



Texture Segmentation via Non-local Non-parametric Active Contours

Jung Miyouun, Gabriel Peyré, Laurent D. Cohen

► To cite this version:

Jung Miyouun, Gabriel Peyré, Laurent D. Cohen. Texture Segmentation via Non-local Non-parametric Active Contours. EMMCVPR, Jul 2011, Saint Petersburg, Russia. pp.74-88. hal-00598164

HAL Id: hal-00598164

<https://hal.science/hal-00598164>

Submitted on 4 Jun 2011

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.

Texture Segmentation via Non-local Non-parametric Active Contours

Miyoun Jung, Gabriel Peyré and Laurent D. Cohen

Ceremade, Université Paris-Dauphine, 75775 Paris Cedex 16, France
{jung,peyre,cohen}@ceremade.dauphine.fr

Abstract. This article introduces a novel active contour model that makes use of non-parametric estimators over patches for the segmentation of textured images. It is based on an energy that enforces the homogeneity of these statistics. This smoothness is measured using Wasserstein distances among discretized probability distributions that can handle features in arbitrary dimension. It is thus usable for the segmentation of color images or other high dimensional features. The Wasserstein distance is more robust than traditional pointwise statistical metrics (such as the Kullback-Leibler divergence) because it takes into account the relative distances between modes in the distributions. This makes the corresponding energy robust and does not require any smoothing of the statistical estimators. To speed-up the computational time, we propose an alternative metric that retains the main qualities of the Wasserstein distance, while being faster to compute. It aggregates 1-D Wasserstein distances over a set of directions, and thus benefits from the simplicity of 1-D statistical metrics while being able to discriminate high dimensional features. We show numerical results that instantiate this novel framework using grayscale and color values distributions. This allows us to segment regions with smoothly varying intensities or colors as well as complicated textures.

1 Introduction

This article considers a variational minimization problem aiming at detecting objects in textured images. It makes use of a comparison principle between pairs of neighboring patches. We thus refer to it as a “non-local” approach, following the terminology initiated in [1]. The resulting method is a general framework to implement a piecewise smooth segmentation model. To illustrate the generality of this concept, we make use of statistical distances between distributions. The resulting method considers that the objects to be segmented are sampled from a non-stationary distribution with parameters that vary smoothly from pixel to pixel. Our approach is general when it comes to the use of distributions, in the sense that it can handle the distributions of general features such as pixel values, small chunks of pixel values to capture the joint dependancies between neighboring pixels, multiscale coefficients (e.g. wavelets, Gabor filter banks, etc) or group of these. We restrict in this article our attention to the distributions of pixel values (intensities or color values).

1.1 Previous Works

Region-based active contours. This paper is focused on the class of region-based approaches to image segmentation. These methods make use of information extracted from inside and outside the region to be segmented. Following the seminal work of Mumford and Shah on optimal partitions [2], several influential works have proposed to perform image segmentation using a variational minimization with a region-based energy that seeks for constant features in each region to be segmented, see for instance [3, 4].

More recently, several models have been introduced to handle a local homogeneity of the features. They require the estimation of a piecewise smooth parameter field, see for instance [5–9]. The non-local active contour method, recently introduced in [10, 11], also makes use of such a local homogeneity principle, but implements the variational minimization using only a pairwise comparison of features, which might for instance depend on patches extracted around each pixel. A chief advantage of this approach is that it does not require the estimation of a piecewise smooth set of parameters, and only requires the design of a metric to compare patches. In this paper, we extend this non-local active contour framework in several directions, to capture locally homogenous statistical distributions.

Non-local segmentation. Non-local methods have been applied in many image processing problems such as denoising [1], general inverse problems [12] and segmentation [13–15], by regularizing the image using comparison of patches in the image. In this article we define an attraction term pulling the contour towards the object boundaries, which is contrast to the existing non-local based segmentation methods that use non-local energy terms only as regularization terms.

Statistical segmentation. Instead of considering a small set of parameters (such as the mean value of the features), more general models make use of statistical distributions to drive the segmentation. These approaches are quite effective for many natural images or textures that contain complicated random fluctuations. The resulting statistical region-based active contours make use of pointwise similarity measures among distributions (such as the Kullback-Leibler divergence) to compare the distributions, in a parametric or non-parametric (using Parzen windows) fashion, see for instance [16–19]. In this paper, we also consider the setting of statistical segmentation, and we propose to use a fully non-parametric estimator that does not require to compute histograms.

Wasserstein metric. Traditional pointwise statistical distances are simple to compute, but suffer from several drawbacks, in particular the difficulty to handle localized distributions. This is because these statistical distances do not take into account the relative positions between the modes of the distributions. In practice, a smoothing of the histogram using Parzen windows is required to make these methods usable on localized distributions. To address these issues, Ni et al. [20] propose to make use of the L^1 Wasserstein distance in order to extend the

segmentation model of Chan and Vese [3]. This work is extended to color image segmentation in [21] using a Wasserstein metric only on the brightness channel, thus resulting in a 1-D optimal transport metric.

The Wasserstein distance is now routinely used in computer vision, for instance as a metric for retrieval, where it improves over more classical pointwise distances between histograms [22]. It is also used for other image processing applications such as warping [23] and texture synthesis [24]. This metric is related to the assignment problem [25]. We use this connection in our work to avoid computing histograms.

1.2 Comparison with Previous Works

Our work draws connexions between recent works in statistical image segmentation, and extends these works in several directions.

We make use of the non-local active contour variational minimization problem proposed recently in [10, 11]. We extend it by introducing a pixel-by-pixel normalization of the energy. This is crucial to obtain a un-biased segmentation result in the case where the size of patches is large. This is typically required for statistical segmentation when the distributions have a large variance (such as for noisy images or for complicated textures). Furthermore, we instantiate this energy using a Wasserstein metric. This defines a notion of locally homogeneous statistical distributions, which might be of independent interest.

We follow the work of [20] which is the first to clearly acknowledge the importance of optimal transport for image segmentation. We develop a different variational model, that does not require a global consistency of the local histograms. This is important to segment images with a slowly varying background, or complicated non-stationary textures. Furthermore, we extend this method to distributions in arbitrary dimension, which leads for instance to a 3-D distance for color distributions. We also make use of the L^p Wasserstein distance using an assignment formulation of the metric, while previous works restrict their attention to $p = 1$. To speed up the evaluation of the Wasserstein distance, we approximate it using a series of 1-D projections, which is a method introduced in [24] for histogram equalization.

1.3 Contributions

Our first contribution is a novel region-based segmentation method that extends the non-local active contour method [10, 11] using a pixel-by-pixel normalization. Our second contribution is the definition of a novel segmentation criteria that requires only a local homogeneity of the statistical distribution of features. This criteria is handled using Wasserstein metrics, which extends the global homogeneity criteria introduced in [20, 21].

2 Non-local Active Contours

Section 2 recalls the non-local active contours energy introduced originally in [10, 11]. We then describe our contribution, which is a normalization of this model that reduces significantly the bias introduced by the use of patches.

2.1 Un-normalized Non-local Active Contours [10, 11]

Our goal is to segment an image $f : [0, 1]^2 \rightarrow \mathbb{R}^d$, where d is dimensionality of the feature space. We make use of a local patch extraction process to design variational energies.

Pairwise Patch Interaction. A patch in some image f around a pixel $x \in [0, 1]^2$ is defined as $p_x(t) = f(x + t)$, $\forall t \in [-\tau/2, \tau/2]^2$.

The non-local interaction between two patches is measured using a metric $d(\cdot, \cdot) \geq 0$ that accounts for the similarity between patches.

The simplest choice, considered in [10], is a weighted L^2 distance

$$d(p_x, p_y) = \int_t G_a(t) \|p_x(t) - p_y(t)\|^2 dt, \quad (1)$$

where a Gaussian weight $G_a(t) = e^{-\frac{\|t\|^2}{2a^2}}$ can be used to give more influence to the central pixel.

More complicated similarity measures can be used, and Section 3 explains how to use statistical distances.

Level set formulation. The segmentation problem corresponds to the computation of some region $\Omega \subset [0, 1]^2$ that should capture the objects of interest.

We represent the segmented region Ω using a level set function $\varphi : [0, 1]^2 \rightarrow \mathbb{R}$ so that $\Omega = \{x \mid \varphi(x) > 0\}$. To simplify the exposition, we make use of a smoothed Heaviside function $H(\varphi) = \frac{1}{2} + \frac{1}{\pi} \text{atan}(\varphi/\varepsilon)$ to introduce variational energies and compute their derivatives. The parameter ε should be chosen small enough to obtain a sharp region boundary, but not too small to avoid numerical instabilities. In the numerical examples, we use $\varepsilon = 1/n$ for a discretized image of $n \times n$ pixels.

A mathematically more rigorous way to derive the corresponding PDE is to make use of the shape derivative machinery, which is formally equivalent to letting ε tend to 0, see for instance [18, 19]. Using such a shape gradient would make the evolution PDE well defined only on the boundary of Ω , and this evolution is then extended to the whole domain by preserving some distance function property on φ .

Non-local segmentation energy. Following [10, 11], we introduce an energy functional $E_0(\varphi)$ enforcing the similarity of features located either inside or outside Ω ,

$$E_0(\varphi) = \iint \rho(H(\varphi(x)), H(\varphi(y))) G_\sigma(x - y) d(p_x, p_y) dx dy. \quad (2)$$

The function ρ restricts the comparison to pairs of patches that are in the same region (inside or outside). Since $H(\varphi(x))$ is close to being a binary function, we use $\rho(u, v) = 1 - |u - v|$ for the numerical experiments (but other similar functionals could be used as well). Note that the parameter $\sigma > 0$ is important since it controls the scale of the local homogeneity one requires for the segmented object.

To enforce the regularity of the boundary of the extracted region, following previous works in active contours, we penalize its length, which is computed as

$$L(\varphi) = \int \|\nabla H(\varphi)(x)\| dx = \int H'(\varphi(x)) \|\nabla \varphi(x)\| dx$$

where $\nabla H(\varphi)(x)$ is the gradient at point x of the function $H(\varphi)$.

The non-local active contour model computes the segmentation as a stationary point of the energy

$$\min_{\varphi} E_0(\varphi) + \gamma L(\varphi) \quad (3)$$

where $\gamma > 0$ is a parameter that should be adapted to the expected regularity of the boundary of the region.

Limitation and Motivation The non-local active contours model works well when the size of patches is small. Figure 1 shows examples of segmentation of piecewise smooth images using the L^2 patch distance (1) with patches of width $\tau = 3/n$ for an image of size $n \times n$ and $a = 0.5/n$. The local homogeneity property of the energy (2) enables the model (3) to correctly detect objects which are only locally homogeneous, and can deal with separated objects with different intensities.

This model however suffers from a segmentation bias. The segmented region is shifted away from the object boundary with an amount proportional to the patch width τ . This becomes problematic when used with large patches, because of the lack of precision of the resulting segmentation. Large patches (and large values of a) are however desirable as the noise level increases, since robustness requires more pixels to evaluate the local homogeneity. Figures 2 and 3 show that increasing the size of patches (the value of a) leads to segmented regions with a smoother boundary, but also reveal that the resulting curve is not located on the exact boundary of the object. Figure 3, first column, shows how the amount of bias increases as the size of patch (or the value a) increases.

2.2 Normalized Non-local Active Contour model

Normalized energy. To reduce the segmentation bias introduced by the non-local active contour energy (2), we define a novel normalized non-local energy

$$E(\varphi) = \int \frac{1}{C(\varphi, x)} \int \rho(H(\varphi(x)), H(\varphi(y))) G_{\sigma}(x - y) d(p_x, p_y) dy dx, \quad (4)$$

where the local normalization factor is $C(\varphi, x) = \int \rho(H(\varphi(x)), H(\varphi(y))) G_{\sigma}(x - y) dy > 0$. Note that the un-normalized energy E_0 defined in (2) is recovered by setting $C(\varphi, x) = 1$.

In practice, the correction factor $1/C(\varphi, x)$ is far from being constant, in particular when the size of patches is large. This normalization is thus crucial to reduce the disparities that increase as a pixel approaches the boundary of the segmented region.

Gradient flow. Our normalized non-local active contour model computes the segmentation as a stationary point of the energy

$$\min_{\varphi} E(\varphi) + \gamma L(\varphi) \quad (5)$$

where L is defined in (2) and $\gamma > 0$ is a regularization parameter.

Introducing an artificial time $t \geq 0$, the gradient flow of (5) reads

$$\frac{\partial \varphi}{\partial t} = -(\nabla E(\varphi) + \gamma \nabla L(\varphi)), \quad (6)$$

for $\varphi(x, t)$ parameterized by space and time.

The gradients are computed as

$$\nabla E(\varphi)(x) = \frac{\partial f(\varphi(x))g(\varphi(x)) - f(\varphi(x))\partial g(\varphi(x))}{(g(\varphi(x)))^2}, \quad (7)$$

$$\nabla L(\varphi)(x) = -\operatorname{div} \left(\frac{\nabla \varphi(x)}{\|\nabla \varphi(x)\|} \right) H'(\varphi(x)), \quad (8)$$

with the notations

$$\begin{aligned} f(u) &:= \int \rho(H(u), H(\varphi(y))) G_{\sigma}(x - y) d(p_x, p_y) dy, \\ g(u) &:= \int \rho(H(u), H(\varphi(y))) G_{\sigma}(x - y) dy, \\ \partial f(u) &:= \int (\partial_1 \rho)(H(u), H(\varphi(y))) G_{\sigma}(x - y) d(p_x, p_y) dy H'(u), \\ \partial g(u) &:= \int (\partial_1 \rho)(H(u), H(\varphi(y))) G_{\sigma}(x - y) dy H'(u). \end{aligned}$$

where ∂_1 is the gradient with respect to the first variable.

Numerical implementation. The segmentation is applied to a discretized image f of $n \times n$ pixels. The length energy (2) is computed using a finite difference approximation of the gradient.

In a preprocessing step, the distance between neighboring patches $d(p_x, p_y)$ is computed. Depending on the numerical application, one might want to use either the weighted L^2 norm (1) or the sliced Wasserstein distance defined in (11).

The gradient flow (6) is then discretized using a gradient descent

$$\varphi^{(\ell+1)} = \varphi^{(\ell)} - \mu_{\ell} (\nabla E(\varphi) + \gamma \nabla L(\varphi)),$$

where $\mu_{\ell} > 0$ is a suitable time step size.

To make all the level sets evolve simultaneously, $H'(\varphi(x))$ appearing in (7) and (8) is replaced by $\|\nabla \varphi(x)\|$. To ensure the stability of the level set evolution (6), one needs to re-initialize it from time to time. This corresponds to replacing φ by the signed distance function to the level set $\{x \mid \varphi(x) = 0\}$.

The width σ of the windowing function $G_{\sigma}(x - y)$ typically depends on the initial curve at time $t = 0$. If the initial curve is far away from the object boundary, a large windowing function might be required.

3 Wasserstein Local Homogeneity

Classical deterministic similarity measures such as the L^2 norm (1) are suitable to segment piecewise smooth images. They can also be used over a transformed domain (such the output of a Gabor filter bank [10]) to characterize some simple geometric textures. To handle complicated images containing textural contents, it is useful to consider similarity measures between statistical distributions.

3.1 Wasserstein Distance

In this article we consider a Lagrangian discretization of distributions, which corresponds to treating a distribution X as a points cloud $X = \{X_i\}_{i=0}^{N-1} \subset \mathbb{R}^d$. In our numerical applications, $d = 1$ for grayscale images and $d = 3$ for color images. This is different from the more traditional Eulerian discretization, that makes use of a fixed set of points (usually a rectangular grid) but where the points are equipped with a weight to reflect the local density of the distribution. These histogram-based (Eulerian) discretizations are at the heart of previous statistical region-based active contours such as [20, 18, 19].

In this Lagrangian setting, for $p \geq 1$, the L^p Wasserstein distance between two distributions $X, Y \subset \mathbb{R}^d$ is defined as

$$W^p(X, Y) = \min_{\sigma \in \Sigma_N} \sum_{i=0}^{N-1} \|X_i - Y_{\sigma(i)}\|^p \quad (9)$$

where Σ_N is the set of all the permutations of N elements. For simplicity we have restricted our attention to distributions having the same number of points, which is the case for our application to segmentation. Note that W should really be understood as being a distance between points clouds, since it is invariant under a permutation of the indexes of the distributions.

The permutation σ minimizing (9) is the optimal assignment between the two points clouds. This optimal assignment problem can be solved using combinatorial optimization schemes in $O(N^{5/2} \log(N))$ operations, see [25].

3.2 Sliced Wasserstein Distance

1-D Wasserstein Distance. In the 1-D case, the optimal assignment σ that solves (9) can be computed in $O(N \log(N))$ operations by ordering the points clouds X and Y

$$X_{\sigma_X(i)} \leq X_{\sigma_X(i+1)} \quad \text{and} \quad Y_{\sigma_Y(i)} \leq Y_{\sigma_Y(i+1)}$$

with two permutations $\sigma_X, \sigma_Y \in \Sigma_N$. The optimal permutation is then $\sigma = \sigma_Y \circ \sigma_X^{-1}$. Equivalently, the Wasserstein distance is the L^p norm of the sorted vectors

$$W^p(X, Y) = \sum_{i=0}^{N-1} |X_{\sigma_X(i)} - Y_{\sigma_Y(i)}|^p. \quad (10)$$

Note the major computational difference between the assignment problem (9) in dimension $d = 1$ and in higher dimensions $d > 1$, where no $O(N \log(N))$ algorithm is available.

Sliced approximation. The numerical complexity of solving (9) in dimension $d > 1$ is prohibitive for imaging applications such as our segmentation problem. To obtain a fast numerical scheme, we follow the work of Rabin et al. [24] that introduces a sliced Wasserstein distance. It is defined as an aggregation of 1-D Wasserstein distances of projected distributions

$$SW^p(X, Y) = \sum_{\theta \in \Theta} W^p(X_\theta, Y_\theta) \quad \text{where} \quad X_\theta = \{\langle X_i, \theta \rangle\}_{i=1}^{N-1}. \quad (11)$$

Here $X_\theta, Y_\theta \subset \mathbb{R}$ are projected 1-D distributions and $\Theta \subset \mathbb{R}^d$ is a discrete set of directions, sampled on the unit sphere.

Evaluating this sliced distance (11) has a complexity of $O(|\Theta|N \log(N))$ operations which is advantageous over the original Wasserstein distance (9) if Θ is not too large. Although there is no mathematical proof of the quality of the approximation of W using SW , numerical observations suggest that SW is a good approximation to solve minimization problems involving the Wasserstein metric, see [24].

3.3 Wasserstein Non-local Active Contours

The sliced approximation (11) is used to measure the similarity between patches to perform statistical region-based segmentation. We thus propose to replace the L^2 norm in (1) by

$$d(p_x, p_y) = W^p([p_x], [p_y])$$

where the operator $[\cdot]$ maps a vector $v = (v_i)_{i=0}^{N-1} \in \mathbb{R}^N$ to a points cloud $[v] = \{v_i\}_{i=0}^{N-1} \subset \mathbb{R}^d$.

4 Experimental Results and Comparisons

This section presents experimental results with synthetic and real images.

Chan-Vese and LBF model incorporated with patches. For grayscale images, we compare our model (5) with extensions of the Chan-Vese (CV) [3] and locally binary fitting (LBF) models [6]. These extensions make use of patches to enable a fair comparison with our method.

They are obtained by replacing intensity features by patches in the original CV or LBF models

$$\min_{p^1, p^2, \varphi} E(p^1, p^2, \varphi) + \gamma L(\varphi) \quad (12)$$

where E is either E_{CV} or E_{LBF} defined as

$$E_{CV} = \lambda_1 \int d(p_x, p^1) H(\varphi(x)) dx + \lambda_2 \int d(p_x, p^2) (1 - H(\varphi(x))) dx,$$




Fig. 1. Results of the un-normalized model (3) with the weighted L^2 distance function (1), and comparison with Chan-Vese model [3]. Patches of width $\tau = 3/n$ (3×3 pixels) and $a = 0.5/n$ are used. 100×100 image and 31×31 windowing function are used.




Fig. 2. Comparisons of our new normalized model (5) with the un-normalized model (3) with the L^2 distance (1). 1st-3rd: the un-normalized model with patches of width $\tau = 1/n$ (1st), $\tau = 3/n$ and $a = \infty$ (2-3rd) with two different parameters $\gamma = 0.3$ (2nd) and 0.5 (3rd). 4th-5th: our new normalized model (5) with patches of width $\tau = 3/n$, $\tau = 5/n$ and $a = \infty$.

$$E_{LBF} = \lambda_1 \iint G_\sigma(x-y) d(p_y, p_x^1) H(\varphi(y)) dy dx \\ + \lambda_2 \iint G_\sigma(x-y) d(p_y, p_x^2) (1 - H(\varphi(y))) dy dx,$$

where $\lambda_1, \lambda_2 > 0$ are parameters, and p^1 and p^2 are updated by iteratively cycling through φ and then (p^1, p^2) . Note that the energy E_{CV} is introduced in [20] in the special case of the L^1 -Wasserstein distance. In the numerical examples, we let $\lambda_1 = \lambda_2 = 1$, and we tried to choose the best smoothness parameters γ for each model.

For color images, we compare our model (5) with the vector-valued Chan-Vese model [26], and the work [21] that uses a Wasserstein metric only on the brightness channel and the vector-valued Chan-Vese model for the chromaticity components.

L^2 Distance Between Patches. Figure 1 shows the results of the un-normalized model (3) tested on synthetic images with spatially varying background and/or




Fig. 3. Comparisons of models with the new distance function (10). 1st-2nd columns: the un-normalized model (3) with two different parameters $\gamma = 0.1, 0.5$ (1st row), $\gamma = 0.2, 0.3$ (2nd row). 3rd-5th columns: our new normalized model (5), CV and LBF models (12) with patches with width $\tau = 15/n$ (top row), $\tau = 11/n$ (bottom).

object, or with several separated objects with different intensities. Due to the local homogeneity property, the model correctly detects objects, while the two-phase Chan-Vese model, requiring a global homogeneity in each region, fails for the correct segmentation.

Figure 2 presents an example where the un-normalized model (3) does not provide satisfactory results with any kinds of patches: patches of width $\tau = 1/n$ (1 pixel), $\tau = 3/n$ (3×3 pixels) with $a = \infty$. By using patches of width $\tau = 1/n$, the un-normalized model (3) produces noisy final curves, and by using patches of width $\tau = 3/n$, it results in smoother final curves that are however not located on the object boundaries in spite of adjusting the smoothness parameter γ . On the other hand, our new normalized model (5) with patches of width $\tau = 5/n$ provides a smooth final curve, located exactly on the boundary. This example also shows the case when a large size of patch is required because our model with smaller size ($\tau = 3/n$) of patch also gives noisy final curves.

Wasserstein Distance Between Patches. Figure 3 presents the examples where the L^2 patch distance (1) cannot be applied because the black and white stripe pattern is a texture that is not homogenous in the pixel domain. Furthermore, these examples require a large patch size to capture the texture statistics. We present comparisons of the un-normalized model (3), our new normalized model (5), CV and LBF models (12), using the 1-D Wasserstein distance (10). The 1st-2nd columns present the results of the un-normalized model with different smoothness parameters γ with patches of $\tau = 15/n$ (1st row) and $11/n$ (2nd row). The un-normalized model produces biased final curves near the boundary, and large values of γ seem to reduce the bias to some extent. However, the bias cannot be reduced completely by the parameter γ , as well as that large values of γ result in too smoothed-out curves and slow convergence. However, our model locates the curves exactly on the boundary. Lastly, the 4th-5th columns show that CV and LBF models also do not locate the final curves on the exact boundary: the




Fig. 4. Results of our new normalized model (5) with the Wasserstein distance (10). 1st-5th columns: initial curve (1st), curve evolution (2nd-3rd), final curve (4th), plot of energy $E(\varphi^{(\ell)})$ vs iteration ℓ (5th). Patches of width $\tau = 3/n$ are used.

curves are located a few pixels far away from the boundaries. Although our model and CV/LBF models have similar behaviors on globally homogeneous textures, this example highlights the importance of our normalization.

Figure 4 shows that our new normalized model (5) with the Wasserstein distance (10) detects objects with smoothly varying distributions of intensities and separated multiple objects with different distributions of intensities. It also shows the curve evolution of our model starting from given initial curves, and displays the convergence of the energy $E(\varphi^{(\ell)})$ as a function of the iteration index ℓ .

Figure 5 presents texture segmentation results of our model (5) and comparison with CV model (12). Again due to the local homogeneity, our model discriminates different textures having different distributions of intensities, while CV fails for the correct discrimination.

In Figure 6, we present color texture segmentation results of our model using a sliced Wasserstein distance (11). We considered only $|\Theta| = 3$ projection directions, i.e. $\Theta = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$, which was enough to obtain satisfactory segmentation in all the given examples. We also compare our model with the vector-valued Chan-Vese model [26], and with the color extension [21] of the original method proposed in [20]. In all the examples, our model correctly detects the boundary of objects and segments separated multiple objects with different distributions of color values, in contrast to the other models that do not locate the curve on the exact boundaries, detect only part of objects or fail to detect objects. Note that the second example was degraded by the random-valued impulsive noise of density 0.3.

Conclusion

In this article, we have proposed a novel non-local energy for the segmentation of textured images, making use of non-parametric estimators over patches. The Wasserstein distance and its sliced approximation can be used as a similarity




Fig. 5. Texture segmentation using our new normalized model (5) with the Wasserstein distance (10), and comparison with CV model (12). 1st-3rd rows: initial curves (1st), results of our model (2nd) and CV model (3rd). 1st, 3rd, 4th columns: patch with width $\tau = 11/n$, 2nd column: patch with width $\tau = 7/n$. In our model, 31×31 windowing function is used.

measure which allows one to segment complicated textural features in arbitrary dimension. Due to the local homogeneity property of the energy, our active contour model is able to detect regions with smoothly spatially varying features. It can also segment several separated regions with different features. All these properties are significant extensions of existing region-based models crucial to solve difficult texture segmentation problems.

References

1. Buades, A., Coll, B., Morel, J.M.: A review of image denoising algorithms, with a new one. *SIAM Mul. Model. and Simul.* **4** (2005) 490–530
2. Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics* **XLII** (1989)
3. Chan, T., Vese, L.: Active contours without edges”. *IEEE Trans. Image Proc.* **10** (2001) 266–277
4. Paragios, N., Deriche, R.: Geodesic active regions: A new framework to deal with frame partition problems in computer vision. *Journal of Visual Communication and Image Representation* **13** (2002) 249–268

5. Tsai, A., Yezzi, A., Willsky, A.S.: Curve evolution implementation of the mumford-shah functional for image segmentation, denoising, interpolation, and magnification. *IEEE Trans. Image Proc.* **10** (2001) 1169–1186
6. Li, C., Kao, C., Gore, J., Ding, Z.: Implicit active contours driven by local binary fitting energy. *Proceedings of the CVPR'07* (2007) 1–7
7. Wang, L., Li, C., Suna, Q., Xia, D., Kao, C.Y.: Active contours driven by local and global intensity fitting energy with application to brain mr image segmentation. *Computerized Medical Imaging and Graphics* **33** (2009) 520–531
8. Wang, L., He, L., Mishra, A., Li, C.: Active contours driven by local gaussian distribution fitting energy. *Signal Processing* **89** (2009) 2435–2447
9. Wang, X., Huang, D., Xu, H.: An efficient local chan–vese model for image segmentation. *Pattern Recognition* **43** (2010) 603–618
10. Jung, M., Peyré, G., Cohen, L.D.: Nonlocal active contours. *Proc. SSVM'11* (2011)
11. Jung, M., Peyré, G., Cohen, L.D.: Nonlocal segmentation and inpainting. *Proc. ICIP'11* (2011)
12. Peyré, G., Bogleux, S., Cohen, L.: Non-local regularization of inverse problems. *ECCV, Part III, LNCS* **5304** (2008) 57–68
13. Gilboa, G., Osher, S.: Nonlocal linear image regularization and supervised segmentation. *SIAM Mul. Model. and Simul.* **6** (2007) 595–630
14. Elmoataz, A., Lezoray, O., Bogleux, S.: Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing. *IEEE Trans Image Process.* **17** (2008) 1047–1060
15. Bresson, X., Chan, T.: Non-local unsupervised variational image segmentation models. *UCLA CAM Report* 08-67 (2008)
16. Kim, J., Fisher, J., Yezzi, A., Cetin, M., Willsky, A.: A nonparametric statistical method for image segmentation using information theory and curve evolution. *IEEE Trans. Image Proc.* **14** (2005) 1486–502
17. Zhu, S., Yuille, A.: Region competition: unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *IEEE Trans. Patt. Anal. and Mach. Intell.* **18** (1996) 884–900
18. Jehan-Besson, S., Herbulot, A., Barlaud, M., Aubert, G.: *Segmentation of Vectorial Image Features Using Shape Gradients and Information Measures*. Springer (2005)
19. Herbulot, A., Besson, S.J., Duffner, S., Barlaud, M., Aubert, G.: Segmentation of vectorial image features using shape gradients and information measures. *Journal of Mathematical Imaging and Vision* **25** (2006) 365–386
20. Ni, K., Bresson, X., Chan, T., Esedoglu, S.: Local histogram based segmentation using the wasserstein distance. *Int. J. Comput. Vis.* **84** (2009) 97–111
21. Bao, Z., Liu, Y., Peng, Y., Zhang, G.: Variational color image segmentation via chromaticity-brightness decomposition. *MMM LNCS* 5916 (2010) 295–302
22. Rubner, Y., Tomast, C., Guibas, L.: The earth mover's distance as a metric for image retrieval. *Int. J. Comput. Vis.* **40** (2000) 99–121
23. Haker, S., Zhu, L., Tannenbaum, A.: Optimal mass transport for registration and warping. *Int. J. Comput. Vis.* **60** (2004) 225–240
24. Rabin, J., Peyré, G., Delon, J., Bernot, M.: Wasserstein barycenter and its application to texture mixing. *Proc. of SSVM 2011* (2011)
25. Burkard, R., Dell'Amico, M., Martello, S.: *Assignment problems*. SIAM (2009)
26. Chan, T., Sandberg, B., Vese, L.: Active contours without edges for vector-valued images. *J. Vis. Comm. Image Repr.* **11** (2000) 130–141




Fig. 6. Texture segmentation using our new normalized model (5) with the sliced Wasserstein distance (11) with $|\Theta| = 3$ fixed directions, and comparison with the vector-valued Chan-Vese model [26] and the extended work [21] of [20]. Patches with width $\tau = 11/n$ (1st), $\tau = 5/n$ (2nd), $\tau = 9/n$ (3rd), $\tau = 7/n$ (4th-5th) are used, and 31×31 windowing functions are used.