L. D. Griffin and M. Lillholm (Eds.), *4th Int. Conf. on Scale Space Theories in Computer Vision*, Isle of Skye, June 2003, LNCS Vol. 2695, 388-400 © Springer.

Towards Recognition-based Variational Segmentation Using Shape Priors and Dynamic Labeling

Daniel Cremers[1], Nir Sochen[2], Christoph Schnörr[3]

[1] Department of Computer Science
University of California at Los Angeles, USA
http://www.cs.ucla.edu/∼cremers

[2] Department of Applied Mathematics
Tel Aviv University, Israel
http://www.math.tau.ac.il/∼sochen

[3] Department of Mathematics and Computer Science
University of Mannheim, Germany
http://www.cvgpr.uni-mannheim.de

Abstract. We propose a novel variational approach based on a level set formulation of the Mumford-Shah functional and shape priors. We extend the functional by a labeling function which indicates image regions in which the shape prior is enforced. By minimizing the proposed functional with respect to both the level set function and the labeling function, the algorithm selects image regions where it is favorable to enforce the shape prior. By this, the approach permits to segment multiple independent objects in an image, and to discriminate familiar objects from unfamiliar ones by means of the labeling function. Numerical results demonstrate the performance of our approach.

1 Introduction

The problem of segmenting an image into its semantically significant components has been eluding researchers in computer vision for over 30 years. In the early days it was believed to be merely a technical problem. The traditional bottom-up approach, which is still very common in the computer vision community, suggests to start with a denoising process and to apply some sort of threshold afterward. Once the image is successfully segmented one can go to higher levels such as 3D reconstruction, motion analysis, classification and recognition.

It was realized by Mumford and Shah in the mid 80's, that denoising and segmentation are two different aspects of the same problem. Indeed a good denoising process should distinguish between a set of significant regions and the border between them. Such a denoising process assumes, implicitly, that the segmentation is known. On the other hand, segmentation approaches based on a threshold process assume that the image is denoised such that jumps in gray value can be attributed to boundaries between objects and are not the result of noise. This line of reasoning culminated in the Mumford-Shah functional [14], which puts the denoised image and the boundary contours on the same footing. Minimizing the functional simultaneously with respects to the dynamic variables results in a denoised image AND the boundaries. The minimization procedure involves the solution of coupled equations for the image and its boundaries.

Despite the success of the Mumford-Shah functional, its many variants [13] and the extensive mathematical analysis that followed its introduction, the problem of segmentation still eludes us. There is no segmentation algorithm that comes near the performance of a 3 year old child.

This failure indicates that a pure bottom-up approach is inappropriate. In fact we believe that higher-level processes related to recognition should participate in the segmentation process. This idea is reflected in the works of several researches. For a recent non-PDE approach starting from such viewpoint see [1].

In this paper, we combine both data-driven and recognition-driven processing in an unbiased way by introducing a dynamic labeling function into a variational segmentation approach with shape priors. In analogy to the reasoning of Mumford and Shah, minimization of the proposed functional with respect to the dynamic variables results in denoising, reconstruction of boundaries *and* the evolution of a decision function which models a higher-level recognition process.

The integration of shape priors into PDE based segmentation methods has been a focus of research in past years (c.f. [10, 21, 17, 6, 12, 20, 8, 7, 16]). Commonly, such approaches introduce a shape prior into the contour evolution in such a way that (ideally) the object of interest is reconstructed and all unfamiliar image structures are suppressed.

More recently, implicit level set based representations of a contour [15] have become a popular framework for image segmentation (cf. [2, 11, 4]). Since the topology of the evolving boundary is not constrained, one can elegantly model topological changes such as splitting and merging. This permits to segment multiply connected objects or several independent objects in a given image.

The question of how to introduce higher-level prior shape knowledge into level set based contour evolutions has been addressed by a number of people in recent years (cf. [12, 20, 5, 16]). In many of these approaches, the shape prior acts on the embedding level set function. As shown e.g. in [12], this approach permits to segment a multiply connected object with a statistical shape prior.

All of these approaches introduce the shape prior in such a way that only familiar structures of *one* given object can be recovered. They do not permit the simultaneous segmentation of several *independent* objects, comprising both familiar and unfamiliar ones. As an example, Figure 1, left side, shows a table scene containing two objects — the cover of a teapot, which is assumed to be familiar, and a pen assumed to be unfamiliar. The same scene is shown on the right, but the familiar object is corrupted: some parts are missing while others are occluded.

The aim of the present work is to introduce a shape prior into the Mumford-Shah functional, in a way which permits the simultaneous segmentation of several

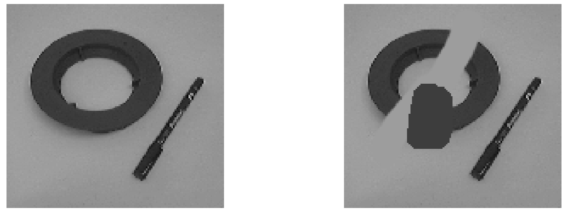

Fig. 1. Left: A table scene showing two objects — the cover of a tea pot which is assumed to be familiar, and a pen which is assumed to be unfamiliar.
Right: Corrupted version of the same image obtained by removing and occluding parts of the familiar object.

objects in one image (each of which may consist of several components). In particular, we will show that this approach permits to reconstruct the familiar object in Figure 1, right side, while leaving the correct segmentation of the unfamiliar object unaffected. To this end, we extend a level set formulation of the Mumford-Shah functional by a labeling function which indicates the regions of the image plane in which a given shape prior is enforced. By simultaneously minimizing the proposed energy functional with respect to the level set function and the labeling function, the labeling function dynamically separates regions of familiar objects from regions of unfamiliar ones.

The organization of this paper is as follows: In Section 2, we briefly review a level set formulation of the piecewise constant Mumford-Shah functional, as proposed in [3]. In Section 3, we augment this variational framework by a shape prior which affects the evolution of the level set function globally. As a consequence, such commonly proposed global shape priors suppress all unfamiliar image structures in the resulting segmentation process. In Section 4, we introduce a static labeling function to explicitly restrict the effect of the shape prior to designated areas of the image plane. This permits to both reconstruct a corrupted version of a known object and segment a novel unknown object in the same input image. In Section 5, we finally propose a variational framework with a dynamic labeling function. Compared to the case of a static labeling, the dynamic one evolves in an unsupervised manner during energy minimization. Numerical results show that this approach permits to segment new objects and reconstruct known ones without specifying the respective image areas beforehand. In Sections 7 and 8, we discuss some limitations and future work and end with a conclusion.

2 Region-based Segmentation with Level Sets

In this section, we will detail a level set method for image segmentation which aims at maximizing the gray value homogeneity in a set of disjoint regions. It is based on a level set formulation of the Mumford-Shah functional proposed by Chan and Vese [3]. This framework will then be extended by shape priors of increasing complexity in the subsequent sections.

2.1 A Level Set Framework for the Mumford-Shah Functional

Mumford and Shah [14] proposed to segment an input image $f : \Omega \to \mathbb{R}$ by minimizing the functional

$$E(u, C) = \frac{1}{2} \int_{\Omega} (f - u)^2 \, dx \; + \; \lambda^2 \, \frac{1}{2} \int_{\Omega - C} |\nabla u|^2 \, dx \; + \; \nu \, |C| \tag{1}$$

simultaneously with respect to the segmenting boundary C and the piecewise smooth approximation u. If the smoothness constraint is further stressed, one obtains for $\lambda \to \infty$ the cartoon limit (or minimal partition problem) in which the input image f is approximated by a piecewise constant segmentation $u = \{u_i\}$:

$$E(u, C) = \frac{1}{2} \sum_i \int_{\Omega_i} (f - u_i)^2 \, dx \; + \; \nu \, |C|. \tag{2}$$

During minimization of (2), the constants $\{u_i\}$ will take on the mean value of f over the set of disjoint regions $\{\Omega_i\}$ which partition the image plane ($\Omega = \bigcup \Omega_i$) and which are separated by the boundary C.

In several papers, Chan and Vese detailed a level set implementation of the Mumford-Shah functional (cf. [3, 4]), which is based on the use of the Heaviside function as an indicator function for the separate phases. The focus of the present work is the modeling of selective shape priors. Therefore, we will restrict ourselves to the case of the piecewise constant Mumford-Shah model and a single level set function $\phi : \Omega \to \mathbb{R}$ to embed the contour: $C = \{x \in \Omega \,|\, \phi(x) = 0\}$.

A piecewise constant segmentation of an input image f with two gray values u_+ and u_- can be obtained by minimizing the functional [3]:

$$E_{CV}(u_+, u_-, \phi) = \int_{\Omega} (f - u_+)^2 H(\phi) + (f - u_-)^2 \big(1 - H(\phi)\big) \, dx + \nu \int_{\Omega} |\nabla H(\phi)|, \tag{3}$$

with respect to the scalar variables u_+ and u_- and the embedding level set function ϕ. Here $H(\phi)$ denotes the Heaviside function:

$$H(\phi) = \begin{cases} 1, & \phi \geq 0 \\ 0, & \text{else} \end{cases} \tag{4}$$

The Euler-Lagrange equation for this functional can be implemented by the following gradient descent:

$$\frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\nu \, \text{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - u_+)^2 + (f - u_-)^2 \right], \tag{5}$$

where the scalars u_+ and u_- are updated in alternation with the level set evolution to take on the mean gray value of the input image f in the regions with $\phi > 0$ and $\phi < 0$, respectively:

$$u_+ = \frac{\int f(x) H(\phi) dx}{\int H(\phi) dx}, \qquad u_- = \frac{\int f(x)(1 - H(\phi)) dx}{\int (1 - H(\phi)) dx}. \tag{6}$$

The implementation in [3] is based on a smooth approximation of the delta function $\delta_\epsilon(s) = H'_\epsilon(s)$, which is chosen to have an infinite support:

$$\delta_\epsilon(s) = \frac{1}{\pi}\frac{\epsilon}{\epsilon^2 + s^2}. \tag{7}$$

In particular, a discretization with a support larger than zero permits the detection of interior contours – for example if one wants to segment a ring-like structure, starting from an initial contour located outside the ring.

2.2 Redistancing

During its evolution according to equation (5), the level set function ϕ generally grows to very large positive values in dark areas of the input image and very large negative values in bright areas of the image (or vice versa). At the zero crossings, it rises steeply, the gradient can become arbitrarily large. In numerical implementations, we found that a very steep slope of the level set function eventually inhibits the flexibility of the boundary to displace.

Many people have advocated the use of a redistancing procedure in the evolution of level set functions to constrain the slope of ϕ to $|\nabla\phi| = 1$, c.f. [9]. In order to reproject the evolving level set function to the space of distance functions, we intermittently iterate several steps of the redistancing equation [19]:

$$\frac{\partial\phi}{\partial t} = \operatorname{sign}(\hat{\phi})\ (1 - |\nabla\phi|)\,, \tag{8}$$

where $\hat{\phi}$ denotes the level set function before redistancing. As pointed out in [3], such a regularization is optional for the above level set model. Yet, we found this to have two favorable properties in our application. Firstly, it improves the convergence of the boundary evolution. And secondly, this normalization facilitates the introduction of shape priors which are encoded in terms of signed distance functions. We found this simple redistancing process to work well for our purpose, therefore we did not revert to more elaborate iterative redistancing schemes such as the one presented in [18].

2.3 Numerical Results

Minimization of the functional (3) is done by alternating the three fractional steps of iterating the gradient descent for the level set function ϕ, as given by equation (3), iterating the redistancing procedure given by equation (8) and updating the mean gray values for the two phases, as given in equation (6).

Figure 2 shows several steps of the evolution of the boundary C obtained by minimizing the functional (3) for the corrupted input image introduced in Figure 1. Due to the implicit representation, the boundary is free to perform splitting and merging. Due to the region-based formulation of the functional, the contour converges to the final segmentation over fairly large distances, while local edge and corner information is well preserved.

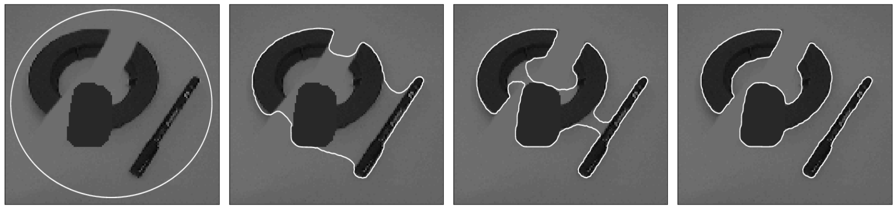

Fig. 2. Segmentation without shape prior. Evolution of the boundary for the Chan-Vese level set formulation of the piecewise constant Mumford-Shah functional (with a single level set function). The contour evolves so as to separate bright and dark areas. Due to the implicit level set representation, the topology is not constrained, which allows for splitting and merging of the boundary.

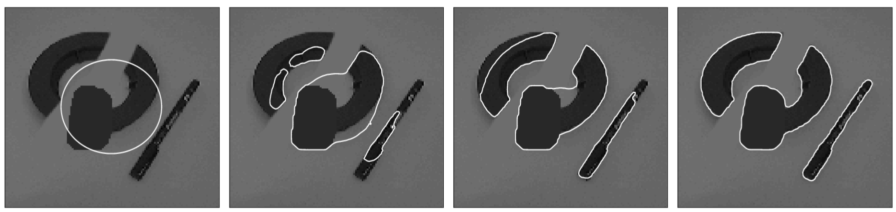

Fig. 3. Segmentation without shape prior. Evolution of the boundary for the Chan-Vese level set formulation with a different initialization as in Figure 2. There is no bias in form of a balloon term, therefore the contour can both expand and shrink for the same parameter value.

Yet, compared to many alternative approaches to level set segmentation, the above approach does not contain a balloon term which induces a bias favoring either contraction or expansion and therefore assumes prior knowledge about whether the objects of interest are inside or outside the initial contour. Figure 3 shows the corresponding contour evolution for a different initial contour. In this case, the contour expands from its initialization to converge to a similar segmentation of the given image.

3 Global Shape Prior in the Level Set Segmentation

In many applications of image segmentation, some knowledge about the shape of expected objects of interest is available. This prior shape information can be introduced into the level set functional in the following way. A number of training shapes is embedded by the signed distance function. The set of associated distance functions is aligned and a statistical model for the level set function is inferred from the training set. This prior is then added either to the evolution equation (5) or directly as a shape energy to the functional (3). Invariance of the prior with respect to similarity transformations of the level set function can be incorporated into these approaches (cf. [12, 16]). Since the focus of this work is selectivity of shape priors, we will not consider such invariances here. Moreover,

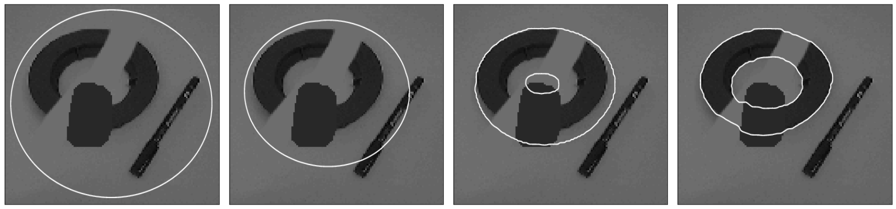

Fig. 4. Segmentation with a global shape prior. Evolution of the boundary for the Chan-Vese level set formulation with a shape prior which favors a ring. Compared to the corresponding segmentation without prior shown in Figure 2, the introduction of the global shape prior has several effects: The ring is reconstructed according to the shape prior, i.e. the partial occlusion is removed and the missing parts are filled in. At the same time, all other image structures – such as the pen on the right side – are removed.

we will only consider a single training shape, but we allow for more than one object in the image. Nevertheless, our approach can in principle be extended to more involved statistical shape priors [7].

A straight-forward extension of the functional (3) with an isotropic Gaussian shape prior is the following:

$$E(u_+, u_-, \phi) = E_{CV}(u_+, u_-, \phi) + \alpha\, E_{shape}(\phi), \tag{9}$$

with

$$E_{shape}(\phi) = \int_{\Omega} \big(\phi(x) - \phi_0(x)\big)^2\, dx, \tag{10}$$

where ϕ_0 is the level set function embedding a given training shape (or the mean of a set of training shapes) and $\alpha \geq 0$ determines the weight of the prior.

Minimizing this functional with respect to ϕ results in an evolution equation of the form:

$$\begin{aligned} \frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\nu \operatorname{div}\left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - u_+)^2 + (f - u_-)^2 \right] \\ - 2\,\alpha\, \big(\phi - \phi_0\big). \end{aligned} \tag{11}$$

Compared to the purely data-driven evolution in (5), we obtain an additional relaxation towards the learned shape ϕ_0.

Applied to the same image as in Figure 2, this generates the contour evolution shown in Figure 4. For sufficiently large weight of the shape prior, all image structures which are *not familiar* will be suppressed from the segmentation. In our example, the training shape consisted of the ring from a tea can. Due to the shape prior, missing parts of this object are recovered and occlusions removed. Moreover, the pen next to the ring is also removed.

4 Selective Shape Prior by Static Labeling

In the example of Figure 4, the shape prior permitted to reconstruct the learned object. In general, however, this object may not be present in a given image. Moreover, a given view of a scene may contain corrupted versions of a known object — the ring in our example — and other unfamiliar objects — in our example the pen next to it. In such cases, it may be desirable to have a *selective* shape prior, which permits to reconstruct the corrupted version of the known object, but which will not affect the segmentation of the unknown objects.

In this paper, we will model such a selective shape prior by a labeling function $L : \Omega \to \mathbb{R}$ which indicates the areas of the image plane in which a given prior should be enforced. This labeling function is to take on the values $+1$ and -1 depending on whether the prior should be enforced or not. For the beginning, we will assume this labeling to be known. This assumption will be removed in the following section.

We propose to segment an input image f by minimizing the functional (9) with a shape prior of the form:

$$E_{shape}(\phi) = \int_{\Omega} \big(\phi(x) - \phi_0(x)\big)^2 \, (L+1)^2 \, dx, \tag{12}$$

where a labeling L defines the parts of the image plane Ω where the shape prior should be active. The gradient descent equation for ϕ is given by:

$$\begin{aligned} \frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\nu \, \mathrm{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - u_+)^2 + (f - u_-)^2 \right] \\ - 2\,\alpha\,(L+1)^2 \, \big(\phi - \phi_0\big). \end{aligned} \tag{13}$$

Compared to the evolution equation with the global prior in (11), the additional relaxation towards the learned shape ϕ_0 is now restricted to the image areas where $L \neq -1$.

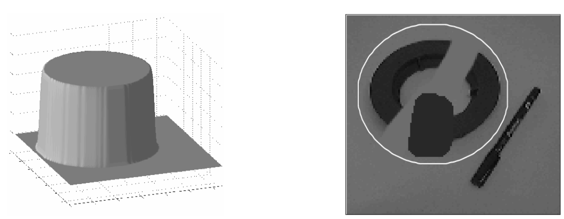

Fig. 5. Left: The labeling function L, shown in a 3D plot, indicates the area around the familiar object in which the prior should be applied ($L = 1$) and the remainder of the image plane, which should be segmented according to the gray value information only ($L = -1$). **Right:** Zero crossing of L superimposed on the input image.

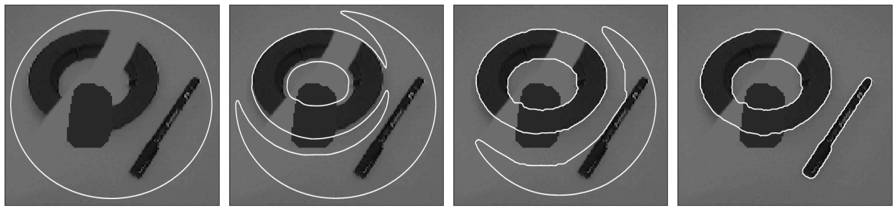

Fig. 6. Segmentation with static labeling function and shape prior. Evolution of the boundary upon minimizing the functional (12) with the static labeling function shown in Figure 5. According to the labeling function, the prior is enforced in an area around the ring only and masked out in the areas away from the ring. Therefore — in contrast to the result in Figure 4 — the ring is reconstructed according to the shape prior while the pen is also segmented as an independent object.

Figure 6 shows the result of minimizing this functional with respect to ϕ for the static labeling shown in Figure 5.

Due to the definition of the labeling, the shape prior is only enforced in a region around the ring. In the remainder of the image plane the labeling function masks out the shape prior. The consequence is that the known object is reconstructed according to the prior, while the correct segmentation of the pen is unaffected by the prior. It remains to be shown how this labeling function L can be determined *from the image data* as well.

5 Selective Shape Prior by Dynamic Labeling

In the previous section, we introduced a labeling function L to indicate regions of the image plane where a given shape prior should be enforced. Yet, this labeling was specified beforehand.

In the present section, we will overcome this limitation by an approach with a *dynamic* labeling function. To this end, we propose to minimize the functional

$$E(u_+, u_-, \phi, L) = E_{CV}(u_+, u_-, \phi) + \alpha \, E_{shape}(\phi), \tag{14}$$

with the shape prior:

$$E_{shape}(\phi, L) = \int (\phi - \phi_0)^2 \, (L{+}1)^2 \, dx + \int \lambda^2 \, (L{-}1)^2 \, dx + \gamma \int |\nabla H(L)| \, dx. \tag{15}$$

Compared to the previous approach of a static labeling function, we now assume the labeling L to be unknown. Instead of specifying the labeling, we simultaneously optimize the above cost functional with respect to both the segmenting level set function ϕ and the labeling function L.

According to the proposed cost functional, the labeling will evolve in an unsupervised manner driven by two criteria:

– The labeling should enforce the shape prior in those areas of the image where the level set function is similar to the prior. In particular, for fixed ϕ, minimizing the first two terms in (15) will generate the following qualitative behavior of the labeling:

$$\begin{array}{ll} L \to 1, & \text{if } |\phi - \phi_0| < \lambda \\ L \to -1, & \text{if } |\phi - \phi_0| > \lambda \end{array}$$

– The boundary separating regions with shape prior from regions without shape prior should have minimal length. This regularizing constraint on the zero crossing of the labeling function – given by the last term in equation (15) — induces a topological “compactness” of the regions with and without shape prior.

6 Evolution of Labeling Function and Level Set Function

We simultaneously minimize the functional (14) with respect to both the labeling function L and the level set function ϕ by iterating the associated gradient descent equations in alternation with updates of the mean gray values u_+ and u_- according to (6).

For fixed ϕ, the gradient descent equation for the labeling function is given by:

$$\frac{\partial L}{\partial t} = -\frac{\partial E}{\partial L} = \alpha \left[2\lambda^2 (1 - L) - 2(\phi - \phi_0)^2 (1 + L) + \gamma \delta_\epsilon(L) \,\mathrm{div}\left(\frac{\nabla L}{|\nabla L|} \right) \right]. \quad (16)$$

The first two terms drive the labeling toward -1 or 1, depending on whether $|\phi - \phi_0|$ is larger or smaller than λ. And the last term in (16) minimizes the length of the zero-crossing of L, thereby enforcing decision regions with minimal boundary.

Conversely, for fixed labeling, the gradient descent equation for the level set function ϕ is given by:

$$\begin{aligned} \frac{\partial \phi}{\partial t} &= -\frac{\partial E}{\partial \phi} \\ &= \delta_\epsilon(\phi) \left[\nu \,\mathrm{div}\left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - u_+)^2 + (f - u_-)^2 \right] \\ &\quad - 2\,\alpha\,(1 + L)^2\,(\phi - \phi_0). \end{aligned} \quad (17)$$

Compared to the purely data-driven evolution in equation (5), we have an additional relaxation toward the learned shape ϕ_0 in all areas of the image where $L \neq -1$.

The resulting evolution of the labeling function L, obtained by minimizing the total energy (14) for the same input image as before, is shown in Figure 7, bottom row: The area around the ring arises in an unsupervised manner as the region where the prior should be applied.

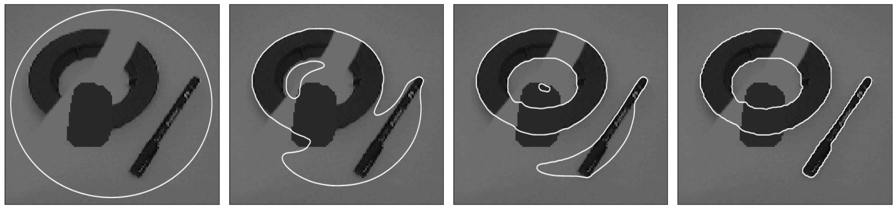

Evolution of the boundary

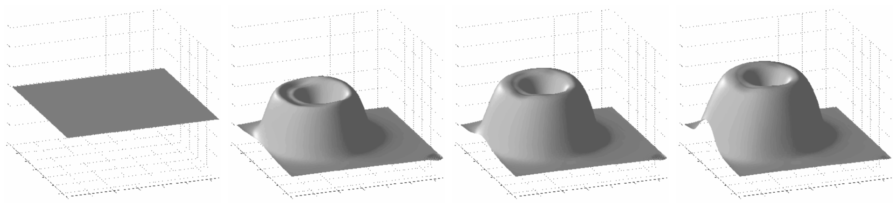

Evolution of the labeling

Fig. 7. Segmentation with dynamic labeling and shape prior. Top row: Evolution of the segmenting boundary obtained by minimizing the functional (14) for the same input image as in Figure 2. The ring object is reconstructed according to the shape prior, whereas the correct segmentation of the pen is unaffected by the prior. The prior affects the embedding level set function ϕ only in the regions indicated by the simultaneously evolving dynamic labeling function. **Bottom row:** Evolution of the dynamic labeling L during energy minimization. The labeling function dynamically selects the region around the ring object to apply the shape prior. Compared to the static labeling shown in Figure 6, the dynamic labeling evolves in an unsupervised manner.

The corresponding evolution of the segmenting boundary given by the zero level set of the function ϕ is shown in Figure 7, top row. These results demonstrate the following favorable properties of our approach:

- Compared to the segmentation without shape prior shown in Figure 2, the ring is reconstructed according to the shape prior, i.e. the missing parts are filled in and the occlusion is removed.
- Compared to the segmentation with a *global* shape prior shown in Figure 4, the correct segmentation of the pen is unaffected by the shape prior. The effect of the shape prior is restricted to the area around the familiar object by the labeling function.
- Compared to the case of a *static* labeling function shown in Figures 5 and 6, one no longer needs to specify beforehand the regions where the prior should be enforced. Instead, the labeling evolves in an unsupervised manner during the minimization of the functional (14). It dynamically selects image regions in which the prior is to be applied.

7 Limitations and Future Work

Segmentation results obtained with the dynamic labeling approach depend on an appropriate choice of the parameters λ and γ which affect the evolution of the labeling. We are currently working on a more rigorous probabilistic formulation of the dynamic labeling approach which, we believe, should result in automatic estimates of these parameters from the data. Moreover, we intend to generalize the proposed approach to *multiple* selective shape priors, and to more elaborate statistical shape priors such as the one introduced in [7]. We will also investigate methods to introduce pose invariance, such as the one presented e.g. in [12, 16], into the proposed approach.

8 Conclusion

We presented a novel variational approach to integrate higher-level shape priors into level set based segmentation methods. In particular, we addressed the problem of applying shape priors selectively, such that a given prior permits the reconstruction of corrupted versions of a familiar object while not affecting the correct segmentation of independent unknown objects.

To this end, we extended the level set approach of Chan and Vese by three shape priors of increasing complexity. The first one is a simple globally active prior which permits the reconstruction of a known object but removes all unfamiliar objects. The second one is a shape prior with a static labeling, which allows to define areas of the input image in which the prior should be applied. This prior permits to reconstruct the known object in the selected area, but does not affect the correct segmentation of independent objects in the remainder of the image plane. Finally, the third shape prior is based on a dynamic labeling function. The latter is not specified beforehand, it rather evolves in an unsupervised manner in order to automatically select the image regions to which the prior should be applied. As a consequence, the familiar object is reconstructed, yet independent unfamiliar objects are correctly segmented. And the decision *which* areas correspond to familiar objects simultaneously evolves with the segmentation during minimization of the proposed functional.

We believe that the results presented in this work demonstrate the capacity of the dynamic labeling approach to model an *unsupervised decision process*. Such a decision process fundamentally extends the applicability of statistical shape priors in segmentation. It permits to combine both data-driven and recognition-driven processing on equal footings in a variational segmentation approach.

Acknowledgements

DC was supported by ONR N00014-02-1-0720 and AFOSR F49620-03-1-0095.

References

1. E. Borenstein and S. Ullman. Class-specific, top-down segmentation. In A. Heyden et al., editors, *Proc. of the Europ. Conf. on Comp. Vis.*, volume 2351 of *LNCS*, pages 109–122, Copenhagen, May 2002. Springer, Berlin.
2. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. IEEE Internat. Conf. on Comp. Vis.*, pages 694–699, Boston, USA, 1995.
3. T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.
4. T. Chan and L. Vese. A level set algorithm for minimizing the Mumford–Shah functional in image processing. In *IEEE Workshop on Variational and Level Set Methods*, pages 161–168, Vancouver, CA, 2001.
5. Y. Chen, S. Thiruvenkadam, H. Tagare, F. Huang, D. Wilson, and E. Geiser. On the incorporation of shape priors into geometric active contours. In *IEEE Workshop on Variational and Level Set Methods*, pages 145–152, Vancouver, CA, 2001.
6. T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. Use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–365, 1994.
7. D. Cremers, T. Kohlberger, and C. Schnörr. Nonlinear shape statistics in Mumford–Shah based segmentation. In A. Heyden et al., editors, *Proc. of the Europ. Conf. on Comp. Vis.*, volume 2351 of *LNCS*, pages 93–108, Copenhagen, May 2002. Springer, Berlin.
8. D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr. Diffusion Snakes: Introducing statistical shape knowledge into the Mumford–Shah functional. *Int. J. of Comp. Vis.*, 50(3):295–313, 2002.
9. J. Gomes and O. D. Faugeras. Level sets and distance functions. In D. Vernon, editor, *Proc. of the Europ. Conf. on Comp. Vis.*, volume 1842 of *LNCS*, pages 588–602, Dublin, Ireland, 2000. Springer.
10. U. Grenander, Y. Chow, and D.M. Keenan. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer, New York, 1991.
11. S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. J. Yezzi. Gradient flows and geometric active contour models. In *Proc. IEEE Internat. Conf. on Comp. Vis.*, pages 810–815, Boston, USA, 1995.
12. M. E. Leventon, W. E. L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proc. Conf. Computer Vis. and Pattern Recog.*, volume 1, pages 316–323, Hilton Head Island, SC, June 13–15, 2000.
13. J.-M. Morel and S. Solimini. *Variational Methods in Image Segmentation*. Birkhäuser, Boston, 1995.
14. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.
15. S. J. Osher and J. A. Sethian. Fronts propagation with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *J. of Comp. Phys.*, 79:12–49, 1988.
16. M. Rousson and N. Paragios. Shape priors for level set representations. In A. Heyden et al., editors, *Proc. of the Europ. Conf. on Comp. Vis.*, volume 2351 of *LNCS*, pages 78–92, Copenhagen, May 2002. Springer, Berlin.
17. L. H. Staib and J. S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 14(11):1061–1075, 1992.
18. M. Sussman and E. Fatemi. An efficient, interface-preserving level set redistancing algorithm and its application to interfacial incompressible fluid flow. *SIAM J. Sci. Comput.*, 20(4):1165–1191, 1999.

19. M. Sussman, Smereka P., and S. J. Osher. A level set approach for computing solutions to incompressible twophase flow. *J. of Comp. Phys.*, 94:146–159, 1994.
20. A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, E. Grimson, and A.. Willsky. Model–based curve evolution technique for image segmentation. In *Conf. on Comp. Vision Patt. Recog.*, pages 463–468, Kauai, Hawaii, 2001.
21. A. Yuille and P. Hallinan. Deformable templates. In A. Blake and A. Yuille, editors, *Active Vision*, pages 21–38. MIT Press, 1992.