

Active-Contour-Based Image Segmentation using Machine Learning Techniques

Patrick Etyngier, Florent Ségonne, and Renaud Keriven

Odyssée Team / Certis - Ecole des Ponts - France - etyngier@certis.enpc.fr

Abstract. We introduce a non-linear shape prior for the deformable model framework that we learn from a set of shape samples using recent manifold learning techniques. We model a category of shapes as a finite dimensional manifold which we approximate using Diffusion maps. Our method computes a Delaunay triangulation of the reduced space, considered as Euclidean, and uses the resulting space partition to identify the closest neighbors of any given shape based on its Nyström extension. We derive a non-linear shape prior term designed to attract a shape towards the shape prior manifold at given constant embedding. Results on shapes of ventricle nuclei demonstrate the potential of our method for segmentation tasks.



Fig. 1. Aligned shape samples of the right ventricle nucleus from different subjects corresponding, from left to right, to 2 young, 2 mid-age, and 2 old subjects; the last shape sample originates from a subject with Alzheimer’s Disease. While shapes appear quite similar, they usually cannot be considered as small deformations around a mean shape.

1 Introduction

1.1 Motivation

Accurate segmentation of anatomical structures from medical images is a fundamental but difficult task that often dictates the outcome of the entire clinical or research analysis (e.g. visualization, neuro-surgical planning, surface-based processing of functional data, inter-subject registration, among others). The challenge is that images are usually corrupted by several artifacts, such as image noise, missing or occluded parts, image intensity inhomogeneity or non-uniformity, and partial volume averaging effect. When dealing with complex images, some prior shape knowledge may be necessary to disambiguate the segmentation process.

In medical imaging, excluding uncommon pathological cases, the overall shape of most macroscopic anatomical structures is prescribed by medical knowledge: it is usually known a priori, does not vary much between individuals, and its observed geometric variability seems to be governed by a small set of unknown parameters. In this paper, we assume that the set of shapes of a specific anatomical structure (e.g. left or right ventricle, hippocampus nucleus) evolves in a low-dimensional, but not necessarily linear, space that we term the *shape prior manifold*. This is clearly exemplified in Fig. 1 which displays some shapes corresponding to the right ventricle nucleus for different subjects at different time points in their life. Even shapes extracted from subjects with Alzheimer’s Disease, a relatively slowly evolving pathology, exhibits some shape similarity with normal controls.

Knowledge of the underlying structure encoding the geometric variability of anatomy provides shape constraints for image segmentation. We propose to discover the structure of the *shape prior manifold* using recent manifold learning techniques, and to exploit it to carefully design a non-linear shape prior integrated into the deformable model framework for the purpose of image segmentation.

1.2 Previous Work

The use of shape prior information in the deformable model framework has long been limited to a smoothness assumption or to simple parametric families of shapes. But a recent and important trend in this domain is the development of deformable models integrating more elaborate shape information.

An important work in this direction is the *active shape model* of Cootes et al. [1]. A principal component analysis (PCA) on the position of some landmark points placed in a coherent way on all the training contours is used to reduce the number of degrees of freedom to the principal modes of variation. Although successfully applied to various types of shapes (hands, faces, organs), the reliance on a parameterized representation and the manual positioning of the landmarks, particularly tedious in 3D images, seriously limits its applicability.

Leventon, Grimson and Faugeras [2] circumvent these limitations by computing parameterization-independent shape statistics within the level set representation [3]. Basically, they perform a PCA on the signed distance functions of the training shapes, and the resulting statistical model is integrated into a geodesic active contour framework. The evolution equation contains a term which attracts the model towards an optimal prior shape as a combination of the mean shape and of the principal modes of variation. Several improvements to this approach have been proposed [4, 5], and in particular an elegant integration of the statistical shape model into a unique MAP Bayesian optimization. Let us also mention another neat Bayesian prior shape formulation, based on a B-spline representation, proposed by Cremers, Kohlberger and Schnörr in [6].

Performing PCA on distance functions might be problematic since they do not define a vector space. To cope with this, Charpiat, Faugeras and Keriven [7] proposed shape statistics based on differentiable approximations of the Hausdorff distance. However, their work is limited to a linearized shape space with small deformation modes around a mean shape. Such an approach is relevant only when the learning set is composed of very similar shapes. Finally, let's also mention the elegant M-reps approach [8], that is restricted to a specific type of deformable models.

1.3 Novelty of our Approach

In this paper, we depart from the small deformation assumption and introduce a new deformable model framework that integrates more general non-linear shape priors. We model a category of shapes as a smooth finite-dimensional sub-manifold of the infinite-dimensional shape space, termed the *shape prior manifold*. This manifold which cannot be represented explicitly is approximated from a collection of shape samples using a recent manifold learning technique called Diffusion maps [9, 10]. Manifold learning, which is already an established tool in object recognition and image classification, has been recently applied to shape analysis [11]. Yet, to our knowledge, such techniques have not been used in the context of image segmentation with shape priors.

Diffusion maps generate a mapping, called an embedding, from the original shape space into a low-dimensional space, which can *advantageously* be considered as Euclidean [9]. We design a shape prior term based on the Nyström extension [12] which provides a sound and efficient framework for extending embedding coordinates to the full infinite dimensional shape space. Motivated by its Euclidean nature, a Delaunay partitioning of the reduced space is used to identify the closest neighbors (in the training set) of any shape in the original infinite dimensional shape space. The neighboring shapes are then integrated into a variational functional designed to attract any given shape towards the *shape prior manifold*.

The remainder of this paper is organized as follows. Section 2 introduces the necessary background in manifold learning: it is dedicated to learning the *shape prior manifold* from a finite set of shape samples using Diffusion maps. Section 3 presents our deformable model framework using non-linear shape priors. Section 4 reports some preliminary numerical experiments which yield promising results with real shapes.

2 Learning the Shape Prior Manifold

In the sequel, we define a *shape* as a simple compact (i.e. closed and non-intersecting) surface, and we denote by \mathbb{S} the (infinite-dimensional) space of such shapes¹. We make the assumption that a category of shapes, i.e. the set of shapes that can be identified with a common anatomical structure, e.g. left or right ventricle, hippocampus nucleus, can be modeled as a finite-dimensional manifold, termed the *shape prior manifold*.

Dimensionality reduction, i.e. the process of recovering the underlying low dimensional structure of a manifold \mathcal{M} embedded into a higher-dimensional space, has enjoyed renewed interest over the past years. Among the most recent and popular techniques are the Locally Linear Embedding (LLE) [13], Laplacian eigenmaps [14], Diffusion maps [9, 15].

In this work, we learn the *shape prior manifold* using Diffusion maps. For the sake of clarity, we present the mathematical formulation for data living in \mathbb{R}^n . An extension to infinite-dimensional shape manifolds is straightforward (We refer the reader to [9, 10, 15] for more detail).

2.1 Manifold Learning and Diffusion Maps

Let \mathcal{M} be a manifold of dimension m lying in \mathbb{R}^n ($m \ll n$). Diffusion maps rely on discrete approximations of the Laplace-Beltrami operator $\Delta_{\mathcal{M}}$ defined on the manifold \mathcal{M} to generate a mapping (called an embedding) $f : \mathcal{M} \rightarrow \mathbb{R}^m$ such that if two points x and z are *close* in \mathcal{M} , so are $f(x)$ and $f(z)$. The optimal mapping is given by the eigen-functions of the Laplace-Beltrami operator corresponding to the m smallest non-zero eigenvalues, where m is the target dimension. Note that the latter dimension can either be known *a priori* or be inferred from the profile of the eigen spectrum [9]. In practice, a discrete counterpart to this continuous formulation must be used since we only have access to a discrete and finite set, denoted Γ , of example shapes in this category.

2.1.a Distance in the Shape Space

The approximation of the Laplace-Beltrami operator requires the choice of a distance

¹ Note that, although this paper only deals with 2-dimensional surfaces embedded in the 3-dimensional Euclidean space, all ideas and results seamlessly extend to higher dimensions.

between shapes. Many different definitions of the distance between two shapes have been proposed in the computer vision literature but there is no agreement on the correct way of measuring shape similarity. The definition used in the experiments presented in Sect. 4 are based on the representation of a surface S in the Euclidean embedding space \mathbb{R}^3 by its signed distance function. In this context, we define the distance between two shapes to be the Sobolev $W^{1,2}$ -norm of the difference between their signed distance functions [7]:

$$d_{W^{1,2}}(S_1, S_2)^2 = \|\bar{D}_{S_1} - \bar{D}_{S_2}\|_{L^2(\Omega, \mathbb{R})}^2 + \|\nabla \bar{D}_{S_1} - \nabla \bar{D}_{S_2}\|_{L^2(\Omega, \mathbb{R}^n)}^2,$$

where \bar{D}_{S_i} denotes the signed distance function of shape S_i ($i = 1, 2$), and $\nabla \bar{D}_{S_i}$ its gradient. Note that to define a distance between shapes that is invariant to rigid displacements (rotations and translations), we first align the shapes using their principal moments before computing distances. Note also that the proposed method is obviously not limited to a specific choice of distance [7].

2.1.b Approximation to the Laplace-Beltrami Operator

Once a distance has been chosen, classical manifold learning techniques can be applied by building an adjacency graph of the learning set of shape examples. Let $\Gamma = \{x_1 \cdots x_p \in \mathbb{R}^n\}$ be p sample points of the m dimensional manifold \mathcal{M} sampled under an unknown density $q_{\mathcal{M}}$. An adjacency matrix $(W_{i,j})_{i,j \in 1, \dots, p}$ is then constructed, the coefficients of which measure the strength of the different edges in the adjacency graph. Typically, $W_{i,j}$, also denoted $w(x_i, x_j)$, is a decreasing function of the distance between x_i and x_j . In this work, $w(x_i, x_j) = \exp(-d^2(x_i, x_j)/2\sigma^2)$, with σ estimated as the median of all the distances between all shapes.

Classical manifold learning methods provide an embedding that combines the information of both the density $q_{\mathcal{M}}$ and the geometry [10, 15]. In order to construct an approximation of the Laplace-Beltrami operator that is independent of the unknown density $q_{\mathcal{M}}$, we renormalize the adjacency matrix $(W_{i,j})$. Briefly, we form the new adjacency matrix $(\tilde{W}_{i,j})$ by $\tilde{w}(x_i, x_j) = \frac{w(x_i, x_j)}{q(x_i)q(x_j)}$, with $q(x) = \sum_{y \in \Gamma} w(x, y)$. We then define the anisotropic transition kernel $(P_{i,j})_{i,j \in 1, \dots, p}$ such that $p(x_i, x_j) = \frac{\tilde{w}(x_i, x_j)}{\tilde{q}(x_i)\tilde{q}(x_j)}$ with $\tilde{q}(x) = \sum_{y \in \Gamma} \tilde{w}(x, y)$. The kernel $(P_{i,j})$ is a density-independent approximation of the operator $\mathbb{1} - \Delta_{\mathcal{M}}$ [9]. From the definition of the adjacency matrix, we find that:

$$p(x_i, x_j) = \frac{w(x_i, x_j)}{\sum_b K_{jb} w(x_i, x_b)} \text{ with } K_{jb} = \frac{q(x_j)}{q(x_b)} = \frac{\sum_{y \in \Gamma} w(x_j, y)}{\sum_{y \in \Gamma} w(x_b, y)}. \quad (1)$$

2.1.c Generating the Embedding using Diffusion Maps

Let's denote $\{\lambda_i\}_{i,j \in 1, \dots, p}$ (with $|\lambda_0| \geq |\lambda_1| \geq \dots$) and Ψ_i the associated eigenvalues and eigenvectors of $(P_{i,j})$. Coifman and coworkers have shown in [9] that the eigenvectors of $(P_{i,j})$ converge to those of the Laplace-Beltrami operator on \mathcal{M} and that a mapping Φ_t that embeds the data into the *Euclidean* space \mathbb{R}^m *isometrically* with respect to a Diffusion distance in the original shape space \mathbb{S} can be constructed:

$$\Phi_t : \Gamma \subset \mathcal{M} \rightarrow \mathbb{R}^m, x_i \mapsto (\lambda_1^t \Psi_1(x_i), \dots, \lambda_m^t \Psi_m(x_i)). \quad (2)$$

Diffusion distance reflects the intrinsic geometry of the data set defined via the adjacency graph in a diffusion process. It was shown to be more robust to outliers than geodesic distances [9], thereby motivating its use to estimate the embedding (Fig. 2-c). In this formulation, t is a time parameter controlling the diffusivity of the adjacency graph and can be chosen arbitrarily. We used $t = 1$ for our experiments (Sect. 4).

2.2 Extending the Embedding based on Nyström Extension

The mapping Φ_t is only defined on the training samples. The Nyström extension method is a popular technique employed for the extension of empirical functions from the training set Γ to new samples. Noticing that every training sample verifies:

$$\forall x \in \Gamma \quad \forall k \in 1, \dots, p \quad \sum_{y \in \Gamma} p(x, y) \Psi_k(y) = \lambda_k \Psi_k(x),$$

the embedding of new data points located outside the set Γ can similarly be computed by extension (Lafon and coworkers define another elegant extension in [15]):

$$\tilde{\Phi}_t : \mathbb{R}^n \rightarrow \mathbb{R}^m, \quad x \mapsto \left(\lambda_1^{t-1} \sum_{y \in \Gamma} p(x, y) \Psi_1(y), \dots, \lambda_m^{t-1} \sum_{y \in \Gamma} p(x, y) \Psi_m(y) \right) \quad (3)$$

3 Image Segmentation using the Shape Prior Manifold

In this section, we propose to use the embedding to carefully design a shape prior term integrated into a deformable model framework for the purpose of image segmentation.

3.1 Image Segmentation as a Variational Problem

Without loss of generality, we cast the segmentation problem as a variational one, where the objective is to find a surface S minimizing a global energy functional E^{ac} . Depending on the segmentation task and the available information, the energy functional E^{ac} can take on different, more or less complex, forms, but, generally, E^{ac} can be written as a combination of image terms, designed to drive the surface towards the searched contour, and regularization terms, enforcing smoothness constraints. Directly finding the global minimum of E^{ac} is usually impossible and one often has to resort to a sub-optimal gradient-descent strategy starting from a guess S_0 . That is we assume that the image segmentation problem amounts to solving the following evolution problem: find the active contour $S : \tau \in \mathbb{R}^+ \mapsto S(\tau) \in \mathbb{S}$ such that $S(0) = S_0, \frac{dS}{d\tau} = -\nabla E^{ac}$.

3.2 Designing and Integrating the Shape Prior Term

We define a shape prior functional E^{sp} designed to attract any given shape S towards the shape prior manifold. Unfortunately, Diffusion maps do not give access to an explicit projection operator onto the reduced manifold. To alleviate this problem, we exploit the Euclidean nature of the reduced space by computing a Delaunay triangulation in \mathbb{R}^m of the training data. The space partition is then used to identify the $m+1$ closest neighbors (in the training set Γ) of the shape S in the \mathbb{S} by computing its embedding coordinates $\tilde{\Phi}(S)$ and finding the corresponding Delaunay triangle formed by $m+1$ -vertices in \mathbb{R}^m . By doing so, we identify the $m+1$ closest neighbors $\mathcal{N} = (S_0, \dots, S_m)$ of S in \mathbb{S} for the Diffusion metric [9]. This neighborhood \mathcal{N} will then be used to attract S towards the manifold.

To this end, we compute the barycentric coordinates $\Theta = (\theta_0, \dots, \theta_m)$ of the shape S in the reduced space \mathbb{R}^m and define the shape prior functional in \mathbb{S} :

$$E_{\mathcal{N}, \Theta}^{sp}(S) = \sum_{i=0}^m \theta_i d^2(S_i, S) \quad \text{with} \quad \tilde{\Phi}(S) = \sum \theta_i \tilde{\Phi}(S_i), \quad \theta_i \geq 0, \quad \sum \theta_i = 1$$

designed to attract the shape S towards a *weighted mean shape* that interpolates between the $m+1$ samples $S_i \in \mathcal{N}$.

Minimization of the energy $E_{\mathcal{N},\theta}^{sp}(S)$ by gradient descent might change the embedding coordinates $\tilde{\Phi}_t(S)$ of the evolving shape S . Therefore, denoting by $\mathbb{S}_x = \tilde{\Phi}_t^{-1}(x)$ the x -level set in \mathbb{S} of the embedding $\tilde{\Phi}_t$ (note that \mathbb{S}_x has codimension m), we define the shape prior term \vec{v}_{sp} as the projection of the velocity field $\vec{v} = -\nabla E_{\mathcal{N},\theta}^{sp}$ onto the tangent space $\mathbb{T}_{\tilde{\Phi}_t(S)}$ of $\mathbb{S}_{\tilde{\Phi}_t(S)}$ at S . Using Eq. 1 and Eq. 3, $\mathbb{T}_{\tilde{\Phi}_t(S)}$ can be expressed by m simple orthogonality conditions in the tangent space $\mathbb{T}_{\mathbb{S}}(S)$ of \mathbb{S} at S :

$$\mathbb{T}_{\tilde{\Phi}_t(S)} = \left\{ \vec{v} \in \mathbb{T}_{\mathbb{S}}(S) \text{ s.t. } \forall k = 1, \dots, m \sum_{S_j \in \Gamma} \langle \nabla_{Sp}(S, S_j) | \vec{v} \rangle_{\mathbb{L}^2} \Psi_k(S_j) = 0 \right\},$$

where $\langle \cdot | \cdot \rangle_{\mathbb{L}^2}$ corresponds to the \mathbb{L}^2 -dot product in the tangent shape space $\mathbb{T}_{\mathbb{S}}(S)$. Projection of the velocity field $-\nabla E_{\mathcal{N},\theta}^{sp}$ onto $\mathbb{T}_{\tilde{\Phi}_t(S)}$ can then be achieved using the orthogonalization Gram-Schmidt process.

Finally, the general deformable model framework corresponds to solving the following evolution problem:

$$S(0) = S_0, \quad \frac{dS}{d\tau} = -\nabla E^{ac} + \alpha \vec{v}_{sp},$$

where α is a weighting parameter. Note that at each step of the evolution, we have to align the shape with the training samples using the principal moments before computing its embedding and deriving the shape prior term \vec{v}_{sp} .

4 Results and Discussion

We illustrate the potential benefits of our approach on a simple segmentation task, the segmentation of the ventricle nucleus from Magnetic Resonance Image (MRI). Training shape samples were obtained from 39 manually segmented images of 10 young, 10 mid-age, and 9 old normal controls and of 11 demented adults (Fig. 1). 39 data points form an insufficiently small data set and more shape samples are desirable to recover a satisfactory embedding. Note also that the artificial nature of the proposed segmentation task is only dedicated to reveal the influence of the shape prior term.

4.1 Estimating the dimension of the shape prior manifold: the dimension m is usually estimated from the profile of the eigenspectrum (Fig. 2-a). Yet, there is not always an obvious choice (especially when the number of data points is insufficient). In




Fig. 2. a) Eigenspectrum profile and degree of separability: on this restricted data set with 39 shapes only, $m = 2$ appears to be the optimal dimension. b) The two-dimensional embedding partitioned by a Delaunay triangulation. c) A manually corrupted shape and its two closest neighbors in \mathbb{S} and in the reduced space: visually, the ones in the reduced space appear more similar.

our case, $m = 2$, $m = 3$, or $m = 4$ appear to be a realistic guess. However, in the case of labeled data, one can disambiguate this choice by also requiring the embedding $\tilde{\Phi}_t$ to separate/cluster “well” the different groups. We simply define the degree of separability $d_{i,j}$ between two groups i and j by the distance $d_{i,j} = \frac{\|\mu_i - \mu_j\|}{\sqrt{\sigma_i^2 + \sigma_j^2}}$, where μ_i and σ_i^2 are the mean and variance in \mathbb{R}^m of data points corresponding to group $\#i$. The degree of separability of the mapping $\tilde{\Phi}_t$ is then $\sum_{i,j} d_{i,j}$. Note that this method can also be used to determine an optimal value for the parameter t . Finally, on this unsatisfactory small data set, we find that the optimal mapping requires $m = 2$ (Fig. 2-a,b).

4.2 Closest neighbors: Diffusion maps embed *advantageously* the data set in the Euclidean space \mathbb{R}^m *isometrically* with respect to a Diffusion distance in \mathbb{S} . This distance was shown to be more robust to outliers than geodesic distances [9]. To illustrate this point, we show in Fig. 2-c a manually corrupted shape with its two closest neighbors in \mathbb{S} and \mathbb{R}^m . At least visually, the identified shapes in \mathbb{R}^m appear more similar to the corrupted shape than the ones in \mathbb{S} .

4.3 Ventricle nucleus segmentation from MRI with occlusion: we consider a simple segmentation task which consists of segmenting the ventricle nucleus from an MRI that was corrupted by white noise and degraded with an artificial occlusion (clearly visible in Fig. 3-a). Motivated by our choice of representing a shape S by its signed distance function \bar{D}_S , our surface deformation is implemented in the level set framework. The level set evolution is guided by a *simple* intensity-based velocity term, a curvature term, and the non-linear shape prior term:

$$\partial_\tau \bar{D}_S(x, \tau) = [\beta(I(x) - T(x)) - \kappa] |\nabla_x \bar{D}_S(x, \tau)| - \alpha \vec{v}_{sp} \cdot \nabla_x \bar{D}_S(x, \tau)$$

where $I(x)$ and κ represents the image intensity and mean curvature respectively at location x , T is a threshold computed locally from image intensities, β and α two weighting coefficients equal to $\beta = 0.1$ and $\alpha = 0.1$. Figure 3 displays our segmentation results. Despite the artificial occlusion, the shape prior term was able to recover the correct shape by attracting the shape onto the *shape prior manifold*. Yet, the final surface is geometrically-accurate because the active contour can evolve freely inside the manifold \mathcal{M} subject to the image term. The red-cross in Fig. 2 locates the final segmented shape in the embedding. Finally, note that, in practice, the shape prior term is not used during the first steps of the evolution (a robust alignment being impossible).




Fig. 3. a) Coronal, horizontal, and sagittal slices of the MRI volume with the final segmentations without (top) and with (bottom) the shape prior. b) Some snapshots of the shape evolution - the shape prior term was not used during the first steps. c) The closest neighbors of the final surface.

5 Conclusion and Future Work

We have proposed a new deformable model framework for image segmentation that incorporates non-linear shape priors by learning a shape prior manifold using recent manifold learning techniques. Our approach exploits carefully the properties of Diffusion maps to derive an innovative shape prior term designed to attract an active contour towards the shape manifold. While preliminary, our segmentation results on shapes of ventricle nucleus demonstrate the potential of our approach.

The proposed method is quite general and is not necessarily restricted to specific 3-dimensional segmentation tasks. In particular, the only requirement is a differentiable kernel. We plan to apply our approach to more general data sets, such as diffusion weighted imaging as well as combined anatomical and functional MRI, in future work.

References

1. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models-their training and application. *Computer Vision and Image Understanding* **61**(1) (1995) 38–59
2. Leventon, M., Grimson, E., Faugeras, O.: Statistical shape influence in geodesic active contours. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (2000) 316–323
3. Osher, S., Sethian, J.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of Computational Physics* **79**(1) (1988) 12–49
4. Rousson, M., Paragios, N.: Shape priors for level set representations. In: *European Conference on Computer Vision*. Volume 2. (2002) 78–92
5. Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging* **22**(2) (2003) 137–154
6. Cremers, D., Kohlberger, T., Schnörr, C.: Nonlinear shape statistics in mumford shah based segmentation. In: *European Conference on Computer Vision*. (2002) 93–108
7. Charpiat, G., Faugeras, O., Keriven, R.: Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics* **5**(1) (2005) 1–58
8. Pizer, S.M., et al.: Deformable M-Reps for 3D medical image segmentation. *International Journal of Computer Vision* **55**(2–3) (2003) 85–106
9. Coifman, R., et al., S.L.: Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS* **102**(21) (2005) 7426–7431
10. Lafon, S., Lee, A.B.: Diffusion maps and coarse-graining: a unified framework for dimensionality reduction, graph partitioning, and data set parameterization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **28**(9) (2006) 1393–1403
11. Charpiat, G., Faugeras, O., Keriven, R., Maurel, P.: Distance-based shape statistics. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. Volume 5. (2006) 925–928
12. Bengio, Y., Vincent, P., et al.: Spectral clustering and kernel pca are learning eigenfunctions. Technical Report 1239, Département d’informatique et recherche opérationnelle, Université de Montréal (2003)
13. Roweis, S., Saul, L.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290** (2000) 2323–2326
14. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* **15**(6) (2003) 1373–1396
15. Lafon, S., Keller, Y., Coifman, R.R.: Data fusion and multicue data matching by diffusion maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(11) (2006) 1784–1797