View Point Tracking of Rigid Objects Based on Shape Sub-manifolds

Christian Gosch¹, Ketut Fundana², Anders Heyden², and Christoph Schnörr¹

¹ Image and Pattern Analysis Group, HCI, University of Heidelberg, Germany

^{2,*} Applied Mathematics Group, School of Technology, Malmö University, Sweden

Abstract. We study the task to infer and to track the viewpoint onto a 3D rigid object by observing its image contours in a sequence of images. To this end, we consider the manifold of invariant planar contours and learn the low-dimensional submanifold corresponding to the object contours by observing the object off-line from a number of different view-points. This submanifold of object contours can be parametrized by the view sphere and, in turn, be used for keeping track of the object orientation relative to the observer, through interpolating samples on the submanifold in a geometrically proper way. Our approach replaces explicit 3D object models by the corresponding invariant shape submanifolds that are learnt from a sufficiently large number of image contours, and is applicable to arbitrary objects.

1 Introduction

Motivation and Contribution. The representation of planar shapes has been a focus of research during the last few years [1, 3, 4, 5]. By mathematically separating similarity transforms and potentially also reparametrisations from other deformations of planar curves, an invariant representation of planar *shapes* is obtained in terms of a smooth manifold embedded in a euclidean space. Furthermore, distances between shapes can be computed that are only sensitive to shape deformations, by determining the geodesic path between the corresponding points of the shape manifold (Fig. 3 below provides an illustration).

In this paper, we adopt this representation and show that it is accurate enough to infer the change in aspect of a given rigid 3D object, represented by a point on the view sphere, just by observing 2D shapes of its silhouette in a given image sequence – see the left panel of Fig. 1 below.

To this end, we assume to be given a collection of silhouettes of any known object, that we represent one-to-one by a corresponding set of points on the view sphere. These data can be acquired off-line by observing the object from different directions. We regard these shapes as samples of an object-specific *submanifold* of the manifold of all planar shapes, that is parametrized by the view sphere. Taking into account the geometry of this submanifold and interpolating

* Funded by the VISIONTRAIN RTN-CT-2004-005439 Marie Curie Action within the EC's FP6.

Fig. 1. Illustration of a view sphere. Right hand: indicated are three sampled contours of an airplane seen from a camera from points on the view sphere. The object is located in the centre of the sphere. Left hand: illustration of the shape sub-manifold. The green lines between sphere and manifold indicate corresponding points, the blue arrow indicates a point that is interpolated using, in this case, three points which are neighbours on the sphere. This specific object was taken from the Princeton 3D shape benchmark [14].

the shape samples accordingly, we show that either the viewpoint of a moving camera, or object pose relative to the observer, can be tracked by observing deformations of the object's silhouette in an image sequence.

We point out the *3D models are not utilized* in this work, besides illustrating graphically various points below. Rather, a sample set of object contours observed from different viewpoints, along with the information to what object they belong, define the input data. Our results are novel and relevant, for instance, for reaching and maintaining a reference position relative to a moving object, through vision based control, in man-made and industrial scenes.

Related work. Related work has been published recently in [2, 10, 11, 12]. Etyngier et al. [10] use Laplacian eigenmaps [16] for embedding a set of training shapes into a low dimensional standard euclidean space. They present a method for projecting novel shapes to the submanifold representing the training samples, in order to model a shape prior for image segmentation. Similarly, Lee and Elgammal [11] use locally linear embedding (LLE) [17] to learn separately a configuration manifold of human gaits and a view manifold corresponding to a circle on the view sphere, based on a tensor factorization of the input data.

While nonlinear euclidean embeddings (Laplacian eigenmap, LLE) of locally connected similarity structures (weighted adjacency graphs) are employed in [10, 11], we use directly the intrinsic manifold of invariant shapes as developed in [1, 5]. Statistical models based on this manifold have been elaborated in [2, 12] for *deformable* objects and shapes of *classes* of rigid objects, respectively, in connection with image segmentation and computational anatomy.

By contrast, we focus on *tracking and pose estimation* of a single rigid object, based on contour deformations and the corresponding shape submanifold. This approach is novel. Our work may be regarded as a *learning-based* approach for

associating views and contours of *arbitrary* objects, that is both more general and easier to apply than earlier work on *model-based* contour representations of *specific* objects [25, 26].

Organization. We describe in Section 2 the mathematical shape and object representation, and the corresponding data acquisition. Section 3 details our approach for pose inference and object tracking on the view sphere. For the sake of completeness, we briefly discuss in Section 4 two major approaches to image segmentation for extracting object contours from images, although this is not the main focus of our present work. We validate our approach by numerical experiments in Section 5 and conclude in Section 6.

2 Shape Model, Object Representation, Learning

We work with the *elastic closed pre-shape space* covering closed regular two-dimensional curves, proposed in [1]. A regular curve $\alpha : [0, 1] \mapsto \mathbb{R}^2$ is represented by $\alpha(t) = \alpha_0 + \int_0^t e^{\Phi(t)} e^{i\Theta(t)} dt$, with the integrand denoting a velocity function along the curve. $e^{\Phi(t)}$ describes the curve speed, while $\Theta(t)$ is the tangent angle relative to the real axis in the complex plane. To achieve invariance under translation, the integral constant α_0 is left out, and shapes are represented by pairs (Φ, Θ) as elements of a vector space of functions denoted by $\mathcal{H}$. To also achieve scale and rotation invariance and to restrict to closed regular curves, further constraints turn $\mathcal{H}$ into the space of pre-shapes $\mathcal{C}$:

$$\mathcal{C} := \left\{ (\Phi, \Theta) \in \mathcal{H} : \begin{array}{ll} \int_0^1 e^{\Phi(t)} e^{i\Theta(t)} dt = 0 & \text{(closure)} \\ \int_0^1 e^{\Phi(t)} dt = 1 & \text{(scale)} \\ \int_0^1 \Theta(t) e^{\Phi(t)} dt = \pi & \text{(rotation)} \end{array} \right\}. \quad (1)$$

So, curves are restricted to length 1 and an angle function average of π . Note that this is an arbitrary choice, we adopted π from [1]. Invariance with respect to reparametrisations is not handled intrinsically, since it would raise a considerable additional computational burden. Instead, shapes are matched by dynamic programming, following [15].

The elastic Riemannian metric [6] used on $\mathcal{C}$ is

$$\langle (p_1, t_1), (p_2, t_2) \rangle_{(\Phi, \Theta)} := a \int_0^1 p_1(t) p_2(t) e^{\Phi(t)} dt + b \int_0^1 t_1(t) t_2(t) e^{\Phi(t)} dt \quad (2)$$

with constants $a, b \in \mathbb{R}$ that weight the amount of stretching and bending, and with tangent vectors $(p_{\{1,2\}}, t_{\{1,2\}})$ at (Φ, Θ) . [1] proposes ways to numerically approximate geodesics on a discrete representation of $\mathcal{C}$, as well as to approximate the inverse exponential map by gradient descent on $\mathcal{C}$. Another recent representation of elastic shape is discussed in [7], also cf. [9], which allows for faster computations. However, rotation invariance is not easy to achieve. [8] introduces an optimisation approach to find minimal geodesics between orbits under the action of rotations and reparametrisations.

View Sphere Sampling. The input data of our approach are given samples on the view sphere $\mathbb{S}^2$ from any object, at known positions (see Fig. 1). These data are acquired off-line and result in a sample set of points in $\mathcal{C}$.

3 Pose Inference and Tracking on the View Sphere

This section describes a model that we use for modelling motion of a point on the sphere that represents the object’s shape in a submanifold of $\mathcal{C}$, as well as a simple scheme to predict positions locally. We also explain how we keep track of points on the view sphere that correspond to shapes measured from images in an image sequence. To avoid confusion with tracking an object *in the image plane*, we call the process of tracking the position on the view sphere *sphere tracking*.

Motion Model. We model a mass point on the sphere as motion in a potential field $V(x) = \mathbf{m} \cdot g \cdot (x - P)^2$, together with a friction component. $\mathbf{m}$ is a constant inertia, g weights the impact of V , and β in Equation (3) weights the impact of friction. The motion is governed by the differential equation

$$\underbrace{-2 \cdot \mathbf{m} \cdot g \cdot (s(t) - P)}_{-\nabla V} \underbrace{-\beta \cdot \dot{s}(t)}_{\text{Stokes friction}} = \mathbf{m} \cdot \ddot{s}(t). \quad (3)$$

This is applied to a point in the tangent space of the group of 3D rotations, i.e. $s(t), P \in T\mathcal{SO}_3$, with rotations representing motions of a point on the sphere $\mathbb{S}^2$. The corresponding exponential and logarithmic maps for $\mathcal{SO}_3$ can be efficiently computed in closed form. The “centre of gravitation” P is updated whenever a new measurement P_k is available. Fig. 2 shows an illustration of the motion model following a path of points P .

Predictions. Given past measurements $p_i \in \mathbb{S}^2$, we would like to predict $s(t)$ locally. Assume to be given a new measurement P_k at time t_k , and the motion model to be at point $s(t_k)$. We then follow the trajectory governed by (3) until the distance $d(s(t_k), P_k)$ has been travelled, say at time t'_k , so that $d(s(t_k), s(t'_k)) =$

Fig. 2. Representing and tracking shape changes as motions on the view sphere. Blue: measurements P_k . Red: path $s(t)$ of the mass point. Magenta: predicted points. The start point of the trajectory is at the far left end. The green grid lines indicate the underlying sphere.

Fig. 3. Illustration of shape interpolation with Karcher means in the closed pre-shape space $\mathcal{C}$. The corners represent the original shapes, the other contours are interpolations weighted with their barycentric coordinates. The corner curves are randomly chosen from the MPEG-7-CE1 shape data base.

$d(s(t_k), P_k)$, and then continue for an additional fixed time period $\Delta t = t'_k - t_k$ to obtain the prediction

$$p_{pred} := s(t'_k + \Delta t). \quad (4)$$

As illustrated in Fig. 2, this simple “mechanical” model can result in rather sensible paths and corresponding predictions of shape changes, as detailed below.

Shape Interpolation. Interpolation of shapes on the view sphere at a point $s \in \mathbb{S}^2$ is realised by Karcher means using a local neighbourhood M of sampled shapes around s . The empirical Karcher mean is

$$\mu = \arg \min_{m \in \mathcal{C}} \sum_{i=1}^{|M|} a_i \cdot d(m, p_i)^2, \quad (5)$$

with $d(\cdot, \cdot)$ the geodesic distance on $\mathcal{C}$, and weights $a_i \geq 0$ with $\sum_i a_i = 1$. μ can in practice be calculated by gradient descent [18]. Fig. 3 illustrates the interpolation of three shapes depicted at the corners of the triangle.

Keeping Track of the Spherical Position. Assume that we know initially a point $c_k \in \mathcal{C}$ and the corresponding position $t_k \in \mathbb{S}^2$. Now, suppose a new shape $q \in \mathcal{C}$ is to be considered, typically delivered by an image segmentation algorithm that tracks an object over a number of frames (see the next section). Fig. 4 illustrates the following problem: We wish to determine a point $c_{k+1} \in \mathcal{C}$ at $t_{k+1} \in \mathbb{S}^2$ on the sub-manifold modeled by the samples p_i from the view sphere at spherical coordinates $t_i \in \mathbb{S}^2$, that is as close as possible to q . That is, we would like to minimise the geodesic distance $d(m, q) = \|\text{Log}_m(q)\|_m$ by minimising

$$F(m, q) = \|\text{Log}_m(q)\|_m^2, \quad (6)$$

Fig. 4. Keeping track of the spherical position: Shape c_k and position t_k are known, as well as a new shape q . What is the (approximate) position t_{k+1} on the view sphere corresponding to q ?

where m results from minimizing (5),

$$m(t) = \arg \min_{\tilde{m} \in \mathcal{C}} \left(\sum_{i=1}^{|M|} a_i(t) \cdot d(\tilde{m}, p_i)^2 \right) \quad (7)$$

with both the neighbourhood M and the weights a_i depending on the spherical position t . We then solve at frame $k + 1$

$$t_{k+1}^* = \arg \min_t F(m(t), q) \quad (8)$$

using non-linear conjugate gradient descent on the view sphere, as follows: choose $b_{\ell,1}, b_{\ell,2} \in \mathbb{R}^3$ to be orthonormal basis vectors of the tangent space $T_{t^\ell}(\mathbb{S}^2)$, and a small constant $\Delta > 0$. Notice that in the following equations, Exp and Log denote the exponential and inverse exponential maps on the sphere $\mathbb{S}^2$, not on the pre-shape space $\mathcal{C}$.

$$\text{trans} : T(\mathbb{S}^2) \times \mathbb{S}^2 \times \mathbb{S}^2 \mapsto T(\mathbb{S}^2), \quad v_2 = \text{trans}(v_1, t_1, t_2) \quad (9)$$

is a function that takes a tangent vector at t_1 and translates it along a geodesic from t_1 to t_2 . Then, let $t^0 = t_k^*$, $\beta_{-1} = 0$, $\tilde{d}_{-1} = 0$, and

$$v_\ell = \sum_{i=1}^2 b_{\ell,i} \cdot \frac{F(m(\text{Exp}_{t^\ell}(\Delta \cdot b_{\ell,i})), q) - F(m(t^\ell), q)}{\Delta} \quad (10)$$

$$d_\ell = -v_\ell + \beta_{\ell-1} \tilde{d}_{\ell-1} \quad (11)$$

$$t^{\ell+1} = \text{Exp}_{t^\ell}(\alpha \cdot d_\ell) \quad (12)$$

$$\tilde{d}_\ell = \text{trans}(d_\ell, t^\ell, t^{\ell+1}) \quad (13)$$

Fig. 5. Experiment tracking the view sphere position using only the segmented contours from a sequence of images. Right: shown are measurements obtained on the view sphere, for the complete sequence. Left: a few images from the sequence are shown, the corresponding interpolated contours from the shape space $\mathcal{C}$ to their right. The initial position $t_0 \in \mathbb{S}^2$ and shape s_0 were given manually. Then for each image, the result from the previous one was used as initialisation. A region based level set segmentation was used, with a curvature regularisation term after [13].

$$\tilde{v}_\ell = \text{trans}(v_\ell, t^\ell, t^{\ell+1}) \quad (14)$$

$$\beta_\ell = \frac{[v_{\ell+1} - \tilde{v}_\ell]^\top v_{\ell+1}}{v_\ell^\top v_\ell}. \quad (15)$$

v_ℓ takes the role of the gradient direction, in the tangent space of $\mathbb{S}^2$ at the current point t^ℓ . d_ℓ is the search direction, computed from the gradient v_ℓ and the previous search direction $\tilde{d}_{\ell-1}$, with factor $\beta_{\ell-1}$ calculated using the Polak-Ribière variant of the non-linear conjugate gradient method in Equation (15), which is more robust than the Fletcher-Reeves variant according to [19]. The rest of the above equations are needed to adapt to the geometry of the sphere. Specifically, we need to translate tangent vectors to the current iterate t^ℓ to be able to combine them, and we need to go back to the sphere using the exponential map.

In order to find a step length $\mathbb{R} \ni \alpha > 0$ for use in Equation (12), we use a standard line search procedure with the Armijo or *sufficient decrease* condition

$$F(m(\text{Exp}_{t^\ell}(\alpha \cdot d_\ell)), q) \leq F(m(t^\ell), q) + c \cdot \alpha \cdot (v_\ell^\top d_\ell), \quad 0 < c < 1. \quad (16)$$

Figures 5 and 6 depict paths on the view sphere.

4 Segmentation and Image Contours

There are several possibilities to obtain contours from actual images, and to track contours while they are moving *in the image plane*. We have so far applied two methods: the well known region based level set method [21] and the related, more recent method from [24]. Since [24] finds a global optimum, it is suitable if the

Fig. 6. Sphere tracking experiment with occlusion. The top row shows the tracked view sphere path on the right (the arrows indicate the direction of motion), and an illustration of the image sequence on the left. The colour coding shows the corresponding contours and view sphere positions. Using the resulting shape from each previous frame to create a prior for the segmentation algorithm enables the sphere tracking to keep going for this sequence, where a small occluding object moves in front of the object. Each row shows the area of interest from 3 subsequent frames with the superimposed segmentation result, followed by the contour representing the shape tracked on the view sphere.

images contain only a more or less homogeneous background and a single object. In more complex scenes containing clutter and heterogeneous background, the level set method that only finds local optima is advantageous. We sketch these two approaches below, and how results from the sphere tracking are used as prior for steering the segmentation process.

Level sets. Our implementation of level set segmentation uses the image energy from [21] and additionally the curvature diffusion regularisation term from [13], replacing the more common mean curvature term in the evolution in all our experiments. We also optionally use a prior energy based on [23] and [22]:

$$E_{\text{shape}} = \frac{1}{2} \int_D \left[H(\Phi(x)) - H(\tilde{\Phi}(s \Gamma x + T)) \right]^2 dx. \quad (17)$$

H denotes the Heaviside function, Φ and $\tilde{\Phi}$ are the embedding functions of the evolving contour and the prior contour, respectively. s, Γ, T are transformation parameters as described further below.

Global Segmentation Method. The variational segmentation model of [21] suffers from the existence of local minima due to the non-convexity of the energy functional. Segmentation results depend on the initialisation. To overcome this limitation, Chan et al. [24] propose algorithms which are guaranteed to find global optima as follows: For a normalised grey scale image $f(x) : D \mapsto [0, 1]$ on the domain D and constants $\lambda, c_1, c_2 \in \mathbb{R}$, a global minimiser u can be found by minimising the *convex* functional

$$\min_{0 \leq u \leq 1} \int_D |\nabla u| dx + \lambda \int_D \{(c_1 - f(x))^2 - (c_2 - f(x))^2\} u(x) dx. \quad (18)$$

It is proved in [24] that if $u(x)$ is a solution of (18), then for almost every $\mu \in [0, 1]$, $\mathbf{1}_{\{x: u(x) > \mu\}}(x)$ is a global minimizer of [21].

In order to segment an object of interest in the image plane, we modify (18) by adding an additional term as shape prior

$$\min_{0 \leq u \leq 1} \int_D |\nabla u| dx + \lambda \int_D \{(c_1 - f(x))^2 - (c_2 - f(x))^2 + (\hat{u}(x) - \tilde{u}(x))\} u(x) dx, \quad (19)$$

where $\hat{u}$ is a 'frozen' u which gets updated after each time step in the numerical implementation, and $\tilde{u}$ is the prior template. We would like (19) to be invariant with respect to euclidean transformations of the object in the 2D image plane. To this end, we add transformation parameters, as in [23], of the fixed $\hat{u}$ with respect to the prior $\tilde{u}$ by minimising

$$E_{\text{shape}} = \int_D [\hat{u}(x) - \tilde{u}(s \Gamma x + T)] u(x) dx \quad (20)$$

for the scale s , translation T , and rotation matrix Γ (rotation by an angle θ). As a result, we obtain

$$\begin{aligned} & \min_{u, s, \Gamma, T} \int_D |\nabla u| dx \\ & + \lambda \int_D \{(c_1 - f(x))^2 - (c_2 - f(x))^2 + (\hat{u}(x) - \tilde{u}(s \Gamma x + T))\} u(x) dx, \end{aligned} \quad (21)$$

which is minimised by gradient descent. This functional is no longer convex in *all* unknowns, but the convexity with respect to u facilitates the computation of the transformation parameters.

Possible Priors. Points on the view sphere predicted by the motion model can be used to provide a prior when segmenting subsequent frames of an image sequence. This can be done in several ways — the most obvious is to take the shape at $p_{\text{pred}} \in \mathbb{S}^2$ from Equation (4) as a template. To incorporate the prior into the segmentation method, it is most appealing to impose a vector field defined on a contour C that drives C along a geodesic in shape space towards the prior; this appears to be a sensible choice and has been proposed amongst others in [2]. Parametric active contour methods seem to be naturally suited for

Fig. 7. Sphere tracking with a real recorded sequence totalling 97 frames. Roughly every 20th is shown, the last three are closer. Indicated in each frame is the segmentation result (green) and aligned interpolated shape (red). Difficult situations where the view tracking goes wrong are indicated in red, yellow are situations which are just ok. The time line on the bottom indicates the situation for the whole 97 frames. The spheres on the right indicate the inferred view positions along the sequence.

this sort of modification, since they work directly on points lying on the contour. For the implicit level set method [20, 21] or the method described in [24], applying a vector field that is defined only on the level set defining the interface is a little more involved. Imposing a flow along a geodesic in the implicit framework for other distance measures has been proposed, e.g., in [27]. The prior we use is a single shape predicted by the motion model on the view sphere. The shape is interpolated using a weighted Karcher mean and converted to a binary image. This binary image is then used as a prior for segmentation.

5 Experiments and Evaluation

Figures 5 and 6 show the results of the following experiment: for a given sequence $\{I_1, \dots, I_n\}$ of images depicting a moving object, the contour c_1 and view sphere position t_1 for the first image were initialised manually. Then, using the methods from Sections 3 and 4, for each subsequent image I_{i+1} the contour c_{i+1} and the respective view sphere point t_{i+1} were updated. The contour c_i from the previous image was used for initialisation and as a weak prior for the segmentation of image I_{i+1} . The segmentation result from I_{i+1} was then used to calculate t_{i+1} , starting at t_i , using the method described in Section 3.

In Fig. 6, an occluding object was added in a different scene, which could be successfully handled by using c_i as prior template for the segmentation algorithm. For these experiments, the level set algorithm was used. The figures depict a few snapshots from the whole sequences, which respectively consist of 100 and 50 frames each. These experiments show that the sphere tracking mechanism is capable of keeping track of the view sphere position fairly well, given a sufficient number of samples on the view sphere for interpolating the shape submanifold corresponding to the object. Fig. 7 shows results for a real recorded sequence.

6 Conclusions and Further Work

We presented a method that combines techniques from elastic shape manifold modelling, segmentation and optimisation, to track the change of pose of a 3D object through tracking its contour. While the given contours of the object are currently sampled more or less uniformly on the view sphere, an adaptive sampling strategy may be investigated in future work: the amount of contour change depends on the position on $\mathbb{S}^2$ and the object in question. Advanced sampling should adapt the density of points in areas of rapid shape change on $\mathbb{S}^2$, thus exploiting the geometry of the shape submanifold already during data acquisition. However, in our experiments sampling 162 points appeared to be sufficient.

Another point concerns initialisation, which is currently done manually. Automatic initialisation may be achieved for example by a voting scheme on the first few frames, for sequences where the first few contours can be extracted well enough by any extraction method.

Regarding the segmentation prior, another option is to investigate weighted combination of a local neighbourhood of shapes around p_{pred} to create a tem-

plate with a “fuzzy” boundary, in order to take more into account inherent uncertainties of the predicted path of shapes.

A last matter worth mentioning is computation speed. Specifically, potential for speed-up is in the numerical calculation of the Log map for $\mathcal{C}$.

References

1. Mio, W., Srivastava, A., Joshi, S.: On Shape of Plane Elastic Curves. *IJCV* 73, 307–324 (2007)
2. Joshi, S.H., Srivastava, A.: Intrinsic Bayesian Active Contours for Extraction of Object Boundaries in Images. *ACCV* (2006)
3. Mio, W., Srivastava, A., Liu, X.: Contour Inferences for Image Understanding. *IJCV* 69(1), 137–144 (2006)
4. Srivastava, A., Joshi, S.H., Mio, W., Liu, X.: Statistical Shape Analysis: Clustering, Learning, and Testing. *IEEE PAMI* 27(4), 590–602 (2005)
5. Klassen, E., Srivastava, A., Mio, W., Joshi, S.H.: Analysis of Planar Shapes Using Geodesic Paths on Shape Spaces. *IEEE PAMI* 26(3), 372–383 (2003)
6. Younes, L.: Computable Elastic Distances Between Shapes. *SIAM J. on App. Math.* 58, 565–586 (1998)
7. Joshi, S.H., Klassen, E., Srivastava, A., Jermyn, I.: An Efficient Representation for Computing Geodesics Between n-Dimensional Elastic Shapes. In: *CVPR* (2007)
8. Joshi, S., Srivastava, A., Klassen, E., Jermyn, I.: Removing Shape-Preserving Transformations in Square-Root Elastic (SRE) Framework for Shape Analysis of Curves. In: Yuille, A.L., Zhu, S.-C., Cremers, D., Wang, Y. (eds.) *EMMCVPR 2007*. LNCS, vol. 4679, pp. 387–398. Springer, Heidelberg (2007)
9. Michor, P.W., Mumford, D., Shah, J., Younes, L.: A Metric on Shape Space with Explicit Geodesics. *ArXiv e-prints*, 706 (2007)
10. Etyngier, P., Segonne, F., Keriven, R.: Shape Priors using Manifold Learning Techniques. In: *ICCV* (2007)
11. Lee, C., Elgammal, A.: Modeling View and Posture Manifolds for Tracking. In: *ICCV* (2007)
12. Davis, B., Fletcher, P.T., Bullitt, E., Joshi, S.: Population Shape Regression From Random Design Data. In: *ICCV* (2007)
13. Delingette, H., Montagnat, J.: Shape and Topology Constraints on Parametric Active Contours. *CVIU* 83, 140–171 (2001)
14. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The Princeton Shape Benchmark (2004)
15. Sebastian, T., Klein, P., Kimia, B.: On Aligning Curves. *PAMI* 25, 116–125 (2003)
16. Belkin, M., Niyogi, P.: Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation* 15, 1373–1396 (2003)
17. Roweis, S.T., Saul, L.K.: Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science* 290, 2323–2326 (2000)
18. Pennec, X.: Probabilities And Statistics On Riemannian Manifolds: Basic Tools For Geometric Measurements. In: *NSIP* (1999)
19. Optimization Technology Center, N. U.: The NEOS Guide, <http://www-fp.mcs.anl.gov/OTC/Guide/>
20. Osher, S., Fedkiw, R.: *Level Set Methods and Dynamic Implicit Surfaces*. Springer, Heidelberg (2003)
21. Chan, T.F., Vese, L.A.: Active Contours Without Edges. *IEEE TIP* (2001)

22. Riklin-Raviv, T., Kiryati, N., Sochen, N.A., Pajdla, T., Matas, J.: Unlevel-Sets: Geometry and Prior-Based Segmentation. In: Pajdla, T., Matas, J.(G.) (eds.) ECCV 2004. LNCS, vol. 3024, pp. 50–61. Springer, Heidelberg (2004)
23. Chen, Y., Tagare, H.D., Thiruvankadam, S., Huang, F., Wilson, D., Gopinath, K.S., Briggs, R.W., Geiser, E.A.: Using Prior Shapes in Geometric Active Contours in a Variational Framework. IJCV 50, 315–328 (2002)
24. Chan, T.F., Esedoglu, S., Nikolova, M.: Algorithms for Finding Global Minimizers of Image Segmentation and Denoising Models. SIAM J. of App. Math. 66, 1632–1648 (2006)
25. Eggert, D.W., Stark, L., Bowyer, K.W.: Aspect Graphs and Their use in Object Recognition. Ann. Math. Artif. Intell. 13(3-4), 347–375 (1995)
26. Vijayakumar, B., Kriegman, D.J., Ponce, J.: Invariant-Based Recognition of Complex Curved 3D Objects from Image Contours. CVIU 72(3), 287–303 (1998)
27. Solem, J.E.: Geodesic Curves for Analysis of Continuous Implicit Shapes. ICPR, 43–46 (2006)