



A higher-order active contour model of a 'gas of circles' and its application to tree crown extraction

P. Horváth^{a,b}, I.H. Jermyn^{a,*}, Z. Kato^b, J. Zerubia^a

^aAriana (joint INRIA/IS3 research group), INRIA, B.P. 93, 06902 Sophia Antipolis, France

^bImage Processing and Computer Graphics Department, P.O. Box 652, University of Szeged, 6701 Szeged, Hungary

ARTICLE INFO

Article history:

Received 22 November 2006

Received in revised form 3 September 2008

Accepted 5 September 2008

Keywords:

Tree
Crown
Extraction
Aerial image
Higher-order
Active contour
Gas of circles
Prior
Shape

ABSTRACT

We present a model of a 'gas of circles': regions in the image domain composed of a unknown number of circles of approximately the same radius. The model has applications to medical, biological, nanotechnology, and remote sensing imaging. The model is constructed using higher-order active contours (HOACs) in order to include non-trivial prior knowledge about region shape without constraining topology. The main theoretical contribution is an analysis of the local minima of the HOAC energy that allows us to guarantee stable circles, fix one of the model parameters, and constrain the rest. We apply the model to tree crown extraction from aerial images of plantations. Numerical experiments both confirm the theoretical analysis and show the empirical importance of the prior shape information.

© 2008 Published by Elsevier Ltd.

1. Introduction

Forestry is a domain in which image processing and computer vision techniques can have a significant impact. Resource management and conservation require information about the current state of a forest or plantation. Much of this information can be summarized in statistics related to the size and placement of individual tree crowns (e.g. mean crown area and diameter, density of the trees). Currently, this information is gathered using expensive field surveys and time-consuming semi-automatic procedures, with the result that partial information from a number of chosen sites frequently has to be extrapolated. An image processing method capable of automatically extracting tree crowns from high resolution aerial or satellite images and computing statistics based on the results would greatly aid this domain.

The tree crown extraction problem can be viewed as a special case of a general image understanding problem: the identification of the region R in the image domain Ω corresponding to some entity or

entities in the scene. In order to solve this problem in any particular case, we have to construct, even if only implicitly, a probability distribution on the space of regions $P(R|I, K)$. This distribution depends on the current image data I and on any prior knowledge K we may have about the region or about its relation to the image data, as encoded in the likelihood $P(I|R, K)$ and the prior $P(R|K)$ appearing in the Bayes' decomposition of $P(R|I, K)$ (or equivalently in their energies $-\ln P(I|R, K)$ and $-\ln P(R|K)$). This probability distribution can then be used to make estimates of the region we are looking for.

In the automatic solution of realistic problems, the prior knowledge K , and in particular prior knowledge about the 'shape' of the region, as described by $P(R|K)$, is critical. The tree crown extraction problem provides a good example: particularly in plantations, R takes the form of a collection of approximately circular connected components of similar size. There is thus a great deal of prior knowledge about the region sought. The question is then how to incorporate such prior knowledge into a model for R . If the model does not include enough prior knowledge, it will be necessary for the user to provide it.

The simplest prior information concerns the smoothness of the region boundary. For example, the Ising model and many active contour models [1–3] use a combination of region boundary length and region area as their prior energies, but curvature can be used too [1]. Such models are integrals over the region boundary of a function of various derivatives of the boundary. In consequence, they

* Corresponding author. Tel.: +33 49238 7683; fax: +33 49238 7643.

E-mail addresses: Peter.Horvath@sophia.inria.fr (P. Horváth), Ian.Jermyn@sophia.inria.fr (I.H. Jermyn), kato@inf.u-szeged.hu (Z. Kato), Josiane.Zerubia@sophia.inria.fr (J. Zerubia)

URLs: <http://www-sop.inria.fr/ariana> (P. Horváth), <http://www-sop.inria.fr/ariana> (I.H. Jermyn), <http://www.inf.u-szeged.hu/~kato/> (Z. Kato), <http://www-sop.inria.fr/ariana> (J. Zerubia).

capture local differential geometric information, corresponding to local interactions between boundary points, but can say nothing more global about the shape of the region. To go further, one must introduce longer range interactions. There are two principal ways to do this: one is to introduce hidden variables, given which the original variables of interest are (more or less) independent. Marginalizing over the hidden variables then introduces interactions between the original variables. Another is to include explicit long-range interactions between the original variables.

The first approach has been much investigated, in the form of template shapes and their deformations. Here a probability distribution or an energy is defined based on a distance measure of some kind between regions. One region, the template, is fixed, while the other is the variable R . Template regions may be learned from examples [4–9] or fixed by hand [10]; similarly the distance function maybe based, for example, on the learned covariance of a Gaussian distribution [5–9], or chosen *a priori* [4,10,11]. The most sophisticated methods use the kernel trick to define the distance as a pull-back from a high-dimensional space, thereby allowing more complex behaviours [12]. Multiple templates may also be used, corresponding to a mixture model [12,13].

These methods assign high probability to regions ‘close’ to certain points in the space of regions. The set of regions with high probability is thus in some sense bounded. As such, it is difficult to construct models of this type that favour regions for which the topology, and in particular the number of connected components, is unknown *a priori*, because the set of regions in this case is unbounded, and cannot be described as variations around one or more templates. There are many problems, however, for which the topology is unknown *a priori*, for example, the extraction of networks, or the extraction of an unknown number of objects of a particular type from astronomical, biological, medical, or remote sensing images. For this type of prior knowledge, a different type of model is needed. Higher-order active contours (HOACs) are one such category of models.

HOACs [14] take the second approach mentioned above. They introduce explicit long-range interactions between region boundary points via energies that contain multiple integrals over the boundary, thus avoiding the use of template shapes. HOAC energies can be made intrinsically Euclidean invariant, and, as required by the above analysis, incorporate sophisticated prior information about region shape without necessarily constraining region topology. As with other methods incorporating significant prior knowledge, it is not necessary to introduce extra knowledge via an initialization close to the target region: a generic initialization suffices, thus rendering the method quasi-automatic. Rochery et al. [14] applied the method to road extraction from satellite and aerial images using a prior which favours network-like objects.

In this paper, we describe a HOAC model of a ‘gas of circles’: the model favours regions composed of an *a priori* unknown number of circles of a certain radius. For such a model to work, the circles must be stable to small perturbations of their boundaries, i.e. they must be local minima of the HOAC energy, for otherwise a circle would tend to ‘decay’ into other shapes. The main theoretical contribution of this paper is an analysis of the stability of local minima of the HOAC energy that allows us to ensure that circles of a given radius are stable. In addition, it allows us to fix one of the model parameters in terms of the others, and to constrain the rest. This type of calculation has wide applicability to other active contour models and to other shapes. For example, it shows that no stable circle is possible using a classical active contour model containing only boundary length and interior area terms. The calculation proceeds by performing a functional Taylor expansion of the HOAC energy around a circle (or more generally, any shape), and then demanding that the first order term be zero for all perturbations, and that the second order term be positive semi-definite. Gradient descent experiments using

the HOAC energy, with parameters fixed using the stability calculations, produce stable circles of the expected radii, thereby demonstrating empirically the coherence between the stability calculations and the numerical computations used in practice to minimize the energy.

The model has many potential applications, to medical, biological, physical, and remote sensing imagery in which the entities to be identified are circular. We choose to apply it to the problem of extracting tree crowns from aerial imagery, using the ‘gas of circles’ model as a prior energy, and an appropriate likelihood. We will see that the extra prior knowledge included in the ‘gas of circles’ model permits the separation of trees that cannot be separated by simpler methods, such as maximum likelihood or classical active contours. We focus on images of plantations and orchards, for which the model is well adapted. The case of general forests is much harder, and will be left for future work.

In the next section, we present a brief introduction to HOACs. In Section 3, we describe the ‘gas of circles’ HOAC model, the stability analysis, and the results of geometric experiments. In Section 4, we apply the new model to tree crown extraction. We describe a likelihood energy for trees, and then present experimental results on synthetic data and on aerial images. We conclude in Section 5, and discuss some open issues with the model.

2. Higher-order active contours

HOAC models, like all active contour models, represent a region R by its boundary, ∂R , a closed 1-chain γ in the image domain Ω ([15] is a useful reference for the following discussion). Although region boundaries correspond to a special subset of closed 1-chains known as domains of integration, active contour energies themselves are defined for general 1-chains. It is convenient to use this more general context to distinguish HOAC energies from classical active contours, because it allows for notions of linearity to be used to characterize the complexity of energy functionals.

Using this representation, HOAC energies can be defined as follows [14]. Let γ be a 1-chain in Ω , and $\text{dom}\gamma$ be its domain. Then $\gamma^n : (\text{dom}\gamma)^n \rightarrow \Omega^n$ is an n -chain in Ω^n . We define a class of $(n-p)$ -forms on Ω^n that are 1-forms with respect to $(n-p)$ factors and 0-forms with respect to the remaining p factors (by symmetry, it does not matter which p factors). These forms can be pulled back to $(\text{dom}\gamma)^n$ by γ^n . The Hodge duals of the p 0-form factors with respect to the induced metric on $\text{dom}\gamma$ can then be taken independently on each such factor, thus converting them to 1-forms, and rendering the whole form an n -form on $(\text{dom}\gamma)^n$. This n -form can then be integrated on $(\text{dom}\gamma)^n$.

In the $(n,p) = (n,0)$ cases, we are simply integrating a general n -form on the image of γ^n in Ω^n , thus defining a linear functional on the space of n -chains in Ω^n , and hence an n th-order monomial on the space of 1-chains in Ω . Taking arbitrary linear combinations of such monomials then gives the space of polynomial functionals on the space of 1-chains. By analogy we refer to the general (n,p) cases as ‘generalized n th-order monomials’ on the space of 1-chains in Ω , and to arbitrary linear combinations of the latter as ‘generalized polynomial functionals’ on the space of 1-chains in Ω . HOAC energies are generalized polynomial functionals. Standard active contour energies are generalized *linear* functionals on 1-chains in this sense, hence the term ‘higher-order’.

The $(1,1)$ case is simply the boundary length in some metric. The $(1,0)$ case gives the region area in some metric. An interesting application of the $(2,2)$ case to topology preservation is described by Sundaramoorthi [16]. We specialize to the $(2,0)$ case. Let F be a 2-form on Ω^n . Using the antisymmetry of F together with the

symmetry of γ^2 , we can write the energy functional in this case as

$$E(\gamma) = \int_{(\partial R)^2} F = \int_{(\text{dom } \gamma)^2} (\gamma \times \gamma)^* F \\ = \iint_{(\text{dom } \gamma)^2} dt dt' \tau(t) \cdot F(\gamma(t), \gamma(t')) \cdot \tau(t'), \quad (2.1)$$

where $F(x, x')$, for each $(x, x') \in \Omega^2$, is a 2×2 matrix, t is a coordinate on $\text{dom } \gamma$, and $\tau = \dot{\gamma}$ is the tangent vector to γ .

By imposing Euclidean invariance on this term, and adding linear terms, Rochery et al. [14] defined the following higher-order active contour prior:

$$E_g(\gamma) = \lambda_C L(\gamma) + \alpha_C A(\gamma) - \frac{\beta_C}{2} \iint dt dt' \tau(t) \cdot \tau(t') \Phi(R(t, t')), \quad (2.2)$$

where L is the boundary length functional, A is the interior area functional and $R(t, t') = |\gamma(t) - \gamma(t')|$ is the Euclidean distance between $\gamma(t)$ and $\gamma(t')$. Rochery et al. [14] used the following interaction function Φ :

$$\Phi(z) = \begin{cases} \frac{1}{2} \left(1 - \frac{z-d}{\varepsilon} - \frac{1}{\pi} \sin \frac{\pi(z-d)}{\varepsilon} \right) & |z-d| < \varepsilon, \\ H(d-z) & \text{else.} \end{cases} \quad (2.3)$$

In this paper, we use this same interaction function with $d = \varepsilon$, but other monotonically decreasing functions lead to qualitatively similar results.

3. The ‘gas of circles’ model

For certain ranges of the parameters involved, the energy in equation (2.2) favours regions in the form of networks, consisting of long narrow arms with approximately parallel sides, joined together at junctions, as described by Rochery et al. [14]. It thus provides a good prior for network extraction from images. This behaviour does not persist for all parameter values, however, and we will exploit this parameter dependence to create a model for a ‘gas of circles’, an energy that favours regions composed of an *a priori* unknown number of circles of a certain radius.

For this to work, a circle of the given radius must be stable, that is, it must be a local minimum of the energy. In Section 3.1, we show that stable circles are indeed possible provided certain constraints are placed on the parameters. More specifically, we expand the energy E_g in a functional Taylor series to second order around a circle of radius r_0 . The constraint that the circle be an energy extremum then requires that the first order term be zero, while the constraint that it be a minimum requires that the operator in the second order term be positive semi-definite. These requirements constrain the parameter values. In Section 3.2, we present numerical experiments using E_g that confirm the results of this analysis.

3.1. Stability analysis

We denote a member of the equivalence class of maps representing the 1-chain defining the circle by γ_0 , and a small perturbation by $\delta\gamma$. To second order,

$$E_g(\gamma) = E_g(\gamma_0 + \delta\gamma) \simeq E_g(\gamma_0) + \left\langle \delta\gamma \left| \frac{\delta E_g}{\delta \gamma} \right| \right\rangle_{\gamma_0} + \frac{1}{2} \left\langle \delta\gamma \left| \frac{\delta^2 E_g}{\delta \gamma^2} \right| \delta\gamma \right\rangle_{\gamma_0}, \quad (3.1)$$

where $\langle \cdot | \cdot \rangle$ is a metric on the space of 1-chains.

Since γ_0 represents a circle, it is easiest to express it in terms of polar coordinates r, θ on Ω . For a suitable choice of coordinate on S^1 , a circle of radius r_0 centred on the origin is then given by $\gamma_0(t) = (r_0(t), \theta_0(t))$, where $r_0(t) = r_0$, $\theta_0(t) = t$, and $t \in [-\pi, \pi]$. We

are interested in the behaviour of small perturbations $\delta\gamma = (\delta r, \delta\theta)$. Because the energy E_g is defined on 1-chains, tangential changes in γ do not affect its value. We can therefore set $\delta\theta = 0$, and concentrate on δr .

On the circle, using the arc length parameterization t , the integrands of the different terms in E_g are functions of $t - t'$ only; they are invariant to translations around the circle. In consequence, the second derivative $\delta^2 E_g / \delta\gamma(t) \delta\gamma(t')$ is also translation invariant, and this implies that it can be diagonalized in the Fourier basis of the tangent space at γ_0 . It is thus easiest to perform the calculation by expressing δr in terms of this basis: $\delta r(t) = \sum_k a_k e^{ir_0 k t}$, where $k \in \{m/r_0 : m \in \mathbb{Z}\}$. Below, we simply state the resulting expansions to second order in the a_k for the three terms appearing in Eq. (2.2). Details can be found in Appendix A.

The boundary length and interior area of the region are given to second order by

$$L(\gamma) = \int_{-\pi}^{\pi} dt |\tau(t)| \simeq 2\pi r_0 \left\{ 1 + \frac{a_0}{r_0} + \frac{1}{2} \sum_k k^2 |a_k|^2 \right\} \quad (3.2)$$

$$A(\gamma) = \int_{-\pi}^{\pi} d\theta \int_0^{r(\theta)} dr' r' \simeq \pi r_0^2 + 2\pi r_0 a_0 + \pi \sum_k |a_k|^2. \quad (3.3)$$

Note that there are no stable solutions using these terms alone. For the circle to be an extremum, we require $\lambda_C 2\pi + \alpha_C 2\pi r_0 = 0$, which tells us that $\alpha_C = -\lambda_C / r_0$. The criterion for a minimum is, for each k , $\lambda_C r_0 k^2 + \alpha_C \geq 0$. We must have $\lambda_C > 0$ for stability at high frequencies. Substituting for α_C , the condition becomes $\lambda_C (r_0 k^2 - r_0^{-1}) \geq 0$. Substituting $k = m/r_0$, gives the condition $m^2 - 1 \geq 0$: the zero frequency perturbation is never stable.

The quadratic term can be expressed to second order as

$$\iint_{-\pi}^{\pi} dt dt' G(t, t') = 2\pi \int_{-\pi}^{\pi} dp F_{00}(p) + 4\pi a_0 \int_{-\pi}^{\pi} dp F_{10}(p) \\ + \sum_k 2\pi |a_k|^2 \left\{ \left[2 \int_{-\pi}^{\pi} dp F_{20}(p) \right. \right. \\ \left. \left. + \int_{-\pi}^{\pi} dp e^{-ir_0 k p} F_{21}(p) \right] \right. \\ \left. - \left[2ir_0 k \int_{-\pi}^{\pi} dp e^{-ir_0 k p} F_{23}(p) \right] \right. \\ \left. + \left[r_0^2 k^2 \int_{-\pi}^{\pi} dp e^{-ir_0 k p} F_{24}(p) \right] \right\}, \quad (3.4)$$

where $G(t', t) = \tau(t') \cdot \tau(t) \Phi(R(t, t'))$. The F_{ij} are functionals of Φ (and hence of d), and functions of r_0 , as well as of p .

Combining Eqs. (3.2)–(3.4), we find, up to second order:

$$E_g(\gamma_0 + \delta\gamma) \simeq e_0(r_0) + a_0 e_1(r_0) + \frac{1}{2} \sum_k |a_k|^2 e_2(k, r_0), \quad (3.5)$$

where

$$e_0(r_0) = 2\pi \lambda_C r_0 + \pi \alpha_C r_0^2 - \pi \beta_C G_{00}(r_0) \\ e_1(r_0) = 2\pi \lambda_C + 2\pi \alpha_C r_0 - 2\pi \beta_C G_{10}(r_0) \\ e_2(k, r_0) = 2\pi \lambda_C r_0 k^2 + 2\pi \alpha_C - 2\pi \beta_C [2G_{20}(r_0) + G_{21}(k, r_0) \\ - 2ir_0 k G_{23}(k, r_0) + r_0^2 k^2 G_{24}(k, r_0)],$$

where $G_{ij} = \int_{-\pi}^{\pi} dp e^{-ir_0(1-\delta(j))kp} F_{ij}(p)$. Note that there are no off-diagonal terms linking a_k and $a_{k'}$ for $k \neq k'$: the Fourier basis diagonalizes the second order term.

3.1.1. Parameter constraints

Note that a circle of any radius is always an extremum for non-zero frequency perturbations (a_k for $k \neq 0$), as these Fourier coefficients do not appear in the first order term (this is also a consequence

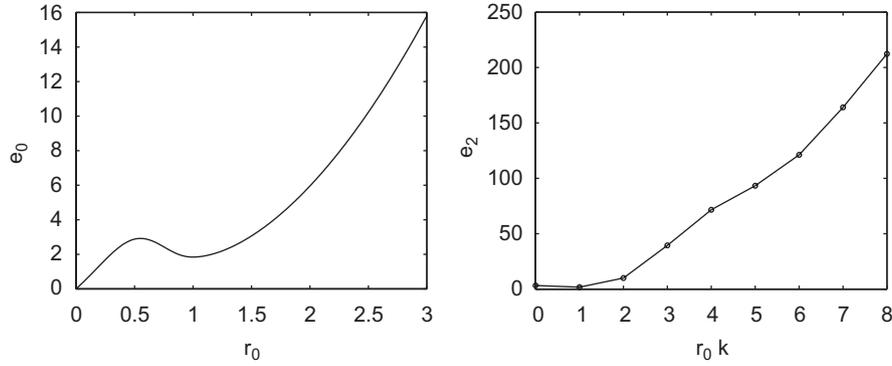


Fig. 1. Plots of e_0 against r_0 and e_2 against $\hat{r}_0 k$. Left: the energy of a circle e_0 plotted against radius r_0 for $\lambda_C = 1.0$, $\alpha = 0.8$, and $\beta_C = 1.39$ calculated from Eq. (3.6) with $\hat{r}_0 = 1.0$. (The parameters of Φ are $d = 1.0$ and $\varepsilon = 1.0$, but note that it is not necessary in general that $d = \hat{r}_0$.) The function has a minimum at $r_0 = \hat{r}_0$ as desired. Right: the second derivative of E_g , e_2 , plotted against $\hat{r}_0 k$ for the same parameter values. The function is non-negative for all frequencies.

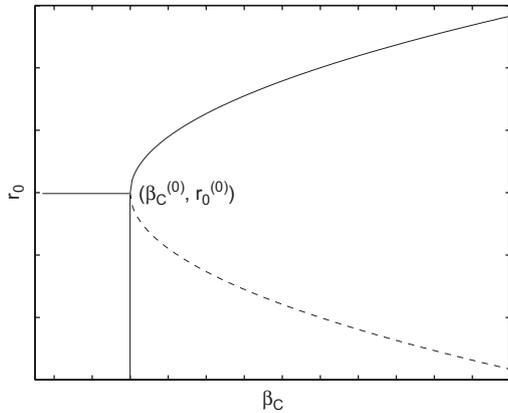


Fig. 2. Schematic plot of the positions of the extrema of the energy of a circle versus β_C .

of invariance to translations around the circle). The condition that a circle be an extremum for a_0 as well ($e_1 = 0$) gives rise to a relation between the parameters:

$$\beta_C(\lambda_C, \alpha_C, \hat{r}_0) = \frac{\lambda_C + \alpha_C \hat{r}_0}{G_{10}(\hat{r}_0)}, \tag{3.6}$$

where we have introduced \hat{r}_0 to indicate the radius at which there is an extremum, to distinguish it from r_0 , the radius of the circle about which we are calculating the expansion (3.1). The left-hand side of Fig. 1 shows a typical plot of the energy e_0 of a circle versus its radius r_0 , with the β_C parameter fixed using Eq. (3.6) with $\lambda_C = 1.0$, $\alpha = 0.8$, and $\hat{r}_0 = 1.0$. The energy has a minimum at $r_0 = \hat{r}_0$ as desired. The relationship between \hat{r}_0 and β_C is not quite as straightforward as it might seem though. As can be seen, the energy also has a maximum at some radius. It is not *a priori* clear whether it will be the maximum or the minimum that appears at \hat{r}_0 . If we graph the positions of the extrema of the energy of a circle against β_C for fixed α_C , we find a curve qualitatively similar to that shown in Fig. 2 (this is an example of a fold catastrophe). The solid curve represents the minimum, the dashed the maximum. Note that there is indeed a unique β_C for a given choice of \hat{r}_0 . Denote the point at the bottom of the curve by $(\beta_C^{(0)}, \hat{r}_0^{(0)})$. Note that at $\beta_C = \beta_C^{(0)}$, the extrema merge and for $\beta_C < \beta_C^{(0)}$, there are no extrema: the energy curve is monotonic because the quadratic term is not strong enough to overcome the shrinking effect of the length and area terms. Note also that the minimum cannot move below $r_0 = r_0^{(0)}$. This behaviour is easily understood qualitatively in terms of the interaction function

in Eq. (2.3). If $2r_0 < d - \varepsilon$, the quadratic term will be constant, and no force will exist to stabilize the circle. In order to use Eq. (3.6) then, we have to ensure that we are on the upper branch of Fig. 2.

Eq. (3.6) gives the value of β_C that provides an extremum of e_0 with respect to changes of radius a_0 at a given \hat{r}_0 ($e_1(\hat{r}_0) = 0$), but we still need to check that the circle of radius \hat{r}_0 is indeed stable to perturbations with non-zero frequency, i.e. that $e_2(k, \hat{r}_0)$ is non-negative for all k . Scaling arguments mean that in fact the sign of e_2 depends only on the combinations $\tilde{r}_0 = r_0/d$ and $\tilde{\alpha}_C = (d/\lambda_C)\alpha_C$. The equation for e_2 can then be used to obtain bounds on $\tilde{\alpha}_C$ in terms of \tilde{r}_0 . (Details of these calculations and bounds can be found in [17].) The right-hand side of Fig. 1 shows a plot of $e_2(k, \hat{r}_0)$ against $\hat{r}_0 k$ for the same parameter values used for the left-hand side, showing that it is non-negative for all $\hat{r}_0 k$.

We call the resulting model, the energy E_g with parameters chosen according to the above criteria, the ‘gas of circles’ model.

3.2. Geometric experiments

To illustrate the behaviour of the ‘gas of circles’ model, in this section we show the results of some experiments using E_g (there are no image terms). Fig. 3 shows the result of gradient descent using E_g starting from various different initial regions. (For details of the implementation of gradient descent for higher-order active contour energies using level set methods, see [14].) In the first column, four different initial regions are shown. The other three columns show the final regions, at convergence, for three different sets of parameters. In particular, the three columns have $\hat{r}_0 = 15.0$, 10.0 , and 5.0 , respectively.

In the first row, the initial shape is a circle of radius 32 pixels. The stable states, which can be seen in the other three columns, are circles with the desired radii in every case. In the second row, the initial region is composed of four circles of different radii. Depending on the value of \hat{r}_0 , some of these circles shrink and disappear. This behaviour can be explained by looking at Fig. 1. As already noted, the energy of a circle e_0 has a maximum at some radius r_{max} . If an initial circle has a radius less than r_{max} , it will ‘slide down the energy slope’ towards $r_0 = 0$, and disappear. If its radius is larger than r_{max} , it will finish in the minimum, with radius \hat{r}_0 . This is precisely what is observed in this second experiment. In the third row, the initial condition is composed of four squares. The squares evolve to circles of the appropriate radii. The fourth row has an initial condition composed of four differing shapes. The nature of the stable states depends on the relation between the stable radius, \hat{r}_0 , and the size of the initial shapes. If \hat{r}_0 is much smaller than an initial shape, this shape will ‘decay’ into several circles of radius \hat{r}_0 .

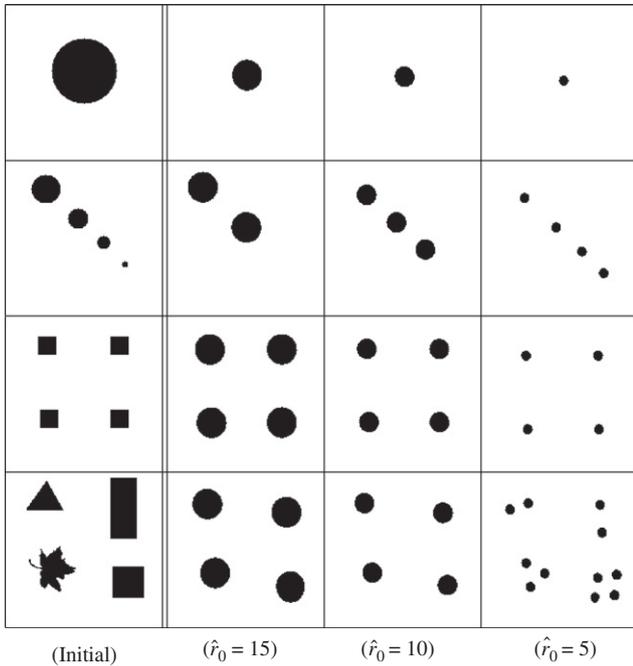


Fig. 3. Experimental results using the geometric term: the first column shows the initial conditions; the other columns show the stable states for various choices of the radius.

4. Likelihood energy and experiments

In this section, we apply the ‘gas of circles’ model to the extraction of trees from aerial images. We give a brief state of the art for tree crown extraction, and then present the likelihood energy we use in Section 4.2. In Section 4.3, we describe tree crown extraction experiments on aerial images and compare the results to those found using a classical active contour model. In Section 4.4, we examine the robustness of the method to noise using synthetic images. This illuminates the principal failure modes of the model, which will be further discussed in Section 5, and which point the way for future work. In Section 4.5, we illustrate the importance of prior information via tree crown separation experiments on synthetic images, and compare the results to those obtained using a classical active contour model.

4.1. Previous work

The problem of locating, counting, or delineating individual trees in high resolution aerial images has been studied in a number of papers. For example, Gougeon [18] observes that trees are brighter than the areas separating them. Local minima of the image are found using a 3×3 filter, and the ‘valleys’ connecting them are then found using a 5×5 filter. The tree crowns are subsequently delineated using a five-level rule-based method designed to find circular shapes, but with some small variations permitted. While the method is quite effective in separating trees, the size of the filters results in significant overestimation of the size of the trees. Larsen [19] concentrates on spruce tree detection using a template matching method. The 3D shape of the tree is modelled using a generalized ellipsoid, while illumination is modelled using the position of the sun and a clear-sky model. Template matching is used to calculate a correlation measure between the tree image predicted by the model and the image data. The local maxima of this measure are treated as tree candidates, and various strategies are then used to eliminate false positives. This method provides 3D information about the trees, but requires

specific models for each species of tree, as well as knowledge of a number of extraneous parameters, for example, illumination. Brandtberg and Walter [20] decompose an image into multiple scales, and then define tree crown boundary candidates at each scale as zero crossings with convex grey-scale curvature. Edge segment centres of curvature are then used to construct a candidate tree crown region at each scale. These are then combined over different scales and a final tree crown region is grown.

The above methods use a series of *ad hoc* steps rather than a single unified model, which makes identifying the assumptions behind the methods difficult. Closer in spirit to the present work is that of [21], which models the collection of tree crowns by a marked point process, where the marks are circles or ellipses. An energy is defined that penalizes, for example, overlapping shapes, and controls the parameters of the individual shapes. Compared to the work described in this paper, the method has the advantage that overlapping trees can be represented as two separate objects, but the disadvantage that the tree crowns are not precisely delineated due to the small number of degrees of freedom for each mark.

4.2. Likelihood energy and gradient descent

In order to couple the region model E_g to image data, we need a likelihood, $P(I|R, K)$. The images we use for the experiments are coloured infrared (CIR) images. Originally they are composed of three bands, corresponding roughly to green, red, and near infrared (NIR). Analysis of the one-point statistics of the image in the region corresponding to trees and the image in the background, shows that the ‘colour’ information does not add a great deal of discriminating power compared to a ‘greyscale’ combination of the three bands, or indeed the NIR band on its own. We therefore model the latter.

The image resolution is ~ 0.5 m/pixel, and tree crowns have diameters of the order of 10 pixels. Little dependence remains between the pixels at this resolution, which means, when combined with the paucity of statistics within each tree crown, that pixel dependencies (i.e. texture) are very hard to use for modelling purposes. We therefore model the interior of tree crowns using a Gaussian distribution with mean μ and covariance $\sigma^2 \delta_R$, where δ_A is the identity operator on images on $A \subset \Omega$.

The background is very varied, and thus hard to model in a precise way. We use a Gaussian distribution with mean $\bar{\mu}$ and variance $\bar{\sigma}^2 \delta_{\bar{R}}$. In general, $\mu > \bar{\mu}$, and $\sigma < \bar{\sigma}$; trees are brighter and more constant in intensity than the background. The boundary of each tree crown has significant inward-pointing image gradient, and although the Gaussian models should in principle take care of this, we have found in practice that it is useful to add a gradient term to the likelihood energy. Our likelihood thus has three factors:

$$P(I|R, K) = Z^{-1} g_R(I_R) g_{\bar{R}}(I_{\bar{R}}) f_{\partial R}(I_{\partial R}),$$

where I_R and $I_{\bar{R}}$ are the images restricted to R and \bar{R} , respectively, and g_R and $g_{\bar{R}}$ are proportional to the Gaussian distributions already described, i.e.

$$-\ln g_R(I_R) = \int_R d^2x \frac{1}{2\sigma^2} (I_R(x) - \mu)^2 \tag{4.1}$$

and similarly for $g_{\bar{R}}$. The function $f_{\partial R}$ depends on the gradient of the image ∂I on the boundary ∂R :

$$-\ln f_{\partial R}(I_{\partial R}) = \lambda_i \int_{\text{dom } \gamma} dt n(t) \cdot \partial I(t), \tag{4.2}$$

where n is the unnormalized outward normal to γ . The normalization constant Z is thus a function of μ , σ , $\bar{\mu}$, $\bar{\sigma}$, and λ_i . Z is also a functional

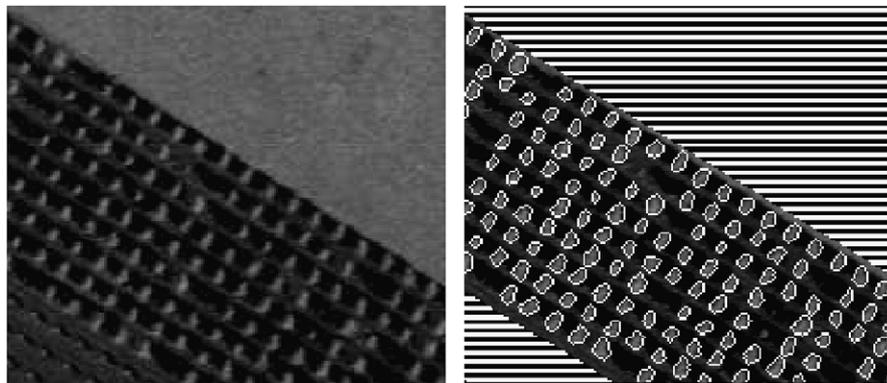


Fig. 4. Left: real image with a planted forest ©IFN (0.3, 0.06, 0.05, 0.05). Right: the result obtained using the 'gas of circles' model (529, 5.88, 5.88, 5.64, 4, 4).

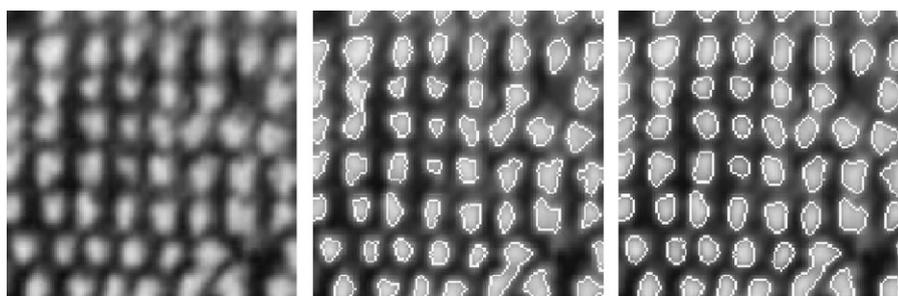


Fig. 5. From left to right: image of poplars ©IFN (0.73, 0.11, 0.23, 0.094); the best result with a classical active contour (880, 13, 73); result with the 'gas of circles' model (100, 6.7, 39, 31, 4.2, 4.2).

of the region R . To a first approximation, it is a linear combination of $L(\partial R)$ and $A(R)$. It thus has the effect of changing the parameters λ_C and α_C in E_g . However, since these parameters are essentially fixed by hand (the criteria described in Section 3.1.1 only allow us to fix β_C and constrain α_C), knowledge of the normalization constant does not change their values, and we ignore it once the likelihood parameters have been learnt.

The full model is then given by $E(R) = E_i(I, R) + E_g(R)$, where

$$E_i(I, R) = -\ln g_R(I_R) - \ln g_{\bar{R}}(I_{\bar{R}}) - \ln f_{\partial R}(I_{\partial R}).$$

The energy is minimized by gradient descent. The functional derivatives of all terms except the quadratic term in E_g are standard. The functional derivative of the quadratic term gives rise to a gradient descent force given by

$$\hat{n} \cdot \frac{\partial \gamma}{\partial \tau}(t) = \beta \int_{\text{dom } \gamma} dt' \hat{R}(t, t') \cdot n(t') \dot{\Phi}(R(t, t')), \quad (4.3)$$

where $\hat{R}(t, t') = (\gamma(t) - \gamma(t')) / |\gamma(t) - \gamma(t')|$. To evolve the region we use the level set framework of [22] extended to the demands of nonlocal forces such as Eq. (4.3) [14].

The computational complexity of the algorithm is unknown without a bound on the number of iterations. However, the complexity of one iteration is easily analysed. Eq. (4.3) involves an integration over the contour for each contour point. The worst case complexity is thus $O(L^2)$, where L is the length of the contour. The implementation however, only integrates over those points within interaction range (i.e. $d + \varepsilon$), and so the complexity depends on the average length l of contour within interaction range of a point, becoming $O(Ll)$. Typically l is a local quantity that does not depend on the size of the image. In our application, L , on the other hand, is proportional to the number of trees, which is in turn proportional to the size of the image, n . So the complexity of one iteration is $O(n)$.

4.3. Tree crown extraction from aerial images

In this section, we present the results of the application of the above model to 50 cm/pixel colour infrared aerial images of poplar stands located in the 'Saône et Loire' region in France. The images were provided by the French National Forest Inventory (IFN). As stated in Section 4.2, we model only the NIR band of these images, as adding the other two bands does not increase discriminating power. The tree crowns in the images are ~ 8 – 10 pixels in diameter, i.e. ~ 4 – 5 m.

In the experiments, we compare our model to a classical active contour model ($\beta_C = 0$). The parameters μ , σ , $\bar{\mu}$, and $\bar{\sigma}$ were the same for both models, and were learned from hand-labelled examples in advance. The classical active contour prior model thus has three free parameters (λ_i , λ_C and α_C), while the 'gas of circles' model has six (λ_i , λ_C , α_C , β_C , d and r_0). We fixed r_0 based on our prior knowledge of tree crown size in the images, and d was then set equal to r_0 . Once α_C and λ_C have been fixed, β_C is determined by Eq. (3.6). There are thus three effective parameters for the HOAC model. In the absence of any method to learn λ_i , α_C and λ_C , they were fixed by hand to give the best results, as with most applications of active contour models. The values of λ_i , λ_C and α_C were not the same for the classical active contour and HOAC models; they were chosen to give the best possible result for each model separately. The initial region in all experiments was a rounded rectangle slightly bigger than the image domain. The image values in the region exterior to the image domain were set to $\bar{\mu}$ to ensure that the region would shrink inwards.

Fig. 4 illustrates the first experiment. On the left are the data, showing a regularly planted poplar stand. The result is shown on the right. We applied the algorithm to the central part of the image only, for reasons that will be explained in Section 5.

Fig. 5 illustrates a second experiment. On the left are the data. The image shows a small piece of an irregularly planted poplar forest.

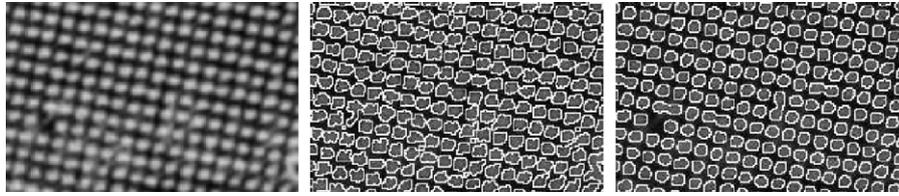


Fig. 6. From left to right: image of poplars ©IFN (0.71, 0.075, 0.18, 0.075); the best result with a classical active contour (24000, 100, 500); result with the 'gas of circles' model (1500, 25, 130, 100, 3.5, 3.5).

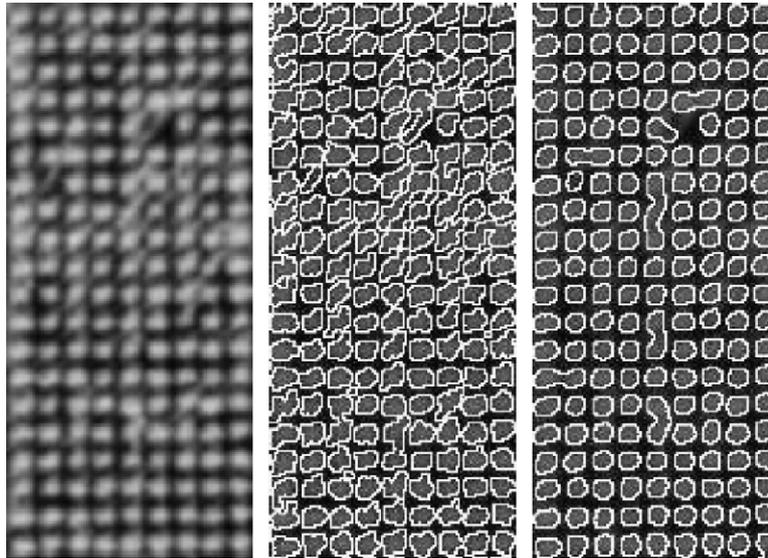


Fig. 7. From left to right: image of poplars ©IFN (0.71, 0.075, 0.18, 0.075); the best result with a classical active contour (35000, 100, 500); result with the 'gas of circles' model (1200, 20, 100, 82, 3.5, 3.5).

The image is difficult because the intensities of the crowns are varied and the gradients are blurred. In the middle is the best result we could obtain using a classical active contour. On the right is the result we obtain with the 'gas of circles' model.¹ Note that in the classical active contour result several trees that are in reality separate are merged into single connected components, and the shapes of trees are often rather distorted, whereas the prior geometric knowledge included when $\beta \neq 0$ allows the separation of almost all the trees and the regularization of their shapes.

Fig. 6 illustrates a third experiment. Again the data is on the left, the best result obtained with a classical active contour model is on the middle, and the result with the 'gas of circles' model is on the right. The trees are closer together than in the previous experiment. Using the classical active contour, the result is that the tree crown boundaries touch in the majority of cases, despite their separation in the image. Many of the connected components are malformed due to background features. The HOAC model produces more clearly delineated tree crowns, but there are still some joined trees. We will discuss this further in Section 5.

Fig. 7 shows a fourth experiment. Again the data is on the left, the best result obtained with a classical active contour model is in the

¹ Unless otherwise specified, in the figure captions the values of the parameters learned from the image are shown when the data is mentioned, in the form $(\mu, \sigma, \tilde{\mu}, \tilde{\sigma})$. The other parameter values are shown when each result is mentioned, in the form $(\lambda_i, \lambda_C, \alpha_C, \beta_C, d, r_0)$, truncated if the parameters are not present. All parameter values are truncated to two significant figures. Unless otherwise specified, images were scaled to take values in $[0, 1]$. The region boundary is shown in white.

Table 1

Results on real images using a classical active contour model (left) and the 'gas of circles' model (right)

Figure	CD %	FP %	FN %	CD %	FP %	FN %
Fig. 5	85	0	15	97	0	3
Fig. 6	96.2	2.8	1.9	97.7	0	2.3
Fig. 7	89.4	5	5.6	95.5	0.6	3.9

CD: correct detections; FP: false positives; FN: false negatives (two joined trees count as one false negative).

middle, and the result with the 'gas of circles' model is on the right. Again, the 'gas of circles' model better delineates the tree crowns and separates more trees, but some joined trees remain also. The HOAC model selects only objects of the size chosen, so that false positives involving small objects do not occur.

Table 1 shows the percentages of correct tree detections, false positives and false negatives (two joined trees count as one false negative), obtained with the classical active contour model and the 'gas of circles' model in the experiments shown in Figs. 5–7. The 'gas of circles' model outperforms the classical active contour in all measures, except in the number of false negatives in the experiment in Fig. 6.

The typical runtime of the 'gas of circles' model in these experiments (image size $O(100)$ pixels) is of the order of 10 minutes on a normal personal computer. In our implementation, this is approximately ten times slower than using classical active contours.

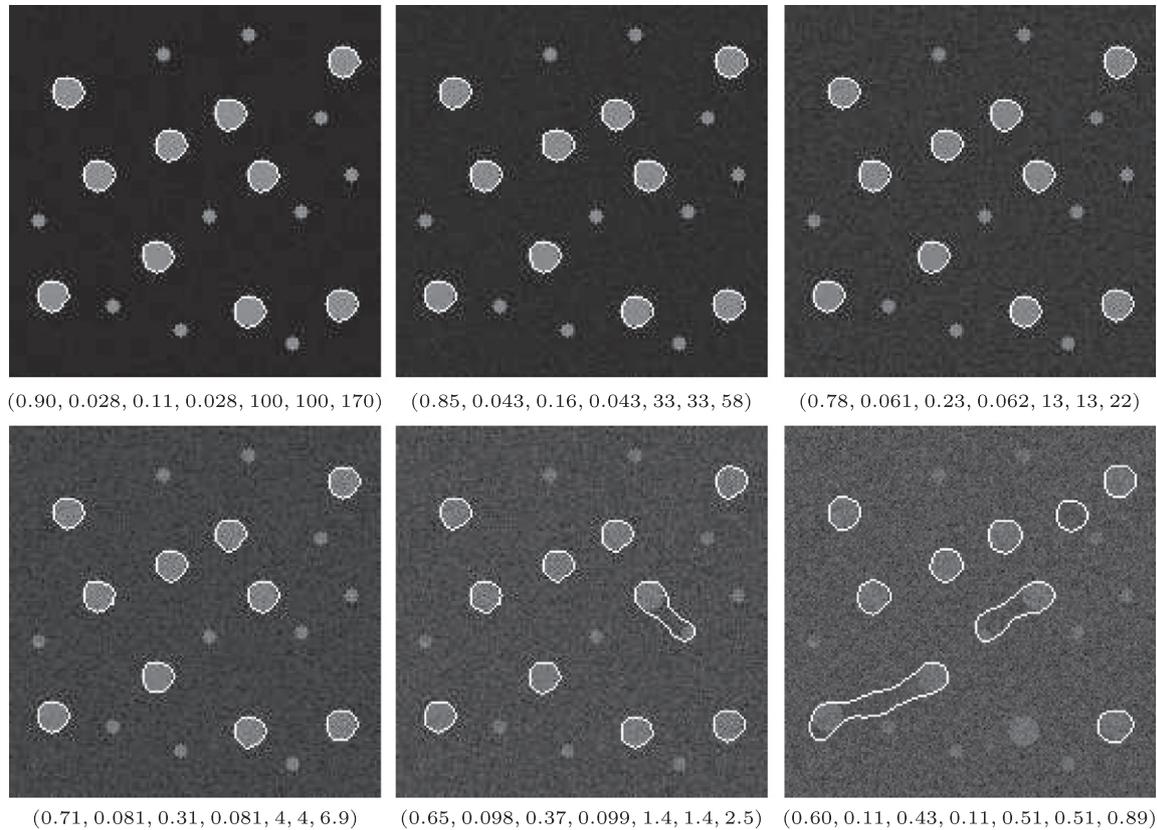


Fig. 8. One of the synthesized images, with six different levels of added white Gaussian noise. Reading from left to right, top to bottom, the image variance to noise power ratios are 20, 15, 10, 5, 0, -5 dB. Parameter values in the form $(\mu, \sigma, \bar{\mu}, \bar{\sigma}, \alpha_C, \beta_C)$ are shown under the six images. The parameters d and r_0 were fixed to 8 throughout.

Table 2
Results on synthetic noisy images

Noise (dB)	20	15	10	5	0	-5
FP %	0	0	0	2	6.4	27.6
FN %	0	0	0	0	4	3.6
J %	0	0	0	0	0	23

FP, FN, J: percentages of false positive, false negative, and joined circle detections, respectively, with respect to the potential total number of correct detections.

4.4. Noisy synthetic images

In this section, we present the results of tests of the sensitivity of the model to noise in the image. Fifty synthetic images were created, each with ten circles with radius 8 pixels and 10 circles with radius 3.5 pixels, placed at random but with overlaps rejected. Six different levels of white Gaussian noise, with image variance to noise power ratios from -5 to 20 dB, were then added to the images to generate 300 noisy images. Six of these, corresponding to noisy versions of the same original image, were used to learn μ , σ , $\bar{\mu}$, and $\bar{\sigma}$. The model used was the same as that used for the aerial images, except that λ_i was set equal to zero. The parameters were adjusted to give a stable radius of 8 pixels.

The results obtained on the noisy versions of one of the 50 images are shown in Fig. 8. Table 2 shows the proportion of false negative and false positive circle detections with respect to the total number of potentially correctly detectable circles ($500 = 50 \times 10$), as well as the proportion of 'joined circles', when two circles are grouped together (an example can be seen in the bottom right image of Fig. 8). Detections of one of the smaller circles (which only occurred a few times even at the highest noise level) were counted as false positives.

The method is very robust with respect to all but the highest levels of noise. The first errors occur at 5 dB, where there is a 2% false positive rate. At 0 dB, the error rate is $\sim 10\%$, i.e. one of the 10 circles in each image was misidentified on average. At -5 dB, the total error rate increases to $\sim 30\%$, rendering the method not very useful.

Note that the principal error modes of the model are false positives and joined circles. There are good reasons why these two types of error dominate. We will discuss them further in Section 5.

4.5. Circle separation: comparison to classical active contours

In a final experiment, we simulated one of the most important causes of error in tree crown extraction, and examined the response of classical active contour and HOAC models to this situation. The errors, which involve joined circles similar to those found in the previous experiment, are caused by the fact that in many cases nearby tree crowns in an image are connected by regions of significant intensity with significant gradient with respect to the background, thus forming a dumbbell shape. Calling the bulbous extremities, the 'bells', and the join between them, the 'bar', the situation arises when the bells are brighter than the bar, while the bar is in turn brighter than the background, and most importantly, the gradient between the background and the bar is greater than that between the bar and the bells.

The first row of Fig. 9 shows a sequence of bells connected by bars. The intensity of the bar varies along the sequence, resulting in different gradient values. We applied the classical active contour and 'gas of circles' models to these images.

The middle row of Fig. 9 shows the best results obtained using the classical active contour model. The model was either unable to

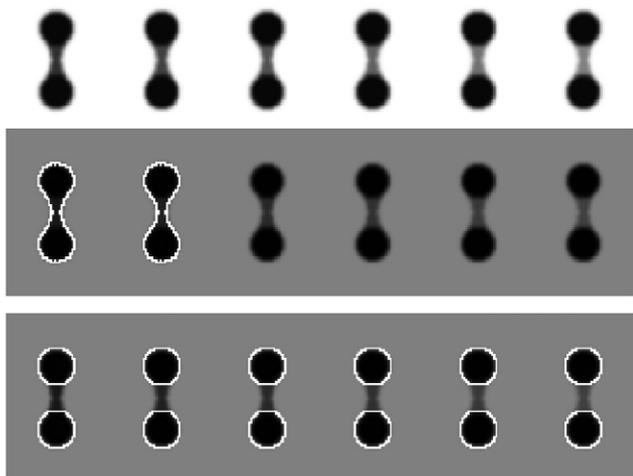


Fig. 9. Results on circle separation comparing the HOAC ‘gas of circles’ model to the classical active contour model. Top: original images. The intensity of the bar takes values equally spaced between 48 and 128 from left to right; the background is 255; the bells are 0. In the middle: the best results obtained using the classical active contour model (8, 1, 1). Either the circles are not separated or the region vanishes. Bottom: the results using the ‘gas of circles’ model (2, 1, 5, 4.0, 8, 8). All the circles are segmented correctly.

separate the individual circles, or the region completely vanished. The intuition is that if there is insufficient gradient to stop the region at the sides of the bar, then there will also be insufficient gradient to stop the region at the boundary between the bar and the bells, so that the region will vanish. On the other hand, if there is sufficient gradient between the bar and the background to stop the region, the circles will not be separated, and a ‘bridge’ will remain between the two circles.²

The corresponding results using the ‘gas of circles’ model are shown in the bottom row of Fig. 9. All the circles were segmented correctly, independent of the grey level of the bar. Encouraging as this is, it is not the whole story, as we indicated in Section 4.4. We make a further comment on this issue in Section 5.

5. Conclusion

Higher-order active contours allow the inclusion of sophisticated prior information in active contour models. HOACs are particularly well adapted to cases in which the topology is unknown *a priori*. In this paper, we have shown via a stability analysis that a HOAC energy can be constructed that describes a ‘gas of circles’, that is, it favours regions composed of an *a priori* unknown number of circles of a certain radius. The requirement that circles be stable, i.e. local minima of the energy, fixes one of the prior parameters and constrains the others.

The ‘gas of circles’ model has many uses in computer vision and image processing. Combined with a suitable likelihood, we have applied it to the problem of tree crown extraction from aerial images of plantations. It performs better than simpler techniques such as maximum likelihood and classical active contours. In particular, it is better able to separate trees that appear joined in the data than is a classical active contour model.

The model is not without its issues, however. First, the computation time is too long. We are currently working on a phase field HOAC [23] version of the ‘gas of circles’ model that we hope will significantly reduce this time. Second, there are two significant

error modes, as shown in the noise experiments of Section 4.4: circles are found where the data do not ostensibly support them (‘phantom circles’), and two circles may be joined into a dumbbell shape and never separated. We discuss these in turn.

The first issue is that of ‘phantom’ circles. Circles of radius \hat{r}_0 are local minima of the prior energy. It is the effect of the data that converts such configurations into global minima. Were we able to find the global minimum of the energy, this would be fine. However, gradient descent finds only a local minimum. This can create problems in areas where the data do not support the existence of circles because a circle, once formed during gradient descent, cannot disappear unless there is an image force acting on it. We thus find that circles can appear and remain even though there is no data to support them.

The second issue is that of joined circles, discussed in Section 4.5. Although the current HOAC model is better able to separate circles than a classical active contour, it still fails to do so in a number of cases, leaving a bridge between the circles. The issue here is a delicate balance between the parameters, which must be adjusted so that the sides of the bridge attract one another, thus breaking the bridge, and so that nearby circles repel one another at close range, so that the bridge does not re-form. Again, this is at least in part an algorithmic issue. Even if the two separated circles have a lower energy than the joined circles, separation may never be achieved due to a local minimum caused by the bridge.

We propose to solve the first problem via a more detailed theoretical analysis of the circle energy that will allow us to remove the local minima causing the problem, and the second via an in-depth analysis of the energy of the dumbbell configuration. Both these studies should lead to further constraints on the parameters, which is a desirable goal in itself.

Acknowledgements

This work was partially supported by EU project IMAVIS (FP5 IHP-MCHT99), EU project MUSCLE (FP6-507752), Egide PAI Balaton (09020XJ), OTKA T-046805, and a Janos Bolyai Research Fellowship of HAS. The authors would like to thank the French National Forest Inventory (IFN) for the aerial images.

Appendix A. Details of stability computations

In this appendix, we give most of the steps involved in reaching Eq. (3.5). The equation of the region boundary is

$$\gamma(t) = \gamma_0(t) + \delta\gamma(t) = (r(t), \theta(t)) = (r_0(t) + \delta r(t), \theta_0(t)), \quad (\text{A.1})$$

where $\gamma_0(t) = (r_0(t), \theta_0(t)) = (r_0, t)$, $\delta r(t) = \sum_k a_k e^{ir_0 kt}$, and $k \in \{m/r_0 : m \in \mathbb{Z}\}$. The components of $\dot{\gamma}$ are

$$\dot{\theta}(t) = 1 \quad \text{and} \quad \dot{r}(t) = \dot{\delta r}(t) = \sum_k a_k i r_0 k e^{ir_0 kt}. \quad (\text{A.2})$$

The tangent vector field is given by

$$\tau(t) = \dot{r}(t)\partial_r + \dot{\theta}(t)\partial_\theta. \quad (\text{A.3})$$

A.1. Linear terms

To compute the length, we need the magnitude of τ to second order. The metric in polar coordinates is $ds^2 = dr^2 + r^2 d\theta^2$, so we have that $|\tau(t)|^2 = \dot{r}(t)^2 + r(t)^2$ by Eqs. (A.2). Substituting from Eqs. (A.1) and (A.2) gives

$$|\tau(t)|^2 = r_0^2 + 2r_0 \sum_k a_k e^{ir_0 kt} + \sum_{k,k'} a_k a_{k'} e^{ir_0(k+k')t} (1 - r_0^2 k k'). \quad (\text{A.4})$$

² ‘Bar’ and ‘bell’ refer to image properties; we use ‘bridge’ and ‘circle’ to refer to the corresponding pieces of a dumbbell-shaped region.

Taking the square root, expanding it as $\sqrt{1+x} \approx 1 + \frac{1}{2}x - \frac{1}{8}x^2$, and keeping terms to second order in the a_k then gives

$$|\tau(t)| = r_0 \left[1 + \sum_k \frac{a_k}{r_0} e^{ir_0 kt} - \frac{1}{2} \sum_{k,k'} a_k a_{k'} k k' e^{ir_0(k+k')t} \right]. \quad (A.5)$$

Using Eq. (A.5), the boundary length is then given to second order by

$$L(\gamma) = \int_{-\pi}^{\pi} dt |\tau(t)| = 2\pi r_0 \left\{ 1 + \frac{a_0}{r_0} + \frac{1}{2} \sum_k k^2 |a_k|^2 \right\},$$

where we have used the reality of δr to set $a_{-k} = a_k^*$, where * indicates complex conjugation, and orthonormality of the Fourier basis elements.

We can write the interior area of the region as

$$A(\gamma) = \int_{-\pi}^{\pi} d\theta \int_0^{r(\theta)} dr' r' = \int_{-\pi}^{\pi} d\theta \frac{1}{2} r^2(\theta)$$

Thus, using Eq. (A.1), and again using orthonormality, we have that

$$A(\gamma) = \pi r_0^2 + 2\pi r_0 a_0 + \pi \sum_k |a_k|^2. \quad (A.6)$$

A.2. Quadratic terms

To compute the expansion of the quadratic term in Eq. (2.2) for E_g , we need the expansions of $\tau(t) \cdot \tau(t')$ and $\Phi(R(t, t'))$.

A.2.1. Inner product of tangent vectors

The tangent vector is given by Eq. (A.3), but we must take care as $\tau(t)$ and $\tau(t')$ live in different tangent spaces, at $\gamma(t)$ and $\gamma(t')$, respectively. It is easiest to convert the tangent vectors to the Euclidean coordinate basis, $\partial_r = \cos(\theta)\partial_x + \sin(\theta)\partial_y$ and $\partial_\theta = -r \sin(\theta)\partial_x + r \cos(\theta)\partial_y$ as these basis vectors are preserved by parallel transport. Taking the inner product then gives

$$\tau \cdot \tau' = \cos(\theta' - \theta) [r_0^2 + r_0 \delta r + r_0 \delta r' + \delta r \delta r' + \dot{\delta r} \dot{\delta r}'] + \sin(\theta' - \theta) [r_0 \dot{\delta r}' - r_0 \dot{\delta r} + \delta r \dot{\delta r}' - \dot{\delta r} \delta r'],$$

where unprimed quantities are evaluated at t and primed quantities at t' .

A.2.2. Interaction function

First, we expand $R(t, t')$. The squared distance between $\gamma(t')$ and $\gamma(t)$ is given by

$$|\gamma(t') - \gamma(t)|^2 = [(r_0 + \delta r') \cos(\theta') - (r_0 + \delta r) \cos(\theta)]^2 + [(r_0 + \delta r') \sin(\theta') - (r_0 + \delta r) \sin(\theta)]^2,$$

which after expansion gives

$$R^2(t, t') = 2r_0^2(1 - \cos(\Delta t)) \left\{ 1 + \frac{1}{r_0}(\delta r + \delta r') + \frac{\delta r^2 + \delta r'^2 - 2 \cos(\Delta t) \delta r \delta r'}{2r_0^2(1 - \cos(\Delta t))} \right\},$$

where $\Delta t = \theta' - \theta = t' - t$. Expanding $\sqrt{1+x} \approx 1 + \frac{1}{2}x - \frac{1}{8}x^2$ to second order and collecting terms, we then find

$$R(t, t') = 2r_0 |\sin(\Delta t/2)| + |\sin(\Delta t/2)|(\delta r + \delta r') + \frac{A(\Delta t)}{4r_0}(\delta r - \delta r')^2, \quad (A.7)$$

where $A(z) = \cos^2(z/2) |\sin(z/2)|^{-1}$.

Expanding $\Phi(z)$ in a Taylor series to second order, and then substituting $R(t, t')$ for z using the approximation in Eq. (A.7), and keeping only terms up to second order in $\delta\gamma$ then gives

$$\begin{aligned} \Phi(R(t, t')) &= \Phi(X_0) + \left| \sin \frac{\Delta t}{2} \right| \Phi'(X_0)(\delta r + \delta r') \\ &+ \frac{1}{4r_0} A(\Delta t) \Phi''(X_0)(\delta r - \delta r')^2 \\ &+ \frac{1}{2} \sin^2 \left(\frac{\Delta t}{2} \right) \Phi''(X_0)(\delta r + \delta r')^2, \end{aligned} \quad (A.8)$$

where $X_0 = 2r_0 |\sin(\Delta t/2)|$.

A.3. Combining terms

Using the above equations gives

$$\begin{aligned} G(t, t') &= \underbrace{r_0^2 \cos(\Delta t) \Phi(X_0)}_{F_{00, \text{even}}} \\ &+ \underbrace{(\delta r + \delta r') r_0 \cos(\Delta t) \left\{ \Phi(X_0) + r_0 \left| \sin \frac{\Delta t}{2} \right| \Phi'(X_0) \right\}}_{F_{10, \text{even}}} \\ &+ \underbrace{(\dot{\delta r}' - \dot{\delta r}) r_0 \sin(\Delta t) \Phi(X_0)}_{F_{11, \text{odd}}} \\ &+ \underbrace{(\delta r^2 + \delta r'^2) r_0 \cos(\Delta t) \left\{ \frac{1}{4} A(\Delta t) \Phi'(X_0) + \frac{1}{2} r_0 \sin^2 \left(\frac{\Delta t}{2} \right) \Phi''(X_0) + \left| \sin \frac{\Delta t}{2} \right| \Phi'(X_0) \right\}}_{F_{20, \text{even}}} \\ &+ \underbrace{(\delta r \delta r') \cos(\Delta t) \left\{ \Phi(X_0) + 2r_0 \left| \sin \frac{\Delta t}{2} \right| \Phi'(X_0) - \frac{1}{2} r_0 A(\Delta t) \Phi'(X_0) + r_0^2 \sin^2 \left(\frac{\Delta t}{2} \right) \Phi''(X_0) \right\}}_{F_{21, \text{even}}} \\ &+ \underbrace{(\delta r' \dot{\delta r}' - \delta r \dot{\delta r}) r_0 \left| \sin \frac{\Delta t}{2} \right| \sin(\Delta t) \Phi'(X_0)}_{F_{22, \text{odd}}} \\ &+ \underbrace{(\delta r \dot{\delta r}' - \delta r' \dot{\delta r}) \sin(\Delta t) \left\{ \Phi(X_0) + r_0 \left| \sin \frac{\Delta t}{2} \right| \Phi'(X_0) \right\}}_{F_{23, \text{odd}}} \\ &+ \underbrace{(\dot{\delta r} \dot{\delta r}') \cos(\Delta t) \Phi(X_0)}_{F_{24, \text{even}}}, \end{aligned}$$

where the F_{ij} denote the functions appearing in the terms of G , and ‘odd’ and ‘even’ refer to parity under exchange of t and t' . Each line, and hence G , is symmetric in t and t' , as it should be. We can now substitute the expressions for δr and $\delta r'$ in terms of their Fourier coefficients, and calculate $\int_{-\pi}^{\pi} dt dt' G(t, t')$. We note that in the terms involving F_{10} , F_{11} , F_{20} , F_{22} , and F_{23} , the presence of the symmetric or antisymmetric factors in δr and $\delta r'$ simply leads to a doubling of the value of the integral for one of the terms in these factors, due to the corresponding symmetry or antisymmetry of the F functions. We therefore only need to evaluate one of these integrals for the relevant terms.

Because the F s depend only on Δt , the resulting integrals can be reduced, via a change of variables $p = \Delta t$, to integrals over p . For F_{00} and F_{10} , we have

$$\int_{-\pi}^{\pi} dt dt' F_{00}(t' - t) = \int_{-\pi}^{\pi} dp dt' F_{00}(p) = 2\pi \int_{-\pi}^{\pi} dp F_{00}(p),$$

and

$$\begin{aligned}
\iint_{-\pi}^{\pi} dt dt' \delta r(t) F_{10}(t'-t) &= \iint_{-\pi}^{\pi} dt dt' \sum_k a_k e^{ir_0 kt} F_{10}(t'-t) \\
&= \sum_k a_k \iint_{-\pi}^{\pi} dp dt' e^{ir_0 k(-p+t')} F_{10}(p) \\
&= \sum_k a_k \int_{-\pi}^{\pi} dt' e^{ir_0 kt'} \int_{-\pi}^{\pi} dp e^{-ir_0 kp} F_{10}(p) \\
&= \sum_k a_k 2\pi \delta(k) \int_{-\pi}^{\pi} dp e^{-ir_0 kp} F_{10}(p) \\
&= 2\pi a_0 \int_{-\pi}^{\pi} dp F_{10}(p).
\end{aligned}$$

The calculations for the other terms proceed in a very similar fashion, using the same change of variable and the orthonormality of the Fourier basis. We merely list the results (full details may be found in [24]):

$$\begin{aligned}
\iint_{-\pi}^{\pi} dt dt' \dot{\delta} r(t) F_{11}(t'-t) &= 0 \\
\iint_{-\pi}^{\pi} dt dt' \delta r^2(t) F_{20}(t'-t) &= 2\pi \sum_k |a_k|^2 \int_{-\pi}^{\pi} dp F_{20}(p) \\
\iint_{-\pi}^{\pi} dt dt' \delta r(t) \delta r(t') F_{21}(t'-t) &= 2\pi \sum_k |a_k|^2 \int_{-\pi}^{\pi} dp e^{-ir_0 kp} F_{21}(p) \\
\iint_{-\pi}^{\pi} dt dt' \delta r(t) \dot{\delta} r(t') F_{22}(t'-t) &= 0 \\
\iint_{-\pi}^{\pi} dt dt' \delta r(t) \delta r(t') F_{23}(t'-t) &= -2\pi \sum_k |a_k|^2 ir_0 k \int_{-\pi}^{\pi} dp e^{-ir_0 kp} F_{23}(p) \\
\iint_{-\pi}^{\pi} dt dt' \dot{\delta} r(t) \dot{\delta} r(t') F_{24}(t'-t) &= 2\pi \sum_k |a_k|^2 r_0^2 k^2 \int_{-\pi}^{\pi} dp e^{-ir_0 kp} F_{24}(p).
\end{aligned}$$

Using these results then gives Eq. (3.4), which in combination with Eqs. (3.2) and (3.3), gives Eq. (3.5).

References

- [1] M. Kass, A. Witkin, D. Terzopoulos, Snakes: active contour models, *Int. J. Comput. Vision* 1 (4) (1988) 321–331.
- [2] L. Cohen, On active contours and balloons, *CVGIP: Image Understanding* 53 (1991) 211–218.
- [3] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, *Int. J. Comput. Vision* 22 (1) (1997) 61–79.

- [4] Y. Chen, H. Tagare, S. Thiruvankadam, F. Huang, D. Wilson, K. Gopinath, R. Briggs, E. Geiser, Using prior shapes in geometric active contours in a variational framework, *Int. J. Comput. Vision* 50 (3) (2002) 315–328.
- [5] D. Cremers, F. Tschh user, J. Weickert, C. Schn orr, Diffusion snakes: introducing statistical shape knowledge into the Mumford–Shah functional, *Int. J. Comput. Vision* 50 (3) (2002) 295–313.
- [6] M. Leventon, W. Grimson, O. Faugeras, Statistical shape influence in geodesic active contours, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Hilton Head Island, SC, USA, 2000, pp. 316–322.
- [7] N. Paragios, M. Rousson, Shape priors for level set representations, in: *Proceedings of the European Conference on Computer Vision (ECCV)*, Copenhagen, Denmark, 2002, pp. 78–92.
- [8] A. Srivastava, S. Joshi, W. Mio, X. Liu, Statistical shape analysis: clustering, learning, and testing, *IEEE Trans. Pattern Anal. Mach. Intell.* 27 (4) (2005) 590–602.
- [9] A. Tsai, A. Yezzi, W. Wells, C. Tempny, D. Tucker, A. Fan, E. Grimson, A. Willsky, Model-based curve evolution technique for image segmentation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Hawaii, USA, 2001.
- [10] A. Foulonneau, P. Charbonnier, F. Heitz, Geometric shape priors for region-based active contours, *Proceedings of the IEEE International Conference on Image Processing (ICIP)* 3 (2003) 413–416.
- [11] D. Cremers, S. Soatto, A pseudo-distance for shape priors in level set segmentation, in: *Proceedings of the 2nd IEEE Workshop on Variational, Geometric and Level Set Methods*, Nice, France, 2003, pp. 169–176.
- [12] D. Cremers, T. Kohlberger, C. Schn orr, Shape statistics in kernel space for variational image segmentation, *Pattern Recognition* 36 (9) (2003) 1929–1943.
- [13] D. Cremers, N. Sochen, C. Schn orr, A multiphase dynamic labeling model for variational recognition-driven image segmentation, *Int. J. Comput. Vision* 66 (1) (2006) 67–81.
- [14] M. Rochery, I.H. Jermyn, J. Zerubia, Higher-order active contours, *Int. J. Comput. Vision* 69 (1) (2006) 27–42.
- [15] Y. Choquet-Bruhat, C. DeWitt-Morette, M. Dillard-Bleick, *Analysis, Manifolds and Physics*, Elsevier Science, Amsterdam, The Netherlands, 1996.
- [16] G. Sundaramoorthi, A. Yezzi, More-than-topology-preserving flows for active contours and polygons, in: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Washington DC, USA, 2005, pp. 1276–1283.
- [17] P. Horvath, I. H. Jermyn, J. Zerubia, Z. Kato, An improved model for tree crown extraction using higher order active contours, *Research report, INRIA, France*, November 2007, to appear.
- [18] F.A. Gougeon, Automatic individual tree crown delineation using a valley-following algorithm and rule-based system, in: D. Hill, D. Leckie (Eds.), in: *Proceedings of the International Forum on Automated Interpretation of High Spatial Resolution Digital Imagery for Forestry*, Victoria, British Columbia, Canada, 1998, pp. 11–23.
- [19] M. Larsen, Finding an optimal match window for Spruce top detection based on an optical tree model, in: D. Hill, D. Leckie (Eds.), *Proceedings of the International Forum on Automated Interpretation of High Spatial Resolution Digital Imagery for Forestry*, Victoria, British Columbia, Canada, 1998, pp. 55–66.
- [20] T. Brandtberg, F. Walter, Automated delineation of individual tree crowns in high spatial resolution aerial images by multiple-scale analysis, *Mach. Vision Appl.* 11 (1998) 64–73.
- [21] G. Perrin, X. Descombes, J. Zerubia, A marked point process model for tree crown extraction in plantations, in: *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, Genova, Italy, 2005.
- [22] S. Osher, J.A. Sethian, Fronts propagating with curvature dependent speed: algorithms based on Hamilton–Jacobi formulations, *J. Comput. Phys.* 79 (1) (1988) 12–49.
- [23] M. Rochery, I.H. Jermyn, J. Zerubia, Phase field models and higher-order active contours, in: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Beijing, China, 2005.
- [24] P. Horvath, I.H. Jermyn, Z. Kato, J. Zerubia, A higher-order active contour model of a ‘gas of circles’ and its application to tree crown extraction, *Research Report 6026, INRIA, France*, November 2006.

About the Author—P ETER HORV ATH received an M.S. (Computer Science) from the University of Szeged, Hungary (2004), a joint Ph.D. (Image Segmentation) from the University of Nice Sophia Antipolis, doing his research at INRIA Sophia Antipolis, France, and from the University of Szeged, Hungary (2007). He is currently a Staff Scientist at ETH Zurich. His research interests include image segmentation, motion detection, and image based classification.

About the Author—IAN JERMYN received a B.A. (Physics) from Oxford University (1986), a Ph.D. (Theoretical Physics) from Manchester University (1991), and a Ph.D. (Computer Science) from the Courant Institute, New York University (2000). He is currently a Senior Research Scientist in the Ariana group at INRIA. His research interests include shape and texture modelling.

About the Author—ZOLTAN KATO received the M.S. degree in Computer Science from the University of Szeged, Hungary in 1990, and the Ph.D. degree from the University of Nice doing his research at INRIA Sophia Antipolis, France in 1994. Currently, he is head of the Department of Image Processing and Computer Graphics at the University of Szeged. His research interests include statistical image models, MCMC methods, and shape modeling.

About the Author—JOSIANE ZERUBIA received an MSc (Electrical Engineering) from ENSIEG (1981), a Doctor of Engineering (1986), a Ph.D. (1988), and an ‘Habilitation’ (1994) all from the University of Nice Sophia Antipolis. She is currently a director of research and head of the Ariana group at INRIA. She is an IEEE Fellow. Her research interests are in image processing using probabilistic models and variational methods.