O. Faugeras, N. Paragios (Eds.), *2nd IEEE Workshop on Variational, Geometric and Level Set Methods in Computer Vision*, Nice, 2003.

A Pseudo-distance for Shape Priors in Level Set Segmentation

Daniel Cremers and Stefano Soatto

Department of Computer Science
University of California, Los Angeles – CA 90049

Abstract

We study the question of integrating prior shape knowledge into level set based segmentation methods. In particular, we investigate dissimilarity measures for shapes encoded by the signed distance function.

We consider extensions and improvements of existing measures. As a result, we propose a novel dissimilarity measure which constitutes a pseudo-distance. Compared to alternative approaches, this measure is symmetric and not biased toward small area. Based on this pseudo-distance, we propose a shape prior for level set segmentation methods which is pose invariant. In numerical experiments, we demonstrate that the resulting prior permits the segmentation of corrupted versions of a familiar object which is independent of the pose of the object. Moreover, we demonstrate the advantage of the symmetric formulation of the dissimilarity measure when segmenting corrupted images of known objects which consist of multiple components.

1. Introduction

Since their introduction by Osher and Sethian [10], level set based contour evolutions have become increasingly popular in image segmentation (cf. [1, 7, 11]). A central research question is how to incorporate higher-level prior shape information into such segmentation processes [8, 13, 4, 12].

Leventon et al. [8] associated each of a set of training shapes with a uniquely defined level set function via the signed distance function, generated a vector representation of each level set function by sampling it on a regular grid, and imposed a prior on the shape of the level set function during the segmentation process. The resulting shape prior was shown to strongly improve segmentation results obtained on 2D and 3D medical images with a geodesic active contour model. In particular, it was shown that imposing the prior on the level set function rather than on the embedded boundary is in accordance with the "philosophy" of the level set approach because it retains the capacity of the evolving boundary to undergo topological changes. Moreover, invariance to the pose of the object is obtained by local optimization of pose parameters. Yet, the above approach had two drawbacks: Firstly, the prior was added to the evolution equation rather than introduced on the variational level consistent with a proper probabilistic formulation. And secondly, the statistical model is defined on a grid of finite size which is not intrinsic to the given problem.

Rousson and Paragios [12] incorporated a shape prior at the variational level. They proposed a shape energy which consists of the squared distance between the evolving level set function and a reference level set function integrated over the part of the image domain where the former is positive. Again the image domain is discretized on a grid and a (spatially independent) statistical variation around the reference shape is permitted. Separate parameters to account for the 2D pose of the object are included in the cost functional. Gradient descent on this shape energy is added to the evolution equation of the geodesic active regions. The resulting segmentation process is capable of segmenting corrupted versions of a learnt object.

The present work builds up on the above approach and tries to overcome the following drawbacks: Firstly, the restriction of the energy integral to the positive part of the level set function induces a bias toward small shapes. The authors in [12] discard the resulting area shrinking term from the evolution equation, arguing that their model already contains such an area shrinking term. Secondly, the underlying shape dissimilarity measure is not symmetric with respect to the compared shapes. Due to the restriction of the shape discrepancy measure to the positive part of the evolving shape, all shape discrepancies outside the evolving shape are neglected. Such priors are therefore not well-suited to encode multi-component shapes.

In particular, we address the following questions:

- Can one formulate a dissimilarity measure on shapes encoded by the signed distance function which does not suffer from a bias toward small area of the shapes?
- What are the limitations of a shape prior which is based on an asymmetric dissimilarity measure?
- Can one formulate a shape dissimilarity measure which fulfills the requirements for a distance measure – in particular symmetry and the triangle inequality?

In the following sections, we will present answers to the above questions. To this end we will briefly review the level set formulation of the Mumford-Shah functional proposed by Chan and Vese [2]. We will then iteratively construct a shape dissimilarity measure which is to incorporate the above requirements. As a result, we obtain a pseudo-distance defined for shapes represented by their signed distance function. By construction, it is symmetric and not biased toward small area. We discuss the advantages of a symmetric formulation for shape priors encoding multi-component objects. We show that the proposed measure does not fulfill the triangle inequality.

Subsequently, the proposed measure is incorporated as a shape prior into the level set formulation of the Mumford-Shah functional. We derive corresponding Euler-Lagrange equations and show numerical results which demonstrate the capacity of the prior to compensate for corrupted image information when segmenting a familiar object. In particular, we show the effect of separately optimizing pose parameters. Moreover, we demonstrate the advantage of the symmetric formulation of the shape dissimilarity measure by segmenting corrupted versions of a known object which consists of several components.

2. The Model of Chan and Vese

In several papers, Chan and Vese (cf. [2]) detailed a level set implementation of the Mumford-Shah functional [9], which is based on the use of the Heaviside function as an indicator function for the separate phases.

Since we are focused on modeling shape priors, 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 segmenting boundary

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

A piecewise constant segmentation of an input image f with two gray values c_1 and c_2 can be obtained by minimizing the functional [2]:

$$\begin{aligned} E_{cv}(c_1, c_2, \phi) &= \int_\Omega (f - c_1)^2 \, H(\phi) \, dx \\ &+ \int_\Omega (f - c_2)^2 \, \big(1 - H(\phi)\big) \, dx \\ &+ \nu \int_\Omega |\nabla H(\phi)| \, dx, \end{aligned} \tag{2}$$

with respect to the scalar variables c_1 and c_2 modeling the gray value of the two phases and with respect to the embedding level set function ϕ modeling the interface separating these phases. $H(\phi)$ denotes the Heaviside function:

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

While the first two terms in (2) aim at maximizing the gray value homogeneity in the two separated phases, the last term aims at minimizing the length of the boundary given by the zero-crossing of ϕ.

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

$$\frac{\partial \phi}{\partial t} = \delta_\epsilon(\phi) \left[\nu \, \text{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - (f - c_1)^2 + (f - c_2)^2 \right], \tag{4}$$

where the scalars c_1 and c_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 with $\phi > 0$ and $\phi < 0$, respectively:

$$c_1 = \frac{\int f(x) H(\phi) dx}{\int H(\phi) dx}, \quad c_2 = \frac{\int f(x)(1 - H(\phi)) dx}{\int (1 - H(\phi)) dx}. \tag{5}$$

The implementation in [2] is based on a smooth approximation of the delta function $\delta_\epsilon(s) = H'_\epsilon(s)$, which is chosen to have an infinite support:

$$\delta_\epsilon(s) = \frac{1}{\pi} \frac{\epsilon}{\epsilon^2 + s^2}. \tag{6}$$

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.

3. A Unique Shape Representation

Two favorable properties of the implicit level set representation (1) over explicit contour representations are its independence of a particular parameterization, and the fact that the topology of the boundary is not constrained, such that merging and splitting of the contour during evolution is facilitated. When integrating a prior on the shape of the boundary C, one would like to retain these properties. To this end, one can impose a prior directly on the level set function ϕ.

As argued in [8], this requires that the level set function ϕ associated with a given contour C be *uniquely defined*. For this reason they suggest to encode each shape by the signed distance function, i.e. setting the value of the function ϕ at a given point to the distance to the nearest contour point, positive inside the contour and negative outside. In particular, this means that $|\nabla \phi| = 1$ almost everywhere.

4. Measuring Shape Dissimilarity

Let us assume that two shapes C_1 and C_2 are given in terms of their signed distance functions ϕ_1 and ϕ_2. How can we measure their dissimilarity $d(\phi_1, \phi_2)$? A natural choice is to measure the deviation in the sense of L_2:

$$d^2(\phi_1, \phi_2) = \int_\Omega (\phi_1 - \phi_2)^2 \, dx. \tag{7}$$

For an extensive discussion of related distance measures we refer to [3]. Yet, the measure (7) has an important drawback: It depends on the choice of the domain Ω. In particular, if one considers the same two shapes but simply extends the domain of integration Ω, then the measure of dissimilarity changes.

In order to remove this undesirable property, Rousson and Paragios [12] suggested to restrict the integration to the part of the image plane, where ϕ_1 is positive:

$$d^2(\phi_1, \phi_2) = \int_\Omega (\phi_1 - \phi_2)^2 \, H(\phi_1) \, dx. \tag{8}$$

For an alternative restriction of the domain of integration driven by a dynamic recognition process, we refer to [5].

The measure (8) of dissimilarity has two drawbacks which will be addressed in this paper, namely it depends on the size of shape ϕ_1, and it is not symmetric.

5. An Area-invariant Measure

The dissimilarity measure (8) depends on the size of the area in which ϕ_1 is positive. Indeed, gradient descent minimization of this measure with respect to ϕ_1 results in the equation:

$$\frac{\partial \phi_1}{\partial t} = -2 \, (\phi_1 - \phi_2) \, H(\phi_1) - (\phi_1 - \phi_2)^2 \, \delta(\phi_1). \tag{9}$$

Whereas the first term moves ϕ_1 toward ϕ_2 (as desired), the second term aims at minimizing the area of the shape which is clearly an undesired property.

A straight-forward remedy is to normalize with respect to the area where ϕ_1 is positive:

$$d^2(\phi_1, \phi_2) = \int_\Omega (\phi_1 - \phi_2)^2 \, h(\phi_1) \, dx, \tag{10}$$

where the normalized Heaviside function is defined by:

$$h(\phi) = \frac{H(\phi)}{\int_\Omega H(\phi) \, dx}. \tag{11}$$

Rather than summing the squared difference of the two level set functions over the area of positive ϕ_1, we propose to *average* it over this area.

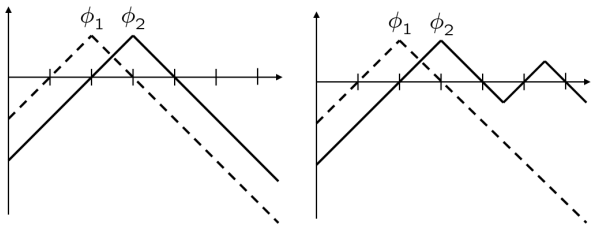

Figure 1: Limitations of asymmetric dissimilarity measures. The measures defined in (8) and (10) are not symmetric with respect to ϕ_1 and ϕ_2. As a consequence, for the above examples, the measures produce exactly the same dissimilarity $d(\phi_1, \phi_2)$. They entirely neglect the second component of ϕ_2 in the right image. Corresponding shape priors are therefore unsuited to encode shapes with multiple components.

Indeed, the corresponding Euler-Lagrange evolution equation is given by:

$$\frac{\partial \phi_1}{\partial t} = -2 \, (\phi_1 - \phi_2) \, h(\phi_1) - \frac{\delta(\phi_1)}{\int H(\phi_1) dx} \left[(\phi_1 - \phi_2)^2 - \int (\phi_1 - \phi_2)^2 \, h(\phi_1) dx \right]. \tag{12}$$

Due to the area-invariant formulation of the cost functional, the area shrinking term in the evolution equation is compensated.

6. A Symmetric Formulation

The measures (8) and (10) are not symmetric, that is in general:

$$d(\phi_1, \phi_2) \neq d(\phi_2, \phi_1). \tag{13}$$

The requirement of symmetry for a shape prior may not appear very relevant at first glance. Yet, one can easily construct examples where the symmetry violation becomes apparent. Figure 1 shows a 1-D example of a shape ϕ_1 which is compared to two other shapes ϕ_2 – one which is simply connected (left) and one which contains two components (right). The measures (8) and (10) give exactly the same for the two cases: They associate the shapes ϕ_1 and ϕ_2 only partially thereby ignoring the second component of ϕ_2.

A remedy is to symmetrize the measure (10):

$$d^2(\phi_1, \phi_2) = \int_\Omega (\phi_1 - \phi_2)^2 \, \frac{h(\phi_1) + h(\phi_2)}{2} \, dx. \tag{14}$$

This symmetrized dissimilarity measure averages the squared deviation of the level set functions not only over the area where $\phi_1 > 0$ but also over the area where $\phi_2 > 0$. In particular, it is therefore capable of estimating shape dissimilarities such as the one in Figure 1, which are due to a variation in the second shape.

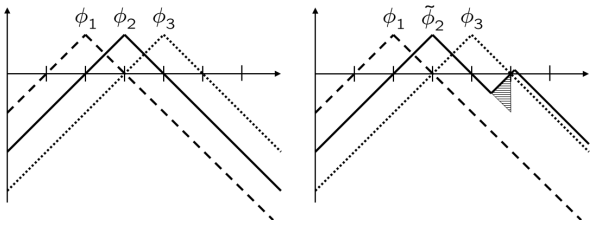

Figure 2: Violation of the triangle inequality. For the distance measure defined in equation (14), the left example fulfills the triangle inequality (exactly), whereas the right example does not. Due to the slight modification of ϕ_2, the distance between ϕ_2 and ϕ_3 becomes smaller (by the shaded area). Therefore this constellation violates the triangle inequality. Consequently, the distance defined in (14) is merely a *pseudo-distance*.

One may argue that the second shape does not vary in an application of shape priors in segmentation. It corresponds to the learnt shape which is fixed, while only the first shape evolves. Yet, the asymmetric shape measures (8) and (10), when integrated as a shape prior in a segmentation process, will only maximize the overlap of the evolving level set function with components of ϕ_2 in areas where $\phi_1 > 0$. All other shape discrepancies with respect to ϕ_2 are completely ignored. As a consequence, shape priors based on the above dissimilarity measures will generally fail to perform well when segmenting known objects with several components. A particular example of segmenting multi-component objects is discussed in Section 10.3 – see Figures 6 and 7. Moreover, a comparison in Figure 8 shows that the asymmetric approaches (8) and (10) are not suited to propagate shape discrepancy information outside the initial shape area.

7. A Pseudo-distance

Let S be the space of all signed distance functions associated with a particular shape on the domain Ω. Then the measure $d : S \times S \to \mathbb{R}$ defined in (14) constitutes a pseudo-distance, i.e. the requirements of positive semi-definiteness and symmetry are fulfilled:

- $d(\phi_1, \phi_2) \geq 0$ and $d(\phi_1, \phi_2) = 0 \Leftrightarrow \phi_1 = \phi_2$.
- $d(\phi_1, \phi_2) = d(\phi_2, \phi_1)$

The requirement that the considered space S only contains those signed distance functions which are associated with a particular shape on Ω guarantees that for each $\phi \in S$, the positive domain $\{x \in \Omega \,|\, \phi(x) > 0\}$ (i.e. the shape) is not empty. This requirement is necessary for the first property to hold: If the two functions are identical in the domain where they are positive, then by the requirement of a distance function they are the same (almost) everywhere.

The last requirement for this measure to qualify as a distance, namely the triangle inequality, is not fulfilled. For simplicity, we give a counterexample in dimension 1. Figure 2, left side, shows the level set functions ϕ_1, ϕ_2 and ϕ_3 associated with three shapes. A simple calculation shows that for the measure d defined in (14) we have:

$$d(\phi_1, \phi_3) = d(\phi_1, \phi_2) \; + \; d(\phi_2, \phi_3). \qquad (15)$$

For this constellation the triangle inequality is still fulfilled. Now let $\tilde{\phi}_2$ be a modification of the shape ϕ_2 which includes an (infinitesimally small) second component – see Figure 2, right side – then the distance $d(\phi_2, \phi_3)$ decreases while the two other distances remain essentially the same. Therefore we have

$$d(\phi_1, \phi_3) > d(\phi_1, \tilde{\phi}_2) \; + \; d(\tilde{\phi}_2, \phi_3). \qquad (16)$$

This violates the triangle inequality and shows that the measure d defined in (14) is merely a *pseudo*distance.

The violation of the triangle inequality is obviously linked to the fact that the proposed dissimilarity measure only takes into account areas of the image plane where one of the level set functions is positive. On the other hand, this restriction to the positive part was introduced in order to make the measure independent of the domain of integration. How one can formulate dissimilarity measures which combine the latter requirement with that of the triangle inequality is an open problem.

8. A Euclidean Invariant Shape Prior

In the following, we will propose a shape prior which is based on the pseudo-distance defined in (14) and which incorporates invariance with respect to the group of Euclidean transformations of a given shape, in a way similar to the one proposed in [8, 4, 12] for similarity transformations.

At this point, we only consider a single training shape. Yet, extensions to incorporate a statistical shape variation such as the ones presented in [8, 12] are conceivable.

Let ϕ_0 be the distance function associated with a given training shape. Then we define a shape energy by

$$E_s(\phi, \mu, \theta) = d^2\Big(\phi(x), \phi_0\left(R_\theta x + \mu\right)\Big), \qquad (17)$$

where we allow for a Euclidean transformation of the reference shape given by a translation $\mu \in \mathbb{R}^2$ and an orthogonal matrix $R_\theta \in \mathbb{R}^{2\times 2}$ which accounts for rotation by an angle $\theta \in [0, 2\pi]$.

In order to obtain a segmentation process which takes into account both the intensity information of the input image and the prior shape knowledge, we propose to minimize the joint energy:

$$E(\phi, c_1, c_2, \mu, \theta) = E_{cv}(\phi, c_1, c_2) \; + \; \alpha\, E_s(\phi, \mu, \theta). \quad (18)$$

This variational integration of shape prior and image information is equivalent to Bayesian a posteriori maximization.

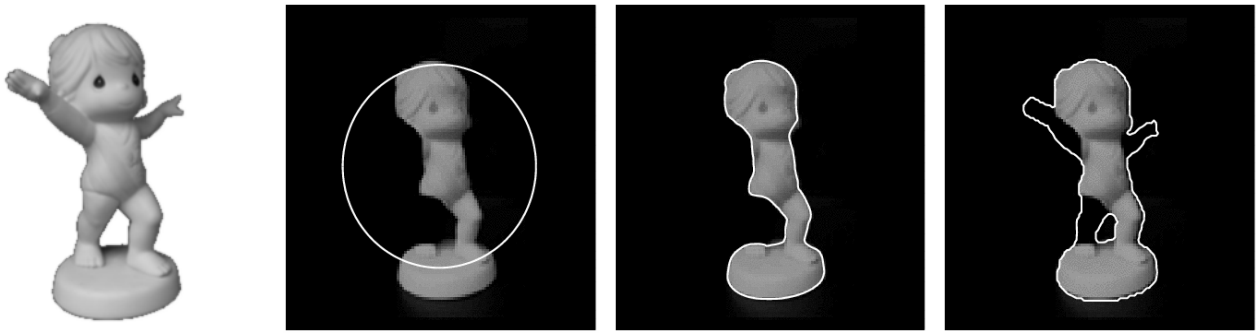

Figure 3: Segmentation of a corrupted input image without and with shape prior. The object of interest is the dancer shown on the left. The second frame shows the input image and initial contour. The frames on the right show the segmentation results without and with shape prior. The shape model permits to reconstruct missing parts of the object.

9. Energy Minimization

In order to minimize the functional (18) with respect to the dynamic variables, we alternate between an update of the constants c_i given in (5) and a gradient descent with respect to the level set function ϕ and the pose parameters μ and θ.

The latter are given by:

$$\frac{\partial \phi}{\partial t} = -\frac{\partial E_{cs}}{\partial \phi} - (\phi - \phi_0)\left(h(\phi) + h(\phi_0)\right) - \frac{\delta(\phi)}{2\int H(\phi)dx}\left[(\phi-\phi_0)^2 - \int (\phi-\phi_0)^2 h(\phi)dx\right], \tag{19}$$

where the first (the data driven) term is given in (4). The optimization of the pose parameters $\rho = \{\mu, \theta\}$ is given by equations of the form:

$$\frac{\partial \rho}{\partial t} = -\int (\phi - \phi_0)\left(h(\phi) + h(\phi_0)\right)\nabla\phi_0^t \frac{\partial g}{\partial \rho}\, dx - \frac{1}{2\int H(\phi)dx}\int\left[(\phi-\phi_0)^2 - \overline{(\phi-\phi_0)^2}\right]\delta(\phi_0)\nabla\phi_0^t\frac{\partial g}{\partial \rho}\, dx, \tag{20}$$

where the expressions $g(x, \theta, \mu) = R_\theta x + \mu$ and

$$\overline{(\phi - \phi_0)^2} = \int (\phi - \phi_0)^2 h(\phi_0)\, dx \tag{21}$$

were introduced for simplification. As above, we approximate the delta function in (19) and (20) by the dilated one given in equation (6).

10. Numerical Results

In this section, we will demonstrate several properties of the proposed shape prior applied to segmenting a known object in a corrupted input image. In particular, we will show the improvements yielded by the normalization and the symmetrization proposed in Sections 5 and 6. We will also demonstrate the effect of the pose optimization.

10.1. Knowledge-driven segmentation

Figure 3 shows the basic property of the shape prior integrated into the variational framework. The left image shows the object of interest. The following images show segmentation results obtained on a corrupted version of the object – the initial contour, the segmentation without prior and the segmentation with prior. The prior compensates for missing or corrupted image information and thereby permits to reconstruct the object of interest.

10.2. Effect of pose optimization

The following examples are to demonstrate the effect of the pose optimization in equation (20). To this end, we compared segmentation results with pose optimization to ones in which the pose optimization was selectively suppressed. The first four images in Figure 4 show the evolution of the boundary with a shape prior which was offset with respect to the object in the image. During energy minimization both the shape of the boundary and its pose are updated. In contrast, if translation optimization is suppressed then the final result will have the same location as the training shape.

Figure 5 shows a similar comparison for the case of a shape prior which is rotated with respect to the correct segmentation. As above, the first four images show the boundary evolution with simultaneous pose optimization. The last image shows the final segmentation when the pose optimization is suppressed: The segmentation fails to capture the pose of the object in the image.

10.3. Encoding multi-component objects

As argued in Section 6, the advantage of the symmetric formulation of the shape dissimilarity measure in (14) permits the corresponding shape prior to encode shapes which contain multiple components. Figure 6 shows the contour evo-

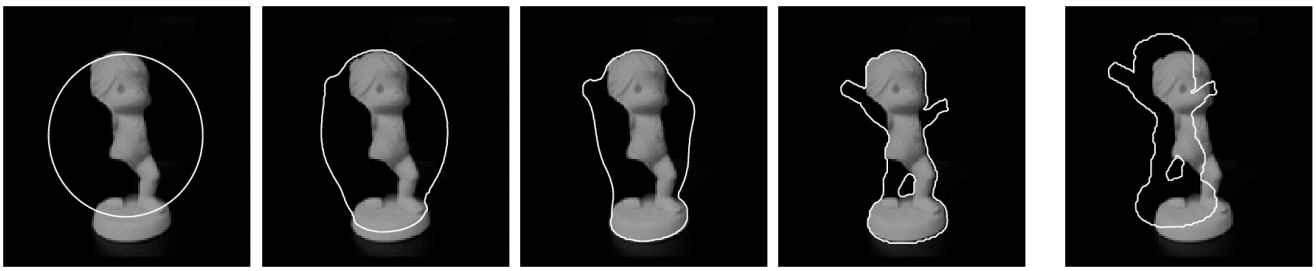

Evolution with translation invariance enabled. **No translation.**

Figure 4: Effect of translation optimization. Due to the simultaneous optimization with respect to the translatory degrees of freedom, the evolving boundary is free to translate while still being restricted to the familiar shape (left images). If the translation optimization is suppressed, then the contour will converge toward the learnt shape in exactly the location of the latter (right image).

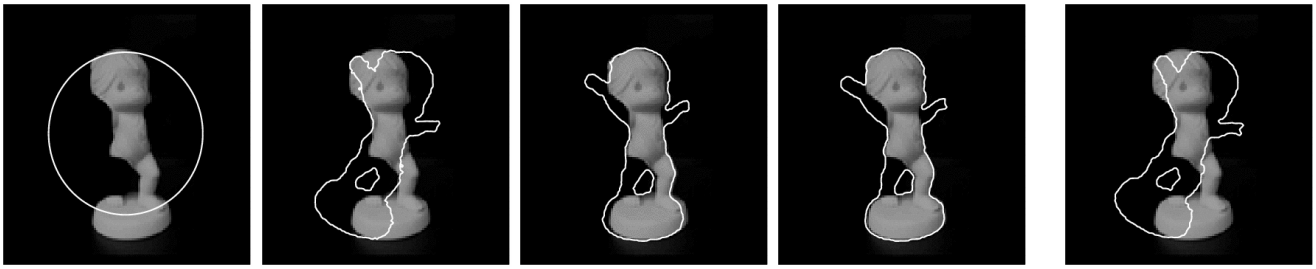

Evolution with rotation invariance enabled. **No rotation.**

Figure 5: Effect of rotation optimization. Due to the simultaneous optimization with respect to a rotation angle, the evolving boundary is also free to rotate while still being restricted to the familiar shape (left images). If the translation optimization is suppressed, then the contour will converge toward the learnt shape in exactly the rotation of the latter (right image).

lution without shape prior obtained on a corrupted version of a text. Although the length prior permits to cope with noise, it cannot handle the cases of occlusion and missing data.

Figure 7 shows the contour evolution obtained for the same input image with a prior which encodes the entire word "shape". The prior permits to compensate for occlusion and missing data. In particular, it permits to reconstruct the letter 'e' which was entirely missing in the input image. As discussed in Section 6, the asymmetric versions of the dissimilarity measure given in (8) and (10) are not capable of solving this task, since they only measure shape discrepancies in areas where the evolving level set function is positive. The symmetrized version (14) on the other hand, takes into account *all* shape discrepancies between the evolving shape and the template. Consequently it permits a complete reconstruction of the object of interest.

10.4. A model comparison

To demonstrate the advantages of the proposed modifications explicitly, we performed the above task of reconstructing a corrupted version of the word "shape" based on the three shape dissimilarity measures discussed in this paper. Figure 8 shows the initial contour and the respective segmentation results obtained for the asymmetric and unnormalized model (8), the asymmetric normalized model (10) and the symmetrized and normalized model (14). For comparison, all parameters were kept constant. The results show that the asymmetric models are not able to propagate the shape discrepancy information to areas outside the initial segmentation (i.e. to areas where $\phi < 0$). Therefore the reconstruction is essentially restricted to the domain where $\phi > 0$ in the initialization – see second and third image in Figure 8. This was discussed in Section 6. Moreover, one can observe a slight area shrinking (in the second image when compared to the third one) due to the lack of normalization – see the discussion in Section 5.

11. Limitations and Future Work

The above framework has several limitations which we intend to overcome in future work. Firstly, the shape prior suggested here merely encodes a single shape. Extensions

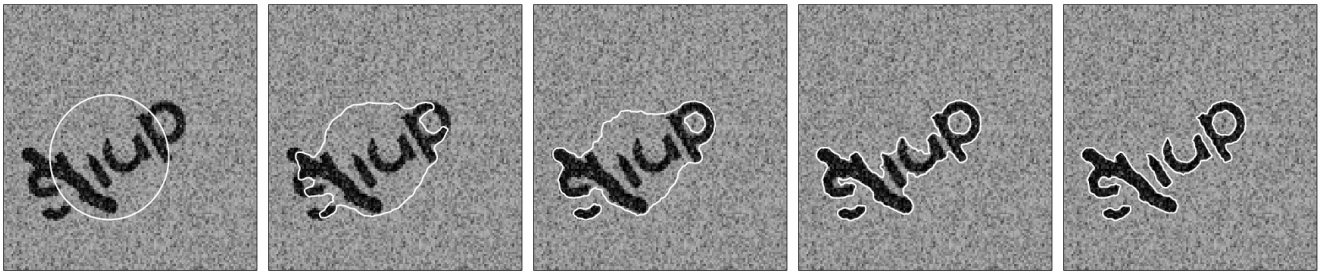

Figure 6: Segmentation of text without shape prior. The length constraint on the boundary can only compensate for small scale noise. It does not handle missing information or occlusions.

Figure 7: Segmentation of text with a multi-component shape prior. The shape prior encodes the entire word "shape". It compensates for missing parts and occlusion and thereby permits a reconstruction of the text. Note in particular, that the letter 'e' is reconstructed although it was entirely missing in the input image. This property can be directly attributed to the symmetric formulation in (14). The asymmetric dissimilarity measures in (8) and (10) only minimize shape discrepancies from the learnt shape in areas where the evolving level set function is positive. Therefore, they are not suited to perform the above task – see Figure 8.

to models encoding multiple training shapes have been proposed in [8, 12]. Present work is focused on deriving probabilistic shape models on the basis of the pseudo-distance introduced in this paper.

Secondly, we want to remark that the pose optimization presented in Sections 8 and 9 and evaluated in Section 10.2 – similar to the one proposed in [8, 4, 12] – is a local optimization scheme. This leads to the following limitation which we observed in numerical implementations: If the pose of the training shape and the object to be segmented are too different (e.g. a relative rotation by an angle π), then the segmentation process will generally fail to converge to the correct solution.

The numerical optimization of explicit pose parameters introduces additional parameters to balance the respective gradient descent equations (19) and (20). This is not only tedious, but also prone to numerical instabilities. A closed-form intrinsic elimination of the pose parameters is preferable. For spline contours, such an approach was proposed by Cremers et al. in [6]. Present work is focused on deriving similar registration schemes for the level set framework.

12. Summary and Conclusions

We proposed a novel approach to integrate higher-level shape priors into level set based variational segmentation methods. To this end, we investigated dissimilarity measures for shapes encoded by the signed distance function. Building up on an approach of Rousson and Paragios [12], we propose a novel dissimilarity measure which improves the latter in several ways:

Firstly, it is not biased toward small shape area. Therefore one no longer needs to remove spurious area-shrinking terms from the evolution equation.

Secondly, it is symmetric in the compared shapes. The corresponding shape prior takes into account shape discrepancies even in areas where the evolving level set function is negative. Shape reconstructions are therefore no longer limited to the initial shape domain. In particular, our formulation permits to segment corrupted versions of known objects which consist of several components.

We show that the constructed dissimilarity measure is in fact a pseudo-distance. In particular, we give an example for which the triangle inequality is *not* fulfilled.

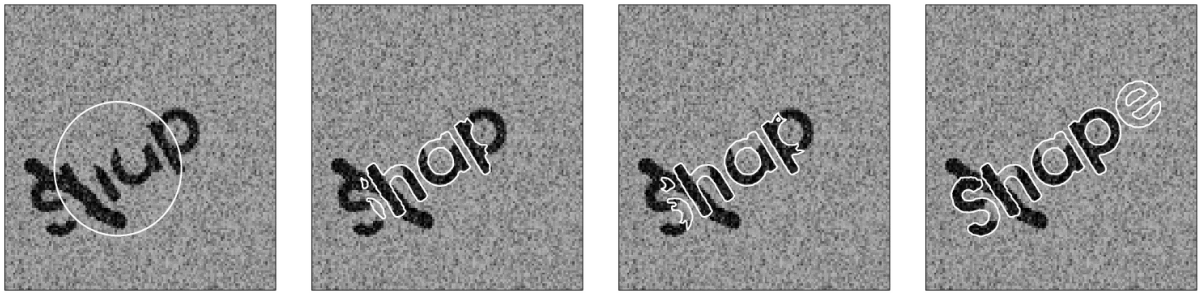

Initial contour | **unnormalized and asymmetric (8)** | **normalized and asymmetric (10)** | **normalized and symmetric (14)**

Figure 8: Model comparison. The image on the left shows the initial contour. The images on the right show the resulting segmentations obtained with a shape prior encoding the entire word "shape", based on the unnormalized and asymmetric dissimilarity (8), the normalized version (10) and the symmetrized and normalized version (14). For comparison, all other parameters were kept fixed during minimization. Neither of the two asymmetric formulations is capable of propagating shape discrepancy information outside the initial shape area ($\phi > 0$) – this was discussed in Section 6. Therefore the reconstructions in the second and third image are essentially restricted to the initial shape domain shown in the left image. Note that the unnormalized version (2nd image) induces a slight area shrinking compared to the normalized version (3rd image) – this was discussed in Section 5.

The proposed shape prior is invariant to 2D pose transformations. This property is demonstrated numerically by comparing segmentation results obtained with and without pose optimization.

Finally, we discuss limitations of the present approach regarding pose invariance and the encoding of multiple training shapes. These are the focus of ongoing research.

Acknowledgments

We thank Paolo Favaro for fruitful discussions. This research was supported by ONR N00014-02-1-0720 and AFOSR F49620-03-1-0095.

References

[1] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. IEEE Internat. Conf. on Comp. Vis.*, pages 694–699, Boston, USA, 1995.

[2] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.

[3] 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.

[4] 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.

[5] 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.

[6] D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr. Diffusion Snakes: Introducing statistical shape knowledge into the Mumford–Shah functional. *Int. J. of Computer Vision*, 50(3):295–313, 2002.

[7] 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.

[8] 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.

[9] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.

[10] 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.

[11] N. Paragios and R. Deriche. Coupled geodesic active regions for image segmentation: a level set approach. In D. Vernon, editor, *Proc. of the Europ. Conf. on Comp. Vis.*, volume 1843 of *LNCS*, pages 224–240. Springer, 2000.

[12] 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.

[13] 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.