SIViP (2014) 8:789–798
DOI 10.1007/s11760-013-0599-y

ORIGINAL PAPER

Combining learning-based intensity distributions with nonparametric shape priors for image segmentation

Abdurrahim Soğanlı · Mustafa Gökhan Uzunbaş · Müjdat Çetin

Received: 1 August 2013 / Revised: 8 December 2013 / Accepted: 17 December 2013 / Published online: 30 December 2013

Abstract Integration of shape prior information into level set formulations has led to great improvements in image segmentation in the presence of missing information, occlusion, and noise. However, most shape-based segmentation techniques incorporate image intensity through simplistic data terms. A common underlying assumption of such data terms is that the foreground and the background regions in the image are homogeneous, i.e., intensities are piecewise constant or piecewise smooth. This situation makes integration of shape priors inefficient in the presence of intensity inhomogeneities. In this paper, we propose a new approach for combining information from shape priors with that from image intensities. More specifically, our approach uses shape priors learned by nonparametric density estimation and incorporates image intensity distributions learned in a supervised manner. Such a combination has not been used in previous work. Sample image patches are used to learn the intensity distributions, and segmented training shapes are used to learn the shape priors. We present an active contour algorithm that takes these learned densities into account for image segmentation. Our experiments on synthetic and real images demonstrate the robustness of the proposed approach to complicated intensity distributions, and occlusions, as well as the improvements it provides over existing methods.

A. Soğanlı (✉) · M. Çetin
Faculty of Engineering and Natural Sciences,
Sabanci University, 34956 Tuzla, Istanbul, Turkey
e-mail: soganli@sabanciuniv.edu

M. Çetin
e-mail: mcetin@sabanciuniv.edu

M. G. Uzunbaş
Computer Science Department, Rutgers University,
Piscataway, NJ 08854, USA
e-mail: uzunbas@cs.rutgers.edu

Keywords Image segmentation · Shape priors · Supervised learning · Level set methods

1 Introduction

Image segmentation is one of the oldest and most profound subjects in computer vision. The goal of segmentation is to partition an image into regions that might be useful for image interpretation. While a wide variety of ideas and frameworks have been explored for image segmentation, of particular interest in this paper are techniques based on active contours. Active contour-based segmentation methods can be grouped into two major categories: edge-based methods and region-based methods. Edge-based models [1,2] use gradient magnitude to identify and separate regions. However, these methods suffer from noise. Some preprocessing steps can be applied for denoising, but these operations weaken edges. Region-based models [3–8] which are preferred in the presence of noise and low-quality data generally assume intensities in the regions are piecewise smooth. This approach is succeeded by several methods such as global [3,6–10] and local [4,5] intensity fitting energy-based methods and information theoretic methods [6,7] that use histograms and probability distribution functions of the input image. Local intensity energy-based models alleviate sensitivity of region-based models to intensity inhomogeneities by integrating a Gaussian kernel into the global model that assigns more weight to intensities near the current estimate of the boundary. Recently, a new model that integrates a Gaussian mixture model with a level set method for natural image segmentation has been proposed [11].

Two other major problems that complicate the segmentation process are occlusion and missing information in the input image. These problems have motivated researchers to

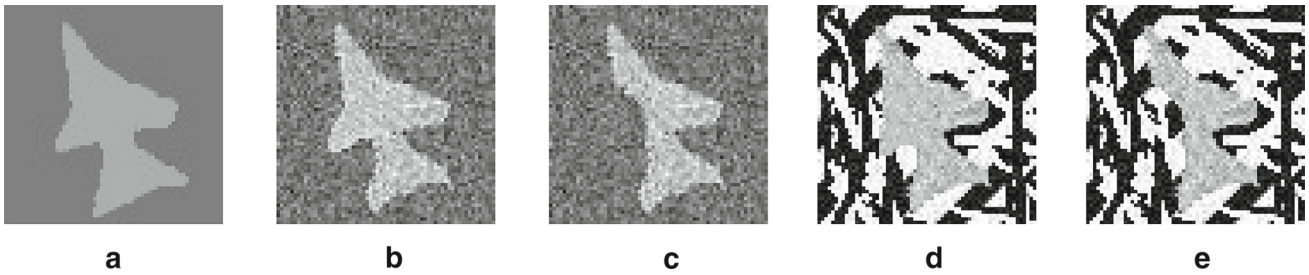

Fig. 1 Synthetically constructed aircraft object with 5 different conditions

use shape statistics in the segmentation process [12–14]. In these approaches, prior knowledge of the shape of the target object is incorporated as a regularization term into an optimization-based segmentation formulation. In [12–14] , the shape model is represented in terms of level sets. These methods assume intensities are piecewise constant and incorporate image data into the segmentation formulation accordingly, and they involve statistical constraints to keep the evolving boundary in the shape space defined by the training shapes. Earlier work [14] on this problem involves the use of principal component analysis (PCA) of the signed distance functions of training shapes in order to capture and represent shape variability. These techniques provide accurate segmentation in the presence of low SNR and missing information. However, PCA-based techniques have two disadvantages. First, these techniques treat the space of signed distance functions as a linear vector space, which it is not. Second, these techniques cannot deal with multimodal shape densities, which involve multiple shape subclasses within the overall shape distribution. Nonparametric shape priors are proposed in order to alleviate some of the problems faced by PCA-based shape modeling approaches. For example, in [12], nonparametric kernel density estimate of the shape distribution is used as the shape prior. In that work, L2 distance and template overlap ratios are used as metrics for the Gaussian-type kernel. In these segmentation frameworks, for the data (region) term, a term based on piecewise constant [3] or piecewise smooth [15] regions was adopted besides the shape term. In segmentation scenarios presented in these works, the shape term regularizes the data term so that the active contours are evolved to the actual boundary of the target object when the object of interest is partially occluded or suffers from noise. However, if the data term falls beyond a certain limit, the shape term may not be able to drive the curve to the correct boundary.

Our key observation that has motivated the work presented in this paper is that powerful shape-driven methods (such as nonparametric shape priors) have so far been used together with only simple data terms involving piecewise constant or piecewise smooth intensity assumptions. There is a need to combine such shape prior methods with flexible, learning-based intensity distributions, such as those proposed in [9,10]. In many real-life scenarios, in the presence of inhomogeneities or characteristic complicated intensity distributions in the regions, existing segmentation methods may fail to capture the target object accurately. Performance of these models can be better explained with Fig. 1. A synthetically constructed aircraft image and different backgrounds are shown in five sample images.

- a. Foreground and background of the object are homogeneous. Segmentation can easily be done by simple edge-based models. Region-based models can also provide accurate segmentation results.
- b. Foreground and background of the object are still homogeneous but there is noise. Edge-based models may have problems. Segmentation can be done by region -based models.
- c. Foreground and background of the object are homogeneous and noisy, and some part of the object is missing. Segmentation can be done by shape-based models [12,13].
- d. Foreground is homogeneous, but background of the object is not homogeneous. There are strong edges in the background, which are not regional boundaries but inner edges. Edge-based models and piecewise smoothness-based region models would fail. Segmentation can be done by using histogram-based models.
- e. The background is not homogeneous, and some part of the object is missing. Shape-based models can be used. However, piecewise smoothness-based data terms commonly used in these models cannot prevent leakages and would prevent the effective exploitation of the shape prior information.

In this paper, our motivation is to present an approach that can handle the types of complication represented by the toy example in Fig. 1e. Classical shape-based models generally fail to provide successful segmentation when faced by complicated intensity distributions because of the simplistic data terms they use. The idea for a remedy, which has motivated the work presented in this paper, could be to combine a data term based on learned intensity distributions with a powerful and versatile approach for incorporating shape priors. Here, the data term proposed in [9] would be a promising

candidate, which is based on the assumption that probability densities of the background and the foreground regions are a priori known. In this work, we propose to combine such a learning-based data term with a nonparametric shape model term to extend the previously proposed shape-based segmentation approach [12] to the case of complicated intensity distributions. We propose an energy functional that incorporates these two pieces, and develop an algorithm for minimizing that energy functional for segmentation. We propose two versions of the data term, one involving learning of both the foreground and the background intensity distributions, and the other considering the case in which only information about the foreground intensities is available. We observe that the proposed approach can handle both complicated intensity variabilities and complicated shape variabilities, a capability that existing methods do not readily exhibit.

This paper is organized as follows. We review the shape-based model proposed in [12] and the data term proposed in [3] in Sect. 2. Next, we explain our proposed model and its extension to color images in Sect. 3. In Sect. 4, we test our proposed approach on both synthetic images and real-world color images. In Sect. 5, we summarize and conclude our paper.

2 Background

The general shape-based segmentation framework that has been used in [12,13] and that we also consider in our work is based on minimizing energy functionals of the following form:

$$E(C) = -\log p(\text{data}|C) - \log p_C(C) \tag{1}$$

In this formulation, C is the segmentation curve, the first term is the data term, and the second one is the shape term capturing statistical prior information about the shape of the object to be segmented. Specific choices for the harmony of these two terms determine the accuracy of segmentation. In this section, we briefly describe the data term used by the shape-based segmentation methods in [12,13]. This data term imposes piecewise constant region intensities and is to be contrasted with the term we will propose in Sect. 3. In this section, we also briefly review nonparametric shape priors [12], which is the shape term we use in our work.

2.1 Data term

The region-based data term proposed in [3] is used in several shape-based methods. Letting I be the grayscale input image, the energy functional of data term is as follows:

$$E_{\text{CV}}(C) = \lambda_1 \int_{\text{inside}(C)} (I(x) - m_{\text{in}})^2 \mathrm{d}x + \lambda_2 \int_{\text{outside}(C)} (I(x) - m_{\text{out}})^2 \mathrm{d}x \tag{2}$$

where λ_1, λ_2 are constants, m_{in} and m_{out} are the mean intensities inside and outside of the current contour, respectively.[1] This term can be adopted to color images as proposed in [16].

2.2 Nonparametric shape priors

Shape-based segmentation methods of interest in this work incorporate shape prior information into a level set-based energy functional as an additional term. In [12], nonparametric shape prior is introduced to capture complex, potentially multimodal shape prior densities. This goes beyond the simpler PCA-based methods, which can only capture "unimodal" shape variability concentrated around a mean shape. The training set consists of n segmented unaligned curves $C_1, \ldots, C_n$ of the target object. These curves are first aligned with respect to translation, scaling, and rotation parameters, so that their remaining variability captures the distribution of the object shape. Aligned curves, $\tilde{C}_1, \ldots, \tilde{C}_n$ are used to learn a shape probability density function. In particular, Parzen density estimation is used within a level set-based formulation as follows:

$$E_{\text{Shape}}(C) = \frac{1}{n} \sum_m k(d_{L_2}(\phi_{\tilde{C}}, \phi_{\tilde{C}_m}), \sigma) \tag{3}$$

where $\phi_{\tilde{C}}$ and $\phi_{\tilde{C}_m}$ are the signed distance functions of the current contour and the training shapes, respectively. Note that signed distance functions for level set formulations encode the distance of any particular point in the image domain to the segmenting boundary, with negative sign for points inside the boundary and with positive sign for points outside the boundary. Through this term, the active contour is constrained by a shape force governed by the training set. Energy formulation of the shape-based approach of [12] can be expressed as:

$$E(C) = E_{\text{CV}}(C) + \beta E_{\text{Shape}}(C) \tag{4}$$

where β is a hyper-parameter that tries to balance data and shape terms. Segmentation is done through minimization of this energy functional by gradient descent. This iterative procedure produces the evolution of the curve from initialization to final segmentation.

This shape-based model provides accurate segmentation results in the presence of missing information or occlusions if the intensity homogeneity assumption of the data term is

[1] The subscript "CV" is used to refer to the first letters of the last names of the authors of [3].

correct for the input image. However, generally, real-world complex images are not homogeneous, and intensities are not slowly varying. Therefore, in the next section, we propose changing the data term in order to obtain accurate segmentation results in the presence of complicated intensity distributions.

3 Proposed approach

We consider the problem of segmentation of an object with an arbitrary intensity distribution and arbitrary shape variability, possibly in a complex background. We are also interested in handling occlusions and missing information. Our model is based on two assumptions to help address this challenging problem:

1. Example contours of the object to be segmented are available for a priori learning of the shape distribution. Note that this assumption is the same with [12] and other shape-based models.
2. Image patches are available to learn the probability density functions of foreground and background (or at least just the foreground) of the target object.

3.1 Foreground and background distributions are known (Model 1)

By using the second assumption above, we can insert a probability density function-based data term to the energy functional in (1). In [10] and later in [9], a data term that is suitable for our objectives was proposed. This term is given by:

$$E_{\mathrm{PD}}(C) = -\int\limits_{\mathrm{outside}(C)} \log p_{\mathrm{out}}(I(x))\mathrm{d}x - \int\limits_{\mathrm{inside}(C)} \log p_{\mathrm{in}}(I(x))\mathrm{d}x \tag{5}$$

In this equation, $p_{\mathrm{in}}(.)$ and $p_{\mathrm{out}}(.)$ are the intensity probability density functions that belong to the foreground and the background regions in the image, respectively. These probability density functions can either be estimated based on foreground and background patches extracted from the test image (if that is feasible in the particular application of interest), or from offline training samples of the type of image to be segmented. One can use parametric or nonparametric density estimation methods for estimating these intensity distributions. In our work presented here, we use nonparametric density estimation on patches extracted from the test image to be segmented. Gradient flow of this data term is given below:

$$\frac{\partial C}{\partial t} = [\log(p_{\mathrm{out}}(I(x))) - \log(p_{\mathrm{in}}(I(x))]\mathbf{N} = \left(\log\frac{p_{\mathrm{out}}(I(x))}{p_{\mathrm{in}}(I(x))}\right)\mathbf{N} \tag{6}$$

$\mathbf{N}$ is inward normal of the segmenting curve. Interpreting this gradient flow can provide some intuition on the behavior of this data term. Assume a pixel at location x in the input image.

- If $p_{\mathrm{in}}(x) > p_{\mathrm{out}}(x)$, the expression before the normal will be negative, which will cause the curve to move outward to include this point in the foreground region.
- If $p_{\mathrm{in}}(x) < p_{\mathrm{out}}(x)$, the expression before the normal will be positive, which will cause the curve to move inward to push this point to the background region.

This model can also be applied on color images by extending data term as in [16]. For an RGB image, forces from the three channels can be calculated separately and summed up. More generally, given a k-channel input image I, the gradient flow of the data term will be as follows:

$$\frac{\partial C}{\partial t} = \left[\sum_{i=1}^{k}\log(p_{\mathrm{out}_i}(I(x))) - \sum_{i=1}^{k}\log(p_{\mathrm{in}_i}(I(x)))\right]\mathbf{N} = \left(\sum_{i=1}^{k}\frac{\log p_{\mathrm{out}_i}(I(x))}{\log p_{\mathrm{in}_i}(I(x))}\right)\mathbf{N} \tag{7}$$

Using this data term based on learning intensity distributions, our energy functional to be used for segmentation becomes:

$$E(C) = E_{\mathrm{PD}}(C) + \beta E_{\mathrm{Shape}}(C) \tag{8}$$

To obtain an iterative curve evolution algorithm to minimize (8), we combine the gradient flow in (6) or (7) with the gradient flow expression for nonparametric shape priors, which can be found in [12].

In this subsection, we assumed the availability of data for learning both the foreground and the background intensities. However, one might also be interested in segmenting an object with a particular intensity distribution on a variety of backgrounds. In that case, one needs an approach that can exploit the foreground intensity distribution without any knowledge about the background. We handle that case in the next subsection.

3.2 Only the foreground distribution is known (Model 2)

We construct a new model in which background probability distribution function is unknown. In this case, a simple idea could be to treat the background as uniformly distributed. Thus, $p_{\mathrm{out}}(I(x))$ becomes nothing but constant with following value:

$$p_{\text{out}} = \frac{1}{L} \tag{9}$$

where L is maximum value of pixel intensities. The gradient flow of the data term under this intensity model becomes:

$$\frac{\partial C}{\partial t} = \left(\log \frac{1/L}{p_{\text{in}}(I(x))}\right)\mathbf{N} \tag{10}$$

This term is a simple comparison between $\frac{1}{L}$ and $p_{\text{in}}(I(x))$. If $p_{\text{in}}(I(x))$ is bigger than $\frac{1}{L}$, then the contour moves to include this point in the foreground. This model can produce successful segmentation results under complex and approximately uniformly distributed backgrounds. The overall gradient flow is given by the linear combination of the flow in (6) due to the data term and the flow resulting from the nonparametric shape priors presented in [12].

The algorithmic structure we use to implement the curve evolution associated with the gradient flows presented in this and the previous subsections in combination with the flow due to the shape term is similar to the structure used in [12]. In particular, we first evolve the curve C using only the data term until convergence and then switch the shape term. This prevents unnecessary evaluations of the nonparametric shape density for the initial and early states of the curve, which are usually very far away from the actual object shape anyway. When the shape term is turned on, both the data and the shape gradient flows act on the curve at each iteration. Before computing the shape force, the current segmenting curve is aligned with respect to the training shapes. Once the shape force is computed, the curve C is updated through the data and the shape forces. This procedure is repeated until the curve converges.

4 Experimental results

We now demonstrate segmentation results of our proposed approach. We first consider segmentation of a synthetically constructed aircraft image, which is similar to the image in Fig. 1e. The aligned training shapes used in nonparametric shape density estimation are shown in Fig. 2. Figure 3 shows the segmentation results on a test image of the aircraft object not included in the training set of shapes. We present results of both versions of our approach. The first version involves learning and using both the foreground and the background intensity distributions (Model 1). The second version learns and uses the foreground intensities, but assumes a uniform density for the background (Model 2). We compare our results with those of [12]. Images in the first column are (a-f-k) initial contours. Second through fourth columns are intermediate states in the evolution process, and the fifth column shows the final segmentation results. Up to the third column, only the data term is activated in each method. So the effect of the shape term can be observed in the last two columns. We observe that the method in [8] fails to segment this image. This is because the simple data term in [12] is not able to capture the complicated foreground and background intensity distributions in this scene. As shown in Fig. 3c, the data term of [12] drives the curve to high-contrast areas in the scene, which do not correspond to the boundary of interest here. The shape prior cannot do much more than trying to fit the best aircraft shape consistent with the training data as well as with the boundary favored by the data term. The final result in Fig. 3e looks like an aircraft, but is a poor segmentation of the object in the scene. The second row presents the results of our Model 1. We observe that the result in Fig. 3j is a successful segmentation. The intermediate result in Fig. 3h demonstrates the effectiveness of the data term used in our approach in handling the complicated intensity distributions in this image. Further steps in the evolution, as shown in Fig. 3i, j, help recover the missing wing of the aircraft through effective incorporation of shape information. The bottom row contains the result of our Model 2. Despite the lack of prior knowledge about the background, this approach is still able to produce a reasonable segmentation result. However, as compared to the results of Model 1, Model 2 produces some artifacts. We can make sense of these artifacts by examining the intensity probability density functions used by the two models, shown in Fig. 4. In particular, the characteristic nature of the background intensity distribution which is not captured effectively by the uniform distribution is that very dark and very bright intensities are highly likely in the background region. Since the uniform distribution does not capture this nature, some of the dark and bright pixels that are not terribly inconsistent with the shape prior are put into the foreground region by the approach based on Model 2.

Next, we present results on real color images. In the first such example, we consider segmentation of the sun in the presence of occlusion. In addition to the potential presence of an occlusions, this problem can also be challenging because of scattering around the sun. The test image and the segmentation results are shown in Fig. 5. While one might argue that the high contrast between the sun and the back-

Fig. 2 Aligned training samples for the aircraft object

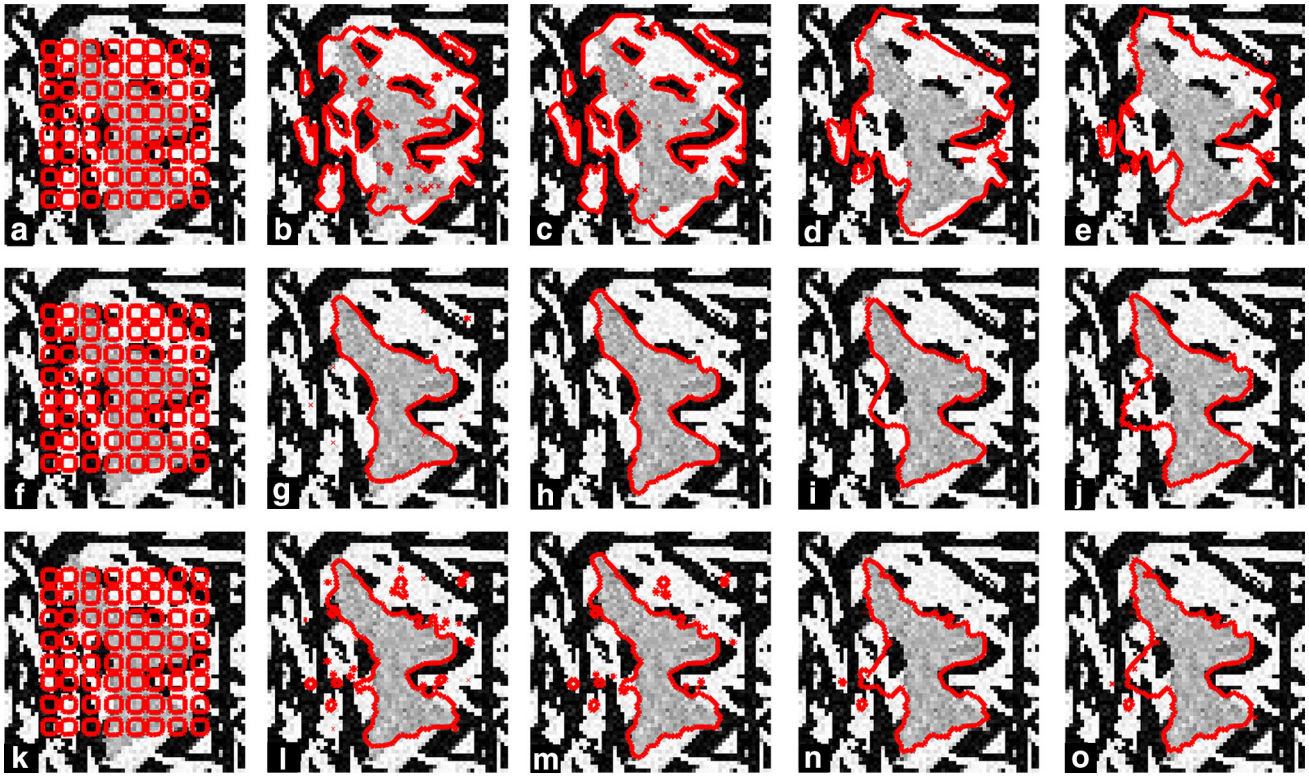

Fig. 3 Segmentation of the aircraft image. *First row* (**a**–**e**) shows the segmentation result of [12]. *Second row* (**f**–**j**) shows the segmentation result of proposed model 1, which is based on learning and using the foreground and the background intensity distributions. *Third row* (**k**–**o**) shows the result of proposed model 2 in which the background is treated as uniformly distributed. *Leftmost and rightmost columns* show the initial curve and the final segmentation result, respectively, whereas the *middle columns* show intermediate states in the curve evolution process

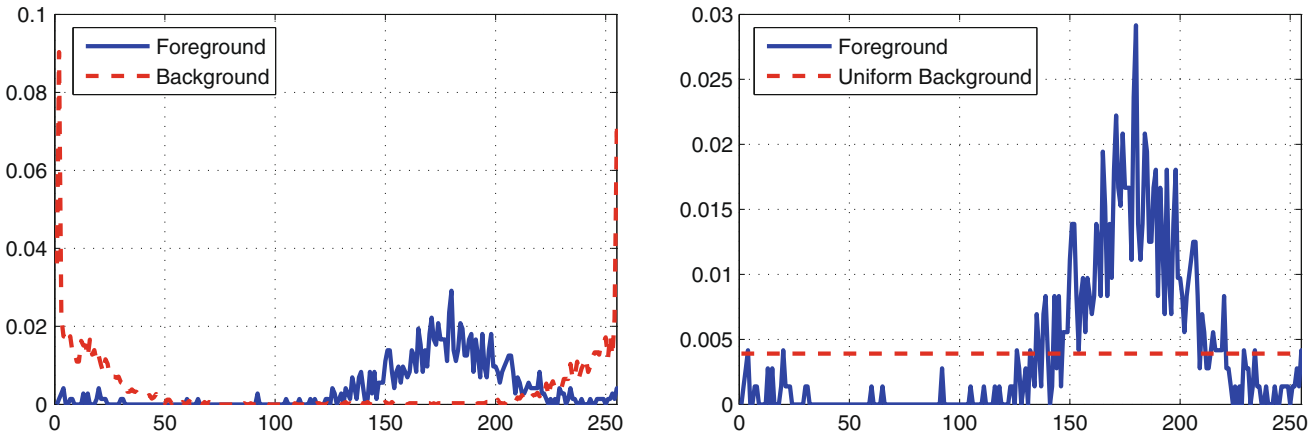

Fig. 4 Probability density functions of foreground and background of aircraft object used in our segmentation approach. Model 1 (*left*), Model 2 (*right*)

ground would enable a simpler segmentation approach to generate a reasonable result, this, of course, is not valid when we have complications such as occlusions, as seen in this example. Prior probability densities are obtained by labeling small foreground and background regions of the input image by supervised learning. In Fig. 5, the first three columns show segmentation results without the shape term, which is added after the third column. Our proposed Model 1 provides a reasonably accurate segmentation result. Model 2 leads to correct localization of the sun but has difficulty recovering all of the occluded region. The model proposed in [12] fails to segment the sun. Segmentation result converges to a cir-

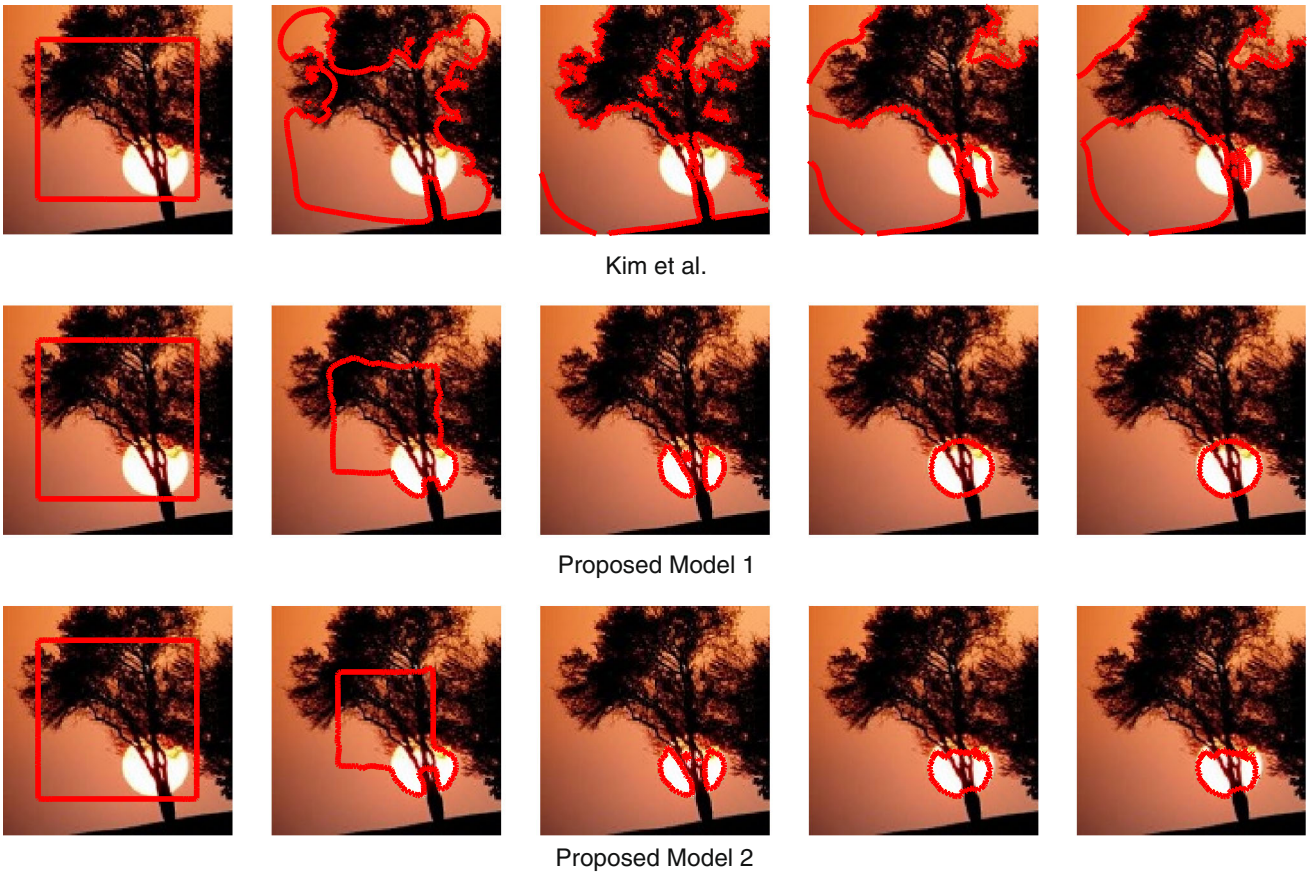

Fig. 5 Segmentation of a circular object. *First row* shows the segmentation result of [12]. *Second row* shows the segmentation result of proposed model 1. *Third row* shows the result of proposed model 2. *Leftmost and rightmost columns* show the initial curve and the final segmentation result, respectively, whereas the *middle columns* show intermediate states in the curve evolution process

Fig. 6 Training samples for birds

Fig. 7 Pieces involved in the bird image segmentation experiment: 7 foreground training samples, 2 background training samples, input image

cular shape because of the activation of the shape term, but the resulting boundary does not correspond to the boundary of the sun.

In the third example, we consider segmentation of images of birds. Shapes of randomly chosen bird species are used in learning the shape priors. The shape training samples used in our experiment are shown in Fig. 6. Although a comprehensive shape prior database would require a significant number of training samples from different bird species, we use 11 training samples for simplicity and show that even such limited data could be very valuable when effectively used in the segmentation process. In this example, prior intensity distributions are not obtained directly from the input image itself. We use distinct sample background and foreground images. In particular, we learn the foreground intensity distribution from 7 different image patches containing birds belonging to the same species as the bird in the input image and learn the background intensity distribution from 2 image patches. Training patches used for learning the foreground and the background intensity distributions are shown in Fig. 7. Final estimated foreground and background intensity distributions are the average distributions of these training patches. One

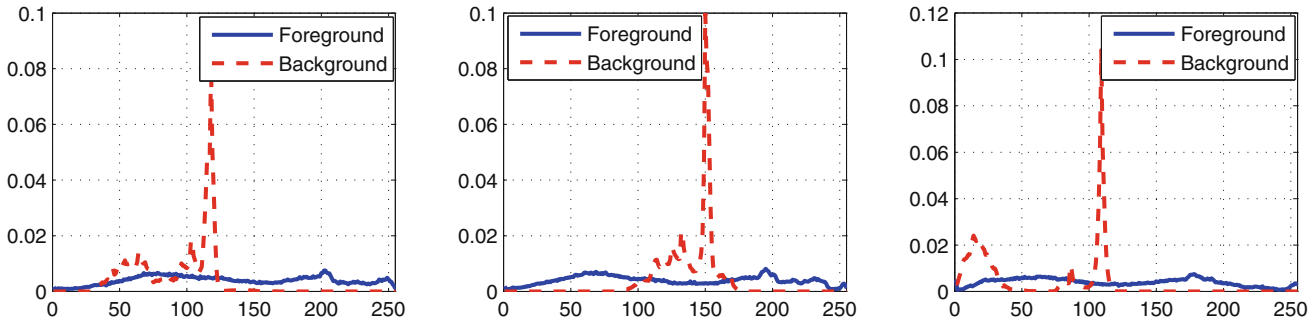

Fig. 8 Prior distributions of foreground and background for the three channels

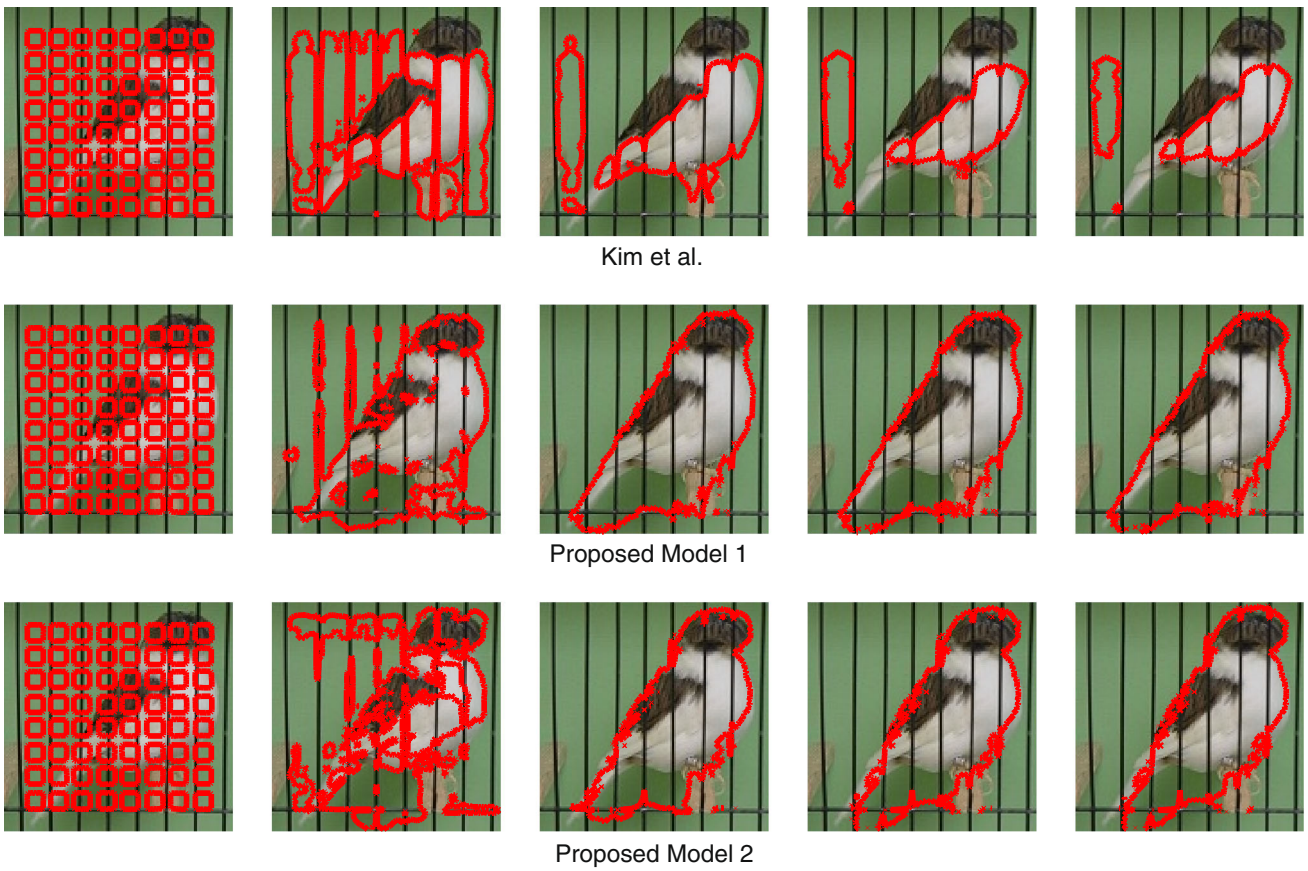

Fig. 9 Segmentation of a bird. *First row* shows the segmentation result of [12]. *Second row* shows the segmentation result of proposed model 1. *Third row* shows the result of proposed model 2. *Leftmost and rightmost columns* show the initial curve and the final segmentation result, respectively, whereas the *middle columns* show intermediate states in the curve evolution process

can also use a weighted average in order to obtain better estimated distributions for a specific scenario. Note that this corresponds to a realistic scenario in which we do not have access to a patch of intensities of the particular object in the input image to be segmented. The estimated probability density functions of the foreground and the background regions for each of the three color channels are shown in Fig. 8. In this example, we use an input image containing a bird inside a birdcage, which causes occlusions as shown in Fig. 7. Most segmentation algorithms would be challenged by such occlusions. Segmentation results are shown in Fig. 9. We observe that the approach in [12] is able to segment only the homogeneous part of the bird. Although activation of the shape term forces the contour to approach the shape of a bird, the piecewise smooth intensity assumption of the data term imposes a strong preference to exclude both the dark and the light parts of the bird inside the contour. Model 1 produces a fairly good result. We observe that the part of the back-

Fig. 10 Segmentation of a car. *First row* shows the segmentation result of [12]. *Second row* shows the segmentation result of proposed model 1. *Third row* shows the result of proposed model 2. *Leftmost and rightmost columns* show the initial curve and the final segmentation result, respectively, whereas the *middle columns* show intermediate states in the curve evolution process

ground that contains the shadow of the bird is included in the segmented foreground region. This is most likely caused by the fact that the darkening of the intensities due to the shadow makes this region unlikely for the background given the learned prior distribution. Nevertheless, considering the challenging nature of this problem, our approach based on Model 1 produces reasonable segmentation results. Model 2 also produces a satisfactory result in this example despite some small mismatches around the boundaries. This can be explained by the mismatch of intensity distributions used in the model with those observed in the input image. In particular, a uniform distribution which is significantly different from the learned distribution shown in Fig. 8 is assumed for the background. Similarly, the foreground intensity distribution learned from 7 image patches on the left of Fig. 7 is a coarse approximation of the actual intensity distribution in the test image. When both of these distributions are inaccurate, some level of imperfection is expected.

Our final example involves the segmentation of car images. Cars can have a variety of color distributions. Transparency of the windows cause the background to be visible through the car. Such complications make the segmentation of cars a challenge. Exploiting prior intensity distribution knowledge on foreground and background may lead to better segmentation results compared to the piecewise smooth assumption in such cases. Segmentation results are shown in Fig. 10. The shape-based model proposed in [12] fails to produce an accurate segmentation result. The data term used in [12] assumes that regions are piecewise smooth, which in this example causes some parts of the car to be treated as background. The final result approaches the shape of a car, but fails to localize the actual car in the scene. Our proposed Model 1 provides significantly better segmentation results, although some parts of the car having a similar intensity distribution to the background are missed. Our Model 2 segments most of the car successfully, but the uniform background model causes some artifacts.

5 Conclusions

In this paper, we have proposed a segmentation approach that involves learning and exploitation of both the intensity distributions of the regions in the image and the shape dis-

tributions of the objects to be segmented. We have learned the shape distributions by using the recently developed idea of nonparametric shape priors. The main contribution of our work is to combine such shape priors with learning-based intensity models in the segmentation process. We have proposed two models, one involving the learning of both the foreground and the background intensity distributions, and the other involving the use of a learning-based foreground density together with an assumed uniform background density. We have shown that our approach can provide improvements over existing shape prior-based segmentation methods, especially when the region intensity distributions are complicated. We have compared the performance of our two models on synthetic and real examples, and we have also shown examples of cases where our approach can fail when the distributions used in the model are inaccurate. We have also demonstrated the robustness of our proposed approach to occlusions. When data to learn the shape and the intensity distributions are available, the proposed approach has the potential to solve very challenging segmentation problems.

Acknowledgments This work was partially supported by the Scientific and Technological Research Council of Turkey (TUBITAK) through a graduate fellowship.

References

1. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: active contour models. Int. J. Comput. Vis. **1**(4), 321–331 (1988)
2. Xu, C., Prince, J.: Snakes, shapes, and gradient vector flow. IEEE Trans. Image Process. **7**(3), 359–369 (1998)
3. Chan, T.F., Vese, L.A.: Active contours without edges. IEEE Trans. Image Process. **10**(2), 266–277 (2001)
4. Li, C., Huang, R., Ding, Z., Gatenby, J.C., Metaxas, D.: A level set method for image segmentation in the presence of intensity inhomogeneities with application to MRI. IEEE Trans. Image Process. **20**(7), 2007–2016 (2011)
5. Zhang, K., Song, H., Zhang, L.: Active contours driven by local image fitting energy. Pattern Recognit. **43**, 1199–1206 (2010)
6. Kim, J., Fisher, J.W., Yezzi, A., Cetin, M., Willsky, A.S.: A nonparametric statistical method for image segmentation using information theory and curve evolution. IEEE Trans. Image Process. **14**, 1486–1502 (2005)
7. Ni, K., Bresson, X., Chan, T., Esedoglu, S.: Local histogram based segmentation using the Wasserstein distance. Int. J. Comput. Vis. **84**, 97–111 (2009)
8. Mumford, M., Shah, J.: Optimal approximation by piecewise smooth functions and associated variational problems. Commun. Pure Appl. Math. **42**, 577–685 (1989)
9. Paragios, N., Deriche, R.: Geodesic active regions and level set methods for supervised texture segmentation. Int. J. Comput. Vis. **46**(3), 223–247 (2002). and 1996
10. Zhu, S., Yuille, A.: Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. **18**, 884900 (1996)
11. Xie, Z., Wang, S., Hu, D.: New insight at level set & Gaussian mixture model for natural image segmentation. Signal Image Video Process. **7**(3), 521–536 (2013)
12. Kim, J., Cetin, M., Willsky, A.S.: Nonparametric shape priors for active contour-based image segmentation. Signal Process. **87**, 3012–3044 (2007)
13. Cremers, D., Soatto, S.: Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. Int. J. Comput. Vis. **69**(3), 335–351 (2006)
14. Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W.E., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. IEEE Trans. Med. Image Imaging **22**(2), 137154 (2003)
15. Cremers, D., Tischhauser, F., Weickert, J., Schnorr, C.: Diffusion snakes: introducing statistical shape knowledge into the Mumford–Shah functional. Int. J. Comput. Vis. **50**(3), 0920–5691 (2002)
16. Tsai, A., Yezzi Jr, A., Willsky, A.S.: Curve evolution implementation of the Mumford–Shah functional for image segmentation, denoising, interpolation, and magnification. IEEE Trans. Image Process. **10**(8), 1169–1186 (2001)