Multiphase Dynamic Labeling for Variational Recognition-Driven Image Segmentation

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

[1] Department of Computer Science
University of California, Los Angeles, USA
[2] Department of Applied Mathematics
Tel Aviv University, Israel
[3] Department of Mathematics and Computer Science
University of Mannheim, Germany

Abstract. We propose a variational framework for the integration multiple competing shape priors into level set based segmentation schemes. By optimizing an appropriate cost functional with respect to both a level set function and a (vector-valued) labeling function, we jointly generate a segmentation (by the level set function) and a recognition-driven partition of the image domain (by the labeling function) which indicates where to enforce certain shape priors. Our framework fundamentally extends previous work on shape priors in level set segmentation by directly addressing the central question of *where* to apply *which* prior. It allows for the seamless integration of numerous shape priors such that – while segmenting both multiple known and unknown objects – the level set process may selectively use specific shape knowledge for simultaneously enhancing segmentation and recognizing shape.

1 Introduction

Image segmentation and object recognition in vision are driven both by low-level cues such as intensities, color or texture properties, and by prior knowledge about objects in our environment. Modeling the interaction between such data-driven and model-based processes has become the focus of current research on image segmentation in the field of computer vision. In this work, we consider prior knowledge given by the shapes associated with a set of familiar objects and focus on the problem of how to exploit such knowledge for images containing multiple objects, some of which may be familiar, while others may be unfamiliar.

Following their introduction as a means of front propagation [13], level set based contour representations have become a popular framework for image segmentation [1,10]. They permit to elegantly model topological changes of the implicitly represented boundary, which makes them well suited for segmenting images containing *multiple* objects. Level set segmentation schemes can be formulated to exploit various low level cues such as edge information [10,2,8], intensity homogeneity [3,18], texture [14] or motion information [6]. In recent years, there has been much effort in trying to integrate prior shape knowledge into level set based segmentation. This was shown to make the segmentation

T. Pajdla and J. Matas (Eds.): ECCV 2004, LNCS 3024, pp. 74–86, 2004.

process robust to misleading low-level information caused by noise, background clutter or partial occlusion of an object of interest (cf. [9,17,5,15]).

A key problem in this context is to ensure that prior knowledge is *selectively* applied at image locations only where image data indicate a familiar object. Conversely, lack of any evidence for the presence of some familiar object should result in a purely data-driven segmentation process. To this end, it was recently proposed to introduce a labeling function in order to restrict the effect of a given prior to a specific domain of the image plane [7] (for a use of a labeling field in a different context see [11]). During optimization, this labeling function evolves so as to select image regions where the given prior is applied. The resulting process segments corrupted versions of a known object in a way that does not affect the correct segmentation of other unfamiliar objects. A smoothness constraint on the labeling function induces the process to distinguish between occlusions (which are close to the familiar object) and separate independent objects (assumed to be sufficiently far from the object of interest).

All of the approaches mentioned above were designed to segment a *single* known object in a given image. But what if there are several known objects? Clearly, any use of shape priors consistent with the philosophy of the level set method should retain the capacity of the resulting segmentation scheme to deal with multiple independent objects, no matter whether they are familiar or not. One may instead suggest to iteratively apply the segmentation scheme with a different prior at each time and thereby successively segment the respective objects. We believe, however, that such a sequential processing mode will not scale up to large databases of objects and that – even more importantly – the parallel use of *competing priors* is essential for modeling the chicken-egg relationship between segmentation and recognition.

In this paper, we adopt the selective shape prior approach suggested in [7] and substantially generalize it along several directions:

- We extend the shape prior by pose parameters. The resulting segmentation process not only selects appropriate regions where to apply the prior, it also selects appropriate pose parameters associated with a given prior. This drastically increases the usefulness of this method for realistic segmentation problems, as one cannot expect to know the pose of the object beforehand.
- We extend the previous approach which allowed one known shape in a scene of otherwise unfamiliar shapes to one which allows two different known shapes. Rather than treating the second shape as background, the segmentation scheme is capable of reconstructing both known objects.
- Finally we treat the general case of an arbitrary number of known and unknown shapes by replacing the scalar-valued labeling by a vector-valued function. The latter permits to characterize up to 2^n regions with different priors, where n is the dimension of the labeling function. In particular, we demonstrate that – through a process of competing priors – the resulting segmentation scheme permits to simultaneously reconstruct three known objects while not affecting the segmentation of separate unknown objects.

In this work, the term *shape prior* refers to fixed templates with variable 2D pose. However, the proposed framework of selective shape priors is easily extended to statistical shape models which would additionally allow certain deformation modes of each template. For promising advances regarding level set based statistical shape representations, we refer to [4].

The outline of the paper is as follows: In Section 2, we briefly review the level set formulation of the piecewise constant Mumford-Shah functional proposed in [3]. In Section 3, we augment this variational framework by a labeling function which selectively imposes a given shape prior in a certain image region. In Section 4, we enhance this prior by explicit pose parameters and demonstrate the effect of simultaneous pose optimization. In Section 5, we extend the labeling approach from the case of one known object and background to that of two independent known objects. In Section 6, we come to the central contribution of this work, namely the generalization to an arbitrary number of known and unknown objects by means of a vector-valued labeling function. We demonstrate that the resulting segmentation scheme is capable of reconstructing corrupted versions of multiple known objects displayed in a scene containing other unknown objects.

2 Data-Driven Level Set Segmentation

Level set representations of moving interfaces, introduced by Osher and Sethian [13], have become a popular framework for image segmentation. A contour C is represented as the zero level set of an embedding function $\phi : \Omega \to \mathbb{R}$ on the image domain $\Omega \subset \mathbb{R}^2$:

$$C = \{x \in \Omega \mid \phi(x) = 0\}. \tag{1}$$

During the segmentation process, this contour is propagated implicitly by evolving the embedding function ϕ. In contrast to explicit parameterizations, one avoids the issues of control point regridding. Moreover, the implicitly represented contour can undergo topological changes such as splitting and merging during the evolution of the embedding function. This makes the level set formalism well suited for the segmentation of *multiple* objects. In this work, we will revert to a region-based level set scheme introduced by Chan and Vese [3]. However, other data-driven level set schemes could be employed.

In [3] Chan and Vese introduce a level set formulation of the piecewise constant Mumford-Shah functional [12]. In particular, they propose to generate a segmentation of an input image f with two gray values μ_1 and μ_2 by minimizing the functional

$$E_{CV}(\mu_1, \mu_2, \phi) = \int_\Omega (f - \mu_1)^2 H(\phi) + (f - \mu_2)^2 \big(1 - H(\phi)\big)\, dx \, + \, \nu \int_\Omega |\nabla H(\phi)|, \tag{2}$$

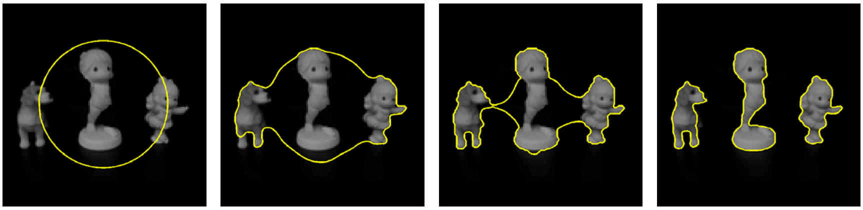

Fig. 1. Purely intensity-based segmentation. Contour evolution generated by minimizing the Chan-Vese model (2) [3]. The central figure is partially corrupted.

with respect to the scalar variables μ_1 and μ_2 and the embedding level set function ϕ. Here H denotes the Heaviside function

$$H(\phi) = \begin{cases} 1, & \phi \geq 0 \\ 0, & \text{else} \end{cases} . \tag{3}$$

The last term in (2) measures the length of the zero-crossing of ϕ.

The Euler-Lagrange equation for this functional is implemented by gradient descent:

$$\frac{\partial \phi}{\partial t} = \delta(\phi) \left[\nu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - \mu_1)^2 + (f - \mu_2)^2 \right], \tag{4}$$

where μ_1 and μ_2 are updated in alternation with the level set evolution to take on the mean gray value of the input image f in the regions defined by $\phi > 0$ and $\phi < 0$, respectively:

$$\mu_1 = \frac{\int f(x) H(\phi) dx}{\int H(\phi) dx}, \quad \mu_2 = \frac{\int f(x)(1 - H(\phi)) dx}{\int (1 - H(\phi)) dx}. \tag{5}$$

Figure 1 shows a representative contour evolution obtained for an image containing three figures, the middle one being partially corrupted.

3 Selective Shape Priors by Dynamic Labeling

The evolution in Figure 1 demonstrates the well-known fact that the level set based segmentation process can cope with multiple objects in a given scene. However, if the low-level segmentation criterion is violated due to noise, background clutter or partial occlusion of the objects of interest, then the purely image-based segmentation scheme will fail to converge to the desired segmentation.

To cope with such degraded low-level information, it was proposed to introduce prior shape knowledge into the level set scheme (cf. [9,17,15]). The basic idea is to extend the image-based cost functional by a shape energy which favors certain contour formations:

$$E_{total}(\phi) = E_{CV}(\mu_1, \mu_2, \phi) + \alpha E_{shape}(\phi) \qquad (\alpha > 0). \tag{6}$$

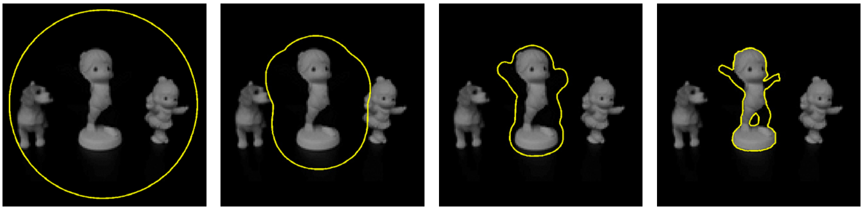

Fig. 2. Global shape prior. Contour evolution generated by minimizing the total energy (6) with a global shape prior of the form (7) encoding the figure in the center. Due to the global constraint on the embedding function, the familiar object is reconstructed while all unfamiliar structures are suppressed in the final segmentation. The resulting segmentation scheme lost its capacity to deal with multiple independent objects.

In general, the proposed shape constraints affect the embedding surface ϕ globally (i.e. on the entire domain Ω). In the simplest case, such a prior has the form:

$$E_{shape}(\phi) = \int_{\Omega} \left(\phi(x) - \phi_0(x)\right)^2 \, dx, \tag{7}$$

where ϕ_0 is the level set function embedding a given training shape (or the mean of a set of training shapes). Uniqueness of the embedding function associated with a given shape is guaranteed by imposing ϕ_0 to be a signed distance function (cf. [9]). For consistency, we also project the segmenting level set function ϕ to the space of distance functions during the optimization [16].

Figure 2 shows several steps in the contour evolution with such a prior, where ϕ_0 is the level set function associated with the middle figure. The shape prior permits to reconstruct the object of interest, yet in the process, all unfamiliar objects are suppressed from the segmentation. The segmentation process with shape prior obviously lost its capacity to handle multiple (independent) objects.

In order to retain this favorable property of the level set method, it was proposed in [7] to introduce a labeling function $L : \Omega \to \mathbb{R}$, which indicates the regions of the image where a given prior is to be enforced. During optimization of an appropriate cost functional, the labeling evolves dynamically in order to select these regions in a recognition-driven way. The corresponding shape energy is given by:

$$E_{shape}(\phi, L) = \int (\phi - \phi_0)^2 \, (L+1)^2 \, dx + \int \lambda^2 \, (L-1)^2 \, dx \, + \, \gamma \int |\nabla H(L)| \, dx, \tag{8}$$

with two parameters $\lambda, \gamma > 0$. The labeling L enforces the shape prior in those areas of the image where the level set function is similar to the prior (associated with labeling $L = 1$). In particular, for fixed ϕ, minimizing the first two terms in (8) induces the following qualitative behavior of the labeling:

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

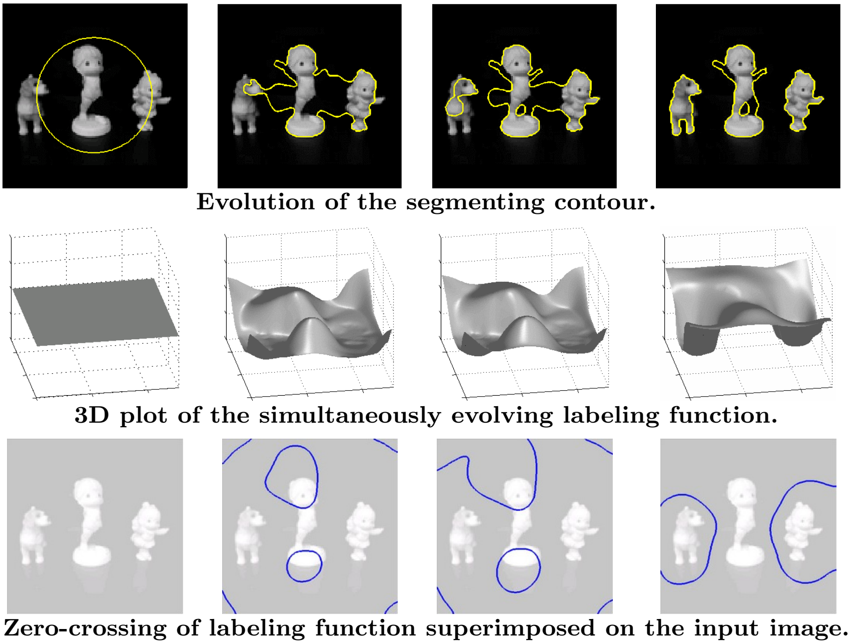

Fig. 3. Selective shape prior by dynamic labeling. Contour evolution generated by minimizing the total energy (6) with a selective shape prior of the form (8) encoding the figure in the center. Due to the simultaneous optimization of a labeling function $L(x)$ (middle and bottom row), the shape prior is restricted to act only in selected areas. The familiar shape is reconstructed, while the correct segmentation of separate (unfamiliar) objects remains unaffected. The resulting segmentation scheme thereby retains its capacity to deal with multiple independent objects. In this and all subsequent examples, labeling functions are initialized by $L \equiv 0$.

In addition, the last term in equation (8) imposes a regularizing constraint on the length of the zero crossing of the labeling, this induces topological "compactness" of both the regions with and without shape prior.

Figure 3 shows the contour evolution generated with the prior (8), where ϕ_0 encodes the middle figure as before. Again the shape prior permits to reconstruct the corrupted figure. In contrast to the global prior (7) in Figure 2, however, the process dynamically selects the region where to impose the prior. Consequently the correct segmentation of the two unknown objects is unaffected by the prior.

4 A Pose-Invariant Formulation

In the above formalism of dynamic labeling, the pose of the object of interest is assumed to be known. In a realistic segmentation problem, one generally does not know the pose of an object of interest. If the object of interest is no longer in the same location as the prior ϕ_0, the labeling approach will fail to generate the desired segmentation. This is demonstrated in Figure 4. While the labeling

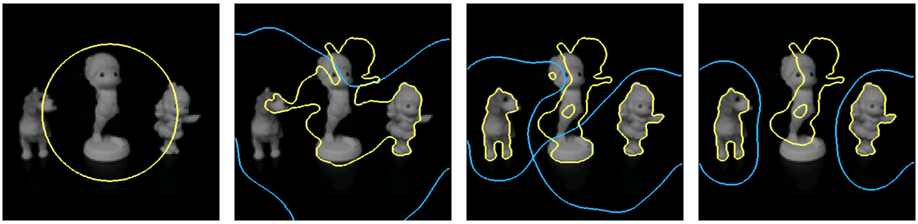

Fig. 4. Missing pose optimization. Evolution of contour (yellow) and labeling (blue) with selective shape prior (8) and a displaced template ϕ_0. Without simultaneous pose optimization, the familiar shape is forced to appear in the displaced position.

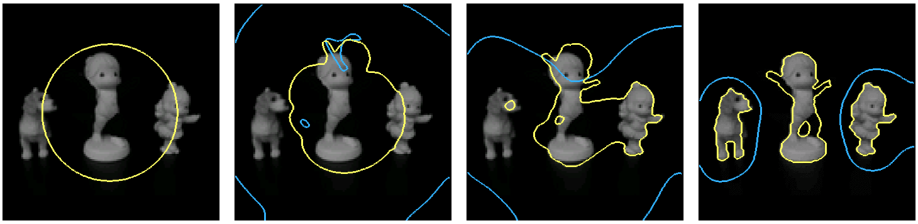

Fig. 5. Effect of pose optimization. By simultaneously optimizing a set of pose parameters in the shape energy (9), one jointly solves the problems of estimating the area where to impose a prior and the pose of the respective prior. Note that the pose estimate is gradually improved during the energy minimization.

still separates areas of known objects from areas of unknown objects, the known shape is not reconstructed correctly, since the pose of the prior and that of the object in the image differ.

A possible solution is to introduce a set of pose parameters associated with a given prior ϕ_0 (cf. [15,5]). The corresponding shape energy

$$E_{shape}(\phi, L, s, \theta, h) = \int \left(\phi(x) - \frac{1}{s}\phi_0(sR_\theta x + h) \right)^2 (L+1)^2 \, dx + \int \lambda^2 \, (L-1)^2 \, dx + \gamma \int |\nabla H(L)| \, dx \qquad (9)$$

is simultaneously optimized with respect to the segmenting level set function ϕ, the labeling function L and the pose parameters, which account for translation h, rotation by an angle θ and scaling s of the template. The normalization by s guarantees that the resulting shape remains a distance function.

Figure 5 shows the resulting segmentation: Again the labeling selects the regions where to apply the given prior, but now the simultaneous pose optimization also allows to estimate the pose of the object of interest.

The main focus of the present paper is to propose selective shape priors. For the sake of simplifying the exposition, we will therefore assume in the following, that the correct pose of familiar objects is known. Moreover, we will drop pose parameters associated with each shape template from the equations, so as to simplify the notation. We want to stress, however, that similar pose invariance can be demonstrated for all of the following generalizations.

5 Extension to Two Known Objects

A serious limitation of the labeling approach in (8) is that it only allows for a *single* known object (and multiple unknown objects). What if there are several familiar objects in the scene? How can one integrate prior knowledge about multiple shapes such as those given by a database of known objects? Before considering the general case, let us first study the case of *two* known objects.

The following modification of (8) allows for two different familiar objects associated with embedding functions ϕ_1 and ϕ_2:

$$E_{shape}(\phi, L) = \frac{1}{\sigma_1^2} \int (\phi - \phi_1)^2 (L+1)^2 \, dx + \frac{1}{\sigma_2^2} \int (\phi - \phi_2)^2 (L-1)^2 \, dx + \gamma \int |\nabla H(L)| \, dx. \quad (10)$$

The terms associated with the two objects were normalized with respect to the variance of the respective template: $\sigma_i^2 = \int \phi_i^2 dx - (\int \phi_i dx)^2$. The resulting shape prior has therefore merely one (instead of two) free parameters. The evolution of the labeling function is now driven by two competing shape priors: each image location will be ascribed to one or the other prior.

Figure 6 shows a comparison: The upper row indicates the contour evolution generated with the shape energy (8), where ϕ_0 encodes the figure on the left. The lower row shows the respective evolution obtained with the shape energy 10, with ϕ_1 and ϕ_2 encoding the left and right figures, respectively. Whereas the object on the right (occluded by a pen) is treated as unknown in the original formulation (upper row), both figures can be reconstructed by simultaneously imposing two competing priors in different domains (lower row).

6 The General Case: Multiphase Dynamic Labeling

The above example showed that the dynamic labeling approach can be transformed to allow for two shape priors rather than a single shape prior and possible background.

Let us now consider the general case of a larger number of known objects and possibly some further independent unknown objects (which should therefore be segmented based on their intensity only). To this end, we introduce a vector-valued labeling function

$$\boldsymbol{L} : \Omega \to \mathbb{R}^n, \qquad \boldsymbol{L}(x) = (L_1(x), \dots, L_n(x)). \quad (11)$$

We employ the $m = 2^n$ vertices of the polytope $[-1, +1]^n$ to encode m different regions, $L_j \in \{+1, -1\}$, and denote by $\chi_i, i = 1, \dots, m$ the indicator function for each of these regions. See [19] for a related concept in the context of multi-region

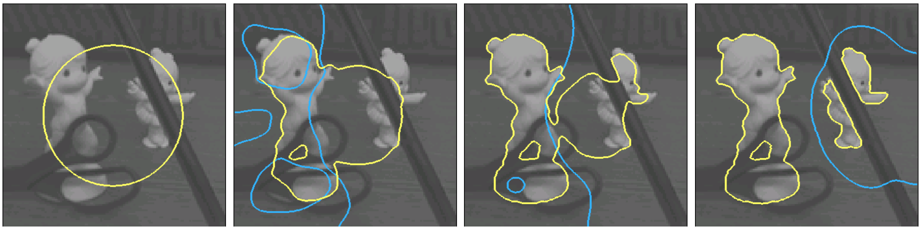

Dynamic Labeling with a single prior and background.

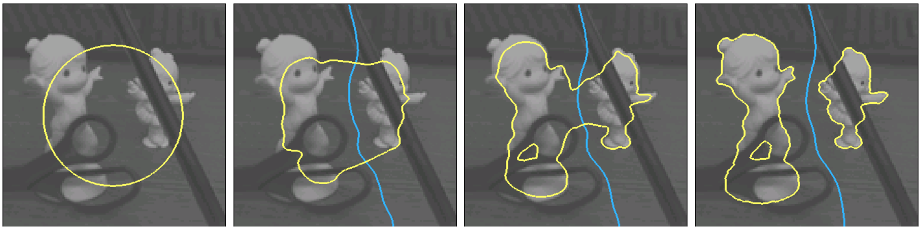

Dynamic Labeling allowing for two competing priors.

Fig. 6. Extension to two priors. Evolutions of contour (**yellow**) and labeling (**blue**) generated by minimizing energy (6) with a selective prior of the form (8) encoding the left figure (**top**) and with a selective prior of the form (10) encoding both figures (**bottom**). In both cases, the left figure is correctly reconstructed despite prominent occlusions by the scissors. However, while the structure on the right is treated as unfamiliar and thereby segmented based on intensities only (top row), the extension to two priors permits to simultaneously reconstruct both known objects (bottom row).

segmentation. For example, for $n = 2$, four regions are modeled by the indicator functions:

$$\chi_1(\boldsymbol{L}) = \tfrac{1}{16}(L_1 - 1)^2\,(L_2 - 1)^2, \qquad \chi_2(\boldsymbol{L}) = \tfrac{1}{16}(L_1 + 1)^2\,(L_2 - 1)^2,$$
$$\chi_3(\boldsymbol{L}) = \tfrac{1}{16}(L_1 - 1)^2\,(L_2 + 1)^2, \qquad \chi_4(\boldsymbol{L}) = \tfrac{1}{16}(L_1 + 1)^2\,(L_2 + 1)^2.$$

In the general case of an n-dimensional labeling function, each indicator function will be of the form

$$\chi_i(\boldsymbol{L}) \equiv \chi_{l_1 \dots l_n}(\boldsymbol{L}) = \frac{1}{4^n}\prod_{j=1}^{n}(L_j + l_j)^2, \quad \text{with} \;\; l_j \in \{+1, -1\}. \tag{12}$$

With this notation, the extension of the dynamic labeling approach to up to $m = 2^n$ regions can be cast into a cost functional of the form:

$$E_{total}(\phi, \boldsymbol{L}, \mu_1, \mu_2) = E_{CV}(\phi, \mu_1, \mu_2) + \alpha E_{shape}(\phi, \boldsymbol{L}), \tag{13}$$

$$E_{shape}(\phi, \boldsymbol{L}) = \sum_{i=1}^{m-1} \int \frac{(\phi - \phi_i)^2}{\sigma_i^2}\chi_i(\boldsymbol{L})dx + \int \lambda^2 \chi_m(\boldsymbol{L})dx + \gamma \sum_{i=1}^{m} \int |\nabla H(L_i)|dx.$$

Here, each ϕ_i corresponds to a particular known shape with its variance given by σ_i.

As mentioned before, we have – for better readability – neglected the pose parameters associated with each template. These can be incorporated by the replacements:

$$\phi_i \longrightarrow \frac{1}{s_i}\phi_i(s_i R_{\theta_i} x + h_i) \quad \text{and} \quad E_{shape}(\phi, \boldsymbol{L}) \longrightarrow E_{shape}(\phi, \boldsymbol{L}, \boldsymbol{p}),$$

where $\boldsymbol{p} = (p_1, \ldots, p_m)$ denotes the vector of pose parameters $p_i = (s_i, \theta_i, h_i)$ associated with each known shape.

7 Energy Minimization

In the previous sections, we have introduced variational formulations of increasing complexity to tackle the problem of multi-object segmentation with shape priors. The corresponding segmentation processes are generated by minimizing these functionals. In this section, we will detail the minimization scheme in order to illuminate how the different components of the proposed cost functionals affect the segmentation process. Let us focus on the case of multiple labels corresponding to the cost functional (13). Minimization of this functional is obtained by alternating the update of the mean intensities μ_1 and μ_2 according to (5) with a gradient descent evolution for the level set function ϕ, the labeling functions L_j and the associated pose parameters p_j. In the following, we will detail this for ϕ and L_j. Respective evolution equations for p_j are straight forward and not our central focus.

For fixed labeling, the evolution of the level set function ϕ is given by:

$$\frac{\partial \phi}{\partial t} = -\frac{\partial E_{total}}{\partial \phi} = -\frac{\partial E_{CV}}{\partial \phi} - 2\,\alpha \sum_{i=1}^{m-1} \frac{\phi - \phi_i}{\sigma_i^2} \chi_i(\boldsymbol{L}). \tag{14}$$

Apart from the image-driven first component given by the Chan-Vese evolution in equation (4), we additionally have a relaxation toward the template ϕ_i in all image locations where $\chi_i = 1$.

Minimization by gradient descent with respect to the labeling functions L_j corresponds to an evolution of the form:

$$\frac{1}{\alpha}\frac{\partial L_j}{\partial t} = -\sum_{i=1}^{m-1} \frac{(\phi - \phi_i)^2}{\sigma_i^2} \frac{\partial \chi_i(\boldsymbol{L})}{\partial L_j} - \lambda^2 \frac{\partial \chi_m(\boldsymbol{L})}{\partial L_j} - \gamma \delta(L_j) \nabla \left(\frac{\nabla L_j}{|\nabla L_j|} \right), \tag{15}$$

where the derivatives of the indicator functions χ_i are easily obtained from (12). The first two terms in (15) drive the labeling $\boldsymbol{L}$ to indicate the template ϕ_i which is most similar to the given function ϕ (or alternatively the background). The last term minimizes the length of the zero crossing of L_j. This has two effects: Firstly, it induces the labeling to decide for one of the possible templates (or the background), i.e. mixing of templates with label values between $+1$ and -1 are

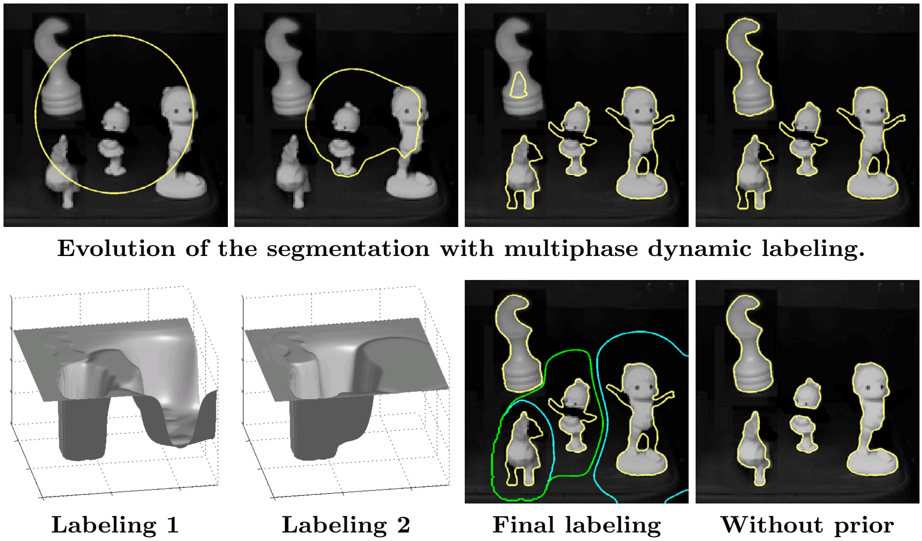

Fig. 7. Coping with several objects by multiphase dynamic labeling. Contour evolution generated by minimizing the total energy (6) with a multiphase selective shape prior of the form (13) encoding the three figures on the left, center and right. The appearance of all three objects is corrupted. Due to the simultaneous optimization of a vector-valued labeling function, several regions associated with each shape prior are selected, in which the given prior is enforced. All familiar shapes are segmented and restored, while the correct segmentation of separate (unfamiliar) objects remains unaffected. The images on the bottom show the final labeling and – for comparison – the segmentation without prior (right).

suppressed. Secondly, it enforces the decision regions (regions of constant label) to be "compact", because label flipping is energetically unfavorable.

Figure 7 shows a contour evolution obtained with the multiphase dynamic labeling model (13) and $n = 2$ labeling functions. The image contains three corrupted objects which are assumed to be familiar and one unfamiliar object (in the top left corner). The top row shows the evolution of the segmenting contour (yellow) superimposed on the input image. The segmentation process with a vector-valued labeling function selects regions corresponding to the different objects in an unsupervised manner and simultaneously applies three competing shape priors which permit to reconstruct the familiar objects. Corresponding 3D plots of the two labeling functions in the bottom rows of Figure 7 show which areas of the image have been associated with which label configuration. For example, the object in the center has been identified by the labeling $\boldsymbol{L} = (+1, -1)$.

8 Conclusion

We introduced the framework of multiphase dynamic labeling, which allows to integrate multiple competing shape priors into level set based segmentation schemes. The proposed cost functional is simultaneously optimized with respect to a level set function defining the segmentation, a vector-valued labeling function indicating regions where particular shape priors should be enforced, and a set of pose parameters associated with each prior. Each shape prior is given by a fixed template and respective pose parameters, yet the extension to statistical shape priors (which additionally allow deformation modes) is straight forward.

We argued that the proposed mechanism fundamentally generalizes previous approaches to shape priors in level set segmentation. Firstly, it is consistent with the philosophy of level sets because it retains the capacity of the resulting segmentation scheme to cope with multiple independent objects in a given image. Secondly, it addresses the central question of where to apply which shape prior.

The selection of appropriate regions associated with each prior is generated by the dynamic labeling in a recognition-driven manner. In this sense, our work demonstrates in a specific way how a recognition process can be modeled in a variational segmentation framework.

References

1. V. Caselles, F. Catté, T. Coll, and F. Dibos. A gemoetric model for active contours in image processing. *Numer. Math.*, 66:1–31, 1993.
2. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. IEEE Intl. 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. G. Charpiat, O. Faugeras, and R. Keriven. Approximations of shape metrics and application to shape warping and empirical shape statistics. Technical Report 4820, INRIA, Sophia Antipolis, May 2003.
5. Y. Chen, H. Tagare, S. Thiruvenkadam, F. Huang, D. Wilson, K. S. Gopinath, R. W. Briggs, and E. Geiser. Using shape priors in geometric active contours in a variational framework. *Int. J. of Computer Vision*, 50(3):315–328, 2002.
6. D. Cremers. A variational framework for image segmentation combining motion estimation and shape regularization. In C. Dyer and P. Perona, editors, *IEEE Conf. on Comp. Vis. and Patt. Recog.*, volume 1, pages 53–58, Madison, June 2003.
7. D. Cremers, N. Sochen, and C. Schnörr. Towards recognition-based variational segmentation using shape priors and dynamic labeling. In L. Griffith, editor, *Int. Conf. on Scale Space Theories in Computer Vision*, volume 2695 of *LNCS*, pages 388–400, Isle of Skye, 2003. Springer.
8. S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. J. Yezzi. Gradient flows and geometric active contour models. In *Proc. IEEE Intl. Conf. on Comp. Vis.*, pages 810–815, Boston, USA, 1995.
9. 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.
10. R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE PAMI*, 17(2):158–175, 1995.
11. J. Marroquin, E.A. Santana, and S. Botello. Hidden markov measure field models for image segmentation. *IEEE PAMI*, 25(11):1380–1387, 2003.

12. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.
13. 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.
14. N. Paragios and R. Deriche. Geodesic active regions and level set methods for supervised texture segmentation. *Int. J. of Computer Vision*, 46(3):223–247, 2002.
15. 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.
16. 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.
17. 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 *Comp. Vision Patt. Recog.*, pages 463–468, Kauai, Hawaii, 2001.
18. A. Tsai, A. J. Yezzi, and A. S. Willsky. Curve evolution implementation of the Mumford-Shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. on Image Processing*, 10(8):1169–1186, 2001.
19. L. Vese and T. Chan. A multiphase level set framework for image processing using the Mumford–Shah functional. *Int. J. of Computer Vision*, 50(3):271–293, 2002.