

Image Segmentation via Convolution of a Level-Set Function with a Rigaut Kernel*

Özlem N. Subakan

Baba C. Vemuri

Department of Computer and Information Science and Engineering

University of Florida, P.O. Box 116120, Gainesville, FL 32611, USA

{ons, vemuri}@cise.ufl.edu

Abstract

Image segmentation is a fundamental task in Computer Vision and there are numerous algorithms that have been successfully applied in various domains. There are still plenty of challenges to be met with. In this paper, we consider one such challenge, that of achieving segmentation while preserving complicated and detailed features present in the image, be it a gray level or a textured image. We present a novel approach that does not make use of any prior information about the objects in the image being segmented. Segmentation is achieved using local orientation information, which is obtained via the application of a steerable Gabor filter bank, in a statistical framework. This information is used to construct a spatially varying kernel called the Rigaut Kernel, which is then convolved with the signed distance function of an evolving contour (placed in the image) to achieve segmentation. We present numerous experimental results on real images, including a quantitative evaluation. Superior performance of our technique is depicted via comparison to the state-of-the-art algorithms in literature.

1. Introduction

Image segmentation techniques exploit several features present in the images while trying to achieve their goal. This area of research has a long history spanning the past several decades. Variational formulation of this problem was popularized by Kass *et al.*, [7] in their seminal work on the so called snakes a.k.a. active contour models. A region-based variational formulation of this problem was proposed earlier by Mumford and Shah [17] and later popularized by Tsai *et al.* in an active contour framework [23]. The snakes model that constitutes a closed curve expressed as an arbitrarily

parameterized curve was primarily designed as an interactive segmentation model and proved to be quite useful and general in this context. An alternative model called the geometric active contour in a level set framework was then proposed in the pioneering works of Malladi *et al.* [12, 13, 14] and Caselles *et al.* [1]. This model involved a closed curve represented in an implicit form that allowed for ease in modeling shapes with arbitrary unknown topologies. Following this, a variational formulation of the geometric active contour model called the geodesic snakes, was independently introduced by Caselles *et al.* [2] and Kichenassamy *et al.* [8]. Over the past decade and a half, there have been several approaches to segmentation some of which are improvements over the geodesic active contours as well as the traditional snakes and some others that are graph-based global optimization approaches. For more on variational formulations of the image segmentation problem that led to improvements of the original proposals of the active contour model and also for graph-based techniques, we refer the reader to [18].

Despite the extensive activity in the computer vision community and the “success” of the aforementioned techniques, segmentation that preserves complex local details has been an elusive goal to achieve. The key innovation in this paper is the presentation of an adaptive, convolution based approach for segmentation which preserves the complicated geometries of the boundaries of objects in real scenes without using any prior information. We compare our method with a state-of-the-art graph theoretic approach presented recently by Schoenemann and Cremers [21]. They proposed an energy minimization framework which employs curvature constraints in a graph-theoretic formulation involving minimum ratio cycles on product graphs. In the context of segmenting textured scenes, we compare our technique with the prominent approach of Rousson *et al.*, who presented a variational formulation in a level-set framework, that incorporates a set of features obtained from the structure tensor [20]. The main limitation of their method is that it is restricted to a 2-class segmentation prob-

*This research was in part supported by NIH EB007082 to BCV and by SM CRSP through a grant from the US Agency of International Development to ONS.

lem. In contrast our method does not have such restrictions. Furthermore, to the best of our knowledge, this is the first time that a convolution based approach is being employed for feature preserving segmentation.

The proposed method consists of two stages, first of which includes the extraction of local orientation information. This information can be obtained via the application of Gabor filters. However, one of the key drawbacks of Gabor representations is the size of the filter bank one needs in order to acquire useful information. This drawback may be overcome by using steerable filters. Recently Kalliomaki and Lampinen [6] presented approximate steerability of Gabor filters in 2D for pattern recognition purposes. 3D steerable Gabor filters were formulated in [22]. Here we employ them to extract orientation information required in the first stage of the proposed approach.

The orientation distribution at each lattice point is then represented by a continuous mixture of oriented Gaussians. Continuous mixture is preferred here over the finite mixture model because one need not specify the number of components in the mixture explicitly. For more details see [5, 22]. The continuous mixture representation is cast as the Laplace transform of the mixing density over the space of covariance (positive definite) matrices. This mixing density is assumed to be in a parameterized form, namely, a mixture of Wisharts, whose Laplace transform evaluates to a closed form expression called the Rigaut type function, a scalar-valued function of the parameters of the Wishart distribution. The weights in the mixture are then computed using a sparse de-convolution technique. The second stage involves iterative local convolution of the signed distance function of an evolving contour placed in the image with the aforementioned kernel function, in a narrow banding algorithm [14]. The remainder of this paper is organized as follows: We briefly expand on the steerable Gabor filters and Rigaut kernel in Section 2. Sufficient description of the level set function is also provided in this section. Then, in Section 3, we present the experimental results along with the quantitative evaluation depicting the merits of the proposed formulation. Lastly, in Section 4 we summarize our contributions.

2. Methodology

2.1. Local Orientation Representation and the Rigaut Kernel

In order to achieve feature (corners and junctions) preserving segmentation, we exploit the local orientation information obtained by Gabor filtering the images. The main advantage of these filters due to their Gaussian envelopes is that they achieve the minimum space-frequency product specified in the uncertainty principle. So they are optimal in terms of space-frequency localization. Additionally, they exhibit the flexibility of being tunable to any frequency or

orientation and they can form a relatively good approximation of a wavelet frame. Such tuning is particularly useful in capturing any locally predominant orientations present in an image. The complex oriented Gabor filter with a non-spherical Gaussian envelope function has the following generic form:

$$G(\boldsymbol{\xi}; \varpi, \boldsymbol{\Sigma}, \mathbf{R}_\nu) = \frac{1}{2\pi\sqrt{|\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{R}_\nu^T \boldsymbol{\xi})^T \boldsymbol{\Sigma}^{-1} \mathbf{R}_\nu^T \boldsymbol{\xi}\right) \exp(i\varpi \mathbf{R}_\nu^T \boldsymbol{\xi}) \quad (1)$$

where $\boldsymbol{\xi}$ is the spatial coordinate vector, ϖ is the center frequency of the filter, $\boldsymbol{\Sigma}$ is a diagonal covariance matrix which determine the frequency bandwidths along the axes in Cartesian coordinates and \mathbf{R}_ν is a rotation matrix whose first column is a unit vector ν . Note that the resulting filter has a constant template ellipsoid determined by $\boldsymbol{\Sigma}$ and is oriented along the orientation determined by ν . To analyze orientations in regions with one or more orientations, we use the steerable Gabor filters. Steerable Gabor filters have been studied extensively in [6, 22].

In order to represent the local image structure, we postulate that at each lattice point there is an underlying probability measure associated with the manifold of $n \times n$ symmetric positive-definite matrices, \mathcal{P}_n . Let $f(\mathbf{K})$ be its density function with respect to some carrier measure $d\mathbf{K}$ on \mathcal{P}_n . (This model has been presented in the context of the diffusion-weighted MR signal attenuation by Jian and Vemuri in [3] and later used in the context of image smoothing by Subakan *et al.* in [22].) We propose to model the orientation distribution by a continuous mixture of Gaussian functions:

$$G(\boldsymbol{\xi}; \mathbf{g})/G_0 = \int_{\mathcal{P}_n} f(\mathbf{K}) \exp[-\mathbf{g}^T \mathbf{K} \mathbf{g}] d\mathbf{K}, \quad (2)$$

where $\boldsymbol{\xi}$ encodes the coordinates, $G(\boldsymbol{\xi}; \mathbf{g})$ is the response of the Gabor filter with an orientation determined by \mathbf{g} a unit direction vector, G_0 denotes the maximal filter response. So, Eq. (2) is a continuous mixture of Gaussian functions with $f(\mathbf{K})$ being a mixing density. This integral can be recognized as the Laplace transform (matrix variable case) of f [16]:

$$G(\boldsymbol{\xi}; \mathbf{g})/G_0 = \int_{\mathcal{P}_n} \exp(-\text{trace}(\mathbf{B}\mathbf{K})) f(\mathbf{K}) d\mathbf{K} = (\mathcal{L}_f)(\mathbf{B}), \quad (3)$$

where \mathcal{L}_f denotes the Laplace transform of a function f which takes its argument as symmetric positive definite matrices from \mathcal{P}_n , and $\mathbf{B} = \mathbf{g}\mathbf{g}^T$. In this expression, we are faced with the problem of recovering a distribution defined on \mathcal{P}_n that best explains the observed orientation data $G(\boldsymbol{\xi}; \mathbf{g})$. This is an ill-posed inverse problem. In general it is intractable without prior knowledge of the mixing density. A Bayesian interpretation may be given to this estimation problem by considering the orientation tensor as a

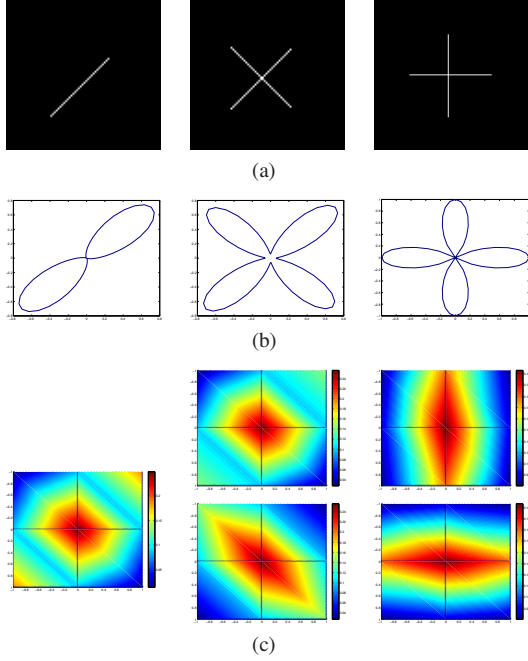


Figure 1. (a) Synthetic images with different orientations (b) Local orientation information at the bifurcation point (c) Rigaut kernel components at the bifurcation point

matrix-valued random variable with a known prior, which allows us to model the uncertainty in the orientation tensor estimation. One can interpret the orientation tensor \mathbf{K} as the concentration matrix (inverse of the covariance matrix) of the Gaussian distribution in the g -space. We impose a Wishart (see definition below) prior on this concentration matrix as is common in multivariate statistics [1] and another reason for this choice being, computational advantage accrued from the fact that the integral can be solved in closed form.

Definition 1 [11] For $\Sigma \in \mathcal{P}_n$ and for p in $(\frac{n-1}{2}, \infty)$, the Wishart distribution $\gamma_{p,\Sigma}$ with scale parameter Σ and shape parameter p is defined as

$$d\gamma_{p,\Sigma}(\mathbf{Y}) = \Gamma_n(p)^{-1} |\mathbf{Y}|^{p-(n+1)/2} |\Sigma|^{-p} e^{-\text{trace}(\Sigma^{-1}\mathbf{Y})} d\mathbf{Y}, \quad (4)$$

where Γ_n is the multivariate gamma function and $|\cdot|$ is the matrix determinant.

This distribution possesses a closed form expression for the Laplace transform, called the Rigaut-type function [19] given by:

$$G(\boldsymbol{\xi}; \mathbf{g}) = G_0 (1 + (\mathbf{g}^T \Sigma \mathbf{g}))^{-p}. \quad (5)$$

where Σ is the scale parameter of the Wishart distribution. Since the Wishart distribution is a single mode distribution, it cannot resolve the orientational heterogeneity. Hence,

it is natural to use a discrete mixture of Wishart distributions for the mixing density in Eq.(3), which is then expressed as $dF = \sum_{i=1}^N w_i d\gamma_{p_i, \Sigma_i}$. Note that the number of components in mixture, N , only depends on the discretization resolution and should not be interpreted as the expected number of bifurcations characterizing the local structure. In order to estimate the numerical scale of the eigenvalues of Σ_i , we first assume a single Gaussian model $G(\boldsymbol{\xi}; \mathbf{g})/G_0 = \exp[-\mathbf{g}^T \Sigma \mathbf{g}]$ and then solve for Σ using linear regression. The trace of the resulting Σ is used as a good estimation for the trace of Σ_i . In practice, we fix the ratio between the larger and the smaller eigenvalues. Hence, the eigenvalues of Σ_i can be determined on a pixel by pixel basis. Furthermore, this rotational symmetry leads to a tessellation where N unit vectors evenly distributed on the unit sphere are chosen as the principal eigenvectors of Σ_i . For M measurements with $\mathbf{g}_j, j = 1, 2, \dots, M$, the mathematical model yields a linear system $\mathbf{A}\mathbf{w} = \mathbf{y}$, where $\mathbf{y} = (G(\boldsymbol{\xi}; \mathbf{g})/G_0)$ contains the normalized measurements, \mathbf{A} is a $M \times N$ matrix with $A_{ji} = (1 + \text{trace}(\mathbf{B}_j \Sigma_i))^{-p}$, and $\mathbf{w} = (w_i)$ is the weight vector to be estimated. This can be cast as a sparse de-convolution problem formulated in a general form as $\mathbf{A}\mathbf{w} = \mathbf{y} + \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ represents certain noise model. We assume that the measurement errors $\boldsymbol{\eta}$ are i.i.d. and normally distributed. Since the maximization of the likelihood function under a conditional Gaussian noise distribution for a linear model is equivalent to minimizing a sum-of-squares error function, we use a non-negative least squares (NNLS) minimization which achieves an accurate and sparse solution for

$$\min \|\mathbf{A}\mathbf{w} - \mathbf{y}\|^2 \quad \text{subject to} \quad \mathbf{w} \geq 0. \quad (6)$$

Jian and Vemuri have shown that this de-convolution method outperforms many other methods in achieving accuracy and sparsity [4, 3]. After \mathbf{w} is estimated for orientations, we locally convolve the signed distance function with the spatially varying Rigaut kernel formulated as

$$R(\mathbf{x}) = \int_{\mathcal{P}_n} f(\mathbf{K}) \exp[-\mathbf{x}^T \mathbf{K} \mathbf{x}] d\mathbf{K}, \quad (7)$$

$dF = f(\mathbf{K})d\mathbf{K}$ is the mixing density described above, i.e. $dF = \sum_{i=1}^N w_i d\gamma_{p_i, \Sigma_i}$. Using Laplace transform of the Wishart distribution, the Rigaut kernel is given as

$$R(\mathbf{x}) = \sum_{i=1}^N w_i (1 + \text{trace}(\mathbf{x}\mathbf{x}^T \Sigma_i))^{-p} \quad (8)$$

When $p \rightarrow \infty$, this model reduces to a mixture of oriented Gaussians with weight vector \mathbf{w} . Also note that since the weight vector w and the Σ_i 's change with regard to the local orientation information, this formulation leads to spatially varying segmentation kernels.

Figure 1 illustrates the Rigaut kernel on synthetic images with different orientations. The orientation informa-

tion obtained by steerable Gabor filters (for 81 different orientations in $[0, 2\pi]$) is projected onto cosine vs. sine of the orientations, as shown in 1(b). The peaks in the plot correspond to the local geometry for the chosen bifurcation point. Figure 1(c) depicts the 2D views of the components in the Rigaut kernel. For example, there are two components in the Rigaut kernel for the X shape in Figure 1. One of these two components has an orientation of 45° as illustrated in the top middle figure in 1(c) and the other has an orientation of 135° as seen in the bottom middle figure in 1(c). Notice that the green color denotes higher values in the Rigaut kernel component compared to the blue. This green color is observed to correctly indicate the 45° orientation which lies along the prongs of the X-shape in the top middle figure of 1(c). Similar reasoning may be applied to the green colored region in the top right figure of 1(c).

2.2. Convolution of a Level-Set Function with the Rigaut Kernel

The key idea of level set methods is to represent an evolving curve C by the zero level set of a Lipschitz continuous function $\phi: \Omega \rightarrow \mathbb{R}$. So, $C = \{(x, y) \in \Omega: \phi(x, y) = 0\}$. We choose ϕ to be positive inside C and negative outside. C is evolved using the described Rigaut kernel convolution; i.e. ϕ is convolved locally with the Rigaut kernel in a narrow banding algorithm. The level-set update equation is simple and given by:

$$\phi_{t+1}(\mathbf{x}) = \phi_t(\mathbf{x}) * R(\mathbf{x}) \quad (9)$$

The update stops when no further changes in the zero level set are observed. The feature/junction preserving property is achieved due to the nature of the characteristic response of the convolution kernel.

3. Experimental Results



Figure 2. Segmentation result (right) of a textured image (left) using our Rigaut kernel-based convolution filter

In this section we present the experimental results obtained from an application of our segmentation technique to some real images containing a variety of junctions, corners and textures. The first test image contains two zebras in a natural scene. The zoomed-in region depicts a low-contrast region where a successful segmentation is quite difficult despite which, our technique clearly segments the feet of the

two zebras while preserving their geometry. Note the presence of junctions in the zoomed region.

As a second example we choose a leopard image which has been experimented with by many other segmentation approaches in the past. Specifically, we compare our technique with the method proposed by Rousson *et al.* [20]. The leopard's tail recovery proves our method to be fully competitive to the recent approaches [9, 20]. Figure 3(a) was taken from [20].




Figure 3. Segmentation results on a test image of a leopard using (a) algorithm by Rousson *et al.* [20] (b) our Rigaut kernel-based convolution filter. The result in (a) was taken from [20]

Figure 4 depicts another example where the results from 4 different segmentation methods are shown, namely, (i) the piecewise constant version of the Mumford-Shah segmentation scheme, (ii) the Mumford-Shah scheme, (iii) elastic ratio technique by Schoenemann *et al.* [21], and (iv) our technique respectively. On careful examination, Figure 4(d) is slightly better in segmentation (see the details around the tail) than 4(c). Techniques in Figures (a) and (b) fail in this difficult scene. (Figures (a), (b) and (c) were reproduced with permission from the authors of [21].) Figure 5(b) depicts another example, note the accurate segmentation of the sling-on strap attached to the case as compared to that obtained by a competing method shown in figure 5(a). Figure 6 contains another set of segmentation results, one performed by humans and another by our technique on some images from the Berkeley Segmentation Dataset [15]. Note the closeness of the two segmentations which depicts the qualitative accuracy of our scheme.

Lastly, we present the precision/recall curves for our approach on the 4 images in Figure 6 in order to evaluate the segmentation quality. Precision and recall are preferred as measures of segmentation quality because they are sensitive to under and over-segmentation. We use human segmentations from the Berkeley Segmentation Dataset as ground truth. Since there are multiple human segmentations, we chose to compare against the union of the segmentations from all human subjects for each image. Matching of the

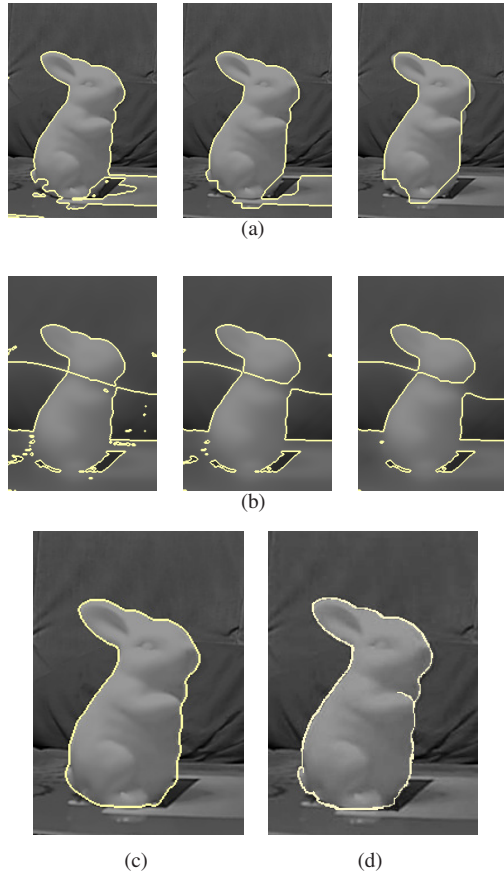


Figure 4. Segmentation results using (a) piecewise constant Mumford-Shah (b) piecewise smooth Mumford-Shah (c) elastic ratio by Schoenemann *et al* [21] (d) our technique. (Figures (a), (b) & (c) were taken from [21])




Figure 5. Segmentation result (a) from Schoenemann and Cremers [21] (b) from our convolution method

boundaries between two segmentations is performed by examining a neighborhood within a radius of $\epsilon = 3$, which is a reasonable choice given the resolution of the images. We tested the effect of the shape parameter p in the Rigaut kernel and the effect of a threshold parameter (for values in $\{0.005, 0.05, 0.1, 0.5\}$) for Gabor filter responses. For each run, we take the average of the precision and recall values for the 4 images above and use these values to generate the tuning curves that characterize the performance of




Figure 6. (first column) Segmentations performed by humans (taken from the Berkeley Segmentation Dataset and Benchmark) (second column) Segmentations obtained by our technique (The images have a size of 481x321 pixels.)




Figure 7. Tuning curves for variations of the shape parameter p for our Rigaut kernel-based segmentation. The points along each curve correspond to variations of a threshold parameter for orientation information obtained by steerable Gabor filters. The threshold varies within $[0.05, 0.5]$.

our method. Experimentation showed that the largest differences between segmentations were obtained by varying the threshold parameter, which determines the level of detail that is desired in the segmentation. A low threshold value results in over-segmentation which is characterized in the curves by high recall but low precision. Moreover, for a given threshold value, the precision/recall values change

only slightly with respect to the changes in p , depicting the insensitivity of the segmentation results to this parameter.

4. Conclusions

In this paper, we treated the problem of feature preserving segmentation. We presented results that depicted accurate segmentation of complex scenes containing a variety of complex local geometries. This was achieved through modeling the local geometry of the image function – extracted as orientation information via the use of steerable Gabor filters – by a continuous mixture of multi-variate Gaussian functions, with the mixing density being assumed to be a mixture of Wishart distributions. This leads to a closed form expression called the mixture of Rigaut functions which serves as the spatially varying convolution filter that when applied iteratively to the signed distance function of an arbitrarily initialized contour in the image to be segmented, yields the desired segmentation.

A sparse de-convolution technique is employed for computation of the weights in the mixture of Wisharts model. The sparseness requirement is justified by the fact that the local image geometries do not in general have more than a small set of maxima, i.e., one does not encounter a large number of spikes at each lattice point in an image. The classic non-negative least squares (NNLS) algorithm developed in [10] is most suitable for our de-convolution problem in achieving sparseness and robustness. To the best of our knowledge, this is the first time that a convolution based approach is being used for feature preserving segmentation. Finally, the experimental results on real images with complex geometries and textures demonstrate that the proposed model provides better overall performance than other state-of-the-art techniques for segmentation.

References

- [1] V. Caselles, F. Catta, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66:1–31, 1993.
- [2] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Intl. Journal of Computer Vision*, 22:61–97, 1997.
- [3] B. Jian and B. Vemuri. Multi-fiber reconstruction from diffusion mri using mixture of wisharts and sparse deconvolution. *Inf. Process Med Imaging*, 20, 2007.
- [4] B. Jian and B. C. Vemuri. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted mri. *IEEE Trans. Med. Imaging*, 26(11):1464–1471, 2007.
- [5] B. Jian, B. C. Vemuri, E. Özarlan, P. Carney, and T. Mareci. A novel tensor distribution model for the diffusion-weighted mr signal. *Neuroimage*, 37(1):164–176, 2007.
- [6] I. Kalliomäki and J. Lampinen. On steerability of gabor-type filters for feature detection. *Pattern Recogn. Lett.*, 28(8):904–911, 2007.
- [7] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Intl. Journal of Computer Vision*, 1(4):321–331, 1988.
- [8] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *IEEE Intl. Conf. on Computer Vision*, 1995.
- [9] J. Kim, J. W. F. III, A. J. Y. Jr., M. Çetin, and A. S. Willsky. Nonparametric methods for image segmentation using information theory and curve evolution. In *ICIP*, pages 797–800, 2002.
- [10] C. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Prentice-Hall, 1974.
- [11] G. Letac and H. Massam. Quadratic and inverse regressions for Wishart distributions. *The Annals of Statistics*, 26(2):573–595, 1998.
- [12] R. Malladi, J. A. Sethian, and B. C. Vemuri. A topology independent shape modeling scheme. In *SPIE Proc. on Geometric Methods in Computer Vision II*, volume 2031, pages 246–256. SPIE, July 1993.
- [13] R. Malladi, J. A. Sethian, and B. C. Vemuri. Evolutionary fronts for topology-independent shape modeling and recovery. In *ECCV (1)*, pages 3–13, 1994.
- [14] R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Anal. Mach. Intell.*, 17(2):158–175, 1995.
- [15] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *IEEE Intl. Conf. on Computer Vision*, volume 2, pages 416–423, July 2001.
- [16] A. M. Mathai. *Jacobians and functions of matrix argument*. World Scientific, 1997.
- [17] D. Mumford and J. Shah. Boundary detection by minimizing functionals, I. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 22–26, 1985.
- [18] N. Paragios, Y. Chen, and O. Faugeras. *Handbook of Mathematical Models in Computer Vision*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
- [19] J. P. Rigaut. An empirical formulation relating boundary lengths to resolution in specimens showing ‘non-ideally fractal’ dimensions. *Journal of Microscience*, 133:41–54, 1984.
- [20] M. Rousson, T. Brox, and R. Deriche. Active unsupervised texture segmentation on a diffusion based feature space. In *CVPR (2)*, pages 699–706, 2003.
- [21] T. Schoenemann and D. Cremers. Introducing curvature into globally optimal image segmentation: Minimum ratio cycles on product graphs. In *IEEE International Conference on Computer Vision*, Rio de Janeiro, Brazil, Oct. 2007.
- [22] Ö. Subakan, B. Jian, B. C. Vemuri, and C. E. Vallejos. Feature preserving image smoothing using a continuous mixture of tensors. In *IEEE International Conference on Computer Vision*, Rio de Janeiro, Brazil, Oct. 2007.
- [23] A. Tsai, A. J. Y. Jr., and A. S. Willsky. Curve evolution implementation of the Mumford-Shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. on Image Processing*, 10(8):1169–1186, 2001.