

Active Contours without Edges for Vector-Valued Images¹

Tony F. Chan,² B. Yezriev Sandberg,³ and Luminita A. Vese⁴

*Department of Mathematics, University of California, Los Angeles, 405 Hilgard Avenue,
Los Angeles, California 90095-1555*

Received November 5, 1999; accepted December 10, 1999

In this paper, we propose an active contour algorithm for object detection in vector-valued images (such as RGB or multispectral). The model is an extension of the scalar Chan–Vese algorithm to the vector-valued case [1]. The model minimizes a Mumford–Shah functional over the length of the contour, plus the sum of the fitting error over each component of the vector-valued image. Like the Chan–Vese model, our vector-valued model can detect edges both with or without gradient. We show examples where our model detects vector-valued objects which are undetectable in any scalar representation. For instance, objects with different missing parts in different channels are completely detected (such as occlusion). Also, in color images, objects which are invisible in each channel or in intensity can be detected by our algorithm. Finally, the model is robust with respect to noise, requiring no a priori denoising step. © 2000 Academic Press

Key Words: vector-valued images; active contours; level sets; segmentation; PDEs; object detection.

0. INTRODUCTION

Active contours are used to detect objects in a given image u_0 using techniques of curve evolution. The basic idea is, starting with an initial curve C , to deform the curve to the boundary of the object, under some constraints from the image u_0 .

The classical approaches use the gradient of the image u_0 to locate the edges. Therefore, an edge-function is used, strictly positive inside homogeneous regions and zero on the edges, to stop the evolving curve on the desired boundary. Using this edge-function, the deformation is usually obtained by minimizing with respect to C a functional, whose (local) minimizer is given by the boundary of the object [2]. The objective functional is the sum of two terms: the first controls the smoothness of the curve, while the second (depending on the

¹ This work was supported in part by ONR Contract N00014-96-1-0277 and NSF Contract DMS-9626755.

² E-mail: chan@math.ucla.edu.

³ E-mail: bsand@math.ucla.edu.

⁴ E-mail: lvese@math.ucla.edu.

gradient of u_0) attracts the curve toward the boundary. For instance, this model cannot handle automatic topology changes of the contour, and also it depends on the parameterization of the curve.

In problems of curve evolution, including snakes and active contours, the level set method of Osher and Sethian [3] has been used extensively because it allows for automatic topology changes, cusps, and corners; moreover, the computations are made on a fixed rectangular grid. Using this approach, geometric active contour models, using a stopping edge-function, have been proposed in [4] and also in [5–7]. These models are based on the theory of curve evolution and geometric flows and in particular on the mean curvature motion of Osher and Sethian [3]. The evolving curve moves by mean curvature, but with an extra factor in the speed, by the stopping edge-function. Therefore, the curve stops on the edges, where the edge-function vanishes. A typical example of edge-function used is given by

$$g(|\nabla u_0|) = \frac{1}{1 + |\nabla(G_\sigma * u_0)|^2},$$

where g is a positive and decreasing function such that $\lim_{t \rightarrow \infty} g(t) = 0$. The image u_0 is first convolved with the Gaussian $G_\sigma(x, y) = \sigma^{-1/2} \exp^{-|x^2+y^2|/4\sigma}$, especially for the cases where u_0 is noisy. But in practice, g is never zero on the edges, and therefore the evolving curve may not stop on the desired boundary. To overcome this problem, a new model has been proposed in [8] as a problem of geodesic computation in a Riemannian space, according to the metric g . The associated Euler–Lagrange equation, in level set formulation, has a new additional gradient term (compared with the previous geometric models [4–7]). This term increases the attraction of the evolving curve toward the boundary of the object and is of special help when the boundary has high variations on its gradient values. For another related approach, see also [9].

These models use the gradient of a smoother version of the image u_0 to detect edges. But, if the image is noisy, the smoothing in the edge-function has to be strong, thus blurring edge features, or a preprocessing has to be implemented, to remove the noise.

In contrast, the Chan–Vese (C-V) active contour model without edges proposed in [1] does not use the stopping edge-function g to find the boundary. Instead, the stopping term is based on Mumford–Shah segmentation techniques. This model has several advantages: it detects edges both with and without gradient (see [10] for a discussion on edges without gradient, called cognitive edges); it automatically detects interior contours; the initial curve does not necessarily have to start around the objects to be detected and instead can be placed anywhere in the image; it gives in addition a partition of the image into two regions, the first formed by the set of the detected objects, while the second one gives the background; finally, there is no need for an a priori noise removal.

The previously described methods have been developed to detect objects in a single image, but what happens when we have several different registered images of the same object? This occurs in multispectral images taken at different wavelengths, in medical images taken by different equipment (i.e., PET, MRI, and CT), in color images, or in textured images. Each image channel may have signal characteristics that can be combined with other channels to enhance contour detection.

Several models of restoration, edge detection, and also active contours have been proposed for vector-valued images. For restoration of color images, we mention the works in [11–13].

In [11, 12], vector edges are computed based on the classical Riemannian geometry, while in [13], the model is based on a particular extension of the total variation of Rudin, Osher, and Fatemi [14] to color images.

Based on the idea of vector edges defined using the classical Riemannian geometry from [11, 12], and on the geodesic active contour model for single-valued images introduced in [8], a color snakes model is introduced in [15, 16]. The notion of vector edges is then used to define the stopping edge-function. This model is also applied to vector-valued images obtained from a textured image.

Again, all these models for vector-valued images are based on the gradient of the image to discriminate edges.

A segmentation of color images based on the Mumford–Shah model and the region growing method is proposed in [17], and a segmentation of vector-valued images for texture discrimination is proposed in [18, 19]. We also refer the reader to a work on curve evolution and segmentation functionals for color images in [20] and to a work on snakes, region growing, and energy–Bayes–MDL for multiband image segmentation in [21].

A related geodesic active region model for texture segmentation, using a multivalued frame analysis, is proposed in [22].

In this paper, the Chan–Vese method [1] is extended to vector-valued images. Our algorithm uses the level set method of Osher and Sethian [3] to determine the boundary of the detected objects. An example of the vector-valued object detector can be seen in Fig. 1. Each channel has a different piece missing, but when the two channels are combined, the complete object is detected. Another example where this algorithm is of particular interest is an occlusion occurring in one channel, while a second channel, complete yet noisier, is available. Another example is RGB images, where intensity detectors and channel by channel boundary detectors fail. All these examples are covered in Section 2 on experimental results.

Other related works on occlusion in single-valued images have been proposed in [23] and more recently in [24].

The outline of the paper is as follows. The following section contains general background of the C-V model and its extension to vector-valued images. Section 2 contains experimental results. Finally, we end the paper by a short concluding section.

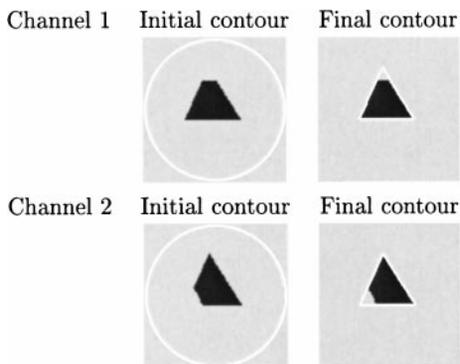


FIG. 1. Each channel has a different part of the same triangle missing. The vector-valued algorithm can detect the full triangle.

1. THE MODEL

1.1. Background: The Scalar Case

Let Ω be a bounded open subset of \mathbb{R}^2 , with $\partial\Omega$ the boundary. Let u_0 be a given image such that $u_0: \bar{\Omega} \rightarrow \mathbb{R}$. Let $C(s): [0, 1] \rightarrow \mathbb{R}^2$ be a piecewise parameterized C^1 curve.

We first recall the C–V model [1], which has the following form:

$$F(c^+, c^-, C) = \mu \cdot \text{Length}(C) + \lambda^+ \int_{\text{inside}(C)} |u_0(x, y) - c^+|^2 dx dy + \lambda^- \int_{\text{outside}(C)} |u_0(x, y) - c^-|^2 dx dy, \quad (1)$$

where c^+ and c^- are constant unknowns representing the average value of u_0 inside and outside the curve, respectively. The parameters $\mu > 0$ and $\lambda^+, \lambda^- > 0$ are weights for the regularizing term and the fitting term, respectively. The minimization problem

$$\inf_{c^+, c^-, C} F(c^+, c^-, C)$$

is considered as a reduction of the Mumford–Shah segmentation model [26]. Minimizing the fitting error in (1), the model looks for the best partition of u_0 taking only two values, namely c^+ and c^- , and with one edge C , the boundary between these two regions, given by $\{u_0 \approx c^+\}$ and $\{u_0 \approx c^-\}$. The object to be detected will be given by one of the regions, and the curve C will be the boundary of the object. The additional length term is a regularizing term and has a scaling role. If μ is large, only larger objects are detected, while for small μ , objects of smaller size are also detected. Because the model does not make use of a stopping edge-function based on the gradient, it can detect edges both with and without gradient.

For curve evolution, the level set method has been used extensively, in particular where the motion is governed by mean curvature, as in [3]. This formulation behaves well even with cusps, corners, and automatic topological changes. The motion by mean curvature [3] is given by

$$\begin{cases} \frac{\partial \phi}{\partial t} = |\nabla \phi| \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right), \\ \phi(0, x, y) = \phi_0(x, y), t \in [0, \infty), (x, y) \in \mathbb{R}^2, \end{cases}$$

where ϕ is the level set function, assumed Lipschitz continuous. By this evolution equation, the level curves of ϕ move by mean curvature, in the normal direction.

Now we can rewrite the original model (1) in the level set formulation. Let the evolving curve C be $C = \{(x, y) \in \Omega : \phi(x, y) = 0\}$, assuming that ϕ has opposite signs on each side of C . Following [27, 1], the energy can be written as:

$$F(c^+, c^-, \phi) = \mu \cdot \text{Length}\{\phi = 0\} + \lambda^+ \int_{\phi \geq 0} |u_0(x, y) - c^+|^2 dx dy + \lambda^- \int_{\phi < 0} |u_0(x, y) - c^-|^2 dx dy.$$

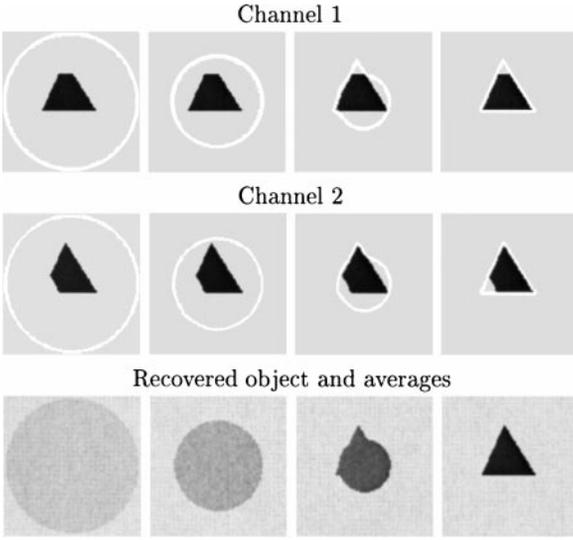


FIG. 2. Each channel has a different part of the same triangle missing. The boundary of the full triangle is found (for instance, the left wing is detected due to channel 1, while the fuselage is detected due to channel 2). This illustrates the ability of our algorithm to detect missing information in each channel and the complete object. The parameters are as follows: $\mu = 0.01 \cdot 211^2$, $\lambda_i^+ = \lambda_i^- = 1$, for $i = 1, 2$.

Using the heaviside function H , defined by

$$H(z) = \begin{cases} 1 & \text{if } z \geq 0, \\ 0 & \text{if } z < 0, \end{cases}$$

and the Dirac delta function $\delta(z) = \frac{d}{dz}H(z)$ (in the sense of distributions), we can rewrite

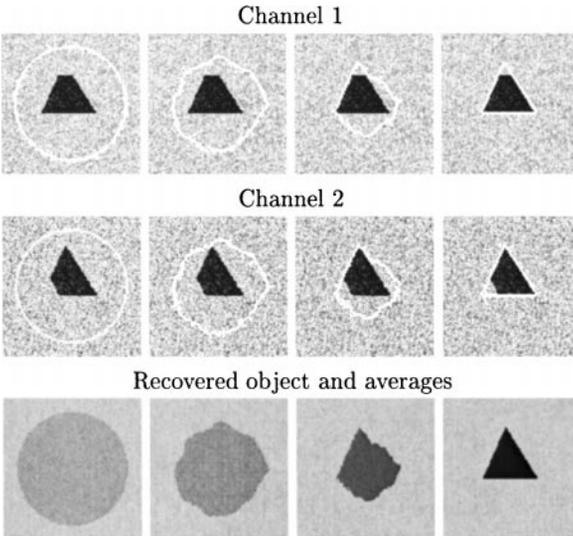


FIG. 3. In this example noise is present in both channels. The boundary of the full triangle is found. This is done without filtering the noise, allowing us to detect the edges without blurring the image, the method typically used by the classical approaches. The parameters are as follows: $\mu = 0.05 \cdot 255^2$, $\lambda_i^+ = 0.7$, $\lambda_i^- = 1$, for $i = 1, 2$.

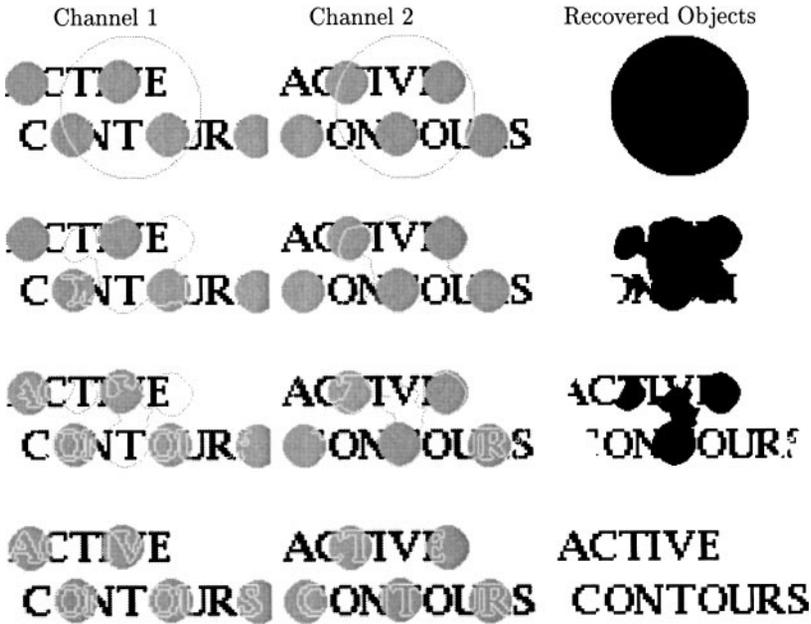


FIG. 4. Example of occlusion, present in both channels. Only using the combined channels, the complete information can be extracted. The parameters are as follows: $\mu = 0.05 \cdot 255^2$, $\lambda_i^+ = 2$, $\lambda_i^- = 1.2$, for $i = 1, 2$.

the energy functional as follows:

$$F(c^+, c^-, \phi) = \mu \int_{\Omega} \delta(\phi(x, y)) |\nabla \phi(x, y)| + \lambda^+ \int_{\Omega} |u_0(x, y) - c^+|^2 H(\phi(x, y)) dx dy \\ + \lambda^- \int_{\Omega} |u_0(x, y) - c^-|^2 (1 - H(\phi(x, y))) dx dy.$$

Minimizing $F(c^+, c^-, \phi)$ with respect to the constants c^+ and c^- , for a fixed ϕ , yields the following expressions for c^+ and c^- , function of ϕ :

$$\begin{cases} c^+ = \text{average}(u_0) \text{ on } \phi \geq 0, \\ c^- = \text{average}(u_0) \text{ on } \phi < 0. \end{cases}$$

Minimizing the energy $F(c^+, c^-, \phi)$ with respect to ϕ , for fixed c^+ and c^- , using a gradient descent method, yields the associated Euler–Lagrange equation for ϕ , governed by the mean curvature and the error terms (see [1] for more details).

The scalar C–V model has many advantages. In particular, it can detect contours without gradient (or cognitive contours) and interior contours automatically. Also, it has a level set formulation. It is therefore natural to extend this model to vector-valued images. This is presented next.

1.2. The Vector-Valued Model

We now present a natural extension of the previous scalar C–V model to the vector case.

Let $u_{0,i}$ be the i th channel of an image on Ω , with $i = 1, \dots, N$ channels, and C the evolving curve. Each channel would contain the same image with some differences, for

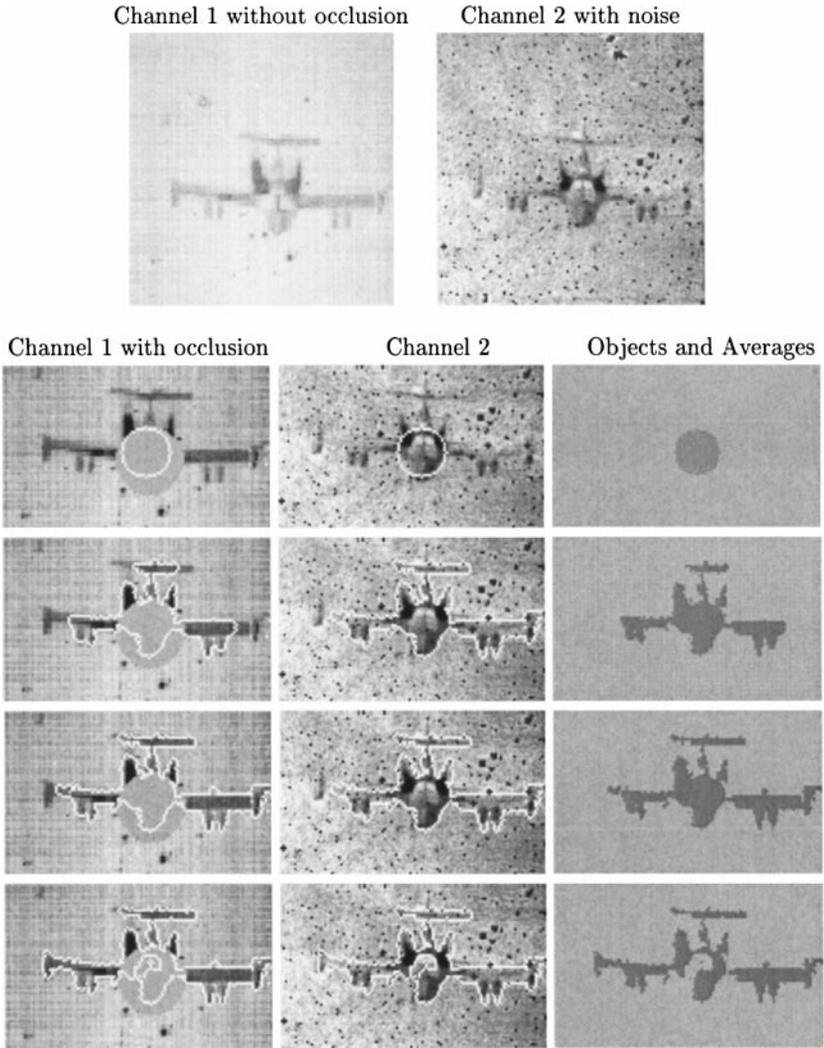


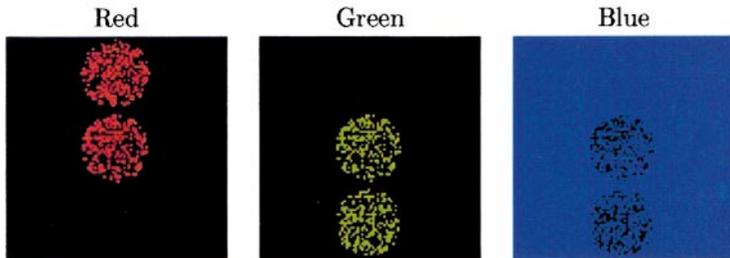
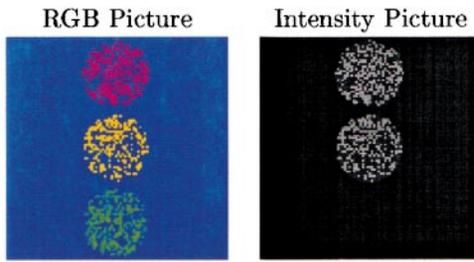
FIG. 5. While the first channel has not much noise, but has an occlusion in it, the second channel is very noisy. From these two pictures, we try to detect as much of the airplane as possible. The parameters are as follows: $\mu = 0.001 \cdot 255^2$, $\lambda_1^+ = \lambda_1^- = 1$, $\lambda_2^+ = \lambda_2^- = 0.55$. In this example, we first performed a renormalization of the channels to $[0, 255]$.

FIG. 6. We give here an example of a color image that has three objects of different colors, while the corresponding gray scale image only shows two of them. Here we use the transformation: Intensity Picture = $0.342 \cdot R + 0.500 \cdot G + 0.158 \cdot B$, the formula used by most image visualization tools. In this example, we have chosen the intensities in such a way that in the gray level picture, the third object has the same intensity as the background.

FIG. 7. The boundary of all the circles is found, while in the gray-scale image the boundary of one of the circles would never be detected. Note that, since this image does not have gradient edges, a gradient-based algorithm would not be able to find the three objects. Also, a channel by channel algorithm would always detect two out of three circles. The parameters are as follows: $\mu = 0.06 \cdot 255^2$, $\lambda_i^+ = \lambda_i^- = 1$, for $i = 1, 2, 3$.

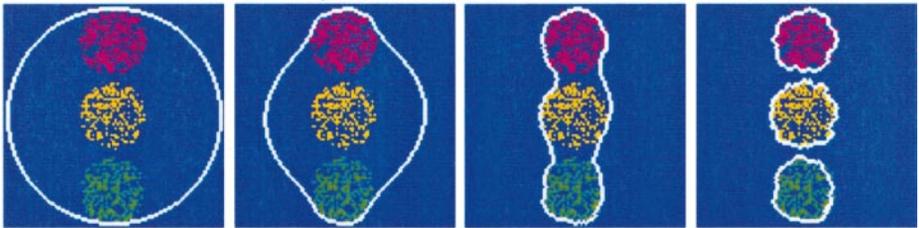
FIG. 8. Results on a real RGB image. The algorithm detects blurred contours, in noisy data. The parameters are as follows: $\mu = 0.03 \cdot 255^2$, $\lambda_i^+ = \lambda_i^- = 1$, for $i = 1, 2, 3$.

6

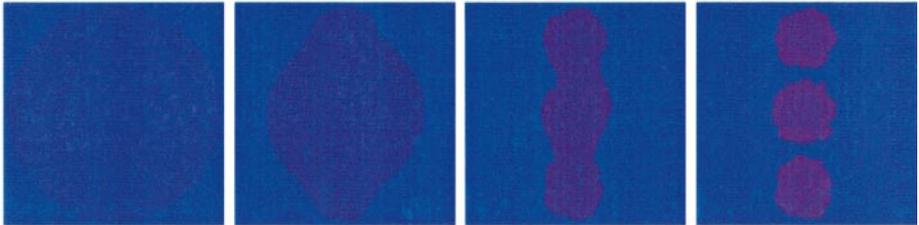


7

Recovered object contours combined in RGB mode

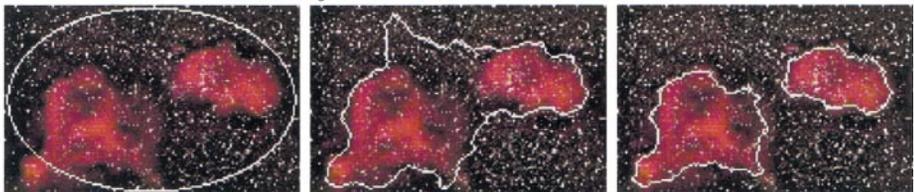


Recovered objects and averages combined in RGB mode

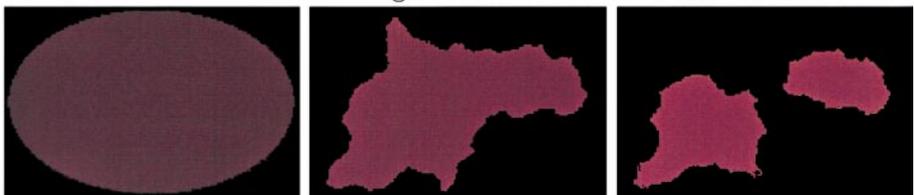


8

Recovered object contours combined in RGB mode



Recovered averages combined in RGB mode



instance different wavelengths at which the image was taken and color images. Let $\overline{c^+} = (c_1^+, \dots, c_N^+)$ and $\overline{c^-} = (c_1^-, \dots, c_N^-)$ be two unknown constant vectors.

The extension of the C–V model to the vector case is

$$F(\overline{c^+}, \overline{c^-}, \phi) = \mu \cdot \text{Length}(C) + \int_{\text{inside}(C)} \frac{1}{N} \sum_{i=1}^N \lambda_i^+ |u_{0,i}(x, y) - c_i^+|^2 dx dy \\ + \int_{\text{outside}(C)} \frac{1}{N} \sum_{i=1}^N \lambda_i^- |u_{0,i}(x, y) - c_i^-|^2 dx dy,$$

where $\lambda_i^+ > 0$ and $\lambda_i^- > 0$ are parameters for each channel.

As in the scalar case, the model looks for the best vector-valued approximation taking only two values, the constant vectors $\overline{c^+}$ and $\overline{c^-}$. The active contour C is the boundary between these two regions. The energy balances the length of the contours in the image, with the fitting of u_0 , averaged over all channels. In this form, our model can detect edges present in at least one of the channels and not necessarily in all channels. We can associate this property with the syntax OR.

Rewriting it in level set form, we obtain

$$F(\overline{c^+}, \overline{c^-}, \phi) = \mu \int_{\Omega} \delta(\phi(x, y)) |\nabla \phi(x, y)| dx dy \\ + \int_{\Omega} \frac{1}{N} \sum_{i=1}^N \lambda_i^+ |u_{0,i}(x, y) - c_i^+|^2 H(\phi(x, y)) dx dy \\ + \int_{\Omega} \frac{1}{N} \sum_{i=1}^N \lambda_i^- |u_{0,i}(x, y) - c_i^-|^2 (1 - H(\phi(x, y))) dx dy,$$

for $i = 1, \dots, N$.

The parameters μ and $\overline{\lambda^{+,-}} = (\lambda_1^{+,-}, \dots, \lambda_N^{+,-})$ are integral to tuning the object detector sensitivity. μ is the weight for the length term, while the $\overline{\lambda}$ coefficients are the weights for the error term. Large μ or small $\overline{\lambda}$ are necessary for the model to filter high frequency noise. Likewise, larger coefficients $\overline{\lambda}$ are necessary to detect objects with fine detail. See Experimental Results for such examples.

Minimizing the energy with respect to the constants c_i^+, c_i^- , for $i = 1, \dots, N$, we obtain:

$$c_i^+ = \frac{\int_{\Omega} u_{0,i}(x, y) H(\phi(x, y)) dx dy}{\int_{\Omega} H(\phi(x, y)) dx dy} \quad (\text{average}(u_{0,i}) \text{ on } \phi \geq 0), \\ c_i^- = \frac{\int_{\Omega} u_{0,i}(x, y) (1 - H(\phi(x, y))) dx dy}{\int_{\Omega} H(\phi(x, y)) dx dy} \quad (\text{average}(u_{0,i}) \text{ on } \phi < 0).$$

We need now to calculate the Euler–Lagrange equation for ϕ . To do so, we will regularize the functions H and δ by C^1 approximations, denoting them by H_{ϵ} and δ_{ϵ} (as $\epsilon \rightarrow 0$), as in [1]. Two examples of such approximations are given by

$$H_{1,\epsilon}(z) = \begin{cases} 1 & \text{if } z > \epsilon, \\ 0 & \text{if } z < -\epsilon, \\ \frac{1}{2} \left[1 + \frac{z}{\epsilon} + \frac{1}{\pi} \sin\left(\frac{\pi z}{\epsilon}\right) \right] & \text{if } |z| \leq \epsilon, \end{cases} \quad \delta_{1,\epsilon}(z) = H'_{1,\epsilon}(z),$$

as proposed in [27], and by

$$H_{2,\varepsilon}(z) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan \left(\frac{z}{\varepsilon} \right) \right), \quad \delta_{2,\varepsilon}(z) = H'_{2,\varepsilon}(z),$$

as proposed in [1].

Then define F_ε as follows:

$$\begin{aligned} F_\varepsilon(\overline{c^+}, \overline{c^-}, \phi) &= \mu \varepsilon t_\Omega \delta_\varepsilon(\phi(x, y)) |\nabla \phi(x, y)| dx dy \\ &+ \varepsilon t_\Omega \frac{1}{N} \sum_{i=1}^N \lambda_i^+ |u_{0,i}(x, y) - c_i^+|^2 H_\varepsilon(\phi(x, y)) dx dy \\ &+ \varepsilon t_\Omega \frac{1}{N} \sum_{i=1}^N \lambda_i^- |u_{0,i}(x, y) - c_i^-|^2 (1 - H_\varepsilon(\phi(x, y))) dx dy. \end{aligned}$$

The regularized minimization problem that we will solve is

$$\varepsilon f_{\overline{c^+}, \overline{c^-}, \phi} F_\varepsilon(\overline{c^+}, \overline{c^-}, \phi).$$

Assuming that $\overline{c^+}$ and $\overline{c^-}$ are constant vectors, and minimizing F_ε with respect to ϕ , yields the following Euler–Lagrange equation for ϕ (parameterizing the descent direction by an artificial time),

$$\frac{\partial \phi}{\partial t} = \delta_\varepsilon \left[\mu \cdot \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) - \frac{1}{N} \sum_{i=1}^N \lambda_i^+ (u_{0,i} - c_i^+)^2 + \frac{1}{N} \sum_{i=1}^N \lambda_i^- (u_{0,i} - c_i^-)^2 \right]$$

in Ω , and with the boundary condition

$$\frac{\delta_\varepsilon(\phi)}{|\nabla \phi|} \frac{\partial \phi}{\partial \vec{n}} = 0$$

on $\partial\Omega$, where \vec{n} denotes the unit normal at the boundary of Ω .

To solve this evolution problem, we use a finite differences scheme, as suggested in [1].

2. EXPERIMENTAL RESULTS

We will use several artificial images and a satellite image to show the benefits of the C–V algorithm used in vector-valued form. The figures are shown as follows. First, a series of images is shown for each channel. Each series contains the original image as well as the progression of the contour. The density within the contour is defined by the norm of the inside averages (c_i^+) and the outside averages (c_i^-), given by $(1/\sqrt{N})\sqrt{(c_1)^2 + \dots + (c_N)^2}$, with corresponding + and – superscripts. In all our numerical results, the initial level set function is positive inside the initial curve (a circle or an ellipse) and negative outside the curve. As seen in the sample pictures, choosing the parameters requires some a priori information about the image. For example, the parameters must be adapted to differentiate between cases where a missing feature is due to an occlusion or an added feature is due to noise. Another example is the case where the noise of one channel is substantially greater than in the other channels.

In Fig. 2, both channels have the same triangle, but with a different piece missing. Each channel on its own is not sufficient to detect the complete triangle. In combination, the complete triangle is easily detected.

Objects are detected even when significant noise is present. When noise is added, see Fig. 3, the objects are detected without first performing noise reduction on the images. This is superior to other noise reduction methods which tend to blur edge details. Parameters were adapted to remove high frequency noise (for instance, taking a larger μ than for the previous example without noise).

In Fig. 4, the first example of occlusion is shown. The occlusion is present in both channels. Only the vector-valued model combining both channels can extract the complete information, including the missing information from each channel. The model has also the ability to automatically detect interior contours, as we can see in this example.

Another example can be found in multispectral images. In Fig. 5, we have an airplane imaged from mid-wave and long-wave infrared channels. One channel is very noisy making it very difficult to detect the edges of the entire airplane, while the other, less noisy, channel has a partial occlusion of the airplane. Each channel is insufficient for determination of the complete contour. However, in combination, most of the features are detected.

The vector-valued C–V model can also be used on color images. By dividing the image into red, green, and blue (RGB) channels, one can detect objects normally undetectable when the color image is transformed to a scalar intensity image. An example of this can be seen in Fig. 6. We can see the stop-light in the RGB image, while the scalar intensity image has the bottom object missing. Channel by channel detection would also be insufficient in this case, since features of the image are not complete in any single channel (see Fig. 7). Our model, however, detects all three features easily. Also note, in this particular example, the algorithm detects edges without gradient.

We end the experimental results with Fig. 8, representing a real RGB image. The algorithm detects blurred edges, in a noisy environment.

3. CONCLUSION

An algorithm was proposed to detect contours in vector-valued images that may or may not contain gradient edges. By using contour information from all the components of the image vector, a highly detailed contour can be found which is superior to contours generated from a single dimension of the image vector. This model has all the benefits of the C–V algorithm, including robustness even with noisy data and automatic detection of interior contours.

ACKNOWLEDGMENTS

We thank Larry Peterson and Gary Hewer at China Lake for providing us with the multispectral infrared data.

REFERENCES

1. T. Chan and L. Vese, Active contours without edges, *IEEE Trans. Image Process.* [UCLA CAM Report 98-53] and An active contour model without edges, in *Proceeding of the Second International Conference, Scale-Space '99. Corfu. Greece, September 26–27, 1999* (M. Nielsen, P. Johansen, O. F. Olsen, and J. Weickert, Eds.). Scale-Space Theories in Computer Vision, 141–151.

2. M. Kass, A. Witkin, and D. Terzopoulos, Snakes: Active contour models, *Internat. J. Comput. Vision* **1**, 1988, 321–331.
3. S. Osher and J. A. Sethian, Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulation, *J. Comput. Phys.* **79**, 1988, 12–49.
4. V. Caselles, F. Catté, T. Coll, and F. Dibos, A geometric model for active contours in image processing, *Numer. Math.* **66**, 1993, 1–31.
5. R. Malladi, J. A. Sethian, and B. C. Vemuri, A topology independent shape modeling scheme, in *Proc. SPIE Conf. on Geometric Methods in Computer Vision II, San Diego, July 1993*, pp. 246–258.
6. R. Malladi, J. A. Sethian, and B. C. Vemuri, Evolutionary fronts for topology-independent shape modeling and recovery, in *Proceedings of the Third European Conference on Computer Vision, Stockholm, Sweden, May 1994*, Lecture Notes in Computer Science, Vol. 800, pp. 3–13, Springer-Verlag, Berlin, 1994.
7. R. Malladi, J. A. Sethian, and B. C. Vemuri, Shape modeling with front propagation: A level set approach, *IEEE Trans. Pattern Anal. Mach. Intell.* **17**(2), 1995, 158–175.
8. V. Caselles, R. Kimmel, and G. Sapiro, Geodesic active contours, *Internat. J. Comput. Vision* **22**(1), 1997, 61–79.
9. S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi, Gradient flows and geometric active contour models, in *Proceedings of ICCV95, Cambridge, June 1995*, pp. 810–815.
10. G. Kanizsa, *La Grammaire du Voir. Essais sur la Perception*, Diderot Éditeur, Arts et Sciences, 1997.
11. G. Sapiro and D. L. Ringach, Anisotropic diffusion of multivalued images, Lecture Notes in Control and Information Sciences, Vol. 219, in *Proceedings of the 12th International Conference on Analysis and Optimization of Systems: Images, Wavelets and PDEs, Paris, June 26–28, 1996*, (M.-O. Berger, R. Deriche, I. Herlin, J. Jaffré, and J.-M. Morel, Eds.), Springer-Verlag, Berlin/New York, 134–140.
12. G. Sapiro and D. L. Ringach, Anisotropic diffusion of multivalued images with applications to color filtering, *IEEE Trans. Image Process.* **5**, 1996, 1582–1586.
13. P. Blomgren and T. Chan, Color TV: Total variation methods for restoration of vector valued images, *IEEE Trans. Image Process.* **7**(3), 1998, 304–309.
14. L. I. Rudin, S. Osher, and E. Fatemi, Nonlinear total variation based noise removal algorithms, *Phys. D* **60**, 1992, 259–268.
15. G. Sapiro, Color snakes, *Comput. Vision Image Understand.* **68**(2), 1997, 247–253.
16. G. Sapiro, Vector (self) snakes: A geometric framework for color, texture, and multiscale image segmentation, in *Proc. IEEE ICIP, Lausanne, September 1996*, Vol. I, pp. 817–820.
17. F. Dibos and G. Koepfler, Color segmentation using a variational formulation, in *Actes du 16me Colloque GRETSI, Grenoble, September 1997*, pp. 367–370.
18. G. Koepfler, C. Lopez, and J. M. Morel, A multiscale algorithm for image segmentation by variational method, *SIAM J. Numer. Anal.* **31**(1), 1994, 282–299.
19. G. Koepfler, C. Lopez, and L. Rudin, Data fusion by segmentation: Application to texture discrimination, in *Actes du 14me Colloque GRETSI, Juan-les-Pins, September 1993*, pp. 707–710.
20. J. Shah, Curve evolution and segmentation functionals: Application to color images, in *Proceedings of the IEEE ICIP 96, Lausanne, September 1996*, Vol. I, pp. 461–464.
21. S. C. Zhu, T. S. Lee, and A. L. Yuille, Region competition: Unifying snakes, region growing, energy/Bayes/MDL for multi-band image segmentation, in *Proceedings of the IEEE 5th ICCV, Cambridge, 1995*, pp. 416–423.
22. N. Paragios and R. Deriche, *Geodesic Active Regions for Texture Segmentation*, INRIA Research Report 3440, 1998.
23. M. Nitzberg, D. Mumford, and T. Shiota, Filtering, segmentation and depth, in *Lecture Notes in Computer Science*, Vol. 662, Springer-Verlag, Berlin, 1993.
24. S. Masnou and J. M. Morel, Level lines based disocclusion, in *Proceedings of DIP'97, Vienna, Austria, 1998*, SPIE, Vol. 3346.
25. Deleted in proofs.
26. D. Mumford and J. Shah, Optimal approximation by piecewise smooth functions and associated variational problems, *Comm. Pure Appl. Math.* **42**, 1989, 577–685.
27. H. K. Zhao, T. Chan, B. Merriman, and S. Osher, A variational level set approach to multiphase motion, *J. Comput. Phys.* **127**, 1996, 179–195.