

Projecting onto a shape prior manifold

Patrick Etyngier
Renaud Keriven
Jean-Philippe Pons

Research Report - CERTIS 06-25
November 2006

Odyssée Project



Ecole des ponts - Certis
6-8 avenue Blaise Pascal
77420 Champs-sur-Marne
France



Inria - Sophia Antipolis
2004 route des Lucioles
06902 Sophia Antipolis
France



Ecole Normale Supérieure - DI
45, rue d'Ulm
75005 Paris
France

Projecting onto a shape prior manifold

Projection sur une variété de formes *a priori*

Patrick Etyngier²
Renaud Keriven^{1,2}
Jean-Philippe Pons²

¹Ecole Normale Supérieure, DI, Paris, France, <http://www.di.ens.fr/OdysseeVision.html>

²CERTIS, Ecole des Ponts, Paris, France, <http://www.enpc.fr/certis/>

Abstract

Incorporating shape priors in image segmentation has become a key problem in computer vision. Most existing work is limited to a linearized shape space with small deformation modes around a mean shape. These approaches are relevant only when the learning set is composed of very similar shapes. Also, there is no guarantee on the visual quality of the resulting shapes. In this paper, we introduce a new framework that can handle more general shape priors. We model a category of shapes as a finite dimensional manifold, the *shape prior manifold*, which we approximate from the shape samples using the Laplacian eigenmap technique. Our main contribution is to properly define a projection operator onto the manifold by interpolating between shape samples using local weighted means, thereby improving over the naive nearest neighbor approach. Our method is stated as a variational problem that is solved using an iterative numerical scheme. We obtain promising results with synthetic and real shapes which show the potential of our method for segmentation tasks.

Résumé

L'utilisation de formes *a priori* dans les processus de segmentation d'images est un problème essentiel en vision par ordinateur. Dans la plupart des travaux existants, l'espace des formes est linéarisé à l'aide de petits modes variations autour d'une forme moyenne. De telles approches sont donc pertinentes lorsque l'ensemble d'apprentissage est composé de formes très similaires. Par ailleurs, la qualité visuelle des formes obtenues ne peut pas être garantie. Dans ce rapport de recherche, nous introduisons un nouveau cadre de travail dont l'objectif est de manier des *a priori* de formes plus généraux.

Notre modèle établit qu'une catégorie de formes est décrite comme une variété de dimension finie plongée dans l'espace des formes. Cette variété que nous désignons par *la variété des formes a priori*, est discrétisée par des échantillons de formes et apprise à l'aide des techniques de graphes laplaciens. Notre contribution principale est de définir un opérateur de projection sur la *variété des formes a priori* en interpolant entre les échantillons à l'aide de moyennes de formes pondérées, améliorant ainsi l'approche naïve par les *k plus proches voisins*. Nous proposons une formulation variationnelle du problème, résolu par un schéma numérique itératif.

Nous obtenons des résultats prometteurs avec des formes synthétiques et réelles, montrant le potentiel de notre approche dans des tâches de segmentation d'images

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Contributions	2
2	Learning the Shape Prior Manifold	3
2.1	Definitions	3
2.2	Distances between shapes	3
2.3	Manifold learning	4
3	Interpolation and Projection for the <i>Shape Prior Manifold</i>	6
3.1	Weighted means as interpolations	6
3.2	Projection onto <i>the shape prior manifold</i>	7
4	Numerical Experiments	9
5	Conclusion & Perspectives	10
	Bibliography	10

1 Introduction

1.1 Motivation

Image segmentation is an ill-posed problem due to various perturbing factors such as noise, occlusions, missing parts, cluttered data, etc. When dealing with complex images, some prior shape knowledge may be necessary to disambiguate the segmentation process. The use of such prior information in the deformable models 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 prior shape information.

An important work in this direction is the *active shape model* of Cootes et al. [6]. This approach performs a principal component analysis (PCA) on the position of some landmark points placed in a coherent way on all the training contours. The number of degrees of freedom of the model is reduced by considering only the principal modes of variation. The active shape model is quite general and has been successfully applied to various types of shapes (hands, faces, organs). However, 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 [9] circumvent these limitations by computing parameterization-independent shape statistics within the level set representation [12, 16, 11]. 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 contours framework. The evolution equation contains a term which attracts the model toward an optimal prior shape. The latter is a combination of the mean shape and of the principal modes of variation. The coefficients of the different modes and the pose parameters are updated by a secondary optimization process. Several improvements to this approach have been proposed [13, 5, 18], 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 [7].

Performing PCA on distance functions might be problematic since they do not define a vector space. To cope with this, Charpiat, Faugeras and Keriven [3] 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.

1.2 Contributions

In this paper, we introduce a new framework that can handle more general shape priors. We model a category of shapes as a smooth finite-dimensional submanifold of the infinite-dimensional shape space. In the sequel, we term this finite-dimensional manifold the *shape prior manifold*. This manifold cannot be represented explicitly. We approximate it from a collection of shape samples using a recent manifold learning technique called Laplacian embedding [2]. Manifold learning is already an established tool in object recognition and image classification. Also, very recently, Charpiat *et al.* [4] have applied the Laplacian eigenmap to a set of fish shapes for the purpose of shape analysis, and obtained promising results. But to our knowledge such techniques have never been used in the context of image segmentation with shape priors.

A Laplacian embedding of the *shape prior manifold* is interesting in itself: it reveals the dimensionality of the shape category and a spatial organization of the associated shape samples. However, this embedding alone does not help to overcome noise, occlusion or other perturbations in a segmentation task. For the *shape prior manifold* to be really useful during a segmentation process, *we need the ability to compute the closest shape of the manifold to some current candidate shape* [13].

Unfortunately, the manifold learning literature does not give a solution to this problem. These approaches are mainly interested in recovering local properties of the manifold by analyzing graph adjacency of samples. They do not focus on recovering information in between samples.

A naive nearest neighbor approach is not an acceptable solution either. First, its answers are limited to the original finite and discrete set of shape samples, which does not account for the smoothness of the *shape prior manifold*. Second, in order to produce an acceptable guess, it would require a very dense sampling of the shape category of interest which is not affordable in practice. Third, it completely disregards the dimensionality and the spatial organization revealed during the manifold learning stage.

Our main contribution is to properly define this projection operator onto the *shape prior manifold*, by interpolating between some carefully selected shape samples using local weighted means. Our method is stated as a variational problem that is solved using an iterative numerical scheme.

The remainder of this paper is organized as follows. Section 2 is dedicated to learning the shape prior manifold from a finite set of shape samples using the Laplacian embedding technique. Section 3 presents a method for interpolation of the shape prior manifold and projection onto it. In Section 4, we report on some numerical experiments which yield promising results with synthetic and real shapes.

2 Learning the Shape Prior Manifold

2.1 Definitions

In the sequel, we define a *shape* as a simple (i.e. non-intersecting) closed curve, and we denote by \mathcal{S} the space of such shapes. Please note that, although this paper only deals with 2-dimensional shapes, all ideas and results seamlessly extend to higher dimensions.

The space \mathcal{S} is infinite-dimensional. We make the assumption that a category of shapes, i.e. the set of shapes that can be identified with a common concept or object, e.g. fish shapes, can be modeled as a finite-dimensional manifold.

In the context of estimating the shape of an object in a known category from noisy and/or incomplete data, we call this manifold the *shape prior manifold*. In practice, we only have access to a discrete and finite set of example shapes in this category. We will assume that this set constitutes a "good" sampling of the *shape prior manifold*, where "good" stands for "exhaustive" and "sufficiently dense" in a sense that will be clarified below.

2.2 Distances between shapes

The notion of regularity involved by the manifold viewpoint absolutely requires to define which shapes are close and which shapes are far apart. However, currently, there is no agreement in the computer vision literature on the right way of measuring shape similarity. Many different definitions of the distance between two shapes have been proposed.

One classical choice is the area of the symmetric difference between the regions bounded by the two shapes:

$$d_{SD}(S_1, S_2) = \frac{1}{2} \int |\chi_{\Omega_1} - \chi_{\Omega_2}|, \quad (1)$$

where χ_{Ω_i} is the characteristic function of the interior of shape S_i . This distance was recently advocated by Solem in [17] to build geodesic paths between shapes.

Another classical definition of distance between shapes is the Hausdorff distance, appearing in the context of shape analysis in image processing in the works of Serra [15] and Charpiat *et al.* [3]:

$$d_H(S_1, S_2) = \max \left\{ \sup_{x \in S_1} \inf_{y \in S_2} \|x - y\|, \sup_{y \in S_2} \inf_{x \in S_1} \|x - y\| \right\}. \quad (2)$$

Another definition has been proposed [9, 13, 3], based on the representation of a curve in the plane, of a surface in 3D space, or more generally of a codimension-1 geometric object in \mathbb{R}^n , by its signed distance function. In this context, the

distance between two shapes can be defined as the L^2 -norm or the Sobolev $W^{1,2}$ -norm of the difference between their signed distance functions. Let us recall that $W^{1,2}(\Omega)$ is the space of square integrable functions over Ω with square integrable derivatives:

$$d_{L^2}(S_1, S_2)^2 = \|\bar{D}_{S_1} - \bar{D}_{S_2}\|_{L^2(\Omega, \mathbb{R})}^2, \quad (3)$$

$$d_{W^{1,2}}(S_1, S_2)^2 = \|\bar{D}_{S_1} - \bar{D}_{S_2}\|_{L^2(\Omega, \mathbb{R})}^2 + \|\nabla \bar{D}_{S_1} - \nabla \bar{D}_{S_2}\|_{L^2(\Omega, \mathbb{R}^n)}^2, \quad (4)$$

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.

2.3 Manifold learning

Once some distance d between shapes has been chosen, classical manifold learning techniques can be applied, by building an adjacency graph of the learning set of shape examples. Let $(S_i)_{i \in 1, \dots, n}$ denote the n shapes of the learning set. Two slightly different approaches can be considered to build the adjacency graph:

ε -neighborhoods: Two nodes S_i and S_j ($i \neq j$) are connected in the graph if $d(S_i, S_j) < \varepsilon$, for some well-chosen constant $\varepsilon > 0$.

k nearest neighbors: Two nodes S_i and S_j are connected in the graph if node S_i is among the k nearest neighbors of S_j , or conversely, for some constant integer k .

The study of advantages and disadvantages of both approaches is beyond the scope of this paper. An adjacency matrix $(W_{i,j})_{i,j \in 1, \dots, n}$ is then designed, the coefficients of which measure the strength of the different edges in the adjacency graph.

Once an adjacency graph is defined from a given set of samples, manifold learning consists in mapping data points into a lower dimensional space while preserving the local properties of the adjacency graph. This dimensionality reduction with minimal local distortion can advantageously be achieved using spectral methods, i.e. through an analysis of the eigen-structure of some matrices derived from the adjacency matrix.

Dimensionality reduction has enjoyed renewed interest over the past years. Among the most recent and popular techniques are the Locally Linear Embedding (LLE) [14], Laplacian eigenmaps [2] and the Locally Preserving Projections (LPP) [8].

Below, we present the mathematical formulation of Laplacian eigenmaps for data living in \mathbb{R}^n . An extension to shape manifolds is straightforward.

Let \mathcal{M} be a manifold of dimension m lying in \mathbb{R}^n ($m \ll n$). For the moment, we take $m = 1$ since generalization to any dimension $m < n$ is immediate.

The dimensionality reduction problem consists in finding a mapping $f : \mathcal{M} \rightarrow \mathbb{R}$ such that if two points x and z are *close* in \mathcal{M} , so are $f(x)$ and $f(z)$. To characterize such an optimal mapping, Belkin and Niyogi [2] proved the following inequality:

$$|f(z) - f(x)| \leq d_{\mathcal{M}}(x, z) \|\nabla f(x)\| + o(d_{\mathcal{M}}(x, z)) \quad (5)$$

where $d_{\mathcal{M}}$ is the geodesic distance on the manifold \mathcal{M} . The optimality condition then writes:

$$f^* = \arg \min_{\|f\|_{L^2(\mathcal{M})}=1} \int_{\mathcal{M}} \|\nabla f\|^2 \quad (6)$$

$$= \arg \min_{\|f\|_{L^2(\mathcal{M})}=1} \int_{\mathcal{M}} \mathcal{L}(f) f \quad (7)$$

where $\mathcal{L} = -\operatorname{div}(\nabla f)$ is the Laplace operator. Note that the equivalence between (6) and (7) is due to the Stokes theorem.

Solving the minimization problem (7) is equivalent to solving the eigen problem $\mathcal{L}(f) = \lambda f$. The optimal mapping is then given by the eigen functions 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.

In practice, a discrete counterpart to this continuous formulation must be used. The discrete approximation of the Laplace operator is given by the matrix $L = W - D$ where D is the diagonal matrix defined by $D_{i,i} = \sum_j W_{i,j}$. Optimal dimensionality reduction is then achieved by finding the eigenvectors of matrix L corresponding to the m smallest non-zero eigenvalues.

Although Laplacian eigenmaps are a powerful dimensionality reduction tool, they do not give access to an explicit projection operator on the low dimensional manifold. In contrast, the Locally Preserving Projection (LPP) [8] defines the mapping f to be locally a linear operator on the high dimensional input data. This is a very interesting property since it allows to project new data points onto the low dimensional representation. Unfortunately, LPP cannot be considered in the context of shapes since it is limited to finite-dimensional problems.

Thus, existing manifold learning techniques are not able to describe the shape prior manifold in between the training shapes. However, in the context of shape estimation from noisy or incomplete images, the two following problems are fundamental:

Problem 1 *How to interpolate the shape prior manifold from shape samples.*

Problem 2 *How to project a shape onto the shape prior manifold.*

The next section deals with these two challenging problems.

3 Interpolation and Projection for the *Shape Prior Manifold*

From now on, we assume that applying the manifold learning techniques described in the previous section allowed us to determine the dimension m of the shape prior manifold \mathcal{M} . Let $\mathcal{N} = (S_0, \dots, S_m)$ be $m + 1$ shapes of \mathcal{M} , close to each other according to the topology of \mathcal{M} . In practice, S_0 is given and S_1, \dots, S_m are its m nearest neighbors in the adjacency graph. We will refer to \mathcal{N} as a *neighborhood system*³ of \mathcal{M} . Our first goal will be to interpolate \mathcal{M} locally, given this neighborhood system.

When $m = 1$, a natural solution would be to choose as an interpolation, some geodesic path between S_0 and S_1 , with respect to distance d . Yet, computing geodesic paths between shapes is still an open problem. See [15, 1, 10, 19, 17]. Moreover, extending this solution to the case $m > 1$ would require some minimal hyper-surface of \mathcal{S} between the S_i . This seems out of reach with our current understanding of shape spaces.

3.1 Weighted means as interpolations

Charpiat *et al.* [3] define the empirical mean \bar{C} of l shapes C_1, \dots, C_l by:

$$\bar{C} = \arg \min_C \sum_{i=1}^l d(C_i, C)^2$$

Following the same path, we propose to use *weighted mean shapes* to locally interpolate \mathcal{M} between our samples S_0, \dots, S_m :

Solution 1 (Problem 1: Local interpolation of the shape manifold) *Let \mathcal{M} be a finite m dimensional shape manifold and $\mathcal{N} = (S_0, \dots, S_m)$ be a neighborhood system as previously defined. Let $\Lambda = (\lambda_0, \dots, \lambda_m)$ with $(\lambda_i \geq 0, \sum \lambda_i = 1)$ be some weights. We call a local interpolation of \mathcal{M} according to \mathcal{N} , the following weighted mean:*

$$\bar{S}_{\mathcal{N}}(\Lambda) = \arg \min_S \sum_{i=0}^m \lambda_i d(S_i, S)^2$$

Λ can be viewed as a local parametrization of \mathcal{M} in the neighborhood system \mathcal{N} . The set covered by $\bar{S}_{\mathcal{N}}(\Lambda)$ for all the possible values of Λ provides a continuous approximation of the manifold *between* the shapes of \mathcal{N} .

³More sophisticated choices are possible but beyond the scope of this paper




Figure 1: Six weighted means. $\lambda = 0, 0.2, 0.4, 0.6, 0.8$ and 1 following the arrows.

As in [3], the interpolation $\bar{S}_{\mathcal{N}}(\Lambda)$ is obtained by a gradient descent, a shape S evolving according to a gradient flow:

$$-\sum_i \lambda_i d(S_i, S) \nabla d(S_i, S) \quad (8)$$

Figure 1 shows an example of such means for two given shapes. Although this involves two shapes only, please note that: (i) the number of shapes is not limited to $m = 1$, and (ii) even when $m = 1$, the path defined by the weighted means is neither a geodesic for some distance, nor a straight gradient descent from S_0 to S_1 . Examples with more than two shapes are given in section 4.

Fitted with a way to locally complete the shape manifold, we can now proceed to the projection problem.

3.2 Projection onto *the shape prior manifold*

Image segmentation methods that take shape priors into account, generally require the projection (in some sense) of a shape candidate onto the set of shape samples. As previously mentioned, this projection is often just the mean of the samples, and sometimes a variation of this mean according to *deformation modes*. Here, we propose a projection based on our local interpolation:

Solution 2 (Problem 2: projection onto the *shape prior manifold*) Let \mathcal{M} be a finite m dimensional shape manifold. Let M be a shape of \mathcal{S} . Let $\mathcal{N}(M) =$




Figure 2: The Snail algorithm : steps are indexed $1, 2, \dots, i_f$

(S_0, \dots, S_m) be a neighborhood system of \mathcal{M} close to M (in practice S_0 is the nearest neighbor of M and S_1, \dots, S_m are chosen as previously). We define the local projection $\Pi_{\mathcal{M}}(M)$ of M onto \mathcal{M} to be the interpolation according to $\mathcal{N}(M)$ that is the closest to M :

$$\begin{aligned} \Pi_{\mathcal{M}}(M) &= \bar{S}_{\mathcal{N}(M)}(\Lambda_{\Pi}) \\ \text{with } \Lambda_{\Pi} &= \arg \min_{\Lambda} d(M, \bar{S}_{\mathcal{N}(M)}(\Lambda)) \end{aligned} \quad (9)$$

While such a projection is clearly better than choosing the nearest neighbor, the energy involved in equation (9) cannot be minimized easily. The variations with respect to Λ of the distance $d(M, \bar{S}_{\mathcal{N}(M)}(\Lambda))$ between the interpolation and shape M are intricate. The gradient of this distance could be written, but, involving second order shape derivatives, it yields a complex minimization scheme that might not be useful for our purpose. Keeping shape priors in mind, it appears that an approximation of the projection $\Pi_{\mathcal{M}}(M)$ is sufficient.

Many algorithms might be designed to get an approximate solution to (9). We suggest an iterative scheme illustrated figure 2 that we call the *snail algorithm*. Although it is not guaranteed to converge, it is fast and proved to give good approximations of the projection of a candidate shape onto the shape prior manifold, in only a few iterations. Actually, we investigated more extensive searches of

the minimum of (9) without any significant improvement. The snail algorithm is defined by:

Solution 3 (Approximation of minimization (9)) *Let \mathcal{M} , M and $\mathcal{N}(M)$ be defined as in solution 2. The snail algorithm proceeds as follows:*

1. Initialization: choose the shapes of the neighborhood system as initial guesses.

For $i = 0, \dots, m$, let $\Lambda^i = (\lambda_0^i, \dots, \lambda_m^i)$ be defined by $\lambda_j^i = \delta_{ij}$

2. Iterations: look for a better projection between the latest estimate and the one computed $m + 1$ steps before.

For $i = m, m + 1, \dots$ until convergence, estimate:

$$\Lambda^{i+1} = \alpha_i \Lambda^i + (1 - \alpha_i) \Lambda^{i-m}$$

with $\alpha_i = \arg \min_{0 \leq \alpha \leq 1} d(M, \bar{S}_{\mathcal{N}(M)}(\alpha \Lambda^i + (1 - \alpha) \Lambda^{i-m}))$ (10)

3. Exit:

Let i_f be the index of last iteration. Approximate the projection by:

$$\Pi_{\mathcal{M}}(M) = \bar{S}_{\mathcal{N}(M)}(\Lambda^{i_f})$$

Note that we still need to design a minimization scheme to estimate the optimal α in (10). Again, a variational method is both too slow for our purpose and useless for an approximation. Computing a small number of interpolations and keeping the best one turns out to be satisfactory. Moreover, because these interpolations are obtained through a gradient descent (see [3]), estimating the interpolations for an increasing series of α is efficient, each interpolation being the initial guess for the next one.

4 Numerical Experiments

In this section, we present some results obtained with both synthetic and real shapes. Our *shape prior manifold* examples are based on a set of rectangles and a set of fishes built in [4]. We reproduce the graph Laplacian obtained in both cases in figure 3. The set of rectangle is randomly chosen such that the distribution of their corners is the uniform law in an authorized area (orientation between $-\frac{\pi}{6}$ and $\frac{\pi}{6}$, length between 2 and time the width). The set of fish is a subset of SQUID database.

In order to show the reliability of the method, we constructed corrupted shapes




Figure 3: Graph Laplacian (Courtesy [4])

from which we extract the neighbor system defined above and computed the best projection onto the *shape prior manifold*. In the toy example, the dimension of the *shape prior manifold* is 2 and thus the interpolation is obviously between 3 shapes. A rectangle is chosen to lie between two angular positions and two different sizes. this rectangle is corrupted in order to move it away from the *shape prior manifold* We show in figure 4 the neighbor system chosen, the corrupted shape and its projection.

We provide also prominent results with the fish example. We have highly corrupted a fish shape M : the head is deformed and the shape suffer from many occlusion. Of course, such a shape does not belong to the set used to build the graph Laplacian. Then, we determined the neighbor system S_0, \dots, S_3 and the projection $\Pi_{\mathcal{M}}(M)$ onto the *shape prior manifold*. Such a projection is clearly better than the nearest neighbor as illustrated in figure 6. Our algorithm overcomes most of shape occlusions and deformations.

5 Conclusion & Perspectives

We proposed a new framework for image segmentation that incorporates more general priors by learning a *shape prior manifold*. We provided a solution to interpolate between shape samples, defined a projection operator onto the *shape prior manifold* and suggested its fast estimation by means of an iterative process. Finally, numerical experiments on synthetic and real images showed promising results and the potential of the method for image segmentation. Incorporating it into a complete segmentation process is actually work in progress.




Figure 4: Toy example : Projection onto *the shape manifold*




Figure 5: Fish example : Projection onto *the shape manifold*




Figure 6: Fish example : Comparison between the nearest neighbor S_0 and the projection $\Pi_{\mathcal{M}}(M)$

References

- [1] M. Faisal Beg, Michael I. Miller, Alain Trounev, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision*, 61(2):139–157, 2005.
- [2] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [3] G. Charpiat, O. Faugeras, and R. Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics*, 5(1):1–58, 2005.
- [4] G. Charpiat, O. Faugeras, R. Keriven, and P. Maurel. Distance-based shape statistics. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 5, pages 925–928, 2006.
- [5] Y. Chen, H. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K. Gopinath, R. Briggs, and E. Geiser. Using prior shapes in geometric active contours in a variational framework. *The International Journal of Computer Vision*, 50(3):315–328, 2002.
- [6] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models—their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.
- [7] Daniel Cremers, Timo Kohlberger, and Christoph Schnörr. Nonlinear shape statistics in Mumford-Shah based segmentation. In *European Conference on Computer Vision*, pages 93–108, 2002.
- [8] X. He and P. Niyogi. Locality preserving projections. In *Advances in Neural Information Processing Systems 16*. MIT Press, 2004.
- [9] M. Leventon, E. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 316–323, 2000.
- [10] Peter W. Michor and David Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.*, 8:1–48, 2006.
- [11] S. Osher and R.P. Fedkiw. Level set methods: an overview and some recent results. *Journal of Computational Physics*, 169(2):463–502, 2001.

- [12] S. Osher and J.A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- [13] M. Rousson and N. Paragios. Shape priors for level set representations. In *European Conference on Computer Vision*, volume 2, pages 78–92, 2002.
- [14] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [15] J. Serra. Hausdorff distances and interpolations. In *International Symposium on Mathematical Morphology and its Applications to Image and Signal Processing*, pages 107–114, 1998.
- [16] J.A. Sethian. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Sciences*. Cambridge Monograph on Applied and Computational Mathematics. Cambridge University Press, 1999.
- [17] J.E. Solem. Geodesic curves for analysis of continuous implicit shapes. In *International Conference on Pattern Recognition*, volume 1, pages 43–46, 2006.
- [18] A. Tsai, A.J. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, W.E. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22(2):137–154, 2003.
- [19] Anthony Yezzi and Andrea Mennucci. Conformal metrics and true “gradient flows” for curves. *iccv*, 1:913–919, 2005.