A Shape-Based Approach to Robust Image Segmentation

Samuel Dambreville, Yogesh Rathi, and Allen Tannenbaum

Georgia Institute of Technology, 30322 Atlanta Georgia USA.

Abstract. We propose a novel segmentation approach for introducing shape priors in the geometric active contour framework. Following the work of Leventon, we propose to revisit the use of linear principal component analysis (PCA) to introduce prior knowledge about shapes in a more robust manner. Our contribution in this paper is twofold. First, we demonstrate that building a space of familiar shapes by applying PCA on binary images (instead of signed distance functions) enables one to constrain the contour evolution in a way that is more faithful to the elements of a training set. Secondly, we present a novel region-based segmentation framework, able to separate regions of different intensities in an image. Shape knowledge and image information are encoded into two energy functionals entirely described in terms of shapes. This consistent description allows for the simultaneous encoding of multiple types of shapes and leads to promising segmentation results. In particular, our shape-driven segmentation technique offers a convincing level of robustness with respect to noise, clutter, partial occlusions, and blurring.

1 Introduction

Segmentation consists of extracting an object from an image, a ubiquitous task in computer vision applications. Such applications range from finding special features in medical images to tracking deformable objects; see [1–5] and the references therein. The active contour framework [6], which utilizes image information to evolve a segmenting curve, has proven to be quite valuable for performing this task. However, the use of image information alone often leads to poor segmentation results in the presence of noise, clutter or occlusion. The introduction of shapes priors in the contour evolution process has proven to be an effective way to circumvent this issue, leading to more robust segmentation performances.

Many different algorithms have been proposed for incorporating shape priors in the active contour framework. For example, various approaches for utilizing shape priors in parameterized representations of contours were proposed by Cootes *et al.* [7], Wang and Staib [8], and Cremers *et al.* [9]. Moreover, Cremers *et al.* [10], recently presented a statistical approach using kernel methods [11], for building shape models. Using this method for parameterized contours, the authors were able to construct shape priors involving various objects, and to obtain convincing segmentation results.

The geometric active contour framework (GAC) (see [12] and the references therein) involves a parameter free representation of contours: contours are represented implicitly by level-set functions (such as signed distance function [13]). Within this framework, Leventon *et al.* [1], proposed an algorithm in which principal component analysis (PCA) was performed on a training set of signed distance functions (SDFs) and the shape statistics thus obtained were used to drive the segmentation process. This statistical approach was shown to be able to convincingly capture the variability in shape of a particular object. This approach inspired other segmentation schemes described in [3, 14], notably, where SDFs were used to learn the shape of an object.

In this paper, we propose to revisit the use of linear PCA to introduce prior knowledge about shapes into the geometric active contour framework. To this end, we present a novel variational approach totally described in terms of shapes. Experimental results are presented to illustrate the robustness of our method: We first demonstrate the ability of our algorithm to constrain the contour evolution in a way that is more faithful to the training sets of shapes than prior work involving linear PCA. Then, we show the possibility, within our framework, to simultaneously encode knowledge about the shape of different objects, while capturing the variability in shape of each particular object.

2 Shape-based Segmentation with level-sets

Level-set representations were introduced in [13] in the field of computational physics and became a popular tool for image segmentation. The idea consists of representing the segmenting contour by the zero level-set of a smooth and

continuous function. The standard choice is to use a signed distance function for embedding the contour. During the segmentation process, the contour is propagated implicitly by evolving the embedding function. Implicit representations present the advantage of avoiding to deal with complex re-sampling schemes of control points. Moreover, the contour represented implicitly can naturally undergo topological changes such as splitting and merging.

In what follows, we consider the problem of segmenting a gray-level given image $I : \Omega \rightarrow \mathbf{R}$, where Ω is a subdomain of the plane. The term *segmentation* in this context will be taken to refer to the process of extracting an object of interest from the background and potential clutter in I. Let Φ denote the corresponding signed distance function to be used for segmentation [12]. We denote the curve corresponding to the zero level-set of Φ as γ. The curve γ separates the image domain Ω into two disjoint regions, γ_1 and γ_2. We assume that for all $(x, y) \in \gamma_1$, we have that $\Phi(x, y) \geq 0$. We denote by $H\Phi$ the Heaviside function defined such as, $H\Phi(x, y) = 1$ if $\Phi(x, y) \geq 0$ and $H\Phi(x, y) = 0$, otherwise. $H\Phi$ is a binary map and will be interpreted as the *shape* associated to the signed distance function Φ, in what follows.

Segmentation using shape priors can be carried using an energy of the form (i.e.: [9])

$$E(\Phi, I) = \beta_1 E_{shape}(\Phi) + \beta_2 E_{image}(\Phi, I). \tag{1}$$

In this expression, E_{shape} is an energy functional embedding shape information, and E_{image} is an energy functional encoding image information available at time t. The minimization of $E(\Phi, I)$ with respect to Φ can be accomplished via the standard gradient descent approach:

$$\frac{d\Phi}{dt} = -\nabla_{\Phi} E(\Phi, I) \quad \text{i.e., } \Phi(t + dt) = \Phi(t) - dt \cdot \nabla_{\Phi} E(\Phi, I) \tag{2}$$

In the rest of this section, we present a method for the introduction of shape priors into level-set framework for purpose of segmentation. We first discuss our approach for building a space of shapes and compare our method with prior work. We then introduce an energy function encoding our knowledge of the shape of the object of interest. Finally, we propose an energy functional involving a shape model that aims at exploiting intensity information in the image.

2.1 Space of Shapes

In what follows, we will assume that we have a training set τ of N binary images $\{\mathbf{I}_1, \mathbf{I}_2, ..., \mathbf{I}_N\}$ (the $\mathbf{I}_i$'s are described by $m \times n$ matrices). The $\mathbf{I}_i$'s represent the possible shapes of objects of interest. Each $\mathbf{I}_i$ has values of 1 inside the object and 0 outside. Such training sets are presented Figure 1. The shapes in τ are supposed to be aligned using an appropriate registration scheme (see, e.g., [3]) in order to discard differences between them due to similarity transformations.

In this paper, we propose to perform PCA directly on the binary images of the training set τ, instead of applying PCA on signed distance functions as advocated in [3, 1]. The method is succinctly presented in what follows. First,

the mean shape μ is computed by taking the mean of the training shapes, $\mathbf{I}_i$'s, $\mu = \frac{1}{N}\sum_{i=1}^{N}\mathbf{I}_i$. Then, the covariance matrix $\mathbf{C}$, representing the variations in shapes, can be computed as follows. The mean shape μ is subtracted from each $\mathbf{I}_i$ to create a mean-offset map $\tilde{\mathbf{I}}_i$. Each such map, $\tilde{\mathbf{I}}_i$ is written as a column vector $\tilde{\mathbf{I}}_i^c$ (The n columns of $\tilde{\mathbf{I}}_i$ are stacked on top of one another to form a large mn-dimensional column vector). Each $\tilde{\mathbf{I}}_i^c$ is then placed as the i^{th} column of a matrix $\mathbf{M}$, resulting in a $(mn) \times N$-dimensional matrix $\mathbf{M} = [\tilde{\mathbf{I}}_1^c, \tilde{\mathbf{I}}_2^c, ..., \tilde{\mathbf{I}}_N^c]$. The covariance matrix is then $\mathbf{C} = \frac{1}{N}\mathbf{M}\mathbf{M}^T$. Finally, using singular value decomposition $\mathbf{C}$ is decomposed as $\mathbf{C} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$ where $\mathbf{U}$ is a matrix whose column vectors represent the set of orthogonal modes of shape variation (principal components) and $\mathbf{\Sigma}$ is a diagonal matrix of corresponding eigenvalues. Each column $\mathbf{u}_i$ of $\mathbf{U}$ can be rearranged as an $m \times n$ matrix by inverting the stacking process involved in the construction of the $\tilde{\mathbf{I}}_i^c$'s. The rearranged column $\mathbf{u}_i$, forms the eigen-shape modes $\mathbf{S}_i$ (corresponding to the i^{th} eigenvalue of $\mathbf{\Sigma}$).

Let $\mathbf{S}$ be any binary map, representing an arbitrary shape. The coordinates α^k of the projection of $\mathbf{S}$ onto the first k components of the space of shapes can be computed as

$$\alpha^k = \mathbf{U}_k^T(\mathbf{S}^c - \mu^c) \tag{3}$$

where $\mathbf{U}_k$ is a matrix consisting of the first k columns of $\mathbf{U}$, $\mathbf{S}^c$ and μ^c are the column vectors obtained by stacking the columns of $\mathbf{S}$ and μ, respectively. Given the coordinates $\alpha^k = [\alpha_1^k, \alpha_2^k, ..., \alpha_k^k]^T$, the projection of $\mathbf{S}$, denoted by $P^k(\mathbf{S})$, can be obtained as $P^k(\mathbf{S}) = \Sigma_{i=1}^k \alpha_i^k.\mathbf{S}_i + \mu$.

2.2 Shape Priors

Shape Energy: To include prior knowledge on the shape, we propose to use $P^k(H\phi)$ as a model and to minimize the following energy:

$$E_{shape}(\Phi) := \|H\Phi - P^k(H\Phi)\|^2 = \int_\Omega [H\Phi - P^k(H\Phi)]^2 \, dxdy. \tag{4}$$

Note that $E_{shape}(\Phi)$ is the squared L_2-distance between $H\Phi$ and the projection of $H\Phi$ onto the shape space. Minimizing E_{shape} amounts to drive the shape $H\Phi$ towards the space of learnt shapes represented by binary maps.

Comparison with Past Work: We believe that the main advantage of using binary maps over SDFs resides in the fact that binary maps have limited support (0 everywhere but inside the object, where it is 1), whereas SDFs can take a wide range of values on the whole domain Ω. As a consequence, linear combinations of SDFs can lead to shapes that are very different from the learned shape. This phenomenon is illustrated Figure 2. Familiar spaces of shapes were constructed for the training sets presented in Figure 1, using either SDFs or binary maps. One of the shapes of the training set was slightly modified to form a new shape $\mathbf{S}$, see Figure 2(b). The projections of $\mathbf{S}$ on both spaces (SDFs and binary maps) are presented in Figure 2(c) and (d). For each of the two cases presented,

the shape obtained from the projection on the space derived from the SDFs (Figure 2(c)) bears little resemblance with the learnt shapes. In contrast, the projection obtained from the space constructed from binary maps, Figure 2(d), is clearly more faithful to the learned shapes. Hence, it can be expected that shape spaces based on binary maps afford more robustness to clutter than shape spaces built from SDFs. This point will be reinforced in the sequel.

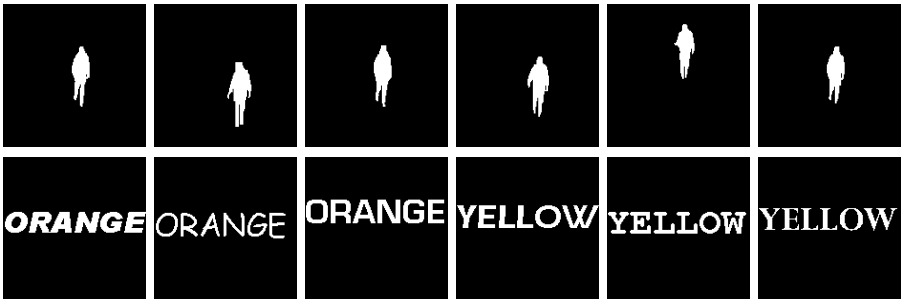

Fig. 1. Two Training sets (before alignment). First row, training set of objects of the same type: human silhouettes. Second row, training set of objects of different aspects: 2 words "ORANGE" and "YELLOW" (3 of the 20 binary maps for each).

2.3 Shape-based Approach to Region-based Segmentation

Different models [15, 16], which incorporate geometric and/or photometric (color, texture, intensity) information, have been proposed to perform region-based segmentation using level-sets. In what follows, we present an region-based based segmentation framework, with a strong shape interpretation. As in Section (2.2), at each step t of the evolution process we build a model of shape to drive segmenting contour towards it. Here, the model of shape is extracted from the image.

Following [16], we assume that the image I is formed by two regions of different mean intensity. The first region can be interpreted as the background, whereas the second region can be interpreted as the object of interest. As presented in Section (2) the zero level-set of Φ separates Ω into two regions γ_1 and γ_2. We compute the mean c_1 and c_2 of the intensity corresponding to pixels located in γ_1 and γ_2, respectively: $c_1 := \frac{\int I(x,y)H(\Phi)dx\ dy}{\int H(\Phi)\ dxdy}$ and $c_2 := \frac{\int I(x,y)(1-H(\Phi))\ dxdy}{\int (1-H(\Phi))\ dxdy}$. In these expressions, c_1 can be interpreted as the best estimate of the mean intensity of the object of interest, while c_2 can be interpreted as the best estimate of the intensity of the background, at time t. We now build

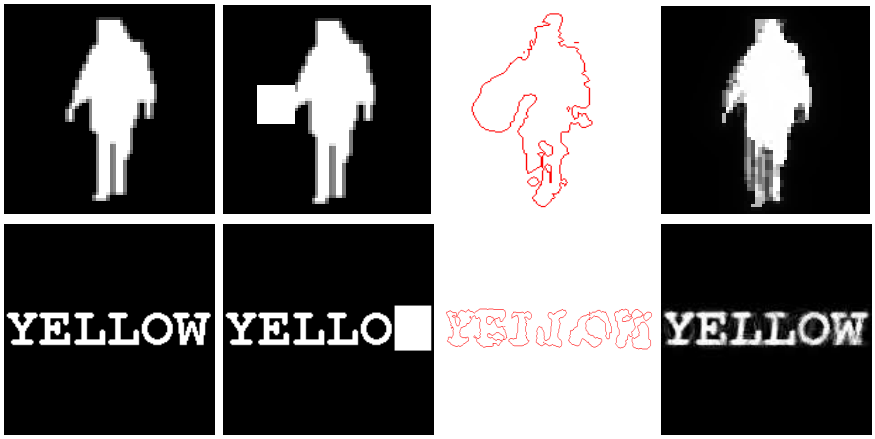

(a) (b) (c) (d)

Fig. 2. Comparison of shape priors built from SDFs and binary maps. (a): Original training shape. (b): Slight modification of the training shape. (c): Projection on space of Shape built from SDFs. (d): Projection on space of Shape built from binary maps. Projections obtained using binary maps are more faithful to the learnt shapes.

the image shape model $G_{[I,\Phi]}$ by thresholding I in the following manner:

$$\begin{array}{lll} \text{if } c_1 > c_2; & G_{[I,\Phi]}(x,y) = 1 & \text{if } I(x,y) \geq \frac{c_1+c_2}{2} \\ & = 0 & \text{otherwise;} \\ \text{if } c_1 \leq c_2; & G_{[I,\Phi]}(x,y) = 0 & \text{if } I(x,y) \geq \frac{c_1+c_2}{2} \\ & = 1 & \text{otherwise.} \end{array} \tag{5}$$

This thresholding insures that $G_{[I,\Phi]}$ is a binary map representing a certain shape. The manner in which the cases are defined makes the pixels in the image whose intensity is closer to c_1 to be set to 1 in the model, while the pixels closer to c_2 are set to 0. $G_{[I,\Phi]}$ can be interpreted as the most likely shape present in the image, knowing Φ.

We next minimize the following energy in order to drive the contour evolution towards the model of shape obtained from thresholding the image I:

$$E_{image}(\Phi, I) := \|H\phi - G_{[I,\Phi]}\|^2 = \int_\Omega (H\phi - G_{[I,\Phi]})^2 \, dxdy. \tag{6}$$

The energy functional amounts again to measuring the L^2 distance between two shapes, represented by $H\Phi$ and $G_{[I,\Phi]}$. Results of this image segmentation approach are presented Figure 3(b) and Figure 5(b) (where contour was evolved minimizing E_{image} only). The consistent description of E_{image} and E_{shape} in terms of binary maps allows for efficient and intuitive equilibration between image cues and shape knowledge.

2.4 Numerical Algorithm

Approximation of Functions: In our implementation of the above framework, we used the following $C^\infty(\bar{\Omega})$ regularizations:

$$H_{\varepsilon_1}\Phi(x,y) := \left(\frac{1}{2} + \frac{1}{\pi}\arctan\frac{\Phi(x,y)}{\varepsilon_1}\right) \quad \text{and} \quad \delta_{\varepsilon_1}\Phi(x,y) := \frac{1}{\pi}\left(\frac{\varepsilon_1}{\Phi^2(x,y)+\varepsilon_1^2}\right) \tag{7}$$

where ε_1, is a parameter such that $H_{\varepsilon_1} \to H$ and $\delta_{\varepsilon_1} \to \delta$ as $\varepsilon_1 \to 0$ (with $\delta = H'$). The function $G_{[I,\Phi]}$ in (5) is regularized as follows:

$$\begin{array}{ll} \text{if } c_1 > c_2; & G_{[I,\Phi,\varepsilon_2]}(x,y) = \frac{1}{2} + \frac{1}{\pi}\arctan\left(\frac{I(x,y)-\frac{c_1+c_2}{2}}{\varepsilon_2}\right) \\ \text{if } c_1 \leq c_2; & G_{[I,\Phi,\varepsilon_2]}(x,y) = \frac{1}{2} - \frac{1}{\pi}\arctan\left(\frac{I(x,y)-\frac{c_1+c_2}{2}}{\varepsilon_2}\right), \end{array} \tag{8}$$

where ε_2, is a parameter such that $G_{[I,\Phi,\varepsilon_2]} \to G_{[I,\Phi]}$ as $\varepsilon_2 \to 0$.

Invariance to Similarity Transformations: Let $\mathbf{p} = [a, b, \theta, \rho] = [p_1, p_2, p_3, p_4]$ be the vector of parameters corresponding to an affine transformation: a and b correspond to translation according to x and y-axis, θ is the rotation angle and ρ the scale parameter. Let us denote by $\hat{I}(\hat{x}, \hat{y})$ the image of I by the affine transformation of parameter $\mathbf{p}$: $\hat{I}(\hat{x}, \hat{y}) = I(\rho(x\cos\theta - y\sin\theta + a), \rho(x\sin\theta + y\cos\theta + b))$. As mentioned above the elements of the training sets are aligned prior to the construction of the space of shapes. Let us suppose that the object of interest in I differs from the registered elements of the training set by an affine transformation. This transformation can be recovered by minimizing $E(\Phi, \hat{I})$ with respect to the p_i's. During evolution, the following gradient descent scheme can be performed:

$$\frac{dp_i}{dt} = -\nabla_{p_i} E(\Phi, \hat{I}) = -\nabla_{p_i} E_{image}(\Phi, \hat{I}) \quad \text{for } i \in [1, 4].$$

Level-set Evolution: Computing the gradients corresponding to E_{shape} (equation (4)) and E_{image} (equation (6)) and accounting for possible affine transformations of the object of interest in I, equation (2) can be written as

$$\frac{d\Phi}{dt} = 2\delta_{\varepsilon_1}\Phi[\beta_1 P^k(H_{\varepsilon_1}\Phi) + \beta_2 G_{[\hat{I},\Phi,\varepsilon_2]} - (\beta_1 + \beta_2)H_{\varepsilon_1}\Phi]. \tag{9}$$

3 Segmentation Results

In this section we present experimental results aimed at testing the segmentation performances of our framework on challenging images.

3.1 Shape Prior Involving Objects of the Same Type: "Swedish Couple":

The Swedish Couple sequence was used as a base of test for many tracking algorithms using active contours [6]. One of the difficulties of performing tracking in this video resides in maintaining the identity of each person: Throughout the sequence, the two people often touch each other and the segmenting contour can leak from one person to the other, leading to a loss of track.

A training set of 7 silhouettes was obtained by manually selecting the contour for the person on the right on 7 different images from the sequence (the corresponding binary maps or SDFs were then completed).

Figure 3(c) presents a segmentation result obtained using shape priors built from applying PCA to SDFs (e.g., [3]). In the image presented, the two persons touch each other. The shape prior is not sufficiently constraining to prevent the contour from leaking from the person on the right (man) to the person on the left (woman). Hence, if no further control is applied to bound the coordinates of the contour in (3), leakage can occur leading to a shape that is very different from the ones in the training set. This phenomenon does not occur when using shape priors built from applying PCA on binary maps as presented in Figure 3(d). Here, the shape prior prevents leakage leading to a satisfying segmentation.

In fact, using shape priors built from the 7 binary maps obtained from the man, we were able to track the entire sequence for both persons (while maintaining their identity). Figure 4 presents the results obtained for a few images from the video. Despite the small number of shapes used, the general posture of each person was convincingly captured in each image. In addition, the final contours are faithful to the training set: No leakage occurred from one person to the other and the bag held by the woman is discarded. Tracking was performed using a very simple scheme. The same initial contour (small square) was used for each image and initially positioned wherever the final contour was in the preceding image. The parameters were set in the following manner: $\beta_1 = \beta_2$ in (1) and $\varepsilon_1 = \varepsilon_2 = .1$ in (7) and (8), respectively. Convincing results were obtained without involved considerations about the system dynamics.

3.2 Shape Priors Involving Objects of Different Types: "Yellow" and "Orange":

The goal of this section is to investigate the ability of our method to deal with objects of different shapes. To this end, we built a training set consisting of two words, "orange" and "yellow", each written using twenty different fonts. The size of the fonts was chosen to lead to words of roughly the same length. The obtained words (binary maps, see Figure 1) were then registered according to their centroid. No further effort such as matching the letters of the different words was pursued. The method presented in Section (2.1) was used to build the corresponding space of shapes for the registered binary maps.

We tested our framework on images where a corrupted version of either the word "orange" or "yellow" was present (Figure 5(a)). Word recognition is a very

challenging task and using geometric active contours to address it may not be a panacea. However, the ability of the level-set representation to naturally handle topological changes was found to be useful for this purpose: In the experiments presented below, evolution led the contour to split and merge a certain number of times to segment the disconnected letters of the different words.

In all the following experiments, we have chosen $\beta_1 = \beta_2$ in (1) and $\varepsilon_1 = \varepsilon_2 = .1$ in (7) and (8), respectively. The same initial contour was used for all the test images. Starting from the same initial contour our framework was able to accurately detect which word was present in the image. This highlights the ability of our method to gather image information throughout evolution and to distinguish between objects of different classes ("yellow" and "orange").

Experiment 1: In this experiment, one of the elements of the training set was used (Figure 1(2b)). A thick line (occlusion) was drawn on the word and a fair amount of gaussian noise was added to the resulting image (Figure 5(1a)). The result of applying our method is presented Figure 5(1d). Despite the noise and the occlusion, a reasonable segmentation is obtained. In particular, the correct font is detected and the thick line almost completely removed. In addition, the final result is smooth as compared to the result obtained without shape prior; see Figure 5(1b). Hence, using binary maps to represent shape priors can have valuable smoothing effects, when dealing with noisy images.

Experiment 2: In this second test, the word "yellow" was written using a *different* font from the ones used to build the training set (visual check was performed to ensure that the length of the word was comparable to the length of the words in the training set). In addition, a "linear shadowing" was used in the background, making the first letter "y" completely hidden. The letter "w" was also replaced by a grey square (Figure 5(2a)). The result of applying our framework is presented in Figure 5(2d). The word "yellow" is correctly segmented. In particular, the letters "y" and "w", were completely reconstructed. Comparing to the results obtained in Figure 5(2b) and (2c) obtained without prior knowledge of the shape or with shape prior built from SDFs, one can notice the effect of our shape prior model in constraining the contour evolution.

Experiment 3: In this experiment, the word "orange" was *handwritten* in capital letters roughly matching the size of letters of the components of the training set. The intensity of the letters was chosen to be rather close to some parts of the background. In addition, the word was blurred and smeared in a way that made its letters barely recognizable (Figure 5 (3a)). This type of blurring effect is often observed in medical images due to patient motion. This image is particularly difficult to segment, since the spacing between letters and the letters themselves are very irregular due to the combined effects of handwriting and blurring. Hence, mixing between classes (confusion between either "yellow" or "orange") can be expected in the final result. The final result of the segmentation process is presented Figure 5(3d). The word "orange" is not only recognized (no mixing) but satisfyingly recovered; in particular, a thick font was obtained to model the thick letters of the word.

4 Conclusion

Our contributions in this paper are twofold. First, we demonstrate that building a space of shapes by applying PCA to binary images (instead of signed distance functions) enables one to constrain the contour evolution in a way that is more faithful to the training set. Secondly, we present a novel region-based segmentation framework, able to separate regions of different intensities in an image. Shape knowledge and image information were encoded into two energies entirely described in terms of shapes. This consistent description allows for intuitive and simple equilibration between both image cues and shape prior.

The method presented allows for the simultaneous encoding of multiple types of shapes and seems to lead to robust segmentation results. In particular, our shape driven segmentation technique was able to cope with noise, clutter, partial occlusion, change in aspect and blurring, in a convincing manner.

In our future work, we aim at comparing segmentation performances obtained with shape priors built from linear and kernel PCA methods. Kernel PCA was proven to be a powerful method to describe data sets. The consistent shape approach, characteristic of our framework, is envisioned to be particularly suitable to deal with the "exotic" norms involved in kernel PCA methods.

(a) (b) (c) (d)

Fig. 3. Comparison of results. (a): base image and initial contour. (b): Segmentation result without shape prior. (c): Result with shape prior obtained from applying PCA on SDFs. (d): Result with shape prior obtained from applying PCA on Binary Maps.

References

1. Leventon, M., Grimson, E., Faugeras, O.: Statistical shape influence in geodesic active contours. In: Proc. CVPR, IEEE (2000) 1316–1324
2. Paragios, N., Deriche, R.: Geodesic active contours and level sets for the detection and tracking of moving objects. In: Transactions on Pattern analysis and Machine Intelligence. Volume 22. (2000) 266–280
3. Tsai, A., Yezzi, T., et al., W.W.: A shape-based approach to the segmentation of medical imagery using level sets. IEEE Trans. on Medical Imaging **22** (2003) 137–153

Fig. 4. Tracking results for the Swedish Couple Sequence, using a training set of 7 silhouettes of the man. First row: Tracking the man. Second row: Tracking the woman.

4. Yezzi, A., Kichenassamy, S., et al., A.K.: A geometric snake model for segmentation of medical imagery. In: IEEE Trans. Medical Imag. Volume 16. (1997) 199–209
5. A. Yezzi, S.S.: Deformotion: Deforming motion, shape average and the joint registration and approximation of structures in images. In: International Journal of Computer Vision. Volume 53. (2003) 153–167
6. Blake, A., Isard, M., eds.: Active Contours. Springer (1998)
7. Cootes, T., Taylor, C., et al., D.C.: Active shape models-their training and application. In: Comput.Vis. Image Understanding. Volume 61. (1995) 38–59
8. Wang, Y., Staib, L.: Boundary finding with correspondance using statistical shape models. In: IEEE Conf. Computer Vision and Pattern Recognition. (1998) 338–345
9. Cremers, D., Kohlberger, T., Schnoerr, C.: Diffusion snakes: introducing statistical shape knowledge into the mumford-shah functional. In: International journal of computer vision. (Volume 50.)
10. Cremers, D., Kohlberger, T., Schnoerr, C.: Shape statistics in kernel space for variational image segmentation. In: Pattern Recognition. Volume 36. (2003) 1292–1943
11. Mika, S., Scholkopf, B., et al., A.S.: Kernel pca and de-noising in feature spaces. In: Advances in neural information processing systems. Volume 11. (1998)
12. Sapiro, G.: Geometric Partial Differential Equations and Image Analysis. Cambridge University Press (2001)
13. Osher, S., Sethian, J.: Fronts propagation with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. Journal of Computational Physics **79** (1988) 12–49
14. Rousson, M., Paragios, N.: Shape priors for level set representations. In: Proceedings of European Conference on Computer Vision. (2002) 78–92
15. Yezzi, A., Tsai, A., Willsky, A.: A statistical approach to snakes for bimodal and trimodal imagery. In: Proc. Int. Conf. Computer Vision. Volume 2. (1999) 898–903
16. Chan, T., Vese, L.: Active contours without edges. In: IEEE Trans. Image Processing. Volume 10. (2001) 266–277

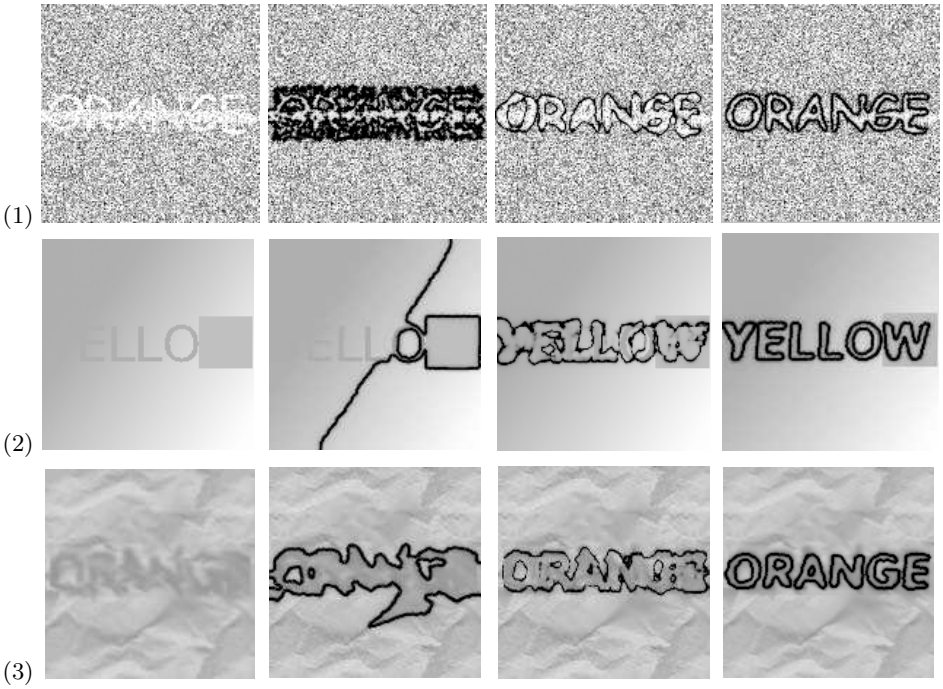

Fig. 5. Each row presents the segmentation results obtained for experiment 1, 2 and 3 respectively. (a): Initial image. (b): Segmentation result using E_{image} only (no shape prior). (c): Typical segmentation result with shape prior built from SDFs. (d): Segmentation result with shape prior built from binary map.