PERGAMON

Pattern Recognition 36 (2003) 475–482

PATTERN RECOGNITION
THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

Growing snakes: active contours for complex topologies

Fernando A. Velasco[a], Jose L. Marroquin[b,*]

[a] *Universidad Michoacana de San Nicolás de Hidalgo, Planta baja Edif C, Ciudad Universitaria, Morelia, Mich., 58000, Mexico*
[b] *Centro de Investigación en Matemáticas, Apdo Postal 402, Guanajuato, Gto., 36000, Mexico*

Received 5 June 2001; received in revised form 6 December 2001; accepted 8 January 2002

Abstract

Snakes are active contours that minimize an energy function. In this paper we introduce a new kind of snakes, called growing snakes. These snakes are modeled as a set of particles connected by thin rods. Unlike the traditional snakes, growing snakes are automatically initialized. They start at the position where the gradient magnitude of an image is largest, and start to grow looking for zones of high gradient magnitude; simultaneously the associated energy function is minimized. Growing snakes can find contours with complex topology, describing holes, occlusions, separate objects and bifurcations. In a post-process the T-junctions are refined looking for the configuration with minimal energy. We also describe a technique that permits one to regularize the field of external forces that act on the Growing Snakes, which allow them to have good performance, even in the case of images with high levels of noise. Finally, we present results in synthetic and real images.

Keywords: Snakes; Active contours; Markov random fields; Image segmentation; Early vision

1. Introduction

The recognition of shapes and contours in a digital image is a big challenge in computational vision. There are many fields of science where the detection of contours and the segmentation of objects in images is an important task, such as in the case of pattern recognition, medical image analysis, visual techniques in strain analysis, etc. Active contours have been an important tool in the solution of the contour detection problem, but their automatic initialization is still an open problem, particularly in the cases of occlusions (T-junctions), bifurcations in the contour and in general in cases when the topology of the contour is not simple. In the current literature one can find two kinds of solutions: the parametric active contours called *snakes*, and the level set approach (e.g., geodesic curves). In 1987, Kass et al. [1] described the snake as a spline guided by external constraint forces and influenced by image forces. Thus, the curve is associated with an energy function, and the minimization of this energy permits one to find the best position of the parametric curve, which usually corresponds to places where the lines and edges in the image are. Snakes are thus parametric curves: $c(q):[0,1] \to \mathbb{R}^2$, on an image $I:[0,a][0,b] \to \mathbb{R}$. The energy of the snake is defined as

$$E(c) = \alpha \int_0^1 |c'(q)|^2 \, \mathrm{d}q + \beta \int_0^1 |c''(q)|^2 \, \mathrm{d}q - \lambda \int_0^1 |\nabla I(c(q))| \, \mathrm{d}q. \quad (1)$$

The first and second terms of Eq. (1) correspond to the internal energy and penalize the length and curvature of the snake, respectively. The third term, where ∇ is the gradient operator, corresponds to the external energy and attracts the curve towards the intensity edges that are present in the image. For the implementation of parametric snakes the curve is viewed as a particle system, and the energy function takes

* Corresponding author.
E-mail addresses: fvelasco@bernoulli.eic.umich.mx (F.A. Velasco), jlm@cimat.mx (J.L. Marroquin).

0031-3203/02/$22.00
PII: S0031-3203(02)00045-6

the form

$$E(p) = \alpha \sum_i |p_i - p_{i-1}|^2 + \beta \sum_i | - p_{i-1} + 2p_i - p_{i+1}|^2 - \lambda \sum_i |\nabla I(p_i)|, \quad (2)$$

where p_i is a 2-vector whose components are the coordinates of the position of the ith particle. Here, the first term penalizes the distance between two contiguous particles, and corresponds to the energy of a spring (with equilibrium length equal to zero) that connects them. The second term penalizes the curvature and corresponds to the energy of a thin rod that connects three contiguous particles. The last term is the gradient magnitude of the image evaluated at the position of the ith particle. The balloon model introduced by Cohen [2,3], uses an additional inflating force, which permits the user to give only an approximate initial position in the interior of the closed contour one is looking for. The inflating force then drives the snake to the desired place, and helps to escape of local minima in case of noise. Although this works well in many cases, it still has limitations: the parameter that controls the inflating force is not easy to set, particularly for high levels of noise; besides, this method can find only simple closed curves. In an attempt to handle more complex situations, McInerney et al. [4] proposed the topologically adaptable snakes, which provide for automatic initialization of the snakes, and permit, in principle, to find contours with complicated topologies. This method, however, often fails in the presence of T-junctions (occlusions) and when the image presents high levels of noise. Other proposals for improving the performance of parametric snakes involve the use of several coupled snakes [5–7], or the regularization of the external force field [8], but they do not address the problem of automatic initialization for complex curve topologies. A different approach, which can deal with automatic topology changes in the contours, is the level set approach, where the active contour is specified as the set of points where a suitably defined surface has a given value (usually zero) [9–11]. This method permits one to find an unspecified number of objects, which might contain holes, in an automatic way; however, these methods may fail when the high gradient zone that attracts the contour has gaps. Also, they cannot automatically find complex contour configurations that include T-junctions. These are partially considered in the work of Tek et al. [12], but these junctions are smoothed out by the method, so that the final configurations are usually not satisfactory (see Section 3). Our work presents a new parametric active contour model called growing snakes. These snakes solve many of the open problems described above. Specifically, they permit an automatic initialization of a set of snakes that may correspond to contours with complex topology (e.g., multiple objects, holes and T-junctions generated by occlusions). They are described in the next section.

2. Growing snakes

Growing snakes are parametric active contours that permit the description of complex topologies such as non-connected contours or bifurcations with a completely automatic initialization. They are based on the following idea: to start the process, "snake seeds" are generated at places of high gradient magnitude. Then, growth and shape smoothing steps are alternated until a final stable configuration is reached. At this point, new seeds are generated, with the constraint that no seed may appear close to the existing snake. At the end of the complete process, the T-junctions are revised to find the minimal energy configuration for the complete snake. We now describe each step in detail.

2.1. Growth step

The algorithm for the growth step is based on the following considerations:

(1) The *seeds* of the snake are generated automatically in those places where the gradient magnitude of the image is largest. Each seed has two particles and its orientation is perpendicular to the image gradient direction in the corresponding point.
(2) The initial (free) particles move in opposite directions following the high-gradient zones.
(3) When the free particles are advanced, the system automatically includes intermediate particles.
(4) When a free particle meets another particle, they both connect, forming either a larger curve or a T-junction.

We now describe the procedure in more detail.

2.1.1. Seed generation

To generate the seeds, first the M image is obtained as the gradient magnitude of the original image I_0; since image differentiation is known to be ill-posed [13], we regularize it by convolving I_0 with the gradient of a Gaussian kernel G_σ with a fixed standard deviation σ (i.e., M is obtained from the Gaussian derivatives of I_0):

$$M = |\nabla G_\sigma * I_0|.$$

The system now looks for the points where the highest value of the M image is attained, excluding those points that are closer than a fixed distance from an existing snake particle (at the beginning there are no existing snake particles) and places a seed for the growing snake at each one of these points (i.e., at the set of points where M is maximal, excluding the inhibitory neighborhood of existing snakes). Each seed consists of two particles, placed at a distance d_0, so that the direction of the line that joins them is perpendicular to the direction of the gradient of M, and the midpoint coincides with the seed point (we use a value of $d_0 = 0.4$ pixels).

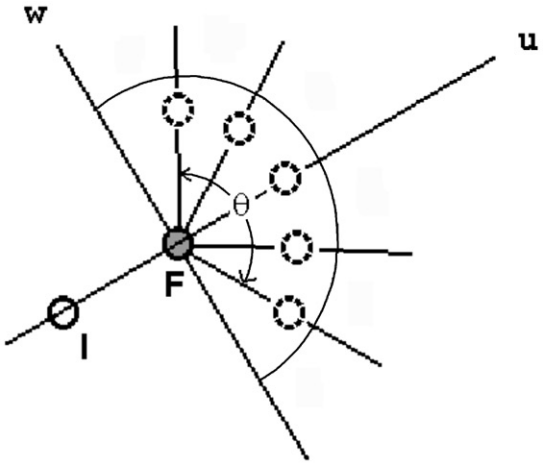

Fig. 1. Candidate new positions for an F particle (see text).

2.1.2. The growth process

The two particles at each seed point are called free (F) particles. The growth system works as follows: For each F particle:

(1) Compute the direction of the line that joins this F particle with its neighboring particle in the snake (note that an F particle has only one neighboring particle); we call this line the u line.
(2) Find a fixed number N of candidate new positions for the F particle. These positions are equispaced along a circular sector of radius h_p and angular width θ, centered at the F particle and oriented along the u line (see Fig. 1).
(3) The new position for the F particle is selected as the candidate position with highest M value. If more than one candidate position has the same value, the system selects the position that is closer to the u line.
(4) The F particle is advanced to the new position, and a new intermediate (I) particle is placed in its previous position.

During the growth process it is possible that two F particles meet (i.e., that their distance becomes smaller than a threshold δ_1). If this happens, the F particles connect and change their status to I particles, so that a longer snake is obtained. If the F particle meets an I particle (i.e., if the distance between the F and I particles is smaller than a threshold δ_2), the system generates a *T-junction* and the F particle changes its status to a T particle; this new kind of particle does not advance even if it is the extreme particle of a snake. An F particle also changes to T if it meets a border of the image. Note that more than one T particle may be joined to a single I particle, so that junctions of order higher than 3 (e.g., "X" junctions) are appropriately handled (see Section 3).

If an F particle meets a T particle, both particles connect and change their status to I, as in the first case.

2.2. Shape smoothing step

In this step, the snake changes its shape under the influence of internal and external forces, as in the case of traditional snakes. These forces are derived from a potential energy of the form

$$E(p) = \beta \sum_i | - p_{i-1} + 2p_i - p_{i+1}|^2 - \lambda \sum_i M(p_i). \quad (3)$$

The first term of Eq. (3) is the internal energy and includes only the thin rod term that controls the curvature of the snake. Note that T particles are connected by thin rods only to previous particles of the snake. The second term represents the external energy and M is computed as explained above. When the position p_i has non-integer values, bilinear interpolation is used to evaluate this term. For the minimization of the energy function, we use the following method [14]: the dynamics of the system is obtained from Newton's second law, considering that the force exerted over each particle is equal to minus the partial derivative of the potential energy with respect to its position. Considering that each particle has unit mass, the equation of motion is

$$\frac{\partial^2 p_i}{\partial t^2} = -\frac{\partial E(p)}{\partial p_i} \doteq v_i(p),$$

where t denotes time and $E(p)$ is given by Eq. (3). To avoid oscillations, one may add a linear friction term with positive coefficient $2k$ to obtain

$$\frac{\partial^2 p_i}{\partial t^2} - v_i(p) + 2k \frac{\partial p_i}{\partial t} = 0. \quad (4)$$

The local minima of Eq. (3) are fixed points of this system. Its practical implementation requires the time discretization: for a given time increment h and for every site i one obtains the Taylor expansions

$$p_i(t+h) = p_i(t) + h \frac{\partial p_i(t)}{\partial t} + \frac{h^2}{2} \frac{\partial^2 p_i(t)}{\partial t^2} + O(h^3), \quad (5)$$

$$p_i(t-h) = p_i(t) - h \frac{\partial p_i(t)}{\partial t} + \frac{h^2}{2} \frac{\partial^2 p_i(t)}{\partial t^2} - O(h^3). \quad (6)$$

Subtracting these equations one gets

$$\frac{\partial p_i(t)}{\partial t} \approx \frac{1}{2h} (p_i(t+h) - p_i(t-h)). \quad (7)$$

Adding Eqs. (5) and (6) one obtains

$$\frac{\partial^2 p_i(t)}{\partial t^2} \approx \frac{1}{h^2} (p_i(t+h) + p_i(t-h) - 2p_i(t)). \quad (8)$$

Substituting Eqs. (7) and (8) in Eq. (4) we find

$$p_i(t+h) = A p_i(t) + B p_i(t-h) + C v_i(p(t)) \quad (9)$$

with $A = 2/(kh+1)$; $B = (kh-1)/(kh+1)$; $C = h^2/(kh+1)$. For $k = 1/h$ the discrete dynamical system (9) corresponds

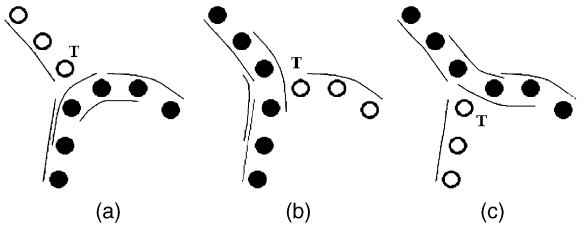

Fig. 2. The 3 possible configurations for a T-junction.

to the gradient descent method with fixed step size. For $k < 1/h$ the system has inertia and moves at higher speed. In our case we use $k = 0.5/h$ and $h = 0.001$.

The shape smoothing step thus consists in applying a fixed number of iterations of system (9) before a new growth step takes place.

2.2.1. The eras and the inhibitory system

When all the F particles are changed either to I or to T particles we say that an "era" has finished; at this point, the system automatically starts looking for new seed positions. All the particles that are already active have an inhibitory neighborhood, which consists of a circular area around each particle where a seed cannot be born. The size of this inhibitory area is fixed and given by the user. The union of these inhibitory neighborhoods is called the inhibitory area of the snake; so, at the start of a new era, the system has to find at least one position out of the inhibitory area of the active snakes where a growing snake can appear. In this way, one prevents the appearance of new snakes parallel and close to the existing ones. User intervention in this method is limited to indicating at the end of which era should the system be stopped, or alternatively, setting a threshold on the maximum value of M (outside the inhibitory area) below which no new seeds may appear. For a particular class of images, this threshold may be set in advance, making the full procedure completely automatic.

2.2.2. Refining the T-junctions

Once the system has finished the last era, it automatically examines all the T-junctions (i.e., places where a T particle is connected to an I particle) and verifies that they adopt the best of their three possible configurations (see Fig. 2). To do this, for each T-junction and each possible configuration, the system is iterated 50 times and the final total energy is computed using Eq. (3). In practice, it is not necessary to update the position of all particles, but only of those that are "close" to the examined T-junction (we update 10 particles along each branch of the T, so that this procedure takes about 0.03 s per T-junction on a 1.2 GHz Pentium 4 machine). Then the configuration associated with the minimal energy is selected for that T-junction. With the final configuration in all the T-junctions, the system is iterated again, to find the final position of the snakes.

2.3. Regularization of the external forces

The external forces that act on each particle in the shape smoothing step are given by the gradient of M evaluated at the particle position. In the ideal case, these forces should be pointing in the direction of the closest edge in the image; in practical cases, due to noise in the image, these forces will not point in the appropriate direction, and must, therefore, be regularized. Note that M itself corresponds to the gradient magnitude of the original image, which was computed using Gaussian derivatives; the σ parameter of the derivative kernels, however, should be kept small, so that the edge locations remain precise; if the noise level is high, this implicit prefiltering is not sufficient to regularize the *gradient* of M, which defines the force field. Xu et al. [8] proposed the use of a classical regularization scheme to smooth out the external force field. In their implementation, the smoothed gradient field f is the minimizer of the cost function:

$$U(f) = \sum_r |\nabla M|^2 \cdot |f(r) - \nabla M(r)|^2 + \lambda \sum_{\langle r,s \rangle} |f(r) - f(s)|^2,$$

where the first sum is taken over all pixels r, and the second one is taken over all nearest neighbor pixel pairs (r, s). If the amount of noise in the image is small, one may select a small value for the λ parameter, and obtain a good regularized field. In noisy situations, however, one must select a relatively large value for λ, which causes the f field to be averaged across the ridges of M (i.e., the edges of the image). This, however is not appropriate, because in the vicinity of a ridge, f points in opposite directions, and averaging these will result in a field which is tangential, instead of perpendicular to the ridge. To avoid this undesired effect, we propose here a different regularization method, based on Bayesian estimation with a prior discrete Markov random field (MRF) model [15].

2.3.1. MRF-based regularization of the external forces

Our goal is to compute a regularized field of directions d for the external forces. To do this, we first discretize the interval $[0, 2\pi)$, where each element may take values, into an appropriate number of intervals (we use 16), so that $d(r) \in \{0, 2\pi/16, \ldots, 15 \cdot 2\pi/16\}$. We take as observations g, the gradient direction of M, i.e., $g(r) = \tan^{-1}(M_y(r)/M_x(r))$, where $M_x(r)$ and $M_y(r)$ denote the partial derivatives of M with respect to x and y, respectively, at pixel r. Since the data also takes values in $[0, 2\pi)$, an appropriate noise model is given by the Von Mises distribution [16], so that the likelihood of the data, given that $d(r) = q_k$ is

$$\Pr(g(r)|d(r) = q_k) = \hat{p}_k(r) = \frac{1}{Z} \exp[\gamma \cos(g(r) - q_k)], \quad (10)$$

where γ is a positive variance parameter and Z is a normalizing constant, chosen so that each $\hat{p}(r)$ is a valid discrete probability distribution, i.e., so that $\sum_k \hat{p}_k(r) = 1$ for every r (note that in our case each $\hat{p}(r)$ is a 16-vector).

The prior distribution should reflect the constraint that the field d should be piecewise constant. To this end, we use a first-order MRF model with Ising potentials [15], so that the prior distribution is Gibbsian:

$$P_d(d) = \frac{1}{Z} \exp\left[-\beta \sum_{\langle r,s \rangle} V(d(r), d(s))\right], \quad (11)$$

where β is a positive parameter and the Ising potentials V are given by

$$V(d(r), d(s)) = \begin{cases} -1 & \text{if } d(r) = d(s), \\ 1 & \text{if } d(r) \neq d(s). \end{cases}$$

Applying Bayes rule, one gets the posterior distribution

$$P_{d|g}(d) = \frac{1}{Z} \exp[-U(d)], \quad (12)$$

where the "energy" U is given by

$$U(d) = -\gamma \sum_r \cos(g(r) - d(r)) + \beta \sum_{\langle r,s \rangle} V(d(r), d(s)).$$

Since the d field is discrete-valued, the optimal estimator may be found by minimizing the expected number of errors, where the expectation is taken with respect to the posterior distribution. It is possible to show that this estimator may be found by finding the modes of the posterior marginal distributions: this is the maximizer of the posterior marginals or MPM estimator (see Refs. [15,17,18] for details). These marginal distributions are defined as:

$$\pi_k(r) = \sum_{d:d(r)=q_k} P_{d|g}|(d),$$

where the sum is taken over all possible discrete direction fields d that have direction q_k at site r. Note that the marginals cannot be directly computed, because the sum in the above equation has too many terms. It is possible to show, however, that a good approximation for these marginals is obtained by a spatial smoothing of the likelihoods, i.e., by the minimization, with respect to a field p of discrete probability distributions, of the function:

$$U_p(p) = \sum_r |p(r) - \hat{p}(r)|^2 + \lambda \sum_{\langle r,s \rangle} |p(r) - p(s)|^2, \quad (13)$$

where each $p(r)$ is a 16-vector, $\hat{p}$ is given by Eq. (10) and λ is a positive parameter (see Ref. [18] for details). Once the p field that minimizes Eq. (13) is found, the optimal estimator for the d field at pixel r is found as the mode of the distribution $p(r)$, i.e., by setting: $d(r) = k_{\max}(r)$, where $k_{\max}(r)$ is such that $p_{k_{\max}(r)}(r) > p_k(r)$ for all $k \neq k_{\max}(r)$.

The resulting field d has now the desired properties. This is because, as one moves across a ridge of M, one does not average the gradient directions, but rather, find the mode of their (marginal) distributions. As one approaches the ridge, these distributions become bimodal (with a distance of π radians between the 2 peaks), and as one passes over the ridge, the location of the largest peak shifts from one position to the other, producing a jump in the regularized d field as desired.To compare this procedure with the one proposed by Xu [8], we performed a series of experiments with a synthetic image of a vertical step edge of height 255, corrupted with additive, white, 0-mean Gaussian noise (with $\sigma = 150$). As a measure of performance, we took the average absolute angular error (AAE) of the regularized gradient direction, with respect to the "ideal" field, in which the gradient vector is normal to the edge and points towards it. This error was averaged over a tubular neighborhood of the edge, 4 pixels wide. The M (gradient magnitude) image was obtained from the Gaussian derivatives of the image, using a Gaussian kernel with $\sigma = 1$. Fig. 3 shows a plot of the AAE vs. the regularization parameter λ for both Xu's and our method. In both cases there is an optimal value for λ ($\lambda = 0.5$ for Xu's, and $\lambda = 5$ for ours). As one can see, however, the error curve is flatter in our case, indicating that the method is less sensitive to the precise setting of this parameter. Also, the error corresponding to the optimal λ is about 20% less in our case. A detail of the regularized field, for different values of λ is shown in Fig. 4.

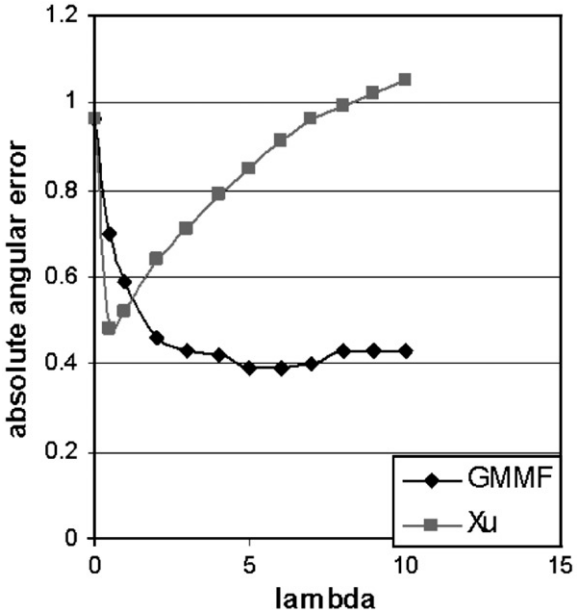

Fig. 3. Mean absolute angular error vs. regularization parameter for Xu's and GMMF methods (see text).

We have found that making the external forces equal to the d field (with unit magnitude everywhere) improves significantly the performance of the growing snakes (and also of the classical ones). This is the procedure we use in the experiments reported below (we used $\gamma = 1$ in Eq. (10) and $\lambda = 10$ in Eq. (13) in all cases). As shown in Ref. [18], the

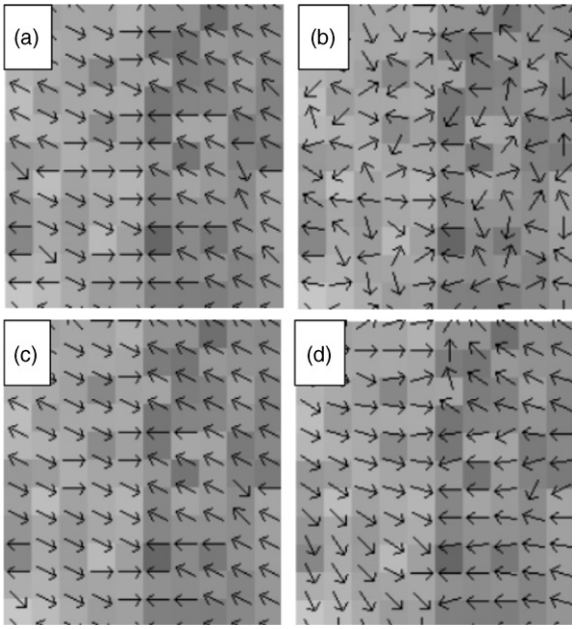

Fig. 4. Detail of the regularized force field obtained by Xu's and GMMF methods for different values of λ: (a) GMMF, $\lambda = 5$. (b) Xu's, $\lambda = 0.5$. (c) GMMF, $\lambda = 10$. (d) Xu's, $\lambda = 2$.

minimization of Eq. (13) may be very efficiently accomplished by solving 16 decoupled linear systems using a fast implementation of the cosine transform. This computation, for a 256×256 image, takes about 2.3 s on a Pentium IV PC running at 1.2 MHz. The complete procedure, including the computation of the regularized force field, the growth and shape smoothing iterations and the refinement of the T-junctions takes about 7 s on the same machine (the precise time depends on the particular image).

The procedure has a number of parameters that must be adjusted. In most cases, however, the method's performance is quite insensitive to their precise value, so that once a good value is found, it may be used in most cases. In *all* the experiments reported here, we have used the following values:

Size of the inhibitory area for new seeds to appear: 7 pixels.

Threshold that defines the meeting of 2 F particles: $\delta_1 = 3$.

Threshold that defines the meeting of F and I particles: $\delta_2 = 1.2$.

Growth step size $h_p = 0.5$.

Angular width for searching candidate growth positions: $\theta = 0.66\pi$.

Number of candidate positions: $N = 5$.

Number of shape-smoothing iterations between growth steps: 5.

Weighting factor for the external force (Eq. (3)): $\lambda = 1$.

Width of the Gaussian kernel for prefiltering the image: $\sigma = 1.0$.

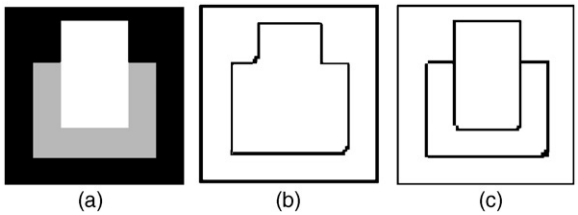

Fig. 5. (a) Synthetic image portraying an occlusion: (b) Solution found by the method of McInerney et al. (c) Solution found by our method.

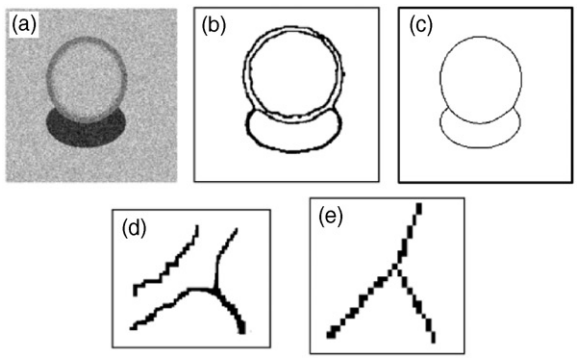

Fig. 6. (a) Synthetic image of a sphere. (b) Contours found by the method of Tek et al. (c) Contours found by our method. (d) Detail of (b) near the right T-junction. (e) Corresponding detail for (c).

Parameter that controls the rigidity of the snake (Eq. (3)) $\beta = 50$.

3. Experiments

In this section, we present some examples that illustrate the performance of the growing snakes (GS) method. Fig. 5 compares the performance of GS with the topologically adaptable snakes (TAS) proposed by McInerney et al. [4] on a synthetic image where the contour topology includes T-junctions (occlusions): panel (b) shows the solution obtained by the TAS's; as one can see, the internal contour is not found. Panel (c) shows the solution with GS. Note that all the significant contours were found and the T-junctions were correctly solved. Fig. 6 shows a synthetic image of a sphere with a cast shadow, corrupted by Gaussian noise; panel (b) shows the solution found by the method of Tek et al. [12]: note that a spurious internal border is found inside the sphere and the contour deformation close to the T-junctions. Panel (c) shows our solution; note the accuracy with which the T-junctions are solved; panels (d) and (e) show in detail the contour configurations at the T-junctions found by both methods.

The next figures illustrate the performance of GSs in some real images. Fig. 7a presents one object occluded by another,

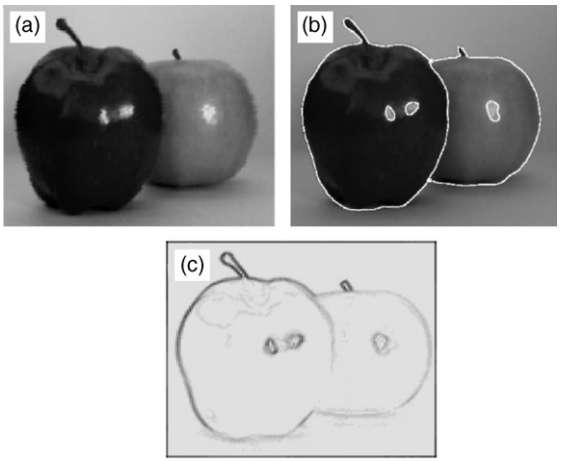

Fig. 7. (a) Real image of 2 apples. (b) Contours found by growing snakes (4 eras) superimposed as white pixels over the original image. (c) Gradient magnitude (coded by gray level) of the image of panel (a).

and Fig. 7b, the contours found by GSs in a completely automatic way. Fig. 7c shows the gradient magnitude of the image of panel (a). Note how the snake grows even across low-gradient zones (lower right of the foreground apple). Note also that there is a single closed contour for the silhouette of the forefront object, whereas the contour of the background object ends up with two T-junctions, as it should.

Fig. 8 presents an example of an image of occluding objects with a high noise level. Note how in this case too the gradient magnitude (panel (b)) presents significant gaps, which does not preclude the GS from finding the correct topology. Once this configuration (shown in panel (c)) is found, if smoother contours are desired, one may iterate the system with a higher value for the parameter β to obtain the final configuration of panel (d).

Fig. 9 presents the contour found by the system in the case of an image with non-Gaussian noise, in this case, the speckle pattern interferogram produced when a mechanical part is illumined by coherent (laser) light. Finally, we present in Fig. 10, the contours found by the system in the case of a complex scene. Note that in this case there are junctions of order higher than 3 (e.g., an X-junction to the left of the center of the image), which are appropriately handled by the system.

4. Conclusions

We have presented a new kind of active contours called "growing snakes". They allow one to find contours with complex topology, so that contours of: occluded and multiple objects; objects with holes, and cases where there are gaps (low gradient zones) in the contour information, are found in a completely automatic way. The pre-processing step that regularizes the external forces that act on the snakes, based on Bayesian estimation with MRF models, is very fast

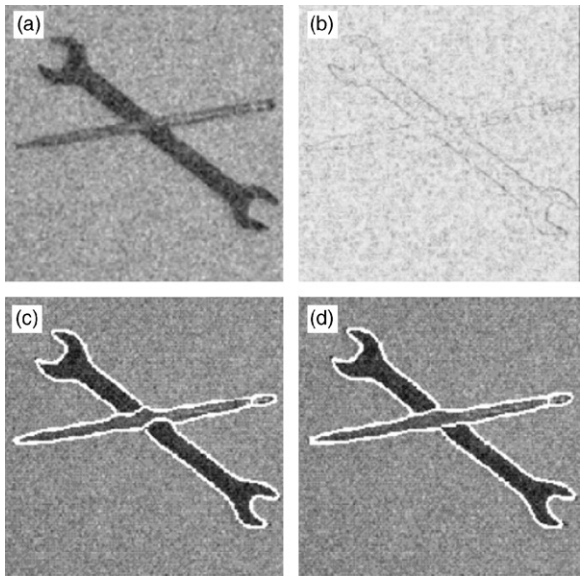

Fig. 8. (a) Noisy image of two objects. (b) Gradient magnitude (coded by gray level) of the image of panel (a). (c) Contours found by the growing snakes (with $\beta = 50$), superimposed as white pixels over the original image. (d) Final configuration, after the system is allowed to iterate with $\beta = 70$.

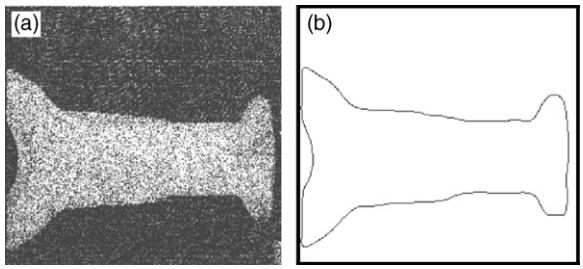

Fig. 9. (a) Laser speckle interferogram of a mechanical part. (b) Contour found by the growing snakes.

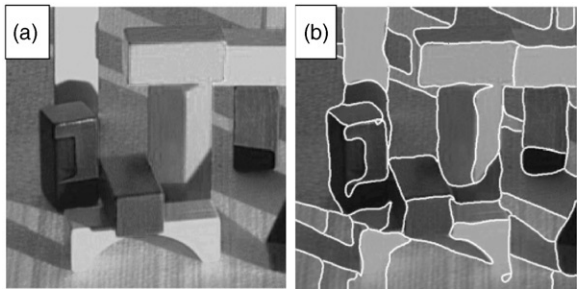

Fig. 10. (a) Wooden blocks scene. (b) Contours found by the growing snakes superimposed to the image (a).

and robust, and permits the snakes to have a good behavior even in images with high levels of noise. We have tried our method in real images with different kinds of noise (Gaussian, shot and speckle), and found a good performance in all cases. The image M that drives the snakes corresponds to the gradient magnitude of the original image convolved with a Gaussian kernel. This convolution may smooth out some fine detail; to recover it one may take the final configuration found by the GSs as the initial configuration of a classical snake system, and allow this system to reach equilibrium, using to generate the external forces a new image M obtained as the gradient magnitude of the original image convolved with a finer kernel (i.e., a Gaussian kernel with smaller σ). Note that this new system is not allowed to grow, since one only wants to refine the final snake positions without altering their topology.

Acknowledgements

The authors wish to thank the anonymous reviewer for his/her comments, which helped to improve the quality of the presentation. The authors were supported in part by Grants I35682-A and 34575-A from the Consejo Nacional de Ciencia y Tecnología, México.

References

[1] M. Kass, A. Witkin, D. Terzopoulos, Snakes, active contours models, Int. J. Comput. Vision 1 (4) (1987) 321–331.
[2] L.D. Cohen, On active contour models and balloons, computer vision, graphics and image processing, Image Understanding 53 (2) (1991) 211–218.
[3] L.D. Cohen, I. Cohen, Finite-element methods for active contour models and balloons for 2-D and 3-D images, IEEE Trans. Pattern Anal. Mach. Intell. 15 (11) (1993) 1131–1147.
[4] T. McInerney, D. Terzopoulos, Topologically adaptable snakes, Proc. ICCV 95 (1995) 840–845.
[5] S.R. Gunn, M.S. Nixon, Improving snake performance via a dual active contour, Lecture Notes in Computer Science, Vol. 970, Springer, Berlin, Heidelberg, 1995, pp. 600–605.
[6] S.R. Gunn, M.S. Nixon, A robust snake implementation: a dual active contour, IEEE Trans. Pattern Anal. Mach. Intell. 19 (1) (1997) 63–68.
[7] F. Velasco, J.L. Marroquin, Sandwich snakes: robust active contours, SPIE Conference in Applications of Digital Image Processing XXI, Vol. 3460, 1998, pp. 226–235.
[8] C. Xu, J.L. Prince, Gradient vector flow: a new external force for snakes, Proc. Comp. Vis. and Pat. Rec. (CVPR 95) IEEE Conference 1997, pp. 66–71.
[9] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, Proc. ICCV 95 (1995) 694–699.
[10] R. Malladi, J.A. Sethian, B.C. Vemuri, Shape modeling with front propagation: a level set approach, IEEE Trans. Pattern Anal. Mach. Intell. 17 (2) (1995) 158–175.
[11] J.A. Sethian, Numerical algorithms for propagating interfaces: Hamilton–Jacobi equations and conservation laws, J. Differential Geom. 31 (1990) 131–161.
[12] H. Tek, B.B. Kimia, Image segmentation by reaction–diffusion bubbles, Proc. Comp. Vis. and Pat. Rec. (CVPR 95) IEEE Conference 1995, pp. 156–162.
[13] T. Poggio, V. Torre, Ill-Posed problems and regularization analysis in early vision, A.I. Memo 773, Massachusetts Institute of Technology, Cambridge, MA, 1984.
[14] J.L. Marroquin, Deterministic interactive particle models for image processing and computer graphics, Graphical Models Image Process. 55 (5) (1993) 408–417.
[15] J. Marroquin, S. Mitter, T. Poggio, Probabilistic solution of ill-posed problems in computational vision, J. Am. Stat. Assoc. 82 (1987) 76–89.
[16] N.I. Fisher, Statistical Analysis of Circular Data, Cambridge University Press, Cambridge, 1993.
[17] J.L. Marroquin, S. Botello, F. Calderon, B.C. Vemuri, The MPM-MAP algorithm for image segmentation, Proceedings of the 15th International Conference in Pattern Recognition (ICPR 2000) Barcelona, Spain, 2000, pp. 303–308.
[18] J.L. Marroquin, F. Velasco, M. Rivera, M. Nakamura, Gauss-Markov measure field models for low-level vision, IEEE Trans. on Pattern Anal. Mach. Intell. 23 (4) (2001) 337–348.

About the Author—FERNANDO A. VELASCO received the B.S. degree in Civil Engineering from the Universidad Michoacana de San Nicolás de Hidalgo, México, in 1984; the M.S. degree in Computer Science from Instituto Tecnológico de Toluca in 1995; and the Ph.D. degree from the Center for Research in Mathematics, Guanajuato, Mexico in 1999. Dr. Velasco is Associate Professor and head of Mathematics Department of Civil Engineering School of the Universidad Michoacana and a member of the Mexican society for computer science (SMCC). His research interests include Computer Vision and Image Processing.

About the Author—JOSE L. MARROQUIN received the B.S. degree in Chemical Engineering in 1968 from the National University of Mexico, and the M.Sc. and Ph.D. degrees in Systems Science in 1985 from the Massachusetts Institute of Technology. He has worked for PEMEX, the Mexican petroleum company, as project leader in geophysical data processing. Currently, he is the head of the Computer Science Department at the Center for Research in Mathematics, Guanajuato, Mexico, and is conducting research related to the computer processing of visual information. Dr. Marroquin is a Fellow of the National Research System of the Mexican Government and of the Mexican Academy of Science.