



A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape

DANIEL CREMERS

Department of Computer Science, University of Bonn, Germany

dcremers@cs.uni-bonn.de

MIKAEL ROUSSON

Department of Imaging and Visualization, Siemens Corporate Research, Princeton, NJ, USA

mikael.rousson@scr.siemens.com

RACHID DERICHE

INRIA, Sophia-Antipolis, France

rachid.deriche@sophia.inria.fr

Received February 14, 2005; Revised April 6, 2006; Accepted April 6, 2006

First online version published in August 2006

Abstract. Since their introduction as a means of front propagation and their first application to edge-based segmentation in the early 90's, level set methods have become increasingly popular as a general framework for image segmentation. In this paper, we present a survey of a specific class of region-based level set segmentation methods and clarify how they can all be derived from a common statistical framework.

Region-based segmentation schemes aim at partitioning the image domain by progressively fitting statistical models to the intensity, color, texture or motion in each of a set of regions. In contrast to edge-based schemes such as the classical Snakes, region-based methods tend to be less sensitive to noise. For typical images, the respective cost functionals tend to have less local minima which makes them particularly well-suited for local optimization methods such as the level set method.

We detail a general statistical formulation for level set segmentation. Subsequently, we clarify how the integration of various low level criteria leads to a set of cost functionals. We point out relations between the different segmentation schemes. In experimental results, we demonstrate how the level set function is driven to partition the image plane into domains of coherent color, texture, dynamic texture or motion. Moreover, the Bayesian formulation allows to introduce prior shape knowledge into the level set method. We briefly review a number of advances in this domain.

Keywords: image segmentation, level set methods, Bayesian inference, color, texture, motion

1. Introduction

The goal of image segmentation is to partition the image plane into meaningful areas, where *meaningful* typically refers to a separation of areas corresponding to

different objects in the observed scene from the area corresponding to the background.

A large variety of segmentation algorithms have been proposed over the last few decades. While earlier approaches were often based on a set of rather

heuristic processing steps (cf. Perkins, 1980), optimization methods have become established as more principled and transparent methods: Segmentations of a given image are obtained by minimizing appropriate cost functionals. Among optimization methods, one can distinguish between *spatially discrete* and *spatially continuous* representations.

In spatially discrete approaches, the pixels of the image are usually considered as the nodes of a graph, and the aim of segmentation is to find cuts of this graph which have a minimal cost. Optimization algorithms for these problems include greedy approaches such as the *Iterated Conditional Modes* (ICM) (Besag, 1986) and continuation methods such as *Simulated Annealing* (Geman and Geman, 1984) or *Graduated Non-convexity* (Blake and Zisserman, 1987). Graph cut algorithms for specific classes of cost functions have gained in popularity with the re-discovery of efficient global optimization methods, which are based on concepts of maximum network flow (Boykov et al., 2001), on spectral methods (Shi and Malik, 1997; Malik et al., 2001), on semidefinite programming techniques (Keuchel et al., 2003), or on random walks and potential theory (Grady, 2006).

In spatially continuous approaches, the segmentation of the image plane $\Omega \subset \mathbb{R}^2$ is considered as a problem of infinite-dimensional optimization. Using variational methods, one computes segmentations of a given image $I : \Omega \rightarrow \mathbb{R}$ by evolving contours in the direction of the negative energy gradient using appropriate partial differential equations (pdes). Such pde-based segmentation methods became popular with the seminal paper on *Snakes* by Kass et al. (1988). In this paper, the contour is implemented by an explicit (parametric) curve $C : [0, 1] \rightarrow \Omega$ which is evolved by locally minimizing the functional

$$E(C) = - \int |\nabla I(C)|^2 ds + \nu_1 \int |C_s|^2 ds + \nu_2 \int |C_{ss}|^2 ds, \quad (1)$$

where C_s and C_{ss} denote the first and second derivative with respect to the curve parameter s . The first term in (1) is the external energy which accounts for the image information, in the sense that the minimizing contour will favor locations of large image gradient. The last two terms—weighted by nonnegative parameters ν_1 and ν_2 —can be interpreted as an *internal energy* of the contour, measuring the length of the contour and its stiffness or rigidity.¹

The Snakes approach had an enormous impact in the segmentation community (with around 4000 citations to date). Yet, it suffers from several drawbacks:

1. The implementation of contour evolutions based on an explicit parameterization requires a delicate re-gridding (or reparameterization) process to avoid self-intersection and overlap of control or marker points.
2. The explicit representation by default does not allow the evolving contour to undergo topological changes such that the segmentation of several objects or multiply-connected objects is not straight-forward.²
3. The segmentations obtained by a local optimization method are bound to depend on the initialization. The Snake algorithm is known to be quite sensitive to the initialization. For many realistic images, the segmentation algorithm tends to get stuck in undesired local minima—in particular in the presence of noise.
4. The Snakes approach lacks a meaningful probabilistic interpretation. Extensions to other segmentation criteria—such as color, texture or motion—are not straight-forward.

In the present paper, we will review recent developments in the segmentation community which aim at resolving the above problems. We will review the level set method for front propagation as a means to handle topological changes of evolving interfaces and to remove the issues of contour parameterization and control point re-gridding. Among the level set methods, we will focus on statistical region-based methods, where the contour is not evolved by fitting to local gradient information (as in the Snakes) but rather by fitting statistical models to intensity, color, texture or motion within each of the separated regions. The respective cost functionals tend to have less local minima for most realistic images. As a consequence, the segmentation schemes are far less sensitive to noise and to varying initialization.

The outline of the paper is as follows: In Section 2, we will review the general idea of level set based boundary propagation and its first applications to image segmentation. In Section 3, we will then review a probabilistic formulation of region-based segmentation. In particular, we will make very explicit what are the assumptions underlying the derivation of appropriate cost functionals. In the subsequent sections, we then detail how to adapt the probabilistic level set framework to different segmentation criteria: In Section 4,

we present probabilistic models which drive the segmentation process to group regions of homogeneous intensity, color or texture. In Section 5, we briefly present extensions of this framework to Diffusion Tensor Images. In Section 6, we discuss a further extension which allows to exploit *spatio-temporal* dynamics to drive a segmentation process, given an entire sequence of images. In particular, this approach allows to separate textures which have identical spatial characteristics but differ in their temporal dynamics. In Section 7 we detail how to integrate motion information as a criterion for segmentation, leading to a partitioning of the image plane into areas of piecewise parametric motion. Finally, in Section 8, we briefly discuss numerous efforts to introduce statistical shape knowledge in level set based image segmentation in order to cope with missing or misleading low-level information.

2. Level Set Methods for Image Segmentation

In the variational framework, a segmentation of the image plane Ω is computed by locally minimizing an appropriate energy functional, such as the functional Eq. (1). The key idea is to evolve the boundary \mathcal{C} from some initialization in direction of the negative energy gradient, which is done by implementing the gradient descent equation:

$$\frac{\partial \mathcal{C}}{\partial t} = -\frac{\partial E(\mathcal{C})}{\partial \mathcal{C}} = F \cdot n, \quad (2)$$

modeling an evolution along the normal n with a speed function F .³

In general, one can distinguish between explicit (parametric) and implicit representations of contours. In explicit representations—such as splines or polygons—a contour is defined as a mapping from an interval to the image domain: $\mathcal{C} : [0, 1] \rightarrow \Omega$. The propagation of an explicit contour is typically implemented by a set of ordinary differential equations acting on the control or marker points. In order to guarantee stability of the contour evolution (i.e. preserve well-defined normal vectors), one needs to introduce certain regridding mechanisms to avoid overlap of control points, for example by numerically resampling the marker points every few iterations, by imposing in the variational formulation a rubber-band like attraction between neighboring points (Cremers et al., 2002c), or by introducing electrostatic repulsion (Unal et al., 2005). Moreover, in order to segment several objects

or multiply connected objects, one needs to introduce numerical tests to enable splitting and remerging of contours during the evolution. Successful advances in this direction were proposed among others by Leitner and Cinquin (1991), McInerney and Terzopoulos (1995), Lachaud and Montanvert (1999) and Delingette and Montagnat (2000).

In implicit contour representations, contours are represented as the (zero) level line of some embedding function $\phi : \Omega \rightarrow \mathbb{R}$:

$$\mathcal{C} = \{x \in \Omega \mid \phi(x) = 0\}. \quad (3)$$

There are various methods to evolve implicitly represented contours. The most popular among these is the level set method (Dervieux and Thomasset, 1979, 1981; Osher and Sethian, 1988), in which a contour is propagated by evolving a time-dependent embedding function $\phi(x, t)$ according to an appropriate partial differential equation. In the following, we will briefly sketch two alternative methods to derive a level set evolution implementing the minimization of the energy $E(\mathcal{C})$.

For a contour which evolves along the normal n with a speed F —see Eq. (2)—one can derive a corresponding partial differential equation for the embedding function ϕ in the following way. Since $\phi(\mathcal{C}(t), t) = 0$ at all times, the total time derivative of ϕ at locations of the contour must vanish:

$$\frac{d}{dt} \phi(\mathcal{C}(t), t) = \nabla \phi \frac{\partial \mathcal{C}}{\partial t} + \frac{\partial \phi}{\partial t} = \nabla \phi F \cdot n + \frac{\partial \phi}{\partial t} = 0. \quad (4)$$

Inserting the definition of the normal $n = \frac{\nabla \phi}{|\nabla \phi|}$, we get the evolution equation for ϕ :

$$\frac{\partial \phi}{\partial t} = -|\nabla \phi| F. \quad (5)$$

By derivation, this equation only specifies the evolution of ϕ (and the values of the speed function F) at the location of the contour. For a numerical implementation one needs to extend the right-hand side of (5) to the image domain away from the contour.

Alternatively to the above derivation, one can obtain a level set equation from a variational formulation (cf. Zhao et al., 1996; Chan and Vese, 2001) Rather than deriving an appropriate partial differential equation for ϕ which implements the contour evolution Eq. (2), one can embed a variational principle $E(\mathcal{C})$ defined on the space of contours by a variational principle $E(\phi)$

defined on the space of level set functions:

$$E(\mathcal{C}) \longrightarrow E(\phi)$$

Subsequently, one can derive the Euler-Lagrange equation which minimizes $E(\phi)$:

$$\frac{\partial \phi}{\partial t} = - \frac{\partial E(\phi)}{\partial \phi}. \quad (6)$$

In both cases, the embedding is not uniquely defined. Depending on the chosen embedding, one can obtain slightly different evolution equations for $\phi(x, t)$.

The first applications of this level set formalism for the purpose of image segmentation were proposed in Caselles et al. (1993) and Malladi et al. (1994a, b). Independently, Caselles et al. (1995) and Kichenassamy et al. (1995) proposed a level set formulation for the Snake energy Eq. (1) given by:

$$\begin{aligned} \frac{\partial \phi}{\partial t} &= |\nabla \phi| \operatorname{div} \left(g(I) \frac{\nabla \phi}{\|\nabla \phi\|} \right) \\ &= g(I) |\nabla \phi| \operatorname{div} \left(\frac{\nabla \phi}{\|\nabla \phi\|} \right) + \nabla g(I) \cdot \nabla \phi, \end{aligned} \quad (7)$$

where the gradient $|\nabla I|$ in functional (1) was replaced by a more general edge function $g(I)$. This approach is known as *Geodesic Active Contours*, because the underlying energy can be interpreted as the length of a contour in a Riemannian space with a metric induced by the image intensity. See Caselles et al. (1995) and Kichenassamy et al. (1995) for details.

Local optimization methods such as the Snakes have been heavily criticized because the computed segmentations depend on the initialization and because algorithms are easily trapped in undesired local minima for many realistic images. In particular in the presence of noise, numerous local minima of the cost functional (1) are created by local maxima of the image gradient. To overcome these local minima and to drive the contour toward the boundaries of objects of interest, researchers have introduced an additional balloon force (Cohen and Cohen, 1993) which leads to either a shrinking or an expansion of contours. Unfortunately this requires prior knowledge about whether the object of interest is inside or outside the initial contour. Moreover, the final segmentation will be biased toward smaller or larger segmentations.

In the following, we will review a probabilistic formulation of the segmentation problem which leads to

region-based functionals rather than edge-based functionals such as the Snakes. Moreover, we will provide numerous experiments which demonstrate that such probabilistic region-based segmentation schemes do not suffer from the above drawbacks. While optimization is still done in a local manner, the respective functionals tend to have few local minima and segmentation results tend to be very robust to noise and varying initialization.

3. Statistical Formulation of Region-Based Segmentation

3.1. Image Segmentation as Bayesian Inference

Statistical approaches to image segmentation have a long tradition, they can be traced back to models of magnetism in physics, such as the Ising model (Ising, 1925), pioneering works in the field of image processing include spatially discrete formulations such as those of Geman and Geman (1984) and Besag (1986), and spatially continuous formulations such as the ones of Mumford and Shah (1989) and Zhu and Yuille (1996).

The probabilistic formulation of the segmentation problem presented in the following extends the statistical approaches pioneered in Leclerc (1989), Zhu and Yuille (1996), Paragios and Deriche (2002) and Tsai et al. (2001). In particular, this extension allows the probabilistic framework to be applied to segmentation criteria such as texture and motion, which will be detailed in subsequent sections. In Leclerc (1989), a segmentation functional is obtained from a Minimum Description Length (MDL) criterion. The link with the Mumford-Shah functional and the equivalence to Bayesian maximum *a posteriori* (MAP) estimation is provided in Zhu and Yuille (1996). Following (Paragios and Deriche, 2002), an optimal partition $\mathcal{P}(\Omega)$ of the image plane Ω (i.e. a partition of the image plane into pairwise disjoint regions) can be computed by maximizing the *a posteriori* probability $p(\mathcal{P}(\Omega) | I)$ for a given image I .⁴ The Bayes rule permits to express this conditional probability as

$$p(\mathcal{P}(\Omega) | I) \propto p(I | \mathcal{P}(\Omega)) p(\mathcal{P}(\Omega)), \quad (8)$$

thereby separating image-based cues (first term) from geometric properties of the partition (second term). The Bayesian framework has become increasingly popular to tackle many ill-posed problems in computer vision. First the conditional probability $p(I | \mathcal{P}(\Omega))$ of an

observation given a model state is often easier to model than the posterior distribution, it typically follows from a generative model of the image formation process. Secondly, the term $p(\mathcal{P}(\Omega))$ in Eq. (8) allows to introduce prior knowledge stating which interpretations of the data are *a priori* more or less likely. Wherever available, such *a priori* knowledge may help to cope with missing low-level information.

One can distinguish between generic priors and object specific priors. Object specific priors can be computed from a set of sample segmentations of an object of interest. In Section 8, we will briefly review a number of recent advances regarding the incorporation of statistically learnt priors into the level set framework.

In this section, we will focus on generic (often called *geometric*) priors. The most commonly used regularization constraint is a prior which favors a short length $|\mathcal{C}|$ of the partition boundary:

$$p(\mathcal{P}(\Omega)) \propto e^{-\nu|\mathcal{C}|}, \quad \nu > 0. \quad (9)$$

Higher-order constraints may be of interest for specific applications such as the segmentation of thin elongated structures (Rochery et al., 2006; Nain et al., 2003).

To further specify the image term $p(I | \mathcal{P}(\Omega))$ in (8), we further assume the image partition to be composed of N regions without correlation between labellings. This gives the simplified expression:

$$p(I | \mathcal{P}(\Omega)) = p(I | \{\Omega_1, \dots, \Omega_N\}) = \prod_{i=1}^N p(I | \Omega_i), \quad (10)$$

where $p(I | \Omega_i)$ denotes the probability of observing an image I when Ω_i is a region of interest.

We will now assume that the observation I consists of a set of feature values $f(x)$ associated with each image location. This feature may be a scalar quantity (such as the image intensity), a vector quantity (such as the pixel color or the spatio-temporal image gradient), or a tensor (such as the structure tensor or a diffusion tensor). This generalization from local intensities to local features will allow us to extend the statistical segmentation framework to texture, motion, etc.

For the features presented in this paper, we make the assumption that the values of f at different locations of the same region can be modeled as independent and identically distributed realizations of the same random process.⁵ Let p_i be the probability density function (*pdf*) of this random process in Ω_i . Expression (10)

then reads

$$p(I | \mathcal{P}(\Omega)) = \prod_{i=1}^N \prod_{x \in \Omega_i} (p_i(f(x)))^{dx}, \quad (11)$$

where the bin volume dx is introduced to guarantee the correct continuum limit. Approximation (11) is not valid in general since image features (such as spatial gradients) are computed on a neighborhood structure and may therefore exhibit local spatial correlations. More importantly, one should expect to find spatial correlations of features when modeling textured regions. However, one can capture certain spatial correlations in the above model by computing appropriate features such as the structure tensor.

Maximization of the *a posteriori* probability (8) is equivalent to minimizing its negative logarithm. Integrating the regularity constraint (9) and the region-based image term (11), we end up with following energy:

$$E(\{\Omega_1, \dots, \Omega_N\}) = - \sum_i \int_{\Omega_i} \log p_i(f(x)) dx + \nu |\mathcal{C}|. \quad (12)$$

In the context of intensity segmentation (i.e. $f = I$), this energy is the basis of several works (Leclerc, 1989; Zhu and Yuille, 1996; Samson et al., 2000; Paragios and Deriche, 2002). The region statistics are typically computed interlaced with the estimation of the boundary \mathcal{C} (Zhu and Yuille, 1996), yet one can also compute appropriate intensity histograms beforehand (Paragios and Deriche, 2002). In this paper, we will focus on the case that distributions and segmentation are computed jointly. Distributions can be either modeled as parametric or non-parametric ones. Upon insertion of parametric representations for p_i with parameters θ_i , the energy (12) takes on the form

$$E(\{\Omega_i, \theta_i\}_{i=1..N}) = - \sum_i \int_{\Omega_i} \log p(f(x) | \theta_i) dx + \nu |\mathcal{C}|. \quad (13)$$

For particular choices of parametric densities, the optimal parameters can be expressed as functions of the corresponding domains and only the regions remain as

unknowns of the new energy

$$\begin{aligned} \hat{E}(\{\Omega_i\}_i) &\equiv \min_{\{\theta_i\}} E(\{\Omega_i, \theta_i\}_i) \\ &= - \sum_i \int_{\Omega_i} \log p(f(x) | \hat{\theta}_i) dx + \nu |\mathcal{C}|, \end{aligned} \quad (14)$$

such that

$$\hat{\theta}_i = \arg \min_{\theta} \left(- \int_{\Omega_i} \log p(f(x) | \theta) dx \right). \quad (15)$$

In this case, the optimal model parameters $\hat{\theta}_i$ typically depend on the regions Ω_i . As pointed out by several authors (Sokolowski and Zolesio, 1991; Schnoerr, 1992; Aubert et al., 2003), this region-dependence can be taken into account in the computation of accurate shape gradients. Exact shape gradients can also be applied with non-parametric density estimation techniques like the Parzen window method (Kim et al., 2002; Rousson et al., 2003; Herbulot et al., 2004). In Rousson and Deriche (2002), it is shown that no additional terms arise in the shape gradient if the distributions p_i are assumed to be Gaussian. And in Heiler and Schnoerr (2003), the authors point out that the additional terms are negligible in the case of Laplacian distributions.⁶ We will therefore neglect higher-order terms in the computation of shape gradients and simply perform an alternating minimization of the energy (13) with respect to region boundaries and region models.

3.2. Two-Phase Level Set Formulation

Let us for the moment assume that the solution to (13) is in the class of binary (two-phase) segmentations, i.e. a partitioning of the domain Ω such that each pixel is ascribed to one of two possible phases. Extending the approach of Chan and Vese (2001), one can implement the functional (13) by:

$$\begin{aligned} E(\phi, \{\theta_i\}) &= \int_{\Omega} -H\phi \log p(f | \theta_1) \\ &\quad - (1 - H\phi) \log p(f | \theta_2) + \nu |\nabla H\phi| dx, \end{aligned} \quad (16)$$

where $H\phi$ denotes the heaviside step function defined as:

$$H\phi \equiv H(\phi) = \begin{cases} 1 & \text{if } \phi \geq 0 \\ 0 & \text{else} \end{cases}. \quad (17)$$

The first two terms in (16) model the areas inside and outside the contour while the last term represents the length of the separating interface.

Minimization is done by alternating a gradient descent for the embedding function ϕ (for fixed parameters θ_i):

$$\frac{\partial \phi}{\partial t} = \delta(\phi) \left(\nu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \log \frac{p(f(x) | \theta_1)}{p(f(x) | \theta_2)} \right). \quad (18)$$

with an update of the parameters θ_i according to (15). In practice, the delta function δ is implemented by a smooth approximation—(cf. Chan and Vese, 2001).

3.3. Multiphase Level Set Formulation

Several authors have proposed level set formulations which can handle a larger number of phases (Zhao et al., 1996; Yezzi et al., 1999; Paragios and Deriche, 2002). These methods use a separate level set function for each region. This clearly increases the computational complexity. Moreover, numerical implementations are somewhat involved since the formation of overlap and vacuum regions needs to be suppressed. By interpreting these overlap regions as separate regions, Vese and Chan derived an elegant formulation which only requires $\log_2(n)$ level set functions to model n regions. Each of the n regions is characterized by the various level set functions being either positive or negative (See Vese and Chan (2002) for details). An efficient variation of this idea was developed in Brox and Weickert (2004).

3.4. Scalar, Vector and Tensor-Valued Images

3.4.1. Scalar Images. Let us consider a scalar image made up of two regions, the intensities of which are drawn from a Gaussian distribution:

$$p(I | \mu_i, \sigma_i^2) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(I-\mu_i)^2}{2\sigma_i^2}}, \quad i = \{1, 2\}. \quad (19)$$

This distribution can be injected in the general bi-partitioning energy (16). Given a partition of the image

plane according to a level set function ϕ , optimal estimates for the mean μ_i and the variance σ_i can be computed analytically:

$$\begin{cases} \mu_1 = \frac{1}{a_1} \int H(\phi)I(x)dx \\ \sigma_1^2 = \frac{1}{a_1} \int H(\phi)(I(x) - \mu_1)^2 dx, \\ \mu_2 = \frac{1}{a_2} \int (1 - H(\phi))I(x)dx \\ \sigma_2^2 = \frac{1}{a_2} \int (1 - H(\phi))(I(x) - \mu_2)^2 dx. \end{cases}$$

where $a_1 = \int H(\phi)dx$ and $a_2 = \int (1 - H(\phi))dx$ are the areas of the inside and outside region. For fixed model parameters, the gradient descent equation for the level set function ϕ —see Eq. (18)—reads

$$\frac{\partial \phi}{\partial t} = \delta(\phi) \left(\nu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \frac{(I - \mu_2)^2}{2\sigma_2^2} - \frac{(I - \mu_1)^2}{2\sigma_1^2} + \log \frac{\sigma_2}{\sigma_1} \right). \quad (20)$$

More details on this derivation when the parameters (μ_i, σ_i) are taken as functions of Ω_i can be found in Rousson and Deriche (2002). We end up with an algorithm that alternates the estimation of the empirical intensity means and variances inside each region and the level set evolution described by Eq. (20). Regarding the complexity, each iteration of the level set evolution is applied only inside a narrow band around the zero-crossing because the Dirac function is equal to zero at other locations. More interesting is that the statistical parameters can also be updated with a similar complexity: new updates are functions of their previous values and of the pixels where the sign of ϕ changes. Assuming the evolving interface to visit each pixel only once, the total complexity is thus linear in the size of the image.

3.4.2. Vector-Valued Images. A direct extension to vector-valued images is to use *multivariate* Gaussian densities as region models. Region *pdfs* are then parameterized by a vector mean and a covariance matrix. Similarly to the scalar case, the optimal statistical parameters are their empirical estimates in the corresponding region. The 2-phase segmentation of an image I of any dimension can thus be obtained through

the following level set evolution—see also (18):

$$\frac{\partial \phi}{\partial t} = \delta(\phi) \left(\nu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + \log \frac{p(I(x) | \mu_1, \Sigma_1^2)}{p(I(x) | \mu_2, \Sigma_2^2)} \right), \quad (21)$$

with:

$$\begin{cases} \mu_i = \frac{1}{|\Omega_i|} \int_{\Omega_i} I(x) dx, \\ \Sigma_i = \frac{1}{|\Omega_i|} \int_{\Omega_i} (I(x) - \mu_i)(I(x) - \mu_i)^\top dx \text{ for } i = 1, 2. \end{cases} \quad (22)$$

As in the scalar case, the estimation of the statistical parameters can be optimized to avoid a full computation over the whole image domain at each iteration. Here, it becomes a bit more technical since cross-components products appear in the covariance matrices but the final complexity is identical to the one obtained in the scalar case (Rousson, 2004).

3.4.3. Tensor-Valued Images. In order to apply the above statistical level set framework to the segmentation of tensor images, one needs to define appropriate distances on the space of tensors. Several approaches have been proposed to define distances from an information theoretic point of view by interpreting the tensors as parameterizations of 0-mean multivariate normal laws. The definition of a distance between tensors is then translated to one of a dissimilarity measure between probability distributions.

The Symmetric KL Divergence. Wang and Vemuri (2004) applied the symmetrized Kulback-Leibler (SKL) divergence—also called J-divergence—to define the region term of the front evolution. For multivariate 0-mean normal laws with covariance matrices J_1 and J_2 , the SKL divergence is given by:

$$\mathcal{D}(J_1, J_2)_{SKL} = \frac{1}{2} \sqrt{\operatorname{trace}[J_1^{-1} J_2 + J_2^{-1} J_1] - 2n}, \quad (23)$$

where n is the dimension of the tensors. This measure has the advantage of being affine invariant and closed form expressions are available for the mean tensors which is particularly interesting to estimate region statistics. Region confidence were also incorporated in Rousson et al. (2004). These works present several promising segmentation experiments on 2D (Wang and

Vemuri, 2004) and 3D (Rousson et al., 2004) real diffusion tensor images.

The Rao Distance. Another distance has been proposed in Lenglet et al. (2004) and Lenglet et al. (2004) with the same idea of considering tensors as covariance matrices of multivariate normal distributions. Following (Skovgaard, 1994), a Riemannian metric is introduced and the geodesic distance between two members of this family is given by

$$\mathcal{D}_G(J_1, J_2) = \sqrt{\frac{1}{2} \sum_{i=1}^2 \log^2(\lambda_i)}, \quad (24)$$

where λ_i denote the eigenvalues of the matrix $J_1^{-1/2} J_2 J_1^{-1/2}$. The same metric was proposed in Pennec et al. (2005) from a different viewpoint. It verifies the basic properties of a distance (positivity, symmetry and triangle inequality) and it is invariant to inversions: $\mathcal{D}(J_1, J_2) = \mathcal{D}(J_1^{-1}, J_2^{-1})$. The above metrics permit to define statistics on sets of SPD matrices which can be used to define the region term of the segmentation. It was also shown in Lenglet et al. (2004) that the (asymmetric) Kullback Leibler divergence 23 is a Taylor approximation of the geodesic distance (24).

In the following sections, we will exploit the statistical level set framework introduced above to construct segmentation schemes for color, texture, dynamic texture and motion. To this end, we will consider different choices regarding the features f —namely intensity values, color values, spatial structure tensors, spatio-temporal image gradients, or features modeling the local spatio-temporal dynamics—and respective sets of model parameters θ_i , modeling color or texture distributions or parametric motion in the separated regions. Moreover, we will consider different choices for the distributions p_i of these model parameters.

4. Intensity, Color and Texture

In the previous section, we considered Gaussian approximations for scalar and vector-values images. These models can be used to segment gray, color and texture images (Chan and Vese, 2001; Rousson and Deriche, 2002; Rousson et al., 2003). In the following, these models are applied to the segmentation of natural images. Curve evolutions are presented to illustrate the gradient descent driving to the segmentation.

4.1. Gray & Color Images

In Fig. 1, we present the curve evolution obtained with the gray-value level set scheme of Section 3.4.1. The curve is initialized with a set of small circles and it successfully evolves toward the expected segmentation. Other initializations may be considered but using tiny circles provides a fast convergence speed and helps to detect small parts and holes. Note that changes of topology during the evolution are naturally handled by the implicit formulation.

We previously argued that the region-based formulation exhibits less local minima than approaches which solely rely on gradient information along the curve. To support this claim, we plotted in Fig. 2 the empirical energy for the segmentation of a 1D slice of an image. In contrast to the edge-based energy, the region-based energy (thick black line) shows a single minimum corresponding to the boundary of the coin.

The region-based approach can directly be extended to color images by applying the vector-valued formulation of Section 3.4.2. The only point to be careful about is the choice of color space for the multivariate Gaussian model to make sense. The RGB space is definitely not the best one since, as can be seen from the MacAdam ellipse, the perception of color difference is nonlinear in this space. The CIE-lab space has been designed to approximate this nonlinearity by




Figure 1. Curve evolution for the segmentation of a gray-level image using Gaussian intensity distributions to approximate region information.

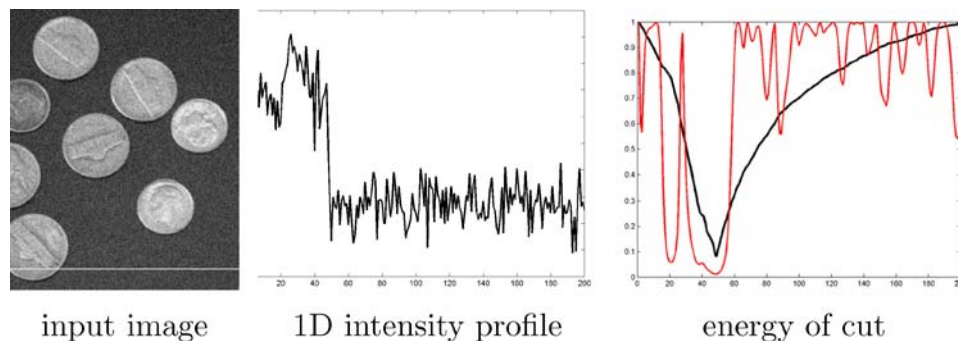


Figure 2. Comparison of edge- and region-based segmentation methods in 1D. For a 1D intensity profile of the coin image (taken along the line indicated in white), we computed the energy associated with a split of the interval at different locations. While the region-based energy exhibits a broad basin of attraction around a single minimum located at the boundary of the coin (thick black line), the energy of the edge-based approach is characterized by numerous local minima (red line). A gradient descent on the latter energy would not lead to the desired segmentation.



Figure 3. Binary segmentation of a color image using multivariate Gaussian distributions as region descriptor (initialization and final segmentation) and multiphase color segmentation obtained with the algorithm developed in Brox and Weickert (2004)

trying to mimic the logarithmic response of the eye. Figure 3 shows a two-phase and a multiphase example of vector-valued segmentation obtained on natural color images using this color space (the algorithm proposed in Brox and Weickert (2004) was used for the multiphase implementation).

4.2. Texture

In gray and color image segmentation, pixel values are assumed to be spatially independent. This is not the case for textured images which are characterized by local correlations of pixel values. In the following, we will review a set of basic features which allow to capture these local correlations. More sophisticated features are conceivable as well (Leung and Malik, 2001).

4.2.1. The Nonlinear Structure Tensor as Texture Feature. While texture analysis can rely on texture samples to learn accurate models (Hassner and Sklansky, 1980; Cross and Jain, 1993; Mallat, 1989; Simoncelli et al., 1992; Zhu et al., 1998), unsupervised image segmentation should learn these parameters on-line. Since high-order texture models introduce

many unknown parameters to be estimated in an unsupervised approach, more compact features are usually favored. Bigün et al. (1991) addressed this problem with the introduction of the structure tensor (also called second order moment matrix) which yields three different feature channels per scale. It has mainly been used to determine the intrinsic dimensionality of images in Bigün and Granlund (1987) and Förstner and Gülch (1987) by providing a continuous measure to detect critical points like edges or corners. Yet, the structure tensor does not only give a scalar value reflecting the probability of an edge but it also includes the texture orientation. All these properties make this matrix a good descriptor for textures. The structure tensor (Forstner and Gulch, 1987; Biguen et al., 1991; Rao and Schunck, 1991; Lindeberg, 1994; Granlund and Knutsson, 1995) is given by the matrix of partial derivatives smoothed by a Gaussian kernel K_σ with standard deviation σ :

$$J_\sigma = K_\sigma * (\nabla I \nabla I^\top) = \begin{pmatrix} K_\sigma * I_{x_1}^2 & K_\sigma * I_{x_1} I_{x_2} \\ K_\sigma * I_{x_1} I_{x_2} & K_\sigma * I_{x_2}^2 \end{pmatrix}. \quad (25)$$

For color images, all channels can be taken into account by summing the tensors of the individual channels (Di Zenzo, 1986).

Despite its good properties for texture discrimination, the structure tensor merely exploits (relative) spatial intensity variation (ignoring the absolute brightness information). In order to segment images with and without texture, a feature vector including the square root of the structure tensor and the intensity was defined in Rousson et al. (2003):

$$f(x) = \left(I, \frac{I_{x_1}^2}{|\nabla I|}, \frac{I_{x_2}^2}{|\nabla I|}, \frac{2I_{x_1}I_{x_2}}{|\nabla I|} \right)^T. \quad (26)$$

The major problem of the classic structure tensor is the dislocation of edges due to the smoothing with Gaussian kernels as shown in Fig. 4. To address this problem, Weickert and Brox (2002) proposed to replace the Gaussian smoothing by nonlinear diffusion, applying nonlinear matrix-valued diffusion schemes introduced in Tschumperlé and Deriche (2001). Applied on the feature vector f , this nonlinear diffusion couples all channels by a joint diffusivity, the information of all channels is used to decide whether an edge is worth to be enhanced or not, leading to the simplification of the data, the removal of outliers, and the closing of structures. Figure 4 shows the features obtained on the *zebra* image.

The feature vectors resulting from the nonlinear diffusion form a vector-valued image which can be

segmented using the vector-valued formulation presented in Section 3.4.2. Figure 5 shows a segmentation of the *zebra* image obtained with this method. For results on a wider range of texture images, we refer to Rousson et al. (2003).

The structure tensor is undoubtedly pertinent for texture discrimination but the approach developed so far still allows for further improvements. We shortly mention two recent extensions of this work in the following two paragraphs.

4.2.2. Scale Introduction Via TV Flow. With the nonlinear structure tensor, one mainly considers a single scale for the whole image. Yet textures often differ from one another with respect to their intrinsic scale. In order to account for varying scale, a straightforward extension is to combine texture features at different scales. While this modification may integrate information at different scales, it also increases dramatically the number of channels and redundant information is introduced, making the second phase—the image partitioning—more difficult. In order to work with a reduced feature space, Brox and Weickert (2004) proposed an elegant and efficient extension of the above framework, which combines similar texture features as above with a local scale measure. By exploiting the linear contrast reduction property of the TV (total variation) flow:

$$\begin{cases} \frac{\partial f}{\partial t} = \operatorname{div} \left(\frac{\nabla f}{\sqrt{|\nabla f|^2 + \epsilon^2}} \right) \\ f(t=0) = I \end{cases} \quad (27)$$



Figure 4. Left: Zebra image and color representation of its structure tensor (the components of the structure tensors are used as RGB components). Right: Intensity and structure tensor of the Zebra image after coupled nonlinear diffusion.




Figure 5. Curve evolution for the segmentation of a zebra image using the nonlinear structure tensor and the smoothed intensity (here, a rectangle is used as initialization but small circles also lead to a similar result).




Figure 6. Segmentation results obtained by running a level set segmentation process on the 5-dimensional feature space given by the features of the structure tensor, the image intensity and a local scale measure computed from the speed of a total variation flow. Images are courtesy of Brox and Weickert (2004).




Figure 7. Segmentation of a textured image with different dissimilarity measures between tensors. The left result was obtained using the Frobenius norm while the right segmentation is based on the Rao distance (see text for details). Images courtesy of de Luis García et al. (2005).

the authors are able to extract a local scale measure computed from the speed of the diffusion process. Upon combining this scale with the intensity and orientation features in (26), one can perform segmentation in a 5-dimensional feature space. Figure 6 shows three representative segmentation results.

4.2.3. Metric Between Tensors. While the previous approaches construct a feature vector from the components of the structure tensor and apply a vector-valued segmentation scheme, one can directly define metrics on the space of structure tensors (de Luís García and Deriche, 2005), for example the metrics defined in Section 3.4.3. The comparison in Fig. 7 shows that appropriate tensor distances lead to drastic improvements in the segmentation.

5. Diffusion Tensor Images

The problem of segmenting tensor-valued data also appears in medical imaging with the relatively new modality of diffusion tensor magnetic resonance images. In these images, a diffusion tensor is measured




Figure 8. Segmentation of corpus callosum from a diffusion tensor image using the geodesic distance in the manifold of multivariate normal distributions. This 3D segmentation was obtained in Lenglet et al. (2004) using the level set formulation presented in Section 2. Being implicit, the level set representation allows a straightforward extension to higher dimensions.

at each voxel. This tensor captures the local motion of water molecules as approximated by a Gaussian law. The metric between tensors described in Section 3.4.3 can be used to segment these images. Figure 8 shows a segmentation of the corpus callosum obtained with the geodesic distance.⁷

6. Dynamic Texture

The texture segmentation framework detailed earlier is based on assigning local texture signatures to each image location. The subsequent integration into a level set framework aims at optimally grouping regions of similar signatures while imposing a length constraint on the separating boundary.

Given a video sequence of temporally varying textures—such as smoke on water—one can extend this concept to the space-time domain and group regions of similar spatio-temporal statistics. The first work addressing this problem was proposed in Doretto et al.

(2003) where the authors made use of recent developments in the modeling of dynamic textures (Doretto et al., 2003). Due to the scope of this paper, we will merely review the key ideas.

Dynamic textures are models of temporally varying textures which assume the image sequence to be generated by a second-order stationary process. Experiments have demonstrated that numerous realistic image sequences, such as water waves, fluttering foliage, smoke and steam can be well synthesized by such Gauss-Markov processes (Doretto et al., 2003).

More specifically, it is assumed that the temporally varying pixel intensities $\{I_i(t)\}_{i=1..m}$ can be approximated by a model $\{y_i(t)\}_{i=1..m}$ which is driven by a random process $r(t) \in \mathbb{R}^n$ as follows (Doretto et al., 2003):

$$\begin{cases} r(t+1) = Ar(t) + \sqrt{Q} v(t); & r(0) = r_0 \\ y(t) = Cr(t) + \sqrt{R} w(t) \end{cases} \quad (28)$$

Here $y(t) \in \mathbb{R}^m$ represents the vector of intensities of all m pixels at time t , $v(t) \in \mathbb{R}^n$ and $w(t) \in \mathbb{R}^m$ are white zero-mean Gaussian processes, $A \in \mathbb{R}^{n \times n}$, $C \in \mathbb{R}^{m \times n}$ are the model parameters, and $Q \in \mathbb{R}^{n \times n}$, $R \in \mathbb{R}^{m \times m}$ are the noise covariance matrices. The model parameters A and C in (28) can be estimated from an image sequence $I(x, t)$ (Doretto et al., 2003).

As suggested in Doretto et al. (2003), we can associate with each image location $x \in \Omega$ a local signature $\xi(x)$ characterizing the spatio-temporal dynamics at this location based on the model parameters A and

C computed in a small spatial window. A meaningful signature cannot be directly defined on these model parameters because—as can be seen from the definition of (28)—there exists an entire equivalence class of model parameters which lead to the same dynamic texture⁸. Instead we define a local signature $\xi(x)$ by:

$$\xi(x) = (\cos \theta_1(x), \dots, \cos \theta_n(x)), \quad (29)$$

where $\{\theta_i\}_{i=1..n}$ are the subspace angles associated with the equivalence classes of the model at location x and some reference model. More precisely, if A and B are two measurement matrices, then $\{\theta_i\}_{i=1..n}$ are given by the principal angles (Decock and DeMoor, 2000) between $\text{range}(A)$ and $\text{range}(B)$. For details on the computation of these angles, we refer to Doretto et al. (2003).

Assuming that the spatio-temporal signatures defined in (29) correspond to two Gaussian distributions, one can apply the vector-valued segmentation scheme introduced in Section 3.4.2 to group areas of similar spatio-temporal dynamics.

Due to the scope of this survey paper, we will merely show two complementary results obtained by the above segmentation scheme. Figure 9 shows the separation based on spatial orientation of a moving texture. The image data shows a water sequence for which we simply rotated two areas by 90 degrees. By construction, intensity characteristics and dynamics of the separated regions are identical, yet due to the different orientation they can be separated. Figure 10 shows a segmentation result which is complementary to the previous one. We




Figure 9. Segmentation by texture orientation: Segmentation of two dynamic textures which differ only in orientation but share the same dynamics and general appearance (intensity values).




Figure 10. Segmentation by changing dynamics: The two dynamic textures are identical in appearance, but differ in the dynamics. This particular segmentation problem is quite difficult, even for human observers. Segmentation is obtained exclusively on the basis of the temporal properties of the textures.

generated a sequence containing regions which only differ with respect to their dynamics (but have identical spatial texture) by overlapping the ocean sequence in the regions corresponding to the disc and square over an ocean sequence slowed down by a factor of 2.

7. Motion

7.1. Motion as a Criterion for Segmentation

The central question underlying the construction of segmentation methods is to identify what properties characterize objects and distinguish them from other objects and from the background. In the previous sections, we reviewed level set methods which exploit low level properties such as color, texture or even dynamical texture. The respective image segmentation algorithms essentially group regions of similar low level properties.

Many objects in our environment are characterized by the fact that they move in a coherent manner. Figure 11, top row, shows the intensity-based segmentation of a single frame taken from an image sequence of two cars driving down the street.⁹ The two cars and the background are moving in different directions. Clearly the individual cars are not homogeneous regarding their intensity or texture. A purely intensity-based segmentation therefore fails to separate the objects from the background.

In the following, we will detail how the statistical segmentation scheme presented above can be adapted to incorporate motion information given two consecutive frames from an image sequence. Minimization of the resulting cost functional leads to a segmentation of the scene in terms of piecewise parametric motion.¹⁰ The present formulation was proposed in Cremers (2003) and Cremers and Soatto (2005) with an earlier (explicit contour) formulation in Cremers and Schnörr (2002). Related approaches were also proposed in Mansouri et al. (2002) and Paragios and Deriche (2005). The central idea is that we do not precompute local motion vectors. The computation of local motion vectors typically requires additional smoothness assumptions which are violated precisely at the motion boundaries we are interested in - see for example Wang and Adelson, 1994; Brox et al., 2003). Instead we jointly estimate the segmentation and the motion models for each of a set of regions by minimizing the proposed functional. In the notation introduced in Section 3, this means that—in contrast to the texture schemes of the previous sections—the model parameters θ_i (the motion models of the separate regions) will not correspond to simple aggregates of the local feature vectors $f(x)$ (the space time gradients), but rather they will be derived quantities. Due to the scope of this article, we will constrain the presentation to the key ideas. For further details and a discussion of related approaches we refer the reader to Cremers and Soatto (2005). For an extension of the proposed framework




Figure 11. Intensity versus motion segmentation. Since cars and background are not well-defined in terms of homogeneous intensity, color or texture, unsupervised low-level segmentation schemes based on a single frame are unable to separate objects and background (top row). By minimizing the motion competition functional (38) with $\nu = 1.5$, one obtains a fairly accurate segmentation of the two cars and an estimate of the motion of cars and background.

to the segmentation of space-time volumes given an entire video sequence, we refer to Cremers and Soatto (2003).

7.2. Motion Competition

Let $I : \Omega \times \mathbb{R} \rightarrow \mathbb{R}$ be a gray value image sequence. Denote the spatio-temporal image gradient of $I(x, t)$ by

$$\nabla_3 I = \left(\frac{\partial I}{\partial x_1}, \frac{\partial I}{\partial x_2}, \frac{\partial I}{\partial t} \right)^\top. \quad (30)$$

Let $v : \Omega \rightarrow \mathbb{R}^3$, $v(x) = (u(x), w(x), 1)^\top$ be the velocity vector at a point x in homogeneous coordinates.¹¹

Let us assume that the intensity of a moving point remains constant throughout time.¹² Expressed in differential form, this gives a relation between the spatio-temporal image gradient and the homogeneous velocity vector, known as the *optic flow constraint*:

$$\frac{dI}{dt} = \frac{\partial I}{\partial t} + \frac{\partial I}{\partial x_1} \frac{dx_1}{dt} + \frac{\partial I}{\partial x_2} \frac{dx_2}{dt} = v^\top \nabla_3 I = 0. \quad (31)$$

For the sake of segmentation, we will assume that the velocity in each of a set of regions can be modeled by a parametric motion of the form

$$v(x) = S(x) \cdot q, \quad (32)$$

with a space dependent matrix S and a parameter vector q . In particular, this includes the case of translational motion where S is the 3×3 unit matrix and $q = (u, w, 1)$ the vector of constant velocity in homogeneous coordinates. The parametric formulation (32) also includes the more general *affine motion model* with:

$$S(x) = \begin{pmatrix} x_1 & x_2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & x_1 & x_2 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \text{ and} \quad (33)$$

$$q = (a, b, c, d, e, f, 1)^\top$$

In the context of segmentation in space *and* time, this parametric formulation can be extended to incorporate temporally varying motion (such as acceleration and deceleration)—for details we refer to Cremers and Soatto (2003). Instead of piecewise parametric, one can also consider piecewise smooth motion (Amiaz and

Kiryati, 2005; Brox et al., 2006). While this allows to apply segmentation to non-rigidly moving regions, the motion estimation becomes a problem of (computationally more intense) infinite-dimensional optimization. From our experience segmentation processes with less restrictive region models typically exhibit more sensitivity with respect to the initialization.

Inserting the parametric model (32) into the optic flow constraint (31) leads to a constraint on the relation between the parameter vector q and the space-time gradient $\nabla_3 I$ at a specific location:

$$\nabla_3 I^\top S(x) q = 0. \quad (34)$$

Neglecting the case that the space-time gradient vanishes, this constraint states that the two vectors q and $S(x)^\top \nabla_3 I(x)$ must be orthogonal. We therefore model the conditional probability to encounter a certain gradient measurement given a velocity model as a function of the angle α between the two vectors:

$$P(\nabla_3 I | q; x) \propto \exp \left(- \frac{(q^\top S^\top \nabla_3 I)^2}{|q|^2 |S^\top \nabla_3 I|^2} \right). \quad (35)$$

This expression is maximal if the two vectors are indeed orthogonal, it is minimal if the two vectors are parallel. Yet it does not depend on the length of the two vectors.¹³ Note that due to the introduction of a spatially parametric model, this conditional probability becomes space-dependent—this is in contrast to the space-independent conditional probabilities considered in the Sections 4, 5 and 6 on color and texture segmentation. Analogous parametric generalizations of intensity-based level set segmentation approaches have been proposed by Vese (2003). And corresponding extensions of the above texture segmentation schemes are certainly conceivable.

Based on the Bayesian formulation introduced in Section 3, we can integrate the conditional probability for a measurement given certain model parameters into a variational framework for segmentation. Inserting Eq. (35) into energy (13), we obtain the functional

$$E(\mathcal{C}, \{q_i\}) = \sum_{i=1}^n \int_{\Omega_i} \frac{q_i^\top T(x) q_i}{|q_i|^2} dx + \nu |\mathcal{C}|, \quad (36)$$

where, for notational simplification, we have introduced the matrix

$$T(x) = \frac{\nabla_3 I S^\top S \nabla_3 I^\top}{|S^\top \nabla_3 I|^2}, \quad (37)$$

The corresponding two-phase level set implementation—cf. Eq. (16)—is given by

$$E(q_1, q_2, \phi) = \int_{\Omega} \frac{q_1^T T q_1}{|q_1|^2} H\phi + \frac{q_2^T T q_2}{|q_2|^2} (1 - H\phi) + \nu |\nabla H\phi| dx, \quad (38)$$

The first two terms in (38) enforce a homogeneity of the estimated motion in the two phases, while the last term enforces a minimal length of the region boundary given by the zero level set of ϕ . As discussed in 3.2, the functional (38) is optimized by alternating the estimate of the motion models q_1 and q_2 and an update of the level set function ϕ defining the motion boundaries.

For fixed level set function ϕ , i.e. fixed regions Ω_i , minimizing this functional with respect to the motion parameters $\{q_i\}$ results in a set of eigenvalue problems of the form:

$$q_i = \arg \min_q \frac{q^T T_i q}{q^T q}, \quad \text{with } T_i = \int_{\Omega_i} T(x) dx. \quad (39)$$

The parametric motion model q_i for each region Ω_i is therefore given by the eigenvector corresponding to the smallest eigenvalue of T_i . It is normalized, such that the third component is 1. Similar eigenvalue problems arise in motion estimation due to normalization with respect to the velocity magnitude (cf., Biguen et al., 1991; Jepson and Black, 1993).

Conversely, for fixed motion models q_i , a gradient descent on the energy (38) for the boundary \mathcal{C} results in an evolution equation—cf. (18)—of the form:

$$\frac{\partial \phi}{\partial t} = \delta(\phi) \left(\nu \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) + e_2 - e_1 \right). \quad (40)$$

where

$$\begin{aligned} e_i &= -\log P(\nabla_3 I | q_i; x) = \frac{q_i^T T q_i}{q_i^T q_i} \\ &= \frac{q_i^T \nabla_3 I S^T S \nabla_3 I^T q_i}{|q_i|^2 |S^T \nabla_3 I|^2} \end{aligned} \quad (41)$$

are the motion energy densities associated with the respective regions.

Note that—as in the previous sections – we have neglected in the evolution Eq. (40) higher-order terms which account for the dependence of the motion parameters q_i on the level set function ϕ . An Eulerian accurate shape optimization scheme as presented for

example in Jehan-Besson et al. (2003) is the focus of ongoing research.

The two terms in the contour evolution (40) have the following intuitive interpretation: The first term aims at minimizing the length of the separating motion boundary. The second term is proportional to the difference of the energy densities e_i in the regions adjoining the boundary: The neighboring regions compete for the boundary in terms of their motion energy density, thereby maximizing the motion homogeneity. Therefore this process is called *Motion Competition*.

7.3. Experimental Results

All image segmentation models are based on a number of more or less explicitly stated assumptions about the properties which define the objects of interest. The motion competition model is based on the assumption that objects are defined in terms of homogeneously moving regions. It extends the Mumford-Shah functional of piecewise constant intensity to a model of piecewise parametric motion. Despite this formal similarity, the segmentations generated by the motion competition framework are very different from those of its gray value analogue. Figure 11, bottom row, shows the boundary evolution obtained by minimizing the motion segmentation functional (38) and the corresponding motion estimates superimposed on the first frame. In contrast to its gray value analogue, the energy minimization simultaneously generates a fairly accurate segmentation of the two cars and an estimate of the motion of cars and background.

Figure 12 shows the contour evolution generated by minimizing functional (38) for two wall paper images with the text region (right image) moving to the right and the remainder of the image plane moving to the left. Even for human observers the differently moving regions are difficult to detect—similar to a camouflaged lizard moving on a similarly-textured ground. The gradient descent evolution superimposed on one of the two frames gradually separates the two motion regions without requiring salient features such as edges or Harris corner points. Figure 13 shows results obtained with a multiphase implementation of the motion competition functional. The static scene filmed by a moving camera is segmented into layers of different depth.

The functional (38) allows to segment piecewise affine motion fields. In particular, this class of motion models includes rotation and expansion/contraction.




Figure 12. Contour evolution obtained with functional (38) for $\nu = 0.06$, superimposed on one of the two input frames. The input images show the text region (right image) of the wallpaper moving right and the remainder moving left. The moving regions are accurately reconstructed, although the input images exhibit little in terms of salient features. The contour evolution took 10 seconds in Matlab.




Figure 13. Multiphase motion segmentation. Contour evolution for a multiphase implementation of motion competition on two consecutive frames from the flower garden sequence. A static scene filmed by a moving camera is partitioned into layers of different depth. See (Cremers and Soatto, 2005) for details.




Figure 14. Piecewise affine motion segmentation. Segmentations obtained by minimizing functional (38) with $\nu = 8 \cdot 10^{-5}$ for two image pairs showing a hand rotating (top) and moving toward the camera (bottom).

Figure 14 shows segmentations obtained for a hand in a cluttered background rotating (in the camera plane) and moving toward the camera. In this example the object of interest can be extracted from a fairly complex background based exclusively on their motion.

8. Statistical Shape Priors for Level Set Segmentation

In the previous sections, we reviewed a number of approaches which allow to drive the level set segmentation based on various low-level assumptions regarding the intensity, color, texture or motion of objects and background. In numerous real-world applications, these approaches may fail to generate the desired segmentations, because the respective assumptions about the low-level properties are either insufficient or even violated. In certain medical images for example, object and background may exhibit very similar intensity char-

acteristics. Moreover, the observed intensity or color of a 3D object may not be uniform due to directional lighting and cast shadows. And finally, misleading low-level information may arise due to noise or partial occlusion of the objects of interest. While the generic constraint Eq. (1) on the length of the segmenting boundary helps to cope with a certain amount of noise, it does introduce a bias toward contours of smaller length, thereby rounding corners or suppressing small scale details.

Beyond simple geometric regularity, the Bayesian formulation of the image segmentation problem allows to introduce higher-level prior knowledge about the shape of expected objects. This idea was pioneered by Grenander et al. (1991). In the following, we will briefly list some of the key contributions in the field of shape priors for level set segmentation.

The first application of shape priors for level set segmentation was developed by Leventon et al. (2000) who propose to perform principal component analysis on a set of signed distance function embedding a set of sample shapes. The distance functions are sampled

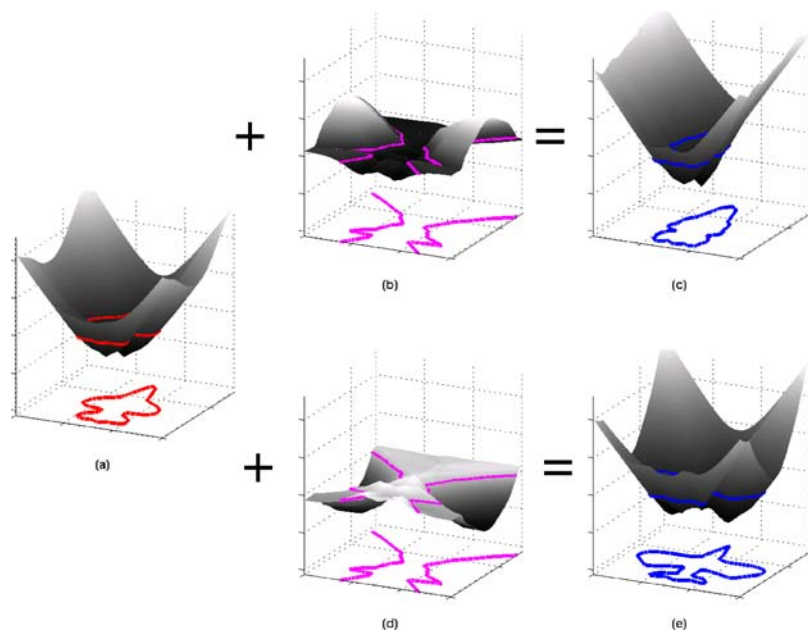


Figure 15. Visualization of principal component analysis on the level set function. The images show the mean level set function (obtained on a set of airplane shapes), and its deformation along the first eigenmode. Image data courtesy of Tsai et al. (2003).

on a regular grid to obtain a vector representation. A term is added to the contour evolution equation to drive the embedding function to the most likely shape of the estimated distribution. Tsai et al. (2001, 2003) proposed a very efficient implementation of shape-driven level set segmentation by directly optimizing in the linear subspace spanned by the principal components. A detailed analysis of various shape distances and statistical shape analysis in the level set formulation can be found in Charpiat et al. (2005). Figure 15 shows the effect of variation along the first principal component on the embedding function and the implicitly represented contour.

The use of principal component analysis to model level set based shape distributions has two limitations: Firstly, the space of signed distance functions is not a linear space, i.e. arbitrary linear combinations of signed distance functions will in general not correspond to a signed distance function. Secondly, while the first few principal components capture (by definition) the most variation on the space of embedding functions, they will not necessarily capture the variation on the space of the embedded contours. As a consequence, one may need to include a larger number of eigenmodes (compared to PCA on explicit contours) in order to capture certain details of the modeled shape. Nevertheless, we found the PCA representation to work

fairly well in practical applications. An alternative linear shape representation on the basis of harmonic embedding has been studied in Duci et al. (2003). (Chen et al., 2002) proposed to impose shape information on the zero crossing (rather than on the level set function). Rousson et al. proposed variational integrations of the shape prior (Rousson and Paragios, 2002; Rousson et al., 2004) based on the assumption of a Gaussian distribution. The use of nonparametric density estimation to model larger classes of level set based shape distributions was developed in Cremers et al. (2006) and Rousson and Cremers (2005). This approach allows to model distributions of shape which are not Gaussian – such as the various views of a 3D object (Cremers et al., 2002) or the silhouettes of a walking person (Cremers et al., 2006). Moreover, in the limit of large sample size, the nonparametric estimator constrains the distribution to the vicinity of the training shapes, such that the distribution favors shapes which are signed distance functions. A method to simultaneously impose shape information about several objects into level set based segmentation and to induce a recognition-driven segmentation through the competition of shape priors was developed in Cremers et al. (2006). Dynamical statistical shape priors for implicit shape representations were proposed in Cremers (2006). The latter approach takes into account that in




Figure 16. Sample segmentations using statistical shape priors. From left to right, the shape priors are a single static shape prior (Rousson and Paragios, 2002), uniformly distributed in the PCA subspace (Rousson, 2004), automatically selected from multiple shape instances (Cremers et al., 2006) and dynamical (Cremers, 2006).

the context of image sequence segmentation, the probability of a contour will depend on which contours have been observed in previous frames. The respective shape models capture the temporal correlations among silhouettes which characterize many deforming shapes.

In Fig. 16, we show a selection of segmentations obtained with some of the above methods. For further details we refer the reader to the respective publications.

9. Conclusion

We presented a survey of the class of region-based level set segmentation methods and detailed how they can be derived from a common statistical framework. The common goal of these approaches is to identify boundaries such that the color, texture, dynamic texture or motion in each of the separated regions is optimally approximated by simple statistical models.

Given a set of features or measurements $f(x)$ at each image location, minimization of the respective cost functionals leads to an estimation of a boundary C and a set of parameter vectors $\{\theta_i\}$ associated with each of the separated regions. Depending on the chosen segmentation criterion, the features f may be the pixel colors, the local structure tensors or the spatio-temporal intensity gradients, while the parameter vectors $\{\theta_i\}$ model distributions of intensity, color, texture or motion. The model parameters $\{\theta_i\}$ can be either simple aggregates of the features (as in the cases of color, texture or dynamic texture presented here) or derived quantities—as in the case of motion which is computed from the aggregated space-time gradients. The boundary $C \subset \Omega$ is implemented as the zero-crossing of an embedding function $\phi: \Omega \rightarrow \mathbb{R}$. Energy minimization leads to a gradient descent evolution of the embedding function interlaced with an update of the parameter vectors $\{\theta_i\}$ modeling the statistical distributions in the separated regions.

In numerous experimental results, we demonstrate that this class of level set methods allows to partition

images into domains of coherent color, texture, dynamic texture or motion. In particular, we show that—in contrast to the traditional edge-based segmentation schemes, these region-based approaches are quite robust to noise and to varying initialization, making them well-suited for local optimization methods such as the level set method. We ended by reviewing some recent advances regarding the introduction of statistical shape knowledge into level set based segmentation schemes.

Acknowledgments

We thank Christoph Schnörr for helpful comments on the manuscript. We also thank Zhizhou Wang for fruitful discussions, and Thomas Brox, Rodrigo de Luís García, Christophe Lenglet, Gianfranco Doretto, Paolo Favaro, Stefano Soatto and Anthony Yezzi for providing image examples of their methods. This work was partly supported by the German Research Foundation (DFG), grant #CR-250/1-1.

Notes

1. From a survey of a number of related publications and from our personal experience, it appears that the rigidity term is not particularly important, such that one commonly sets $\nu_2 = 0$.
2. It should be pointed out that based on various heuristics, one can successfully incorporate regriding mechanisms and topological changes into explicit representations—(cf. McInerney Terzopoulos, 1995; Lachaud and Montanvert, 1999; Delingette and Montagnat, 2000; Cremers et al., 2002).
3. Most meaningful contour evolutions do not contain any tangential component as the latter does not affect the contour, but only the parameterization.
4. In the following, I can refer to a single image or to an entire image sequence.
5. In Section 7, we will consider a generalization in which the underlying random processes are assumed to be space-varying. The distributions p_i in (11) then contain an explicit space dependency $p_i(f(x), x)$ which allows to model spatially varying statistical distributions of features.
6. For a recent study of various noise models on level set segmentation (see Martin et al., 2004).
7. This result was obtained in Lenglet et al. (2004) with data provided by J. F. Mangin and J. B. Poline.

8. Substituting in (28) A with TAT^{-1} , C with CT^{-1} , Q with TQT^{-1} , and choosing the initial condition $Tr(0)$, where $T \in GL(n)$ is any invertible $n \times n$ matrix generates the same output covariance sequence.
9. http://i21www.ira.uka.de/image_sequences/
10. In this paper, we are only concerned with 2D motion models. Such 2D motion models allow in particular to separate the different depth layers of a static scene filmed by a moving camera (cf. Cremers and Soatto, 2003). In terms of 3D motion, such a static scene instead corresponds to a *single* motion model.
11. Since we are only concerned with two consecutive frames from a sequence, we will drop the time coordinate in the notation of the velocity field.
12. To allow for variation of the global illumination, one can alternatively assume constancy of higher-order derivatives (cf. Brox et al., 2004).
13. For an alternative derivation of a similar likelihood we refer to Cremers and Yuille (2003).

References

- Amiaz, T. and Kiryati, N. 2005. Dense discontinuous optical flow via contour-based segmentation. In *Proc. IEEE Int. Conf. Image Processing*, Lausanne, vol 4, pp. 1264–1267.
- Aubert, G., Barlaud M., Faugeras, O., and Jehan-Besson, S. 2003. Image segmentation using active contours: Calculus of variations or shape gradients? *SIAM Journal of Applied Mathematics*, 63(6):2128–2154.
- Besag, J. 1986. On the statistical analysis of dirty pictures. *J. Roy. Statist. Soc., Ser. B.*, 48(3):259–302.
- Bigün, J. and Granlund G. 1987. Optimal orientation detection of linear symmetry. In *Proceedings of the 1st International Conference on Computer Vision*, London, England, IEEE Computer Society Press, pp. 433–438.
- Bigün, J., Granlund, G.H., and Wiklund, J. 1991. Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 13(8):775–790.
- Blake, A. and Zisserman, A. 1987. *Visual Reconstruction*. MIT Press.
- Boykov, Y., Veksler O., and Zabih, R. 2001. Fast approximate energy minimization via graph cuts. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 23(11):1222–1239.
- Brox, T., Bruhn, A., Papenberg, N., and Weickert, J. 2004. High accuracy optical flow estimation based on a theory for warping. In T. Pajdla and V. Hlavac, (Eds.), *European Conf. on Computer Vision*, Vol. 3024 of LNCS, Prague Springer, pp. 25–36.
- Brox, T., Bruhn, A., and Weickert, J. 2006. Variational motion segmentation with level sets. In A. Leonardis, H. Bischof and A. Pinz (Eds.), *European Conference on Computer Vision (ECCV)*, Graz, Austria, Springer, LNCS, Vol. 3951, pp. 471–483.
- Brox, T., Rousson, M., Deriche, R., and Weickert, J. 2003. Unsupervised segmentation incorporating colour, texture, and motion. In N. Petkov and M. A. Westenberg (Eds.), *Computer Analysis of Images and Patterns*, vol. 2756 of LNCS, Groningen, The Netherlands Springer, pp. 353–360.
- Brox, T. and Weickert, J. 2004. Level set based image segmentation with multiple regions. In H. Bülthoff, M. Giese, and B. Schölkopf (Eds.), *In Pattern Recognition, Springer LNCS 3175, C.-E. Rasmussen 26th DAGM*, Tübingen, Germany, pp. 415–423.
- Brox, T. and Weickert, J. 2004. A TV flow based local scale measure for texture discrimination. In T. Pajdla and V. Hlavac, (eds.), *European Conf. on Computer Vision*, Vol. 3022 of LNCS, Prague, Springer, pp. 578–590.
- Caselles, V., Catté, F., Coll, T., and Dibos, F. 1993. A geometric model for active contours in image processing. *Numer. Math.*, 66:1–31.
- Caselles, V., Kimmel, R., and Sapiro, G. 1995. Geodesic active contours. In *Proc. IEEE Intl. Conf. on Comp. Vis.*, Boston, USA, pp. 694–699.
- Chan, T.F. and Vese, L.A. 2001. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277.
- Charpiat, G., Faugeras, O., and Keriven, R. 2005. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Journal of Foundations of Computational Mathematics*, 5(1):1–58.
- Chen, Y., Tagare, H., Thiruvankadam, S., Huang, F., Wilson, D., Gopinath, K. S., Briggs, R.W., and Geiser, E. 2002. Using shape priors in geometric active contours in a variational framework. *Int. J. of Computer Vision*, 50(3):315–328.
- De Cock, K. and De Moor, B. 2000. Subspace angles between linear stochastic models. In *Int. Conf. on Decision and Control*, vol 2, pp. 1561–1566.
- Cohen, L.D. and Cohen, I. 1993. Finite-element methods for active contour models and balloons for 2-d and 3-d images. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 15(11):1131–1147
- Cremers, D. 2003. A variational framework for image segmentation combining motion estimation and shape regularization. In C. Dyer and P. Perona (Eds.), *IEEE Conf. on Comp. Vis. and Patt. Recog.*, vol 1, pp. 53–58.
- Cremers, D. 2006. Dynamical statistical shape priors for level set based tracking. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 28(8):1262–1273. .
- Cremers, D., Kohlberger, T., and Schnörr, C. 2002. Nonlinear shape statistics in Mumford–Shah based segmentation. In A. Heyden et al. (Eds.), *Europ. Conf. on Comp. Vis.*, vol. 2351 of LNCS, Copenhagen Springer, pp. 93–108.
- Cremers, D., Osher, S.J., and Soatto, S. 2006. Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int. J. of Computer Vision*, 69(3):335–351.
- Cremers, D. and Schnörr, C. 2002. Motion Competition: Variational integration of motion segmentation and shape regularization. In L. van Gool (Ed.), *Pattern Recognition*, vol. 2449 of LNCS, Zürich, Springer, pp. 472–480.
- Cremers, D. and Soatto, S. 2003. Variational space-time motion segmentation. In B. Triggs and A. Zisserman (Eds.), *IEEE Int. Conf. on Computer Vision*, Nice, vol. 2, pp. 886–892.
- Cremers, D., and Soatto, S. 2005. Motion Competition: A variational framework for piecewise parametric motion segmentation. *Int. J. of Computer Vision*, 62(3):249–265.
- Cremers, D., Sochen, N., and Schnörr, C. 2006. A multiphase dynamic labeling model for variational recognition-driven image segmentation. *Int. J. of Computer Vision*, 66(1):67–81.
- Cremers, D., Tischhäuser, F., Weickert, J., and Schnörr, C. 2002. Diffusion Snakes: Introducing statistical shape knowledge into the Mumford–Shah functional. *Int. J. of Computer Vision*, 50(3):295–313.
- Cremers, D. and Yuille, A.L. 2003. A generative model based approach to motion segmentation. In B. Michaelis and G. Krell (Eds.), *Pattern Recognition*, vol. 2781 of LNCS, Magdeburg, Springer, pp. 313–320.

- Cross, G. and Jain, A. 1983. Markov random field texture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5:25–39.
- de Luis García, R., Deriche, R., Rousson, M., and Alberola-López, C. 2005. Tensor processing for texture and colour segmentation. *Scandinavian Conference on Image Analysis 2005*: 1117–1127.
- Delingette, H., and Montagnat, J. 2000. New algorithms for controlling active contours shape and topology. In D. Vernon, (Ed.), *Proc. of the Europ. Conf. on Comp. Vis.*, vol. 1843, of *LNCS*, Springer, pp. 381–395.
- Dervieux, A. and Thomasset, F. 1979. A finite element method for the simulation of Raleigh-Taylor instability. *Springer Lect. Notes in Math.*, 771:145–158.
- Dervieux, A. and Thomasset, F. 1981. Multifluid incompressible flows by a finite element method. *Lecture Notes in Physics*, 11:158–163.
- Doretto, G., Chiuso, A., Wu, Y.N., and Soatto, S. 2003. Dynamic textures. *Int. Journal of Computer Vision*, 51(2):91–109.
- Doretto, G., Cremers, D., Favaro, P., and Soatto, S. 2003. Dynamic texture segmentation. In B. Triggs and A. Zisserman (Eds.), *IEEE Int. Conf. on Computer Vision*, Nice, vol 2, pp. 1236–1242.
- Duci, A., Yezzi, A., Mitter, S., and Soatto, S. 2003. Shape representation via harmonic embedding. In *ICCV*, pp. 656–662.
- Förstner, M.A., and Gülch, E. 1987. A fast operator for detection and precise location of distinct points, corners and centers of circular features. In *Proceedings of the Intercommission Workshop of the International Society for Photogrammetry and Remote Sensing*, Interlaken, Switzerland.
- Geman, S. and Geman, D. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 6(6):721–741.
- Grady, L. 2006. Random walks for image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, to appear.
- Granlund, G.H. and Knutsson, H. 1995. *Signal Processing for Computer Vision*, Kluwer Academic Publishers.
- Grenander, U., Chow, Y., and Keenan, D.M. 1991. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer, New York.
- Hassner, M. and Sklansky, J. 1980. The use of Markov random fields as models of texture. *Computer Graphics and Image Processing*, 12:357–370.
- Heiler, M. and Schnoerr, C. 2003. Natural image statistics for natural image segmentation. In *IEEE Int. Conf. on Computer Vision*, pp. 1259–1266.
- Herbulot, A., Jehan-Besson, S., Barlaud, M., and Aubert, G. 2004. Shape gradient for multi-modal image segmentation using mutual information. In *Int. Conf. on Image Processing*.
- Ising, E. 1925. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 23:253–258.
- Jehan-Besson, S., Barlaud, M., and Aubert, G. 2003. DREAM2S: Deformable regions driven by an eulerian accurate minimization method for image and video segmentation. *Int. J. of Computer Vision*, 53(1):45–70.
- Jepson, A., and Black, M.J. 1993. Mixture models for optic flow computation. In *Proc. IEEE Conf. on Comp. Vision Patt. Recog.*, pp. 760–761.
- Kass, M., Witkin, A., and Terzopoulos, D. 1988. Snakes: Active contour models. *Int. J. of Computer Vision*, 1(4):321–331.
- Keuchel, J., Schnörr, C., Schellewald, C., and Cremers, D. 2003. Binary partitioning, perceptual grouping, and restoration with semidefinite programming. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 25(11):1364–1379.
- Kichenassamy, S., Kumar, A., Olver, P.J., Tannenbaum, A., and Yezzi, A. J. 1995. Gradient flows and geometric active contour models. In *IEEE Int. Conf. on Computer Vision*, pp. 810–815.
- Kim, J., Fisher, J.W., Yezzi, A., Cetin, M. and Willsky, A. 2002. Nonparametric methods for image segmentation using information theory and curve evolution. In *Int. Conf. on Image Processing*, vol. 3, pp. 797–800.
- Lachaud, J.-O. and Montanvert, A. 1999. Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction. *Medical Image Analysis*, 3(2):187–207.
- Leclerc, Y.G. 1989. Constructing simple stable description for image partitioning. *The International Journal of Computer Vision*, 3(1):73–102.
- Leitner, F. and Cinquin, P. 1991. Complex topology 3d objects segmentation. In *SPIE Conf. on Advances in Intelligent Robotics Systems*, Boston, vol. 1609.
- Lenglet, C., Rousson, M., and Deriche, R. 2004. Toward segmentation of 3D probability density fields by surface evolution: Application to diffusion MRI. INRIA Research Report.
- Lenglet, C., Rousson, M., Deriche, R., and Faugeras, O. 2004. Statistics on multivariate normal distributions: A geometric approach and its application to diffusion tensor MRI. INRIA Research Report.
- Leung, T. and Malik, J. 2001. Representing and recognizing the visual appearance of materials using three-dimensional textures. *Int. J. of Computer Vision*, 43(1):29–44.
- Leventon, M., Grimson, W., and Faugeras, O. 2000. Statistical shape influence in geodesic active contours. In *Int. Conf. on Computer Vision and Pattern Recognition*, Hilton Head Island, SC, vol. 1, pp. 316–323.
- Lindeberg, T. 1994. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers.
- Malik, J., Belongie, S., Leung, T., and Shi, J. 2001. Contour and texture analysis for image segmentation. *Int. J. of Computer Vision*, 43(1):7–27.
- Malladi, R., Sethian, J.A., and Vemuri, B.C. 1994a. Evolutionary fronts for topology-independent shape modeling and recovery. In *Europ. Conf. on Computer Vision*, vol. 1, pp. 3–13.
- Malladi, R., Sethian, J.A., and Vemuri, B.C. 1994b. A topology independent shape modeling scheme. In *SPIE Conf. on Geometric Methods in Comp. Vision II*, vol. 2031, pp. 246–258.
- Mallat, S. 1989. Multiresolution approximations and wavelet orthonormal bases of $L^2(R)$. *Trans. Amer. Math. Soc.*, 315:69–87.
- Mansouri, A., Mitiche, A., and Feghali, R. 2002. Spatio-temporal motion segmentation via level set partial differential equations. In *Proc. of the 5th IEEE Southwst Symposium on Image Analysis and Interpretation (SSIAI)*, Santa Fe.
- Martin, P., Refregier, P., Goudail, F., and Guerault, F. 2004. Influence of the noise model on level set active contour segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 26(6):799–803.
- McInerney, T. and Terzopoulos, D. 1995. Topologically adaptable snakes. In *Proc. 5th Int. Conf. on Computer Vision*, Los Alamitos, California, IEEE Comp. Soc. Press, pp. 840–845.
- Mumford, D., and Shah, J. 1989. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685.
- Nain, D., Yezzi, A., and Turk, G. 2003. Vessel segmentation using a shape driven flow. In *Intl. Conf. on Medical Image Computing and Comp. Ass. Intervention (MICCAI)*, pp. 51–59.

- Osher, S.J., and Sethian, J.A. 1988. Fronts propagation with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *J. of Comp. Phys.*, 79:12–49.
- Paragios, N. and Deriche, R. 2002. Geodesic active regions: a new paradigm to deal with frame partition problems in computer vision. *Journal of Visual Communication and Image Representation*, 13(1/2):249–268.
- Paragios, N. and Deriche, R. 2005. Geodesic active regions and level set methods for motion estimation and tracking. *Computer Vision and Image Understanding*, 97(3):259–282.
- Pennec, X., Fillard, P., and Ayache, N. 2005. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 65(1).
- Perkins, W.A. 1980. Area segmentation of images using edge points. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 2(1):8–15.
- Rao, A.R., and Schunck, B.G. 1991. Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing*, 53:157–185.
- Rochery, M., Jermyn, I., and Zerubia, J. 2006. Higher order active contours. *Int. J. of Computer Vision*, 69(1): 1573–1405.
- Rousson, M. 2004. *Cues Integrations and Front Evolutions in Image Segmentation*. PhD thesis, Université de Nice-Sophia Antipolis.
- Rousson, M., Brox, T., and Deriche, R. 2003. Active unsupervised texture segmentation on a diffusion based feature space. In *Proc. IEEE Conf. on Comp. Vision Patt. Recog.*, Madison, WI, pp. 699–704.
- Rousson, M. and Cremers, D. 2005. Efficient kernel density estimation of shape and intensity priors for level set segmentation. In *Intl. Conf. on Medical Image Computing and Comp. Ass. Intervention (MICCAI)*, vol. 1, pp. 757–764.
- Rousson, M. and Deriche, R. 2002. A variational framework for active and adaptative segmentation of vector valued images. RR 4515, INRIA.
- Rousson, M. and Deriche, R. 2002. A variational framework for active and adaptive segmentation of vector valued images. In *Proc. IEEE Workshop on Motion and Video Computing*, Orlando, Florida, pp. 56–62.
- Rousson, M., Lenglet, C., and Deriche, R. 2004. Level set and region based surface propagation for diffusion tensor MRI segmentation. In *Computer Vision Approaches to Medical Image Analysis (CVAMIA) and Mathematical Methods in Biomedical Image Analysis (MMBIA) Workshop*, Prague.
- Rousson, M. and Paragios, N. 2002. Shape priors for level set representations. In A. Heyden et al. (Eds.), *Europ. Conf. on Comp. Vis.*, Springer, vol. 2351 of LNCS, pp. 78–92.
- Rousson, M., Paragios, N., and Deriche, R. 2004. Implicit active shape models for 3d segmentation in MRI imaging. In *Intl. Conf. on Medical Image Computing and Comp. Ass. Intervention (MICCAI)*, Springer, vol. 2217 of LNCS, pp. 209–216.
- Samson, C., Blanc-Feraud, L., Aubert, G., and Zerubia, J. 2000. A variational model for image classification and restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(5):460–472.
- Schoenerr, C. 1992. Computation of discontinuous optical flow by domain decomposition and shape optimization. *Int. J. of Computer Vision*, 8(2):153–165.
- Shi, J. and Malik, J. 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905.
- Simioncelli, P., Freeman, W., Adelson, H., and Heeger, J. 1992. Shifttable multiscale transforms. *IEEE trans. on Information Theory*, 38:587–607.
- Skovgaard, L.T. 1984. A Riemannian geometry of the multivariate normal model. *Scand. J. Statistics*, 11:211–223.
- Sokolowski, J. and Zolesio, J.P. 1991. Introduction to shape optimization. *Computational Mathematics*. Springer Verlag.
- Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, E., and Willsky, A. 2001. Model-based curve evolution technique for image segmentation. In *Comp. Vision Patt. Recog.*, Kauai, Hawaii. pp. 463–468.
- Tsai, A., Yezzi, A.J., and Willsky, A.S. 2001. Curve evolution implementation of the Mumford-Shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. on Image Processing*, 10(8):1169–1186.
- Tsai, A., Yezzi, A. J., and Willsky, A. S. 2003. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Trans. on Medical Imaging*, 22(2):137–154.
- Tschumperlé, D. and Deriche, R. 2001a. Diffusion tensor regularization with constraints preservation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Kauai Marriott, Hawaii.
- Tschumperlé, D. and Deriche, R. 2001b. Regularization of orthonormal vector sets using coupled PDE's. In *IEEE Workshop on Variational and Level Set Methods*, Vancouver, Canada, pp. 3–10.
- Unal, G., Krim, H., and Yezzi, A.Y. 2005. Information-theoretic active polygons for unsupervised texture segmentation. *Int. J. of Computer Vision*.
- Vese, L.A., 2003. Multiphase object detection and image segmentation. In S. J. Osher and N. Paragios (Eds.), *Geometric Level Set Methods in Imaging, Vision and Graphics*, New York, Springer, pp. 175–194.
- Vese, L.A., and Chan, T.F. 2002. A multiphase level set framework for image segmentation using the Mumford and Shah model. *The International Journal of Computer Vision*, 50(3):271–293.
- Wang, J.Y.A., and Adelson, E.H. 1994. Representing moving images with layers. *IEEE Trans. on Image Processing*, 3(5):625–638.
- Wang, Z. and Vemuri, B.C. 2004. An affine invariant tensor dissimilarity measure and its application to tensor-valued image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, Washington, DC.
- Weickert, J. and Brox, T. 2002. Diffusion and regularization of vector and matrix-valued images. In *Contemporary Mathematics*, vol. 313, pp. 251–268.
- Yezzi, A., Tsai, A., and Willsky, A. 1999. A statistical approach to snakes for bimodal and trimodal imagery. In *Proceedings of the 7th International Conference on Computer Vision*, Kerkyra, Greece, vol. II, pp. 898–903.
- Di Zenzo, S. 1986. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing*, 33:116–125.
- Zhao, H.-K., Chan, T., Merriman, B., and Osher, S. 1996. A variational level set approach to multiphase motion. *J. of Comp. Phys.*, 127:179–195.
- Zhu, S.C., Wu, Y., and Mumford, D. 1998. Filters, random fields and maximum entropy (frame). *The International Journal of Computer Vision*, 27(2):1–20.
- Zhu, S.C. and Yuille, A. 1996. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 18(9):884–900.