

Variational Segmentation using Fuzzy Region Competition and Local Non-Parametric Probability Density Functions

Benoît Mory, Roberto Ardon, Jean-Philippe Thiran

▶ To cite this version:

Benoît Mory, Roberto Ardon, Jean-Philippe Thiran. Variational Segmentation using Fuzzy Region Competition and Local Non-Parametric Probability Density Functions. IEEE 11th International Conference on Computer Vision, 2007. ICCV 2007., Oct 2007, Rio de Janeiro, Brazil. pp.1-8, 10.1109/ICCV.2007.4408959. hal-00438812

HAL Id: hal-00438812 https://hal.science/hal-00438812

Submitted on 7 Dec 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Variational Segmentation using Fuzzy Region Competition and Local Non-Parametric Probability Density Functions

Benoit Mory Roberto Ardon

Philips Medical Systems Research Paris, 51 rue Carnot, B.P. 301, F-92156 Suresnes Cedex, France {benoit.mory,roberto.ardon}@philips.com

Jean-Philippe Thiran

Signal Processing Institute (ITS), Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland

jp.thiran@epfl.ch

Abstract

We describe a novel variational segmentation algorithm designed to split an image in two regions based on their intensity distributions. A functional is proposed to integrate the unknown probability density functions of both regions within the optimization process. The method simultaneously performs segmentation and non-parametric density estimation. It does not make any assumption on the underlying distributions, hence it is flexible and can be applied to a wide range of applications. Although a boundary evolution scheme may be used to minimize the functional, we choose to consider an alternative formulation with a membership function. The latter has the advantage of being convex in each variable, so that the minimization is faster and less sensitive to initial conditions. Finally, to improve the accuracy and the robustness to low-frequency artifacts. we present an extension for the more general case of local space-varying probability densities. The approach readily extends to vectorial images and 3D volumes, and we show several results on synthetic and photographic images, as well as on 3D medical data.

1. Introduction

Image segmentation is such an essential component of modern computer vision applications that it remains a major topic of research, despite considerable efforts made over the last two decades in terms of theory and algorithms. In particular, variational principles have greatly helped the design of consistent frameworks. The common assumption is that the expected partitioning can be obtained by minimizing an appropriate objective functional. The performance of such segmentation models mainly depends on the relevance of the functional for specific homogeneity requirements. Statistical criteria on low-level features such as intensity, color, motion and texture have proved suitable to discriminate between image regions. With the success of active contours, many recent attempts to embed such region-based statistics into a variational formulation have relied on boundary evolution. Among those, we shall here distinguish between *parametric* and *non-parametric* approaches.

Using bayesian principles, the *Region Competition* algorithm [18] has unified earlier works and paved the way for subsequent efforts along the same line [6, 10, 16]. A review of these parametric methods can be found in [7]. They can incorporate complex multivariate texture and color cues [3] and have in common to (a) derive a statistical criterion from the maximization of the posterior probability of the segmentation, given the observed image, and (b) make strong assumptions about the distributions in the form of *parametric* models, so that only a small set of statistical parameters are optimized. The choice of a specific model, often Gaussian, restricts the applicability to the limited set of images that satisfy the underlying assumptions.

To overcome this limitation, *non-parametric* statistical boundary evolution algorithms have emerged for segmentation and tracking [8, 11, 12]. Using pure intensity distributions [12], complex multivariate texture [1] or motion information [9], these methods follow a common methodology: (a) derive a minimization criterion from informationtheoretic measures on the region distributions, and (b) use the Parzen window technique [17] to estimate the unknown densities. Typical measures use entropy, mutual information or Kullback-Leibler distance between distributions.

The aforementioned variational approaches have two practical shortcomings. First, the minimization is based on boundary evolution schemes, whose convergence is relatively slow and sensitive to initial conditions. Second, the criteria assume that each region can be statistically represented by a single global distribution. There are many situations where this global perspective is too simplistic to achieve a precise delineation of the boundary, while a local analysis of the distributions would be more discriminative. In this paper, we propose a novel method that has none of these drawbacks and keeps both the robustness of *nonparametric* approaches and the simplicity of *Region Competition*. We focus on the two-phase case, already covering a wide range of optimal separation problems.

We shall present our contributions as follows: In Section 2, we revisit the Region Competition and propose a novel minimization criterion. It can be used in classical boundary evolution schemes in order to discriminate regions according to their global feature statistics, without any prior parametric model. In Section 3, an alternative convex formulation is developed, optimizing membership functions instead of boundaries. In Section 4, we expose an extension for the more general case of space-varying, smooth, non-parametric probability density functions. In Section 5, we briefly describe a stable and efficient minimization strategy. The flexibility of the method is shown with various 2D and 3D segmentation examples of synthetic, photographic and medical images.

2. Non-Parametric Competition Model

The *Region Competition* algorithm [18] has inspired many subsequent works on variational region-based image partitioning. The principle is to minimize the sum of suitably defined error functions in each phase and a regularization term. When a two-phase partition of an image I is considered over the domain $\Omega \subset \mathbb{R}^n$, a general form of the functional is

$$F_0(\Sigma, \alpha_1, \alpha_2) = \int_{\partial \Sigma} ds + \int_{x \in \Sigma} r_1(\alpha_1, x) + \int_{x \in \Sigma^c} r_2(\alpha_2, x), \quad (1)$$

where $\Sigma \subset \Omega$ is the foreground region, $\Sigma^c = \Omega \setminus \Sigma$ the background and $\partial \Sigma$ their common boundary. The first term is a classical regularization penalizing the length of the boundary. Functions $r_i : \Omega \to \mathbb{R}$, are *a priori* given error functions that encode the underlying model of each region. In general, these error functions depend on some unknown parameters α_i , typically a small set of scalars. The usual strategy for the minimization of F_0 is to perform successive steps on the partition Σ and on the region parameters α_i , alternatively.

When the region parameters are considered fixed, the minimization step on Σ is classically carried out using a gradient-descent scheme. Based on the Euler-Lagrange equation derived from F_0 , it consists in evolving the boundary $\Gamma = \partial \Sigma$, through the equation

$$\frac{\partial \Gamma}{\partial t} = \left\{ \kappa - \left(r_1 - r_2 \right) \right\} \boldsymbol{n} \tag{2}$$



Figure 1. Domain notations, distributions p_i and kernel K.

where κ and n stand for the curvature and outward-pointing normal of Γ . The mean curvature term κ ensures the *smoothness* of the boundary, and the term $r = r_1 - r_2$ is the *competition* that drives the evolution of Γ toward the expected segmentation. Thus, the choice of appropriate functions r_i and parameters α_i is critical to control the result.

Bayesian inference and maximum-likelihood principles are often used to determine error functions capable of modeling image regions with statistical analysis [7, 16, 18]. The corresponding general expression of r_i is given by

$$r_i(\alpha_i, x) = -\lambda \, \log(P_i(I(x))|\alpha_i), \tag{3}$$

where λ is a parameter balancing smoothness and competition. P_i are known probability densities, only their unknown parameters α_i have to be periodically optimized. A classical assumption considers densities P_i to be Gaussian $\mathcal{N}(m_i, \sigma_i)$. In this case, the optimal parameters (m_i^*, σ_i^*) are at each step the mean and the variance of the image in each region. If the variances are known, (3) becomes $r_i(m_i, x) = \lambda_i (I(x) - m_i)^2$ and (1) boils down to a widely used method [6] for piecewise-constant images.

When reliable *a priori* information about the expected distributions is available, these parametric models perform remarkably well. Unfortunately, the specific choice of a model can also limit their applicability to a restricted class of images. For sake of generality, we seek a non-parametric formulation that would consider the probability density functions themselves to be the unknown region parameters. To that end, we propose a novel error term whose expression is justified by the study of the corresponding optimality conditions and the competition function. If *I* is a *m*-component vectorial image taking its values in $\mathcal{A} \subset \mathbb{R}^m$, we define the following error functions:

$$r_i(p_i, x) = \lambda \int_{\mathcal{A}} \left(p_i(a) - K(I(x) - a) \right)^2 da \qquad (4)$$

where K is a symmetric m-dimensional kernel such that:

$$\int_{\mathcal{A}} K(a) da = 1 \quad \text{and} \quad K(a) > 0 \quad \forall a \in \mathcal{A} \quad (5)$$

The expression of the non-parametric competition functional is obtained by plugging these errors functions in (1):

$$F(\Sigma, p_1, p_2) = \int_{\partial \Sigma} ds + \lambda \iint_{x \in \Sigma, a \in \mathcal{A}} (p_1(a) - K(I(x) - a))^2 + \lambda \iint_{x \in \Sigma^c, a \in \mathcal{A}} (p_2(a) - K(I(x) - a))^2$$
(6)

Let us now study the optimality conditions of F with respect to functions p_i . Standard calculus of variations leads to the optimal expressions

$$p_1^*(a) = \frac{1}{|\Sigma|} \int_{\Sigma} K(I(x) - a) dx$$

$$p_2^*(a) = \frac{1}{|\Sigma^c|} \int_{\Sigma^c} K(I(x) - a) dx$$
(7)

They correspond to continuous versions of non-parametric Parzen window density estimations [17] in the foreground and the background. Thus, the periodic update of the distributions in the alternate minimization scheme naturally involves a well-established method for the estimation of the probability densities in each region. Note that as soon as K satisfies (5), then $\int_{\mathcal{A}} p_i^* = 1$, without any additional explicit constraint in F.

For a complete understanding, it is essential to study the effects of the new errors terms on the boundary evolution. As already mentioned, the optimal position of the boundary is driven by the competition $r(x) = r_1(p_1, x) - r_2(p_2, x)$, whose expression is now given by

$$r(x) = 2 \int_{a \in \mathcal{A}} (p_2(a) - p_1(a)) K(I(x) - a) + \int_{a \in \mathcal{A}} p_1(a)^2 - p_2(a)^2$$
(8)

The first term measures the difference between the densities in a window K around I(x) (see Fig. 1). It is a local likelihood test on I(x) driving the boundary toward the most likely region. If this difference vanishes, no decision can be made locally from I(x), and the second term is a bias driving the boundary toward the region of smaller density energy. A possible interpretation is that the final steady state will define regions whose intensity distributions tend to be well-separated and equally compact.

As most related approaches, functional (6) is defined over the set of regions or equivalently their boundaries. For optimization purpose, the unstructured nature of this set, in particular its non-convexity, is a drawback and would require elaborated minimization strategies to avoid local minima. In the following section, we rely on an alternative formulation that ensures the convexity with respect to each variable.

3. Fuzzy Membership Formulation

We recently proposed in [15] to perform the minimization of any functional of the form (1) by considering a closely related convex problem that does not involve boundary evolution. This *Fuzzy Region Competition* formulation, inspired by [5], has computational advantages and provides solutions that are in practice less sensitive to initial conditions. The idea is to replace in (1) the region Σ by a fuzzy membership function u, and minimize instead

$$\int_{\Omega} |\nabla u| + \int_{x \in \Omega} u(x) r_1(\alpha_1, x) + \int_{x \in \Omega} (1 - u(x)) r_2(\alpha_2, x), \quad (9)$$

where u belongs to the convex set of bounded variation functions between 0 and 1. It represents the membership to the foreground and can be seen as a fuzzy version of its characteristic function. The regularization term is the total variation of u, *i.e.* the sum of the perimeters of its level sets. This problem is convex in u and the set of its solutions proves to be stable under thresholding [15]. Thus, to any solution corresponds a thresholded binary characteristic function that still minimizes (9) for given α_1 and α_2 . This defines a partition of the image that is also optimal with respect to (1). Minimization of (9) most often leads to a unique binary solution, making the thresholding step superfluous. Moreover, one can use stable numerical schemes based on total variation [2, 4] that offer notably faster convergence than gradient-descent.

We propose to apply the same principle to the two-phase region competition using non-parametric density functions introduced in the previous section. Replacing in (6) the foreground Σ by a fuzzy membership u constrained to take its values in [0, 1], we shall now minimize:

$$F_{G}(u, p_{1}, p_{2}) = \int_{\Omega} |\nabla u| + \lambda \int_{x \in \Omega} u(x) \int_{a \in \mathcal{A}} (p_{1}(a) - K(I(x) - a))^{2} + \lambda \int_{x \in \Omega} (1 - u(x)) \int_{a \in \mathcal{A}} (p_{2}(a) - K(I(x) - a))^{2}$$
(10)

The optimal p_1 and p_2 naturally depend in this case on the function u, and it is easy to show that (7) becomes:

$$p_{1}^{*}(a) = \frac{1}{\|u\|_{1}} \int_{\Omega} u(x) K(I(x) - a) dx$$

$$p_{2}^{*}(a) = \frac{1}{\|1 - u\|_{1}} \int_{\Omega} (1 - u(x)) K(I(x) - a) dx$$
(11)

where $||u||_1 = \int_{\Omega} u(x) dx$. These optimal functions now correspond to continuous versions of *weighted* Parzen density estimates. The contribution of each point x to the estimation of the foreground distribution is weighted by its membership u(x). Hence during minimization, each pixel



Figure 2. Synthetic images: In rows, we present 3 synthetic 2D experiments where foreground and background have been artificially generated by a global intensity density function. In each row, from left to right: original image; initial membership function *u*; two intermediate states of *u*; result at convergence; corresponding boundary overlaid on the image; the initial (dashed) and final (solid) distributions of the foreground (white) and background (black). First row shows uni-modal densities of distinct mean and variance, a Gaussian foreground over a uniform background. Second row shows a uni-modal Gaussian foreground over a bi-modal Gaussian background, both distributions having the exact same mean and variance. Last row shows a tri-modal Gaussian foreground on a bi-modal uniform background, with same mean and same variance, with additional spatial correlations.

may contribute to both regions according to the *certainty* of its current classification.

Observe that the expression of the functional F_G in (10) is convex with respect to each variable u, p_1 and p_2 . Thus it is convenient to employ an alternate minimization scheme, successively considering each variable to be the only argument while the others are kept fixed. At each step, p_1 and p_2 are updated with their optimal expression (11). As for the membership function u, a standard gradient-descent scheme could be used, implementing the evolution equation:

$$\frac{\partial u}{\partial t} = \operatorname{div}\left(\frac{\nabla u}{|\nabla u|}\right) - (r_1 - r_2) \tag{12}$$

Nonetheless, faster convergence and improved stability can be obtained by considering a suitable approximation of F, avoiding the singularities of the curvature term. This technique will be further detailed in Section 5. As exposed in the original Parzen's paper on the nonparametric estimation of probability densities [17], there are a number of valid choices for the kernel K. For sake of simplicity and efficiency, we use a m-dimensional Gaussian

$$K(a) = \frac{1}{(2\pi)^{m/2}} \exp\left(-\frac{1}{2}a^T M^{-1}a\right).$$
 (13)

M, the covariance matrix, is chosen diagonal in our experiments. However, in the general case, it may be adjusted to reflect the dependence between channels.

In Fig. 2, we show the results of the method on synthetic images, where background and foreground have been generated by various gray-value density functions. The method is able to cope with non-Gaussian, multi-modal, overlapping distributions. In the last two examples, both distributions have identical mean and variance, making the foreground virtually invisible.

In Fig. 4, realistic foreground segmentations of photographic images illustrate the multivariate case, using the CIEL*a*b* color space.

The synthetic results are qualitatively similar to those reported in [12], although our method uses neither information theory nor curve evolution. Moreover, this new formulation can be extended to the more challenging case of local and space-varying non-parametric density functions. This extension is the focus of the next section.

4. Extension to Local Densities

Like previous segmentation approaches with global region distributions, our model still has several practical limitations. First, in real images, it is seldom possible to discriminate between foreground and background with a single probability density for each region. Especially critical are the cases of cluttered and heterogeneous backgrounds and the presence of low-frequency artifacts such as illumination changes. Second, it may be difficult to obtain a precise delineation of the boundary since the local contributions of nearby pixels from both sides are diluted in the global estimation of the densities. Third, simple image operations such as cropping to a region of interest can seriously affect the background estimation and the final result. Those limitations could be significantly reduced if the distributions p_i were defined locally and allowed to vary in space.

Our purpose is to formulate a generalization of the previous non-parametric region competition model that would lead to a local estimation of the probability densities. To that end, we introduce in the functional a sliding window in order to localize the error. In (10), the global contribution of the foreground to the total error is:

$$E = \int_{x \in \Omega} u(x) \int_{a \in \mathcal{A}} \left(p_1(a) - K(I(x) - a) \right)^2 \qquad (14)$$

We now consider the same error locally, around a point $y \in \Omega$: the contribution of each point is multiplied by a symmetrical, positive and smooth window $W : \Omega \to \mathbb{R}$,

$$e(y) = \int_{x \in \Omega} W(x-y)u(x) \int_{a \in \mathcal{A}} (p_1(y,a) - K(I(x)-a))^2,$$
(15)

where the density $p_1(y, a)$ is now dependent on the position. The total contribution of the foreground is obtained by integrating the local error e(y) in the whole domain. Adding a similar contribution for the background and switching the order of integration leads to the new functional to minimize:

$$F_{L}(u, p_{1}, p_{2}) = \int_{\Omega} |\nabla u| + \lambda \int_{x \in \Omega} u(x) \int_{y \in \Omega}^{\Omega} W(x-y) \int_{a \in \mathcal{A}} (p_{1}(y, a) - K(I(x)-a))^{2} + \lambda \int_{x \in \Omega} (1-u(x)) \int_{y \in \Omega} W(x-y) \int_{a \in \mathcal{A}} (p_{2}(y, a) - K(I(x)-a))^{2}$$
(16)

which corresponds to the definition of the following error function:

$$r_i(p_i, x) = \lambda \int_{y \in \Omega} W(x-y) \int_{a \in \mathcal{A}} (p_i(y, a) - K(I(x) - a))^2$$
(17)

The competition $r(x) = r_1(p_1, x) - r_2(p_2, x)$ has a similar form as in (8) and the classification is still driven by likelihood tests on I(x). Here, the local likelihood differences are measured using a generalized neighborhood in both intensity and space, with a window K around I(x) and W around x. The study of the optimality conditions of (16) leads to the optimal local density functions:

$$p_{1}^{*}(y,a) = \frac{\int_{\Omega} W(x-y)u(x)K(I(x)-a)dx}{\int_{\Omega} W(x-y)u(x)dx}$$
$$p_{2}^{*}(y,a) = \frac{\int_{\Omega} W(x-y)(1-u(x))K(I(x)-a)dx}{\int_{\Omega} W(x-y)(1-u(x))dx}$$
(18)

Defining function $f_a(x) = K(I(x)-a)$, this can be rewritten with convolutions,

$$p_{1}^{*}(y,a) = \frac{[W*(u.f_{a})](y)}{[W*u](y)}$$

$$p_{2}^{*}(y,a) = \frac{[W*((1-u).f_{a})](y)}{[W*(1-u)](y)}.$$
(19)

Even though no spatial regularity constraint on the local probability densities is explicit in (16), the solutions are as regular as the window W. Indeed, for a given image value a we recognize in (19) normalized convolutions of f_a with W, the fuzzy membership functions u and (1 - u) being the certainty measures. The theory of normalized convolution, introduced in [13], is a simple and useful extension of convolution that takes into account uncertain or missing image samples. Here, for each value a, the local probability densities are obtained by a spatial smoothing of f_a . This smoothing is selective: contributions to p_1 are weighted by



Figure 3. Local Competition on the cameraman image. Left column, results of the global model, original image (top), global distributions p_i (center) and final membership function u (bottom). Right column, results of the local model, solid final boundary contours with a dashed representation of the window W at two pixel locations (top), local distributions p_i (center) and final u (bottom).

the membership u while contributions to p_2 are weighted by 1 - u. The minimization process eventually provides a binary solution, so that no smoothing occurs across the boundary at convergence. Therefore, a solution (u^*, p_1^*, p_2^*) can be interpreted as a piecewise-smooth approximation of the local probability density of the image.

A suitable choice for the window function W is the normalized isotropic n-dimensional Gaussian kernel

$$W(x) = (2\pi\sigma^2)^{-n/2} \exp(-|x|^2/2\sigma^2).$$
(20)

The standard deviation σ explicitly provides the model with an intrinsic spatial scale, related to the expected density variations. When $\sigma \to \infty$, we end up with the global model of section 3. Furthermore, being positive, non-compactly supported and \mathcal{C}^{∞} , the Gaussian window guarantees the regularity of p_i functions everywhere in the domain Ω .

Fig. 3 illustrates on the cameraman image the advantages of using local probability densities. The global intensity distribution is mainly composed of 3 modes of increasing mean value, corresponding to the cameraman, the grass and the sky, respectively. On the left, the global model is easily capable of roughly extracting the first mode, but fails to precisely capture the boundary. On the right, we check that the local model has more powerful discrimination capabilities by carrying out the minimization of (16), using the global model as initialization. Observe how the legs and the camera are now precisely extracted from the background, thanks to the use of local densities. The local equilibrium between foreground and background densities implicitly selects optimal spatially-adaptive thresholds, which is visible on the distributions shown on the right. Although the background seems to be approximately Gaussian, locally, we do not need to make that assumption explicit.

Fig. 4 shows a detailed 3D medical segmentation experiment using exclusively our non-parametric models. The advantages of the local densities as a refinement tool are emphasized for the vascular case.

5. Minimization

We now describe a possible strategy to carry out the minimization of the component-wise convex functionals (10) and (16). We focus on the minimization of F_L , F_G being a particular case. As already mentioned, we follow an alternate scheme where u, p_1 and p_2 are considered successively. For the membership function u, a possible way would be to rely on the gradient-descent scheme (12) derived from the Euler-Lagrange equation. This involves the computation of the curvature term, known to cause stability issues and limited convergence speed. Instead, we choose to follow the strategy proposed by *Bresson et al.* [2] in a related context. The crux is to introduce an auxiliary variable v and consider the following approximation of F_L :

$$\int_{\Omega} |\nabla u| + \frac{1}{2\theta} \int_{\Omega} |u - v|^2 + \int_{\Omega} v r_1 + \int_{\Omega} (1 - v) r_2$$
(21)

where r_1 and r_2 are given by (17) and θ is chosen to be small enough so that the two components of any minimizing couple (u^*, v^*) are almost identical. In that form, the dependence on u is restricted to the first two terms, which are exactly the terms of the minimization problem solved by *Chambolle* with a dual approach in the context of denoising [4]. Thus his fast and remarkably stable projection algorithm can be used to minimize w.r.t. u while the other variables are kept fixed. Now, we only need to find optimal solutions for p_1 , p_2 and v taken independently. It turns out that those solutions can be directly obtained, without additional iterative schemes. Indeed, optimal p_1^* and p_2^* are the normalized convolutions (19), replacing u by v. The opti-



Figure 4. Color images: the global non-parametric model in the multivariate case. The images are samples from the Berkeley segmentation database [14]. Feature space \mathcal{A} is CIE L*a*b*, chosen for its ability to mimic the logarithmic response of the eye and linearize the perception of color differences. Simple euclidean distance can be used in \mathcal{A} , hence the covariance M for the kernel K can be set diagonal. In all cases, initialization consists in setting initial u to a square region partly encompassing the foreground. The algorithm takes advantage of all channels and discriminates between regions that have a compact and multi-modal distribution in the 3D feature space (L,a,b).



Figure 5. A medical application: the hierarchical segmentation of a 3D Computerized Tomography volume. The global model is used to extract the boundaries of the lung from the multi-modal background, here composed of all surrounding organs. Initial u is set to a small cube inside the lungs. At convergence, the exact same process is recursively applied to the foreground, to separate pulmonary vessels from tissues. If distributions are expected to vary in space, which is a valid assumption for the vascular tree, the local competition model can provide more precise boundaries. Thus it is finally used to refine the result. On the left, coronal and axial views of the volume, with overlaid final parenchyma and vessel contours. Corresponding triangular meshes are rendered on the right: in the first column, the outer surface, alone (top) and mixed with the final vessels (bottom); in the second column, the vessels before (top) and after (bottom) local refinement.

mal v^* is given by [2]:

$$v^*(x) = \min\left\{\max\{0, u(x) - \theta r(x)\}, 1\right\}$$
(22)

where $r = r_1 - r_2$ is the competition function.

6. Conclusion

We introduced a novel variational approach for the partitioning of *n*-dimensional vectorial images into two regions based on the non-parametric estimation of local feature probability densities in each region. The crux is to use a modified formulation of the region competition functional whose arguments are not only the partition but also the probability density functions themselves, without any parametric representation. In other words, the new functional simultaneously controls both the non-parametric estimation and the segmentation. The optimality conditions for the unknown distributions naturally justify the use of continuous non-parametric Parzen density estimation techniques.

To overcome the limitations of the minimization over the non-convex set of boundaries or regions, we also presented an alternative componentwise-convex formulation over the set of fuzzy membership functions. The minimization is less sensitive to initial conditions and can be carried out using fast and stable algorithms.

Finally, we developed an extension of the same principles when the unknown probability density functions are allowed to vary in space. The segmentation is expected to be more accurate and robust to low-frequency artifacts such as illumination changes. The optimal local probability density functions involve convolutions, hence spatial smoothing ensures the regularity of the solutions. Moreover, this smoothing is selective and averages only values of the same region, so that a piecewise-smooth approximation of the local image intensity distribution can be obtained.

The wide range of applications and the performance of these new methods have been demonstrated on various images of different nature: synthetic intensity, natural color images and 3D medical volumes.

References

- S. P. Awate, T. Tasdizen, and R. T. Whitaker. Unsupervised texture segmentation with nonparametric neighborhood statistics. In *European Conference on Computer Vision*, *Proceedings*, pages 494–507, 2006.
- [2] X. Bresson, S. Esedoglu, P. Vandergheynst, J.-P. Thiran, and S. Osher. Fast global minimization of the active contour/snake model. *Jour. of Math. Imaging and Vision*, 2007.
- [3] T. Brox, M. Rousson, R. Deriche, and J. Weickert. Unsupervised segmentation incorporating colour, texture, and motion. *Computer Analysis of Images and Patterns. Lecture Notes in Computer Science*, 2756:353–360, 2003.

- [4] A. Chambolle. An algorithm for total variation minimization and applications. *Jour. of Math. Imaging and Vision*, 20(1-2):89–97, 2004.
- [5] T. Chan, S. Esedoglu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. UCLA CAM Report 04-54, 2004.
- [6] T. Chan and L. Vese. Active contours without edges. *IEEE Trans. on Image Proc.*, 10(2):266–277, February 2001.
- [7] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *Int. J. Comput. Vision*, 72(2):195– 215, 2007.
- [8] A. Herbulot, S. Jehan-Besson, M. Barlaud, and G. Aubert. Shape gradient for multi-modal image segmentation using mutual information. *Proc. of IEEE Int. Conference on Image Processing (ICIP)*, 10(8):2729–2732, October 2004.
- [9] A. Herbulot, S. Jehan-Besson, S. Duffner, M. Barlaud, and G. Aubert. Segmentation of vectorial image features using shape gradients and information measures. *Jour. of Math. Imaging and Vision*, 25(3):365–386, October 2006.
- [10] S. Jehan-Besson, M. Barlaud, and G. Aubert. Dream2s: Deformable regions driven by an eulerian accurate minimization method for image and video segmentation. *Int. Journal* of Computer Vision, 53(1):45–70, 2003.
- [11] T. Kadir and M. Brady. Unsupervised non-parametric region segmentation using level sets. 2:1267–1274, 2003.
- [12] J. Kim, J. Fisher, A.Yezzi, M. Cetin, and A. Willsky. A nonparametric statistical method for image segmentation using information theory and curve evolution. *IEEE Trans. on Image Processing*, 14:1486–1502, 2005.
- [13] H. Knutsson and C. Westin. Normalized and differential convolution: Methods for interpolation and filtering of incomplete and uncertain data. In *Proc. of Computer Vision and Pattern Recognition*, pages 515–523, 1993.
- [14] D. R. 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. pages 416–425, 2001.
- [15] B. Mory and R. Ardon. Fuzzy region competition: A convex two-phase segmentation framework. In *International Conference on Scale-Space and Variational Methods in Computer Vision, Proceedings*, pages 214–226, 2007.
- [16] N. Paragios and R. Deriche. Geodesic active regions and level set methods for supervised texture segmentation. *Int. Journal of Computer Vision*, 46(3), 2002.
- [17] E. Parzen. On the estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33:1065–1076, 1962.
- [18] S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *IEEE Trans. On Pattern Analysis and Machine Intelligence*, 18(9):884–900, September 1996.