

Shadow Segmentation in SAS and SAR Using Bayesian Elastic Contours

Darshan Bryner and Anuj Srivastava
 Department of Statistics
 Florida State University, Tallahassee, FL
 {dbryner, anuj}@stat.fsu.edu

Abstract

We present a variational framework for naturally incorporating prior shape knowledge in guidance of active contours for boundary extraction in images. This framework is especially suitable for images collected outside the visible spectrum, where boundary estimation is difficult due to low contrast, low resolution, and presence of noise and clutter. Accordingly, we illustrate this approach using the segmentation of synthetic aperture sonar (SAS) and synthetic aperture radar (SAR) images. The shadows produced from these imaging modalities often times offer more consistent pixel values with clearer contrast to the background than the targets pixels themselves, and thus we focus on the extraction of shadow boundaries rather than target boundaries. Since shadow shapes can vary under approximately affine transformation with different target range and aspect angle, we incorporate an affine-invariant, elastic shape prior based on the shape analysis techniques developed in [2] to the active contour model. We show experimental results on both a simulated SAS and a simulated SAR image database in three segmentation scenarios: without shape prior, with similarity-invariant shape prior, and with affine-invariant shape prior.

1. Introduction

An object of interest in an image can be characterized to some extent by the shape of its external boundary. It is therefore important to develop procedures for boundary extraction in problems of detection, tracking, and classification of objects in images. Active contour algorithms have become an important tool in image segmentation for object detection [5, 7]. As segmentation algorithms become more sophisticated, they are tested in more difficult imaging environments of real-world scenarios where images either do not have enough contrast to provide sharp boundaries, some occlusion of the target occurs, or there exists target-like clutter or noise. One example of this scenario is when images are collected in a spectrum outside the visible

domain. Thus, it is of increasing importance that boundary extraction algorithms make use of prior knowledge about the expected target class in order to help compensate for the lack of clear data. This is accomplished by influencing the contour evolution in part with a *shape prior*, a statistical model derived from a set of known training shapes, in a *Bayesian active contour* approach [4, 9, 7].

1.1. Past Work on Prior-Driven Active Contours

There are two broad categories of active contour methods: *parametric* methods that evolve an explicitly defined parameterized curve, and *geometric* methods that evolve implicitly defined zero-level sets of higher-dimensional functions. Therefore, in a parametric model the shape prior is a statistical model on closed or open contours in \mathbb{R}^2 ; whereas, in a geometric model the shape prior is a statistical model on signed-distance functions in \mathbb{R}^3 or higher. Past Bayesian active contour methods, i.e. those that incorporate shape prior, have been applied almost exclusively to geometric models where most efforts follow the ideas presented in Leventon *et al.* [7] and improvements thereafter [9, 6]. The major limitation to these geometric Bayesian approaches is that they cannot apply state-of-the-art statistical shape models to form a shape prior. In order to apply more sophisticated and computationally efficient shape analysis techniques to our Bayesian active contour model, we select a parametric approach.

There have been a few Bayesian active contour models in the past that take a parametric approach. Early approaches make use of “landmark-based” shape analysis [3] to impose a shape prior. While this method involves formal shape spaces and their geometries, it is not ideally suited to active contours since it was developed primarily for landmarks denoting salient points on a shape. Joshi *et al.* [4] use ideas from elastic shape analysis of planar curves to create a shape prior from an intrinsic density on shape space to help evolve a curve in a parametric active contour model. Although the method presented in [4] incorporates intrinsic shape statistics, it uses an older shape representation, and is only limited to similarity-invariant shape analysis. *We are unaware*

of the existence of any Bayesian active contour model that incorporates affine-invariant shape statistics of parametric curves.

1.2. Our Approach

Our goal is to use techniques from elastic shape analysis to develop a method to incorporate prior shape information into an active contour model especially suited for segmentation of images beyond the visible spectrum. Using the theory and statistical tools developed in [11] for similarity-invariant shapes and in [2] for affine-invariant shapes, we formulate a prior shape model on both shape spaces, denoted \mathcal{S}_{sim} and \mathcal{S}_{aff} respectively. The variational framework for contour evolution remains the same as that of earlier methods, except that the shape prior term now comes from either of the two elastic shape models. The main contributions of this paper are: (1) develop a similarity-invariant and an affine-invariant prior shape model for planar curves, (2) incorporate these statistical shape models in driving Bayesian active contours, and (3) demonstrate this framework using segmentation of shadow boundaries in SAS and SAR images.

The organization of the rest of the paper is as follows. Section 2 provides an overview of the synthetic aperture imaging technique, Section 3 outlines our Bayesian active contour model, Section 4 shows experimental results on SAS/SAR datasets, and Section 5 is the conclusion.

2. Difficulties of Boundary Extraction in Synthetic Aperture Imaging

The particular application of interest to us here is boundary extraction in the synthetic aperture imaging modalities such as SAS and SAR, although our procedure can be applied to a wide variety of beyond visible spectrum applications such as medical diagnosis and infrared surveillance. Synthetic aperture imagery is created in the following manner. An autonomous vehicle travels in a straight line and emits a sound chirp (in the case of sonar) or an electromagnetic chirp (in the case of radar) at regular time intervals. The emitted chirp is designed to cover a predetermined range of frequencies suitable for the imaging environment. Each chirp in the series reflects off the sea floor (in the case of sonar) or the ground (in the case of radar), and the system detects and stores these return signals that contain information describing the environment. A single aperture sonar or radar would form one image from one return signal, but a synthetic aperture imaging system uses sophisticated post-processing algorithms to combine the raw data from many return signals into one single image. Although this technique allows for a much larger imaging footprint with greater contrast and a finer resolution than a standard sonar or radar, synthetic aperture images are still notoriously dif-

icult for standard boundary detection algorithms to achieve high performance.

The task of automatically extracting *target* boundaries in synthetic aperture imagery is challenging due to the following reasons. (1) Even though a synthetic aperture system improves on standard aperture systems, the images can still exhibit low contrast. That is, some target pixels can be quite similar in intensity level to some background pixels, and thus boundaries are not clear. (2) Since the SAS and SAR imagery used here comes from a side scan system, one target side faces away from the sonar or radar. The target is therefore partially occluded in shadow, and its shape can vary widely with its aspect angle. (3) The resolution of these images is often much lower than those obtained in the visible spectrum, resulting in relatively fewer pixels on targets. (4) The imaging environments normally contain high speckle clutter due to rough backgrounds.

By focusing instead on segmenting the boundaries of the target *shadows* rather than the targets themselves, many of the above difficulties, although not completely eliminated, are alleviated to some extent. In regards to difficulty (1), shadow signatures are always consistently dark, and there is often clearer contrast between shadow and background than between target and background. The severity of difficulty (2) is diminished since there is no partial occlusion of the shadow like there is with the target. Difficulty (3) is reduced because shadows are typically larger and contain more pixels than the targets themselves. Lastly, difficulty (4) is not reduced, as background speckle noise remains the major difficulty in segmenting target shadows. Fig. 2 shows a few examples of SAS (top row) and SAR (bottom row) imagery. The SAS imagery is of an underwater cylinder target at varying aspect angles and ranges, and the SAR imagery is of a ground-based armored vehicle at varying aspect angles.




Figure 1. Synthetic aperture images of a target at different aspect angles and ranges. Top row: SAS images of a cylinder. Bottom row: SAR images of an armored vehicle.

3. Active Contour Model

We pose the problem of boundary extraction as a MAP estimation [4]. In this formulation we seek the closed, parameterized curve $\hat{\beta}(t)$ that minimizes an energy functional in the following manner: $\hat{\beta} = \operatorname{argmin}_{\beta \in \mathbb{L}^2(\mathbb{S}^1, \mathbb{R}^2)} E_{total}(\beta)$, where E_{total} is defined as

$$E_{total}(\beta) = \lambda_1 E_{image}(\beta) + \lambda_2 E_{smooth}(\beta) + \lambda_3 E_{prior}(\beta). \quad (1)$$

This total energy functional is the summation of three separate energy functionals weighted by the coefficients λ_i . The image energy E_{image} is defined solely from the pixel data in the image, the smoothing energy E_{smooth} is a penalty that maintains a certain amount of smoothness in the active contour, and E_{prior} is the prior shape energy defined from a probability density on shape space. In order to minimize this expression, we introduce a time variable s and seek a solution of the differential equation

$$\begin{aligned} \frac{\partial}{\partial s} \beta(t, s) &= -\lambda_1 \nabla E_{image}(\beta)(t, s) \\ &\quad - \lambda_2 \nabla E_{smooth}(\beta)(t, s) - \lambda_3 \nabla E_{prior}(\beta)(t, s). \end{aligned} \quad (2)$$

That is, we search for a local minimization of Eqn. 1 via gradient descent. The weights λ_i greatly affect the performance of the contour evolution, and appropriate values can vary widely from application to application. For this reason we leave their selection up to the discretion of the user. In the paragraphs that follow, we summarize our formulation of each of the three energy terms.

Image Energy: Our image energy term is based on the work presented in [12, 9]. In these papers, the authors propose an image energy that is not based on the gradient of a Gaussian smoothed edge map, as in for example EPGVF in [8], but rather instead a region-based energy term that is more robust to noise and diffuse object boundaries. Here, one obtains *a priori* two estimated probability densities, f and g , of pixel values in a region containing only shadow pixels and a region containing only background pixels, respectively. These two prior densities will be used to calculate E_{image} during an active contour evolution on a test image.

Let $I(x)$ be the pixel intensity value of the test image I at location x . It can be shown that for a closed curve β defining two regions Ω_{in} and Ω_{out} , the minimization of the energy functional $E_{image}(\beta) = -\int_{\Omega_{in}} \log(f(I(x))) dx - \int_{\Omega_{out}} \log(g(I(x))) dx$ is equivalent to the maximization of the image-based *a posteriori* segmentation probability. Using functional differentiation and Green's Theorem, as shown in [12], the gradient of the image energy is calculated as $\nabla E_{image}(\beta) = -\log\left(\frac{f(I(\beta))}{g(I(\beta))}\right) \mathbf{n}$, where \mathbf{n} is the outward unit normal vector field to the curve β . Define

$\ell_\beta \equiv \log\left(\frac{f(I(\beta))}{g(I(\beta))}\right)$. Notice that contour evolution according to the negative gradient will be along the outward normal direction if $\ell_\beta > 0$ and along the inward normal direction if $\ell_\beta < 0$. This evolution will therefore push any part of the contour more likely lying in the target out towards its most likely boundary, and it will pull any part lying outside of the target in towards its most likely boundary.

In the SAS and SAR imaging environments, the proposed region-based image energy gradient has many advantages over other commonly used edge detection methods such as EPGVF [8]. First, evolution along the normal direction provides a simpler flow compared to that which is given by a gradient vector field on the image directly. Such edge detection techniques can produce erratic, swirling flows because it will find edges throughout the entirety of a noisy SAS/SAR background. Also, capture range is only limited to the size of the image; whereas, the capture range given by [8] is extremely limited, especially in noisy SAS and SAR images. Finally, computing a gradient vector field on a single closed curve is computationally more efficient than computing a vector field across an entire image domain. Furthermore, the smoothing gradient described next is also formulated as a scalar times the unit normal, so these two energy terms can be combined easily.

Smoothing Energy: For regulating contour smoothness we follow a common approach from geometric active contours based on the idea of Euclidean heat flow. Define the smoothing energy functional as $E_{smooth}(\beta) = \int_0^1 |\dot{\beta}(t)| dt$, which is equal to the length of the curve and is naturally invariant to any re-parameterization. The gradient of E_{smooth} is given by the Euclidean heat flow equation $\nabla E_{smooth}(\beta) = \kappa_\beta \mathbf{n}$, where κ_β is the curvature at each point of β . It is well known that this particular penalty on a curve's length automatically leads to smoothing of a curve by forcing the curve to become convex over time, and, eventually as the evolution time goes to infinity, the curve evolves to a circle and shrinks to a point.

Shape Prior Energy: The shape prior energy term E_{prior} is based on a Gaussian-type probability model that is defined on a shape space of elastic curves. This probability model is estimated *a priori* from given training shapes, and we can choose to define it either on similarity-invariant elastic shape space \mathcal{S}_{sim} [11] or on affine-invariant elastic shape space \mathcal{S}_{aff} [2]. As mentioned in Section 2, with SAS and SAR shadows there is certain amount of shape variability that can be attributed to an approximately affine transformation, so a prior shape model on \mathcal{S}_{aff} is desirable.

Let $q(t) = \dot{\beta}(t) / \sqrt{|\dot{\beta}(t)|}$ be the square-root velocity function of a parameterized curve $\beta \in \mathbb{L}^2(\mathbb{S}^1, \mathbb{R}^2)$. Recall from [11, 2] that the elastic shape spaces \mathcal{S}_{sim} and \mathcal{S}_{aff} are

comprised of SRVF's of appropriate sets of curves rather than the curves themselves. Notice that E_{prior} in our active contour model is a function of β , but the statistical models built on \mathcal{S}_{sim} and \mathcal{S}_{aff} depend on q . Therefore, we first must calculate the prior energy as a function of q and then perform a numerical approximation of the change of variable to obtain $E_{prior}(\beta)$ and the subsequent gradient vector field on β . Here, we describe how to form E_{prior} on \mathcal{S}_{aff} , though the procedure is the same for \mathcal{S}_{sim} .

Given a set of prior training shapes $\{[q_i], i = 1, \dots, n\}$ in \mathcal{S}_{aff} , let us assume that we have computed their Karcher mean $[\mu]$ and covariance K . We define the prior shape density as a *truncated wrapped-normal density*, which is estimated from the data as follows. First, obtain the singular value decomposition of K as $[U, S, V] = \text{svd}(K)$, and let U_m be the m -dimensional principal subspace of $T_{[\mu]}(\mathcal{S}_{aff})$ defined as the first m columns of U . The prior density is given as the exponential mapping of $f_m(v)$, the multivariate Gaussian density on U_m , where

$$f_m(v) = \frac{1}{Z} e^{-\frac{1}{2}(v_{\parallel}^T S_m^{-1} v_{\parallel} + \|v_{\perp}\|^2 / \delta^2)} \mathbf{1}_{\|v\| < \pi}, \quad (3)$$

$v = \exp_{\mu}^{-1}(q)$, $v_{\parallel} = U_m^T v$ is the projection of v into U_m , $v_{\perp} = v - U_m v_{\parallel}$, S_m is the diagonal matrix containing the first m singular values, and Z is the normalizing constant. The scalar value δ is chosen to be less than the smallest singular value in S_m .

Suppose now that we have a test shape q that represents the SRVF of the active contour at a particular evolution time. Let $v = \exp_{\mu}^{-1}(q)$, the shooting vector from μ to q . Now define $E_{prior}(q)$ to be the negative of the exponent in $f_m(v)$ given by Eqn. 3. That is, define $E_{prior}(q) = \frac{1}{2} v^T (U_m S_m^{-1} U_m^T) v + \frac{1}{2\delta^2} \|v - U_m U_m^T v\|^2$. Minimizing this functional is therefore equivalent to maximizing the likelihood of q . The gradient vector is equal to $w = Av$, where A is the matrix $A = U_m S_m^{-1} U_m^T + (I - U_m U_m^T) / \delta^2$. Notice that w is defined on the tangent space at μ rather than at q , so the final step is to parallel translate w from μ to q . Denote this parallel translation as $\bar{w} = \nabla E_{prior}(q)$. An evolution of q along the negative gradient direction will result in an energy minimization precisely at the mean μ . The shooting vector \bar{w} is converted through numerical approximation to a gradient vector field on β that represents $\nabla E_{prior}(\beta)$.

4. Experimental Results

We test the performance of our Bayesian active contour model on two datasets of imagery collected beyond the visible spectrum. First, we segment the shadows of a cylinder target in underwater SAS imagery, and second, we segment the shadows of an armored ground-based vehicle in SAR imagery. The SAS images were created from the Shallow Water Acoustics Toolkit (SWAT), a program developed

by the Naval Surface Warfare Center Panama City Division (NSWC PCD) that synthesizes SAS imagery of various targets in seabed environments [10]. The SWAT simulator is considered accurate to reality and is widely used to test automatic target detection and recognition algorithms in place of real SAS data. The SAR images are a part of the Moving and Stationary Target Acquisition and Recognition (MSTAR) public dataset [1], which consists of X-band SAR images with 1×1 foot resolution. The dataset was collected in September of 1995 at the Redstone Arsenal, Huntsville, AL by the Sandia National Laboratory (SNL) SAR sensor platform and was jointly sponsored by DARPA and the Air Force Research Laboratory. For our purposes here, we test our method on a 100 image subset of each dataset. Each subset consists of imagery of the same target at different aspect angles and, in the case of the SAS imagery, also at different ranges. This yields in each case a shadow signature that exhibits a within-class shape variability that can be modeled to some extent by an affine transformation. Again, Fig. 2 shows a few examples taken from these SAS and SAR datasets.

In order to evaluate the accuracy of any segmentation result on these datasets, we compare the estimated boundary to its respective ground truth curve via two metrics: $d_{geod}(\cdot, \cdot)$ the geodesic shape distance on \mathcal{S}_{sim} (see [11] for more details), and $d_{bin}(\cdot, \cdot)$ a binary image metric that measures the area of non-overlapping regions. The distance d_{bin} is defined in the following manner. If \hat{B} is the binary image obtained by the segmentation and B is the ground truth binary image, the binary image distance is defined as $d_{bin}(\hat{B}, B) = \text{area}(\hat{B} \cup B - \hat{B} \cap B) / \text{area}(\hat{B} \cup B)$. The values of these two metrics together show how accurately our segmentation result matches the correct shape as well as the correct location, orientation, and scale in the image.

Assuming that the ground truth shadow boundary curves are available (e.g. using hand segmentation), we perform the following experiment on each dataset. First, we select an appropriate set of parameter values $(\lambda_1, \lambda_2, \lambda_3)$. If $\lambda_3 = 0$, it represent no prior. If $\lambda_3 > 0$, then we select either our similarity-invariant prior or our affine-invariant prior for the segmentation. We then proceed to run 3 cross-validation iterations where in each iteration we select 10 images at random for training, form the shape prior density on either \mathcal{S}_{sim} or \mathcal{S}_{aff} from the corresponding ground truth curves, and segment the remaining 90 test images with the influence of that shape prior. After each segmentation, we calculate the values of d_{geod} and d_{bin} to ground truth. Tables 1 and 2 tabulate the averages of these two distance values (shown as the pair d_{geod}, d_{bin} in each case) after performing the above procedure with three sets of parameter values. Table 1 shows results from the SAS dataset, and Table 2 shows results from the SAR dataset.

Figs. 2 and 3 give both similarity and affine shape mod-

$(\lambda_1, \lambda_2, \lambda_3)$	No Prior	Similarity	Affine
(0.1, 0.5, 0.005)	0.46, 0.21	0.25, 0.25	0.22, 0.18
(0.1, 1, 0.005)	0.28, 0.14	0.24, 0.31	0.20, 0.16
(0.1, 1, 0.001)	0.28, 0.14	0.21, 0.12	0.18, 0.11

Table 1. Evaluation of SAS image segmentation: Average values of (d_{geod}, d_{bin}) relative to ground truth curve under different shape priors.

$(\lambda_1, \lambda_2, \lambda_3)$	No Prior	Similarity	Affine
(0.01, 0.075, 0.005)	0.57, 0.27	0.50, 0.24	0.48, 0.20
(0.01, 0.15, 0.005)	0.36, 0.22	0.32, 0.17	0.31, 0.16
(0.01, 0.2, 0.008)	0.32, 0.23	0.30, 0.20	0.28, 0.18

Table 2. Evaluation of SAR image segmentation: Average values of (d_{geod}, d_{bin}) relative to ground truth curve under different shape priors.

els for a random selection of 10 ground truth curves in the SAS and SAR databases, respectively. Notice that in the SAS case in Fig. 2, most of the shape variability is removed when affine-invariance is introduced. This is because the 3D cylinder target shape is very simple and without any irregular features that may be hidden or diminished under certain aspect angles. The models on \mathcal{S}_{sim} and \mathcal{S}_{aff} exhibit quite different shape samples because the model on \mathcal{S}_{sim} must still account for affine transformation in its variability. Some of the remaining shape variability in the model on \mathcal{S}_{aff} is due to the fact that the shadow shapes are more closely related by projective transformation, but in this case, an affine approximation is valid. On the other hand, the shape model for the SAR vehicle shadow in Fig. 3 exhibits a great deal more shape variability beyond that of an affine transformation. This is because an armored vehicle is not a simple shape and consists of many features that may only show up in the shadow under certain aspect angles. Also, contrary to the SAS data, the SAR imagery is from a constant range, which produces shadow lengths that are consistent. The affine shape model visually may not show such a stark contrast to the similarity model, but its benefits are still apparent from the results presented here.

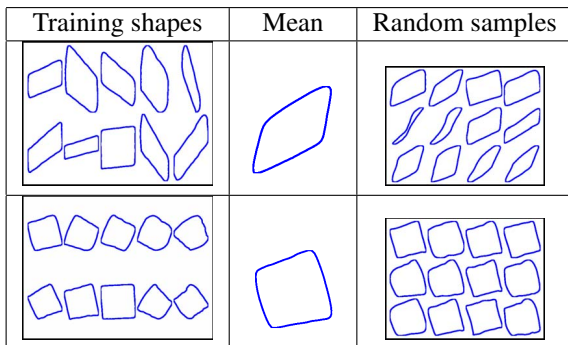


Figure 2. SAS cylinder shadow shape models. Top row: Similarity invariant shape prior. Bottom row: Affine-invariant shape prior.




Figure 3. SAR vehicle shadow shape models. Top row: Similarity invariant shape prior. Bottom row: Affine-invariant shape prior.

From Table 1 we see that with respect to all 6 distance values, segmentation of the cylinder shadow with our affine prior beats the similarity prior, and furthermore, in 5 out of the 6 distance values, segmentation with the affine prior beats segmentation with no prior. From Table 2 we see that segmentation of the vehicle shadow with our affine prior beats the similarity prior with respect to all 6 distance values, and furthermore, in all 6 cases, segmentation with the similarity prior beats segmentation with no prior. Figs. 4 and 5 show a few selections from the cross-validation segmentation experiments that contributed to the results in the tables above. The columns from left to right represent the test image, segmentation without shape prior, segmentation with similarity prior, segmentation with affine prior, and the ground truth segmentation. We remind the reader that although the distance values in the above tables seem small and relatively similar, since $d_{geod} \in [0, \pi/2]$ and $d_{bin} \in [0, 1]$, small differences in distance values can translate to impactful visual differences as seen in these figures. The benefits of using the affine prior over the similarity prior are much more pronounced in the cylinder segmentation results of Fig. 4 than the vehicle segmentation results of Fig. 5 because the cylinder shadow exhibits much more affine variability than the vehicle shadow. In summary, although a similarity-invariant, elastic shape prior is still useful in some cases, using our affine-invariant, elastic shape prior will more consistently yield improved segmentation results over that without shape prior in shadow boundary extraction in synthetic aperture imaging modalities such as SAS and SAR.

5. Conclusion

Bayesian active contour models have become an increasingly popular technique for segmenting objects in images and are especially suited for segmentation of images taken in a spectrum beyond that of visible light. Typically, such images present difficulties for segmentation algorithms due to low resolution, high background noise, and excess clutter. In order to combat these difficulties, we have developed a




Figure 4. Shadow segmentations of various SAS cylinder test images under different shape priors.

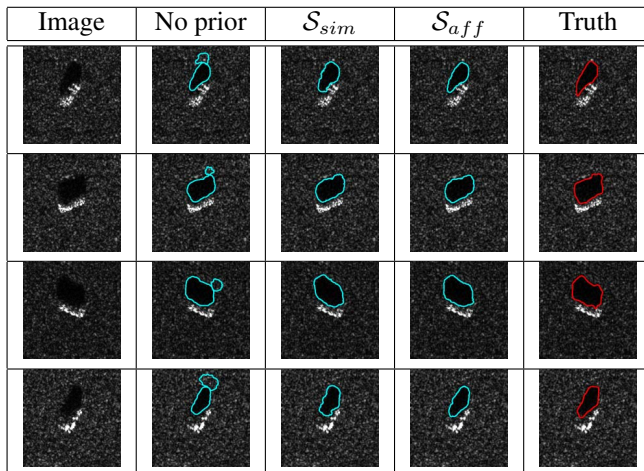


Figure 5. Shadow segmentations of various SAR armored vehicle test images under different shape priors.

parametric active contour model that employs a prior shape model based on state-of-the-art elastic shape statistics. Our framework allows for the use of either a similarity or affine-invariant shape prior. The affine-invariant shape prior is especially useful for segmenting shadow boundaries, and we test our segmentation model with both priors on SAS and SAR shadow imagery. We show an increase in shadow seg-

mentation accuracy when using the affine shape prior compared to the similarity prior and without shape prior, and, in most cases, we show increase in accuracy with the similarity prior over segmentation without shape prior.

Acknowledgements

This work was supported in part by the grants ONR N00014-09-1-0665 and NSF DMS 09-15003 to AS. The authors would also like to acknowledge the NSWC PCD ILIR Program and the support of Dr. Kerry Commander and Dr. Quyen Huynh.

References

- [1] Moving and stationary target acquisition and recognition (mstar) public dataset. <https://www.sdms.afrl.af.mil/datasets/mstar/>. 4
- [2] D. Bryner, E. Klassen, and A. Srivastava. Affine-invariant, elastic shape analysis of planar contours. *IEEE Conference on CVPR*, pages 390–397, 2012. 1, 2, 3
- [3] I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. New York: Wiley, 1998. 1
- [4] S. H. Joshi and A. Srivastava. Intrinsic bayesian active contours for extraction of object boundaries in images. *International Journal of Computer Vision*, 81(3):331–355, March 2009. 1, 3
- [5] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988. 1
- [6] J. Kim, M. Çetin, and A. S. Willsky. Nonparametric shape priors for active contour-based image segmentation. *Signal Processing*, 87(12):3021–3044, Dec. 2007. 1
- [7] M. Leventon, W. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proc. of IEEE Conference on CVPR*, volume 1, pages 316–323, 2000. 1
- [8] C. Li, J. Liu, and M. Fox. Segmentation of edge preserving gradient vector flow: an approach toward automatically initializing and splitting of snakes. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 162 – 167, June 2005. 3
- [9] M. Rousson and D. Cremers. Efficient kernel density estimation of shape and intensity priors for level set segmentation. In *Proceedings of the 8th MICCAI*, pages 757–764, 2005. 1, 3
- [10] G. Sammelmann, J. Christoff, and J. Lathrop. Synthetic images of proud targets. *Proc. of MTS/IEEE Oceans 2006 Conference*, pages 1–6, Sept 2006. 4
- [11] A. Srivastava, E. Klassen, S. Joshi, and I. Jermyn. Shape analysis of elastic curves in euclidean spaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(7):1415–1428, July 2011. 2, 3, 4
- [12] S. Zhu, T. Lee, and A. Yuille. Region competition: unifying snakes, region growing, energy/bayes/mdl for multi-band image segmentation. In *Proceedings., Fifth ICCV*, pages 416–423, June 1995. 3