

NON-LOCAL SEGMENTATION AND INPAINTING

Miyoun Jung, Gabriel Peyré and Laurent D. Cohen

Ceremade, Université Paris-Dauphine, 75775 Paris Cedex 16, France

ABSTRACT

This article introduces a new variational image segmentation method that makes use of non-local comparisons between pairs of patches in the image and is robust to missing data (e.g. damaged pixels or large missing regions). The resulting segmentation is at the heart of a novel inpainting algorithm that also uses a non-local regularization. This segmentation and inpainting approach only requires a local homogeneity of the features inside and outside the region to be segmented. In contrast to existing region-based segmentation methods, it allows us to segment regions with smoothly varying intensity as well as multiple objects with different intensities. This comparison principle is also less sensitive to initialization than edge-based approaches.

Index Terms— segmentation, active contour, inpainting, non-local energy, level set function

1. INTRODUCTION

Region-based and edge-based segmentation. Image segmentation refers to the process of partitioning an image into several regions or locating objects and boundaries. Many existing methods segment an image according to edge information and/or region information. Edge-based methods such as geodesic active contours of Caselles et al. [1] use edge detection functions depending on the image gradient and evolve contours towards sharp gradients of pixel intensity. Chan and Vese [2] proposed an active contours without edges model, which is also based on a level set implementation [3], but the gradient-based information is replaced by a criterion related to region homogeneity. This model approximates an image by a two-phase piecewise constant function. This was also extended to vector valued images [4]. Sagiv et al. [5] used a hybrid model that incorporates a more general weighted arc-length in the active contours without edges model, for the problem of texture segmentation. In this article, we propose an active contour model with a novel energy functional that allows to only constrain the local homogeneity, in contrast to the Chan-Vese approach. The local homogeneity property allows our model to capture regions with features that vary spatially in a smooth way, as well as to segment several separated objects with different features.

Variational inpainting. Image inpainting is the process of reconstructing lost or corrupted parts of an image. Image inpainting has received considerable interest since the work by Masnou and Morel [6] who proposed variational principles for image disocclusion. A recent wave of interest in inpainting has also started from [7] of Bertalmio et al. These authors proposed nonlinear partial differential equations for non-texture inpainting. In this work, we incorporate the segmentation model with inpainting problem so that the interactions between pairs of patches inside and outside the segmented domain can improve the result of inpainting.

Non-local segmentation and inpainting. Nonlocal image processing regularizes the image using comparison of patches in the image. This was originally proposed by Buades et al [8] for denoising, and has been extended to regularize inverse problems [9] and regularize partition functions for semi-supervised and supervised segmentation [10, 11, 12].

Contributions. We introduce a novel non-local framework for unsupervised segmentation and inpainting of images. Previous non-local segmentation methods only use patch interactions to regularize the contours of the segmentation. In contrast, we propose a novel energy that compares pairs of patches within the segmented regions, thus extending region based method to a non-local setting. This is particularly useful to segment object and background with smoothly varying intensity and color. We also take into account missing data to enable a robust segmentation of damaged images, thus enabling an inpainting process as a post-processing.

2. NON-LOCAL ACTIVE CONTOURS

Given some missing region $D \subset [0, 1]^2$, we aim at performing a segmentation and then an inpainting (filling in the image inside D) from the noisy damaged observation $f = \Phi u_0 + w$. Here u_0 is the original high resolution image to recover, w is an additive noise, and Φ is the damaging operator. It is a masking operator: $(\Phi u)(x) = u(x)$ for $x \notin D$ and $(\Phi u)(x) = 0$ if $x \in D$. If $D = \emptyset$, then $\Phi = \text{Id}$ and one is only interested in segmentation.

2.1. Non-local Segmentation

Pairwise Patch Interaction. A patch in some image f around a pixel $x \in [0, 1]^2$ is defined as

$$\forall t \in [-\tau/2, \tau/2]^2, \quad p_x^f(t) = f(x + t). \quad (1)$$

The non-local interaction between two patches is measured using a weighted L^2 distance

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

with $G_a(t) = e^{-\frac{\|t\|^2}{2a^2}}$. The Gaussian weight is used to give more influence to the central pixel. Note that missing pixels in patches are not accounted for the computation of the distance.

Non-local segmentation energy. The segmented region Ω is represented using a level set function $\varphi : [0, 1]^2 \rightarrow \mathbb{R}$ so that $\Omega = \{x \mid \varphi(x) > 0\}$. The integration inside and outside the domain is carried over using a smoothed Heaviside function $H(\varphi) = \frac{1}{2} + \frac{1}{\pi} \text{atan}(\varphi/\varepsilon)$. 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.

We introduce an energy functional $E(\varphi)$ enforcing the similarity of features located either inside or outside Ω ,

$$E(\varphi) = \iint \rho(H(\varphi(x)), H(\varphi(y))) G_\sigma(x - y) \cdot d(p_x^f, p_y^f) T(x, y) dx dy \quad (3)$$

The function ρ restricts the comparison of pairs of patches that are in the same region (inside or outside). Since $H(\varphi(x))$ takes values 0 or 1 most of the time, we use $\rho(u, v) = 1 - |u - v|$ for the numerical experiments (but other similar functionals could be used as well). The function $T(x, y)$ excludes patches whose center is located in the damaged region from the comparison, and we use $T(x, y) = 0$ if $x \in D$ or $y \in D$, and $T(x, y) = 1$ otherwise. 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 extracted region, following previous works in active contours, we penalize the length of the boundary, 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 $H(\varphi)$.

Non-local segmentation algorithm. Our robust non-local active contour method computes the segmentation as a stationary point of the energy

$$\min_{\varphi} E(\varphi) + \gamma L(\varphi) \quad (4)$$

where $\gamma > 0$ is a parameter that should be adapted to the expected regularity of the boundary of the region. Note that within the missing region only the term $L(\varphi)$ regularizes the contours, which is thus expected to be a straight line segment in each connected component. Penalization using higher order curve derivatives could be used to allow curvilinear contours completion, such as those proposed in [6, 7].

We compute a local minimizer of (4) using a gradient descent, that corresponds to a discretization of the following evolution equation with an artificial time parameter $t > 0$

$$\frac{\partial \varphi_t}{\partial t} = -(\nabla E(\varphi_t) + \gamma \nabla L(\varphi_t)), \quad (5)$$

where $\varphi_t(x)$ is the level set function at time t and the gradient $\nabla E(\varphi_t)$ is computed as

$$\begin{aligned} \nabla E(\varphi_t)(x) &= \int (\partial_1 \rho)(H(\varphi_t(x)), H(\varphi_t(y))) G_\sigma(x - y) \\ &\quad \cdot d(p_x^f, p_y^f) T(x, y) dy H'(\varphi_t(x)), \\ \nabla L(\varphi_t)(x) &= -\text{div} \left(\frac{\nabla \varphi_t(x)}{\|\nabla \varphi_t(x)\|} \right) H'(\varphi_t(x)). \end{aligned}$$

The evolution equation (5) for φ_t is solved with an explicit scheme, and $H'(\varphi_t)$ is replaced by $\|\nabla \varphi_t\|$ to evolve all the level sets of φ_t in parallel. To ensure the stability of the level set evolution (5), one needs to re-initialize it from time to time. This corresponds to replacing φ_t by the signed distance function to the level set $\{x \mid \varphi_t(x) = 0\}$.

2.2. Non-local Inpainting

Once the evolution (5) has converged to some final segmentation function, $\varphi_t \rightarrow \varphi_*$, we perform the inpainting by minimizing the non-local energy with respect to the image to fill in the missing values.

Non-local inpainting energy. Inpainting is achieved by solving the following convex (quadratic) minimization problem over the image u

$$\min_u \frac{\lambda}{2} \|f - \Phi u\|^2 + F(u) \quad (6)$$

where $\lambda > 0$ is a parameter that should be adapted to the noise level $\|w\|$, and where the non-local interaction functional reads

$$F(u) = \iint \rho(H(\varphi_*(x)), H(\varphi_*(y))) G_\sigma(x - y) \cdot d(p_x^u, p_y^u) dx dy. \quad (7)$$

Non-local inpainting algorithm. The minimization (6) is obtained using a conjugate gradient descent to solve

$$\lambda \Phi^*(\Phi u - f) + \nabla F(u) = 0,$$

$$\text{where } \nabla F(u)(x) = \int \rho(H(\varphi_*(x)), H(\varphi_*(y))) \cdot G_\sigma(x - y) \nabla d(p_x^u, p_y^u) dy.$$

Parameters for segmentation and inpainting. It is possible to use a different patch width τ in the definition (1) of the patches for the segmentation energy (3) and for the inpainting energy (7). In the numerical examples, we used a segmentation patch size of $\tau = 3$ or 1 and an inpainting size of $\tau = 1$ pixels. The width σ of the windowing function $G_\sigma(x - y)$ depends on the initial curve: if the initial curve is far away from the object boundaries, then a large windowing function may be required. Here, 31×31 or 41×41 are used with a fixed $\sigma = 10$ for 100×100 or 200×200 images.

3. EXPERIMENTAL RESULTS AND COMPARISONS

This section presents experimental results with synthetic and real images.

Segmentation Results. We compare our approach with both region-based and edge-based active contour segmentation models, by minimizing a hybrid energy of the form


$$\min_{\varphi, p} \alpha E_r(\varphi, p) + (1 - \alpha) E_c(\varphi) + \gamma L(\varphi) \quad (8)$$

where α weights the influence of the region term E_r and the edge term E_c :

$$\begin{aligned} E_r(\varphi, p) &= \lambda_1 \int H(\varphi(x)) d(p_x, p_1) dx \\ &\quad + \lambda_2 \int (1 - H(\varphi(x))) d(p_x, p_2) dx, \\ E_c(\varphi) &= \mu \int \|\nabla H(\varphi(x))\| g(x) dx, \end{aligned}$$

with positive parameters $\lambda_1, \lambda_2, \mu$ and a positive edge function g , and where p represents the expected constant value of the features inside and outside the object. In particular, we consider (A) the geodesic active contour model ($\alpha = 0$, GAC model [13]) with the balloon force term $\eta g \|\nabla \varphi\|$ that helps to avoid poor local minima by forcing moving the curve forward/outward (depending on the sign of η): $\frac{\partial \varphi}{\partial t} = \mu \|\nabla \varphi\| \operatorname{div} \left(g \frac{\nabla \varphi}{\|\nabla \varphi\|} \right) + \eta g \|\nabla \varphi\|$, (B) the region-based model ($\alpha = 1$) of Chan and Vese [2], and (C) the integrated region/edge based model ($\alpha = 1/2$) of Sagiv et al. [5], called IAC model. Note that, in practice, we use an edge function $g(x) = \frac{1}{\delta^2 + (G_b^* \|\nabla f\|)(x)}$ with $\delta^2 = 0.1$ and $b = 0.5$, and then we normalize it from 0 to 1. And we let $\lambda_1 = \lambda_2 = 1$.

In Fig. 1, our model detects objects with spatially varying background or with small gradients on the boundary, as well as multiple separated objects with different intensities, unlike two phase Chan-Vese model. This is due to the local homogeneity property of our model, while the two-phase Chan-Vese model [2] requires a global homogeneity in each region. We note that our model needs a small number of iterations (around 150 iterations) to obtain final curves, even with an explicit scheme.



Top: initial curve, Middle: Our model, Bottom: Chan-Vese model

Fig. 1. Detection of objects with spatially varying background or object, and comparison with Chan-Vese model ($\alpha = 1$) in (8). 100×100 image and 31×31 windowing function is used. Our model: final curves are obtained at $k = 150$ (1st, 2nd), $k = 80$ (3rd), $k = 120$ (4th).




Fig. 2. Final curves of models given in (8): $\alpha = 0$ with balloon force term (GAC), $\alpha = 1/2$ (IAC) with two different but close parameters μ . GAC: $\mu = 1, \eta = -0.3$. IAC: (top) $\mu = 3.6$ and 3.5 , (bottom) $\mu = 1$ and 0.9 . IAC used initial curves given in Fig. 1.

Fig. 2 presents the results of other existing models, given in (8): $\alpha = 0$ with balloon force term $\eta g(x) \|\nabla \varphi(x)\|$ (geodesic model), $\alpha = 1/2$ (integrated active contour model). For the IAC model, two final curves are shown with two different but close parameters μ ($\mu_1 > \mu_2$). Because μ is a balancing term between the region-based and edge-based energies, when $\mu > \mu_1$ (or $\mu < \mu_2$), the model tends to act like the geodesic snake model (or Chan-Vese model). Thus, with the given initial curves, all the models fail to detect the correct object boundaries. This shows that our model is less sensitive to the choice of initial curves than edge-based active contour models.

In Fig. 3, by using an initial curve near the boundary of object(s) and a small windowing function, our model could detect the boundary of non-homogeneous object(s). The segmentation result is fairly good, compared with vector-valued




Fig. 3. Real color images. Initial Curve, final curves of our model, vector-valued Chan-Vese model [4], and integrated active contour model (IAC).




Fig. 4. Segmentation in images with missing pixels, and post-processing inpainting. 1st: segmentation result, 2nd: inpainting with φ obtained by the segmentation, 3rd: inpainting using (6) but without ρ in $F(u)$.

Chan-Vese model [4] and IAC model [5] that only capture part of object(s). On the other hand, these examples also show a limitation of our model: in order to detect the boundary of non-homogeneous objects, the initial curve needs to be located near the object boundary so that a small windowing function can be used.

Inpainting Results. In Fig. 4, the first example has large missing regions and the second one has 50% of missing pixels. Our segmentation model detects the boundaries of objects in the inpainting area, which helps inpainting process by restricting the inpainting area to either inside or outside objects. The edges of objects are also kept well after inpainting process. These are well shown by the comparison with the inpainting results using our model (6) but without ρ in the energy $F(u)$ (without using the segmentation results), especially in the first example.

4. CONCLUSION

In this article, we have proposed a novel non-local energy for segmentation and inpainting of images. The local homogeneity property allows us to detect regions with smoothly spatially varying features and segment several separated objects with different features. It is also less sensitive to the choice of initial curves than edge-based segmentation mod-

els. Furthermore, the segmentation model is robust to missing data, which enables incorporating segmentation with inpainting to improve the inpainting process.

5. REFERENCES

- [1] V. Caselles, F. Catté, T. Coll, and F. Dibos, "A geometric model for active contours in image processing," *Numerische Mathematik*, vol. 66, no. 1, pp. 1–31, 1993.
- [2] T. Chan and L. Vese, "Active contours without edges," *IEEE Trans. Image Proc.*, vol. 10, no. 2, pp. 266–277, 2001.
- [3] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, pp. 12–49, 1988.
- [4] T. Chan, B. Sandberg, and L. Vese, "Active contours without edges for vector-valued images," *J. Vis. Comm. Image Repr.*, vol. 11, no. 2, pp. 130–141, 2000.
- [5] C. Sagiv, N. A. Sochen, and Y. Y. Zeevi, "Integrated active contours for texture segmentation," *IEEE Trans. Image Proc.*, vol. 15, no. 6, pp. 1633–1646, 2006.
- [6] S. Masnou and J.M. Morel, "Level lines-based disocclusion," *IEEE Int. Conf. Image Processing*, vol. 3, pp. 259–263, 1998.
- [7] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," *Siggraph*, pp. 417–424, 2000.
- [8] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *SIAM Mul. Model. and Simul.*, vol. 4, no. 2, pp. 490–530, 2005.
- [9] G. Peyré, S. Bougleux, and L. Cohen, "Non-local regularization of inverse problems," *ECCV, Part III, LNCS*, vol. 5304, pp. 57–68, 2008.
- [10] G. Gilboa and S. Osher, "Nonlocal linear image regularization and supervised segmentation," *SIAM Mul. Model. and Simul.*, vol. 6, no. 2, pp. 595–630, 2007.
- [11] A. Elmoataz, O. Lezoray, and S. Bougleux, "Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing," *IEEE Trans Image Process.*, vol. 17, no. 7, pp. 1047–1060, 2008.
- [12] X. Bresson and T. Chan, "Non-local unsupervised variational image segmentation models," *UCLA CAM Report 08-67*, 2008.
- [13] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International Journal of Computer Vision*, vol. 22, no. 1, pp. 61–79, 1997.