

AN ACTIVE REGIONS APPROACH FOR THE SEGMENTATION OF 3D BIOLOGICAL TISSUE

Juan Cardelino, Gregory Randall

IIE, Facultad de Ingeniería
Universidad de la República
Montevideo, Uruguay
{juanc,randall}@fing.edu.uy

Marcelo Bertalmío

Departament de Tecnologia
Universitat Pompeu Fabra
Barcelona, Spain
marcelo.bertalmio@upf.edu

ABSTRACT

Some of the most successful algorithms for the automated segmentation of images use an Active Regions approach, where a curve is evolved so as to maximize the disparity of its interior and exterior. But these techniques require the manual selection of several parameters, which make impractical the work with long image sequences or with a very dissimilar set of sequences. Unfortunately this is precisely the case with 3D biological image sequences. In this work we improve on previous Active Regions algorithms in two aspects: by introducing a way to compute and update the optimum weights for the different channels involved (color, texture, etc.) and by estimating if the moving curve has lost any object so as to launch a re-initialization step. Our method is shown to outperform previous approaches. Several examples of biological image sequences, quite long and different among themselves, are presented.

1. INTRODUCTION

Segmentation, reconstruction and visualization of 3D structures from series of 2D image slices is a widely used technique in medical and biological research. These series have tens to hundreds of slices, so the images may vary a lot along the same sequence. Also, different image acquisition techniques (e.g. electron, optical or confocal microscopy) and different biological data produce very dissimilar sequences. Our goal (part of an ongoing effort started with [1]) is to devise a general technique, independent of the acquisition technique and the type of cellular tissue involved that, once the biologist has selected in the first frame of the sequence the structures in which he/she is interested, is able to automatically follow these structures along the sequence, segmenting them as they evolve and also detecting and segmenting any new objects of the same kind that may appear.

Some of the most successful algorithms for the automated segmentation of images use an Active Regions approach ([2, 3, 4, 5, 6, 7, 8], where a curve is evolved so as

to maximize the disparity of its interior and exterior. But these techniques require the manual selection of several parameters, which make impractical the work with long image sequences or with a very dissimilar set of sequences.

Our contribution is the following: based on the remarkable Geodesic Active Regions segmentation algorithm of Rousson et al. [7], we increase its robustness (by estimating if the moving curve has lost any object so as to launch a re-initialization step) and make automatic the crucial choice of the optimum weights for the different channels involved (color, texture, etc).

This paper is organized as follows. In section 2 we present the general framework. Our algorithm is described in section 3. Section 4 shows the experimental results. Finally, in section 5 we give some conclusions.

2. GENERAL FRAMEWORK

Our starting point is the region-based active contour framework proposed by Rousson et al. [7]. For a single (color) image $I : \Omega \rightarrow R^3$, with $\Omega \subset R^2$, we compute a corresponding *feature image* $U : \Omega \rightarrow R^N$ where N is the number of feature channels (like color, texture, optical flow, etc).

To solve the segmentation problem, we seek for a curve Γ which gives a partition of the space Ω in two (non necessarily connected) regions Ω_1 and Ω_2 : the object of interest and the background. This is accomplished by finding the curve that divides the feature image in such a way that maximizes the likelihood of each region probability distribution with respect to the observed previous frame. This process is repeated for each frame, taking as ground truth the segmentation of the previous one. The algorithm takes as input the sequence of images and an user provided ground truth for the first image (a set of manually segmented objects of interest).

Under the hypothesis that both regions are *a-priori* equiprobable, following [5] and using a level-set formula-

tion we arrive to the following Partial Differential Equation:

$$\frac{d\phi}{dt} = \left(\alpha\kappa + \delta(\phi) \sum_{j=1}^N w_j \log \frac{p_1^j(\mathbf{x})}{p_2^j(\mathbf{x})} \right) |\nabla\phi| \quad (1)$$

where $\phi(t)$ is an embedding function for the evolving curve, p_i^j is the Probability Density Function (PDF) of feature j in region i (usually computed with Gaussian estimation or a Generalized Gaussian Density) and κ is the Euclidean curvature. The parameters that need to be chosen and updated for each frame are the w_j , the weights for the different channels.

The desired curve Γ is obtained computing the zero level-set of ϕ when $t \rightarrow \infty$.

3. OUR CONTRIBUTION

3.1. Weighting algorithm

We propose a method for the automatic choice and updating of the weights w_j . For the first frame the user supplies a segmentation which is taken as ground truth. For each feature channel j we compute the PDF in Ω_1 , call it p_1^j , and the PDF in Ω_2 , call it p_2^j . Then we measure the distance d_j between p_1^j and p_2^j , using a symmetrized version of the Kullback-Leibler divergence [9]. A big value for d_j means that the channel j provides good discrimination, so w_j should be big as well. We choose:

$$w_j = \frac{d_j}{\sum_{i=1}^N d_i} \quad (2)$$

For the second and following frames, the procedure is the same, only that the PDF's are computed for the segmentation of the previous frame.

3.2. Reinitialization

As in many tracking algorithms, we take as initial estimation for the curve Γ in the $(n+1)$ -th frame the final result for Γ in the n -th frame. Assuming that the difference between two consecutive frames is not too big, this choice of initial condition speeds up the convergence of equation 1. It also minimizes the risk of falling into a local minima and not adequately segment our object of interest, something that would happen often if our starting condition were an arbitrary curve (a set of circles covering the image, for instance).

The downside of this choice of initial condition is that, if a new object of interest appears at frame $n+1$ that was not there at frame n , it will most probably not be detected. See for instance the sequence in fig. 1, where the objects of interest are the small dark round shapes, which generally keep

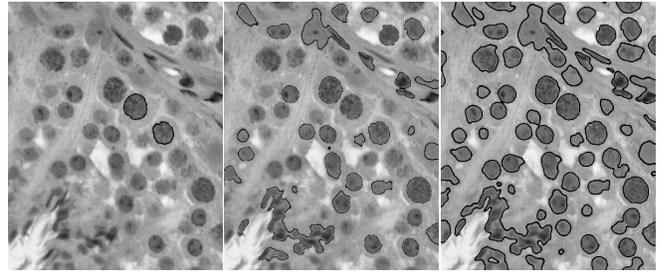


Fig. 1. Blobs detected in the 9th frame of the *ram* sequence. Left: no re-initialization. Center: initializing with evenly spaced circles. Right: our re-initialization (see text.)

their position (so they are easy to track) but keep appearing and disappearing along the sequence. For this sequence, choosing as initial condition the curve(s) found in the previous frame would imply that only the objects present in the first frame of the sequence would be tracked, they would gradually disappear and after just a few frames the segmentation would yield an empty set, although there are always many objects of interest in each one of the frames of this sequence.

To solve this problem, we propose a method to detect if the curve Γ is missing any object, and if it does we launch a re-initialization step. This mechanism is composed of two stages: decision and re-initialization.

The decision mechanism works as follows. As the evolution reaches a minimum, we assume that we have a considerable part of the object detected. Thus we can obtain a roughly accurate estimation of the object PDF in the inner region (Ω_1). In the outer region, we have the background and the missing part of the object. So we can model the distribution of Ω_2 as a mixture model $p_2(x) = P_a.p(x|\theta_a) + P_b.p(x|\theta_b)$ where P_i are the *prior* probabilities of each distribution and θ_i its parameters. We use the Expectation Maximization algorithm to learn P_i and θ_i . After the mixture is trained, the distance between $p_1(x)$ and $p(x|\theta_i) \forall i$ (computed as in 3.1) is measured and if it is lower than a threshold (d_{max}), we say that they are the same object. Once a match is found, we check if the area of the found object is bigger than a threshold ($P_i > a_{min}$), and if so, we start the re-initialization process.

For the re-initialization, we modify the front in such a way that we get two curves: the current curve and another one that will evolve outwards to find the lost objects. This is achieved simply by displacing the *provisional* curve Γ outwards a few pixels, inverting the front, and intersecting with the original curve. Figure 1 illustrates the importance of such a mechanism.

An important property of the re-initialization algorithm is that it is devised in such a way that the segmentation previously achieved cannot be hindered and the detected ob-



Fig. 2. 14th frame of the *crypt* sequence. Left: fixed weights. Right: using automatic weighting.

jects cannot be lost. However this step has two major drawbacks: evolving outwards we can add spurious objects to the previously achieved segmentation and it increases the (time consuming) iterations for each frame treatment. For these reasons it is important to use the decision mechanism in order to re-initialize only when it is needed.

4. RESULTS

4.1. Implementation

The framework permits the combination of different features in the segmentation process, and the automatic weighting mechanism is responsible for the relative importance of each channel in the segmentation of a given sequence. In these experiments we use the *YCbCr* components for color images or the intensity channel in the case of gray level images. In order to produce texture feature channels, we process the intensity channel with the wavelet packet transform (assuming that all texture information is contained in the *Y* channel).

In these experiments we use six feature channels: intensity *Y*, *Cb* and *Cr* for the color images, and the horizontal, vertical and diagonal detail for the first level of wavelet decomposition.

The values chosen for the parameters needed by the algorithm are fixed and the same for all the experiments: the weight of the regularization term α , the minimum size of the detected object a_{min} , the maximum distance between the PDFs d_{max} , and the stopping condition for the curve evolution.

Our algorithm was tested with four sequences of images of real biological tissue¹. The series *spider* (560x460x50) and *ram* (1000x800x19) are obtained from optical microscopy (OM). *Crypt* (2200x1000x32) is a collage of OM images and the mosaic effect is quite apparent. The *Echinococcus* sequence (512x768x50) is produced by confocal microscopy (CM). Figure 1 shows some results over the *ram* sequence. The biological material corresponds to slices of seminiferous cells of rams. This sequence is composed by many small objects, which tend to appear only

¹Some of this sequences are in color, a more detailed color version can be found at <http://iie.fing.edu.uy/juanc/publ/icip05.pdf>

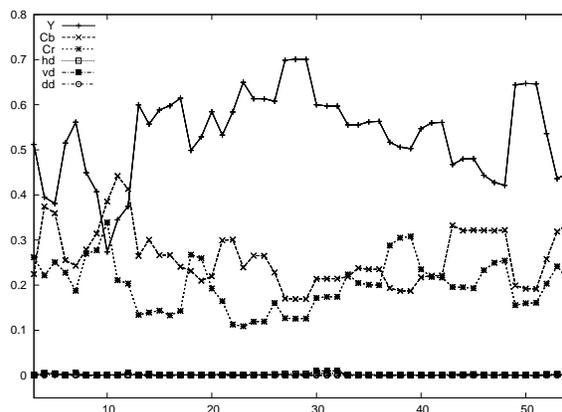


Fig. 3. Evolution of the weights for the *spider* sequence.

for a few frames. Sometimes the decision algorithm fails, and no re-initialization is performed, so we lose some objects. However, the overall number of lost objects reduces dramatically using this mechanism, as the figure shows. In order to be sure that no object is lost, the re-initialization can be forced in all frames.

Figure 4 shows the results on the *spider* sequence. The objects are successfully tracked, despite the important topology changes taking place over the 50 frames sequence. However if a similar color structure appears close to the boundary, the algorithm detects it as part of the object (e.g. frames 30 and 50).

Figure 3 illustrates the evolution of the computed parameters w_j for this sequence. Some features become more relevant (for the segmentation process) in different frames, compare for instance the weights at frames 10 and 30. In all the tested sequences the weighting mechanism outperforms the fixed-weights approach, even if we choose manually the channels that contain relevant information. In addition, the computed weights are close to the manual choice. However, their variability over the sequence shows the importance of using an adaptive strategy rather than a fixed approach. Figure 2 shows the results on the *crypt* sequence. Here, the interesting structure is the homogeneous region which forms a cavity that must be reconstructed in order to analyze its shape. The images are built with a collage of several OM images and some mosaic artifacts can be seen. The structure is well segmented in general even if some regions are lost.

Figure 5 shows the results on the *echinococcus* sequence. This is a very different sequence, obtained by Confocal Microscopy and characterized by the absence of strong edges. The structures marked by the biologist in the first frame are well tracked along the sequence. After the 40th frame, the point density becomes too low and the algorithm quickly loses the structure of interest.

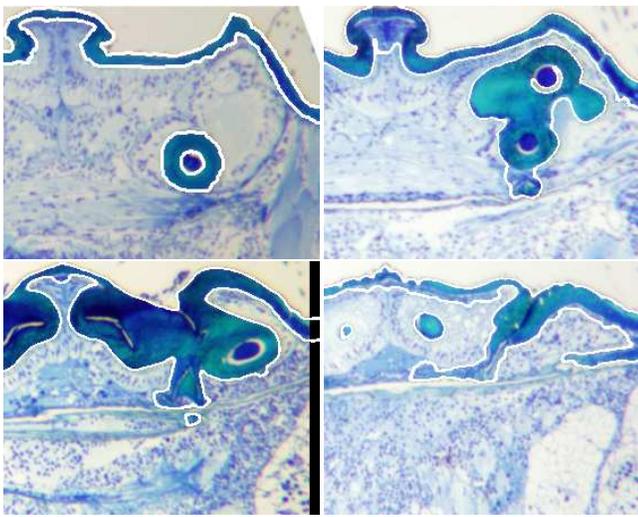


Fig. 4. *spider*: frames 1,10,30,50

5. CONCLUSION AND FUTURE WORK

The Active Regions framework allows for a combination of cues in the segmentation process. Critical issues are the initialization criteria and the estimation of the relative weights of the features. We have added two novel mechanisms: automatic re-initialization and feature weighting. Our experiments show that these mechanisms are essential when dealing with long and variable sequences. We are currently addressing the improvement of our technique regarding several aspects: adding some *inertia* to the computation of the weights, taking more than one frame into account, refining the results by running the sequence *backwards*, and introducing some sort of *shape* information into the procedure.

6. ACKNOWLEDGMENTS

This work was partially supported by AECI-PCI 2003, RACINE European Commission Project, SpeedFX IST-2001-34337 Project, Ramón y Cajal Program and PNP GC (BFM2003-02125).

7. REFERENCES

- [1] G. Randall; A. Fernández; O. Trujillo; F. Malmierca; P. Morelli; G. Apelbaum; M. Bertalmío and L. Vázquez, "Neuro3d: an interactive 3d reconstruction system of serial sections using automatic registration," *SPIE BiOS98, California*, vol. 22, pp. 61–79, March 2002.
- [2] R. Ronfard, "Region-based strategies for active contour models," *Int. J. Comp. Vision*, vol. 2, pp. 229–251, 13 1994.

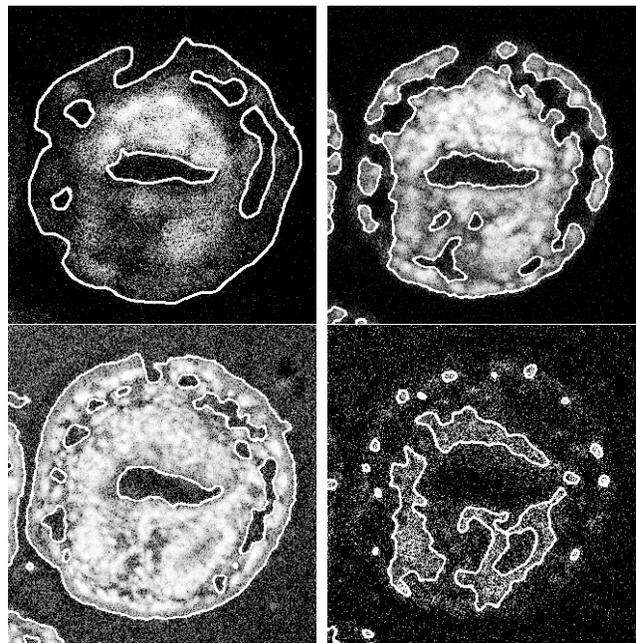


Fig. 5. *echinococcus*: frames 15,20,30,40

- [3] L. Cohen, "On active contour models and balloons," *Comp. Vision, Graphics and Image Processing: Im. Understanding*, no. 53, pp. 211–218, 1991.
- [4] T. Chan; L. Vese, "Active contour without edges," *IEEE Trans. Image Proc.*, vol. 2, no. 10, pp. 266–277, 2001.
- [5] N. Paragios; R. Deriche, "Geodesic active regions: A new framework to deal with frame partition problems in computer vision," *J. of Visual Comm. and Image Representation*, vol. 13, no. 1, pp. 249–268, June 2002.
- [6] S. C. Zhu; A. Yuille, "Region competition: Unifying snakes, region growing and bayes/mdl for multi-band image segmentation," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. 79, pp. 12–49, September 1996.
- [7] T. Brox; M. Rousson; R. Deriche and J. Weickert, "Un-supervised segmentation incorporating colour, texture, and motion," *10th Int. Conf. on Comp. Analysis of Images and Patterns*, vol. 2756, pp. 353–360, August 2003.
- [8] J-F. Aujol; G. Aubert; L. Blanc-Féraud, "Wavelet-based level set evolution for classification of textured images," *IEEE Trans. Image Proc.*, vol. 12, no. 12, December 2003.
- [9] D. Johnson; S. Sinanovic, "Symmetrizing the kullback-leibler distance," *submitted to IEEE Trans. on Information Theory*, 2001.