Designing Spatially Coherent Minimizing Flows for Variational Problems Based on Active Contours

Guillaume Charpiat Renaud Keriven Jean-Philippe Pons Olivier Faugeras
Odyssée Laboratory
ENPC / ENS Paris / INRIA Sophia-Antipolis, France
Jean-Philippe.Pons@sophia.inria.fr

Abstract

This paper tackles an important aspect of the variational problems involving active contours, which has been largely overlooked so far: the optimization by gradient flows. Classically, the definition of a gradient depends directly on the choice of an inner product structure. This consideration is largely absent from the active contours literature. Most authors, overtly or covertly, assume that the space of admissible deformations is ruled by the canonical L^2 inner product. The classical gradient flows reported in the literature are relative to this particular choice. In this paper, we investigate the relevance of using other inner products, yielding other gradient descents, and some other minimizing flows not deriving from any inner product. In particular, we show how to induce different degrees of spatial coherence into the minimizing flow, in order to decrease the probability of getting trapped into irrelevant local minima. We show with some numerical experiments that the sensitivity of the active contours method to initial conditions, which seriously limits its applicability and its efficiency, is alleviated by our application-specific spatially coherent minimizing flows.

1. Introduction

Many problems in computer vision can advantageously be cast in a variational form, i.e. as a minimization of an energy functional. In this paper, we focus on variational methods dedicated to the recovery of contours. In this case, the problem amounts to finding a contour which corresponds to a global minimum of the energy. Unfortunately, in most cases, the exact minimization of the energy functional is computationally unfeasible due to the huge number of unknowns.

The *graph cuts* method is a powerful energy minimization method which allows to find a global minimum or a strong local minimum of an energy. In the last few years, this method has been successfully applied to several problems in computer vision, including stereovision [10] and image segmentation [2]. However, it has a severe limitation: it cannot be applied to an arbitrary energy function [11], and, when applicable, is computationally expensive.

Hence, in most cases, a suboptimal strategy must be adopted. A common minimization procedure consists in evolving an initial contour, positioned by the user, in the direction of steepest descent of the energy. This approach, known in the literature as *active contours* or *deformable models*, was pioneered by Kass. et al. in [9] for the purpose of image segmentation. Since, it has been applied in many domains of computer vision and image analysis (image segmentation [3], surface reconstruction [22], stereo reconstruction [6, 8, 18], etc.).

However, due to the highly non-convex nature of most energy functionals, a gradient descent flow is very likely to be trapped in a local minimum. Also, this local minimum depends on the position of the initial contour. If the latter is far from the expected final configuration, the evolution may be trapped in a completely irrelevant state. This sensitivity to initial conditions seriously limits the applicability and efficiency of the active contours method.

In the following we note Γ a codimension one contour in $\mathbb{R}^n$ and $E(\Gamma)$ the energy functional to be minimized. In order to define the gradient of the energy functional, the first step is to compute its Gâteaux derivatives in all directions, i.e. for all admissible velocity fields v:

$$\delta E(\Gamma, v) \stackrel{def}{=} \lim_{\epsilon \to 0} \frac{E(\Gamma + \epsilon v) - E(\Gamma)}{\epsilon} \,. \qquad (1)$$

Then, we would like to pick the gradient as the direction of steepest descent of the energy. However, it is not yet possible at this stage: to be able to assess the steepness of the energy, the deformation space needs additional structure, namely an inner product introducing the geometrical notions of angles and lengths. This consideration is largely absent from the active contours literature: most authors, overtly or covertly, assume that the deformation space is

ruled by the canonical L^2 inner product on Γ:

$$\langle u, v \rangle_{L^2} = \int_\Gamma u(\mathbf{x}) \cdot v(\mathbf{x}) d\mathbf{x} \ ,$$

where $d\mathbf{x}$ is the area element of the contour.

Here, for sake of generality, we model the space of admissible deformations as an inner product space $(F, \langle , \rangle_F)$. If there exists a vector $u \in F$ such that

$$\forall v \in F, \ \delta E(\Gamma, v) = \langle u, v \rangle_F \ ,$$

then u is unique, we call it the gradient of E *relative to the inner product* $\langle , \rangle_F$, and we note $u = \nabla_F E(\Gamma)$. Clearly, each choice of inner product yields its own gradient. This is often neglected and most authors improperly refer to *the* gradient of the energy. Thus, the classical gradient flows reported in the literature (mean curvature flow, geodesic active contours [3, 7, 19], etc.) are relative to the L^2 inner product.

A slightly different definition of the gradient, based on a representation of the space of admissible surfaces as a differential manifold, is proposed in [19]. However, this definition requires the contours and the deformations to be smooth, the energy functional to be differentiable, and the deformation space to be a separable Hilbert space. Our definition is more general since it only demands the existence of the directional derivatives.

The gradient descent method consists in deforming an initial contour Γ_0 in the opposite direction of the gradient.

$$\left\{ \begin{array}{l} \Gamma(0) = \Gamma_0 \\ \dfrac{d\Gamma}{dt} = -\nabla_F E(\Gamma) \end{array} \right. \tag{2}$$

The problem of the existence and the uniqueness of this minimizing flow is out of the scope of this article. Indeed, it is highly dependent on the properties of each particular energy functional. If this evolution exists, it decreases the energy:

$$\frac{dE(\Gamma)}{dt} = - \left\| \nabla_F E(\Gamma) \right\|_F^2 \leq 0 \ .$$

The standard choice for F is the Hilbert space of square integrable velocity fields $L^2(\Gamma, \mathbb{R}^n)$ equipped with its canonical inner product. Very few authors in the active contours area have considered using other inner products, whereas this is an established technique in image registration [20]. Very recently, in the context of shape representation and analysis, [12, 21] have shown that slightly modifying the L^2 inner product allows to build well-behaved metrics in the space of curves.

Minimizing flows not deriving from any inner product, that is to say evolutions that decrease the energy, without any gradient interpretation, have also been overlooked so far. Note that any evolution fulfilling the condition

$$\frac{dE(\Gamma)}{dt} = \left\langle \nabla_F E(\Gamma), \frac{d\Gamma}{dt} \right\rangle_F \leq 0 \tag{3}$$

is a candidate to solve the minimization problem. This idea, proposed in [19], is applied by the same authors to the alignment of curve in images in [14]: a complicated term in the gradient is safely neglected after checking that the evolution still decreases the energy.

The spirit of our work is different. We do not focus either on a specific inner product or on a particular energy functional. We rather explore general procedures to build some new inner products and to compute the associated gradients. We also address the design of non-gradient minimizing flows.

Our motivation is also different. Our primary concern in this work is the sensitivity of the active contours method to initial conditions. There are essentially two ways of dealing with this problem: positioning the initial contour very close to the expected final configuration, or using a multiresolution coarse-to-fine strategy, in other words running the optimization on a series of smoothed and subsampled contours and input data. In this paper, we pioneer a third way to tackle the problem of unwanted local minima: the careful design of the minimizing flow.

We do not modify the energy, hence the relief of the energy landscape and in particular the "number" of local minima remains unchanged. But by using an evolution that favors certain types of directions, we expect to decrease the probability of falling into unwanted energy basins.

Typically, in many applications, spatially coherent motions are to be preferred over erratic evolutions. For example, in the tracking problem, the object of interest is likely to have similar shapes in consecutive frames. So if we init the contour with the result of the previous frame, it makes sense to encourage the motions which preserve its overall appearance. This way, it may be easier to dodge unexpected local low-energy configurations. A traditional L^2 gradient descent definitely does not have this desirable property since the L^2 inner product completely disregards the spatial coherence of the velocity field.

The rest of this paper is organized as follows. In Section 2, we carry out an abstract study of gradient and non-gradient minimizing flows. In Section 3, we propose some particular flows that yield different degrees of spatial coherence. Finally, in Section 4, we show with some numerical experiments that the robustness of the active contours method to local minima is improved by our application-specific spatially coherent minimizing flows.

2. Abstract study

2.1. Designing new inner products

We suppose that the space F of admissible deformations is initially equipped with the inner product $\langle , \rangle_F$. Then, for any symmetric positive definite linear operator $L : F \to F$,

a new inner product can be defined by

$$\langle u, v \rangle_L = \langle Lu, v \rangle_F \ . \tag{4}$$

The following observation is central to our work: if $\nabla_F E(\Gamma)$ exists and if L is invertible, then $\nabla_L E(\Gamma)$ also exists and we have

$$\nabla_L E(\Gamma) = L^{-1} \nabla_F E(\Gamma) \ . \tag{5}$$

The above procedure is of great practical interest because it allows to upgrade any existing L^2 gradient flow. However, it is not completely general in the sense than all inner products cannot be expressed in this form. This construction is illustrated in Subsections 3.1 and 3.2 by some particular inner products yielding spatially coherent gradient flows.

2.2. Designing new minimizing flows

In this subsection, we follow the inverse approach. Instead of working on the inner product, we apply a linear operator $L : F \to F$ to the gradient, and we study the properties of the resulting flow:

$$\frac{d\Gamma}{dt} = -L \ \nabla_F E(\Gamma) \ . \tag{6}$$

This naturally sets up a hierarchy among different types of operators:

- if L is positive, the energy is non-increasing along the flow (6). Indeed,

$$\frac{dE(\Gamma)}{dt} = - \langle \nabla_F E(\Gamma), L \nabla_F E(\Gamma) \rangle_F \leq 0 \ .$$

- if L is positive definite, the energy strictly decreases along the flow (6) until a critical point of the original gradient flow (2) is reached.

- if L is symmetric positive definite and invertible, the flow (6) coincides with a gradient descent relative to the inner product $\langle , \rangle_{L^{-1}}$, as defined in equation (4).

The third case is contained in Subsection 2.1. The second case is illustrated in Subsection 3.3 by a Gaussian smoothing of the gradient along the contour, in order to generate a smoother minimizing flow.

3. Spatially coherent minimizing flows

In this section, we propose some minimizing flows yielding different degrees of spatial coherence. We insist on the fact that this spatial coherence has nothing to do with an eventual regularity term in the energy functional. We do not modify the energy, so the regularity constraint on the contour remains unchanged. We modify the trajectory of the minimizing flow, by favoring spatially coherent motions, but this does not condition the regularity of the final contour.

In the following, we make an intense use of differential geometry. We refer the reader to [5] for the basic notions.

3.1. The H^1 gradient flow

A first way to introduce a notion of spatial coherence is to use an inner product that penalizes not only the length of the velocity field, but also its variations along the contour. To this end, we consider the canonical inner product of the Sobolev space $H^1(\Gamma, \mathbb{R}^n)$ of square integrable velocity fields with square integrable derivatives:

$$\langle u, v \rangle_{H^1} = \int_\Gamma u(\mathbf{x}) \cdot v(\mathbf{x}) d\mathbf{x} + \int_\Gamma \nabla_\Gamma u(\mathbf{x}) \cdot \nabla_\Gamma v(\mathbf{x}) d\mathbf{x} \ ,$$

where ∇_Γ denotes the intrinsic gradient on the contour. The H^1 inner product is related to the L^2 inner product by equation (4) with $L(u) = u - \Delta_\Gamma u$, where Δ_Γ denotes the intrinsic Laplacian operator on the contour, often called the Laplace-Beltrami operator.

As a result, $\nabla_{H^1} E$ is a smoothed version of $\nabla_{L^2} E$ which can be obtained either by solving an intrinsic heat equation with a data attachment term:

$$\Delta_\Gamma u = u - \nabla_{L^2} E \ , \tag{7}$$

or by finding the optimum of:

$$\arg\min_u \int_\Gamma |\nabla_\Gamma u(\mathbf{x})|^2 d\mathbf{x} + \int_\Gamma |u(\mathbf{x}) - \nabla_{L^2} E(\mathbf{x})|^2 d\mathbf{x} \ .$$

To sum up, using the H^1 inner product instead of the L^2 inner product leads to a smoother gradient flow.

3.2. Motion decomposition

Another simple and useful procedure to design new inner products yielding spatially coherent flows, is to decompose the deformation space into a sum of several mutually orthogonal linear subspaces, and to apply different penalty factors to the different types of motions. Typically, the subspaces are chosen according to an application-specific hierarchy of the motions. For example, rigid/non-rigid, affine/non-affine, etc.

We suppose that such an orthogonal (with respect to $\langle , \rangle_F$) decomposition of the deformation space F into N closed linear subspaces is available:

$$F = F_1 \perp \cdots \perp F_N \ .$$

Then a new inner product is derived from $\langle , \rangle_F$ by applying the procedure of Subsection 2.1 with

$$L = \bigoplus_{i=1}^{N} \lambda_i \, \mathrm{Id}_{F_i} \ ,$$

where $\forall i, \ \lambda_i > 0$. The lower is λ_i, the shorter is the length of the velocity fields of F_i, and the stronger will be this type of motion in the new gradient flow.

Obviously, L is symmetric positive definite and invertible. If $\nabla_F E$ exists, so does $\nabla_L E$ and

$$\nabla_L E = \sum_{i=1}^{N} \frac{1}{\lambda_i} \Pi_{F_i} \nabla_F E \ , \tag{8}$$

where Π_{F_i} denotes the orthogonal projection on the i^{th} subspace. Of course, if all λ_i are equal to 1, $\nabla_L E$ coincides with $\nabla_F E$.

We apply this idea to two useful cases. In the first case, we decompose the velocity field into a translation, an instantaneous rotation, a rescaling motion and a non-rigid residual. In the second case, we isolate the instantaneous affine motion.

In the following, we note $M(\Gamma) = \int_\Gamma d\mathbf{x}$ the mass of the contour (its length in 2D, its area in 3D, and so on) and $\overline{f} = \left(\int_\Gamma f(\mathbf{x}) d\mathbf{x}\right) / M(\Gamma)$ the average of a scalar or vector quantity on Γ. With this notation in hand, the centroid of the contour writes $\overline{\mathbf{x}}$.

3.2.1 Translation, rotation and scaling

In this paragraph, we focus on the two-dimensional and three-dimensional cases. The expressions below are for the 3D case, but can easily be adapted to 2D.

We note T, R and S the subspaces of the translations, the instantaneous rotations around the centroid, and the scaling motions centered on the centroid, respectively:

$$\begin{aligned} T &= \left\{ u : \mathbf{x} \mapsto \mathbf{t} \mid \mathbf{t} \in \mathbb{R}^3 \right\} \ , \\ R &= \left\{ u : \mathbf{x} \mapsto (\mathbf{x} - \overline{\mathbf{x}}) \wedge \omega \mid \omega \in \mathbb{R}^3 \right\} \ , \\ S &= \left\{ u : \mathbf{x} \mapsto s(\mathbf{x} - \overline{\mathbf{x}}) \mid s \in \mathbb{R} \right\} \ . \end{aligned}$$

These subspaces are mutually orthogonal for the L^2 inner product. We suppose that they are included in the space of admissible deformations F, and that the latter is ruled by the L^2 inner product. We note G the orthogonal complement of these subspaces: $F = T \perp R \perp S \perp G$. The orthogonal projection of a velocity field u on T, R and S writes:

$$\begin{aligned} \Pi_T u(\mathbf{x}) &= \overline{u} \ , \\ \Pi_R u(\mathbf{x}) &= (\mathbf{x} - \overline{\mathbf{x}}) \wedge \frac{\int_\Gamma u(\mathbf{x}) \wedge (\mathbf{x} - \overline{\mathbf{x}})}{\int_\Gamma \|\mathbf{x} - \overline{\mathbf{x}}\|^2} \ , \\ \Pi_S u(\mathbf{x}) &= \frac{\int_\Gamma u(\mathbf{x}) \cdot (\mathbf{x} - \overline{\mathbf{x}})}{\int_\Gamma \|\mathbf{x} - \overline{\mathbf{x}}\|^2} (\mathbf{x} - \overline{\mathbf{x}}) \ . \end{aligned}$$

The new gradient is deduced from the L^2 gradient by equation (5) with

$$L^{-1} = \mathrm{Id} + \left(\frac{1}{\lambda_T} - 1\right) \Pi_T + \left(\frac{1}{\lambda_R} - 1\right) \Pi_R + \left(\frac{1}{\lambda_S} - 1\right) \Pi_S \ .$$

The weights λ_T, λ_R and λ_S are adapted to the user's needs in each particular application. For example:

- Boost rigid+scaling motions: $\lambda_T, \lambda_R, \lambda_S \ll 1$,
- Boost rigid motions: $\lambda_T, \lambda_R \ll 1, \lambda_S = 1$,
- Boost translations: $\lambda_T \ll 1, \lambda_R = \lambda_S = 1$.

3.2.2 Affine motion

We can apply this same idea to the subspace A of instantaneous affine motions:

$$A = \left\{ u : \mathbf{x} \mapsto \mathbf{A}\mathbf{x} + \mathbf{b} \mid \mathbf{A} \in \mathbb{R}^{n \times n}, \ \mathbf{b} \in \mathbb{R}^n \right\} \ .$$

The L^2 orthogonal projection on this subspace writes:

$$\Pi_A u(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b} \ ,$$

where

$$\begin{aligned} \mathbf{A} &= \left[\int_\Gamma u(\mathbf{x}) (\mathbf{x} - \overline{\mathbf{x}})^T \right] \left[\int_\Gamma (\mathbf{x} - \overline{\mathbf{x}}) (\mathbf{x} - \overline{\mathbf{x}})^T \right]^{-1} \ , \\ \mathbf{b} &= \overline{u} - \mathbf{A}\overline{\mathbf{x}} \ . \end{aligned}$$

3.3. Intrinsic Gaussian smoothing

We apply the procedure of Subsection 2.2 to design a smoothed version of the L^2 gradient flow. To some extent, it resembles the H^1 gradient flow of Subsection 3.1. However, here, we apply an *ad hoc* procedure to the L^2 gradient, a Gaussian smoothing along the contour, without resorting to an inner product.

We define a linear intrinsic smoothing operator which may be seen as the counterpart on the contour of Gaussian smoothing in $\mathbb{R}^{n-1}$, by considering the solution $\tilde{u}$ of the intrinsic heat equation on Γ with initial condition u:

$$\left\{ \begin{array}{l} \tilde{u}(., 0) = u \\ \dfrac{\partial \tilde{u}}{\partial \tau} = \Delta_\Gamma \, \tilde{u} \end{array} \right. \ . \tag{9}$$

We then note $L_\tau \, u$ its solution $\tilde{u}(., \tau)$ at time $\tau \geq 0$. We use a flow (6) based on L_τ to drive the contour. In other words, to get the new flow, we diffuse the L^2 gradient isotropically on the surface during a time τ. The larger is τ, the smoother is the flow.

On the one hand, L_τ is symmetric positive. In particular, the new flow decreases the energy. But the other hand, the inversion of L_τ for $\tau > 0$ is an ill-posed anti-diffusive process. So a gradient interpretation is not available.

4. Numerical experiments

The approach presented in this paper can be applied to virtually any active contour evolution. Below, we show some particular applications which demonstrate its interest.

The content of this paper is not specific to a particular implementation of the contour evolution. In our experiments, we have used the level set framework [13], motivated by its numerical stability and its ability to handle topological changes automatically. The implicit framework also offers an elegant expression of the Laplace-Beltrami operator [1] and of the average of a quantity along the contour [16].

The additional computational cost of our approach depends on the type of minimizing flow we consider. The extra time is barely noticeable for the rigid plus scaling and affine flows of paragraphs 3.2.1 and 3.2.2. The latter only require to compute a handful of integrals on the contour. The smooth minimizing flows of Subsections 3.1 and 3.3 are more demanding. In 2D, the implicit diffusion equations (7) and (9) are equivalent to some convolutions with respect to the curvilinear coordinate on Γ. In 3D and more, they must be solved with some iterative methods, for each time step. Due to space limitations, no experiment with these flows is presented in this paper.

4.1. Shape warping

We illustrate our approach in the problem of shape warping. In this context, the energy functional to be minimized is a measure of dissimilarity between the evolving contour and a target contour. The study of shape metrics is still an active research area and there are many candidates for the dissimilarity measure. In this paper, we use a differentiable approximation of the well-known Hausdorff distance, as proposed in [4], to warp the contours of two different hands.

Figure 1 compares the evolution of the contour when using the L^2 gradient descent (top row) and a modified gradient descent favoring rigid plus scaling motions (bottom row) as in paragraph 3.2.1. Both evolutions achieve a perfect warping. However, despite the similarity of the two input shapes, the L^2 gradient flow goes through some states of completely different appearances. The trajectory followed by this flow looks particularly inefficient and unnatural, because the notion of length contained in the L^2 inner product is very far from our intuition. In contrast, the behavior of our gradient flow is natural and visually pleasing.

In Figure 2, we show a three-dimensional warping example from a Teddy bear to Hayao Miyazaki's character Totoro. Once again, a modified gradient descent favoring rigid plus scaling motions yields better results than the L^2 gradient descent.

This suggests that our approach can infer relevant correspondences between the two contours, as a byproduct of the warping process. This point-to-point matching is obtained by tracking the points along the evolution. It does not make much sense with a L^2 gradient flow, because the latter yields a strictly normal velocity field. But when using our approach, the velocity field has a meaningful tangential part. Maintaining point correspondences during the evolution is straightforward in an implementation with meshes. It is also feasible in a level set implementation, with an extension proposed in [17].

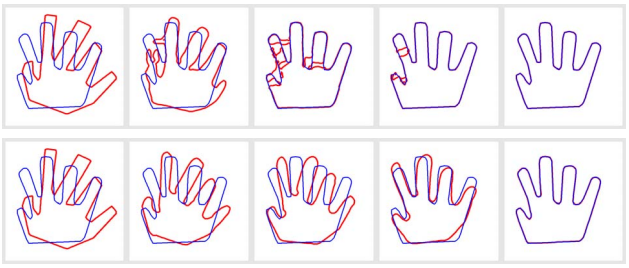

Figure 1. Shape warping with the L^2 gradient descent (top) and with a modified gradient descent favoring rigid plus scaling motions (bottom): $\lambda_T = \lambda_R = \lambda_S = 0.025$.

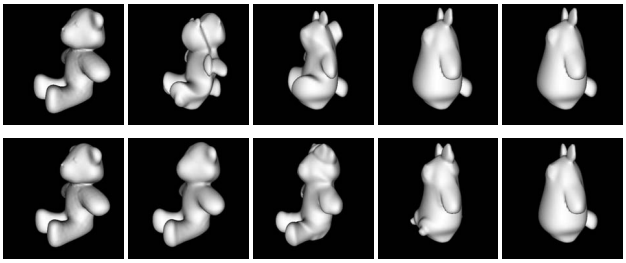

Figure 2. 3D shape warping with the L^2 gradient descent (top) and with a modified gradient descent favoring rigid plus scaling motions (bottom): $\lambda_T = \lambda_R = \lambda_S = 0.025$.

4.2. Tracking

We now illustrate the better robustness of spatially coherent minimizing flows to local minima, in the problem of tracking an object in a monocular video sequence. We have used the contour-based energy of the original geodesic active contours method [3], to track a moving hand. Note that a region-based approach [15] would give better results on our particular test sequence. However, our concern here are not the results themselves but rather the improvements brought by our approach.

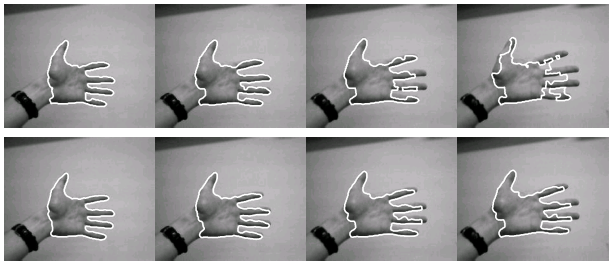

Figure 3. Tracking a hand in a video sequence with the L^2 gradient descent (top) and with a modified gradient descent favoring affine motions (bottom): $\lambda_A = 0.025$.

Figure 3 compares the evolution of the contour when using the L^2 gradient descent (top row) and a modified gradient descent favoring affine motions (bottom row) as in paragraph 3.2.2. Due to large displacements between consecutive frames, the L^2 gradient flow fails and the contour finally locks between two fingers, whereas our gradient flow manages to dodge this unwanted low-energy configuration.

5. Conclusion

The impact of the inner product structure of the deformation space on the behavior of the active contours method has been overlooked so far. In this paper, we have explored several families of inner products, as well as some minimizing flows not deriving from any inner product, which introduce different degrees of spatial coherence in the evolution of the contour. We have shown in some numerical experiments that these evolutions, as they better fit our intuitive notion of deformation cost, and as they mimic the behavior of the objects of interest, are at the same time more pleasing visually and more robust to local minima.

References

[1] M. Bertalmío, L. Cheng, S. Osher, and G. Sapiro. Variational problems and partial differential equations on implicit surfaces. *Journal of Computational Physics*, 174(2):759–780, 2001.

[2] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *International Conference on Computer Vision*, volume 1, pages 26–33, 2003.

[3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *The International Journal of Computer Vision*, 22(1):61–79, 1997.

[4] 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, Feb. 2005.

[5] M. P. DoCarmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976.

[6] O. Faugeras and R. Keriven. Variational principles, surface evolution, PDE's, level set methods and the stereo problem. *IEEE Transactions on Image Processing*, 7(3):336–344, 1998.

[7] B. Goldlücke and M. Magnor. Weighted minimal hypersurfaces and their applications in computer vision. In *European Conference on Computer Vision*, volume 2, pages 366–378, 2004.

[8] H. Jin, S. Soatto, and A. Yezzi. Multi-view stereo reconstruction of dense shape and complex appearance. *The International Journal of Computer Vision*, 63(3):175–189, 2005.

[9] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *The International Journal of Computer Vision*, 1(4):321–331, 1987.

[10] V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. In *European Conference on Computer Vision*, volume 3, pages 82–96, 2002.

[11] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004.

[12] P. Michor and D. Mumford. Riemannian geometries of space of plane curves. Preprint, 2005.

[13] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.

[14] N. Overgaard and J. Solem. An analysis of variational alignment of curves in images. In *International Conference on Scale Space and PDE Methods in Computer Vision*, 2005.

[15] N. Paragios and R. Deriche. Geodesic active regions and level set methods for motion estimation and tracking. *Computer Vision and Image Understanding*, 97(3):259–282, 2005.

[16] D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang. A PDE-based fast local level set method. *Journal of Computational Physics*, 155(2):410–438, 1999.

[17] J.-P. Pons, G. Hermosillo, R. Keriven, and O. Faugeras. How to deal with point correspondences and tangential velocities in the level set framework. In *International Conference on Computer Vision*, volume 2, pages 894–899, 2003.

[18] J.-P. Pons, R. Keriven, and O. Faugeras. Modelling dynamic scenes by registering multi-view image sequences. In *International Conference on Computer Vision and Pattern Recognition*, volume 2, pages 822–827, 2005.

[19] J. Solem and N. Overgaard. A geometric formulation of gradient descent for variational problems with moving surfaces. In *International Conference on Scale Space and PDE Methods in Computer Vision*, 2005.

[20] A. Trouvé. Diffeomorphisms groups and pattern matching in image analysis. *The International Journal of Computer Vision*, 28(3):213–21, 1998.

[21] A. Yezzi and A. Mennucci. Metrics in the space of curves. Preprint, 2005.

[22] H. Zhao, S. Osher, B. Merriman, and M. Kang. Implicit and non-parametric shape reconstruction from unorganized points using a variational level set method. *Computer Vision and Image Understanding*, 80(3):295–314, 2000.