

Chapter 7

Shape Priors for Image Segmentation

Daniel Cremers

7.1 Image Analysis and Prior Knowledge

The segmentation of images into meaningful regions is among the most studied problems in image analysis. The term *meaningful* typically refers to a semantic partitioning where the computed regions correspond to individual objects in the observed scene. Unfortunately, generic purely low-level segmentation algorithms often do not provide the desired segmentation results, because the traditional low level assumptions like intensity or texture homogeneity and strong edge contrast are not sufficient to separate objects in a scene.

To stabilize the segmentation process with respect to missing and misleading low-level information, researchers have proposed to impose prior knowledge into low-level segmentation methods. In the following, we will review methods which allow to impose knowledge about the *shape* of objects of interest into segmentation processes.

In the literature there exist various definitions of the term *shape*, from the very broad notion of shape of Kendall [37] and Bookstein [5] where shape is whatever remains of an object when similarity transformations are factored out (i.e., a geometrically normalized version of a gray value image) to more specific notions of shape referring to the geometric outline of an object in 2D or 3D. In this work, we will adopt the latter view and refer to an object's silhouette or boundary as its shape. Intentionally we will leave the exact mathematical definition until later, as different representations of geometry actually imply different definitions of the term *shape* and will require different algorithms. In fact, we will see that the question of how to *represent* shapes is closely coupled to the question of finding efficient algorithms for shape optimization.

One can distinguish three kinds of shape knowledge:

D. Cremers (✉)

Departments of Computer Science & Mathematics, TU Munich, Garching, Germany
e-mail: cremers@tum.de

Table 7.1 Shapes can be represented explicitly or implicitly, in a spatially continuous or a spatially discrete setting. More recently, researchers have adopted hybrid representations [67], where objects are represented both in terms of their interior (implicitly) and in terms of their boundary (explicitly)

	Spatially continuous	Spatially discrete	
Explicit	polygons [15, 74], splines [3, 26, 36]	edgels labeling & dyn. progr. [1, 53, 60, 64, 66]	hybrid repres. LP relaxation [67]
Implicit	level-set methods [27, 51], convex relaxation [12, 23]	graph cut methods [6, 33]	

- *Low-level shape priors* typically favor shorter boundary length, that is, curves with shorter boundary have lower shape energy [4, 6, 33, 36, 48].
- *Mid-level shape priors* characterize a certain class of shapes without specifying their exact shape. For example, thin and elongated structures can be preferred to facilitate the segmentation of roads in satellite imagery or of blood vessels in medical imagery [30, 49, 55]. Similarly one can impose a prior on the low-order shape moments without otherwise constraining the shape [41].
- *High-level shape priors* favor similarity to previously observed shapes, such as hand shapes [15, 26, 34], silhouettes of humans [18, 21] or medical organs like the heart, the prostate, the lungs or the cerebellum [42, 58, 59, 71].

Among a wealth of works on shape priors for image segmentation we will focus in this chapter on high-level shape priors. Specifically, we will present a range of representative works—with many of the examples taken from the author’s own work—and discuss their advantages and shortcomings.

7.2 Explicit versus Implicit Shape Representation

Among mathematical representations of shape, one can distinguish between *explicit* and *implicit* representations. In the former case, the boundary of the shape is represented explicitly as a mapping from a chart into the embedding space. Alternatively, shapes can be represented implicitly in the sense that points in the ambient space are labeled as part of the interior or the exterior of the object. In the spatially continuous setting, the optimization of such implicit shape representations is solved by means of partial differential equations. Among the most popular representatives are the level-set method [27, 51] or alternative convex relaxation techniques [11, 12]. In the spatially discrete setting, implicit representations have become popular through the graph cut methods [7, 33]. More recently, researchers have also advocated hybrid representations where objects are represented both explicitly and implicitly [67]. Table 7.1 provides an overview of a few representative works on image segmentation using explicit and implicit representations of shape.

Figure 7.1 shows examples of shape representations using an explicit parametric representation by spline curves (spline control points are marked as black boxes),




Fig. 7.1 Examples of shape representations by means of a parametric spline curve (1st image), a signed distance function (2nd image), a binary indicator function (3rd image), and an explicit discrete representation (4th image)




Fig. 7.2 The linear interpolation of the signed distance functions associated with two human silhouettes also gives rise to intermediate shapes, yet it does not constrain the shape’s topology. The interpolation of signed distance functions is generally no longer a signed distance function

implicit representations by a signed distance function or a binary indicator function and an explicit discrete representation (4th image).

Both explicit and implicit shape representations can be used for *statistical shape learning* where one can generalize a family of plausible shapes from a few sample shapes—see Fig. 7.2.

In the following, we will give an overview of some of the developments in the domain of shape priors for image segmentation. In Sect. 7.3, we will discuss methods to impose statistical shape priors based on explicit shape representations. In Sect. 7.4, we discuss methods to impose statistical shape priors in level-set based image segmentation including the concept of *dynamical shape priors* to learn temporal models of shape evolution as priors for image sequence segmentation. And lastly, in Sect. 7.5, we will present a method to compute polynomial-time optimal segmentations with elastic shape priors.

7.3 Statistical Shape Priors for Explicit Shape Representations

Over the last decades Bayesian inference has become an established paradigm to tackle the problem of image segmentation—see [22, 76], for example. Given an input image $I : \Omega \rightarrow \mathbb{R}$ on a domain $\Omega \subset \mathbb{R}^2$, a segmentation C of the image plane Ω can be computed by maximizing the posterior probability $\mathcal{P}(C|I) \propto \mathcal{P}(I|C) \mathcal{P}(C)$, where $\mathcal{P}(I|C)$ denotes the data likelihood for a given segmentation C and $\mathcal{P}(C)$ denotes the prior probability which allows to impose knowledge about which segmentations are a priori more or less likely.

Maximizing the posterior distribution can be performed equivalently by minimizing its negative logarithm given by a cost function of the form

$$E(C) = E_{\text{data}}(C) + E_{\text{shape}}(C), \quad (7.1)$$

where $E_{\text{data}}(C) = -\log \mathcal{P}(I | C)$ and $E_{\text{shape}}(C) = -\log \mathcal{P}(C)$ are typically referred to as *data fidelity term* and *regularizer* or *shape prior*. By maximizing the posterior, one aims at computing the most likely solution given data and prior.

Over the years various data terms have been proposed. In the following, we will simply use a piecewise-constant approximation of the input intensity I [48]. More sophisticated data terms based on color likelihoods [8, 40, 50, 75] or texture likelihoods [2, 22] are conceivable.

7.3.1 Linear Shape Priors

Among the most straightforward ways to represent a shape is to model its outline as a parametric curve, for example a spline curve of degree k [14, 26, 29, 46]. For $k = 1$, we simply have a polygonal shape [74]. Such parametric representations are quite *compact* in the sense that very detailed silhouettes can be represented by a few control points. This representation can be made invariant to translation, rotation and scale by appropriate normalizations often called *procrustes analysis* [28].

With this contour representation, the image segmentation problem boils down to computing an optimal spline control point vector for a given image. The segmentation process can be constrained to familiar shapes by imposing a statistical shape prior computed from the set of training shapes. The most popular shape prior is based on the assumption that the training shapes are Gaussian distributed—see for example [15, 26, 38]. One can define a shape prior that is invariant to similarity transformations (translation, rotation and scaling) by applying the Gaussian assumption to the similarity-normalized control point vector [26]. Since the space of similarity-normalized shapes is no longer a vector-space, however, the resulting distribution will not be exactly Gaussian.

Figure 7.3 shows several intermediate steps in a gradient descent evolution on the energy (7.1) combining the piecewise constant intensity model with a Gaussian shape prior constructed from a set of sample hand shapes. Note how the similarity-invariant shape prior constrains the evolving contour to hand-like shapes without constraining its translation, rotation or scaling. We refer to this as a *linear* shape prior since admissible shapes are linear combinations of respective eigen-shapes.

Figure 7.4 shows the gradient descent evolution with the same shape prior for an input image of a partially occluded hand. Here the missing part of the silhouette is recovered through the statistical shape prior. The curve converges to the desired segmentation over rather large spatial distance.




Fig. 7.3 Evolution of a parametric spline curve during gradient descent on the energy (7.1) combining a piecewise constant intensity data term model with a Gaussian shape prior constructed from a set of sample hand shapes [26]. Since the shape prior is by construction invariant to similarity transformations, the contour easily undergoes translation, rotation and scaling during energy minimization



Fig. 7.4 Gradient descent evolution of a parametric curve with similarity invariant shape prior. The statistical shape prior permits a reconstruction of the hand silhouette in places where it is occluded

7.3.2 Nonlinear Shape Priors

In general, a given set of shapes—say the various projections of a 3D object observed from different view points or the various silhouettes of a walking person—will not be Gaussian-distributed. There are many ways to go beyond the Gaussian distribution—using mixtures of Gaussians, kernel density estimators or manifold learning techniques. Alternatively one can introduce nonlinearity by means of Mercer kernel methods. In [20], it was proposed to model the shape prior not by a Mahalanobis distance in the input space (arising from the Gaussian model), but by a corresponding distance upon a transformation $\psi : \mathbb{R}^n \rightarrow Y$ of the control point vector $z \in \mathbb{R}^n$ to some generally higher-dimensional *feature space* Y . This gives rise to a Mahalanobis distance of the form:

$$E(z) = (\psi(z) - \psi_0)^t \Sigma_\psi^{-1} (\psi(z) - \psi_0) \quad (7.2)$$

with \hat{z} being the similarity-normalized control point vector. Here ψ_0 and Σ_ψ denote the mean and covariance matrix computed for the transformed shapes:

$$\psi_0 = \frac{1}{m} \sum_{i=1}^m \psi(z_i), \quad \Sigma_\psi = \frac{1}{m} \sum_{i=1}^m (\psi(z_i) - \psi_0)(\psi(z_i) - \psi_0)^\top. \quad (7.3)$$

As shown in [20], the energy $E(z)$ above can be evaluated without explicitly specifying the nonlinear transformation ψ . It suffices to define the corresponding




Fig. 7.5 Tracking a familiar object over a long image sequence with a nonlinear statistical shape prior constructed from a set of sample silhouettes. In contrast to commonly used Gaussian shape priors, the nonlinear prior allows the emergence of a multitude of familiar shapes which are not merely a linear combination of familiar shapes

Mercer kernel [17, 47]:

$$k(x, y) := \langle \psi(x), \psi(y) \rangle, \quad \forall x, y \in \mathbb{R}^n, \quad (7.4)$$

representing the scalar product of pairs of transformed points $\psi(x)$ and $\psi(y)$. A popular choice of k is a Gaussian kernel function: $k(x, y) \propto \exp(-\frac{1}{2\sigma^2} \|x - y\|^2)$. It was shown in [20], that the resulting energy is related to the classical Parzen-Rosenblatt density estimators. As shown in Fig. 7.5, this nonlinear shape prior allows the emergence of multiple very different shapes and therefore better preserves small-scale shape details.

7.4 Statistical Priors for Level-Set Representations

Parametric representations of shape such as those presented above have numerous favorable properties. In particular, they allow the representation of rather complex shapes with few parameters, resulting in low memory requirements and low computation time. Nevertheless, the explicit representation of shape has several drawbacks: Firstly, explicit shapes require a specific choice of curve (or surface) parameterization. To factor out this dependency in the representation and in respective algorithms gives rise to computationally challenging problems of regridding or reparameterization. This becomes particularly difficult for higher-dimensional shapes. Secondly, parametric representations are difficult to adapt to varying topology of the represented shape. Numerically topology changes require sophisticated splitting and remerging procedures. Thirdly, the commonly used energies are not convex with respect to a parametric boundary representation. Gradient descent algorithms will therefore only determine locally optimal solutions.

A mathematical representation of shape which is independent of parameterization was pioneered in the analysis of random shapes by Fréchet [31] and in the school of mathematical morphology founded by Matheron and Serra [45, 70]. The level-set method [27, 51] provides a means of propagating contours C (independent of parameterization) by evolving associated embedding functions ϕ via partial

differential equations—see Fig. 7.2 for a visualization of the level-set function associated with a human silhouette. It has been adapted to segment images based on numerous low-level criteria such as edge consistency [10, 39, 44], intensity homogeneity [13, 73], texture information [9, 35, 52, 57] and motion information [24].

7.4.1 Nonparametric Shape Priors

For level-set based shape representations, researchers have fit a linear sub-space to the sampled signed distance functions [43, 59, 72]. These approaches were shown to capture some shape variability. Yet, they exhibit two limitations: Firstly, they rely on the assumption of a Gaussian distribution which is not well suited to approximate shape distributions encoding more complex shape variation—see above. Secondly, they work under the assumption that shapes are represented by signed distance functions. Yet, the space of signed distance functions is not a linear space. Therefore, in general, neither the mean nor the linear combination of a set of signed distance functions will correspond to a signed distance function.

In the following, we will propose an alternative approach for generating a statistical shape dissimilarity measure for level-set based shape representations. It is based on classical methods of (so-called non-parametric) kernel density estimation and overcomes the above limitations.

Given a set of training shapes $\{\phi_i\}_{i=1,\dots,N}$, one can introduce a nonparametric shape prior on the space of signed distance functions [25] by means of a Parzen-Rosenblatt kernel density estimator [54, 56]:

$$\mathcal{P}(\phi) \propto \frac{1}{N} \sum_{i=1}^N \exp\left(-\frac{1}{2\sigma^2} d^2(\phi, \phi_i)\right), \quad (7.5)$$

with an appropriate distance d to measure the dissimilarity of two given level-set functions. The kernel density estimator is among the theoretically most studied density estimation methods. In the finite-dimensional case, it was shown to converge to the true distribution in the limit of infinite samples (and $\sigma \rightarrow 0$).

As in the case of parametric curves, segmentation can be cast as a problem of maximum a posteriori inference which boils down to an energy minimization problem of the form

$$E(\phi) = E_{\text{data}}(\phi) + E_{\text{shape}}(\phi), \quad (7.6)$$

with $E_{\text{shape}}(\phi) = -\log \mathcal{P}(\phi)$ and an appropriate data term E_{data} .

Figure 7.6 shows a direct comparison of a level-set segmentation process without and with the non-parametric shape prior in (7.5). The shape prior permits the accurate reconstruction of an entire set of fairly different shapes. Since the shape prior is defined on the level-set function ϕ , it can easily handle topological changes of the represented curve.

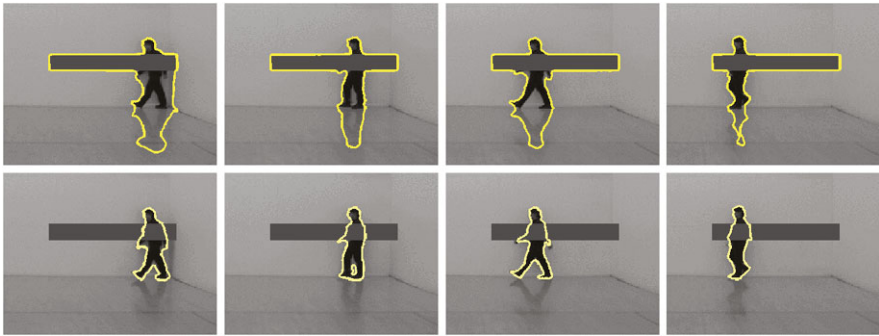


Fig. 7.6 By extending a purely data driven level set segmentation (*top row*) with a nonparametric shape prior (*bottom row*) the resulting segmentation method is robust to misleading low-level information such as shadows or partial occlusion

7.4.2 Dynamical Shape Priors for Implicit Shapes

Although the above shape priors can be applied to tracking objects in image sequences, they are not suited for this task, because they neglect the *temporal coherence of silhouettes* which characterizes many deforming shapes. In the following, we will present temporal statistical shape models for implicitly represented shapes that were first introduced in [18]. At any given time, the shape probability depends on the shapes observed at previous time steps. The integration of such dynamical shape models into the segmentation process can be formulated within a Bayesian framework for image sequence segmentation: Let $I_t : \Omega \rightarrow \mathbb{R}$ denote the input image at time t and let $\hat{\phi}_{1:t-1} := (\hat{\phi}_1, \dots, \hat{\phi}_{t-1})$ denote the segmentations obtained for the previous frames. Under the assumption that these segmentations are correct and that no knowledge about future data is available, the most likely segmentation at time t can be computed as follows:

$$\hat{\phi}_t = \arg \max_{\phi_t} \mathcal{P}(\phi_t | I_t, \hat{\phi}_{1:t-1}) = \arg \max_{\phi_t} \mathcal{P}(I_t | \phi_t) \mathcal{P}(\phi_t | \hat{\phi}_{1:t-1}). \quad (7.7)$$

Under certain assumptions, it is even possible to reinterpret the past observations in closed form [61]. The intuition is then to find the segmentation which best partitions the current image and all past images (when propagated backward in time with the dynamical model). Similarly one could take into account future observations (if available) by propagating the model forward in time.

Again, one can equivalently minimize the negative logarithm of the above expression. Gradient descent induces an evolution of the level set function which is driven both by the intensity information of the current image as well as by a dynamical shape prior which relies on the segmentations obtained for the preceding frames. Experimental evaluation demonstrates that the resulting segmentations are not only similar to previously learned shapes, but they are also consistent with the temporal correlations estimated from sample sequences. The resulting segmentation




Fig. 7.7 Variational image sequence segmentation with a dynamical shape prior on noisy and partially occluded data. 90 % noise means that nine out of ten intensity values were replaced by a random intensity. The statistically learned dynamical model allows for reliable segmentation results despite large amounts of noise (*above*) and prominent occlusion (*below*)

process can cope with large amounts of noise and occlusion because it exploits prior knowledge about *temporal* shape consistency and because it aggregates information from the input images over time (rather than treating each image independently).

As in the case of static shape priors, one can consider linear [18] or nonlinear [19] dynamical shape priors. As shown in Fig. 7.7, a linear dynamical shape prior allows reliable tracking of a walking person in an image sequence degraded by large amounts of noise and prominent occlusion.

7.5 Parametric Representations Revisited: Combinatorial Solutions for Segmentation with Shape Priors

In previous sections, we saw that shape priors improve the segmentation and tracking of familiar deformable objects, biasing the segmentation process to favor familiar shapes or even familiar shape evolutions. Unfortunately these approaches are based on locally minimizing the respective energies via gradient descent. Since these energies are generally non-convex, the computed locally optimal solutions typically depend on an initialization and may be suboptimal in practice. One exception based on implicit shape representations as binary indicator functions and convex relaxation techniques was proposed in [23]. Yet, the linear interpolation of shapes represented by binary indicator functions will generally not give rise to plausible intermediate shapes: For example, linearly interpolating two human silhouettes with one arm in different locations will fade out the arm in one location and make it emerge again in the other location. It will not translate the arm from one location to the other which would be desirable. In this sense, there is no generalization to plausible intermediate shapes.

Moreover, while implicit representations like the level-set method circumvent the problem of computing correspondences between points on either of two shapes, it is well-known that the aspect of point correspondences plays a vital role in human notions of shape similarity. For matching planar shapes, there is abundant literature




Fig. 7.8 A polynomial-time solution for matching shapes to images: matching a template curve $C : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ (left) to the image plane $\Omega \subset \mathbb{R}^2$ is equivalent to computing an orientation-preserving cyclic path $\Gamma : \mathbb{S}^1 \rightarrow \Omega \times \mathbb{S}^1$ in the product space spanned by the image domain and the template domain. The latter problem can be solved in polynomial time—see [66] for details

on how to solve this correspondence problem in polynomial time using dynamic programming techniques [32, 62, 69].

Similar concepts of dynamic programming can be employed to localize deformed template curves in images. Coughlan et al. [16] detected open boundaries by shortest path algorithms in higher-dimensional graphs. Felzenszwalb et al. used dynamic programming in chordal graphs to localize shapes, albeit not on a pixel level.

Polynomial-time solutions for localizing deformable closed template curves in images using minimum ratio cycles or shortest circular paths were proposed in [66], with a further generalization presented in [65]. There the problem of determining a segmentation of an image $I : \Omega \rightarrow \mathbb{R}$ that is elastically similar to an observed template $cc : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ is computed as a cycle

$$\Gamma : \mathbb{S}^1 \rightarrow \Omega \times \mathbb{S}^1 \quad (7.8)$$

of minimum ratio in the product space spanned by the image domain Ω and template domain \mathbb{S}^1 . See Fig. 7.8 for a schematic visualization. All points along this circular path provide a pair of corresponding template point and image pixel. In this manner, the matching of template points to image pixels is equivalent to the estimation of orientation-preserving cyclic paths, which can be solved in polynomial time using dynamic programming techniques such as ratio cycles [63] or shortest circular paths [68].

Figure 7.9 shows an example result obtained with this approach: The algorithm determines a deformed version (right) of a template curve (left) in an image (center) in globally optimal manner. An initialization is no longer required and the best conceivable solution is determined in polynomial time.

Figure 7.10 shows further examples of tracking objects: Over long sequences of hundreds of frames the objects of interest are tracked reliably—despite low contrast, camera shake, bad visibility and illumination changes. For further details, we refer to [66].




Fig. 7.9 Segmentation with a single template: despite significant deformation and translation, the initial template curve is accurately matched to the low-contrast input image. The globally optimal correspondence between template points and image pixels is computed in polynomial time by dynamic programming techniques [66]

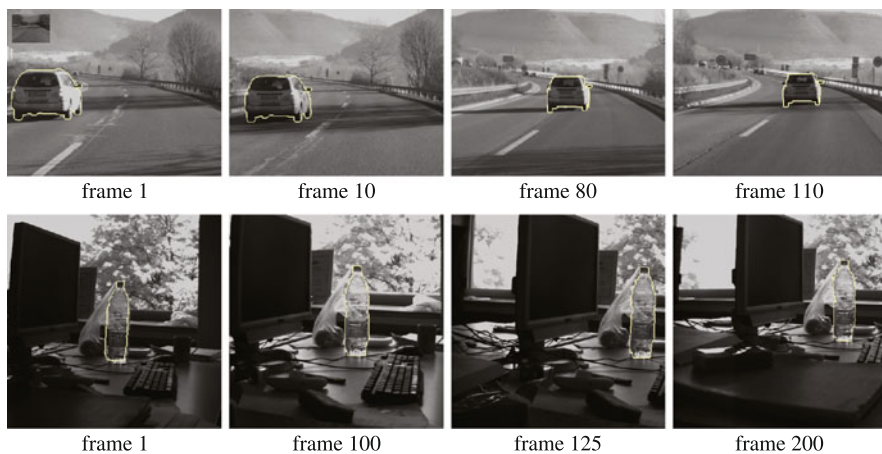


Fig. 7.10 Tracking of various objects in challenging real-world sequences [66]. Despite bad visibility, camera shake and substantial lighting changes, the polynomial-time algorithm allows to reliably track objects over hundreds of frames. Image data taken from [66]

7.6 Conclusion

In the previous sections, we have discussed various ways to include statistical shape priors in image segmentation methods. We have made several observations:

- By imposing statistically learned shape information one can generate segmentation processes which favor the emergence of familiar shapes—where familiarity is based on one or several training shapes.
- Statistical shape information can be elegantly combined with the input image data in the framework of Bayesian maximum a posteriori estimation. Maximizing the posterior distribution is equivalent to minimizing a sum of two energies representing the data term and the shape prior. A further generalization allows to impose

dynamical shape priors so as to favor familiar deformations of shape in image sequence segmentation.

- While linear Gaussian shape priors are quite popular, the silhouettes of typical objects in our environment are generally not Gaussian distributed. In contrast to linear Gaussian priors, nonlinear statistical shape priors based on Parzen–Rosenblatt kernel density estimators or based on Gaussian distributions in appropriate feature spaces [20] allow to encode a large variety of rather distinct shapes in a single shape energy.
- Shapes can be represented explicitly (as points on the object’s boundary or surface) or implicitly (as the indicator function of the interior of the object). They can be represented in a spatially discrete or a spatially continuous setting.
- The choice of shape representation has important consequences regarding the tractability of the resulting optimization problem. Moreover, different notions of shape similarity and shape interpolation are more easily expressed with respect to one or the other shape representation. As a result, there is no single ideal representation of shape. In fact, a good compromise between *desirable* and *tractable* cost functions may be obtained using hybrid representations such as the one proposed in [67]. It is an overcomplete shape representation which combines an explicit (albeit not parametric) and an implicit representation coupled via linear constraints. As a consequence, properties of both the object’s interior and its boundary can be directly accessed in the respective cost function. If this cost function is linear then LP relaxation can provide minimizers of bounded optimality.

Acknowledgements The work described here was done in collaboration with numerous researchers. The author would like to thank T. Schoenemann, F.R. Schmidt, C. Schnoerr, S. Soatto, N. Sochen, T. Kohlberger, M. Rousson and S.J. Osher for their support.

References

1. Amini AA, Weymouth TE, Jain RC (1990) Using dynamic programming for solving variational problems in vision. *IEEE Trans Pattern Anal Mach Intell* 12(9):855–867
2. Awate SP, Tasdizen T, Whitaker RT (2006) Unsupervised texture segmentation with nonparametric neighborhood statistics. In: European conference on computer vision (ECCV), Graz, Austria, May 2006. Springer, Berlin, pp 494–507
3. Blake A, Isard M (1998) *Active contours*. Springer, London
4. Blake A, Zisserman A (1987) *Visual reconstruction*. MIT Press, Cambridge
5. Bookstein FL (1978) The measurement of biological shape and shape change. *Lect notes in biomath*, vol 24. Springer, New York
6. Boykov Y, Kolmogorov V (2003) Computing geodesics and minimal surfaces via graph cuts. In: *IEEE int conf on computer vision*, Nice, pp 26–33
7. Boykov Y, Kolmogorov V (2004) An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans Pattern Anal Mach Intell* 26(9):1124–1137
8. Brox T, Rousson M, Deriche R, Weickert J (2003) Unsupervised segmentation incorporating colour, texture, and motion. In: Petkov N, Westenberg MA (eds) *Computer analysis of images and patterns*, Groningen, The Netherlands, August 2003. LNCS, vol 2756. Springer, Berlin, pp 353–360

9. Brox T, Weickert J (2004) A TV flow based local scale measure for texture discrimination. In: Pajdla T, Hlavac V (eds) European conf. on computer vision, Prague. LNCS, vol 3022. Springer, Berlin, pp 578–590
10. Caselles V, Kimmel R, Sapiro G (1995) Geodesic active contours. In: Proc IEEE intl conf on comp vis, Boston, USA, pp 694–699
11. Chambolle A, Cremers D, Pock T (2012) A convex approach to minimal partitions. *SIAM J Imaging Sci* 5(4):1113–1158
12. Chan T, Esedoğlu S, Nikolova M (2006) Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM J Appl Math* 66(5):1632–1648
13. Chan TF, Vese LA (2001) Active contours without edges. *IEEE Trans Image Process* 10(2):266–277
14. Cipolla R, Blake A (1990) The dynamic analysis of apparent contours. In: IEEE int. conf on computer vision. Springer, Berlin, pp 616–625
15. Cootes TF, Taylor CJ, Cooper DM, Graham J (1995) Active shape models—their training and application. *Comput Vis Image Underst* 61(1):38–59
16. Coughlan J, Yuille A, English C, Snow D (2000) Efficient deformable template detection and localization without user initialization. *Comput Vis Image Underst* 78(3):303–319
17. Courant R, Hilbert D (1953) *Methods of mathematical physics*, vol 1. Interscience, New York
18. Cremers D (2006) Dynamical statistical shape priors for level set based tracking. *IEEE Trans Pattern Anal Mach Intell* 28(8):1262–1273
19. Cremers D (2008) Nonlinear dynamical shape priors for level set segmentation. *J Sci Comput* 35(2–3):132–143
20. Cremers D, Kohlberger T, Schnörr C (2003) Shape statistics in kernel space for variational image segmentation. *Pattern Recognit* 36(9):1929–1943
21. Cremers D, Osher SJ, Soatto S (2006) Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int J Comput Vis* 69(3):335–351
22. Cremers D, Rousson M, Deriche R (2007) A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *Int J Comput Vis* 72(2):195–215
23. Cremers D, Schmidt FR, Barthel F (2008) Shape priors in variational image segmentation: convexity, Lipschitz continuity and globally optimal solutions. In: IEEE conference on computer vision and pattern recognition (CVPR), Anchorage, Alaska, June 2008
24. Cremers D, Soatto S (2005) Motion Competition: a variational framework for piecewise parametric motion segmentation. *Int J Comput Vis* 62(3):249–265
25. Cremers D, Sochen N, Schnörr C (2006) A multiphase dynamic labeling model for variational recognition-driven image segmentation. *Int J Comput Vis* 66(1):67–81
26. Cremers D, Tischhäuser F, Weickert J, Schnörr C (2002) Diffusion Snakes: introducing statistical shape knowledge into the Mumford–Shah functional. *Int J Comput Vis* 50(3):295–313
27. Dervieux A, Thomasset F (1979) A finite element method for the simulation of Raleigh-Taylor instability. *Springer Lect Notes in Math*, vol 771. pp 145–158
28. Dryden IL, Mardia KV (1998) *Statistical shape analysis*. Wiley, Chichester
29. Farin G (1997) *Curves and surfaces for computer-aided geometric design*. Academic Press, San Diego
30. Franchini E, Morigi S, Sgallari F (2009) Segmentation of 3d tubular structures by a pde-based anisotropic diffusion model. In: Intl. conf. on scale space and variational methods. LNCS, vol 5567. Springer, Berlin, pp 75–86
31. Fréchet M (1961) Les courbes aléatoires. *Bull Inst Int Stat* 38:499–504
32. Geiger D, Gupta A, Costa LA, Vlontzos J (1995) Dynamic programming for detecting, tracking and matching deformable contours. *IEEE Trans Pattern Anal Mach Intell* 17(3):294–302
33. Greig DM, Porteous BT, Seheult AH (1989) Exact maximum a posteriori estimation for binary images. *J R Stat Soc B* 51(2):271–279
34. Grenander U, Chow Y, Keenan DM (1991) *Hands: a pattern theoretic study of biological shapes*. Springer, New York
35. Heiler M, Schnörr C (2003) Natural image statistics for natural image segmentation. In: IEEE int. conf. on computer vision, pp 1259–1266

36. Kass M, Witkin A, Terzopoulos D (1988) Snakes: active contour models. *Int J Comput Vis* 1(4):321–331
37. Kendall DG (1977) The diffusion of shape. *Adv Appl Probab* 9:428–430
38. Kervrann C, Heitz F (1999) Statistical deformable model-based segmentation of image motion. *IEEE Trans Image Process* 8:583–588
39. Kichenassamy S, Kumar A, Olver PJ, Tannenbaum A, Yezzi AJ (1995) Gradient flows and geometric active contour models. In: *IEEE int. conf. on computer vision*, pp 810–815
40. Kim J, Fisher JW, Yezzi A, Cetin M, Willsky A (2002) Nonparametric methods for image segmentation using information theory and curve evolution. In: *Int. conf. on image processing*, vol 3, pp 797–800
41. Klodt M, Cremers D (2011) A convex framework for image segmentation with moment constraints. In: *IEEE int. conf. on computer vision*
42. Kohlberger T, Cremers D, Rousson M, Ramaraj R (2006) 4d shape priors for level set segmentation of the left myocardium in SPECT sequences. In: *Medical image computing and computer assisted intervention*, October 2006. LNCS, vol 4190, pp 92–100
43. Leventon M, Grimson W, Faugeras O (2000) Statistical shape influence in geodesic active contours. In: *Int. conf. on computer vision and pattern recognition*, Hilton Head Island, SC, vol 1, pp 316–323
44. Malladi R, Sethian JA, Vemuri BC (1995) Shape modeling with front propagation: a level set approach. *IEEE Trans Pattern Anal Mach Intell* 17(2):158–175
45. Matheron G (1975) *Random sets and integral geometry*. Wiley, New York
46. Menet S, Saint-Marc P, Medioni G (1990) B–snakes: implementation and application to stereo. In: *Proc. DARPA image underst workshop*, April 6–8, pp 720–726
47. Mercer J (1909) Functions of positive and negative type and their connection with the theory of integral equations. *Philos Trans R Soc Lond A* 209:415–446
48. Mumford D, Shah J (1989) Optimal approximations by piecewise smooth functions and associated variational problems. *Commun Pure Appl Math* 42:577–685
49. Nain D, Yezzi A, Turk G (2003) Vessel segmentation using a shape driven flow. In: *MICCAI*, pp 51–59
50. Nieuwenhuis C, Cremers D (2013) Spatially varying color distributions for interactive multi-label segmentation. *IEEE Trans Pattern Anal Mach Intell* 35(5):1234–1247
51. Osher SJ, Sethian JA (1988) Fronts propagation with curvature dependent speed: algorithms based on Hamilton–Jacobi formulations. *J Comp Phys* 79:12–49
52. Paragios N, Deriche R (2002) Geodesic active regions and level set methods for supervised texture segmentation. *Int J Comput Vis* 46(3):223–247
53. Parent P, Zucker SW (1989) Trace inference, curvature consistency, and curve detection. *IEEE Trans Pattern Anal Mach Intell* 11(8):823–839
54. Parzen E (1962) On the estimation of a probability density function and the mode. *Ann Math Stat* 33:1065–1076
55. Rochery M, Jermyn I, Zerubia J (2006) Higher order active contours. *Int J Comput Vis* 69(1):27–42
56. Rosenblatt F (1956) Remarks on some nonparametric estimates of a density function. *Ann Math Stat* 27:832–837
57. Rousson M, Brox T, Deriche R (2003) Active unsupervised texture segmentation on a diffusion based feature space. In: *Proc. IEEE conf. on comp. vision patt. recog*, Madison, WI, pp 699–704
58. Rousson M, Cremers D (2005) Efficient kernel density estimation of shape and intensity priors for level set segmentation. In: *MICCAI*, vol 1, pp 757–764
59. Rousson M, Paragios N, Deriche R (2004) Implicit active shape models for 3d segmentation in MRI imaging. In: *MICCAI*. LNCS, vol 2217. Springer, Berlin, pp 209–216
60. Rosenfeld A, Zucker SW, Hummel RA (1977) An application of relaxation labeling to line and curve enhancement. *IEEE Trans Comput* 26(4):394–403
61. Schmidt FR, Cremers D (2009) A closed-form solution for image sequence segmentation with dynamical shape priors. In: *Pattern recognition (Proc. DAGM)*, September 2009

62. Schmidt FR, Farin D, Cremers D (2007) Fast matching of planar shapes in sub-cubic runtime. In: IEEE int. conf. on computer vision, Rio de Janeiro, October 2007
63. Schoenemann T, Cremers D (2007) Globally optimal image segmentation with an elastic shape prior. In: IEEE int. conf. on computer vision, Rio de Janeiro, Brasil, October 2007
64. Schoenemann T, Cremers D (2007) Introducing curvature into globally optimal image segmentation: minimum ratio cycles on product graphs. In: IEEE int conf on computer vision, Rio de Janeiro, October 2007
65. Schoenemann T, Cremers D (2008) Matching non-rigidly deformable shapes across images: a globally optimal solution. In: IEEE conference on computer vision and pattern recognition (CVPR), Anchorage, Alaska, June 2008
66. Schoenemann T, Cremers D (2009) A combinatorial solution for model-based image segmentation and real-time tracking. *IEEE Trans Pattern Anal Mach Intell*
67. Schoenemann T, Kahl F, Masnou S, Cremers D (2012) A linear framework for region-based image segmentation and inpainting involving curvature penalization. *Int J Comput Vis* 99:53–68
68. Schoenemann T, Schmidt FR, Cremers D (2008) Image segmentation with elastic shape priors via global geodesics in product spaces. In: British machine vision conference, Leeds, UK, September 2008
69. Sebastian T, Klein P, Kimia B (2003) On aligning curves. *IEEE Trans Pattern Anal Mach Intell* 25(1):116–125
70. Serra J (1982) *Image analysis and mathematical morphology*. Academic Press, London
71. Tsai A, Wells W, Warfield SK, Willsky A (2004) Level set methods in an EM framework for shape classification and estimation. In: MICCAI
72. Tsai A, Yezzi A, Wells W, Tempany C, Tucker D, Fan A, Grimson E, Willsky A (2001) Model-based curve evolution technique for image segmentation. In: *Comp vision patt recog, Kauai, Hawaii*, pp 463–468
73. Tsai A, Yezzi AJ, Willsky AS (2001) Curve evolution implementation of the Mumford-Shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans Image Process* 10(8):1169–1186
74. Unal G, Krim H, Yezzi AY (2005) Information-theoretic active polygons for unsupervised texture segmentation. *Int J Comput Vis*, May
75. Unger M, Pock T, Cremers D, Bischof H (2008) TVSeg—interactive total variation based image segmentation. In: British machine vision conference (BMVC), Leeds, UK, September 2008
76. Zhu SC, Yuille A (1996) Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans Pattern Anal Mach Intell* 18(9):884–900