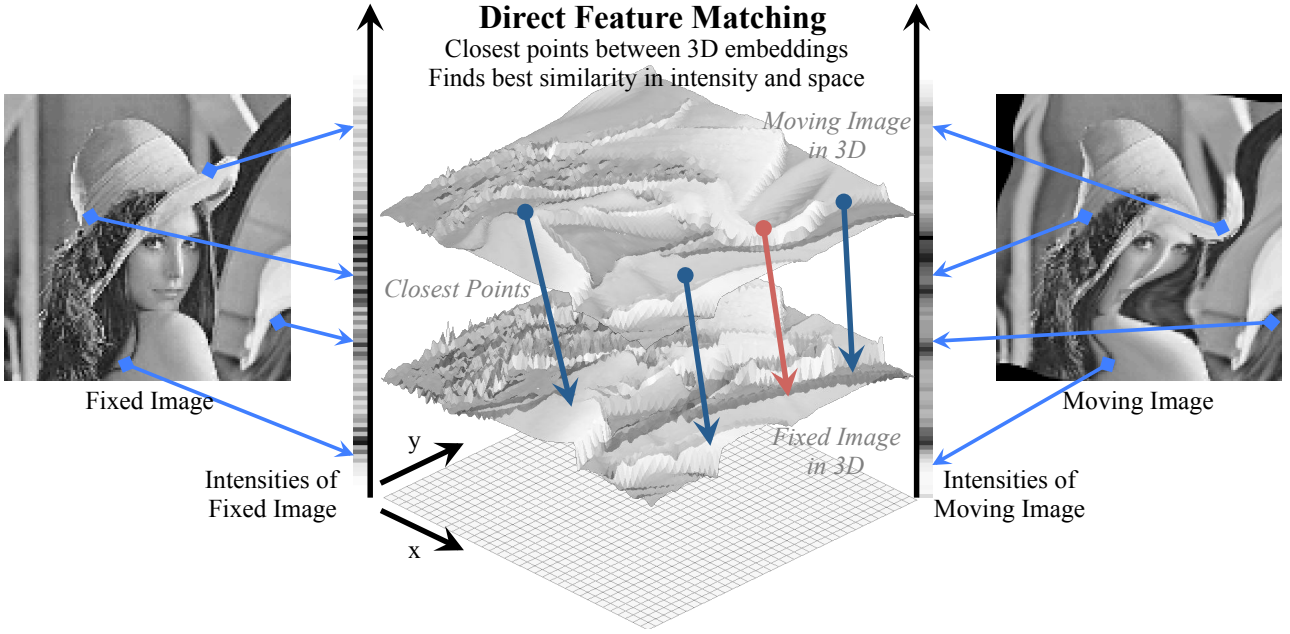


Fig. 1 Direct Feature Matching – Correspondences are established with simple and fast nearest neighbor searches between multidimensional embeddings. Here, both images are represented in 3D (intensity channel and Euclidean coordinates x, y). Closest points reveal in fact the best compromise between *intensity* and *spatial location*. Note that any additional features can be used, for instance, spectral coordinates (resulting with the best compromise between *intensity*, *space*, and *geometry*)

noise and may be set with $\sigma = \text{mean}\{|I(i) - I(j)|\}_{(i,j) \in \mathcal{E}}$. 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 $\mathcal{L} = G^{-1}(D - W)$ (Grady and Polimeni, 2010), which is a $|\mathcal{V}| \times |\mathcal{V}|$ sparse matrix where G is the diagonal node weighting matrix, e.g., $G = D$.

2.2.2 Spectral Coordinates

The graph spectrum (Chung, 1997) computed from the decomposition of the general Laplacian $\mathcal{L} = \mathbf{X}^T \Lambda \mathbf{X}$ comprises the eigenvalues, in increasing order, $\Lambda = \text{diag}(\lambda_0, \lambda_1, \dots, \lambda_{|\mathcal{V}|})$ and their associated eigenvectors $\mathbf{X} = (\mathcal{x}^{(0)}, \mathcal{x}^{(1)}, \dots, \mathcal{x}^{(|\mathcal{V}|)})$. Here, \mathbf{X} is a $|\mathcal{V}| \times |\mathcal{V}|$ sparse matrix where each column $\mathcal{x}^{(\cdot)}$ is an eigenvector. The first eigenvector $\mathcal{x}^{(0)}$ is the stationary distribution, which is also 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 for the shape depicted by I_Ω with free ends. We thus prefer the term eigenmode since they are effectively functions over I_Ω . They are visualized as images, as shown in Fig. 2. 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, which are lines where the eigenmodal value is 0. They reside on prominent demarcations of the shape geometry. For instance, graph-based segmentation methods (e.g., (Shi and Malik, 2000; Meila and Shi, 2000)) rely on the nodal set of $\mathcal{x}^{(1)}$, called the Fiedler vector (Chung, 1997), to find a binary partition of an image. Additional nodal sets may also be considered (Ding and He, 2004). The number 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 $\mathcal{x}^{(1..k)}$, which correspond to 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 $\mathcal{x}(i) = (\mathcal{x}^{(1)}(i), \mathcal{x}^{(2)}(i), \dots, \mathcal{x}^{(k)}(i))$. In other words, $\mathcal{x}(i)$ is a truncated line of matrix \mathbf{X} . The spectral representation \mathcal{x} has the strong property of being quasi invariant to isometry: 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, which are weighted with α_g . A nearest-neighbor search between

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}

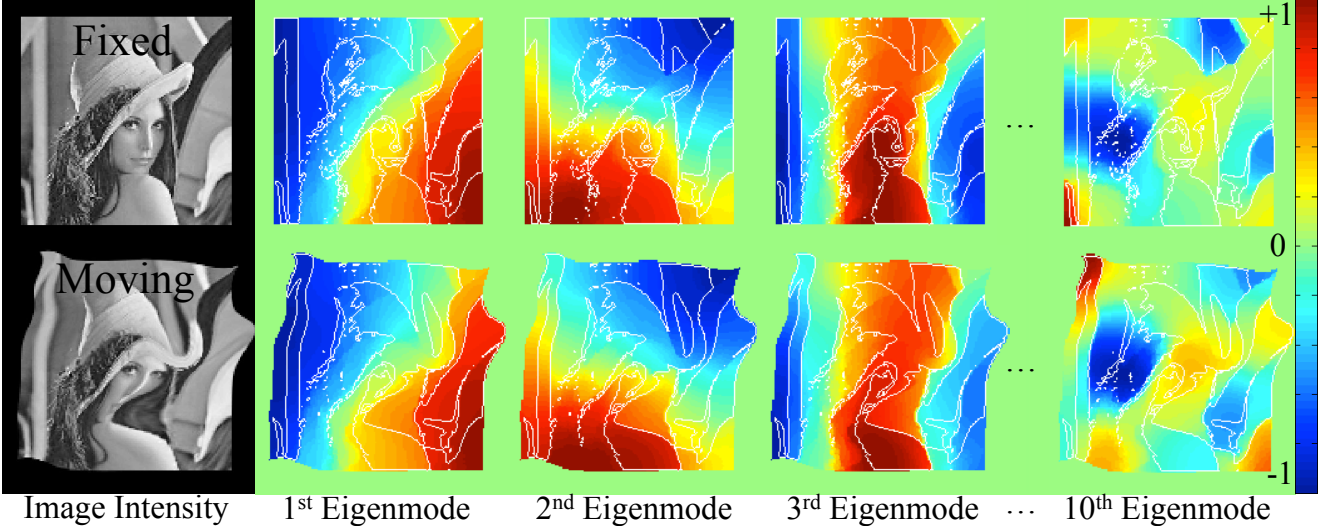


Fig. 2 Spectral Coordinates – The eigenmodes of the graph Laplacian are used as geometric descriptors, comparable to a set of unique shape coordinates. Each column shows one component, or dimension, of the spectral coordinates, and remains stable under complex deformations (invariant to isometry). The lower eigenmodes describe coarse geometric properties, while higher eigenmodes describe finer details in the images. Note the similarity of the coordinate values between corresponding points (the image contours are overlaid in white for comparison).

$\mathbf{F} = (\alpha_i I_F, \alpha_s \mathbf{x}_F, \alpha_g \mathcal{X}_F)$ 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:

$$\begin{aligned} \text{Sim}(F, M, \phi) = & \|I_F - I_{M \circ \phi}\|^2 \\ & + \frac{\alpha_s^2}{\alpha_i^2} \|\mathbf{x}_F - \mathbf{x}_{M \circ \phi}\|^2 \\ & + \frac{\alpha_g^2}{\alpha_i^2} \|\mathcal{X}_F - \mathcal{X}_{M \circ \phi}\|^2, \end{aligned} \quad (3)$$

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 (Thusty, 2010) which is in this case the multiplicity of the Fiedler vector n_{λ_1} and is also related (Thusty, 2010) to the maximal dimension of a space in which the graph \mathcal{G} can be mapped. 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.

2.2.3 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 $-\alpha \mathcal{x}^{(\cdot)}$.

This 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, since 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, in fact, 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 (Mateus et al, 2008; Lombaert et al, 2011a), minimizes the following 2D dissimilarity matrix:

$$\begin{aligned} C(u, v) = & \sqrt{\frac{1}{|I_\Omega|} \sum_{i \in I_\Omega} \left(\mathcal{x}_F^{(u)}(i) - \mathcal{x}_M^{(v)}(i) \right)^2} \\ & + \sqrt{\sum_{i, j} \left(h_F^{(u)}(i, j) - h_M^{(v)}(i, j) \right)^2} \end{aligned} \quad (4)$$

The first term is the difference in eigenmodal values between the images and, the second term measures

the dissimilarities between the joint histograms $h(i, j)$, which is a 2D matrix where an 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 $Q(u, v) = \min\{C(u, v), C(u, -v)\}$, where $C(u, -v)$ is the cost of pairing eigenvector $\mathcal{X}^{(u)}$ with the negative $-\mathcal{X}^{(v)}$. 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 is capable of providing unique geometric features in our new direct feature matching technique, which effectively add quasi invariance to changes in pose and to shape deformations when establishing correspondences. Small perturbations in shape isometry, which may be caused by small local distortions, are responsible for differences in spectral coordinates between shapes. Fortunately, these differences become smaller when images are warped together. The direct feature matching technique may also generate a many-to-one mapping between points and leave points with unassigned correspondences. Diffeomorphism is, therefore, not guaranteed since invalid folding of space may be created by the correspondence map. We now explain how these issues can be addressed in order to register images and construct atlases.

2.3 Pairwise Registration

In a pairwise registration, one moving image M is non-rigidly aligned to a fixed image F , and establishing correspondences is crucial in the registration process. The direct feature matching technique described earlier is able to capture very large deformations, but is however not able to guarantee a diffeomorphic transformation. Conventional registration methods and our direct feature matching technique may benefit from each other in many aspects. *Firstly*, conventional registration methods may be improved in order to capture very large deformations by simply replacing their update schemes, which are typically based on image gradients, with our direct feature matching technique, which uses spectral correspondence. The updates derived from our approach have, therefore, a *global* scope instead of a typical *local* scope. They capture global geometrical similarities rather than directions defined within local a neighborhood. *Secondly*, our technique may generate better correspondences if the alignment between images is refined through iterations. Both images would ultimately become identical in an optimal registration method, which would, therefore, result in identical spectral coordinates. For these reasons, we modify one con-

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
 $W_{ij} = \exp\left(-(I(i) - I(j))^2 / 2\sigma^2\right) / \|\mathbf{x}(i) - \mathbf{x}(j)\|^2$
 $D_{ii} = \sum_j W_{ij}$,
 - Compute first k eigenmodes of Laplacians
 - Reorder \mathcal{X}_M with respect to \mathcal{X}_F (Eq. 4)
 - 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 $\mathbf{F}(i) \mapsto \mathbf{M}(c(i))$
-

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 General 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 $M \circ \exp(v)$
and updates $u_{M \rightarrow F}$ mapping M to $F \circ \exp(-v)$
(for Conventional *Log-Demons*: use Eq. 6)
(for *Spectral Log-Demons*: use Alg. 1)
 - 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
-

ventional registration method in order to perform spectral correspondence. As an example, we enhance the established *Log-Demons* algorithm (Vercauteren et al, 2007), which is simple and intuitive. Note, however, that our approach could be applied to other algorithms.

2.3.1 Log-Demons

The *Log-Demons* algorithm is now briefly reminded. From the theory of Lie groups, a diffeomorphic transformation ϕ , which resides on a Lie group structure, is related to the exponential map of a velocity field v , its

associated Lie algebra: $\phi = \exp(v)$. In the case of stationary velocity fields, a practical and fast approximation is possible with the scaling-and-squaring method (Vercauteren et al, 2007) (Alg. 2). Additionally, the inverse of the transformation is simply $\phi^{-1} = \exp(-v)$.

The *Log-Demons* framework alternates 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, that is the correspondences c . This allows a small error between alternations, e.g., $\text{dist}(c, \phi) = \|c - v\|$. The analogy with Maxwell’s demons is made by considering small demons, sitting on each pixel, that pull and push forces in order to deform the moving image (Pennec et al, 1999). Additionally, invariance to the order of the input images is possible with the symmetric extension of the algorithm (Vercauteren et al, 2008, 2009). The energy of the symmetric *Log-Demons* may be written:

$$\begin{aligned} E(F, M, \exp(c), \exp(v)) = & \\ & + \frac{1}{2} \alpha_i^2 (\text{Sim}(F, M \circ \exp(c)) + \text{Sim}(F \circ \exp(-c), M)) \\ & + \alpha_x^2 \text{dist}(c, v)^2 \\ & + \alpha_T^2 \text{Reg}(v), \end{aligned} \quad (5)$$

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 (Vercauteren et al, 2007) for more details):

$$u_{F \rightarrow M} = - \frac{I_F - I_{M \circ \phi}}{\|\nabla I_{M \circ \phi}\|^2 + \alpha_x^2 \|I_F - I_{M \circ \phi}\|^2} \nabla I_{M \circ \phi}. \quad (6)$$

It can be shown (Cachier et al, 2003) that the alternating minimization of Eq. 5 can be achieved at relatively low cost with convolution kernels $K_{\text{fluid}, \text{diffusion}}$, e.g., Gaussian with standard deviations σ_{fluid} and $\sigma_{\text{diffusion}}$. In a first step, the transformation $\phi = \exp(v)$ is fixed and the correspondence update u is computed with Eq. 6. In the second step, the stationary velocity field v of the transformation is computed via the intermediate regularization of the correspondence update $u \leftarrow K_{\text{fluid}} \star u$ and of the updated velocity field $v \leftarrow K_{\text{diffusion}} \star (v + u)$. The *Log-Demons* regularization can also be formulated as a Tikhonov problem, in which an optimal velocity field can be obtained by filtering correspondences in the Fourier domain (Mansi et al, 2011). The weighting α_i, α_x and α_T are in practice dictated (Mansi et al, 2011) by these regularization parameters σ_{fluid} and $\sigma_{\text{diffusion}}$. The general symmetric diffeomorphic Demons framework is summarized in Alg. 3.

2.3.2 Spectral Forces and Spectral Log-Demons

The correspondence update u is, however, limited by the local scope of the image gradient. This limitation is in fact also true for most conventional registration algorithms that base their iterative scheme on the image gradient, such as all current Demons variants. They typically circumvent this limitation by using a multiresolution scheme, even though this does not change the local scope of image gradients in the coarser resolutions.

The conventional *Log-Demons* algorithm is, therefore, extended in order to perform spectral correspondence. The so-called *Spectral Log-Demons* algorithm (Lombaert et al, 2012a) takes advantage of the efficient general diffeomorphic *Log-Demons* framework and finds the correspondences between images F and M using spectral correspondence. To be more precise, the first two steps of Alg. 3, which originally compute updates from image gradients with Eq. 6, now perform spectral correspondence between images F and $M \circ \exp(v)$. The correspondence update $u_{F \rightarrow M}$ is simply defined as the displacement vectors created by the output of Alg. 1. It drives the registration with a different strategy, and is, therefore, referred as *Spectral Forces* in the *Spectral Log-Demons* algorithm. The symmetric correspondence between M and $F \circ \exp(-v)$ is computed similarly by setting $u_{M \rightarrow F}$ with the output of Alg. 1. 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, provided that they are invariant to isometry, between images as well as a faster convergence of the algorithm.

The energy underlying the *Spectral Log-Demons* has the form of the general Eq. 5 where its similarity term is expressed with Eq. 3:

$$\begin{aligned} E(F, M, \exp(c), \exp(v)) = & \\ & + \frac{\alpha_i^2}{2} (\|I_F - I_{M \circ \exp(c)}\|^2 + \|I_{F \circ \exp(-c)} - I_M\|^2) \\ & + \frac{\alpha_s^2}{2} (\|\mathbf{x}_F - \mathbf{x}_{M \circ \exp(c)}\|^2 + \|\mathbf{x}_{F \circ \exp(-c)} - \mathbf{x}_M\|^2) \\ & + \frac{\alpha_g^2}{2} (\|\mathcal{X}_F - \mathcal{X}_{M \circ \exp(c)}\|^2 + \|\mathcal{X}_{F \circ \exp(-c)} - \mathcal{X}_M\|^2) \\ & + \alpha_x^2 \text{dist}(c, v)^2 \\ & + \alpha_T^2 \text{Reg}(v), \end{aligned} \quad (7)$$

where the parameters $\alpha_{i,s,g}$ control the consistency in intensity, space and geometry, and $\alpha_{x,T}$ are the traditional Demons parameters controlling the step size and the regularization. In practice, $\alpha_i = 0.8$ is general enough for a broad range of images with intensities

ranging between $[0; 1]$. The spatial and geometrical consistency, $\alpha_s = 0.15$ and $\alpha_g = 0.05$, prevents too small or too large displacements during a single iteration, which produces, in practice, a faster warping of images toward their registration. We show, later in Sec. 3.1.3, that best combinations actually have ratios of parameters between $1/10 < \alpha_g/\alpha_s < 1/3$.

2.3.3 Multilevel Scheme

A multilevel scheme could be motivated with the assumption that large deformations are often related with coarse geometric information. The correspondence update u can, therefore, be computed with spectral correspondence in the lower levels of resolution, while finer details and local deformations can be computed more efficiently with the conventional gradient-based updates in the higher levels of resolutions. The eigenmodes are consequently computed only in the coarser levels of resolution, on smaller images, and therefore, in a faster time. On the same note, the computation of the eigenmodes can be used with the efficient Lanczos method, which is used by Matlab. It has a running time of $O(n\sqrt{n}) + O(n^2)$ (Shi and Malik, 2000), 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 $O(\log n)$.

2.4 Groupwise Registration

In a groupwise registration, a set of N images $\{I_i\}_{i=1..N}$ is registered in parallel with a reference shape \tilde{I} that evolves until a convergence is reached on an average shape \tilde{I} . In order to further illustrate the benefits of using *Spectral Forces*, within conventional registration methods, we extend the symmetric *Log-Demons* algorithm (Vercauteren et al, 2008) in order to perform groupwise registration (Lombaert et al, 2012b) by using either classical gradient-based updates (in the *Groupwise Log-Demons*, or *GL-Demons*), or *Spectral Forces*, (in the *Groupwise Spectral Log-Demons*, or *GSL-Demons*).

2.4.1 Groupwise Log-Demons

Our new shape averaging framework is based on Guimond's *et al.* approach (Guimond et al, 2000) where they construct the average image \tilde{I} *sequentially* by alternating between pairwise registrations (fixing a reference image) and updates of the average image (transforming the reference image). Our novelty is to directly compute \tilde{I} *in parallel* with the simultaneous, or groupwise, registrations (illustrated in Fig. 3). To do so, in a first step, the symmetric *Log-Demons* algorithm (Alg. 3)

is slightly modified such that both the fixed and moving images, F and M , are deformed at the same time until they both converge toward an average shape. The original Eq. 5 is, therefore, changed for:

$$\begin{aligned} E(F, M, c, v) = & \alpha_i^2 \text{Sim}(F', M') \\ & + \alpha_x^2 \text{dist}(c, v)^2 \\ & + \alpha_T^2 \text{Reg}(v), \end{aligned} \quad (8)$$

where $\text{Sim}(F', M') = \|F' - M'\|^2$, $\text{dist}(c, v) = \|c - v\|$, and $\text{Reg}(v) = \|\nabla v\|^2$. Both images F' and M' are effectively mutually converging toward an average shape $\tilde{I} = F \circ \phi^{-1} + M \circ \phi$, in a similar approach than (Avants and Gee, 2004; Bossa et al, 2007). The similarity term incorporates diffeomorphism and symmetry with $F' = F \circ \exp(-c)$ and $M' = M \circ \exp(+c)$.

In a second step, Eq. 8 is extended to incorporate N velocity fields that warp all images $\{I_i \circ \exp(c_i)\}_{i=1..N}$ toward the dynamic reference image \tilde{I} . The new groupwise framework is summarized in Alg. 4 and its general energy is:

$$E(\tilde{I}, \{I_i, c_i, v_i\}) = \frac{1}{N} \sum_{i=1}^N \left(\begin{aligned} & \alpha_i^2 \text{Sim}(\tilde{I}, I_i \circ \exp(c_i)) \\ & + \alpha_x^2 \text{dist}(c_i, v_i)^2 \\ & + \alpha_T^2 \text{Reg}(v_i) \end{aligned} \right) \quad (9)$$

The reference image can be optionally generated with weighted contributions from all images. For instance, weights can be different than $1/N$, for instance, to remove outliers. The minimization of all similarity terms, $\{\text{Sim}(\tilde{I}, I'_i)\}$, causes all warped images to become similar to the reference image and the sum of all velocity fields is brought to a minimal value at convergence. Similar to the convergence of (Guimond et al, 2000), the Groupwise Demons framework effectively brings the reference image to the barycenter of all images. The dynamic reference image is simply generated with $\tilde{I} = \frac{1}{N} \sum_{i=1}^N I_i \circ \exp(c_i)$.

2.4.2 Groupwise Spectral Log-Demons

The Groupwise Demons framework, which has been so far described as using the conventional gradient-based updates, can also be extended to use *Spectral Forces*. Similarly to the *Spectral Log-Demons*, the *Groupwise Spectral Log-Demons* algorithm (*GSL-Demons*) uses the same Alg. 4 with a different mapping function, where updates are computed with Alg. 1. *GSL-Demons* enable, therefore, large jumps during the computation of the average shape, and allow points to be registered toward their isometric equivalents even if they are far away in the image. The use of this spectral approach




Fig. 3 Groupwise Demons – Simultaneous registration of 4 images (blue circles) toward a reference image that evolves in the space of diffeomorphisms (colored manifold). The reference image is computed in parallel and converges to the average shape (middle red circle).

Algorithm 4 Groupwise Demons Framework

Input: N images with initial reference (e.g., $\tilde{I} = I_1$)
Output: Transformations $\phi_i = \exp(v_i)$ mapping \tilde{I} to I_i
Average shape is $\tilde{I} = \frac{1}{N} \sum_{i=1}^N I_i \circ \exp(v_i)$

repeat
for $i = 1 \rightarrow N$ **do**

- Find updates $u_i \leftarrow \text{mapping}(\tilde{I}, I_i \circ \exp(v_i))$.
(mapping() differs in *GL* and *GSL-Demons*)
(using respectively Eq. 6 or Alg. 1)
- Smooth updates: $u_i \leftarrow K_{\text{fluid}} \star u_i$.
(convolution of a Gaussian kernel on u_i)
- Update velocity fields: $v_i \leftarrow \log(\exp(v_i) \circ \exp(u_i))$
(approximated with $v_i \leftarrow v_i + u_i$).
- Smooth velocity fields: $v_i \leftarrow K_{\text{diff}} \star v_i$.

end for

- Get reference update: $u_{\text{ref}} = -\frac{1}{N} \sum_{i=1}^N v_i$
- Update velocity fields: $v_i \leftarrow v_i + u_{\text{ref}}$.
- Update reference: $\tilde{I} \leftarrow \frac{1}{N} \sum_{i=1}^N I_i \circ \exp(v_i)$.

until convergence

permits the construction of atlases with images presenting very large deformations, and with fewer iterations. Typically 5 iterations are sufficient. The energy has the same form of Eq. 9 and uses the similarity term of Eq. 3:

$$E(\tilde{I}, \{I_i, c_i, v_i\}) = \frac{1}{N} \sum_{i=1}^N \left(\begin{aligned} &\alpha_i^2 \|\tilde{I} - I_i \circ \exp(c_i)\|^2 \\ &+ \alpha_s^2 \|\mathbf{x}_{\tilde{I}} - \mathbf{x}_{I_i \circ \exp(c_i)}\|^2 \\ &+ \alpha_g^2 \|\mathcal{X}_{\tilde{I}} - \mathcal{X}_{I_i \circ \exp(c_i)}\|^2 \\ &+ \alpha_x^2 \text{dist}(c_i, v_i)^2 \\ &+ \alpha_T^2 \text{Reg}(v_i) \end{aligned} \right) \quad (10)$$

Additionally, a multilevel approach is also possible where large and complex deformations can be captured in a low resolution level with *GSL-Demons*, improving thus the processing time, and where the remaining

small and local deformations can be recovered with *GL-Demons* in higher resolutions.

3 Results

The evaluation focuses on illustrating the properties and advantages of using our spectral approach in conventional registration methods. In controlled experiments, the full power of our *Global Spectral Forces*, is demonstrated in a pairwise and groupwise configurations (*Spectral Log-Demons* and *GSL-Demons*) where images with drastic deformations are used. We also explore the impact of changing the algorithm parameters on registration. The improvement in registration accuracy, robustness to deformation and noise are all accurately assessed with ground truth data. Furthermore, we present practical applications with the construction of medical atlases of organs exhibiting a large shape variability.

3.1 Pairwise Registration

In the first controlled experiment, we evaluate the fundamental difference of our novel update scheme within the same registration framework, by comparing the *Log-Demons* and *Spectral Log-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 parameters α_i, α_x and α_T , which are in practice defined by $\sigma_{\text{fluid}, \text{diffuse}}$, were first chosen by obtaining an optimized registration with the conventional *Log-Demons* algorithm. All algorithms use subsequently the same parameters: $\sigma_{\text{fluid}, \text{diffuse}} = 1, \alpha_x = 1$. For now, we use $k = 2$ eigenmodes with the *Spectral Log-Demons*, as well as the following weighting parameters: $\alpha_g = 0.05, \alpha_s = 0.15, \alpha_i = 0.8$, which favor consistency in geometry and intensity, while avoiding too large displacements in a single iteration. We register the images on




Fig. 4 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.

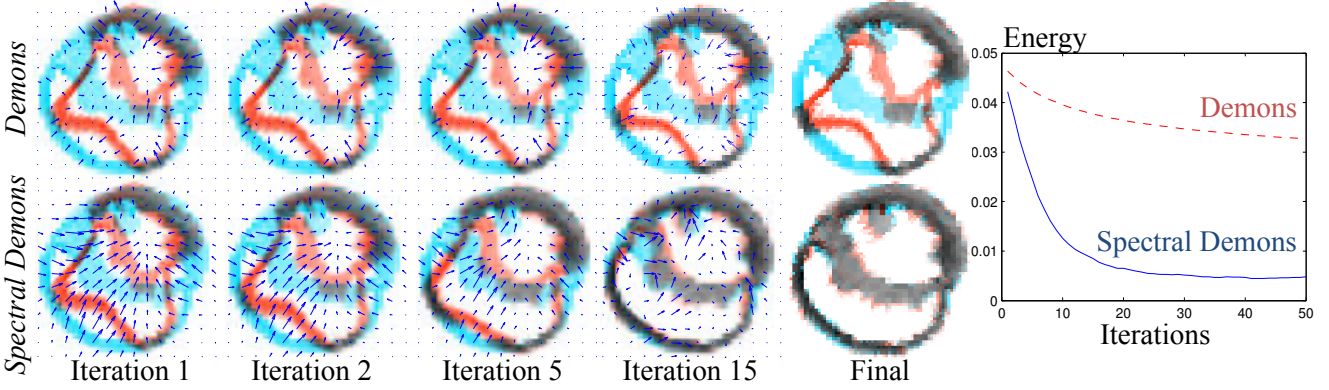


Fig. 5 Comparison of the iterations in the *Log-Demons* and the *Spectral Log-Demons* (our method) within the same resolution level. The updates (indicated with the arrows and scaled for visualization) have a local scope with the *Log-Demons* and a global scope with the *Spectral Log-Demons*. This global scope allows a faster convergence while the *Log-Demons* remains in a local minimum.

Fig. 4: 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 at control points randomly located, and diffused across the image with a 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. 4) severely deformed and almost fully collapsed. The muscle forms, however, a dent in the image and provides a signature that *Spectral Log-Demons* can understand.

The iterations of the *Log-Demons* and *Spectral Log-Demons* are compared in Fig. 5. It shows that within the same level of resolution, the updates computed with spectral correspondence are coherent spatially and geometrically, i.e., points move toward their geometric equivalent, however the updates 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 approach allows the *Log-Demons* to capture larger deformations, as shown in the second row of Fig. 6, but does not change the inherent local scope of its update scheme. For instance,

the *Log-Demons*, even with 4 levels of resolution, ultimately fails in recovering the extreme deformations on the anterior side of the heart, while the *Spectral Log-Demons*, without a multilevel scheme, can successfully register the whole myocardium with a 71% improvement in performance. This is observed with the decrease of the mean squared differences (or MSE) of intensities with the ground truth from 19.9×10^{-3} to 5.7×10^{-3} . The performance is further improved when 4 levels are used. In fact, the MSE decreases down to 2.4×10^{-3} . This is 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 Log-Demons* with respectively 73% improvement over *Log-Demons* in Lena’s image and 63% in the baseball image.

For comparison purposes, we also registered these images with the LDDMM algorithm (Beg et al, 2005). We used the implementation of (Risser et al, 2011), with 32 time steps and kernel bandwidth $\sigma = 10$ for Lena and the baseball player, and $\sigma = 3$ for the cardiac image. These settings led to the registrations shown on the top row of Fig. 6. The LDDMM based similar observations: areas of very large deformations were mismatched. By varying the kernel bandwidth from $\sigma = 0$ to $\sigma = 20$, Lena’s hat, the thinner myocardial wall, and the baseball player’s arm could not be correctly regis-

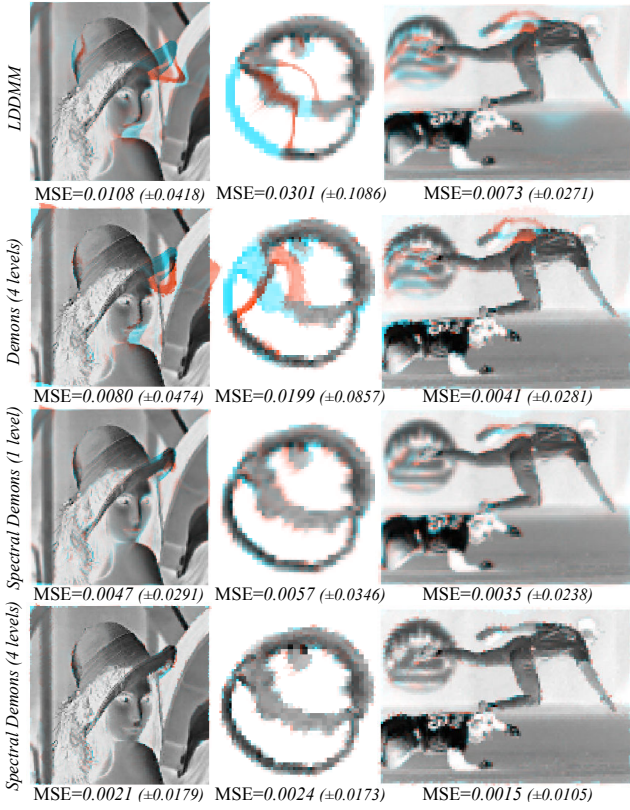


Fig. 6 Final registrations for (1st row) LDDMM, (2nd row) Multilevel *Log-Demons* with 4 levels, (3rd row) *Spectral Log-Demons* with 1 level, (4th row) Multilevel *Spectral Log-Demons* with 4 levels. The fixed image is in blue and misalignments of the registered image are in red. The mean squared differences of intensities are reported along their standard deviations. *Log-Demons* is limited in areas of high deformations, while *Spectral Log-Demons* can capture these large deformations.

tered. One reason is that the underlying optimization in LDDMM also relies on the direction of image gradients. Again, both *Log-Demons* and LDDMM are inherently limited by the local scope of their gradient-based update schemes, which may be inadequate to drive the registration within areas of very large deformations. However, the *Spectral Log-Demons* could naturally capture such large deformations since its underlying update scheme exploits global geometric characteristics rather than locally defined image gradients. This experiment demonstrates the effectiveness of using a spectral approach for very large deformations.

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* and LDDMM result in registration maps, shown in Fig. 7, with larger errors in high deformation areas, such as in Lena’s hat or neck, whereas the *Spectral Log-Demons* results in a smoother registration




Fig. 7 Differences (in pixels) between computed transformations and ground truth for (1st row) LDDMM, (2nd row) Multilevel *Log-Demons*, (3rd row) *Spectral Log-Demons* and, (4th row) Multilevel *Spectral Log-Demons*. The fixed images are overlaid on the error maps. The Multilevel *Spectral Log-Demons* decreases the error by 60% from the Multilevel *Log-Demons*.

map with significantly less errors, with a 62% decrease in these same areas.

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

3.1.1 Robustness to Deformation

The robustness to deformation of the *Spectral Log-Demons* algorithm is evaluated on Lena’s image by exaggerating the previous synthetic transformation from a zero displacement, $\phi_0 = \exp(0v)$, to $\phi_2 = \exp(2v)$, which creates a maximal displacement of 40 pixels (see samples on Fig. 8). The performance is evaluated with the trans-




Fig. 8 Robustness to deformation and noise. *a)* Deformations of Fig. 4 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 Log-Demons* (blue) than with *Log-Demons* (red). *b)* Gaussian noise is used (samples around the hat area). *Spectral Log-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 Log-Demons* continues to recover large deformations.

formation differences, in pixels, from the ground truth. Both *Log-Demons* and *Spectral Log-Demons* perform with sub-pixel accuracy on deformations below 20 pixels, however they differ with larger deformations where our method shows a greater robustness. With a deformation of 40 pixels, which is more than 30% of the image size, the average transformation error is 5.9 pixels with *Log-Demons* and 1.6 pixels with *Spectral Log-Demons*. That is a 73% decrease over the conventional *Log-Demons* algorithm.

3.1.2 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. 8). *Spectral Log-Demons* performs better with noise $\sigma < 0.075$, however the comparison with control images, which are the unregistered noisy images, reveals that *Log-Demons* stops transforming the images when noise is increased. The registration is almost immediately trapped in a local minimum. Fig. 8 shows that when using *Log-Demons* with noise $\sigma = 0.13$, the registered image is similar to its initial state. On the contrary, when using the *Spectral Log-Demons*, the registration recovers the large deformations, even though the average displacement error is larger (6.8 pixels for *Spectral Log-Demons* versus 3.7 pixels *Log-Demons*). This is noticeable in the hat area on Fig. 8. With noise $\sigma > 0.2$, the corrupted images become problematic for *Spectral Log-Demons*, whereas *Log-Demons* is almost immediately trapped in a local minimum. With such noise levels, graph edge

weights are almost null and may need a different heuristic weighting.

3.1.3 Weighting Parameters

The *Spectral Log-Demons* algorithm has several regularization parameters: $\alpha_{i,s,g,x,T}$. The latter parameters, $\alpha_{x,T}$, come from the conventional *Log-Demons* algorithm and control smoothing of transformations. They correspond to the regularization terms in Eq. 5 and Eq. 7, and are in practice dictated by the Gaussian widths σ_{fluid} and $\sigma_{\text{diffusion}}$ as described in (Mansi et al, 2011). The data term in Eq. 7 is controlled by $\alpha_{i,s,g}$, which respectively correspond to a weighting on intensity, spatial coherence and geometrical coherence.

The effect of these three new parameters $\alpha_{i,s,g}$ is studied by sweeping their values from 0 to 1 and measuring the registration accuracy in terms of intensity MSE between the registered and fixed image image. Noting that any combination of these three parameters could be expressed with an equivalent normalization $\alpha_i^2 + \alpha_s^2 + \alpha_g^2 = 1$, the parameter sweep is in fact forming the positive quadrant of a sphere with radius 1, as shown in Fig. 9. For instance, the combination used previously, $\alpha_{i,s,g} = (0.8; 0.15; 0.05)$, is equivalent to $\alpha_{i,s,g} = (0.98, 0.18, 0.06)$, up to a scaling factor, which leads to an $MSE = 4.7 \times 10^{-3}$. The coloring on Fig. 9 shows the registration accuracy for all possible combinations of $\alpha_{i,s,g}$ on Lena’s image. The worst and best registrations are shown with their corresponding combinations of $\alpha_{i,s,g}$ and their resulting registered images. This experiment shows that the weighting on spatial coherence should be low but non-zero. Having a stronger weighting actually degrades the registration




Fig. 9 Registration Accuracy with different combinations of $(\alpha_i, \alpha_s, \alpha_g)$. Coloring indicates the mean squared differences of intensities between registered and fixed Lena images for various ratios of parameters α . Blue area shows that best combinations should have ratios of $\frac{1}{10} < \alpha_s/\alpha_g < \frac{1}{3}$.

accuracy, since α_s penalizes large deformations. Best parameter combinations are shown to have ratios between $1/10 < \alpha_s/\alpha_g < 1/3$.

3.1.4 Registration of Medical Images

The performance of the *Spectral Log-Demons* is evaluated in a medical application with the registration of brain MR images from the Internet Brain Segmentation Repository (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 (Fischl et al, 1999; Yeo et al, 2010; Reuter, 2009; Lombaert et al, 2011a), 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. See for instance Fig. 10 where 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 Log-Demons*, however the comparison of the

mean squared differences of intensities reveals a 4% improvement in accuracy and precision when using *Spectral Log-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 Log-Demons*). This experiment showed that *Spectral Log-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 Log-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.

3.1.5 Note on Image Segmentation

The Laplacian eigenmodes have been demonstrated to have important properties in the field of spectral graph theory (Chung, 1997; Shi and Malik, 2000; Luxburg, 2007; Grady and Polimeni, 2010) by providing a probabilistic foundation (Meila and Shi, 2000; Robles-Kelly, 2005) for graph-based segmentation methods. In particular, the normalized cut problem (Shi and Malik, 2000) finds a segmentation x by minimizing $\frac{x^T L x}{x^T D x}$ (revealed by the Fiedler vector of the normalized Laplacian $D^{-\frac{1}{2}} L D^{-\frac{1}{2}}$ (Meila and Shi, 2000)). *Spectral Log-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. Here, the 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 top of Fig. 11 with overlaid image contours, was obtained with the nodal sets of the first and second eigenmodes. The positive values of the Fiedler vector gave warmer colors, negative values gave cooler colors (see also 1st and 2nd eigenmodes on Fig. 2). 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.




Fig. 10 Registration in 3D between two brains from healthy subjects using *Log-Demons* and *Spectral Log-Demons* (both with 4 levels). While observations show apparently similar results (ventricles are circled), the mean squared differences of intensities (MSE, less is better) between the fixed and registered images reveal an increase of 4% in accuracy and precision when using *Spectral Log-Demons*. The Dice metrics of the white and gray matter (measuring segmentation overlaps, more is better) also increase with *Spectral Log-Demons*.




Fig. 11 Implicit image segmentation by partitioning the images with the nodal sets of the first and second eigenmodes. Image contours are overlaid for visualization.

3.2 Groupwise Registration

In this second set of experiments, we evaluate the advantage of using spectral our new *Global Spectral Forces*, for constructing atlases. We verify that the groupwise registration produces a valid average shape, and that it is capable of handling highly complex deformations. Similarly to the previous experiments, we first choose the parameters by using the conventional approach, here, with the *GL-Demons*, and compare the registration accuracy when using *GSL-Demons* with the same set of parameters: $\sigma_{\text{fluid,diff}} = 1, \alpha_x = 1, k = 5, \alpha_g = 0.1, \alpha_s = 0.2, \alpha_i = 0.7$ in 2D and $\sigma_{\text{fluid,diff}} = 0.75, \alpha_x = 1, k = 5, \alpha_g = 0.25, \alpha_s = 0.35, \alpha_i = 0.4$ in 3D. Examples with real cardiac images are provided in order to demonstrate the advantage of using *Spectral Forces*, when organs exhibit a high shape variability.

3.2.1 Synthetic deformations

Convergence and capture of large deformations are now evaluated. $N/2$ velocity fields v are generated randomly using 15 control points with random locations in the image and random displacements of at most 15 pixels, or 20% of the image size, that are diffused over the image. Their forward and background transformations, $\exp(v)$ and $\exp(-v)$, are applied to an initial image I_0 , holding thus the average shape to I_0 . This establishes our ground truth. Again, since we compare the convergence and its rate, and not the final performance, the multilevel scheme is purposely not applied. Fig. 12 shows the groupwise registrations of 10 random hearts, 2D 75×75 images, through 100 trials. That is a total of 1000 hearts. The average Dice metric, which measures the overlap between all computed average shapes and I_0 , as well as the intensity errors (MSE) reveal that the reference shape, defined arbitrarily as one of the 10 images, evolves toward the ground truth. In other words, the Dice metric increases while the MSE decreases. Moreover, the N deformation fields become closer to the ground truth during registration. The striking difference in the convergence rates shows the full power of *GSL-Demons*, less than 5 iterations are required, while *GL-Demons* might not converge with such large deformations. We stopped the algorithms after 200 iterations. Time-wise, 35 iterations takes 194 seconds with *GSL-Demons*, and 53 seconds with *GL-Demons* using unoptimized Matlab code on a 2.53GHz Core 2 Duo. *GSL-Demons* shows a better performance with high deformations than *GL-Demons*.




Fig. 12 Groupwise registration of 10 images deformed randomly (100 trials, 1 sample on top row, with known ground truth) using *GL-Demons* and *GSL-Demons*, (Left) Best and worst atlases (based on Dice metric among 100 trials) demonstrating the capability of the *GSL-Demons* to handle large deformations, a) Average Dice metric with ground truth, b) Intensity difference between average shape and ground truth, c) transformation error with ground truth. *GSL-Demons* converges faster toward the average shape.

3.2.2 Cardiac Atlases

We now evaluate the construction of atlases with organs of high shape variability. *Ex vivo* hearts are particularly challenging to register as they present a high variability in fixture poses due to flabby ventricular walls. The human *ex vivo* DTMRI dataset (Rapacchi et al, 2010; Lombaert et al, 2011b, 2012c) provides good candidates to evaluate our algorithms. We use four hearts ($b = 0$ images of size 64^3) that were excluded in the construction of the human atlas (Lombaert et al, 2012c, 2011c) due to their hypertrophy and highly deformed shapes (see Fig. 13). *GL-Demons*, with 4 resolution levels, fail in recovering the shapes of the right ventricles, while *GSL-Demons* successfully constructs the atlas with only 1 level of resolution. Images were down-sampled at size 28^3 . As a comparison, 35 iterations takes 40 minutes in Matlab with *GSL-Demons* and 9 minutes with *GL-Demons*. Using *GSL-Demons* with 4 resolution levels reduce the intensity error (MSE) by half, from 10.8 to 5.08. Moreover, the Jacobian determinants of the transformation fields show that large deformations are successfully captured with the spectral-based update scheme and produces high and smooth Jacobian, as shown in Fig. 13 b. Local deformations are captured with the gradient-based update scheme in the higher levels of *GSL-Demons* as shown in Fig. 13 c.

A second cardiac atlas (Fig. 14) is constructed from a 3D+t MRI sequence, 11 frames of 64^3 with high systolic deformations on frame 6, as an example to show that *GL* and *GSL-Demons* can be successfully used on denser images and not only on segmented hearts.

4 Discussion and Conclusions

Finding an accurate dense pointwise correspondence between images becomes particularly challenging when a region of interest undergoes a large and complex deformation. The typical response found in the state-of-the-art consists of using a multilevel scheme in order to capture larger jumps in coarser resolutions. Although sufficient for most applications, a multilevel scheme does not change the inherent local scope of the updates derived from image gradients and quickly reaches its limit when image deformations become too strong. Other approaches utilize region-based descriptors (Liu et al, 2008, 2011; Brox et al, 2009; Brox and Malik, 2011) that are specialized for finding correspondences between images with fast motion. The resulting transformations are however not diffeomorphic, and rely on a hierarchical matching of local patches that are characterized with histograms of local orientations, like in SIFT (Lowe, 2004).

In this paper, we propose a fundamentally new update scheme based on a simple 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 nearest-neighbor searches in a multidimensional space that embeds pixel intensity, Euclidean coordinates and spectral coordinates. Closest points between such embeddings find the best compromise between these three previous properties. The full power of this simple direct feature matching technique is demonstrated when it serves as a new original update scheme in classical frameworks for image registration and atlas construction. Larger deforma-

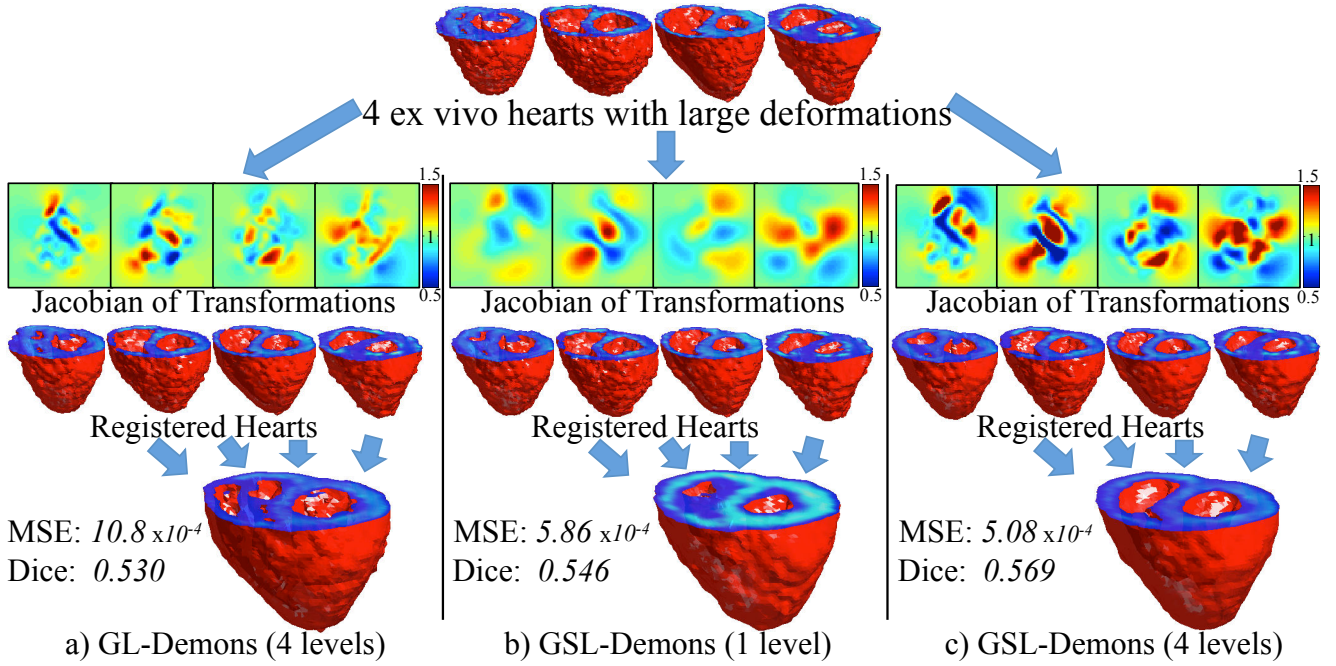


Fig. 13 Atlas of *ex vivo* hearts (isosurfaces are shown) using a) *GL-Demons* (4 levels, showing failure in the right ventricle), b) *GSL-Demons* (1 level), c) and *GSL-Demons* (4 levels, with correct right ventricle). *GSL-Demons* capture successfully large deformations. Jacobian determinants (axial planes) show that spectral matching capture smooth and large deformations while gradient-based updates capture local deformations.

tions can effectively be captured through spectral representations, which are in turn refined throughout the improved registration process. The slight differences in spectral representations, which are due to small perturbation in isometry, diminish when images are warped to one another during the registration process. Among the existing frameworks, the *Log-Demons* algorithm is simple, intuitive, and provides symmetry and diffeomorphism of the transformation. It can additionally be extended to perform groupwise registration with the *GL-Demons*, which computes the average shape during the registration process. The *Log-Demons* algorithm and its groupwise extension are effective benchmark frameworks for comparing image registration and atlas construction between conventional and spectral approaches.

Other registration algorithms could have been enhanced with *Spectral Forces*, however, our evaluation and experiments focused on illustrating the properties and advantages of a spectral approach. It is likely that *Spectral Forces* may similarly improve other registration algorithms that are based on gradient-based updates, particularly on images that exhibit very large and complex deformations. One promising candidate could be the LDDMM, which was also shown to have difficulty in recovering very large deformations. Its elegant modeling of large deformations via time-dependent velocity fields could indeed benefit from using *Spec-*

tral Forces, which could bring the algorithm optimization to farther away local minima that are unreachable with gradient-based, locally defined, update schemes. Other candidates also includes recent advances within the Demons framework, such as the Polyaffine Demons (Seiler et al, 2012), or the Local Correlation Coefficient Demons (Lorenzi et al, 2013).

Spectral Log-Demons may further benefit from other alternatives approaches since the general formulation in Eq. 2 could embed in fact any type of extra information that could potentially improve the matching accuracy. It may, for instance, incorporate local SIFT descriptors (Liu et al, 2008, 2011) as extra information in the multidimensional embeddings. However, rather than seeking for the best combination of extra features, or determining which registration method could best benefit from *Spectral Forces*, our experiments evaluated the improvements in registration accuracy that were due to the *Spectral Forces*. Their global scope effectively brings the optimization to a basin of attraction that is closer to an optimal registration. This also motivates the use of a hybrid multilevel scheme where global shape characteristics are present on coarser levels. Once large deformations are captured on a coarser level, finer details are efficiently registered using a conventional, faster, gradient-based update scheme.

The new pairwise registration method, called *Spectral Log-Demons*, successfully captured large and com-




Fig. 14 Atlas (average shape) from 11 frames of a 3D+t cardiac MRI sequence. (Left) Original 3D frames before and after groupwise registration (contraction/expansion is corrected). (Right) Computed atlas with the average shape (only the *GSL-Demons* version is shown). The intensity error (between all registered frames and the atlas) is reduced by 30% with *GSL-Demons*.

plex deformations. The spectral approach actually demonstrated a 73% improvement in accuracy over the conventional *Log-Demons* algorithm when registering images with displacements of more than 30% of the image size. Furthermore, the new groupwise framework, called *GSL-Demons*, has demonstrated a drastic increase in the rate of convergence of the shape averaging process, typically, less than 5 iterations is necessary. A better robustness is also observed for large image deformations. In fact, while the conventional approach gets easily trapped in local minima, the spectral approach is able to successfully generate average shapes from 1000 hearts with deformations as large as 20% of the image size.

Our spectral approaches illustrate the feasibility of registering images and constructing atlases of organs having a large shape variability. As real applications, we constructed atlases with medical images exhibiting extreme deformations, in particular, with hearts having significant changes in shape configurations. The use of *Spectral Forces*, actually enabled the construction of atlases that were not possible with conventional approaches. However, there are important limitations to mention. From a practical point of view, we assumed that registered images must have a similar topology with no occlusions, holes, or missing parts. Images must have the same global shape. 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, but this can be a limiting factor in specific applications, e.g., organs with missing structures between images. This may for instance improve automatic organ localization by enhancing adapted registration algorithms with *Spectral Forces* (Ranjan, 2011). From an engineering point of view, there is still room for improvement as our main limitation currently resides in the computational cost of the spectral decomposition,

which is applied on a $N \times N$ matrix, where N is the number of pixels.

Future work will aim at improving this computational burden, for instance by precomputing the eigenmodes before registration and updating them via differential calculus (Magnus, 1985) in order to avoid the costly spectral decompositions during iteration. We will also investigate the use of the Nyström approximation (Drineas and Mahoney, 2005) by sampling the Laplacian matrix. Moreover, the groupwise frameworks could further gain performance by benefiting from parallel computing, e.g., GPU. Besides practical improvements, the strong links between our approach and image segmentation (Meila and Shi, 2000; Robles-Kelly, 2005; Weiss, 1999) will also be studied in more depth, since this may lead to a global optimization framework where image registration, segmentation and shape averaging may be computed within the same joint 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 and atlas construction. More precisely, our correspondence map respects the *global* Riemannian structure of the image on top of the differentiable structure. *Isomorphism* is brought by *spectral coordinates* while *diffeomorphism* comes from state-of-the-art methods. Our new improved spectral approach can improve applications that employ objects or study organs undergoing very large and complex deformations.

Acknowledgments

The authors wish to thank Pierre Croisille for the *ex vivo* cardiac images as well as Marco Lorenzi and Christof Seiler for constructive comments. We are also grateful to the anonymous reviewers for their suggestions.

The project was supported financially by the National Science and Engineering Research Council of Canada (NSERC) and the Fonds de Recherche du Québec (FQRNT).

References

- Allasonnière S, Amit Y, Trouvé A (2007) Towards a coherent statistical framework for dense deformable template estimation. *Royal Statistical Society* 69(1):3–29 [1](#)
- Avants B, Gee JC (2004) Geodesic estimation for large deformation anatomical shape averaging and interpolation. *NeuroImage* 23:139–150 [1](#), [8](#)
- Beg MF, Khan A (2006) Computing an average anatomical atlas using LDDMM and geodesic shooting. In: *ISBI*, pp 1116–1119 [1](#)
- Beg MF, Miller MI, Trouvé A, Younes L (2005) Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International Journal of Computer Vision* 61(2):139–157 [1](#), [10](#)
- Bhatia KK, Hajnal JV, Puri BK, Edwards AD, Rueckert D (2004) Consistent groupwise non-rigid registration for atlas construction. In: *ISBI*, pp 908–911 [1](#)
- Bossa M, Hernandez M, Olmos S (2007) Contributions to 3D diffeomorphic atlas estimation: application to brain images. In: *MICCAI*, pp 667–674 [1](#), [8](#)
- Brox T, Malik J (2011) Large displacement optical flow: Descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(3):500–513 [2](#), [15](#)
- Brox T, Bregler C, Malik J (2009) Large displacement optical flow. In: *CVPR*, pp 41–48 [2](#), [15](#)
- Cachier P, Bardin E, Dormont D, Pennec X, Ayache N (2003) Iconic feature based nonrigid registration: the PASHA algorithm. *CVIU* 89(2-3):272–298 [7](#)
- Carcassoni M, Hancock E (2003) Spectral correspondence for point pattern matching. *Pattern Recognition* 36(1):193–204 [2](#)
- Carcassoni M, Hancock ER (2000) Point pattern matching with robust spectral correspondence. In: *CVPR*, pp 649–655 [2](#)
- Chui H, Rangarajan A (2000) A new algorithm for non-rigid point matching. In: *CVPR*, pp 44–51 [1](#)
- Chung FRK (1997) *Spectral Graph Theory*. AMS [2](#), [3](#), [4](#), [13](#)
- Crum WR, Hartkens T, Hill DL (2004) Non-rigid image registration: theory and practice. *British Journal of Radiology* 77(2):140–153 [1](#)
- Ding C, He X (2004) K-means clustering via principal component analysis. In: *ICML* [4](#)
- Drineas P, Mahoney MW (2005) On the Nyström method for approximating a Gram matrix for improved Kernel-Based learning. *Journal Machine Learning Research* 6:2153–2175 [17](#)
- Durrleman S, Fillard P, Pennec X, Trouvé A, Ayache N (2011) Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage* 55(3):1073–1090 [1](#)
- Egozi A, Keller Y, Guterman H (2010) Improving shape retrieval by spectral matching and meta similarity. *IEEE Transactions on Image Processing* 19(5):1319–1327 [2](#)
- Fischl B, Sereno M, Tootell R, Dale A (1999) High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* 8(4):272–284 [13](#)
- Glocker B, Sotiras A, Komodakis N, Paragios N (2011) Deformable medical image registration: Setting the state of the art with discrete methods. *Annual Review of Biomedical Engineering* 13(1):219–244 [2](#)
- Grady L, Polimeni JR (2010) *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer [2](#), [3](#), [4](#), [13](#)
- Guimond A, Meunier J, Thirion JP (2000) Average brain models: a convergence study. *Computer Vision and Image Understanding* pp 192–210 [2](#), [8](#)
- Jain V, Zhang H (2006) Robust 3D shape correspondence in the spectral domain. In: *Int. Conf. on Shape Modeling and App.* [2](#), [3](#)
- Joshi S, Davis B, Jomier M, Gerig G (2004) Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23:151–160 [1](#)
- van Kaick O, Zhang H, Hamarneh G, Cohen-Or D (2011) A survey on shape correspondence. *Eurographics* 30(6):1681–1707 [2](#)
- Konukoglu E, Glocker B, Criminisi A, Pohl KM (2012) WESD - weighted spectral distance for measuring shape dissimilarity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(9):2284–2297 [2](#)
- Liu C, Yuen J, Torralba A, Sivic J, Freeman WT (2008) SIFT flow: Dense correspondence across different scenes. In: *ECCV*, pp 28–42 [2](#), [15](#), [16](#)
- Liu C, Yuen J, Torralba A (2011) SIFT flow: Dense correspondence across scenes and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(5):978–994 [2](#), [15](#), [16](#)
- Lombaert H, Grady L, Polimeni JR, Cheriet F (2011a) Fast brain matching with spectral correspondence. In: *IPMI*, pp 660–670 [2](#), [3](#), [5](#), [13](#)
- Lombaert H, Peyrat JM, Croisille P, Rapacchi S, Fanton L, Clarysse P, Delingette H, Ayache N (2011b) Statistical analysis of the human cardiac fiber architecture from DT-MRI. In: *FIMH*, vol 6666, pp 171–179 [2](#), [15](#)

- Lombaert H, Peyrat JM, Fanton L, Cheriet F, Delingette H, Ayache N, Clarysse P, Magnin I, Croisille P (2011c) Statistical atlas of human cardiac fibers: Comparison with abnormal hearts. In: MICCAI STACOM, pp 207–213 [15](#)
- Lombaert H, Grady L, Pennec X, Ayache N, Cheriet F (2012a) Spectral demons - image registration via global spectral correspondence. In: ECCV [2](#), [7](#)
- Lombaert H, Grady L, Pennec X, Peyrat JM, Ayache N, Cheriet F (2012b) Groupwise spectral Log-Demons framework for atlas construction. In: MICCAI MCV [2](#), [8](#)
- Lombaert H, Peyrat JM, Croisille P, Rapacchi S, Fanton L, Cheriet F, Clarysse P, Magnin I, Delingette H, Ayache N (2012c) Human atlas of the cardiac fiber architecture: Study on a healthy population. *IEEE Transactions on Medical Imaging* 31:1436–1447 [2](#), [15](#)
- Lombaert H, Grady L, Polimeni JR, Cheriet F (2013a) FOCUSR: Feature oriented correspondence using spectral regularization - a method for accurate surface matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35:2143–2160 [2](#)
- Lombaert H, Sporring J, Siddiqi K (2013b) Diffeomorphic spectral matching of cortical surfaces. In: IPMI, pp 376–389 [2](#)
- Lorenzi M, Ayache N, Frisoni GB, Pennec X, Alzheimer’s Disease Neuroimaging Initiative (ADNI) (2013) LCC-demons: a robust and accurate symmetric diffeomorphic registration algorithm. *NeuroImage* 81:470–483 [16](#)
- Lowe D (2004) Distinctive image features from Scale-Invariant keypoints. *International Journal of Computer Vision* 60(2):91–110 [15](#)
- Luxburg U (2007) A tutorial on spectral clustering. *Statistics and Computing* 17(4):395–416 [13](#)
- Magnus JR (1985) On differentiating eigenvalues and eigenvectors. *Econometric Theory* 1(2):179–191 [17](#)
- Mansi T, Pennec X, Sermesant M, Delingette H, Ayache N (2011) iLogDemos: A Demons-Based registration algorithm for tracking incompressible elastic biological tissues. *International Journal of Computer Vision* 92(1):92–111 [7](#), [12](#)
- Marsland S, Twining CJ, Taylor CJ (2003) Groupwise non-rigid registration using polyharmonic Clamped-Plate splines. In: MICCAI, pp 771–779 [1](#)
- Mateus D, Horaud R, Knossow D, Cuzzolin F, Boyer E (2008) Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In: CVPR, pp 1–8 [2](#), [3](#), [5](#)
- Meila M, Shi J (2000) Learning segmentation by random walks. In: NIPS [2](#), [4](#), [13](#), [17](#)
- Miller MI, Trouné A, Younes L (2002) On the metrics and Euler-Lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering* 4(1):375–405 [1](#)
- Pennec X, Cachier P, Ayache N (1999) Understanding the demon’s algorithm: 3D non-rigid registration by gradient descent. In: MICCAI, pp 597–605 [7](#)
- Peyrat JM, Sermesant M, Pennec X, Delingette H, Xu C, McVeigh ER, Ayache N (2007) A computational framework for the statistical analysis of cardiac diffusion tensors: application to a small database of canine hearts. *IEEE Transactions on Medical Imaging* (11):1500–1514 [2](#)
- Ranjan SR (2011) Organ localization through anatomy-aware non-rigid registration with atlas. In: AIPR, pp 1–5 [17](#)
- Rapacchi S, Croisille P, Pai V, Grenier D, Viallon M, Kellman P, Newton N, Wen H (2010) Reducing motion sensitivity in free breathing DWI of the heart with localized Principal Component Analysis. In: ISMRM [15](#)
- Reuter M (2009) Hierarchical shape segmentation and registration via topological features of Laplace-Beltrami eigenfunctions. *International Journal of Computer Vision* 89(2):287–308 [2](#), [13](#)
- Risser L, Vialard F, Wolz R, Murgasova M, Holm DD, Rueckert D (2011) Simultaneous multi-scale registration using large deformation diffeomorphic metric mapping. *IEEE Transactions on Medical Imaging* 30(10):1746–1759 [10](#)
- Robles-Kelly A (2005) Segmentation via Graph-Spectral methods and Riemannian geometry. In: CAIP, pp 661–668 [13](#), [17](#)
- Rueckert D, Sonoda LI, Hayes C, Hill DL, Leach MO, Hawkes DJ (1999) Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging* 18(8):712–721 [1](#)
- Scott GL, Longuet-Higgins HC (1991) An algorithm for associating the features of two images. *Royal Society Biological Sciences* 244(1309):21–26 [2](#), [3](#)
- Seiler C, Pennec X, Reyes M (2012) Capturing the multiscale anatomical shape variability with polyaffine transformation trees. *Medical Image Analysis* 16(7):1371–1384 [16](#)
- Shapiro LS, Brady JM (1992) Feature-based correspondence: an eigenvector approach. *Image and Vision Computing* 10(5):283–288 [2](#), [3](#)
- Shekhovtsov A, Kovtun I, Hlavac V (2007) Efficient MRF deformation model for Non-Rigid image matching. In: CVPR, pp 1–6 [2](#)
- Shi J, Malik J (2000) Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8):888–905 [2](#), [4](#), [8](#), [13](#)

- Studholme C, Cardenas V (2004) A template free approach to volumetric spatial normalization of brain anatomy. *Pattern Recognition Letters* 25:1191–1202 [1](#)
- Thursty T (2010) 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 [5](#)
- Umeyama S (1988) An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10(5):695–703 [2, 3](#)
- Vercauteren T, Pennec X, Perchant A, Ayache N (2007) Non-parametric diffeomorphic image registration with the demons algorithm. In: *MICCAI*, pp 319–326 [2, 6, 7](#)
- Vercauteren T, Pennec X, Perchant A, Ayache N (2008) Symmetric Log-Domain diffeomorphic registration: A Demons-Based approach. In: *MICCAI*, pp 754–761 [2, 7, 8](#)
- Vercauteren T, Pennec X, Perchant A, Ayache N (2009) Diffeomorphic demons: efficient non-parametric image registration. *NeuroImage* 45:61–72 [1, 7](#)
- Weiss Y (1999) Segmentation using eigenvectors: a unifying view. In: *ICCV*, pp 975–982 [2, 17](#)
- Wilson RC, Hancock ER, Pekalska E, Duin RPW (2010) Spherical embeddings for non-Euclidean dissimilarities. In: *CVPR*, pp 1903–1910 [2](#)
- Wu G, Jia H, Wang Q, Shen D (2011) SharpMean: groupwise registration guided by sharp mean image and tree-based registration. *NeuroImage* 56(4):1968–1981 [2](#)
- Yeo BT, Sabuncu MR, Vercauteren T, Ayache N, Fischl B, Golland P (2010) Spherical demons: Fast diffeomorphic Landmark-Free surface registration. *IEEE Transactions on Medical Imaging* 29(3) [13](#)
- Zhang H, Van Kaick O, Dyer R (2010) Spectral mesh processing. *Eurographics* 29(6) [2](#)
- Zikic D, Glocker B, Kutter O, Groher M, Komodakis N, Kamen A, Paragios N, Navab N (2010) Linear intensity-based image registration by markov random fields and discrete optimization. *Medical Image Analysis* 14(4):550–562 [2](#)
- Zikic D, Baust M, Kamen A, Navab N (2011) A general preconditioning scheme for difference measures in deformable registration. In: *ICCV*, pp 49–56 [2](#)
- Zollei L, Learned-Miller E, Grimson E, Wells W (2005) Efficient population registration of 3D data. In: *ICCV CVBIA* [1](#)