

UNIVERSITY OF SZEGED - HUNGARY
Doctoral School in Mathematics and Computer Science

UNIVERSITY OF NICE - SOPHIA ANTIPOLIS - FRANCE
Information and Communication Sciences & Technologies Doctoral School

THESIS

for the degree of

DOCTOR OF PHILOSOPHY

by

Péter HORVÁTH

THE ‘GAS OF CIRCLES’ MODEL
AND ITS APPLICATION TO TREE CROWN EXTRACTION

Supervised by Ian H. JERMYN², Zoltán KATÓ¹ and Josiane ZERUBIA²
and prepared at the University of Szeged¹ and at
INRIA Sophia Antipolis in the ARIANA² project

2007

Dedication

To Krisztina, who not just had her faith within me but shown me that
nothing is impossible,

and

In Memory of my Grandfather, HORVÁTH József.

*“Two roads diverged in a wood, and I –
I took the one less traveled by,
And that has made all the difference.”*

*“Az erdőben egy útelágazáshoz értem, s én –
Én a kevésbé jártat választottam,
S ez volt minden különbség.”*

Robert FROST – *The Road Not Taken* (1920)

Acknowledgements

First of all I would like to say thank you to a great person, researcher, adviser, and friend who always motivated me and shown me what science and dedication is, this person is Ian. I would like to express my appreciation to Josiane and Zoltán for the huge amount of advice they gave me; the work we had done together was sometimes very hard but always the best way to follow for me.

To my parents Erzsébet and József, and my sister Linda, who were by my side even in the most disappointing situation as well as in my successful moments, supported my studies above all sense. I feel grateful to the rest of my family, and all of my friends (Bazsi, Pisti, Zsolti, Atus, Tanti, Döme Busi, Ricsi, ...), who always understood me and helped me to relax when I needed the most.

Thanks to my colleges who had been working with me at Ariana, and in the University of Szeged. I also appreciate them for the nice time spent together out of work (Avik, Enrique, Marie, Guillaume, Giuseppe, Mats, Corinne, Laure, Xavier, Dan, Ting, Pierre, Alexy, Alexander, Maria, Caroline, Praveen, Zotya, Ájven, Erick, Aymen, Olivier, Csabi, Santanu). Thank you to the collaborators of Ariana and University of Szeged for giving me new ideas and listening my talks with great interest (Prof. Anuj Srivastava, Prof. Joseph Francos, Prof. Nick Kingsbury, Prof. Attila Kuba). Thank you to Prof. Pierre Couteron for inviting me to Pondicherry, India and organizing my stay and lecture.

Thank you to the French Forest Inventory (IFN) and the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) for providing aerial forestry images. Thank you to EU project IMAVIS (FP5 IHPMCHT99), EU project MUSCLE (FP6-507752), PAI “Balaton”, INRIA Sophia Antipolis, the French Ministry of Foreign Affairs, the CROUS Nice, especially Mme. Ghislaine Rodriguez, the Hungarian Scholarship Board (MÖB), the French Institute in Budapest, especially to M. Bob Kaba-Loemba, the Mecénatúra Scholarship Agency, and last but not least to the Doctoral School of the University of Szeged for financially supporting my studies, stays, and trips.

Contents

| | |
|---|------------|
| Acknowledgements | iii |
| Notations | xi |
| 1 Introduction | 1 |
| 1.1 Prior knowledge and beyond, probabilistic view | 1 |
| 1.2 State of the art: from image segmentation to the ‘gas of circles’ model . . . | 3 |
| 1.3 The ‘gas of circles’ model for tree crown detection | 8 |
| 1.4 Road map | 11 |
| 2 The higher-order active contour ‘gas of circles’ model | 13 |
| 2.1 The higher-order active contour model | 14 |
| 2.2 The ‘gas of circles’ model | 17 |
| 2.2.1 Stability analysis | 18 |
| Parameter constraints | 20 |
| 2.2.2 Geometric and synthetic experiments | 22 |
| The geometric energy | 22 |
| Noisy synthetic images | 23 |
| Circle separation: comparison to classical active contours | 24 |
| 2.3 The inflection point ‘gas of circles’ model | 27 |
| Determination of d | 28 |
| 2.4 Minimization of the energy, implementation issues | 31 |
| 2.4.1 Minimization of the higher-order active contour energy | 31 |
| 2.4.2 Implementation of the higher-order active contour evolution | 32 |
| 2.4.3 Implementation of the parameter estimations | 33 |
| Determination of β_C parameter: | 33 |
| The inflection point parameters: | 34 |
| Determination of d : | 34 |
| 3 The phase field ‘gas of circles’ model | 35 |
| 3.1 The higher-order phase field model | 37 |
| 3.2 The ‘gas of circles’ phase field model | 40 |
| 3.2.1 Recipe: determination of the parameters | 41 |
| 3.2.2 Synthetic experiments | 41 |

| | | |
|----------|---|-----------|
| 3.3 | Inflection point in the energy function | 42 |
| 3.3.1 | Experimental verification | 43 |
| 3.4 | Minimization of the phase field energy | 43 |
| 3.4.1 | Minimization of the local term, topology | 44 |
| 3.4.2 | The non-local term | 45 |
| 3.4.3 | Initialization | 45 |
| 3.4.4 | Implementation details | 46 |
| | The interaction operator | 46 |
| | The higher-order term | 46 |
| | Boundary condition | 47 |
| 4 | Tree crown extraction | 49 |
| 4.1 | History of tree crown extraction from aerial images | 50 |
| 4.1.1 | Data acquisition and the aerial images | 50 |
| 4.1.2 | Individual tree crown delineation | 52 |
| 4.2 | Data terms for tree crown extraction | 54 |
| 4.2.1 | Model for single band images | 55 |
| | Energy minimization | 57 |
| | Phase field data term | 57 |
| 4.2.2 | Multispectral data model | 57 |
| | Energy minimization | 61 |
| 4.3 | Experimental results | 61 |
| 4.3.1 | The ‘gas of circles’ HOAC model | 62 |
| 4.3.2 | The inflection point ‘gas of circles’ model | 63 |
| 4.3.3 | The ‘gas of circles’ phase field model | 67 |
| 4.3.4 | The inflection point ‘gas of circles’ phase field model | 68 |
| 4.3.5 | Experiments on color infrared images | 77 |
| 5 | Conclusion, unsolved problems | 81 |
| 5.1 | Summary | 81 |
| 5.2 | Unsolved problems | 83 |
| 5.3 | Possible directions | 84 |
| A | Details of stability computations | 87 |
| A.1 | Length | 88 |
| A.2 | Area | 88 |
| A.3 | Quadratic energy | 88 |
| A.3.1 | Inner product of tangent vectors | 88 |
| A.3.2 | Distance between two points | 89 |
| A.3.3 | Interaction function | 89 |
| A.3.4 | Combining terms | 90 |

| | |
|---|------------|
| B Polynomial approximation | 95 |
| B.1 Scaling analysis of the circle energy | 95 |
| B.2 The Taylor-polynomial of the energy function | 96 |
| B.2.1 The interaction function | 96 |
| B.2.2 The derivatives of G_{00} | 97 |
| B.2.3 The polynomial coefficients | 98 |
| B.2.4 Polynomial approximation of G_{00} , G_{10} and \tilde{G} | 99 |
| Publications and scientific activities of the author | 101 |
| Bibliography | 103 |

List of Figures

| | | |
|------|---|----|
| 1.1 | Level-set and contour evolution. | 6 |
| 1.2 | Real image with planted forest ©IFN. | 9 |
| 2.1 | The interaction function. | 16 |
| 2.2 | Two situations of point pairs. | 16 |
| 2.3 | Examples of gradient descent using the HOAC energy. | 17 |
| 2.4 | Road detection on a SPOT satellite image ©CNES. | 17 |
| 2.5 | Curve evolution using different HOAC parameters. | 18 |
| 2.6 | Plots of E_0 against r_0 and E_2 against $\hat{r}_0 k$ | 21 |
| 2.7 | Schematic plot of the positions of the extrema of the energy of a circle versus β_C , for α_C fixed. | 22 |
| 2.8 | Experimental results using the geometric term. | 23 |
| 2.9 | Synthesized images with six different levels of added white Gaussian noise. | 25 |
| 2.10 | Results on circle separation comparing the HOAC ‘gas of circles’ model to the classical active contour model. | 26 |
| 2.11 | Results on regularly planted poplars ©IFN | 28 |
| 2.12 | α and β against d near the critical domain. | 29 |
| 2.13 | Plot of E_0 against r for α and β values resulting inflection point. | 30 |
| 2.14 | Contour tracing algorithm. | 32 |
| 2.15 | Configurations in the contour tracking algorithm | 33 |
| 3.1 | The V term of the linear energy. | 37 |
| 3.2 | Simplified model to approximate the parameters. | 38 |
| 3.3 | Evolution of the region using the simple phase field model. | 39 |
| 3.4 | Region evolution using the non-local phase field energy. | 40 |
| 3.5 | Comparison of the HOAC and phase field evolutions. | 42 |
| 3.6 | Phase field evolution. | 43 |
| 3.7 | Graphs of the energy of a circle. | 43 |
| 3.8 | Preservation of the inflection point. | 44 |
| 3.9 | Phase field evolution with different initial conditions, showing topological freedom. | 45 |
| 3.10 | The Fourier transform of the interaction function. | 46 |
| 3.11 | Mapping of a forestry image onto a torus. French Forest Inventory (IFN) | 47 |
| 4.1 | Plot of vegetation reflectance against wavelength. | 51 |

| | | |
|------|---|----|
| 4.2 | False colouring technique. | 52 |
| 4.3 | Real color infrared (CIR) image with forest and urban area ©IFN. | 52 |
| 4.4 | Three typical cases of segmentation of individual tree crowns. | 53 |
| 4.5 | Infrared spectral band of a CIR image. | 55 |
| 4.6 | Magnitude of the gradient of the infrared spectral band. | 56 |
| 4.7 | CIR image and its ground truth segmentation. | 58 |
| 4.8 | Histograms of pixel values in the three bands of a CIR image. | 58 |
| 4.9 | Maximum likelihood classifications of figure 4.7 (left) using the different models trained on the same image. | 59 |
| 4.10 | The same classification trained on figure 4.7 (middle). | 60 |
| 4.11 | Image of poplars ©IFN. | 63 |
| 4.12 | Image with a planted poplar forest ©IFN. | 64 |
| 4.13 | Image with a regularly planted forest ©IFN. | 64 |
| 4.14 | Result obtained in figure 4.11. | 65 |
| 4.15 | Regularly planted poplars ©IFN. | 65 |
| 4.16 | Bigger slice of planted forest ©IFN. | 66 |
| 4.17 | Regularly planted poplars ©IFN. | 67 |
| 4.18 | Experimental result with the phase field GOC model on figure 4.13 | 68 |
| 4.19 | Result with the phase field GOC model on figure 4.11 | 69 |
| 4.20 | Result on figure 4.17. | 69 |
| 4.21 | Image of regularly planted poplar stands with a less regularly planted trees at the upper part © IFN. | 70 |
| 4.22 | Image of regularly planted poplars with different fields on the right © IFN | 72 |
| 4.23 | Real image with sparsely planted trees © CAO, FA. | 73 |
| 4.24 | Real image with regularly planted poplars next to a farm © CAO, FA. | 74 |
| 4.25 | Rare irregular separated tree crowns © CAO, FA. | 74 |
| 4.26 | Regularly planted pine forest © CAO, FA. | 75 |
| 4.27 | Real image with a forest © CAO, FA. | 76 |
| 4.28 | Regularly planted poplar forest with fields in the top and bottom corners © CAO, FA. | 76 |
| 4.29 | Results obtained on the image shown in figure 4.7. | 77 |
| 4.30 | Results obtained on the image shown in figure 4.7 | 77 |
| 4.31 | Results on a CIR image © IFN. | 78 |
| 4.32 | Real results on a CIR image © IFN. | 79 |
| 4.33 | Results with the multispectral GOC model on CIR image © IFN. | 80 |
| 5.1 | Overview and connection between the methods we used and developed. | 82 |
| 5.2 | Illustration of the attraction arising between circles. | 83 |
| 5.3 | Proposed configurations to determine the repulsive force between circles. | 84 |

Notations

| | |
|-----------------------------|---|
| α | the parameter of the phase field energy |
| α_C | the area parameter of the contour |
| $A(R)$ | the area of the region R |
| β_C | the parameter of the quadratic HOAC term |
| CAC | Classical Active Contour |
| CIR | Colour Infrared Image |
| d | the interaction range parameter of the interaction function |
| ϵ | the parameter of the interaction function |
| E_g | the geometric or prior energy |
| E_i | the data or image energy |
| E_{NL} | the non-local phase field energy |
| γ | the contour |
| $\square\gamma$ | the domain of γ |
| GOC | Gas of Circles |
| HOAC | Higher-Order Active Contour |
| λ | the parameter of the phase field energy |
| λ_C | the length parameter of the contour |
| $L(R)$ | the length of the region boundary |
| μ_{in}, μ_{out} | the mean values on the background and foreground |
| M_{in}, M_{out} | the mean vectors on the background and foreground |
| Ω | the image domain |
| PDE | Partial Differential Equation |
| PDF | Probability Distribution Function |
| ϕ | the phase field function |
| Ψ | the interaction function |
| $R(t, t')$ | the Euclidean distance between the contour points |
| R | the region |
| ∂R | the boundary of the region R |
| \mathcal{R} | the space of regions |
| $\sigma_{in}, \sigma_{out}$ | the variances on the background and foreground |
| $\Sigma_{in}, \Sigma_{out}$ | the covariance matrices on the background and foreground |
| SNR | Signal-to-Noise Ratio |
| $\dot{\gamma}$ | the tangent vector of γ |
| w | the width of the phase field's interface region |

Chapter 1

Introduction

We have been developing intelligent observation systems to take over the duties of human eyes, in this way helping and relieving man's labour as well as opening up new prospects. The efficiency and completeness of the eyes and brain is unparalleled in comparison with any piece of apparatus or instrumentation ever invented. However, some specific tasks such as the accurate extraction of tree crowns in most aerial images, which is the main purpose of this thesis, is already possible. We provide a probabilistic framework for our approach, and review the literature pertinent to our work, finally we describe the relevance of the work in this thesis to forestry management.

WE present a probabilistic framework to our work in section 1.1, which assist us in classifying image segmentation and energy minimization methods. In section 1.2 we discuss previously published work in the area, ranging from image processing to the models introduced in this thesis. We give an introduction to image processing as applied in forestry management, as well as presenting the results we achieved in this area in section 1.3. Finally in section 1.4 we review the structure of this thesis.

1.1 Prior knowledge and beyond, probabilistic view

In this section we have two purposes. One is to introduce the image segmentation models from a probabilistic point of view. The other is to provide an interpretation of the methods presented in the state of the art section and our methods in the thesis.

Note, the space of possible hypotheses \mathcal{H} , let H be a variable over \mathcal{H} and h a point in this space. Similarly \mathcal{D} denotes the space of all available data while D is a variable and d is a point in \mathcal{D} . If $P(D = d|H = h)$, our knowledge of how the model represented by the hypothesis generates the data is known, for given $h \in \mathcal{H}$ and $d \in \mathcal{D}$, and our prior knowledge of the likelihood for a given h is $P(H = h)$. Then, we can compute the posterior probability of the correctness of our h for a given d and our prior knowledge, using the Bayes theorem, as:

$$P(H = h|D = d) = \frac{P(D = d|H = h)P(H = h)}{P(D = d)} .$$

In image processing using the notations \mathcal{R} the space of regions, \mathcal{J} the space of images, and \mathcal{K} the space of *a priori* knowledge. With the Bayes theorem we can infer R from a given I and K as follows:

$$P(R|I, K) = \frac{P(I|R, K)P(R)}{P(I|K)} . \quad (1.1.1)$$

If the task is to find one R from \mathcal{R} , according to some defined criteria (*e.g.* most probable R for given I and K), we need to define a *loss function* defined on $\mathcal{H} \times \mathcal{H}$ or in the case of image segmentation $\mathcal{R} \times \mathcal{R}$ and called L . $L(\hat{h}, h)$ gives a loss measure, assuming the true hypothesis is \hat{h} when it is in fact h . Different uses of the available I , K , and R may involve different loss functions. In general we want to find the \hat{R} that minimizes the value:

$$\langle L(\hat{R}) \rangle = \int dR L(\hat{R}, R)P(R|I, K) . \quad (1.1.2)$$

Apart from some special cases, if one can use special loss functions (Jermyn, 2000), in general in the absence of any information except the preconditions for probability theory, the delta function is the only one obvious loss function, $L(\hat{R}, R) = -\delta(\hat{R}, R)$. Substituting this function into equation 1.1.2, gives that $\langle L(\hat{R}) \rangle = \int dR -\delta(\hat{R}, R)P(R|I, K) = -P(\hat{R}|I, K)$. Consequently, the region minimizing the mean loss is given by:

$$\hat{R} = \arg \max_{R \in \mathcal{R}} P(R|I, K) . \quad (1.1.3)$$

This is the *maximum a posteriori* (MAP) estimate. The MAP estimate can be rewritten by minimizing the negative logarithm of the probability. The negative logarithm of the probability in statistical physics is known as the *energy*. In the rest of this section we concentrate on this expression.

Notice that we look for the region \hat{R} that maximizes the probability defined in equation 1.1.1 for a given image I and prior knowledge K , therefore $P(I|K)$ in the denominator is constant. We can rewrite the probabilities in the nominator as:

$$P(I|R, K) = \frac{1}{Z_i(R, K)} e^{-E_i(I, R, K)} \quad \text{and}$$

$$P(R|K) = \frac{1}{Z_P(K)} e^{-E_P(R, K)} ,$$

where E_i and E_P are the image and prior energies respectively and Z_i and Z_P are normalizing functions. Notice that Z_P has no dependency on the region \hat{R} . We can determine now the optimal region by maximizing the probability:

$$R_{\text{opt}} = \arg \max_{R \in \mathcal{R}} \frac{e^{-E_i(I, R, K)}}{Z_i(R, K)} e^{-E_P(R, K)}$$

Finally, as we discussed above, rather than maximizing this probability we can minimize its negative logarithm, and so we define a very important approach in image segmentation, the *energy minimization* methods:

$$R_{\text{opt}} = \underbrace{E_i(I, R, K)}_{\text{image energy}} + \underbrace{\ln(Z_i(R, K)) + E_p(R, K)}_{\text{prior energy}} .$$

In this section, we have presented a probabilistic interpretation of our problem: how it is possible to find the region(s) in a given image that fits best with our prior knowledge. This framework is very general and all image segmentation models can be classified by using it. In the second half of the section, we described one of the possible ways to compute MAP estimates.

1.2 State of the art: from image segmentation to the ‘gas of circles’ model

This section gives an overview starting from image segmentation focusing on *variational methods*, to the ‘gas of circles’ higher-order active contour (HOAC) model. We follow the ‘general to specific’ direction, but sometimes diverge if necessary.

The main goal of *image segmentation* models is to divide the image into regions that hopefully corresponds to structural units in the scene or distinguished objects of interest. In the book of Sonka et al. (1999), or in Gonzalez and Woods (2002), we find the fundamentals of image segmentation, such as edge-based models, thresholding, region growing, *etc.* Even in the case of the simplest segmentation models, we can observe the presence of prior knowledge, as described in section 1.1, albeit very simple. The knowledge K is the threshold level in the case of thresholding, the definition of edge in edge detection case, *etc.* The hypothesis space is the set of all possible regions \mathcal{R} , and the best segmentation R_{opt} is one element of this set. To find R_{opt} is usually very easy, because this type of knowledge does not suggest non-local pixel dependencies. In other words, it is easy to classify a given pixel (w.r.t. the model parameters) without any information about nearby pixel values. Classical methods cannot deal very well with high noise, a cluttered background or occlusions, and they strongly depend on the model parameters. As a result the idea of incorporating some prior knowledge about the shape of the region has been considered by many researchers. Early shape priors were quite generic, enforcing some kind of homogeneity and contour smoothness (Geman and Geman, 1984, Blake and Zisserman, 1987, Kass et al., 1988, Mumford and Shah, 1989, Cohen, 1991, Kato et al., 1996, Caselles et al., 1997a). For example, Geman and Geman (1984), Kato et al. (1996), and Szirányi and Zerubia (1997) use a Markovian smoothness prior (basically a Potts model (Potts, 1952, Baxter, 1990)); Blake and Zisserman (1987) use a line process to control the formation of region boundaries; and active contour model (Kass et al., 1988) use elasticity, rigidity, contour length, and balloon or area minimizing forces (Cohen, 1991, Caselles et al., 1997a) in order to favor smooth closed curves. In spite of their simplicity, these methods proved to be very efficient in dealing with noisy images. From the probabilistic point of view, we add extra

knowledge to K . To find R_{opt} is more difficult, usually MAP estimate coupled with various minimization approaches have been used (*e.g.* simulated annealing, gradient descent).

We concentrate on *variational methods* (for a detailed summary see Leymarie (2003), Bresson (2005)) and two very important families within this approach, namely *active contours* (Kass et al., 1988) and the *Mumford-Shah functional* (Mumford and Shah, 1989).

The goal of the Mumford-Shah model is to divide the image into distinct homogeneous regions.¹ To reach the criterion they introduced a functional and its minimization in Mumford and Shah (1985, 1989). The functional gives a piecewise smooth approximation of the image (for a detailed explanation, see (Aubert and Kornprobst, 2002)). The functional is defined as:

$$E_{MS} = \lambda^2 \int_{\Omega} dx (u_0 - u)^2 + \int_{\Omega \setminus \gamma} dx |\nabla u|^2 + \mu |\gamma| ,$$

where $\Omega \subset \mathbb{R}^2$ and u_0 is the original image defined on Ω , u is the piecewise smooth approximation of the image, γ is the boundary between the regions and $|\gamma|$ is its length. λ and μ are the strategic parameters of the functional. The first term considers that the smoothed image should be close to the original image, the second enforces the smoothness, and the third minimizes the contour length. The minimization of the functional is not easy, as explained by Chan and Vese (2002) and Aubert and Kornprobst (2002). The numerical approximation is usually based on the Euler-Lagrange equations, but the length of the contour is not continuous in any compact topology which makes it difficult to use the calculus of variations. Hence, many authors have proposed to approximate the functional by a sequence of regular functionals, *e.g.* Ambrosio and Tortorelli (1990), Chambolle (1995), and Chan and Vese (2001).

Active contour models capture the closest object to their initialization state. They were first introduced by Kass et al. (1988). The model is able to capture sharp image intensity variations by deforming, according to this behaviour sometimes it is called as a *snake* model. The energy function is defined as:

$$E_{AC} = \alpha \int_0^1 ds |\dot{\gamma}(s)|^2 + \beta \int_0^1 ds |\ddot{\gamma}(s)|^2 + E_i(\gamma, u_0) , \quad (1.2.1)$$

where $\gamma(s)$ is a parametrization of the contour $s \in [0, 1]$, and α and β are weighting parameters that control the snake's tension and rigidity. The external energy $E_i(\gamma, u_0)$ is derived from the image so that it takes on smaller values at the features of interest. E_{AC} is a non-convex function, with possibly many local minima, one of them can be reached by solving an Euler-Lagrange equation (Kass et al., 1988). A fast numerical algorithm has been proposed but in the case of closed curves it does not allow changes of topology, since the final curve has the same topology as the initial one. To overcome the limitation of the changes of topology, Osher and Sethian (1988) proposed the level-set method, where the curve γ is implicitly represented by a higher dimensional function (see also in Sethian (1999), and Osher and Paragios (2003)).

¹We note that this requirement is not well defined goal because 'homogenous' is not defined.

One of the main directions along which the active contour idea had been further developed, was introduced by Caselles et al. (1993), and called *geometric active contours*. In Caselles et al. (1997a) they proposed the *geodesic active contours*, a new curve parametrization invariant energy function based on the model in equation 1.2.1:

$$E_{GAC}(\gamma) = \int_0^1 dp f(|\nabla u_0(\gamma(p))|) |\gamma| = \int_0^{|\gamma|} ds f(|\nabla u_0(\gamma(s))|) ,$$

where the function f is the edge function similar to the E_i in equation 1.2.1. Caselles *et al.* also proved that the curve, which minimizes this energy is a geodesic in a Riemannian space. Cohen (1991) introduced a new constant artificial force term into the geodesic active contour energy called, the *balloon force*. The advantages of this model are the possibility of detecting non-convex objects, and faster convergence to the steady state solution. Geodesic active contours were extended to more than two dimensions in Kichenassamy et al. (1995), and Caselles et al. (1997b).

To eliminate the problems of active contour-like models during topological changes, Osher and Sethian (1988) proposed the *level-set* method. The method has been used in many applications from physics to image processing, such as fluid dynamics, control, ray tracing, image restoration, image segmentation, *etc.* (Malladi et al., 1995, Sethian, 1999). The general evolution equation for the problem of curve evolution is given by the partial differential equation (PDE):

$$\begin{cases} \frac{\partial \phi}{\partial t} = |\nabla \phi| \operatorname{div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \\ \phi(0, x) = \phi_0(x) \end{cases} .$$

The evolution governs the motion by the curvature. $\phi(t, x)$, $t \in [0, \infty)$, $x \in \mathbb{R}^2$, where x parameterizes the level-set surface and t parameterizes the evolution. The main idea of the level set method is to implicitly represent an interface γ in \mathbb{R}^n as a level set of a function ϕ . For details of and numerical implementation, see Sethian (1999), Chan and Vese (2001), and Osher and Paragios (2003).

We illustrate the evolution of a curve using the level-set framework in figure 1.1. In the top two rows we present the contour evolution, starting from the left-hand side. The bottom two rows contain the corresponding level-set functions of every second contour (*i.e.* the first surface is the level-set function corresponding to the first contour, while the second one corresponds to the third contour, *etc.*).

Chan and Vese (2001) introduced a new data term of the piecewise constant approximation of the Mumford-Shah functional, using the level-set framework. The energy functional is defined as:

$$E_{CV}(c_1, c_2, R) = \mu L(\partial R) + \nu A(R) + \lambda_1 \int_R dx dy (u_0(x, y) - c_1)^2 + \lambda_2 \int_R dx dy (u_0(x, y) - c_2)^2 ,$$

where R is a region in the image domain and ∂R is its boundary; L and A are the boundary length and the region area functionals respectively. c_1 and c_2 are the average intensity values





Figure 1.1: Evolution of a curve (top) and the related level-set function (bottom). However the curve goes through topological changes, there is no topological change on the level-set surface. The level-set functions correspond to the first, third, fifth and seventh contours respectively.

inside and outside of R , updated dynamically. The parameters control the strength of the terms.

Paragios (2000), and Paragios and Deriche (2002a) proposed the *geodesic active region* model, based on the combination of a classical snake model and the region growing algorithm (Zhu and Yuille, 1996). This model unifies boundary and region based knowledge in a variational and statistical framework. The information depends on a statistical estimation of the image histogram using a mixture of Gaussian distributions, each of them represents one of the desired regions. Paragios and Deriche (2002b) extended the framework to supervised

texture segmentation.

Samson et al. (2000) presented a variational supervised image segmentation technique based on known intensity means and variances for each region to be segmented. They used the multi phase level-set model of Zhao et al. (1996).

A new level-set energy function was proposed by de Rivaz and Kingsbury (2000) for which it is possible to obtain a good estimate for the optimum step size in a gradient descent and which therefore produces much faster convergence. The surface is represented by the coefficients of a complex wavelet transform.

The above presented models have some common properties, they minimize a functional, which is a composition of energies coming from the image and the properties of the contour (length, tension, rigidity, *etc.*). It is possible to introduce higher level knowledge into the geometric part of the energy function as:

$$E_g = E_0 + E_P ,$$

where the prior energy E_P is a functional, which has a minimum if the current segmentation is similar to a given template shape or a mean computed from a set of shapes allowing small variances over the contour. E_0 is one of the classical models described above.

Leventon et al. (2000b) presented a method using PCA to extract and embed the main variation of a training set and remove the redundant information. They represent the shape information by embedding the template shape into the level-set framework. The PCA method is applied not to the curve directly, but to the signed distance function. This approach is robust to slight differences between the template and the segmented shapes. Notice that Cootes and Taylor (1992) used a similar technique for parametric active contours.

Leventon et al. (2000a) defined a joint distribution between intensity value and distance to the boundary, given by a non-parametric density estimation based distance function. Tsai et al. (2001) integrated this model with a region-based energy functional to segment organs in 3-D medical images.

One of the most important methods was introduced by Paragios and Rousson (2002). They defined a new level-set based shape representation and registered the sample shape with the actual segmentation using rigid transformations. They minimized the sum of squared distances. The model is invariant w.r.t. translation, rotation, and scale. Rousson and Paragios (2002) combined the *active shape model* of Cootes and Taylor (1992) with level set methods using the PCA based representation framework of Leventon et al. (2000b). Huang et al. (2004) presented a registration method using distance functions and using a global registration criterion.

Cremers et al. (2001, 2002) introduced *diffusion snakes*. They modified the Mumford-Shah functional and its cartoon limit to segment known object types. They transformed the functional from non-convex to convex using a quadratic shape energy term. They used quadratic B-spline during the numerical implementation. They applied the proposed method with success to delineating noisy or occluded shapes. They also extended the linear shape model to a nonlinear space using a *kernel PCA* method.

Cremers and Soatto (2003) published a technique called *pseudo-distance*. This novel approach represents the shape as a signed distance function. They improved the presented

model to be invariant w.r.t. translation and scale. Rotation invariance is also possible. Cremers et al. (2004, 2006) proposed a novel approach using more than one template shape, and they partially solved one of the main challenges of shape representation methods, which is the segmentation of different known shapes at the same time. They use additional level-set functions to distinguish the different shapes. The number of template shapes increases the computational complexity, even in the case of multiple instances of a fixed object.

Foulonneau et al. (2003) presented a translation and scale invariant *Legendre-moments* based model, they minimize the sum of the square distances between shape moments of a template shape and the actual segmented region. Foulonneau et al. (2006) introduced affine-invariant priors allowing rotation as well. As regularization, they combined the prior term with the Chan and Vese model (Chan and Vese, 2001).

Rousson and Paragios (2007) introduced a level set method for shape-driven object extraction, using a voxelwise probabilistic level set formulation. The objects are represented in an implicit form.

Riklin-Raviv et al. (2005) presented a method invariant to general projective transformations. Riklin-Raviv et al. (2007) introduced a novel shape similarity measure and embed the projective homography between the prior shape and the image to segment within a region-based segmentation functional. They combined their method with the Chan and Vese framework.

Rochery et al. (2003, 2005a, 2006) introduced a method using a new generation of active contours called *higher-order active contours* (HOACs). They proposed a novel energy term composed of multiple integrals over the contour and creating non-local interactions between contour points far from each other. With this energy they are able to describe families of shapes without using templates, and multiple instances of similar shapes do not increase the complexity of the model. They used the proposed model to detect road networks in aerial images. They extended the above model with a more complex prior term to close gaps on the road structure (Rochery et al., 2004). In Rochery et al. (2005b), they introduced a *phase field* model, which is a level-set based modelling technique, well known in physics. They proposed an equivalent higher-order phase field model to the HOAC, leading to faster energy minimization.

1.3 The ‘gas of circles’ model for tree crown detection

In the last few decades computers have been playing an essential part in everyday life, since with the help of these electronic devices work can be carried out more quickly and efficiently, and requires less human interaction. One of the most important areas, in which computers are used, is called image processing: analyzing and evaluating images. Nowadays, thanks to the development of software and hardware components, faster and more accurate image acquisition equipment, advanced algorithms and faster computers, image processing is applied in more and more areas of life. Economically as well as ecologically one of the most significant areas to which image processing is applied is agriculture. In modern agriculture, image processing methods, which are used from microbiological applications to aerial and satellite imagery, are just a few among other things. We will apply

image processing to forestry management.

In fact computer-based methods are widely used by forestry inventories (Shao and Reynolds, 2006). Some of the most important applications are: active remote-sensing; creating digital elevation models; terrain-mapping; radar and LIDAR (Light Detection and Ranging) images on forested areas; forestry information systems for managing huge databases; forest change and growth forecasting and planning; spatial data visualization; computer-aided decision making using artificial intelligence; and last but not least it is possible to analyze forests on the scale of individual trees, by solving the ‘tree crown extraction’ problem, which is the main purpose of this thesis.

Information on the size of individual trees enables forestry managers to study forests at the scale of trees. In this way, it is possible to get a more accurate evaluation of resources, *e.g.* biomass, average tree crown size, number of trees, their size, density of the stands, the changes of population, growth etc. At present there are two ways to gather this type of data: 1.) field surveys, which provides accurate information, though it is time-consuming and very costly. 2.) methods to extract tree crowns from aerial images. Different approaches exist, ranging from very expensive manual segmentation to semi-automatic methods. The semi-automatic methods require user interaction and only a few are able to delineate the shapes of the trees.

In this thesis, we propose a method for extracting tree crowns. We assume that tree crowns are approximately circular-shaped, and are of about the same size. We evaluate our method using Colour Infrared (CIR) and grayscale panchromatic aerial images provided by the French Forest Inventory (IFN) and the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) . A typical image can be seen in figure 1.2. In the image, planted poplar forests can be seen in different sizes and shades.




Figure 1.2: Real image with planted forest ©IFN.

To solve the problem of individual tree crown delineation, we propose a model of a ‘gas of circles’ (GOC) (Horváth et al., 2006a, b), the ensemble of regions in the image domain consisting of an unknown number of circles with approximately fixed radius and short range repulsive interactions. The method is based on the ‘higher-order active contour’ (HOAC) framework, introduced by Rochery et al. (2003, 2005c, 2006), which incorporates long-range interactions between contour points, and thereby include prior geometric information without using a template shape. This makes them ideal when looking for multiple instances of an entity in an image. 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. This is a non-trivial requirement. We impose it by performing a functional Taylor expansion of the HOAC energy around a circle, and then demanding that the first order term be zero for all perturbations, and that the second order term be positive semi-definite. These conditions allow us to fix one of the model parameters in terms of the others, and constrain the rest. The energy is minimized using gradient descent algorithm, and implemented using the level-set method. Experiments using the HOAC energy demonstrate empirically the coherence between these theoretical considerations.

The general ‘gas of circles’ model has many potential applications in varied domains, but it suffers from a drawback: such local minima can trap the gradient descent algorithm used to minimize the energy, thus producing phantom circles even with no supporting data. The model as such is not at fault: an algorithm capable of finding the global minimum would not produce phantom circles. This suggests two approaches to tackling the difficulty. One is to find a better algorithm. The other is to compromise with the existing algorithm by changing the model to avoid the creation of local minima, while keeping intact the prior knowledge contained in the model. We choose this second approach (Horváth et al., 2006c). We solve the problem of phantom circles by calculating, via a Taylor expansion of the energy, parameter values that make the circles into inflection points rather than minima. In addition, we find that this constraint halves the number of model parameters, and severely constrains one of the two that remain, while improving the empirical success of the model.

Although the HOAC ‘gas of circles’ model is an effective tool to model circular shapes, there are some difficulties. It is complicated to express the space of regions in the contour representation, and consequently difficult to work with a probabilistic formulation. In addition, from the algorithmic point of view, the current model does not allow enough topological freedom, and the implementation of the HOAC model is difficult and computationally expensive. But it is possible to create an alternative formulation of HOAC models, based on the ‘phase field’ framework much used in physics to model regions and interfaces. The standard phase field model is, to a good approximation, equivalent to a classical active contour (CAC) model with energy given by boundary length. Rochery et al. (2005b) showed how to extend the basic phase field energy with extra, non-local terms that produce phase field models equivalent to higher-order active contours. The new formulation has several important advantages:

- There is more topological freedom during the gradient descent evolution than with other methods, which is important when the topology is unknown *a priori* as it is

the case in the tree crown extraction problem. In addition, more topological freedom means less chance of becoming stuck in local minima.

- The implementation of the phase field version of CACs, and in particular of HOACs, is much simpler than the equivalent contour or level set implementation. Gradient descent consists of a single partial differential equation derived directly from the model energy with no need for reinitialization or regularization. HOAC terms consist of convolutions and can be evaluated in Fourier space, causing two orders of magnitude increase in computation time.
- Phase field models provide a neutral initialization for gradient descent—no initial region is needed—meaning that bias, caused by the initialization, is reduced.

In Horváth and Jermyn (2007b) we address the tree crown extraction problem by constructing a phase field model of a ‘gas of circles’. We compute, as a function of the HOAC energy parameters, the phase field energy parameters that produce an equivalent model. This means that we can adjust the phase field parameters to ensure stable circles of a given radius also.

We extend the phase field ‘gas of circles’ model in Horváth and Jermyn (2007a) to the case of having an inflection point instead of a minimum in the circle energy at the desired radius. With this model we benefit from all the advantages both of the phase field model and of the model presented in Horváth et al. (2006c).

In Horváth et al. (2006b, c) we also define a data model. This model describes the behaviour of only one, the most significant, infrared band of the three available bands in the CIR images. The model is Gaussian, with the values at different pixels independent, and with different means and variances for tree crowns and the background. While successful, this model, even with the strong region prior, it is not capable of extracting accurately the borders of all the trees. Some trees are simply too similar to the background. To solve this problem in Horváth (2007) we construct a new data model that makes use of all three bands in the CIR images. We study the improvement or otherwise of the extraction results produced by modelling the three bands as independent or as correlated. As we will see, even at the level of maximum likelihood, the inclusion of ‘colour’ information, and in particular, interband correlations, can improve the results, and in conjunction with the region prior, the full model is considerably better than that based on one band alone.

The above models are not restricted to forest management. They can also be applied to detection of other circular objects. A few other examples: in nanotechnology to detect various particles and microarrays in electron-microscopic images; in biology to segment circular cells and molecules, or to detect pollen grains; in medical image processing, to detect 2D circular objects; in remote sensing, in the processing of aerial and satellite images, for meteorological, military, and agricultural management.

1.4 Road map

In chapter 2, we describe general higher-order active contour energies, and introduce the quadratic energy model we use. We present the stability analysis of this energy, to create

a set of stable circles with approximately fixed radius, the so called ‘gas of circles’ model. To demonstrate the prior knowledge contained in the model, and the empirical correctness of the stability analysis, we present experimental results using the new energy. To solve a drawback of the model, which is producing phantom circles even with no supporting image data, we determine the parameters so that the energy has an inflection point rather than a minimum. This constraint also halves the number of parameters and improves the empirical success of the model.

In chapter 3 we describe a novel representation of the HOAC energy, the phase field model. We present that the basic phase field energy is equivalent with the classical active contour energy, and define a new phase field model, which is approximately equivalent with the HOAC model, but offers very important advantages. For the tree crown extraction problem, the computation time is two orders of magnitude faster, provides topological freedom and a less complex implementation. We introduce an algorithm to convert the ‘gas of circles’ HOAC parameters to the phase field model. We examine the energy with an inflection point at the desired radius and present experimental results to confirm its empirical success.

To produce a successful tree crown segmentation model, we couple our prior model with real data in chapter 4. We propose two different energies. First, we describe a Gaussian model taking into account the mean values and variances of the image intensities in the tree crowns and also in the background. This model uses the most representative spectral channel of the images. Our second model uses all three spectral bands of the images. In the second half of the chapter, we present experimental results obtained by combining the prior geometric models with the proposed data models.

In chapter 5, we summarize the contributions and novelties of the thesis. We propose ideas for future work and suggest concepts for tackling unresolved difficulties.

In appendices A and B, we describe the detailed steps of the stability analysis of the HOAC model and the parameter settings for the inflection point model.

Chapter 2

The higher-order active contour ‘gas of circles’ model

We present models of a ‘gas of circles’, the ensemble of regions consisting of an unknown number of circles with approximately the same radius and short-range repulsive interactions. The methods use higher-order active contours (HOACs), which incorporate long-range interactions between contour points, and thereby include prior geometric information without using a template shape. This makes them ideal when looking for multiple instances of an entity in an image. First, we present the existing HOAC model for networks, and show via a stability calculation that circles stable to perturbations are possible for constrained parameter sets. Second, we solve a drawback of the energy function, i.e. that it creates phantom shapes on homogenous areas, choosing the parameters so that without image support the shape vanishes.

THE aim of this chapter is to develop a model to describe an unknown number of similar circles with short-range of inter-circle interaction. We use the recently introduced higher-order active contour framework.

Higher-order active contours (Rochery et al., 2003, 2005c, 2006) introduce long-range interactions between boundary points not via the intermediary of a template region or regions to which the actual segmentation is compared, but directly, by using energy terms that involve multiple integrals over the boundary. The integrands of such integrals thus depend on two or more, perhaps widely separated, boundary points simultaneously, and can thereby impose relations between tuples of points. Euclidean invariance of such energies can be imposed directly on the energy, without the necessity to estimate a transformation between the boundary sought and the template. More importantly, because there is no template, the topology of the region needs not be constrained, a factor that is critical when the topology is not known *a priori*. A short introduction to HOAC models can be found in section 2.1.

The ‘gas of circles’ (GOC) model is a HOAC model favouring 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. This is a non-trivial requirement. We impose it by performing a functional Taylor expansion of the HOAC energy around a circle, and then demanding that the first order term be zero for all perturbations, and that the second order term be positive semi-definite. These conditions allow us to fix one of the model parameters in terms of the others, and constrain the rest. Synthetic experiments using the HOAC energy demonstrate empirically the coherence between these theoretical considerations and the gradient descent algorithm used in practice to minimize the energy. The computation is described in section 2.2.

Although the GOC model is useful in many applications, it suffers from a drawback: such local minima can trap the gradient descent algorithm used to minimize the energy, thus producing phantom circles even with no supporting data. The model as such is not at fault: an algorithm capable of finding the global minimum would not produce phantom circles. This suggests two approaches to tackling the difficulty. One is to find a better algorithm. The other is to compromise with the existing algorithm by changing the model to avoid the creation of local minima, while keeping intact the prior knowledge contained in the model. In section 2.3, we take this second approach. We solve the problem of phantom circles by calculating, via a Taylor expansion of the energy, parameter values that make the circles into inflection points rather than minima. In addition, we find that this constraint halves the number of model parameters, and severely constrains one of the two that remain, while improving the empirical success of the model.

The minimization of classical active contour energies is computationally fast, but difficult to implement, and it is hard to handle changes in contour topology. The level-set framework handles topology change easily, but computationally is more expensive, the contour is represented implicitly and can be extracted meshing the surface at the zero level-set. Most of the time this thresholding approach is sufficient, but in the case of higher-order active contours it is not enough to extract the contour points by meshing, because the quadratic energy is defined on the contour as a closed curve using properties which are not available in this case. Moreover, we need to compute the tangent vectors. The way to achieve this is by contour extraction. Rochery et al. (2005c, 2006) presented an implementation of the HOAC using the level-set framework combined with contour extraction. In section 2.4, we give a short overview of this technique. We also present the detailed numerical approximation used to find the ‘gas of circles’ parameters, described in sections 2.2 and 2.3.

2.1 The higher-order active contour model

As with all active contour models, a region R is represented by its boundary, ∂R . There are various ways to think of the boundary of a region. If the region has only one connected component, which is also simply-connected, then a boundary is an equivalence class of embeddings of the circle S^1 under the action of orientation-preserving diffeomorphisms of S^1 . When more, possibly multiply-connected components are included, however, things get complicated. First, the number of embeddings of S^1 that are required depends on the topology, and second, there are constraints on the orientations of different components if

they are to represent regions with handles.

An alternative is to view ∂R as a closed 1-chain γ in the image domain Ω (Choquet-Bruhat et al. (1996) 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 (Rochery et al., 2005c, 2006). Let γ be a 1-chain in Ω , and $\square\gamma$ be its domain, the energy functional is:

$$E(\gamma) = \iint_{(\square\gamma)^2} dt dt' \dot{\gamma}(t) \cdot F(\gamma(t), \gamma(t')) \cdot \dot{\gamma}(t'), \quad (2.1.1)$$

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

By imposing Euclidean invariance on this term, and adding linear terms, Rochery et al. (2003) defined the following higher-order active contour prior:

$$E_{\mathbb{g}}(\gamma) = \lambda_C L(\gamma) + \alpha_C A(\gamma) - \frac{\beta_C}{2} \iint dt dt' \dot{\gamma}(t) \cdot \dot{\gamma}(t') \Psi(R(t, t')), \quad (2.1.2)$$

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. (2003) used the following interaction function (figure 2.1) Ψ :

$$\Psi(z) = \begin{cases} 1 & z < d - \epsilon, \\ \frac{1}{2} \left(1 - \frac{z-d}{\epsilon} - \frac{1}{\pi} \sin \frac{\pi(z-d)}{\epsilon} \right) & d - \epsilon \leq z < d + \epsilon, \\ 0 & z \geq d + \epsilon. \end{cases} \quad (2.1.3)$$

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

Figure 2.2 illustrates two situations, where contour points are located on the same and on the opposite side of the contour. Using Ψ described in equation 2.1.3, if the points are close enough to each other, then their tangential dot product determines the arising energy between them. In the left image this dot product is strongly positive, while in the right it is negative. This means, that situation on the left part of the image is more favoured than the right one.

For certain ranges of the parameters involved, the energy (2.1.2) 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. (2003, 2005c, 2006). Figure 2.2 gives an explanation on the reason of appearing network shapes. Without the quadratic term the region tries to be as small as possible, both in the sense of its area and boundary length. With the quadratic term, if we have a situation presented in the left side of figure 2.2, there is no energy change coming from the quadratic term if the region becomes smaller. But,




Figure 2.1: The interaction function.




Figure 2.2: Two situations of point pairs. In the left image case the dot product of the tangent vectors is positive, while in the second case (right image) this product is negative, causing opposite effect on the energy.

if the region turns into so small that its opposite boundary is closer than $d + \epsilon$ then the situation illustrated on the right image can be observed. But this situation is not favoured, so the region tends to avoid it. The possible way to do so, is to stop somewhere, where this unfavourable energy is not strong, in other words stop at the equilibrium, where the quadratic energy is as strong as the rest. This equilibrium can be a road network with a given road width or as we will discuss in the next section, it can also be a stable circle, with radius influenced by the interaction function. Two examples can be seen in figure 2.3. In both cases the initial regions were circles, and the minimized energy was $E_g(\gamma)$ with different parameters. In the lower row we can observe a more dense network system, with smaller width ($d = 3$), meanwhile in the top row the width is bigger ($d = 7$), but the density of the arms is less.

The energy E_g thus provides a good prior for network extraction from images. Rochery et al. (2003) used the proposed model to detect road structures in remotely-sensed, single-




Figure 2.3: Examples of gradient descent using the HOAC energy with different parameter settings. Left: Initial region, three rightmost: evolution after 100, 500, and 1000 iterations respectively (Rochery et al., 2005c).



Figure 2.4: Road detection on a SPOT satellite image. Left: Original image; four rightmost: evolution. Image ©French Space Agency (CNES) (Rochery et al., 2005c).

band images. An example can be seen in figure 2.4. In the high resolution SPOT image ©CNES the roads are perfectly extracted. The data term was based on the gradient of the image.

2.2 The ‘gas of circles’ model

The creation of network structures 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. To illustrate this, figure 2.5 shows the evolution of the contour using different β_C parameters, the initial region is a rounded rectangle slightly smaller than the image, $\alpha_C = 5.8$, and $d = \epsilon = 15$. In the first row, the contour completely vanishes. This is because the quadratic parameter is not strong enough to compensate the thinning effect of the area parameter. The second row illustrates the behaviour with a relatively strong quadratic term. The contour

grows arms. In the last row, we can observe the most interesting evolution related to our aims, the contour forms stable circles with the same radius. The β_C parameter was slightly stronger than in the first case. We analyze the energy creating a ‘gas of circles’.




Figure 2.5: Curve evolution using different β_C parameters ($\alpha_C = 5.8$, $d = \epsilon = 15$). Left: Initial region, rightmost images: curve evolution. For different parameter settings we find different contour behaviour. First row: the contour completely vanishes ($\beta_C = 0.5$). Second row: The contour grows arms ($\beta_C = 3.8$). Third row: The final regions are stable circles ($\beta_C = 0.8$).

For this to work, a circle of the given radius, hereafter denoted r_0 , must be stable, that is, it must be a local minimum of the energy. In section 2.2.1, we conduct a stability analysis of a circle, and discover 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 subsection 2.2.2, we present numerical experiments using E_g that confirm the results of this analysis.

2.2.1 Stability analysis

We want to expand the energy E_g around a circle of radius r_0 . We denote a member of the equivalence class of maps representing the 1-chain defining the circle by γ_0 . The energy E_g is invariant to diffeomorphisms of $\square\gamma_0$, and thus is well-defined on 1-chains. To second order:

$$E_g(\gamma) = E_g(\gamma_0 + \delta\gamma) = E_g(\gamma_0) + \langle \delta\gamma | \frac{\delta E_g}{\delta\gamma} \rangle_{\gamma_0} + \frac{1}{2} \langle \delta\gamma | \frac{\delta^2 E_g}{\delta\gamma^2} | \delta\gamma \rangle_{\gamma_0}, \quad (2.2.1)$$

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

Since γ_0 represents a circle, the easiest is 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 $r_0(t) = r_0$, $\theta(t) = t$, and $t \in [-\pi, \pi)$. We are interested in the behaviour of small perturbations $\delta\gamma = (\delta r, \delta\theta)$. The first thing to notice is that 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 thus turns out to be easiest to perform the calculation by expressing δr in terms of this basis:

$$\delta r(t) = \sum_k a_k e^{i r_0 k t}, \quad (2.2.2)$$

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 equation (2.1.2). Details can be found in appendix A.

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

$$L(\gamma) = \int_{-\pi}^{\pi} dt |\dot{\gamma}(t)| = 2\pi r_0 \left\{ 1 + \frac{a_0}{r_0} + \frac{1}{2} \sum_k k^2 |a_k|^2 \right\} \quad (2.2.3)$$

$$A(\gamma) = \int_{-\pi}^{\pi} d\theta \int_0^{r(\theta)} dr' r' = \pi r_0^2 + 2\pi r_0 a_0 + \pi \sum_k |a_k|^2. \quad (2.2.4)$$

Note the k^2 in the second order term for L . This is the same frequency dependence as the Laplacian, and shows that the length term plays a similar smoothing role for boundary perturbations as the Laplacian does for functions. In the area term, by contrast, the Fourier perturbations are ‘white noise’.

It is also worth mentioning, 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 , $\lambda_C r_0 k^2 + \alpha_C \geq 0$. Note that 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 $m^2 - 1 \geq 0$. Two points are worth noting. The first is the one we have already made: the zero frequency perturbation is not stable. Zero frequency means changes of radius, *i.e.* the energy thus defined is the maximum w.r.t. changes of radius, and hence is not stable. The second is that the $m = 1$ perturbation is marginally stable to second order, that is, such changes require no energy to this order. To fully analyse them, we must therefore go to higher order in the Taylor series.

The quadratic term can be expressed to second order as

$$\begin{aligned}
\iint_{-\pi}^{\pi} dt dt' \dot{\gamma}(t') \cdot \dot{\gamma}(t) \Psi(R(t, t')) &= 2\pi \int_{-\pi}^{\pi} dp F_{00}(p) + 4\pi a_0 \int_{-\pi}^{\pi} dp F_{10}(p) \\
&+ \sum_k 2\pi |a_k|^2 \left\{ \left[2 \int_{-\pi}^{\pi} dp F_{20}(p) + \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{21}(p) \right] \right. \\
&\left. - \left[2ir_0k \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{23}(p) \right] + \left[r_0^2 k^2 \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{24}(p) \right] \right\}, \quad (2.2.5)
\end{aligned}$$

The F_{ij} are functionals of Ψ (hence functions of d and ϵ for Ψ given by equation (2.1.3)), and functions of r_0 , as well as functions of the dummy variable p .

Combining equations (2.2.3), (2.2.4), and (2.2.5), we find the energy functional (2.1.2) up to the second order:

$$\begin{aligned}
E_g(\gamma_0 + \delta\gamma) &= E_0(r_0) + a_0 E_1(r_0) + \frac{1}{2} \sum_k |a_k|^2 E_2(k, r_0) \\
&= \left\{ 2\pi\lambda_C r_0 + \pi\alpha_C r_0^2 - \pi\beta_C G_{00}(r_0) \right\} \\
&\quad + a_0 \left\{ 2\pi\lambda_C + 2\pi\alpha_C r_0 - 2\pi\beta_C G_{10}(r_0) \right\} \\
&\quad + \frac{1}{2} \sum_k |a_k|^2 \left\{ 2\pi\lambda_C r_0 k^2 + 2\pi\alpha_C \right. \\
&\quad \left. - 2\pi\beta_C [2G_{20}(r_0) + G_{21}(k, r_0)] \right. \\
&\quad \left. - 2ir_0k G_{23}(k, r_0) + r_0^2 k^2 G_{24}(k, r_0) \right\}, \quad (2.2.6)
\end{aligned}$$

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

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 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:

$$\beta_C(\lambda_C, \alpha_C, \hat{r}_0) = \frac{\lambda_C + \alpha_C \hat{r}_0}{G_{10}(\hat{r}_0)}, \quad (2.2.7)$$

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 (2.2.1). Figure 2.6 left shows a typical plot of the energy E_0 of a circle versus its radius r_0 , with the β_C parameter fixed using the equation (2.2.7) 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 from figure 2.6 left, 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 figure 2.7 (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 equation (2.1.3). If $2r_0 < d - \epsilon$, the quadratic term will be constant, and no force will exist to stabilize the circle. In order to use equation (2.2.7) then, we have to ensure that we are on the upper branch of figure 2.7.




Figure 2.6: 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 equation (2.2.7) with $\hat{r}_0 = 1.0$. (The parameters of Ψ are $d = 1.0$ and $\epsilon = 1.0$, 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 as in the left plot. The function is non-negative for all frequencies.

Equation (2.2.7) 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 . Figure 2.6 right shows a plot of $E_2(k, \hat{r}_0)$ against $\hat{r}_0 k$ for the same parameter values used in figure 2.6 left, 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




Figure 2.7: Schematic plot of the positions of the extrema of the energy of a circle versus β_C , for α_C fixed.

above criteria, the ‘gas of circles’ model.

2.2.2 Geometric and synthetic experiments

In order to illustrate the behaviour of the prior energy E_g with parameter values fixed according to the above analysis, in this section we show the results of some experiments using it. In the first case we show examples using only the geometric term, while in the rest we use image terms as well, computed using pixel intensity values.

The geometric energy

Figure 2.8 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 section 2.4 or Rochery et al. (2005c, 2006), Rochery (2005).) 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 figure 2.6 left. 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 and shape 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 .




Figure 2.8: 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.

Noisy synthetic images

Here, we present the results of tests of the sensitivity of the model to noise in the image. The energy used is $E = E_g + E_i$, where E_i is the energy coming from the image, based on the pixel intensity values and the image gradient. We model the background and foreground with Gaussian distributions and fix the parameters by learning them. The region, which minimizes the image term has pixels close to the mean foreground intensity and its boundary is on places where the gradient is high (for details see section 4.2.1). Fifty synthetic images were created, each with ten circles with radius 8 pixels and ten circles with radius 3.5 pixels, placed at random but with overlaps rejected. Six different levels of white Gaussian noise,

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

Table 2.1: Results on synthetic noisy images. FP, FN, J: percentages of false positive, false negative, and joined circle detections respectively, with respect to the potential total number of correct detections.

with signal-to-noise ratios from -5 dB to 20 dB¹, 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_{in} , σ_{in} , μ_{out} , and σ_{out} , the mean and variance values on the foreground and background regions. The model used was the same as that used for the aerial images, except that the image gradient parameter 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 fifty images are shown in figure 2.9. Table 2.1 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 figure 2.9). 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 ten 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 chapter 5.

Circle separation: comparison to classical active contours

In our second 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

¹Noise is usually measured in Signal-to-noise ratio (SNR), its unit is the decibel (dB) (Kato).

SNR (dB) = $10 \log_{10} \left(\frac{\sigma_{image}^2}{\sigma_{noise}^2} \right)$, where σ_{image} and σ_{noise} are the variances of the image and noise respectively.




Figure 2.9: 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 $(\mu_{in}, \sigma_{in}, \mu_{out}, \sigma_{out}, \lambda_C, \alpha_C, \beta_C)$ are shown under the six images. The parameters d and r_0 were fixed to 8 throughout.

‘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 image term E_i we used during this experiment is only the image gradient.

The first row of figure 2.10 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 HOAC ‘gas of circles’ models to these images.




Figure 2.10: Results on circle separation comparing the HOAC ‘gas of circles’ model to the classical active contour model. Top: the 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 HOAC ‘gas of circles’ model (2, 1, 5, 4.0, 8, 8). All the circles are segmented correctly. Parameter values are shown in the form $(\lambda_i, \lambda_C, \alpha_C)$ and $(\lambda_i, \lambda_C, \alpha_C, \beta_C, d, r_0)$ respectively.

The middle row of figure 2.10 shows the best results obtained using the classical active contour model. The model was either unable to 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 HOAC ‘gas of circles’ model can be seen in the bottom row of figure 2.10. All the circles were segmented correctly, independent of the gray level of the junction. Encouraging as this is, it is not the whole story, as we indicated in section 2.2.2. We make a further comment on this issue in chapter 5.

2.3 The inflection point ‘gas of circles’ model

Figure 2.6 left shows a plot of the energy of a circle versus radius for parameter values selected according to the energy minimum criteria at the desired radius. This curve has just the form we require: circles of radius r_0 are metastable (*i.e.* local minima), with an energy that is low but nevertheless higher than that of the empty region. In the absence of supporting data, the global minimum will thus be the empty region, the correct behaviour. A gradient descent algorithm, however, cannot escape from these local minima, meaning that circles of radius r_0 , once formed during gradient descent, cannot disappear, even if the data does not support their existence. In practice such circles sometimes do form, which is clearly undesirable. The best solution to this problem would be an algorithm capable of finding the global minimum of the energy. A slightly less ambitious approach, which we take here, involves making a compromise with the algorithm, changing the model to avoid the creation of these local minima, while preserving as much of the prior knowledge as possible.

The idea we will pursue is to adjust the parameters so that the minimum of the curve on the left of figure 2.6 is replaced by a broad, approximately flat area, as shown in the plots in figure 2.13. Such an energy means that in the absence of image data, a circle will shrink and disappear, whereas small amounts of image data will be sufficient to create a minimum in the flat area, thus producing a stable circle. The natural method to achieve such a broad flat region is to create an energy function that has a single inflection point. If necessary the parameters can then be tweaked to ensure that the gradient of energy w.r.t. radius is positive rather than simply non-negative. It is, however, a nontrivial exercise to find parameter values that result in inflection points.

In figure 2.11, we present a small piece of real aerial image and the segmentation results to demonstrate the importance of our novel model. On the left side we can see the original image. The image contains a poplar plantation and a field in the right top corner. The energy we minimize is $E = E_g + E_i$, where E_i is the data energy described in more details in section 4.2.1. The image is difficult because the field has a homogenous intensity, very similar to the tree crowns. The ‘gas of circles’ HOAC model creates phantom objects in this area, as can be seen in the middle, while the new model using an inflection point provides a better result without phantom shapes (right-hand image).

We address this problem via further analysis of the energy 2.1.2. We still require that a circle of radius r_0 be stable to sinusoidal perturbations with $k > 0$, but now we also require that such a circle be an inflection point with respect to perturbations with $k = 0$, that is,

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




Figure 2.11: Regularly planted poplars ©French Forest Inventory (IFN) (left). Result with the ‘gas of circles’ model in section 2.2 (middle). Result with the monotonic energy function (right).

changes of radius. We will see that these demands are sufficient to fix the prior energy E_g up to an overall multiplicative constant and a small range of values for d . More precisely, we still require that $E_1(r_0) = 0$ and $E_2(k, r_0) > 0$ for $k > 0$, but we now require that $E_2(0, r_0) = 0$ too. The first condition gives equation 2.2.7. The second condition, which follows from the last term of equation 2.2.6, also relates α and β :

$$\alpha(r_0) = \beta(r_0)\tilde{G}(r_0) , \quad (2.3.1)$$

where $\tilde{G}(r_0) = 2G_{20}(r_0) + G_{21}(0, r_0)$. We can solve equations 2.2.7 and 2.3.1 for α and β , giving:

$$\alpha(r_0) = \frac{\lambda\tilde{G}(r_0)}{G_{10}(r_0) - r_0\tilde{G}(r_0)} \quad \text{and} \quad \beta(r_0) = \frac{\lambda}{G_{10}(r_0) - r_0\tilde{G}(r_0)} . \quad (2.3.2)$$

These equations fix α and β as functions of r_0 and d . Since r_0 is fixed by the application, the only remaining parametric degrees of freedom are the value of d , and the overall strength of the prior term, represented by λ . Recall, however, that we also require α and β to be positive. The question is then how to find values of d for a given r_0 so that $\alpha(r_0) > 0$ and $\beta(r_0) > 0$.

Determination of d

To illustrate the behaviour we want to understand, figure 2.12 shows plots of α and β against d for fixed r_0 , in this case $r_0 = 5$. There are two critical points, d_{\min} and d_{\max} . Only for the range $d_{\min} < d < d_{\max}$ are both α and β positive. Our goal is therefore to find d_{\min} and d_{\max} as functions of r_0 .

From equations (2.3.2), it can be seen that d_{\max} arises from a zero in the denominator, while d_{\min} arises from a zero in the numerator. It is, therefore sufficient to find these zeros in order to find d_{\min} and d_{\max} . To proceed, we first note a scaling property of G_{00} . The function G_{00} is given by the following integral (see Appendix A and Horváth et al. (2006b)):

$$G_{00}(r) = \int_{-\pi}^{\pi} dp \cos(p) r^2 \Psi\left(2r \left|\sin \frac{p}{2}\right|\right) . \quad (2.3.3)$$




Figure 2.12: Top left: α plotted against d ; Top right: β plotted against d ; Bottom: enlarged plot of α near the critical domain. ($r_0 = 5.0$ throughout)

Since $\Psi(z)$ is a function of z/d only, by pulling d^2 out of the integral we can write G_{00} as $G_{00}(r) = d^2 \hat{G}_{00}(r/d)$. Now recall that $G_{10} = \frac{1}{2} \partial_r G_{00}$ and $\tilde{G} = \partial_r G_{10}$ (Horváth et al., 2006c). We then find that:

$$\tilde{G}(r_0) = \hat{\tilde{G}}(r_0/d) \quad \text{and} \quad G_{10}(r_0) - r_0 \tilde{G}(r_0) = d(\hat{G}_{10}(r_0/d) - \frac{r_0}{d} \hat{\tilde{G}}(r_0/d)), \quad (2.3.4)$$

where $\hat{G}_{10}(z) = \frac{1}{2} \partial_z \hat{G}_{00}(z)$ and $\hat{\tilde{G}}(z) = \partial_z \hat{G}_{10}(z)$. Thus both numerator and denominator of equations (2.3.2) can be written, up to multiplication by positive coefficients, as functions of r_0/d . Now, $f(r, d) = \hat{f}(r/d)$ and $f(r, d_0) = 0$ imply $f(ar, ad_0) = 0$ for all $a \in \mathbb{R}$; thus if we determine d_{\min} and d_{\max} for one value of r_0 , we know their values for any r_0 .

To determine d_{\min} and d_{\max} while avoiding iterative numerical procedures to find these points, we use a polynomial approximation to G_{00} (for more detailed computations see appendix B):

$$G_{00}(r) = \sum_{n=0}^{\infty} b_n r^n.$$

It is easy to show that:

$$b_m = \begin{cases} 0 & m < 2, \\ \frac{1}{(m-2)!} \int_{-\pi}^{\pi} dp \cos(p) Y^{(m-2)}(0) & m \geq 2, \end{cases} \quad (2.3.5)$$

where $Y(r) = \Psi(2r|\sin(p/2)|)$. The derivatives of Y evaluated at zero are:

$$\frac{Y^{(m)}(0)}{(2|\sin(p/2)|)^m} = \Psi^{(m)}(0) = \begin{cases} 1 & m = 0, \\ 0 & m = 1 \text{ or } m \text{ even}, \\ (-1)^{\frac{m-1}{2}} \frac{1}{2d} \left(\frac{\pi}{d}\right)^{m-1} & m \geq 3 \text{ and } m \text{ odd}. \end{cases}$$

Substituting into equation (2.3.5) gives b_m :

$$b_m = \begin{cases} 0 & m < 5 \text{ or } m \text{ even}, \\ (-1)^{\frac{m-1}{2}} \frac{4(2\pi)^{m-3}}{m!!(m-4)!!} \frac{1}{d^{m-2}} & m \geq 5 \text{ and } m \text{ odd}. \end{cases}$$

We can then derive expressions for \tilde{G} and $G_{10} - r\tilde{G}$:

$$\tilde{G}(r) = 2 \sum_{\substack{m \geq 3 \\ m \text{ odd}}} \frac{(-1)^{\frac{m+1}{2}} (2\pi)^{m-1} (m+1)}{m!!(m-2)!!} \left(\frac{r}{d}\right)^m$$

$$G_{10}(r) - r\tilde{G}(r) = 2d \sum_{\substack{m \geq 4 \\ m \text{ even}}} \frac{(-1)^{\frac{m-2}{2}} (2\pi)^{m-2}}{[(m-3)!!]^2} \left(\frac{r}{d}\right)^m.$$

We computed the roots of these polynomials including terms up to $m = 49$. The smallest positive roots furnish the values of d_{\min} and d_{\max} . The result is that $d_{\min} \simeq 1.2776r_0$ and $d_{\max} \simeq 1.4499r_0$. The graphs in figure 2.13 show plots of E_0 against r for $r_0 = 5$, with d values chosen from the domain $d_{\min} < d < d_{\max}$. From left to right d is increasing. This influences the shape of the energy function, but the inflection point is correctly located at the desired radius in each case.




Figure 2.13: Plot of E_0 against r for $r_0 = 5.0$ with α and β determined by equations (2.3.2). First, $d = 6.4$; second, $d = 6.8$; third, $d = 7.2$. For this value of r_0 , $d_{\min} = 6.3880$, $d_{\max} = 7.2495$.

2.4 Minimization of the energy, implementation issues

We minimize the geometric energy using gradient descent. In section 2.4.1 we approach the process from a more theoretical point of view. More precisely we present the descent equation and give a short overview on the level-set framework. Section 2.4.2 contains the implementation details proposed and used by Rochery et al. (2005c). to realize the HOAC model. This approach is based on a method using both contour extraction as well as level-set functions. Finally, in section 2.4.3 the implementation of the ‘gas of circles’ parameter computations and the inflection point parameters are discussed.

2.4.1 Minimization of the higher-order active contour energy

The energy is minimized using gradient descent. Thus, the contour evolution is determined by:

$$\partial_t \gamma(t) = -\frac{\delta E}{\delta \gamma(t)}, \quad (2.4.1)$$

where $\delta E/\delta \gamma$ is the functional derivative of the energy with respect to the contour. The descent equation can be written as:

$$\hat{\mathbf{n}} \cdot \partial_t \gamma(t) = \kappa(t) - \alpha_C + 2\beta_C \int dt' \hat{r}(t, t') \cdot \mathbf{n}(t') \Psi'(r(t, t')), \quad (2.4.2)$$

where $\hat{r}(t, t') = (\gamma(t) - \gamma(t'))/|\gamma(t) - \gamma(t')|$, \mathbf{n} is the outward pointing normal vector and $\hat{\mathbf{n}}$ is its normalized version.

Although, the linear space is useful for constructing functionals, the minimization of the energy takes place over the space of region boundaries, not over the space of closed 1-chains. To minimize the energy we use gradient descent, but we embed the contour into a higher dimensional space using the level-set framework (Osher and Sethian, 1988) to evolve it. If the contour propagates along its outward normal direction with speed F , namely $\hat{\mathbf{n}} \cdot \partial_t C(t) = F[C](t)$, then the changes in the level-set function can be written as:

$$\partial_t \phi = -\nabla \phi \cdot F \hat{\mathbf{n}} = F \nabla \phi \cdot \phi / |\nabla \phi| = F |\nabla \phi|, \quad (2.4.3)$$

During the evolution ϕ appears as a signed distance function, which is usually initialized in this way, but during the evolution it is difficult to guarantee. However, since the exact evolution of ϕ off the contour is of no consequence provided it is well enough behaved, it can be chosen for convenience, reinitializing every so often if necessary to restore the signed distance function. A typical choice is to apply the expression for F to each level-set, and evolve the function ϕ accordingly. However, as can be seen from equation 2.4.2, the evolution equations derived from quadratic energies contain nonlocal terms, and this creates new difficulties. Following the procedure of applying the expression for F to every level-set is impractical, since it means extracting the level-set belonging to each point of the discretized version of Ω and integrating over it. In order to construct the speed at all points of Ω from the speed on the contour, we therefore use the technique of ‘extension velocities’ (Adalsteinsson and Sethian, 1999).

Figure 2.14: Contour tracing algorithm.

-
1. Choose a starting point A in a set of points R . Set current point $C = A$ and search direction $S = 6$.
 2. While C is different from A or $first = 1$, do steps 3 to 9.
 3. $found = 0$.
 4. While $found = 0$, do steps 5 to 8, at most 3 times.
 5. If B , the neighbour $(S - 1)$ of C is in R ; $C = B$, $S = S - 2$, $found = 1$.
 6. Else, if B , the neighbour S of C is in R , $C = B$ and $found = 1$.
 7. Else, if B , the neighbour $(S + 1)$ of C is in R , $C = B$ and $found = 1$.
 8. Else $S = S + 2$.
 9. $first = 0$.
-

2.4.2 Implementation of the higher-order active contour evolution

The evolution of the level-set is done in four steps. The first step is the reinitialization of the variable ϕ , then we compute the zero level-set and the speed of the contour. The third step is to extend the speed of the contour to Ω . Finally we update ϕ . In this section we briefly overview the drawbacks and crucial steps of the implementation.

To (re)initialize ϕ as a signed distance function an approach from Sussman et al. (1994) was used. This classical method caused a drawback, namely area loss, which was solved in Sussman and Fatemi (1997).

One of the most important steps of the implementation is the very precise contour extraction. To find the exact intersections of the level-set surface with the grid, Rochery et al. (2005c) used ENO interpolation by Siddiqi et al. (1997). After interpolation, the boundary is extracted using the contour tracking algorithm shown in figure 2.14 by Pavlidis (1982). At each step, one starts from the current point and considers six possible directions for the next point. Figure 2.15 shows these directions are adapted to the different possible configurations. We obtain an ordered set of points $\{C(p_i); i = 1, \dots, n\}$ representing the boundary.

The bottom right image of figure 2.15 shows an example that some of the configurations are ambiguous, creating complicated situations. To deal with these, it is necessary to adopt a convention: either the interior or the exterior, but not both, can have subcellular width. We choose the former.

Having extracted the boundary, and after interpolating the necessary values from the grid, we compute the speed F for each extracted point by performing a numerical integration over the contour.

We need to extend the speed to the domain Ω . Rochery et al. (2005c) used the ‘extension velocities’ method of Adalsteinsson and Sethian (1999).

In practice, it is not necessary to compute the evolution of the level-set function over the whole of Ω . Computational efficiency can be increased by restricting computation to a band around the zero level-set, known as the ‘Narrow Band’ (Sethian, 1999), defined by




Figure 2.15: Top row and bottom left: configurations encountered in the contour tracking algorithm. Bottom right: an ambiguous configuration. Figure courtesy of M. Rochery (Rochery, 2005).

$|\phi(x, y)| < t$, where t is a threshold. When the zero level-set comes too close to the edge of the Narrow Band, the level-set function is reinitialized as described above, and the Narrow Band is reconstructed.

2.4.3 Implementation of the parameter estimations

Determination of β_C parameter:

In section 2.2 we described the ‘gas of circles’ model, which fixes one of the HOAC parameters, assuming that the rest are given. We recall equation 2.2.7:

$$\beta_C(\lambda_C, \alpha_C, r_0) = \frac{\lambda_C + \alpha_C r_0}{G_{10}(r_0)}, \quad (2.4.4)$$

i.e., α_C , λ_C , and r_0 are given, and we determine β_C . All but $G_{10}(r_0)$ are known. To determine the value of the integral we used numerical approximation. We applied the recursive adaptive Simpson quadratic form to integrate, and we chose the accuracy of the approximation to be 10^{-6} . Since the function is symmetric, it is sufficient to evaluate the integral on the $[0, \pi]$ domain and double the value. Using the above described parameter values one can successfully create a ‘gas of circles’. Going further, determining the zeroth and second order energies, we can follow the described procedure and approximate the G_{00} , G_{20} , G_{21} , G_{23} , and G_{24} functionals.

The inflection point parameters:

We need to compute the right-hand sides of equations 2.3.2:

$$\alpha(r_0) = \frac{\lambda \tilde{G}(r_0)}{G_{10}(r_0) - r_0 \tilde{G}(r_0)} \quad \text{and} \quad \beta(r_0) = \frac{\lambda}{G_{10}(r_0) - r_0 \tilde{G}(r_0)}. \quad (2.4.5)$$

To compute the value of \tilde{G} , we use the same scheme as in the G_{10} case.

Determination of d :

In practice to determine the zero crossings of the α and β functions, we created polynomials up to a given order, and used the classical root finder method of MATLAB based on companion matrix. In our experiments we used 49th order polynomials, which provided us with a very accurate approximation.

Chapter 3

The phase field ‘gas of circles’ model

Phase-field models are widely used in physics in solving diffusion equations for heat and solute without explicitly tracking the liquid-solid interface. Considering phase fields as a level-set framework, it is possible to create phase field models equivalent to the higher-order active contour models discussed in previous chapter. Phase fields provide more topological freedom, easier and faster implementation, easier probabilistic formulation and on the other hand all the opportunities of the HOAC model. In this chapter, we present how to transform the ‘gas of circles’ parameters into the phase field model. We then describe a phase field version of the inflection point model.

PHASE field models are well-studied and popular in physics during the last decades (Boettinger et al., 2002, Sun and Beckermann, 2007). During the solidification of different materials, they create different sometimes complicated and practically important shapes. The analysis of these shapes was the interest of many researchers in the recent years. The problem of accurately calculating the solid-liquid interface, using phase fields, became popular due to the increasing power of computers. The phase field variable is a function of time and position which describes whether the material is solid or liquid according to a threshold value. The behavior of the variable is determined by an energy function which describes the heat and solute transport. The interface domain between the solid and liquid states is a smooth but highly localized change of the variable. The non-interface regions are characterized by fixed values of the variable, (usually 0 and 1; or -1 and +1 as solid and liquid stages respectively). This approach emerges from the mathematically difficult problem of applying boundary conditions at an interface whose location is part of the unknown solution (the so-called Stefan problem, see Gupta (2003) for details). No boundary conditions are required and, as discussed above, the location of the interface is determined by thresholding the phase field variable (this threshold value is usually 0 or 1/2 but sometimes determined by the energy function’s parameters). The method is also powerful because it treats topology changes easily, unlike the methods dealing with contours. Therefore, phase fields are widely used in physics for modeling dendritic-growth in pure materials, dendritic-, eutectic-,

and peritectic-growth in alloys, solute trapping during rapid solidification, *etc.*

Although, the model has the above mentioned advantages its, application in image processing is not frequent yet. It was used for image inpainting (Grossauer and Scherzer, 2003), and to detect codimension 2 objects in the image domain (Aubert et al., 2004), while in Samson et al. (2000), a model with multiple minima was used with Gamma convergence to construct piecewise constant approximations to an image. In this chapter we present a phase field model equivalent to classical active contours, and then extend it to higher-order active contours. The reason for being interested in an equivalent higher-order active contour model is that the phase field formulation has the following advantages:

- Phase field models provide a neutral initialization for gradient descent, no initial region is needed.
- The implementation of the phase field version of classical active contours, and in particular of HOACs, is much simpler than the equivalent contour or level-set implementation. Gradient descent consists of a single partial differential equation derived directly from the model energy with no need for reinitialization or regularization. HOAC terms consist of convolutions and can be evaluated in Fourier space.
- There is more topological freedom during the gradient descent evolution than with other methods, which is important when the topology is unknown *a priori* as is the case in the tree crown extraction problem. In addition, more topological freedom means less occurrence of becoming stuck in local minima.
- The computation time does not depend on the complexity of the region, and for the tree crown extraction problem, is two order of magnitude faster for HOACs than the level-set implementation.

Rochery et al. (2005b) applied the phase field version of their HOAC network model to the road network extraction problem. In Horváth and Jermyn (2007a, b) we extended Rochery's phase field model to our 'gas of circles' model, and applied it to tree crown extraction.

In section 3.1, we present a framework based on phase field. We present how to define an equivalent active contour model minimizing the area and the length of the segmented objects. Defining the normal vector of the contour as the normal of the phase field surface at the threshold level, we also describe the higher-order phase field model, which is approximately equivalent with the HOAC model.

Rochery et al. (2005b) defined the transformation of the phase field parameters into the HOAC parameters. In the previous chapter we described the parameter settings of the 'gas of circles' model. We give an algorithm to convert the existing HOAC 'gas of circles' parameters into higher-order phase field parameters in section 3.2.

Section 3.3 presents the settings of the phase field parameters which create an inflection point in the energy function at the desired radius. We justify the accuracy of the inflection point 'gas of circles' model with synthetic examples.

In section 3.4, we describe how to minimize the geometric phase field energy and discuss implementation issues.

3.1 The higher-order phase field model

The phase field function ϕ is a level-set function on Ω . For a given threshold z , ϕ defines region in the space of all possible regions \mathcal{R} with a map $\zeta_z(\phi) \in \mathcal{R} = \{x \in \Omega : \phi(x) > z\}$. In contrast to many other segmentation methods, which use representation spaces isomorphic to \mathcal{R} (e.g. snakes, distance functions), there is no constraint on ϕ : it can be any element of a linear function space, and thus many functions correspond to the same region. One of the simplest phase field energies can be formulated as:

$$E_0(\phi) = \int_{\Omega} d^2x \left\{ \frac{D}{2} \partial\phi \cdot \partial\phi + \underbrace{\lambda \left(\frac{1}{4} \phi^4 - \frac{1}{2} \phi^2 \right) + \alpha \left(\phi - \frac{1}{3} \phi^3 \right)}_{V(\phi)} \right\}. \quad (3.1.1)$$

We call the sum of the last two terms $V(\phi)$. If $\alpha = 0$, the model is known as the Ginzburg-Landau model (Ginzburg and Landau, 1950). If $\lambda > |\alpha|$ the V term of $E_0(\phi)$ has two minima, at ± 1 . For positive α values, the lower of the two minima is at -1 . We will show that using this configuration we can give a good estimation for the region boundary and the area of the interior. One of these possible V functions can be seen in figure 3.1.




Figure 3.1: The V term of equation 3.1.1 for $\lambda = 10$ and $\alpha = 2$.

Without the gradient term, minimizing the energy E_0 while keeping the region R fixed would create a discontinuity in ϕ at ∂R with values -1 and $+1$ on the exterior and interior areas respectively. However, with the gradient term, the minimized ϕ function has a soft transition region instead of sharp jump at the boundary of R . The variable ϕ defines a region energy which approximates a linear combination of the region's interior area $A(R)$ and the length of the boundary $L(R)$. The usage of this connection is novel in image processing, first having been introduced by Samson et al. (2000) and Rochery et al. (2005b). However, we notice that for $\lambda \rightarrow \infty$ and on a discrete grid this is known as the Ising model (Ising,

1925). We briefly overview the connection between a simple active contour model $E_{C,0}(R)$, minimizing the area and the boundary length, and $E_0(R)$.

Rochery et al. (2005b) used a simplified model ($\tilde{\phi} \approx \phi$) to show the equivalence of the models and the relation between the parameters. Divide the domain into three pieces: R_γ , R_+ , and R_- . R_γ is a narrow interface region around ∂R characterized by its width w . The interior region $R_+ = R \setminus R_\gamma$, and the exterior region $R_- = \bar{R} \setminus R_\gamma$. The approximated phase field function is ± 1 in R_\pm respectively, while on R_γ it changes between ± 1 . The described model $\tilde{\phi}$ can be seen in figure 3.2.




Figure 3.2: Simplified model to approximate the parameters.

If we consider a very small arc on ∂R , we can suppose that ϕ changes linearly and can estimate $\tilde{\phi}$ in a polar coordinate system as¹:

$$\tilde{\phi}(r, \theta) = -\frac{2}{w}(r - r_0) + z . \quad (3.1.2)$$

If one substitutes this form back into equation 3.1.1, one finds that the energy, up to an additive constant, is:

$$E_0(\tilde{\phi}) = \int_0^{L(R)} ds \left\{ \frac{2D}{w} + w \left(\frac{2\lambda K}{15} + \frac{3\alpha}{40} w \kappa \right) \right\} + \frac{4\alpha}{3} A(R) , \quad (3.1.3)$$

where the integral is parameterized by the arc length and $K = 1 + 5(\alpha/\lambda)^2$. Minimizing the energy w.r.t. the width of the interface w , the optimal width w_* can be determined, supposing that $w\kappa \ll 1$, as $w_*^2 = (15D)/(\lambda K)$. Therefore, one can approximate the classical active contour model energy as:

$$E_{C,0}(R) \cong E_0(\tilde{\phi}) = \lambda_C L(R) + \alpha_C A(R) . \quad (3.1.4)$$

The parameters of the active contour model are given approximately by:

¹Here we suppose that the centre of the polar coordinates is the centre of curvature.

$$\alpha_C = \frac{4\alpha}{3} \quad \text{and} \quad \lambda_C^2 = \frac{16D\lambda K}{15}. \quad (3.1.5)$$

Rochery et al. (2005b) noticed that this approximation gave parameter values, which were very close to those obtained by more sophisticated methods. Notice, that this is true in the case of road network or in general, in situations where the curvature is relatively small, however we need to be careful in the cases of large curvatures around the contour. Using the phase field model instead of the active contour model in the case of the ‘gas of circles’, we have observed that this accuracy gives satisfactory results even in the case of the inflection point model.




Figure 3.3: Evolution of the region using the simple phase field model in equation 3.1.1 ($D = 0.6$, $\lambda = 1.5$, and $\alpha = 0.01$).

To illustrate the behaviour of the region under the energy E_0 , figure 3.3 shows the evolution of an initial region containing arms and indentations. From left to right and from top to bottom, notice that the evolution fills the indentations and smooths the arms, so that finally we end up having a circle, which vanishes. This phenomenon is due to the minimization of the length of the region boundary and its area.

Rochery et al. (2005b) shown how it is possible to create a phase field term equivalent to a HOAC’s quadratic term. Our geometric energy takes the form $E_g = E_0 + E_{NL}$, where E_{NL} is a non-local term, similarly to the quadratic term of the HOAC model. HOAC energies are integral(s) over the contour, while phase field energies are defined as integrals on Ω . The non-local phase field energy is defined as:

$$E_{NL}(\phi) = -\frac{\beta}{2} \int_{\Omega^2} d^2x d^2x' \partial\phi(x) \cdot \mathbf{G}(x, x') \cdot \partial\phi(x'), \quad (3.1.6)$$

where \mathbf{G} is defined in term of Ψ by:




Figure 3.4: Region evolution using the non-local phase field energy. In the course of the evolution the region grows arms with a given width. Upper left: initial region, from left to the right: evolution, bottom right: final result ($D = 0.6$, $\lambda = 1.1$, $\alpha = 0.1$, $\beta = 0.1$, and $d = \epsilon = 5.0$).

$$\mathbf{G}(x, x') = \Psi(|x - x'|)\mathbf{I} \quad (3.1.7)$$

where \mathbf{I} is the 2×2 unit matrix.

Using the simplified approximation of the phase field variable, described in equation 3.1.2, one can determine the connection between β_C and β . The regions R_+ and R_- contain zero gradient, so all the contribution comes from the double integral over the interface region R_γ . The gradient of ϕ in R_γ is given by $\partial\phi = -(2/w)\hat{n}$, where \hat{n} is the outward normal to the boundary, extended to R_γ . The integrals across the interface in (3.1.6) are thus easy to perform, if we assume that \mathbf{G} is roughly constant over distances $\sim w$. Each integral simply contributes a factor w , which cancels that in the corresponding gradient. We are left with two boundary integrals of E_Q , with $\beta_C = 4\beta$.

Figure 3.4 shows an example of the evolution of a region using the non-local phase field energy ($E_g = E_0 + E_{NL}$). We initialized with a circle and with parameters creating road network structure. The steps of the evolution can be seen in the figure, starting from the upper left side.

3.2 The ‘gas of circles’ phase field model

In the previous chapter and in Horváth et al. (2006a, b), we analyzed the HOAC model and showed how to set the parameters so that the energy favours regions consisting of sets of circles with a prescribed radius. The aim of this section is to present an algorithm converting the existing ‘gas of circles’ parameters into phase field parameters. We will present synthetic experiments starting with different initial conditions and creating circles with the given radius, and present a comparison between the HOAC and phase field ‘gas of circles’ models.

We need to invert the expressions for the phase field parameters in terms of the HOAC parameters. Equation 3.1.5 gives α and λ , while at the end of section 3.1, we defined β in the terms of the HOAC parameters. In the case of the α and β parameters it is clear that

$\alpha = 3\alpha_C/4$ and $\beta = \beta_C/4$. It is slightly more difficult to compute λ . It turns out to be simplest to use the interface width w as a parameter instead of D (recall that $w = 4D/\lambda_C$). There are in principle two solutions, given by:

$$\tilde{\lambda}_{\pm} = \frac{\lambda_{\pm}}{\lambda_C} = \frac{15}{8w} \left[1 \pm \sqrt{1 - 4\tilde{\alpha}_C^2 w^2 / 5} \right], \quad (3.2.1)$$

where $\tilde{\cdot}$ denotes division by λ_C . In practice, we choose λ_+ . Note that this equation imposes a constraint on α_C for there to be a solution: $\tilde{\alpha}_C \leq \sqrt{5}/(2w)$. Finally, $D = \lambda_C w/4$. Note that for fixed $\tilde{\alpha}_C$ and $\tilde{\beta}_C$, the effect of λ_C is simply to scale all the phase field parameters.

3.2.1 Recipe: determination of the parameters

We can take the following steps to create a phase field ‘gas of circles’ model for a prescribed radius r_0 . (We assume that units have been chosen so that $d = 1$. If other units are chosen, *e.g.* pixel units, the parameters must be rescaled.)

- Choose w . It cannot be too small, or a subpixel discretization will be needed for gradient descent, and it cannot be too large or the phase field model will not be a good approximation to the HOAC model (Rochery et al., 2005b). We have found that $w = 3$ or $w = 4$ works well.
- Equation (3.2.1) now determines an upper bound on $\tilde{\alpha}_C$. Choose $\tilde{\alpha}_C$ in this range.
- Determine the $\tilde{\beta}_C$ parameter corresponding to r_0 and $\tilde{\alpha}_C$ using the method in Horváth et al. (2006b).
- Set $\tilde{\lambda} = \tilde{\lambda}_+$, $\tilde{\alpha} = 3\tilde{\alpha}_C/4$, $\tilde{\beta} = \tilde{\beta}_C/4$, and $\tilde{D} = w/4$.
- Choose λ_C appropriately for the application and multiply \tilde{D} , $\tilde{\lambda}$, $\tilde{\alpha}$, and $\tilde{\beta}$ by the chosen value.

3.2.2 Synthetic experiments

In our first experiment we compare the evolution of the HOAC and phase field ‘gas of circles’ models, initialized with a number of different shapes. The first row of figure 3.5 shows the evolution produced by $E_{C,G}$ with parameters chosen so that the contour creates circles with radius 5 pixels. Time runs from left to right. In the second row the initial value of ϕ was equivalent to that used in the experiment shown in the first row. The phase field parameters were derived using the algorithm given in section 3.2.1. The evolution is very similar to the HOAC evolution.

In figure 3.6 the initialization was the neutral initialization plus Gaussian noise of very small amplitude (for a more detailed explanation please see section 3.4.3). As it can be seen, the phase field evolves towards a set of circles with the prescribed radius.




Figure 3.5: Comparison of the HOAC and phase field evolutions. First row: gradient descent using $E_{C,G}$ ($\lambda_C = 10$, $\alpha_C = 1$, $\beta_C = 2.3137$, $r_0 = 5$). Second row: gradient descent using E_g , with parameters equivalent to those used in the first row ($\lambda = 9.064$, $\alpha = 0.75$, $\beta = 0.5784$, $D = 10$, $r_0 = 5$). The initialization used was equivalent to that used in the first row. Time runs from left to right.

3.3 Inflection point in the energy function

Now, we present how to choose the phase field ‘gas of circles’ parameters so that the circle energy would have an inflection point rather than a minimum at the desired radius. This is a non-trivial requirement. Rochery et al. (2005b) presented the relations between the phase field and contour parameters using an approximate “ansatz” for ϕ_R in section 3.1. For the ‘gas of circles’ model, the approximations are not expected to be important, since small errors in the parameters will produce small changes in behaviour. However, an inflection point represents a set of measure zero in the parameter space. It is important to see whether these approximate parameter relations preserve the inflection point behaviour.

In the previous section, we chose $\tilde{\alpha}_C$ according to the constraint $\tilde{\alpha}_C \leq \sqrt{5}/(2w)$. In section 2.3 we described how a given value of r fixes $\tilde{\alpha}_C$ and $\tilde{\beta}_C$.

To determine the parameters of the new phase field inflection point ‘gas of circles’ model, we therefore take the following steps:

1. Choose a d value satisfying $1.2776 \leq d \leq 1.4499$. This fixes $\tilde{\alpha}_C$ and $\tilde{\beta}_C$.
2. Choose w using the above criterion.
3. Determine $\tilde{\lambda}$, $\tilde{\alpha}$, $\tilde{\beta}$, and \tilde{D} as before.
4. Multiply these parameters by λ_C .




Figure 3.6: Gradient descent using E_g , with parameters ($\lambda = 9.064$, $\alpha = 0.75$, $\beta = 0.5784$, $D = 10$, $r_0 = 5$). The initialization was the neutral initialization with a small amount of added Gaussian noise. Time runs from left to right.

3.3.1 Experimental verification

To test that the parameter transformations work in practice, we fixed a set of contour parameters corresponding to an inflection point at $r_0 = 10$, and then translated these parameters to give an equivalent phase field model, the corresponding energy functions can be seen in figure 3.7. We then performed three gradient descent experiments using E_g . One used the value of β corresponding to the inflection point, β^* , while the other two used β values 4% above and below β^* . The region was initialized to a circle of radius r_0 . Figure 3.8 shows the results. In the first row, with $\beta < \beta^*$, the region shrinks and disappears. In the second row, with $\beta > \beta^*$, the region grows until it reaches the corresponding energy minimum, at a radius of 12 pixels. In the third row, with $\beta = \beta^*$, the circle is almost stable, grows only very slightly, to a radius of 10.5 pixels. This is because we initialize exactly in the inflection point, and using gradient descent there is no arising force because the gradient is zero. The inflection point behaviour is therefore preserved to a very good approximation.




Figure 3.7: Graphs of the energy of a circle.

3.4 Minimization of the phase field energy

In this section, we describe the minimization of the phase field energy terms. We show different choices for the initialization. We also describe useful implementation details.




Figure 3.8: Preservation of the inflection point.

3.4.1 Minimization of the local term, topology

The energy of the phase field is minimized by gradient descent. The functional derivative of the local energy term E_0 is:

$$\frac{\delta E_0}{\delta \phi} = -D\partial^2\phi + \lambda(\phi^3 - \phi) + \alpha(1 - \phi^2). \quad (3.4.1)$$

Gradient descent thus follows the equation, given by Allen and Cahn (1979). If we guarantee that the interface width is relatively small compare to the discretization grid size, the implementation can be performed using finite differences. The main advantage of the model compared to those that use a distance function representation (Osher and Sethian, 1988, Sethian, 1999, Osher and Paragios, 2003), is that no reinitialization and *ad hoc* regularization are required.

As we mentioned above, the new framework provides greater topological freedom than previous explicit or implicit contour based models. In the standard level-set framework, topology change can happen only by merging or splitting of the regions. In the phase field case regions can appear or disappear. In figure 3.9 we give an example to illustrate the topological freedom provided by the phase field model. We used E_0 as a geometric energy and a data energy which takes its minimum when the region boundary corresponds to edges in the image (for more details, see chapter 4). The underlying data was a black annulus on a white background. In the first row the initial region was a disk in the size of the annulus, while in the second row we used the neutral initialization (see section 3.4.3). Both the experiments segmented the desired shape correctly.

In addition, in the standard level-set framework, the number of handles of a region can only change by wrapping or separating. Phase fields again show more topological freedom: handles can appear and disappear within an existing region. This is illustrated in the top row of figure 3.9. The data and energy are the same as for the bottom row, but now the initialization is a circle. Note the formation of the hole in the centre.




Figure 3.9: Phase field evolution with different initial conditions, showing topological freedom. The image used was a black annulus corresponding to a white background, while the energy was the non-local phase field energy combined with the image gradient. In the top row the initial region was a circle while in the bottom row, we initialized with the neutral initialization. After the evolution the stable state was the region on the original annulus.

3.4.2 The non-local term

The functional derivative of the non-local phase field energy (equation 3.1.6) can be written as:

$$\frac{\delta E_{NL}}{\delta \phi}(x) = \beta \int_{\Omega} dx' \partial^2 G(x - x') \phi(x'). \quad (3.4.2)$$

We see explicitly the particular advantage of the phase field formulation for HOAC energies. Instead of the complex evaluation of the force due to E_Q described by Rochery et al. (2006), involving, at each iteration, contour tracing, contour integration, and force extension, the equivalent force arising from E_{NL} can be computed with a simple convolution. It can, for example, be evaluated in the Fourier domain. Implementation is thereby made much easier, and execution much faster. Execution time for the HOAC formulation scales as the square of the boundary length, which in turn scales as the number of trees, which in turn scales as the number of pixels. Thus execution time for the HOAC formulation can be expected to scale as the number of pixels squared. In contrast, execution time for the phase field formulation scales as the number of pixels. For large images, then, the advantage of the phase field formulation in practical terms is obvious.

3.4.3 Initialization

As we discussed before, the unconstrained nature of the phase field function offers greater topological freedom than previous methods. It also gives us a nice opportunity to choose the initial conditions.

Using the gradient descent method, we can initialize the phase field variable using a constant value $\phi_{\text{init}} = \alpha/\lambda$. Using this constant and choosing the threshold value of the map function ζ to be $z = \alpha/\lambda$ this initial condition is completely neutral, where neutral means that this constant corresponds to the maximum of the potential V , so it is not biased towards one or the other potential minimum, nor it is biased towards the inside or outside

area, because both are empty. Neutral initialization means that the phase field variable can create regions in the image domain, where necessary, so that we do not need complex initialization techniques or to initialize a region covering the whole image.

We note that two other initial conditions can be applied successfully. During our geometric experiments we used Gaussian noise, with a mean of $\mu = \alpha/\lambda$, and very small variance. This initialization creates random interior and exterior regions and during the evolution, the geometric terms control their behaviour. An example for this initialization can be seen in figure 3.6. The model preserves the opportunity to initialize the phase field variable with a predefined region. In this case we choose the interior $R_+ = \alpha/\lambda + c$ and the exterior $R_- = \alpha/\lambda - c$, where c is an arbitrary constant value, however in practice values $R_+ \approx 1$ and $R_- \approx -1$ were the most successful. An example can be seen in the bottom row of figure 3.5.

3.4.4 Implementation details

As we mentioned above, the functional derivative of the quadratic term can be computed easily in the Fourier domain. Here we briefly describe one way to perform this computation.

The interaction operator

First we need to determine the Fourier transform of the interaction operator. This can be done by constructing the interaction function on the grid, and applying the Fourier transform. An example can be seen in figure 3.10. For visualization purposes, we have taken its logarithm.




Figure 3.10: The logarithm of the amplitude of the Fourier transform of the interaction function. The grid size was 128×128 , $d = \epsilon = 4$.

The higher-order term

We need to compute the Fourier transform of the phase field. This we should multiply point-by-point with the interaction operator and computing its inverse Fourier transform.

Boundary condition

Using the Fourier transform we need to take into account its periodicity. In other words, we should picture the image as defined on the surface of a torus. This creates problems in detecting tree crowns on the boundaries, producing some false positive and negative detections.




Figure 3.11: Mapping of a forestry image onto a torus. Image by ©French Forest Inventory (IFN) .

Figure 3.11 illustrates the situation. On the left we can see an aerial image, while on the right the same image mapped onto a torus. It is clear that where the boundaries meet, there will be jumps in the image function, which creates problems during the segmentation. To solve this problem we enlarged the boundary with d pixels in all directions and set the pixel intensities to the value of μ_{out}^2 . This simple method was appropriate in our case for two reasons: first, in practice using the enlarged images produced satisfactory results; and second, since the tree crown size is relatively small compared to the image size, the approach is computationally efficient, because the size of the image not changing significantly. If we need longer range interactions or more accurate boundary conditions, we would need to use more complex methods *e.g.* Reeves (2005).

²We note that one could also solve this problem by duplicating or mirror the boundaries.

Chapter 4

Tree crown extraction

Herein we give a short overview of the history of the application of aerial imagery to forestry and of the recent publications on tree crown extraction. We introduce the multispectral and panchromatic aerial images we use. We develop two image data terms, one based on the intensity and the gradient of the most representative spectral band of the images, while the other uses up all three bands. Finally, we present our experimental results combining the different prior terms studied above with the data models.

FORESTRY applications, involving tree crown delineation in aerial images is a popular and well-studied image processing area. However, this area has a long history. Most of the models proposed are semi-automatic, and require quite a lot of human-interaction before or during the segmentation, and have difficult parameter setting procedures. There are only a few recent models that offer more automatic segmentation and need less human-interaction. The methods use only the most representative spectral band of the images or a very simplified model based on the multispectral bands. In section 4.1 we review the history of forestry remote sensing applications and previous publications in the field. The image data we use was provided by the French Forest Inventory (IFN) and the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) .

The images are color infrared (CIR) and panchromatic images and have three or one single spectral bands. To integrate the data into our model, we present two image terms (Horváth et al., 2006b, Horváth, 2007). The first one uses only one band. In the CIR case this is the infrared band, in which the reflectance of the vegetation is the highest. This model is based on two observations: one is that there is a significant change of intensity (gradient) between the tree crowns and the background, while the other is that the intensity levels both of the background and of the tree crowns follow normal distributions. Based on these observations, we define our first energy model so that a segmentation corresponding to the constraints minimizes it. The second model is based on the fact that by using the other bands of the CIR images, in many cases we can significantly improve the quality of the result. We present four different ways to incorporate multispectral data into the segmentation process

and analyze them in two ways: accuracy on a known image and generalization ability on an unknown image. We show that the best choice is Gaussian distributions with full covariance matrices. We present the minimization of both terms. Our data terms can be found in section 4.2. In section 4.3 we present an exhaustive series of experiments combining the different geometric prior models with the different data terms.

4.1 History of tree crown extraction from aerial images

Aerial imaging is more than 150 years old as a discipline. The first aerial image was taken from an air-balloon in the late 1850's in Paris by Tournachor. However the quality and the altitude can not be compared to current techniques. During the beginning of the 20th century and the First and Second World Wars, aerial imaging really developed. Balloons were replaced by airplanes and better cameras were used. Nowadays the field benefits from modern aircraft and space technology, make use of modern navigation techniques, and the development of computers to process the acquired data faster and with sophisticated methods.

The need to use aerial images in forestry grew at the end of the first half of the last century (Spurr, 1948) by recognizing their possibilities in resource management, statistics, *etc.* This idea can be considered the parent of the worldwide ground gathering automated interpretation of high spatial resolution digital imagery in silviculture. The main goal is to partly or fully exchange the human interpreter for an intelligent computer vision system, able to make decisions, and which requires minimal human interaction during the image analysis.

4.1.1 Data acquisition and the aerial images

In this section, we review the most important types of existing aerial data. In recent decades, the quality of the most common film types has been developed, and recently fine tuning of film characteristics is a focus of research (Brandtberg, 1999, Shao and Reynolds, 2006). We distinguish four main categories of film for aerial photography:

- black-and-white panchromatic
- colour
- black-and-white infrared
- colour infrared (CIR).

The CIR film type has the ability to capture the reflectance characteristics of vegetation. Vegetation has a unique spectral signature which enables it to be readily distinguished from other types of land cover in an optical/near-infrared image. The reflectance, as can be seen in figure 4.1, is low in both the blue and red regions of the spectrum, due to absorption by chlorophyll for photosynthesis. It has a peak at the green region, which gives rise to the green colour of vegetation. In the near infrared (NIR) region, the reflectance is much higher

than that in the visible band, due to the cellular structure in the leaves. Hence, vegetation can be identified by high NIR but generally low visible reflectance. This property was used in early reconnaissance missions in war for “camouflage detection”.




Figure 4.1: Plot of vegetation reflectance against wavelength.

For our experiments, we used CIR images provided by the French Forest Inventory (IFN) and panchromatic images provided by the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) taken by analogue cameras. Herein, we introduce the main properties of the images. The CIR image capturing time is usually close to 12⁰⁰ PM, and in most cases the exact time is available. The flight height is ~ 4000 m and the focal length of the camera is $f = 300$ mm. The film has three different spectral layers in the 520-900 nm wavelength domain. The film layers have three typical pigment colours, *i.e.* cyan, magenta and yellow. Each of them is dedicated to capturing the primary colours of light. The film layers are tuned so that they extract the 520-900 nm, 520-700 nm and 520-600 nm, respectively. This technique is called the false colouring infrared representation, which translates the green into blue, the red into green and the infrared colours into red, respectively (see figure 4.2). The film has an interesting feature, namely that the second and the third layers have sensitivity almost three times higher than the first layer, to compensate for the fact that the reflectance of healthy vegetation is significantly stronger in the near-infrared domain. The panchromatic images provided by the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) are grayscale images. The capturing times were between 9⁰⁰ AM and 12⁰⁰ PM, during the summer of 2002. The magnification is 1 : 10000, the flight height is ~ 1680 m, and the camera focal length is 153 mm.

The last step before digital image processing is digitization. We need to take into account the Shannon sampling theorem (Shannon, 1948), namely the sampling interval should be chosen to be less than half the size of the smallest interesting detail in the image. In our case, the images have ~ 0.5 m/pixel resolution. This implies that the size of the smallest details we are able to handle is 1 m. The usual size of the tree crowns that we are interested in, is between 3 m and 9 m, so this image resolution is satisfactory. In figure 4.3, we can see a typical CIR image.




Figure 4.2: False colouring technique.



Figure 4.3: Real color infrared (CIR) image with forest and urban area ©IFN.

4.1.2 Individual tree crown delineation

If we observe a forest in an aerial image, it contains trees in different positions next to each other. We can distinguish three different types of configurations, as shown in figure 4.4, for which different image features can be used by visual systems. In the first case (left), each tree stands individually. To extract the exact crown-shape it might be sufficient to threshold the image. In second case (middle), the local maximum might be useful to detect the bright tree top; contour extraction or valley detection (Eberly et al., 1994) might also be able to separate the crowns. In the third case (right), a group of overlapped trees without grey-level valleys, the detection of the individual trees must be based on the visible edges or contours.

Now we overview the most important results published on this topic. Three fundamental approaches exist to delineate individual tree crowns. The methods in the first group are based on the detection of local intensity extrema as the top of the crowns; methods in the second group are based on the contours between the crowns and background; while the third group matches templates and looks for the best local correlations. Finally, we present new, interesting works published very recently. Exhaustive overviews can be found in Brandtberg




Figure 4.4: Three typical cases of segmentation of individual tree crowns.

(1999), Erikson (2004), Perrin (2006), and Shao and Reynolds (2006).

The simplest group of tree crown delineation models uses local intensity maxima to identify the most illuminated, *i.e.* the brightest, part of the crowns, which is usually the tree top. Local maxima are pixels which have greater intensity than all the other pixels in a defined neighbourhood. Blazquez (1989) introduced a method for finding citrus trees using local maxima to identify tree crown centres, combined with crown perimeter, mean intensity and area. He also presented a method for separating joined trees, using shadows and radiometric normalization between photographs. The aerial images are often affected by noise, resulting in bright pixels and misleading the local maximum search algorithms. Notice that the difference between this situation and tree crowns, which is a bright circular area with usually 10–30 pixel diameter, is that in the crown case a bigger bright area is present. Therefore, one idea for handling this situation is to convolve the image with a Gaussian-kernel. Dralle and Rudemo (1996) presented a model to find the appropriate parameters. The model is based on smoothing the images with different filters, *i.e.* different scales, and comparing the result with field plot data of the number of stems present.

The second group of methods is based on contours. The term ‘contour’ in image processing is used as a delimiter between homogenous areas, but we note that in tree crown detection we distinguish contours between tree crowns and the background and valleys taking place between individual tree crowns. Gougeon (1995) presented an approach to tree crown delineation. He separated the crowns from the background vegetation and from each other. The tree crowns are subsequently delineated using a five-level rule-based method designed to find circular shapes, but with some small variations permitted. Using MEIS-II images of coniferous plantations with a resolution of 31 cm/pixel, 81 % of the crowns are the same as those obtained by visual interpretation of the imagery. Brandtberg and Walter (1998) decompose an image into multiple scales (Lindeberg, 1998), and then define tree crown boundary candidates at each scale as zero crossings with convex greyscale 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.

Using templates to model tree crowns is a well studied and successful part of the field. Although template matching methods are known as a classical searching technique, their success in the field is due to all the necessary information being available about aerial images to create a very precise tree model. The first vision system capable of recognizing individual tree crowns, based on matching of a synthetic tree crown image model with an

aerial image, was developed by Pollock (1996). The system was tested on monocular high spatial resolution image data in Ontario and Alberta (Pollock, 1998). The procedure is based on a model of the image formation process at the scale of an individual tree. Natural variation of the tree crown is considered, as is the species. Larsen (1998, 1999) concentrated on spruce tree detection using a template matching method. The main difference between these two papers is the use of multiple templates in the second. 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. Reflectance is modelled using single reflections, with the branches and needles acting as scatterers, while the ground is treated as a Lambertian surface. 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.

More recently, new approaches have been proposed using novel ideas and mixing the three previously used image data. Erikson (2003) used a region-growing method to separate crowns of mixed forest in high-resolution aerial images in central Sweden. The method starts from single (seed) pixels as representative pixels on the crown and spreads over neighbouring pixels that satisfy spatial and spectral requirements. In Erikson (2003), the region-growing part was completed with a random walk. The methods described so far use multiple steps rather than a unified model. Closer in spirit to the present work is that of Perrin et al. (2004, 2005), who model 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. Reversible Jump Markov chain Monte Carlo and simulated annealing are used to estimate the tree crown configuration. Compared to our model, 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 Data terms for tree crown extraction

In this section, we present two data terms designed to incorporate the information from aerial images into the global energy. The first is based on the following observations: the infrared spectral band of the vegetation images, and hence the tree crowns, is the most representative (see subsection 4.1.1); thus we can use this band on its own¹; the intensity of the tree crowns is different from non-forested areas, and the distribution of the intensity values can be characterized by a normal distribution; the gradient of the intensity values is significantly bigger at the transition between the crowns and the background. Using these facts, in subsection 4.2.1 we present a data term based on the intensity and gradient of the infrared spectral band.

Although, the reflectance of the vegetation is highest in the infrared band, using all the available bands we can improve the segmentation. We analyzed four different possible candidates to incorporate multispectral information and concluded that modeling the distributions of the crowns and the background with normal distributions using full covariance

¹in the case of panchromatic images we have only one spectral band

matrices gives the best results: generalization ability is strong, few parameters are needed and the learning process is simple. The details can be found in section 4.2.2.

4.2.1 Model for single band images

The images have a resolution $\sim 0.5\text{m/pixel}$, and tree crowns have diameters of the order of ten pixels. Very little, if any, 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. Therefore we choose a lower level image representation. In the upper left image of figure 4.5, we see the infrared spectral band of a typical CIR image containing poplar stands, while on the right is its ground truth segmentation.




Figure 4.5: Infrared spectral band of a CIR image. Upper left: real image with poplar forest ©French Forest Inventory (IFN) . Upper right: ground truth segmentation of the image. Bottom: Histogram of the intensity values based on the ground truth segmentation and the corresponding normal distribution functions. Black: non forest areas, blue: tree crowns.

Our data energy contains two parts, one based on the intensities in the foreground and background, the other on the gradient. The bottom image in figure 4.5 shows the histograms of the intensities corresponding to the ground truth segmentation. Since, both the histograms can be approximated with Gaussian distributions, we model the interior of the tree crowns with mean μ_{in} and covariance $\sigma_{in}^2 \delta_R$, where δ_A is the identity operator on images on $A \subset \Omega$,

and R is the region of our interest. In most cases, *e.g.* in figure 4.5, the background also follows a Gaussian distribution. However it can be very varied, and thus hard to model in a precise way. We use a Gaussian distribution with mean μ_{out} and variance $\sigma_{\text{out}}^2 \delta_{\bar{R}}$.




Figure 4.6: Magnitude of the gradient of the infrared spectral band.

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. Figure 4.6 shows how significant is the magnitude of the gradient at the boundary of the crowns.

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.*

$$-\ln g_R(I_R) = \int_R d^2x \frac{1}{2\sigma_{\text{in}}^2} (I_R(x) - \mu_{\text{in}})^2$$

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

$$-\ln f_{\partial R}(I_{\partial R}) = \lambda_i \int_{\square_\gamma} dt \mathbf{n}(t) \cdot \partial I(t)$$

where \mathbf{n} is the unnormalized outward normal to γ . The normalization constant Z is thus a function of μ_{in} , σ_{in} , μ_{out} , σ_{out} , and λ_i . Z is also a functional 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 the prior energy. However, these parameters are essentially fixed by hand, knowledge of the normalization constant does not change their values, and we ignore it once the likelihood parameters have been learnt.

The data model is then given by:

$$E_{C,I}(I, R) = -\ln P(I|R, K) + \ln Z = -\ln g_R(I_R) - \ln g_{\bar{R}}(I_{\bar{R}}) - \ln f_{\partial R}(I_{\partial R}) .$$

Energy minimization

The energy is minimized by gradient descent. The gradient descent equation for $E_{C,I}$ is

$$\hat{\mathbf{n}} \cdot \frac{\partial \gamma}{\partial s}(t) = -\partial^2 I(\gamma(t)) - \frac{(I(\gamma(t)) - \mu_{\text{in}})^2}{2\sigma_{\text{in}}^2} + \frac{(I(\gamma(t)) - \mu_{\text{out}})^2}{2\sigma_{\text{out}}^2} , \quad (4.2.1)$$

where s is the descent time parameter. As already mentioned, to evolve the region we use the level-set framework of Osher and Sethian (1988) extended to the demands of nonlocal forces such as equation (2.4.2) (Rochery et al., 2005c, 2006).

Phase field data term

It is easy to reformulate active contour likelihoods as phase field models via the dictionary: normalized inward-pointing boundary normal vector $\partial\phi/2$; boundary characteristic function $|\partial\phi|^2$; region characteristic function $\phi_+ = (1 + \phi)/2$; region complement characteristic function $\phi_- = (1 - \phi)/2$. Using these definitions, the phase field data term equivalent to the previous contour data term is:

$$E_i(I, R) = \int_{\Omega} d^2x \left\{ \lambda_i \partial I \cdot \partial \phi + \frac{(I - \mu_{\text{in}})^2}{2\sigma_{\text{in}}^2} \phi_+ + \frac{(I - \mu_{\text{out}})^2}{2\sigma_{\text{out}}^2} \phi_- \right\} .$$

4.2.2 Multispectral data model

In this section, we study data models that use all three bands of the CIR image data.

The left of figure 4.7 shows the falsely coloured CIR version of figure 4.5. It is of a poplar stand. To see how colour can help, note that the bright pixels in the spaces between the trees are light grey in the colour image, while the trees are red. In the greyscale image, they have roughly the same intensity, making the separation of trees and background difficult. Although, the prior model helps to disambiguate these situations, it is not always successful, and it makes sense to consider a data model that uses the available information to the full.

We want to construct a data model for the observed CIR image, given that region R corresponds to tree crowns. We can divide the image I (a three-component quantity) into two pieces, I_R and $I_{\bar{R}}$ corresponding to the tree crowns and the background. Then we have $P(I|R, K) = P(I_R, I_{\bar{R}}|R, K)$. Without further information, I_R and $I_{\bar{R}}$ are not independent given R . However, we may introduce parameters for the two regions, θ_R and $\theta_{\bar{R}}$, so that the two pieces of the image become independent given these parameters. We note, that similarly to the previous case using only the infrared spectral band, due to the small size of



Figure 4.7: (left): typical CIR aerial image of a poplar plantation; (middle): another CIR image; (right): corresponding ground truth. Images ©French Forest Inventory (IFN) .




Figure 4.8: Histograms of pixel values in the three bands of figure 4.7 (left), based on the manual labelling shown in figure 4.5 (upper right). Green is background; blue is tree crown.

the tree crowns we cannot define meaningful texture features. Thus, we will assume that the image values at different pixels are independent. The data model then takes the form

$$P(I|R, \theta_R, \theta_{\bar{R}}, K) = \prod_{x \in R} P(I_R(x)|\theta_R, K) \prod_{x \in \bar{R}} P(I_{\bar{R}}(x)|\theta_{\bar{R}}, K) .$$

To help us design the model for individual pixels, we examine the statistics of the pixel values in the different bands. Figure 4.8 shows histograms of the pixel values in figure 4.7 (left) for all three bands, separated into tree crown and background based on a manual labelling shown on figure 4.5 (upper right). As expected, the infrared band shows the largest separation. Can adding the other two bands help?

To test this idea, we performed four different types of maximum likelihood classification, based on four different estimates of the probability distributions for individual pixels of each class. Two of these estimates use raw histograms with different bin sizes. Of these, one is constructed as a product of the individual histograms for each band (independent bands), called HI for short, while the other uses the colour histogram (HC). The other two estimates use Gaussian models, either with covariances diagonal in colour space (independent bands), called GI, or with full covariances (G3D). The models parameters were learned from figure 4.7 (left) and figure 4.7 (middle), based on the manual labellings. The resulting models were then used to classify the image in figure 4.7 (left).




Figure 4.9: Maximum likelihood classifications of figure 4.7 (left) using the different models trained on the same image.

The results of maximum likelihood classification on the same image are shown in figure 4.9. The images have four different colours: black and dark grey correspond to correct classification of background and tree crowns respectively, while light grey and white correspond to incorrect classifications in which tree crowns were classified as background and vice-versa respectively. Table 4.1 (top) shows the resulting classification error rates.

Naturally, the results using HC are almost perfect. The number of bins is very large, and this means that there are unlikely to be more than one or two pixels in each bin. Consequently, any given pixel is very likely to have zero probability to be in the incorrect class. Equally clearly, the results using HI are poor: the different bands are not independent. This is confirmed by the result for GI. G3D, however, produces a reasonable performance, second only to the HC results. Bearing in mind that G3D has $3 + 6 = 9$ parameters, while HC has the same number of parameters as bins, this is encouraging. These conclusions are confirmed by the label images, which clearly show the inferior classifications produced by the models with independent bands.

To test the generalization ability of the models, we used a different image to learn the model parameters, and used them to classify figure 4.7 (left). The new training image is figure 4.7 (middle), along with a manual labelling. Figure 4.10 shows the results, while table 4.1 (bottom) shows the error rates.

It is not a surprise that the error rates are larger. The histogram-based methods do not generalize well, and produce more errors than both Gaussian models. The Gaussian results are naturally not as good as in the previous test, but are adequate in the absence of a prior energy. The model with dependent bands performs considerably better than the independent




Figure 4.10: The same classification trained on figure 4.7 (middle).

| Method | B→F | F→B | error (%) |
|----------------------------|------|------|-----------|
| HI (64 bins) | 446 | 404 | 9.64 |
| HI (128 bins) | 446 | 399 | 9.58 |
| HC (64 ³ bins) | 121 | 214 | 3.8 |
| HC (128 ³ bins) | 19 | 97 | 1.32 |
| GI | 470 | 426 | 10.16 |
| G3D | 256 | 328 | 6.62 |
| Method | B→F | F→B | error (%) |
| HI (64 bins) | 748 | 242 | 11.22 |
| HI (128 bins) | 752 | 253 | 11.39 |
| HC (64 ³ bins) | 1028 | 490 | 17.21 |
| HC (128 ³ bins) | 1747 | 1277 | 34.52 |
| GI | 1106 | 123 | 13.93 |
| G3D | 841 | 85 | 10.5 |

Table 4.1: Error rates for the maximum likelihood classification of figure 4.7 (left), using models trained on the same image (upper table) and on figure 4.7 (middle) (bottom table).

band model in both cases. In particular, the independent band models, whether histogram-based or Gaussian, consistently confuse certain types of inter-tree background with the tree crown foreground.

Energy minimization

Our full energy functional for tree crown extraction is a combination of the energy associated to the likelihood, $E_{C,I}(\gamma, I) = -\ln P(I|R, \theta_R, \theta_{\bar{R}}, K)$, and the HOAC ‘gas of circles’ prior geometric term E_g : $E(\gamma, I) = E_g(\gamma) + E_{C,I}(\gamma, I)$, θ_R and $\theta_{\bar{R}}$ are the mean and variance parameters on the foreground and background respectively. In the previous section, we established that the Gaussian model with full covariance provides the best compromise between precision and generalization. Here we describe this data term and how we minimize the energy.

The parameters of $E_{C,I}$ are learnt from samples of each class using maximum likelihood, and then fixed. We denote the mean vectors inside and outside as M_{in} and M_{out} and the covariance matrices Σ_{in} and Σ_{out} . We define the energy as above:

$$E(\gamma) = E_g(\gamma) - \int_R d^2x \ln \left[\det^{-1/2}(\Sigma_{in}/2\pi) e^{-\frac{1}{2}(I(x)-M_{in})^T \Sigma_{in}^{-1} (I(x)-M_{in})} \right] \\ - \int_{\bar{R}} d^2x \ln \left[\det^{-1/2}(\Sigma_{out}/2\pi) e^{-\frac{1}{2}(I(x)-M_{out})^T \Sigma_{out}^{-1} (I(x)-M_{out})} \right].$$

The energy is minimized using gradient descent. The descent equation is

$$\hat{\mathbf{n}} \cdot \frac{\partial \gamma}{\partial s}(t) = \frac{1}{2} \ln \left(\frac{\det(\Sigma_{in})}{\det(\Sigma_{out})} \right) \\ - \frac{1}{2} \left\{ (I(\gamma(t)) - M_{in})^T \Sigma_{in}^{-1} (I(\gamma(t)) - M_{in}) - (I(\gamma(t)) - M_{out})^T \Sigma_{out}^{-1} (I(\gamma(t)) - M_{out}) \right\},$$

To evolve the contour we use the level-set framework (Osher and Sethian, 1988) extended to cope with the nonlocal forces arising from higher-order active contours (Rochery et al., 2005c, 2006).

4.3 Experimental results

In this section, we present an exhaustive series of experiments combining our prior energy with the data terms we defined in this chapter. We will split this section into five different subparts. In each of them we present one of our models, and compare with previous results where appropriate.

The models and the list of the parameters we used during the experiments are the following:

- The ‘gas of circles’ (GOC) model combined with the single spectral data model in subsection 4.3.1. $(\lambda_i, \lambda_C, \alpha_C, d, r_0)$

- The inflection point ‘gas of circles’ (AGOC) model with the single spectral model in 4.3.2. (λ_i, d, r_0)
- The higher-order phase field ‘gas of circles’ (PF GOC) model with single spectral data term in 4.3.3. $(\lambda_i, \lambda, \alpha, \beta, D, r_0)$
- The phase field inflection point ‘gas of circles’ (PF AGOC) model in subsection 4.3.4. (λ_i, d, r_0)
- And the inflection point HOAC ‘gas of circles’ (CAGOC) model combined with multi-spectral data term in 4.3.5. (λ_i, d, r_0)

The parameters, learned from the image, are shown when the data is mentioned, in the form $(\mu_{\text{in}}, \sigma_{\text{in}}, \mu_{\text{out}}, \sigma_{\text{out}})$. Unless otherwise specified, images were scaled to take values in $[0, 1]$. Sometimes we compare our models with the classical active contour, the contour parameters are in the form $(\lambda_i, \lambda_C, \alpha_C)$.

We present our results on aerial images provided by the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) and the French Forest Inventory (IFN). Details of the acquisition can be found in section 4.1.1. The panchromatic images were provided by the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA), taken in Hungary near to the cities of Kecskemét and Nagybjom. The images contain mainly different individual tree groups and regularly planted poplar forests. The color infrared images were provided by the French Forest Inventory (IFN) from the region Saône et Loire in France. The images represent mostly regular tree stands and some of them irregular tree groups. In most cases (except in subsection 4.3.5), we used only the infrared spectral channel of the images. If it is not essential or not ambiguous, we will not especially highlight this fact. During the experiments, for visualization purposes sometimes we set the images darker.

4.3.1 The ‘gas of circles’ HOAC model

The tree crowns in the images are $\sim 8\text{--}10$ pixels in diameter, *i.e.* $\sim 4\text{--}5\text{m}$. We compare our model to a classical active contour model ($\beta_C = 0$). The parameters $\mu_{\text{in}}, \sigma_{\text{in}}, \mu_{\text{out}},$ and σ_{out} were the same for both models, and were learned from hand-labelled examples in advance. 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 has been fixed, β_C is determined by equation (2.2.7). There are thus three effective parameters for the HOAC model. In the absence of any method to learn λ_i, α_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 real 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 μ_{out} to ensure that the region would shrink inwards.

Figure 4.11 illustrates a first experiment. On the left is the data. The image shows a small piece of an irregularly planted poplar forest. The image is difficult because the

intensities of the crowns are varied and the gradients are blurred. In the middle there is the best result we could obtain by using a classical active contour. On the right there is the result we obtain with the HOAC ‘gas of circles’ model. 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.




Figure 4.11: 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, 4.2, 4.2).

Figure 4.12 illustrates a second experiment. Again, the data is on the left, the best result obtained with a classical active contour model is in the middle, and the result with the HOAC ‘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.

Figure 4.13 shows a third experiment. The data is on the left, the best result obtained with a classical active contour model is in the middle, and the result with the HOAC ‘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.

4.3.2 The inflection point ‘gas of circles’ model

We compare the AGOC model to the GOC model containing an energy minimum. Note, that the AGOC model has only three free parameters, λ_i , α_C and d , since the other likelihood parameters are fixed by training, while the other prior parameters are fixed once r_0 is known. As in the previous subsection, the initial contour in all experiments, except that in figure 4.16, was a rounded rectangle slightly bigger than the image domain. The image




Figure 4.12: 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, 3.5, 3.5).




Figure 4.13: 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, 3.5, 3.5).

values in the region exterior to the image domain were set to μ_{out} to ensure that the contour would shrink inwards.

Figure 4.14 shows the result we obtained on the left image of figure 4.11. Note, that the parameter values for the model, although fixed, nevertheless produce a good result. One tree on the border is missing, but on the other hand, two trees are separated that were merged by the GOC model in figure 4.11 right.




Figure 4.14: Result obtained in figure 4.11 (left) with the AGOC model (1125, 6, 4.16).

Figure 4.15 shows three images. On the left is the data; in the middle is the result obtained with the GOC model; on the right is the result obtained with the AGOC model. Despite its fixed parameters, the AGOC model produces a better result, finding a tree missed by the GOC model, and again separating trees that were merged by the GOC model.




Figure 4.15: Left to right: regularly planted poplars ©IFN; result with the GOC model (40, 0.05, 5, 3.47, 3.47); result with the AGOC model (1285, 5, 3.47).

Figure 4.16 shows two images. On the left is the data, while on the right is the result obtained. The initial contour in this experiment was the red line. With a couple of exceptions, the trees are separated and the extraction is accurate.




Figure 4.16: Left: bigger slice of planted forest ©IFN; right: result using the AGOC model (2250, 5, 3.47). The contour was initialized to the red line.

For the experiment in figure 4.17, we used an α value slightly larger than that given by equations (2.3.2), in order to make E_1 slightly positive for all r . This ensures that in the absence of image data, circles will disappear. The resulting E_0 is shown on the upper left in the figure. Next comes the data. The aim of the experiment is to detect the older, larger radius trees in the upper part of the plantation area. Bottom left is the best result using the GOC model. Note, the phantom regions generated as the contour becomes trapped in local energy minima (the phantom regions in the bright exterior area are also reinforced by the image term). On the bottom right is the result using the AGOC model. With one exception, the phantom regions are eliminated, while the level of error elsewhere is comparable to the GOC model.




Figure 4.17: Top left: energy of a circle with α slightly greater than the value given by equation (2.3.2), to create slightly positive gradient everywhere; top right: regularly planted poplars ©IFN; bottom left: result with the GOC model (15, 0.008, 4.5, 2.51, 2.51); bottom right: result with the AGOC model (40, 3.6, 2.51), $\alpha_C = 6 > 5.40$.

4.3.3 The ‘gas of circles’ phase field model

In this section, we compare the PF GOC model with a classical active contour (CAC) model, and with the GOC model. The CAC, GOC and AGOC code is in C++, while both the PF GOC and the PF AGOC code is in Matlab. This should be born in mind when comparing

execution times.

To count the number of free parameters for PF GOC models, we note that: w is fixed *a priori*; r_0 is fixed by the application; we always choose $d = r_0$; β_C is determined from α_C and r_0 ; the phase field parameters are determined once the contour parameters are chosen. This means, that the only truly free parameters are λ_i , α_C and λ_C (or equivalently λ_i , D , and α). In addition, for the phase field model, α_C is constrained. Thus, both the GOC and PF GOC models have three effectively free parameters, the same number as the CAC model. The free parameters for each model, in common with most variational and many other methods, were fixed empirically (separately for each model) to give good results.

Figure 4.18 shows the result we obtained on the image in figure 4.13 (105×236). With the CAC model the segmentation is poor: there are many misclassified and fused objects. The best result obtained using the GOC model, although significantly better than the CAC one (see figure 4.13), required execution time was 152 minutes. The best result of the PF GOC model is still not perfect, but is an improvement over the GOC result, showing fewer misclassified tree crowns, and the execution time was less than 1 minute.




Figure 4.18: Experimental result with the PF GOC model on figure 4.13 left ($1200, 239, 22.5, 16.8, 150, 3.5$).

Figure 4.19 shows the result on figure 4.11 left (128×128). This is a less regularly-planted poplar stand. The challenge of this image is that the tree crowns sometimes appear connected, and their crowns have varied intensities. In the best result obtained using a CAC model, several crowns are merged together, and the boundary is rather noisy. The best result obtained using the GOC model took 96 minutes to compute (see figure 4.11). Figure 4.19 shows the best result obtained using the PF GOC model. Again, it is an improvement on the GOC result, with fewer fused tree crowns, while it took only 15 seconds to compute.

4.3.4 The inflection point ‘gas of circles’ phase field model

We compare the PF AGOC model with the AGOC model and with the PF GOC model.

We made our first experiment on the upper right image of figure 4.17 (200×200) with a regularly-planted poplar forest. In the top right and bottom left there are fields, while in the middle, two different sizes of poplars. The aim is to extract the larger trees. The result with the AGOC model is good, but the method found two small trees, and there is another




Figure 4.19: Experimental result with the PF GOC model on figure 4.11 (left) (100, 24.9, 5.63, 2.59, 18.8, 4.2).

false positive in the bottom left of the image (see figure 4.17). The execution time was 89 minutes. Figure 4.20 left shows the result using the PF GOC model. There were no false negatives, but the model found false positives in the fields. Figure 4.20 right shows the result with the PF AGOC model. All but two somewhat smaller circles were successfully found. For the phase field models, the execution time was less than 3 minutes.




Figure 4.20: Left: result on figure 4.17 with the PF GOC model (1111, 278, 177.2, 55.8, 30, 2.5); right: result with the PF AGOC model (1160, 3.57, 2.5).

Figure 4.21 left shows an image (133×271) of a planted forest. Figure 4.21 middle shows the result with the PF GOC model. The result is good, with only a few joined tree crowns. Figure 4.21 right shows the result with the PF AGOC model. There are fewer joined tree crowns. Both results were obtained in less than 2 minutes.

The upper row of figure 4.22 left shows a difficult image (129×139) to analyse. It has two fields with different intensities on the right. The result with the PF GOC model is shown in figure 4.22 middle. This result clearly demonstrates the disadvantage of the non-inflection point model: phantom objects are created in the homogenous areas. Figure 4.22




Figure 4.21: Left: an image of regularly planted poplar stands with a less regularly planted trees in the upper part © IFN (0.8, 0.06, 0.43, 0.2, 3.5); middle: result with the PF GOC model (500, 50, 34.2, 9.3, 5.2, 3.5); right: result with the PF AGOC model (500, 4.73, 3.5).

right shows the result with the PF AGOC model. The result is very good, with only one false positive. Both results were obtained in less than 1 minute. In the bottom image of figure 4.22 the corresponding phase field surface is illustrated meshed at the threshold level with the original image.

Figure 4.23 top shows an image (646×639) with agricultural areas and some planted trees. The result obtained with the PF AGOC model can be seen in the bottom image. The aim was to find all the crowns. The image is a difficult task, because the intensity of the tree crowns is very similar to the field intensity, and between the different fields we can observe significant gradient. The result, however is good, all the crowns were successfully detected, while the execution time was 4 minutes.

In the left image of figure 4.24, a real image (323×174) is presented with a regularly planted poplar grove next to a farm. The goal of the segmentation was to find all the tree crowns. The result we obtained can be seen in the right image. We find most of the crowns but some are missing, and a false positive error also can be seen near to a building. The run time was 5 minutes.

Figure 4.25 left shows an image (376×349) with a field and irregularly grown trees with slightly different crown sizes. Our aim was to detect all the crowns. The result is on the right. The algorithm found all but one tree, which is relatively small, in 7 minutes.

In the upper image of figure 4.26, we show a very difficult, however regularly planted pine forest (430×301). The challenge of the image is that the size of the trees varies and the branches can be seen, which reduces the circular shapes of the crowns. Our result can be observed in the bottom image. Beyond two merged regions, the method found very accurately all the crowns in 2 minutes.

In figure 4.27 planted trees with different crown size can be seen (368×486). The aim of the segmentation was to find as many crowns as possible, although their size is quite varied. The result, as we can see in the right image is quite good, but one tree is missing, the exact crown shape of the bigger trees is not extracted and in the middle of the bottom part one big crown was segmented as two trees. The execution time was 5 minutes.

Finally, in the upper image of figure 4.28 (577×268) we can see a regularly planted poplar forest with very similar crown size. The challenge in the image is to successfully avoid false positive detections in the fields in the upper left and bottom right. Our method was able to detect all the crowns except some very young trees, and made only one false positive error. This can be seen in the bottom image. The result was obtained in 4 minutes.




Figure 4.22: Upper row; left: an image of regularly planted poplars with different fields on the right © IFN (0.8, 0.06, 0.43, 0.2); middle: result with the PF GOC model (500, 50, 34.2, 9.3, 5.2, 3.5); right: result with the PF AGOC model (500, 4.73, 3.5). Bottom: The corresponding phase field surface, thresholded with the original image.




Figure 4.23: Top: image with sparsely planted trees © CAO, FA. (0.41, 0.14, 0.66, 0.07)
Bottom: result with the PF AGOC model (60, 12.15, 9).




Figure 4.24: Left: real image with regularly planted poplars next to a farm © CAO, FA. (0.27, 0.1, 0.65, 0.1) Right: result obtained with the PF AGOC model (1000, 5.4, 4).




Figure 4.25: Left: separated tree crowns © CAO, FA. (0.4, 0.09, 0.8, 0.1) Right: result obtained with the PF AGOC model (1000, 9.45, 7).




Figure 4.26: Top: regularly planted pine forest © CAO, FA. (0.72, 0.05, 0.34, 0.13) Bottom: result obtained with the PF AGOC model (400, 4.05, 3).




Figure 4.27: Left: Real image with trees of different sizes © CAO, FA. (43, 0.12, 0.68, 0.09) Right: result with the PF AGOC model (250, 10.8, 8).




Figure 4.28: Top: regularly planted poplar forest with fields in the top and bottom corners © CAO, FA. (0.25, 0.14, 0.66, 0.05) Right: result obtained with the PF AGOC model (225, 5.4, 4).

4.3.5 Experiments on color infrared images

We compare three models: the CAGOC model, which uses the multispectral data term with the ‘gas of circles’ prior; the AGOC model, which uses only the infrared band of the CIR image with the ‘gas of circles’ prior; and a classical active contour model, which uses the multispectral data model, but only the length and area terms of E_g , *i.e.* $\beta_C = 0$. There is thus no prior shape information in this third model. In all experiments.




Figure 4.29: Results obtained on the image shown in figure 4.7, using the CAGOC model (left) (14.25, 3.58, 2.5), the AGOC model (middle) (0.25, 3.58, 2.5), and the classical active contour model combined with the multispectral data term (right). © IFN

Figure 4.29 shows the results obtained on the image shown in figure 4.7 left, using the CAGOC model, the AGOC model, and the classical active contour model respectively. The CAGOC model is the most successful, separating trees that are not separated by the other models.




Figure 4.30: Results obtained on the image shown in figure 4.7 middle, using the CAGOC model (left) (17.5, 3.58, 2.5), the AGOC model (middle) (0.2, 3.58, 2.5), and the classical active contour model (right). © IFN

Figure 4.30 shows the results obtained on the image shown in figure 4.7 (middle). None of the results is perfect, all the models failing to separate some trees, but the CAGOC model detects several trees that are not detected by the AGOC model. The classical active contour model was not able to separate all the crowns, and found a large connected area at the bottom right, due to the missing prior shape information.




Figure 4.31: Top left: a CIR image; top right: result with the CAGOC model (30, 3.58, 2.5); bottom left: result with the AGOC model (0.3, 3.58, 2.5); bottom right: result with classical active contour model combined with the multispectral data term.

Figure 4.31 (top left) shows a difficult image with a field at the top, and strong shadowing. The result with the CAGOC model, shown in figure 4.31 (top right), is good, detecting all the trees and ignoring the field and shadows. The result with the AGOC model, shown in figure 4.31 (bottom left), is not as good. Some trees are missed, but more importantly, the fact that the field has a similar IR response to the tree crowns means that a large incorrect region is produced. The result with the classical active contour model, shown in figure 4.31 (bottom right), avoids this error thanks to the multispectral information, but the lack of prior shape information means that some trees are merged.

Figure 4.32 (top left) shows a different type of image, of isolated trees in fields. The result with the CAGOC model, shown in figure 4.32 (top right), is correct, ignoring the field, for example. The result with the AGCO model is not as good, with one large false positive, and smaller errors on each of the detected trees, due to confusion between the field and parts of the road and the tree crowns (figure 4.32 (bottom left)). Figure 4.32 (bottom right) shows the result obtained using the multispectral data term combined with a classical active contour model. The result is almost as good as the CAGOC model, except that the contours are slightly less smooth, and there is a small false positive area in the upper right




Figure 4.32: Top left: a CIR image; top right: result with CAGOC model (25, 9.5, 7); bottom left: result with the AGOC model (0.3, 9.5, 7); bottom right: result with classical active contour model ($\beta_C = 0$) combined with the multispectral data term. © IFN

corner, which was not detected by the CAGOC model, presumably because it is not circular.

Figure 4.33 (left) shows another CIR image with fields and some sparse trees. It is a difficult image, because some of the fields have a similar colour to the trees. The result with the CAGOC model, shown in the second image of figure 4.33, is good, detecting all the trees, and only merging two of them. The result with the AGOC model shown in the third image of figure 4.33, is not as good. The greyscale level between some of the trees is too similar to the tree crowns to be separated, despite the prior shape information, meaning that several trees are merged. In addition, some non-tree objects were detected as tree crowns, again due to similarity of grey scale. The result obtained with the classical active contour and multispectral data model is slightly better, but due to the missing prior shape information several tree crowns are merged and a small non-tree area was detected.




Figure 4.33: From left to right: a CIR image; result with the CAGOC model; result with the AGOC model (stable radius $r_0 = 4.0$); result with classical active contour model ($\beta_C = 0$) combined with the multispectral data term. © IFN

Chapter 5

Conclusion, unsolved problems

In this chapter we give a summary of the aims and methods. We present the unsolved problems raised in this thesis and propose possible solutions to them. We discuss future directions for the methods, and possible ways to improve the data term, the prior model, and the optimization algorithms speed.

In section 5.1, we describe our initial goal and the way we approached it; the process we applied; the results we achieved. In section 5.2 we present the insufficiencies of our method and provide some possible solutions. In section 5.3, we set some tasks for the future that need to be dealt with in order to improve the modelling capacity of the method, accuracy, and complexity.

5.1 Summary

The initial aim of our work was to create a computer-based model for the extraction of individual tree-crowns, which combine shape prior information with available image features. We used a new generation of active contour (AC) model, the higher-order active contours (HOAC), initially created to detect road networks. However, we have shown that with the proper setting of the parameters, the same model can be used to produce stable circles. We analyzed its energy; and introduced the ‘gas of circles’ (GOC) model. While this model had resulted in a great improvement over classical methods (*e.g.* active contours), some problems still arose. The model formed phantom circles in those areas, where there were no edge and no difference from the tree crowns in intensity. The model had too many free parameters, and its complexity was large compared to the classical methods. We solved the problem of phantom circles by setting the values of the parameters in such way that the energy had an inflection point at the desired radius. If there is no information coming from the image then the circles vanishes, but a very small amount of image information can create a minimum in the energy function (AGOC). This model also helped in reducing the number of parameters. The inflection point criteria halve the number of parameters. To solve the

question of computational complexity, we modelled the energy within the phase field (PF) framework. The results were two orders of magnitude faster runtime, more topological freedom, a less difficult implementation, and other modelling and algorithmic opportunities. With the aid of the phase field model, it is possible to create a higher-order phase field (HOPF) framework similar to HOAC; we showed how to set the parameters of this model so as to result in a model (PF GOC) approximately equivalent to the ‘gas of circles’ model. Finally, we introduced the inflection point version of the model (PF AGOC). The derivation tree of the evolution and logical connection of the methods can be seen in figure 5.1.




Figure 5.1: Overview and connection between the methods we used and developed. The methods we developed can be found under the red curve. **AC**: active contour model; **PF**: phase field model; **HOAC**: higher-order active contour model; **HOPF**: higher-order phase field model; **GOC**: ‘gas of circles’ model; **PF GOC**: phase field ‘gas of circles’ model; **AGOC**: inflection point ‘gas of circles’ model; **PF AGOC**: phase field inflection point ‘gas of circles’ model

The development of our model has been viewed from a modelling viewpoint so far, but we also used different data models to be able to combine the information content of the image with the prior model. First, we presented a data model based on the gradient and intensity values of the most significant spectral-band of the images. Taking advantage of working on multispectral images, we then introduced another data model based on the mean intensity and covariance matrix of tree crown and background statistics.

5.2 Unsolved problems

We describe two unsolved problems of our model in this section. Using the current parameter settings, those circles that get too close to each other tend to merge: first they form a dumbbell shape, then an ellipse shape and finally a circle. This process is illustrated in figure 5.2. In the presence of an image term that is not too strong, circles situated too close to each other fuse and form into a dumbbell shape. We do not offer a solution to the problem, but introduce two possible models which might provide a useful approach.




Figure 5.2: Illustration of the attraction arising between circles. The evolution goes from left to right. Starting from two separated circles with the desired radius, they merge and create a dumbbell shape and later a stable circle.

In the current model, the fusion of the circles generates a decrease in energy. We propose to find parameter settings for which approaching circles would generate an increase in energy, *i.e.* the force between two circles is repulsive. A possible way is to define an extra parameter w , the distance between the circles. Then, adjust the parameters so that the change of global energy w.r.t. w satisfies the requirements. In figure 5.3, we present two simple configurations where the change of the energy can be determined. The upper model illustrates two, not yet fused circles. This ideal configuration rarely occurs in practice, but mathematically can be handled easily. The HOAC energy of the system can be defined rather simply. On separate circles which have the desired radius r_0 , the arising force is zero, otherwise the circles would deform. The distance between the two circles does not affect the area and the length, so the only force on the points of the two circles is from the quadratic energy. This energy can be defined as a function of w . In figure 5.3, the lower image illustrates already fused circles. Besides the value w , the model can be completed with a parameter r_{out} , that describes the radius of the arcs between the centres of the two circles in the dumbbell shape. The local components of the energy (*i.e.* area, region boundary) can easily be determined analytically given the parameters. It is more challenging to calculate the non-local energy but notice that we can divide the contour into four arcs and determine the interactions between the individual arcs and their self-energies. The computations can be simplified using symmetries. We can further simplify the model by supposing that the outer radius is $r_{\text{out}} = r_0$.

The other problem we experienced during our work that although the forest as a coherent region has a very characteristic texture. The size of the individual tree crowns is too small to have useful texture. Therefore, to detect the crowns we use the gradient and the color or intensity.




Figure 5.3: Proposed configurations to determine the repulsive force between circles.

5.3 Possible directions

We can divide the possible directions to follow into three separate groups: the improvement of the data model, the development of the shape prior, and new approaches for modeling and optimization.

In the future, one could expand our data term with a new term, based on color gradient, and incorporate it into the active contour and phase field models. This should offer a more precise segmentation result on those regions where the color information itself is insufficient.

We used a simplification to determine the parameters of the ‘gas of circles’ phase field model (see figure 3.2). For the parameter range we used, we did not meet any problem, but in the case of smaller radii or sharper edges difficulties may occur. In the future one could analyze the parameters more accurately with a more sophisticated approximation taking into account changes of curvature.

The phase field framework improves over active contours from a modeling and as well as from an algorithmic point of view. In the future, it is possible to introduce a new representation using Markov random fields or wavelets as in de Rivaz and Kingsbury (2000).

Presently, the HOAC framework provides ways to model networks and circular objects, but hopefully the modeling of other shape families is also possible. Another important task is to determine inter-circle interactions so that the circles organize into different desired structures. We need to emphasize that while changing the model parameters can result in a slightly different behaviour of the shapes, in many cases it may be better to choose an other interaction function. On the other hand, note that we can achieve stable circles not only with the interaction function defined in equation 2.1.3, but also with other monotonically decreasing functions (*e.g.* exponential or Bessel-type functions).

Finally, we discuss the extension of the model to 3D applications. While it is easy to extend the local terms to higher-dimensions, the extension of the non-local term can create difficulties (*e.g.* extraction of the exact surface). The phase field model can be helpful also in this case. The other difficulty is the stability analysis of spheres and other 3D objects.

Appendix A

Details of stability computations

In this appendix, starting from the equation for the circle and the expression for the radial perturbation in terms of its Fourier coefficients,

$$\gamma(t) = \gamma_0(t) + \delta\gamma(t) = (r(t), \theta(t)) = (r_0(t) + \delta r(t), \theta_0(t)) \quad (\text{A.0.1a})$$

where

$$\gamma_0(t) = (r_0(t), \theta_0(t)) = (r_0, t) \quad (\text{A.0.1b})$$

and

$$\delta r(t) = \sum_k a_k e^{ir_0 k t}, \quad (\text{A.0.1c})$$

with $k \in \{m/r_0 : m \in \mathbb{Z}\}$, we give most of the steps involved in reaching the expression, equation 2.2.6, for the expansion to second order of E_g around a circle.

The derivative of γ is given by

$$\dot{\theta}(t) = 1 \quad (\text{A.0.2a})$$

$$\dot{r}(t) = \dot{\delta r}(t) = \sum_k a_k i r_0 k e^{ir_0 k t}. \quad (\text{A.0.2b})$$

The tangent vector field is given by

$$\dot{\gamma}(t) = \dot{r}(t)\partial_r + \dot{\theta}(t)\partial_\theta. \quad (\text{A.0.3})$$

We need the magnitude of this vector to second order. The metric in polar coordinates is given by $ds^2 = dr^2 + r^2 d\theta^2$, so we have that $|\dot{\gamma}(t)|^2 = \dot{r}(t)^2 + r(t)^2$ by equation (A.0.2a). Substituting from equations (A.0.1) and (A.0.2b) gives

$$|\dot{\gamma}(t)|^2 = r_0^2 + 2r_0 \sum_k a_k e^{ir_0 k t} + \sum_{k,k'} a_k a_{k'} e^{ir_0(k+k')t} (1 - r_0^2 k k'). \quad (\text{A.0.4})$$

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

$$|\dot{\gamma}(t)| = r_0 \left\{ 1 + \sum_k \frac{a_k}{r_0} e^{ir_0 k t} - \frac{1}{2} \sum_{k,k'} a_k a_{k'} k k' e^{ir_0(k+k')t} \right\}. \quad (\text{A.0.5})$$

A.1 Length

Using equation (A.0.5), the boundary length is then given to second order by

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

where we have used the fact that

$$\int_{-\pi}^{\pi} dt e^{ir_0 k t} = 2\pi \delta(k), \quad (\text{A.1.1})$$

and that $a_{-k} = a_k^*$, where $*$ indicates complex conjugation, because δr is real.

A.2 Area

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 equations (A.0.1), and again using equation (A.1.1) to integrate Fourier basis elements, we have that

$$A(\gamma) = \pi r_0^2 + 2\pi r_0 a_0 + \pi \sum_k |a_k|^2. \quad (\text{A.2.1})$$

A.3 Quadratic energy

To compute the expansion of the quadratic term in equation (2.1.2) for E_g , we need the expansions of $\dot{\gamma}(t) \cdot \dot{\gamma}(t')$ and $\Psi(R(t, t'))$.

A.3.1 Inner product of tangent vectors

The tangent vector is given by equation A.0.3, but we must take care as $\dot{\gamma}(t)$ and $\dot{\gamma}(t')$ live in different tangent spaces, at $\gamma(t)$ and $\gamma(t')$ respectively. Since parallel transport does not preserve the coordinate basis vectors ∂_r and ∂_θ , it will change the components of $\dot{\gamma}(t')$, say,

when we parallel transport it to the tangent space at $\gamma(t)$. It is easiest to convert the tangent vectors to the Euclidean coordinate basis,

$$\begin{aligned}\partial_r &= \cos(\theta)\partial_x + \sin(\theta)\partial_y \\ \partial_\theta &= -r\sin(\theta)\partial_x + r\cos(\theta)\partial_y ,\end{aligned}$$

as these basis vectors are preserved by parallel transport. Doing so, and then taking the inner product gives

$$\begin{aligned}\dot{\gamma} \cdot \dot{\gamma}' &= \cos(\theta' - \theta)[r_0^2 + r_0\delta r + r_0\delta r' + \delta r\delta r' + \dot{\delta r}\dot{\delta r}'] \\ &\quad + \sin(\theta' - \theta)[r_0\dot{\delta r}' - r_0\dot{\delta r} + \delta r\dot{\delta r}' - \dot{\delta r}\delta r'] .\end{aligned}$$

where unprimed quantities are evaluated at t and primed quantities at t' . Note that when $t = t'$, the expression reduces to equation (A.0.4).

A.3.2 Distance between two points

The squared distance between $\gamma(t')$ and $\gamma(t)$ is given by

$$\begin{aligned}|\gamma(t') - \gamma(t)|^2 &= (x(t') - x(t))^2 + (y(t') - y(t))^2 \\ &= [(r_0 + \delta r')\cos(\theta') - (r_0 + \delta r)\cos(\theta)]^2 \\ &\quad + [(r_0 + \delta r')\sin(\theta') - (r_0 + \delta r)\sin(\theta)]^2 ,\end{aligned}$$

which after expansion gives

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

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

$$\begin{aligned}R(t, t') &= |\gamma(t') - \gamma(t)| = \\ &= 2r_0|\sin(\Delta t/2)| + |\sin(\Delta t/2)|(\delta r + \delta r') + \frac{A(\Delta t)}{4r_0}(\delta r - \delta r')^2 , \quad (\text{A.3.2})\end{aligned}$$

where $A(z) = \left(\frac{\cos^2(\frac{z}{2})}{|\sin \frac{z}{2}|} \right)$.

A.3.3 Interaction function

Expanding $\Psi(z)$ in a Taylor series to second order, and then substituting $R(t, t')$ for z using the approximation in equation (A.3.2), and keeping only terms up to second order in $\delta\gamma$

then gives

$$\begin{aligned} \Psi(R(t, t')) &= \Psi(X_0) + \left| \sin \frac{\Delta t}{2} \right| \Psi'(X_0)(\delta r + \delta r') \\ &\quad + \frac{1}{4r_0} A(\Delta t) \Psi'(X_0)(\delta r - \delta r')^2 + \frac{1}{2} \sin^2 \left(\frac{\Delta t}{2} \right) \Psi''(X_0)(\delta r + \delta r')^2, \end{aligned} \quad (\text{A.3.3})$$

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

A.3.4 Combining terms

Now let $G(t, t') = \dot{\gamma}(t) \cdot \dot{\gamma}(t') \Psi(R(t, t'))$. Combining the expressions already derived, we have

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

where we have introduced the notation $F_{00} \dots F_{24}$ for the functions appearing in the terms of G , and ‘odd’ and ‘even’ refer to parity under exchange of t and t' . Note that the F are functionals of Ψ , and functions of r_0 and $t' - t$ (but not t and t' separately). Note also that each line, and hence G , is symmetric in t and t' .

The integral in the quadratic energy term is now given by $\iint_{-\pi}^{\pi} dt dt' G(t, t')$. We can now substitute the expressions for δr and $\dot{\delta} r$ in terms of their Fourier coefficients, $\delta r(t) = \sum_k a_k e^{ir_0 kt}$ and $\dot{\delta} r(t) = \sum_k a_k i r_0 k e^{ir_0 kt}$. Due to the dependence of the F on $t - t'$ only, the resulting integrals can be reduced, via a change of variables $p = t' - t$, to integrals over p . 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. For example,

$$\iint_{-\pi}^{\pi} dt dt' (\delta r \dot{\delta} r' - \dot{\delta} r \delta r') F_{23}(t' - t) = 2 \iint_{-\pi}^{\pi} dt dt' \delta r \dot{\delta} r' F_{23}(t' - t) .$$

We therefore only need to evaluate one of these integrals for the relevant terms. Below we list the calculations for all the F integrals for completeness:

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

which survives because F_{00} is symmetric;

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

which survives because F_{10} is symmetric;

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \dot{\delta}r(t) F_{11}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k a_k i r_0 k e^{i r_0 k t} F_{11}(t' - t) \\
&= \sum_k a_k i r_0 k \iint_{-\pi}^{\pi} dp dt' e^{i r_0 k (-p+t')} F_{11}(p) \\
&= \sum_k a_k i r_0 k \int_{-\pi}^{\pi} dt' e^{i r_0 k t'} \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{11}(p) \\
&= \sum_k a_k i r_0 k 2\pi \delta(k) \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{11}(p) \\
&= 0 ;
\end{aligned}$$

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \delta r^2(t) F_{20}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k \sum_{k'} a_k a_{k'} e^{i r_0 (k+k') t} F_{20}(t' - t) \\
&= \sum_k \sum_{k'} a_k a_{k'} \iint_{-\pi}^{\pi} dp dt' e^{i r_0 (k+k') (-p+t')} F_{20}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} \int_{-\pi}^{\pi} dt' e^{i r_0 (k+k') t'} \int_{-\pi}^{\pi} dp e^{-i r_0 (k+k') p} F_{20}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} 2\pi \delta(k+k') \int_{-\pi}^{\pi} dp e^{-i r_0 (k+k') p} F_{20}(p) \\
&= \sum_k a_k a_{-k} 2\pi \int_{-\pi}^{\pi} dp F_{20}(p) \\
&= 2\pi \sum_k |a_k|^2 \int_{-\pi}^{\pi} dp F_{20}(p) ,
\end{aligned}$$

which survives because F_{20} is symmetric;

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \delta r(t) \delta r(t') F_{21}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k \sum_{k'} a_k a_{k'} e^{ir_0(kt+k't')} F_{21}(t' - t) \\
&= \sum_k \sum_{k'} a_k a_{k'} \iint_{-\pi}^{\pi} dp dt' e^{ir_0k(-p+t')} e^{ir_0k't'} F_{21}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} \int_{-\pi}^{\pi} dt' e^{ir_0(k+k')t'} \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{21}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} 2\pi \delta(k+k') \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{21}(p) \\
&= \sum_k a_k a_{-k} 2\pi \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{21}(p) \\
&= 2\pi \sum_k |a_k|^2 \int_{-\pi}^{\pi} dp e^{-ir_0kp} F_{21}(p) ;
\end{aligned}$$

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \delta r(t) \dot{\delta} r(t) F_{22}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k \sum_{k'} a_k a_{k'} ir_0 k e^{ir_0(k+k')t} F_{22}(t' - t) \\
&= \sum_k \sum_{k'} a_k a_{k'} ir_0 k \iint_{-\pi}^{\pi} dp dt' e^{ir_0(k+k')(-p+t')} F_{22}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} ir_0 k \int_{-\pi}^{\pi} dt' e^{ir_0(k+k')t'} \int_{-\pi}^{\pi} dp e^{-ir_0(k+k')p} F_{22}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} ir_0 k 2\pi \delta(k+k') \int_{-\pi}^{\pi} dp e^{-ir_0(k+k')p} F_{22}(p) \\
&= 0 ,
\end{aligned}$$

because with $k + k' = 0$ from the delta function, the integral becomes one over F_{22} only, which vanishes due to the antisymmetry of F_{22} ;

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \delta r(t) \dot{\delta r}(t') F_{23}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k \sum_{k'} a_k a_{k'} i r_0 k' e^{i r_0 (k t + k' t')} F_{23}(t' - t) \\
&= \sum_k \sum_{k'} a_k a_{k'} i r_0 k' \iint_{-\pi}^{\pi} dp dt' e^{i r_0 (k(-p+t') + k' t')} F_{23}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} i r_0 k' \int_{-\pi}^{\pi} dt' e^{i r_0 (k+k') t'} \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{23}(p) \\
&= \sum_k \sum_{k'} a_k a_{k'} i r_0 k' 2\pi \delta(k+k') \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{23}(p) \\
&= -2\pi \sum_k |a_k|^2 i r_0 k \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{23}(p) ;
\end{aligned}$$

$$\begin{aligned}
& \iint_{-\pi}^{\pi} dt dt' \dot{\delta r}(t) \delta r(t') F_{24}(t' - t) \\
&= \iint_{-\pi}^{\pi} dt dt' \sum_k \sum_{k'} a_k a_{k'} i^2 r_0^2 k k' e^{i r_0 (k t + k' t')} F_{24}(t' - t) \\
&= - \sum_k \sum_{k'} a_k a_{k'} r_0^2 k k' \iint_{-\pi}^{\pi} dp dt' e^{i r_0 (k(-p+t') + k' t')} F_{24}(p) \\
&= - \sum_k \sum_{k'} a_k a_{k'} r_0^2 k k' \int_{-\pi}^{\pi} dt' e^{i r_0 (k+k') t'} \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{24}(p) \\
&= - \sum_k \sum_{k'} a_k a_{k'} r_0^2 k k' 2\pi \delta(k+k') \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{24}(p) \\
&= 2\pi \sum_k |a_k|^2 r_0^2 k^2 \int_{-\pi}^{\pi} dp e^{-i r_0 k p} F_{24}(p) .
\end{aligned}$$

Using these results then gives equation 2.2.5, which in combination with equations 2.2.3 and 2.2.4, gives equation 2.2.6.

Appendix B

Polynomial approximation

We present the scaling analysis of the G_{00} , G_{10} and \tilde{G} functions which allow us to generalize our approximation to the complete domain of r , and we compute a polynomial approximation of the energy function using a Taylor expansion around $r = 0$.

B.1 Scaling analysis of the circle energy

G_{00} , G_{10} and \tilde{G} are functions of r and d . We can write these functions in an equivalent form so that they only depend upon the ratio r/d . Since d_{\min} and d_{\max} are the first roots of G_{00} and \tilde{G} respectively, using the scaling property, it is sufficient to determine the roots for a single value of r and it generalizes to the complete domain of radii.

The interaction function in equation 2.1.3 is a function of x/d , while in G_{00} the variable of the interaction function is a function of r and a parameter, so we can set

$$\Psi\left(2r \sin\left|\frac{p}{2}\right|\right) = \hat{\Psi}\left(\frac{r}{d}\right). \quad (\text{B.1.1})$$

We define a new \hat{G}_{00} function from G_{00} which is also just a function of $\frac{r}{d}$.

$$G_{00}(r) = d^2 \int_{-\pi}^{\pi} dp \cos(p) \left(\frac{r}{d}\right)^2 \hat{\Psi}\left(\frac{r}{d}\right) = d^2 \hat{G}_{00}\left(\frac{r}{d}\right). \quad (\text{B.1.2})$$

If $G_{00}(r, d) = 0$ for a given r and d than $G_{00}(ar, ad) = 0$ for all $a \in \mathbb{R}$. We require \tilde{G} and $G_{10} - r\tilde{G}$ for further analysis. \tilde{G} can be transformed

$$\tilde{G}(r) = \frac{1}{2} \left(\frac{\partial}{\partial r}\right)^2 G_{00}(r) = \frac{d^2}{2} \left(\frac{\partial}{\partial r}\right)^2 \hat{G}_{00}\left(\frac{r}{d}\right) = \hat{\tilde{G}}\left(\frac{r}{d}\right), \quad (\text{B.1.3})$$

where $\hat{\tilde{G}}(z) = \frac{1}{2} \left(\frac{\partial}{\partial z}\right)^2 \hat{G}_{00}(z)$. G_{10} can be written as

$$G_{10}(r) = \frac{1}{2} \frac{\partial}{\partial r} G_{00}(r) = \frac{d^2}{2} \frac{\partial}{\partial r} \hat{G}_{00}\left(\frac{r}{d}\right) = d \hat{G}_{10}\left(\frac{r}{d}\right), \quad (\text{B.1.4})$$

where $\hat{G}_{10}(z) = \frac{1}{2} \frac{\partial}{\partial z} \hat{G}_{00}(z)$. Finally

$$G_{10}(r) - r\tilde{G}(r) = d\hat{G}_{10}\left(\frac{r}{d}\right) - r\hat{G}\left(\frac{r}{d}\right) = d\left(\hat{G}_{10}\left(\frac{r}{d}\right) - \left(\frac{r}{d}\right)\hat{G}\left(\frac{r}{d}\right)\right). \quad (\text{B.1.5})$$

B.2 The Taylor-polynomial of the energy function

We present a polynomial approximation to the above functions and determine the roots for one r value; using the scaling property we can than determine them for any other radii. To find the critical points d_{\min} and d_{\max} we determine the zero-crossings of G_{10} , and \tilde{G} , which are the derivatives of G_{00} . We use a polynomial approximation around zero. Here we give the approximation of the function G_{00} , which can be written as:

$$G_{00}(r) = \int_{-\pi}^{\pi} dp \cos(p) r^2 \Psi\left(2r \left|\sin \frac{p}{2}\right|\right), \quad (\text{B.2.1a})$$

$$= \sum_{n=0}^{\infty} b_n r^n \quad (\text{B.2.1b})$$

To determine the coefficients b_n , we use

$$\left(\frac{\partial}{\partial r}\right)^m G_{00}(r) \Big|_{r=0} = b_m \cdot m!. \quad (\text{B.2.2})$$

Thus,

$$b_m = \frac{1}{m!} \left(\frac{\partial}{\partial r}\right)^m G_{00}(r). \quad (\text{B.2.3})$$

In the following subsections we will give the derivatives of G_{00} . For this we define the derivatives of the interaction function, the argument of the interaction function and finally G_{00} itself. We mention that if we use a new interaction function, we need to compute the derivatives of the function and substitute, following the steps below.

B.2.1 The interaction function

We compute the derivatives of the interaction function in equation 2.1.3 on the $x < 2d$ domain

$$\Psi^{(1)}(x) = \frac{1}{2d} \left(-1 + \cos \frac{\pi x}{d}\right). \text{ In general, if } i > 1: \quad (\text{B.2.4})$$

$$\Psi^{(i)}(x) = \begin{cases} \frac{1}{2d} \left(\frac{\pi}{d}\right)^{i-1} \sin \frac{\pi x}{d} & \text{if } (i \bmod 4) = 0, \\ \frac{1}{2d} \left(\frac{\pi}{d}\right)^{i-1} \cos \frac{\pi x}{d} & \text{if } (i \bmod 4) = 1, \\ -\frac{1}{2d} \left(\frac{\pi}{d}\right)^{i-1} \sin \frac{\pi x}{d} & \text{if } (i \bmod 4) = 2, \\ -\frac{1}{2d} \left(\frac{\pi}{d}\right)^{i-1} \cos \frac{\pi x}{d} & \text{if } (i \bmod 4) = 3. \end{cases}$$

The derivatives of the variable X_0 of the interaction function w.r.t. r are

$$\begin{aligned} X_0(r) &= 2r \left| \sin \frac{p}{2} \right| \\ X_0^{(1)}(r) &= 2 \left| \sin \frac{p}{2} \right| \\ X_0^{(n)}(r) &= 0, \quad \text{if } n \geq 2. \end{aligned} \quad (\text{B.2.5})$$

B.2.2 The derivatives of G_{00}

The derivatives of G_{00} defined in B.2.1b

$$\left(\frac{\partial}{\partial r} \right)^m G_{00}(r) = \int_0^{2\pi} dp \cos(p) \left(\frac{\partial}{\partial r} \right)^m [r^2 \Psi(X_0(r))] . \quad (\text{B.2.6})$$

To determine the derivatives, we define the derivative of the product of two functions in general and substitute the factors of G_{00} into the expression. Let X and Y the functions, the derivatives of their product can be defined as:

$$\begin{aligned} \left(\frac{\partial}{\partial r} \right)^m [X(r)Y(r)] &= \sum_{i=0}^m \binom{m}{i} \left(\frac{\partial}{\partial r} \right)^i X(r) \left(\frac{\partial}{\partial r} \right)^{m-i} Y(r) \\ &= X^{(0)}(r)Y^{(m)}(r) + mX^{(1)}(r)Y^{(m-1)}(r) + \\ &\quad \frac{m(m-1)}{2} X^{(2)}(r)Y^{(m-2)}(r) + \dots \end{aligned} \quad (\text{B.2.7})$$

We can use this formula substituting $X(r) = r^2$ and $Y(r) = \Psi(X_0(r))$. The derivatives of $X(r)$ are computed by definition. The derivatives of $Y(r)$ function are:

$$\begin{aligned} Y(r) &= \Psi(X_0(r)) , \\ Y^{(1)}(r) &= \Psi^{(1)}(X_0(r))X_0^{(1)}(r) , \\ Y^{(2)}(r) &= \Psi^{(2)}(X_0(r)) \left[X_0^{(1)}(r) \right]^2 , \\ &\vdots \qquad \qquad \qquad \vdots \end{aligned} \quad (\text{B.2.8})$$

$$Y^{(n)}(r) = \Psi^{(n)}(X_0(r)) \left[X_0^{(1)}(r) \right]^n , \quad (\text{B.2.9})$$

where $\Psi^{(n)}$ and $X_0^{(m)}$ are defined in equations B.2.4 and B.2.5. Using B.2.7 we can write the derivatives of the product as

$$\left(\frac{\partial}{\partial r} \right)^m [r^2 \Psi(X_0(r))] = r^2 Y^{(m)}(r) + 2mr Y^{(m-1)}(r) + m(m-1) Y^{(m-2)}(r) . \quad (\text{B.2.10})$$

After substituting back into the product, we can write G_{00} using equation B.2.6 as

$$\left(\frac{\partial}{\partial r}\right)^m G_{00}(r) = \int_0^{2\pi} dp \cos(p) \left[r^2 Y^{(m)}(r) + 2mr Y^{(m-1)}(r) + m(m-1) Y^{(m-2)}(r) \right]. \quad (\text{B.2.11})$$

We then find that

$$\left(\frac{\partial}{\partial r}\right)^m G_{00}(r) \Big|_{r=0} = \int_0^{2\pi} dp \cos(p) m(m-1) Y^{(m-2)}(0), \quad (\text{B.2.12})$$

where $Y^{(m-2)}(0)$ can be defined substituting $r = 0$ back to the form given in equation B.2.8.

$$\begin{aligned} Y^{(m-2)}(0) &= \Psi^{(m-2)}(X_0(0)) [X_0^{(1)}(0)]^{m-2} \\ &= \Psi^{(m-2)}(0) \left[2 \left| \sin \frac{p}{2} \right| \right]^{m-2} \\ &= \begin{cases} 0 & \text{if } m \text{ even,} \\ \left[2 \left| \sin \frac{p}{2} \right| \right]^{m-2} \cdot (-1)^{\frac{m-3}{2}} \frac{1}{2d} \left(\frac{\pi}{d} \right)^{m-3} & \text{if } m \text{ odd.} \end{cases} \end{aligned} \quad (\text{B.2.13})$$

B.2.3 The polynomial coefficients

In this subsection, we give a closed form for the polynomial coefficients we defined in equation B.2.3. Using this equation and the result of B.2.12 we can express the coefficients as:

$$b_m = \frac{1}{m!} \left(\frac{\partial}{\partial r}\right)^m G_{00}(r) \Big|_{r=0} = \left(\frac{m(m-1)}{m!} \right) \int_0^{2\pi} dp \cos(p) Y^{(m-2)}(0). \quad (\text{B.2.14})$$

In equation B.2.13, we defined the evaluation of $Y^{(i)}$ at 0, using this we can give the form for b_m for odd m values as

$$\begin{aligned} b_m &= \frac{m(m-1)}{m!} \int_0^{2\pi} dp \cos(p) (-1)^{\frac{m-3}{2}} \frac{1}{2d} \left(\frac{\pi}{d} \right)^{m-3} \left[2 \left| \sin \frac{p}{2} \right| \right]^{m-2} \\ &= (-1)^{\frac{m-3}{2}} \frac{m(m-1)(2\pi)^{m-3}}{m! \cdot d^{m-2}} \int_0^{2\pi} dp \cos(p) \left| \sin \frac{p}{2} \right|^{m-2}. \end{aligned} \quad (\text{B.2.15})$$

We note, that we can use $\sin \frac{p}{2}$ instead of its absolute value. The integral is given by

$$\begin{aligned}
\int_0^{2\pi} dp \cos(p) \left(\sin \frac{p}{2}\right)^{m-2} &= \int_0^{2\pi} dp \left(1 - 2 \sin^2 \frac{p}{2}\right) \left(\sin \frac{p}{2}\right)^{m-2} \\
&= \int_0^{2\pi} dp \left(\sin \frac{p}{2}\right)^{m-2} - 2 \int_0^{2\pi} dp \left(\sin \frac{p}{2}\right)^m \\
\frac{p}{2} = q \text{ and } dp &= 2dq \\
&= 4 \left\{ \int_0^{\frac{\pi}{2}} dq (\sin q)^{m-2} - 2 \int_0^{\frac{\pi}{2}} dq (\sin q)^m \right\} \\
\text{using that } \int_0^{\frac{\pi}{2}} dq (\sin q)^m &= \frac{(m-1)!!}{m!!} \\
&= 4 \left\{ \frac{(m-3)!!}{(m-2)!!} - 2 \frac{(m-1)!!}{m!!} \right\} \\
&= \frac{4}{m!!} \{m(m-3)!! - 2(m-1)!!\} . \tag{B.2.16}
\end{aligned}$$

The closed form for b_m is then

$$b_m = (-1)^{\frac{m-1}{2}} \frac{4(2\pi)^{m-3}}{m!!(m-4)!!} \frac{1}{d^{m-2}} . \tag{B.2.17}$$

B.2.4 Polynomial approximation of G_{00} , G_{10} and \tilde{G}

Using equations B.2.1b and B.2.17 the polynomial approximation of the G_{00} function can be written as

$$G_{00}(r) = 4d^2 \sum_{\substack{m \geq 5 \\ m \text{ odd}}} \frac{(-1)^{\frac{m-1}{2}} (2\pi)^{m-3}}{m!!(m-4)!!} \left(\frac{r}{d}\right)^m . \tag{B.2.18}$$

$G_{10} = \frac{1}{2} \frac{\partial}{\partial r} G_{00}$ is then given by

$$G_{10}(r) = 2d \sum_{\substack{m \geq 4 \\ m \text{ odd}}} \frac{(-1)^{\frac{m-3}{2}} (2\pi)^{m-1}}{m!!(m-4)!!} \left(\frac{r}{d}\right)^{m-1} , \tag{B.2.19}$$

while $\tilde{G} = \frac{\partial}{\partial r} G_{10}$ is

$$\tilde{G}(r) = 2 \sum_{\substack{m \geq 3 \\ m \text{ odd}}} \frac{(-1)^{\frac{m+1}{2}} (2\pi)^{m-1} (m+1)}{m!!(m-2)!!} \left(\frac{r}{d}\right)^m . \tag{B.2.20}$$

We use the above functions to determine the coefficients of the polynomials which approximates $\tilde{G}(r)$ and $G_{10}(r)p - \tilde{G}(r)$, and then we compute their roots.

Publications and scientific activities of the author

International journals

- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. A higher-order active contour model of a ‘gas of circles’ and its application to tree crown extraction. *Submitted to Pattern Recognition*, November 2006d

International conferences

- P. Horváth and I. H. Jermyn. A ‘gas of circles’ phase field model and its application to tree crown extraction. In *Proc. European Signal Processing Conference (EUSIPCO)*, Poznan, Poland, September 2007b
- P. Horváth and I. H. Jermyn. A new phase field model of a ‘gas of circles’ for tree crown extraction from aerial images. In *Proc. International Conference on Computer Analysis of Images and Patterns (CAIP)*, Lecture Notes in Computer Science, Vienna, Austria, August 2007a
- P. Horváth. A multispectral data model for higher-order active contours and its application to tree crown extraction. In *Proc. Advanced Concepts for Intelligent Vision Systems*, Lecture Notes in Computer Science, Delft, Netherlands, August 2007
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. An improved ‘gas of circles’ higher-order active contour model and its application to tree crown extraction. In *Proc. Indian Conference on Vision, Graphics and Image Processing (ICVGIP)*, Lecture Notes in Computer Science, Madurai, India, December 2006c
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. A higher-order active contour model for tree detection. In *Proc. International Conference on Pattern Recognition (ICPR)*, Hong Kong, China, August 2006b
- P. Horváth, A. Bhattacharya, I. H. Jermyn, J. Zerubia, and Z. Kato. Shape moments for region-based active contours. In *Proc. Hungarian-Austrian Conference on Image Processing and Pattern Recognition*, Veszprém, Hungary, May 2005

Research reports

- P. Horváth, I. H. Jermyn, Z. Kato, and 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 2006a

National conferences

- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. Circular object segmentation using higher-order active contours. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 133–141, Debrecen, Hungary, January 2007. In Hungarian
- P. Horváth and Z. Kato. Optical flow computation using an energy minimization approach. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 125–130, Miskolc-Tapolca, Hungary, January 2004b. In Hungarian, non-refereed
- P. Horváth and Z. Kato. Color, texture and motion segmentation using gradient vector flow. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 131–137, Miskolc-Tapolca, Hungary, January 2004a. In Hungarian, non-refereed

Invited talks

- Tree Crown Segmentation, Hungarian Forest Inventory (new name: Hungarian Central Agricultural Office, Forestry Administration (CAO, FA)), Budapest, Hungary, May 2006.
- Shape Priors for Variational Image Segmentation - Higher-Order Active Contour Model for Tree Detection, Pondicherry University and French Institute of Pondicherry, Pondicherry, India, December 2006.

Software licence

- PHASECIRCLE v1.0 – software deposited to the APP (Agence pour la Protection des Programmes) in 2007 under the number IDDN-FR-001-280029-000-S-C-2007-000-21000, and transferred to the Joint Research Center (JRC) of the European Union in Ispra, Italy and to the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) in Budapest, Hungary.

Bibliography

- D. Adalsteinsson and J. A. Sethian. The fast construction of extension velocities in level set methods. *Journal of Computational Physics*, 148:2–22, 1999.
- S. Allen and J. W. Cahn. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metallurgica*, 27:1084–1095, 1979.
- L. Ambrosio and V. M. Tortorelli. Approximation of functional depending on jumps by elliptic functional via t-convergence. *Communications on Pure and Applied Mathematics*, 43:999–1036, 1990.
- G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing*. Springer-Verlag, 2002.
- G. Aubert, J.F. Aujol, and L. Blanc-Féraud. Detecting codimension-two objects in an image with Ginzburg-Landau models. Research Report 5254, INRIA, France, July 2004.
- R. J. Baxter. *Exactly Solved Models in Statistical Mechanics*. Academic Press, 1990.
- A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, MA, U.S.A., 1987.
- C. H. Blazquez. Computer-based image analysis and tree counting with aerial color infrared photography. *Journal of Imaging Technology*, 15:163–168, 1989.
- W. J. Boettinger, J. A. Warren, C. Beckermann, and A. Karma. Phase-field simulation of solidification. *Annual Review of Materials Research*, 32:163–194, 2002. doi: 10.1146/annurev.matsci.32.101901.155803.
- T. Brandtberg. *Automatic Individual Tree-Based Analysis of High Spatial Resolution Remotely Sensed Data*. PhD thesis, Swedish University of Agricultural Sciences, 1999.
- T. Brandtberg and F. Walter. Automated delineation of individual tree crowns in high spatial resolution aerial images by multiple-scale analysis. *Machine Vision and Applications*, (11):64–73, 1998.
- X. Bresson. *Image Segmentation with Variational Active Contours*. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2005.

- V. Caselles, F. Catte, T. Coll, and F. Dibos. A geometric model for active contours. *Numerische Mathematik*, 66:1–31, 1993.
- V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997a.
- V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. A three dimensional segmentation approach. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19(4):394 – 398, 1997b.
- A. Chambolle. Image segmentation by variational methods: Mumford and Shah functional and the discrete approximations. *SIAM Journal of Applied Mathematics*, 55(3):827–863, June 1995. ISSN 0036-1399 (print), 1095-712X (electronic).
- T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. Image Processing*, 10(2):266–277, 2001.
- T. F. Chan and L. A. Vese. A multiphase level set framework for image segmentation using the Mumford and Shah model. *International Journal of Computer Vision*, pages 271 – 293, 2002.
- Y. Choquet-Bruhat, C. DeWitt-Morette, and M. Dillard-Bleick. *Analysis, Manifolds and Physics*. Elsevier Science, Amsterdam, The Netherlands, 1996.
- L.D. Cohen. On active contours and balloons. *CVGIP: Image Understanding*, 53:211–218, 1991.
- T.F. Cootes and C.J. Taylor. Active shape models - smart snakes. In *Proc. British Machine Vision Conference (BMVC)*, 1992.
- D. Cremers and 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*, pages 169–176, Nice, France, 2003.
- D. Cremers, C. Schnorr, and J. Weickert. Diffusion-snakes: combining statistical shape knowledge and image information in a variational framework. In *Proc. IEEE Workshop Variational, Geometric and Level Set Methods in Computer Vision*, pages 137–144, Vancouver, BC, Canada, 2001.
- D. Cremers, F. Tischhäuser, J. Weickert, and C. Schnörr. Diffusion snakes: Introducing statistical shape knowledge into the Mumford-Shah functional. *International Journal of Computer Vision*, 50(3):295–313, 2002.
- D. Cremers, N. Sochen, and C. Schnörr. Multiphase dynamic labeling for variational recognition-driven image segmentation. In V. Hlavac, editor, *Proc. European Conference on Computer Vision (ECCV)*, volume 3024 of *LNCS*, pages 74–86, Prague, 2004. Springer.

- D. Cremers, N. Sochen, and C. Schnörr. A multiphase dynamic labeling model for variational recognition-driven image segmentation. *International Journal of Computer Vision*, 66(1):67–81, 2006.
- P. de Rivaz and N. Kingsbury. Fast segmentation using level set curves of complex wavelet surfaces. In *Proc. IEEE International Conference on Image Processing (ICIP)*, pages Vol III: 592–595, 2000.
- K. Dralle and M. Rudemo. Stem number estimation by kernel smoothing of aerial photos. *Canadian Journal of Forest Research*, 26:1228–1236, 1996.
- D. Eberly, R. Gardner, B. Morse, S. Pizer, and C. Scharlach. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4(4):353–373, 1994.
- M. Erikson. Segmentation of individual tree crowns in colour aerial photographs using region growing supported by fuzzy rules. *Canadian Journal of Forest Research*, 33: 1557–1563(7), 2003.
- M. Erikson. *Segmentation and Classification of Individual Tree Crowns*. PhD thesis, Swedish University of Agricultural Sciences, 2004.
- A. Foulonneau, P. Charbonnier, and F. Heitz. Geometric shape priors for region-based active contours. *Proc. IEEE International Conference on Image Processing (ICIP)*, 3:413–416, 2003.
- A. Foulonneau, P. Charbonnier, and F. Heitz. Affine-invariant geometric shape priors for region-based active contours. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(8):1352–1357, 2006.
- S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- V. Ginzburg and L. Landau. *On the theory of superconductivity*. Zh. Eksper. Teo. Fiz., 1950.
- R. C. Gonzalez and R. E. Woods. *Digital image processing*. Addison-Wesley, 2002.
- F. A. Gougeon. A crown-following approach to the automatic delineation of individual tree crowns in high spatial resolution aerial images. *Canadian Journal of Remote Sensing*, 21(3), pages 274–284, 1995.
- H. Grossauer and O. Scherzer. Using the complex Ginzburg-Landau equation for digital inpainting in 2D and 3D. In *Proc. Scale-Space*, pages 225–236. LNCS 2695, Springer-Verlag, 2003.
- S.C. Gupta. *The Classical Stefan Problem*. North-Holland Series in Applied Mathematics and Mechanics. Elsevier Science, Amsterdam, The Netherlands, 2003.

- P. Horváth. A multispectral data model for higher-order active contours and its application to tree crown extraction. In *Proc. Advanced Concepts for Intelligent Vision Systems*, Lecture Notes in Computer Science, Delft, Netherlands, August 2007.
- P. Horváth and I. H. Jermyn. A new phase field model of a ‘gas of circles’ for tree crown extraction from aerial images. In *Proc. International Conference on Computer Analysis of Images and Patterns (CAIP)*, Lecture Notes in Computer Science, Vienna, Austria, August 2007a.
- P. Horváth and I. H. Jermyn. A ‘gas of circles’ phase field model and its application to tree crown extraction. In *Proc. European Signal Processing Conference (EUSIPCO)*, Poznan, Poland, September 2007b.
- P. Horváth and Z. Kato. Color, texture and motion segmentation using gradient vector flow. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 131–137, Miskolc-Tapolca, Hungary, January 2004a. In Hungarian, non-refereed.
- P. Horváth and Z. Kato. Optical flow computation using an energy minimization approach. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 125–130, Miskolc-Tapolca, Hungary, January 2004b. In Hungarian, non-refereed.
- P. Horváth, A. Bhattacharya, I. H. Jermyn, J. Zerubia, and Z. Kato. Shape moments for region-based active contours. In *Proc. Hungarian-Austrian Conference on Image Processing and Pattern Recognition*, Veszprém, Hungary, May 2005.
- P. Horváth, I. H. Jermyn, Z. Kato, and 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 2006a.
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. A higher-order active contour model for tree detection. In *Proc. International Conference on Pattern Recognition (ICPR)*, Hong Kong, China, August 2006b.
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. An improved ‘gas of circles’ higher-order active contour model and its application to tree crown extraction. In *Proc. Indian Conference on Vision, Graphics and Image Processing (ICVGIP)*, Lecture Notes in Computer Science, Madurai, India, December 2006c.
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. A higher-order active contour model of a ‘gas of circles’ and its application to tree crown extraction. *Submitted to Pattern Recognition*, November 2006d.
- P. Horváth, I. H. Jermyn, Z. Kato, and J. Zerubia. Circular object segmentation using higher-order active contours. In *Conference of the Hungarian Association for Image Analysis and Pattern Recognition*, pages 133–141, Debrecen, Hungary, January 2007. In Hungarian.

- R. Huang, V. Pavlovic, and D. N. Metaxas. A graphical model framework for coupling MRFs and deformable models. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, Washington, DC, USA, 2004.
- E. Ising. Beitrag zur Theorie der Ferromagnetismus. *Z. Physik*, 31:253–258, 1925.
- I. H. Jermyn. *On the Use of Functionals on Boundaries in Hierarchical Models of Object Recognition*. PhD thesis, New York University, Department of Computer Science, 2000.
- M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- Z. Kato. *Modélisations markoviennes multirésolutions en vision par ordinateur. Application à la segmentation d’images SPOT*. PhD thesis, Université de Nice-Sophia-Antipolis, December . Available in French and English.
- Z. Kato, M. Berthod, and J. Zerubia. A hierarchical Markov random field model and multi-temperature annealing for parallel image classification. *Computer Vision, Graphics and Image Processing: Graphical Models and Image Processing*, 58(1):18–37, January 1996.
- S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 810–815, Boston, MA, USA, 1995.
- M. Larsen. Finding an optimal match window for Spruce top detection based on an optical tree model. In D.A. Hill and D.G. Leckie, editors, *Proc. of the International Forum on Automated Interpretation of High Spatial Resolution Digital Imagery for Forestry*, pages 55–66, Victoria, British Columbia, Canada, February 1998.
- M. Larsen. Individual Tree Top Position Estimation by Template Voting. In *Proc. of the Fourth International Airborne Remote Sensing Conference and Exhibition / 21st Canadian Symposium on Remote Sensing*, volume 2, pages 83–90, Ottawa, Ontario, June 1999.
- M. Leventon, O. Faugeras, W. E. L. Grimson, and W. M. Wells. Level set based segmentation with intensity and curvature priors. In *Mathem. Meth. In Biomed. Image Anal.*, 2000a.
- M. Leventon, W. E. L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 316–322, Hilton Head Island, South Carolina, USA, 2000b.
- F. Leymarie. *Three-Dimensional Shape Representation via Shock Flows*. PhD thesis, Brown University, Providence, RI, USA, May 2003.
- T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):117–154, 1998.

- R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17:158–175, 1995.
- D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 22–26, 1985.
- D. Mumford and J. Shah. Optimal approximation by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–684, 1989.
- S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging, Vision, and Graphics*. Springer, New York, 2003.
- S. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- N. Paragios. *Geodesic Active Regions and Level Set Methods: Contributions and Applications in Artificial Vision*. PhD thesis, Université de Nice-Sophia-Antipolis, 2000.
- N. Paragios and R. Deriche. Geodesic active regions: A new framework to deal with frame partition problems in computer vision. *Journal of Visual Communication and Image Representation*, 13:249–268, 2002a.
- N. Paragios and R. Deriche. Geodesic active regions and level set methods for supervised texture segmentation. *International Journal of Computer Vision*, 46:223–247, 2002b.
- N. Paragios and M. Rousson. Shape priors for level set representations. In *Proc. European Conference on Computer Vision (ECCV)*, pages 78–92, Copenhagen, Denmark, 2002.
- T. Pavlidis. *Algorithms for Graphics and Image Processing*, chapter 7. Springer Verlag, Computer Science Press, Inc., 1982.
- G. Perrin. *Etude du couvert forestier par processus ponctuels marqués*. PhD thesis, Ecole Centrale de Paris, October 2006.
- G. Perrin, X. Descombes, and J. Zerubia. Tree crown extraction using marked point processes. In *Proc. European Signal Processing Conference (EUSIPCO)*, September 2004.
- G. Perrin, X. Descombes, and J. Zerubia. A marked point process model for tree crown extraction in plantations. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Genova, Italy, September 2005.
- R. Pollock. *The automatic recognition of individual trees in aerial images of forests based on a synthetic tree crown image model*. PhD thesis, University of British Columbia, Vancouver, Canada, 1996.

- R. Pollock. Individual Tree Recognition based on Synthetic Tree Crowns Image Model. In D.A. Hill and D.G. Leckie, editors, *Proc. of the International Forum on Automated Interpretation of High Spatial Resolution Digital Imagery for Forestry*, pages 25–34, Victoria, British Columbia, Canada, February 1998.
- R. B. Potts. Some generalized order-disorder transformations. In *Cambridge Philosophical Society*, volume 48, pages 106–109, 1952.
- S. J. Reeves. Fast image restoration without boundary artifacts. *IEEE Trans. Image Processing*, 14(10):1448–1453, 2005.
- T. Riklin-Raviv, N. Kiryati, and N. Sochen. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, Beijing, China, 2005.
- T. Riklin-Raviv, N. Kiryati, and N. Sochen. Prior-based segmentation and shape registration in the presence of perspective distortion. *International Journal of Computer Vision*, 72(3):309–328, 2007.
- M. Rochery. *Contours actifs d'ordre supérieur et leur application à la détection de linéiques dans des images de télédétection*. PhD thesis, Université de Nice, Sophia Antipolis, September 2005.
- M. Rochery, I. H. Jermyn, and J. Zerubia. Higher order active contours and their application to the detection of line networks in satellite imagery. In *Proc. IEEE Workshop Variational, Geometric and Level Set Methods in Computer Vision*, Nice, France, October 2003.
- M. Rochery, I. H. Jermyn, and J. Zerubia. Gap closure in (road) networks using higher-order active contours. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Singapore, October 2004.
- M. Rochery, I. H. Jermyn, and J. Zerubia. New higher-order active contour energies for network extraction. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Genoa, Italy, September 2005a.
- M. Rochery, I. H. Jermyn, and J. Zerubia. Phase field models and higher-order active contours. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, Beijing, China, October 2005b.
- M. Rochery, I. H. Jermyn, and J. Zerubia. Higher order active contours. Research Report 5656, INRIA, France, August 2005c.
- M. Rochery, I.H. Jermyn, and J. Zerubia. Higher order active contours. *International Journal of Computer Vision*, 69(1):27–42, August 2006.
- M. Rousson and N. Paragios. Shape priors for level set representations. In *Proc. European Conference on Computer Vision (ECCV)*, Lecture Notes in Computer Science, pages 78–92, Copenhagen, Denmark, 2002. Springer.

- M. Rousson and N. Paragios. Prior knowledge, level set representations & visual grouping. *International Journal of Computer Vision*, 2007. In press.
- C. Samson, L. Blanc-Féraud, G. Aubert, and J. Zerubia. A variational model for image classification and restoration. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(5):460–472, May 2000.
- J. A. Sethian. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Geometry Fluid Mechanics, Computer Vision and Materials Science*. Cambridge University Press, 1999.
- C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.
- G. Shao and K. M. Reynolds. *Computer Applications in Sustainable Forest Management*. Springer-Verlag, 2006.
- K. Siddiqi, B.B. Kimia, and C.W. Shu. Geometric shock-capturing ENO schemes for sub-pixel interpolation, computation and curve evolution. *Graphical Models and Image Processing*, 59:278–301, 1997.
- M. Sonka, V. Hlavac, and R. Boyle. *Image Processing: Analysis and Machine Vision*. O’Reilly, 1999.
- S. H. Spurr. *Aerial Photographs in Forestry*. The Ronald Press Co., 1948.
- Y. Sun and C. Beckermann. Sharp interface tracking using the phase-field equation. *J. Comput. Phys.*, 220(2):626–653, 2007.
- M. Sussman and E. Fatemi. An efficient, interface-preserving level set redistancing algorithm and its application to interfacial incompressible fluid flow. *SIAM Journal on Scientific Computing*, 20(4):1165–1191, 1997.
- M. Sussman, P. Smereka, and S. Osher. A level set approach for computing solutions to incompressible two-phase flow. *Journal of Computational Physics*, 114:146–159, 1994.
- T. Szirányi and J. Zerubia. Markov random field image segmentation using cellular neural network. *IEEE Trans. on Circuits And Systems II*, 44:86–89, 1997.
- A. Tsai, A. Yezzi, W. Wells, C. Tempny, D. Tucker, A. Fan, W. E. Grimson, and A. Willsky. Model-based curve evolution techniques for image segmentation. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 463 – 468, 2001.
- H. K. Zhao, T.F. Chan, B. Merriman, and S. Osher. A variational level set approach to multiphase motion. *Journal of Computational Physics*, 127:179–195, 1996.
- S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multi-band image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18(9):884–900, 1996.

Abstract

In this thesis, we present the ‘gas of circles’ (GOC) model, which is a tool to describe a set of circles with an approximately fixed radius. The model, when combined with image features, is able to capture circular objects in the image plane without using a template shape. We use the model, showing its empirical success, for a forestry application: the detection of the exact crown shape of individual trees in remotely sensed images. The algorithm is automatic, providing cheaper and faster statistics for forestry management (*e.g.* the density of trees, the mean crown area and diameter, *etc.*).

The ‘gas of circles’ model is based on the recently introduced ‘higher-order active contour’ (HOAC) framework, which incorporates long-range interactions between contour points, and thereby includes prior geometric information without using a template shape. The geometric properties of the HOAC model are controlled by an interaction function and the model parameters. In general the model creates long arms; this property makes it ideal for detecting road networks in satellite images, as was presented in previous work. For certain ranges of the parameters, the model creates stable circles with an approximately fixed radius instead of networks. We show now to determine this set of parameters, thereby creating the ‘gas of circles’ model. 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. This is a nontrivial requirement. We impose it by performing a functional Taylor expansion of the HOAC energy around a circle, and then demanding that the first order term be zero for all perturbations, and that the second order term be positive semi-definite. These conditions allow us to fix one of the model parameters in terms of the others, and constrain the rest. The energy is minimized using a gradient descent algorithm, and implemented using the level-set method. We illustrate with experiments that the geometric model itself creates stable circles with the desired radius, and is able to distinguish between circles with other radii even in very noisy images. We also demonstrate how the model behaves in a difficult synthetic case, where two circles are deformed to a dumbbell shape, and classical models are not able to separate them.

The general ‘gas of circles’ model has many potential applications in varied domains, but it suffers from a drawback: the local minima corresponding to circles can trap the gradient descent algorithm, thus producing phantom circles even with no supporting data. The model as such is not at fault: an algorithm capable of finding the global minimum would not

produce phantom circles. This suggests two approaches to tackling the difficulty. One is to find a better algorithm. The other is to compromise with the existing algorithm by changing the model to avoid the creation of local minima, while keeping intact the prior knowledge contained in the model. We take this second approach. We solve the problem of phantom circles by calculating, via a Taylor expansion of the energy, parameter values that make the circles into inflection points rather than minima. The result is that without supporting data the circles vanish, but even the presence of a small amount of image information can create an energy minimum at the desired radius. In addition, we find that this constraint reduces the parameters to one overall weight, and severely constrains one of the two that remain, while improving the empirical success of the model. We describe the details of how the parameters are determined.

Although, the HOAC ‘gas of circles’ model is an effective tool for modeling circular shapes, there are some difficulties. It is complicated to express the space of regions in the contour representation, and consequently difficult to work with a probabilistic formulation. In addition, from the algorithmic point of view, the current model does not allow enough topological freedom, and the implementation of the HOAC model is difficult and computationally expensive. It is possible to create an alternative formulation of HOAC models, based on the ‘phase field’ framework much used in physics to model regions and interfaces. The standard phase field model is, to a good approximation, equivalent to a classical active contour (CAC) model. Rochery *et al.* showed how to extend the basic phase field energy with extra, nonlocal terms that produce phase field models equivalent to higher-order active contours. The new formulation has several important advantages:

- There is more topological freedom during the gradient descent evolution than with other methods, which is important when the topology is unknown *a priori* as is the case in the tree crown extraction problem. In addition, more topological freedom means less chance of becoming stuck in local minima.
- The implementation of the phase field version of CACs, and in particular of HOACs, is much simpler than the equivalent contour or level set implementation. Gradient descent consists of a single partial differential equation derived directly from the model energy with no need for reinitialization or regularization. HOAC terms consist of convolutions and can be evaluated in Fourier space, causing two orders of magnitude increase in computation time.
- Phase field models provide a neutral initialization for gradient descent—no initial region is needed—meaning that bias caused by the initialization is reduced.

We address the tree crown extraction problem by constructing a phase field model of a ‘gas of circles’. We compute, as a function of the HOAC energy parameters, the phase field energy parameters that produce an equivalent model. This means that we can adjust the phase

field parameters to ensure stable circles of a given radius also. Synthetic experimental results show the consistency between the calculations and practice, *i.e.* starting from different initial conditions the gradient descent stops the regions at circles with the desired radius. We make a comparison between the HOAC and phase field evolution showing that the models are approximately equivalent.

We also extend the phase field ‘gas of circles’ model to the case of an inflection point rather than a minimum in the circle energy at the desired radius. With this model we benefit from all the advantages both of the phase field model and the inflection point GOC model. The correct parameter transformation is empirically demonstrated by using parameter values slightly different from the ones resulting in an inflection point, and showing that the desired behaviour is produced. We present the minimization of the standard phase field energy, and its HOAC extension. We describe the topological freedom offered by the model, and finally the implementation details of the framework.

To apply the model to image data, we need to develop an energy term that has a minimum value in case the segmented regions are on tree crowns, and that can be combined with the prior models presented above. The available images are color-infrared (CIR) images from the French Forest Inventory (IFN) , containing mainly regularly planted poplar stands and some separated trees, in the Saône et Loire region in France. We also tested our models on panchromatic images provided by the Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) , which contain mostly agricultural areas with individual trees. We introduce two data models. The first describes the use of only one, in the CIR case the most significant infrared, band of the three available bands. The model is based on the image gradient and on Gaussian distributions, with the values at different pixels independent, and with different means and variances for tree crowns and the background. While successful, this model, even with the strong region prior, is not capable of extracting accurately the borders of all the trees. Some trees are simply too similar to the background. To solve this problem our second data model makes use of all three bands in the CIR images. We study the improvement or otherwise of the extraction results produced by modelling the three bands as independent or as correlated. As we will see, even at the level of maximum likelihood, the inclusion of ‘colour’ information, and in particular, inter-band correlations, can improve the results, and in conjunction with the region prior, the full model is considerably better than that based on one band alone.

Combining the data terms with the prior models, we analyze the empirical success of five possible combinations. We make comparisons among them and describe the execution time and the parameter settings. In the panchromatic case, the phase field model proves to be the best, and if there is more than one spectral channel available, then the color model shows great improvement over the others.

We use the above models to extract tree crowns on aerial images, but the models are not restricted to forest management. They can also be applied to the detection of other circular objects. A few examples: in nanotechnology to detect various particles and microarrays

| | 1 | 2 | 3 | 4 |
|----------------------------|---|---|---|---|
| Horváth et al. (2006b) | * | | | * |
| Horváth et al. (2006c) | | * | | |
| Horváth and Jermyn (2007b) | | | * | * |
| Horváth and Jermyn (2007a) | | | * | |
| Horváth (2007) | | | | * |

Table B.1: Connection between the theses and the Author’s publications

in electron-microscopic images; in biology to segment circular cells and molecules, or to detect pollen grains; in medical image processing, to detect circular objects; in remote sensing, in the processing of aerial and satellite images, for meteorological, military, and agricultural management.

In conclusion, we present the evolution of our models starting from the HOAC and phase field frameworks. We describe one weakness of the method, which is that circles that are too close to each other tend to merge, and propose two ways to overcome this problem. Finally, we describe some possible ways to improve the ‘gas of circles’ model.

Summary of the Author’s contributions

We summarize the results of the Author arranging them into four theses. Table B.1 illustrates how the thesis points are covered by the publications of the Author.

1. The Author presented the stability analysis of the higher-order active contour model for a ‘gas of circles’, defined the required conditions and consequent parameter constraints.
2. The Author presented new parameter settings of the HOAC energy that creates an inflection point in the energy function rather than a minimum, so that circles vanish without supporting image data, but a small amount of image information can create an energy minimum at the given radius. This criteria fixes all parameters of the prior except one.
3. The Author showed how to transform the previously determined ‘gas of circles’ parameters into the phase field model so that it creates stable circles with the desired radius and determined the constraints on the parameter conversion. The Author presented the steps of the transformation and the constraints converting the inflection point ‘gas of circles’ parameters to the phase field framework.
4. The Author introduced two different data models. First, the HOAC model was completed with a term based on the image gradient and Gaussian-based intensity terms

coming from a single-spectral image. Second, the previous intensity term was extended to multispectral images using a multi-dimensional Gaussian model.

Kivonat

Ebben a tézisben bemutatjuk a ‘gas of circles’ (GOC) modellt, amely megközelítőleg azonos sugarú körök leírására használható. Ha a modellt képi információkkal kombináljuk, akkor template alakzat felhasználása nélkül is képes kör alakú objektumok detektálására. A modell gyakorlati használhatóságát egy erdészeti alkalmazáson keresztül mutatjuk be: különálló fák koronájának pontos meghatározása légi felvételeken. Mivel az algoritmus automatikus, ezért olcsóbb és gyorsabb statisztikák készíthetők az erdészeti hivatalok részére (*pl.*: fák sűrűsége; átlagos korona átmérő, terület, *stb.*).

A ‘gas of circles’ modell a korábban bemutatott ‘higher-order active contour’ (HOAC) keretrendszeren alapul, amely a kontúrponatok közötti hosszabb távú interakciókkal előzetes geometriai információt ír le, template alakzat felhasználása nélkül. A modell geometriai tulajdonságait egy interakciós függvény és a modell paraméterei határozzák meg. Általános esetekben a modell hosszú karokat képez, mely tulajdonsága alkalmassá teszi úthálózatok detektálására műholdképeken, ahogy azt korábbi munkákban már bemutatták. Bizonyos paraméterek esetén azonban, a modell stabil, megközelítőleg azonos sugarú körök halmazát hozza létre hálózat helyett. Megmutatjuk, hogyan határozható meg ezen paraméterek egy halmaza, így létrehozva a ‘gas of circles’ modellt. Egy ilyen modell előállításához a köröknek stabilan kell viselkedniük a határaikon fellépő perturbációkkal szemben, azaz a körök HOAC energia lokális minimumai kell hogy legyenek, különben átalakulnának más alakzatokká. Ez egy nem triviális feltétel, a problémát HOAC energia kör körüli Taylor-sorfejtésével oldjuk meg, majd az első rendű tagtól megköveteljük, hogy nulla legyen minden perturbáció esetére, a másodrendű tagtól pedig, hogy pozitív szemidefinit legyen. Ezen feltételek lehetőséget adnak a modell egyik paraméterének rögzítésére a többi ismeretében. Az energiát gradiens módszerrel minimalizáljuk, ennek megvalósításához a level-set keretrendszert használjuk. Kísérletekkel igazoljuk, hogy a geometriai modell önmagában stabil, meghatározott sugarú köröket képez, és képes megkülönböztetni az ettől eltérő sugarú köröket, még nagyon zajos képeken is. Bemutatjuk a modell viselkedését egy másik bonyolult esetben is, amikor két egymáshoz közeli kör súlyzó alakká deformálódik, és a klasszikus módszerek nem képesek szétválasztani azokat.

Az általános ‘gas of circles’ modell több területen is alkalmazható, de van egy hátránya: a körnek megfelelő lokális minimum képes megállítani a gradiens módszert, így fantom körök képződnek még akkor is, ha ezt az adatkifejezés nem támogatja. Ez nem jelenti azt,

hogya a modell rossz. Ha rendelkezésünkre állna egy olyan algoritmus, mely megtalálja a globális minimumot, akkor nem keletkeznének fantom alakzatok. Ez a probléma két módon oldható meg. Az egyik, találni egy jobb algoritmust. A másik, beérni a jelenlegi algoritmus-sal, és úgy megváltoztatni a modellt, hogy az ne tartalmazzon lokális minimumot, de megmaradjon az az előzetes tudás, melyet eddig beleintegráltunk. Mi a második megközelítést választjuk. A problémát a paraméterek olyan választásával oldjuk meg, melyek használatával a körök az energiának nem lokális minimumai, hanem inflexiós pontjai lesznek. Ehhez az energiát Taylor-sorba fejtjük. Mindezek eredményeképpen azt tapasztaljuk, hogy képi információ nélkül a körök eltűnnek, de már kicsiny információ is képes energiaminimumot képezni a kívánt sugárban. Megmutatjuk még, hogy így a modell paramétereinek száma felére csökken, és a megmaradó paraméterek egyikére korlátozó feltételeket kapunk; és ezzel együtt a modell gyakorlati alkalmazhatósága is javul. Részletesen ismertetjük a paraméterek meghatározását.

Habár a HOAC ‘gas of circles’ modell hatékonyan használható kör alakú alakzatok modellezésére, használata során mégis felmerül néhány nehézség. Kontúr reprezentáció felhasználásával bonyolult kifejezni a régiók terét és ebből következően nehéz valószínűségi formulákkal dolgozni. Algoritmikai szempontból nézve a jelenlegi modell nem kínál topológiai szabadságot, és a HOAC modell implementációja bonyolult és számításigényes. Előállítható a HOAC modell egy alternatív formulája, mely a fizikában a régiók és interfészek modellezésére használt phase field keretrendszeren alapul. Az alap phase field modell ekvivalens a klasszikus aktív kontúr modellel. Rochery és tsai. bemutatták a phase field modell egy nem-lokális taggal való kiterjesztését, mely így ekvivalens a HOAC keretrendszerrel. Az új modellnek van néhány igen kedvező tulajdonsága, ezek:

- Nagyobb topológiai szabadságot kínál a gradiens módszer során mint más modellek, amely igen fontos, ha a topológia nem ismert előre, mint például a fák koronájának szegmentálása esetén. Ráadásul, a nagyobb topológiai szabadság azt is jelenti, hogy kisebb eséllyel kerülünk lokális minimumba.
- A klasszikus aktív kontúr (CAC) és különösképpen a HOAC modellek phase field implementációja sokkal egyszerűbb, mint a velük ekvivalens kontúr vagy level-set implementáció. A gradiens módszer egyenlete csupán egy egyszerű parciális differenciál egyenlet, mely közvetlenül számolható a modell energiájából, nincs szükség újra inicializálásra és regularizálásra sem. A HOAC tagok egy konvolúciót tartalmaznak, mely a Fourier-térben számolva két nagyságrenddel gyorsabb futásidőt eredményez.
- A phase field modell esetén lehetséges a gradiens módszer teljesen neutrális inicializációja — nincs szükség kiindulási régióra — amely azt jelenti, hogy megszüntetni az inicializálásból származó zavarokat.

A fakoronák szegmentálására bemutatjuk a phase field ‘gas of circles’ modellt. A HOAC energia paramétereinek függvényében meghatározzuk azon phase field paramétereket, mely-

ek ekvivalens modellt eredményeznek. Ez azt jelenti, hogy meg tudunk határozni olyan phase field paramétereket melyek stabil, adott sugarú köröket eredményeznek. Megmutatjuk a számításaink és a gyakorlati alkalmazás közötti konzisztenciát azaz, hogy különböző inicializációk esetén a gradiens módszer kívánt sugarú köröket képez. Összehasonlítjuk a HOAC és phase field modellek evolúcióját és megmutatjuk, hogy azok megközelítőleg ekvivalensek.

Kiterjesztjük a phase field ‘gas of circles’ modellünket arra az esetre is, amikor az energiának inflexiós pontja van a kívánt sugarú kör esetén, nem pedig minimuma. Ez a modell egyesíti mind a phase field, mind az inflexiós pontot tartalmazó GOC modell előnyeit. A paraméter transzformáció helyességét kísérlettel is igazoljuk, az inflexiós pontot eredményező beállításoktól kissé eltérő értékek választásával, melyek esetén a modell elvárásainknak megfelelően viselkedik. Bemutatjuk, hogyan minimalizálható a phase field energia és annak HOAC kiterjesztése. Áttekintjük a modell kínálta topológiai szabadságot és a keretrendszer implementációjának részleteit.

Ha a modellt képi információval együtt szeretnénk használni, ki kell fejlesztenünk olyan energiagombokat, melyek akkor veszik fel minimumukat, ha a szegmentált régiók a fák koronáin találhatóak, és kombinálni tudjuk a fent bemutatott előzetes alakzatinformációt leíró modellekkel. A rendelkezésünkre álló képek színes infravörös (CIR) légifelvételek, melyeket a French Forest Inventory (IFN) biztosított számunkra, többségében szabályosan ültetett nyárfákkal, és néhányon különálló egyéb fákkal. A felvételek a franciaországi Saône et Loire térségben készültek. Modelljeinket kipróbáltuk pankromatikus képeken is, melyeket a Hungarian Central Agricultural Office, Forestry Administration (CAO, FA) bocsájtott rendelkezésünkre és főként mezőgazdasági területeken található különálló fákat tartalmaznak.

Szintén bevezetünk két adatmodellt. Az első azt írja le, hogyan használható a rendelkezésünkre álló színes infravörös képeknek egyetlen, a legjellemzőbb csatornája. Ez az infravörös csatorna. Ez a modell a képgradiensén és Gauss-eloszláson alapul, kihasználva a fakoronák és a háttér pixel intenzitásainak különböző szórását és középértékét. Habár, a modell sikeresnek bizonyult, még jelentős előzetes régió információ esetén sem képes pontosan meghatározni az összes fakorona körvonalát. Ennek megoldására bemutatjuk a CIR képek mindhárom csatornáját felhasználó második adatmodellünket. Megvizsgáljuk az eredményeket abban az esetben amikor feltételezzük, hogy a három csatorna független és amikor azok nem függetlenek. Mint azt látni fogjuk, már a maximum likelihood esetében is sokat javul az eredmény a ‘szín’ információ felhasználásával, és ez különösképp igaz, ha a csatornák között korreláció van. Következésképpen, az előzetes alakzat-információval kombinálva a modellünk jelentős javulást eredményez az egycsatornás-modellhez képest.

Öt különböző, a geometriai és az adatgombok kombinálásával kapott modell gyakorlati működését teszteljük. Összehasonlítjuk őket, ismertetjük futásidejüket és az alkalmazott paraméterbeállításokat. A pankromatikus képek esetében a phase field modell mutatkozik legjobbnak, míg ha több spektrális csatorna is rendelkezésünkre áll, akkor a színes modell

| | 1 | 2 | 3 | 4 |
|----------------------------|---|---|---|---|
| Horváth et al. (2006b) | * | | | * |
| Horváth et al. (2006c) | | * | | |
| Horváth and Jermyn (2007b) | | | * | * |
| Horváth and Jermyn (2007a) | | | * | |
| Horváth (2007) | | | | * |

Table B.2: Kapcsolat a tézispontok és a szerző publikációi között.

eredményez jelentős javulást.

Modellünket fakoronák körvonalának detektálására használjuk légi felvételeken, habár más, nem erdészeti jellegű alkalmazások is lehetségesek. Néhány más példa: nanotechnológiában különböző részecskék és microarray-ek detektálására elektronmikroszkóppal készült képeken; biológiában kör alakú sejtek és molekulák szegmentálására, vagy pollen részecskék detektálására; orvosi képfeldolgozásban, kör alakú objektumok szegmentálására; távérzékelésben, légi és műholdképek feldolgozása során, meteorológiai, harcászati vagy mezőgazdasági alkalmazásokban.

Bemutatjuk modelljeink fejlődését, a HOAC és a phase field keretrendszerektől kezdődően. Ismertetjük modellünk egy gyenge pontját, nevezetesen azt az eshetőséget, amikor az egymáshoz túl közel kerülő körök összeolvadnak. Majd, két lehetséges megoldást javasolunk a problémára. Végül a ‘gas of circles’ lehetséges fejlesztésének új irányzatait mutatjuk be.

Az eredmények tézisszerű összefoglalása

A szerző eredményeit négy tézispontba rendezve foglaljuk össze. A B.2. táblázat bemutatja, hogy az egyes tézispontokat a szerző mely publikációiban találhatjuk.

1. A szerző ismertette a higher-order active contour modell stabilitásának analízisét, elkészítve a ‘gas of circles’ modellt, és definiálta a szükséges feltételeket és az ezekből következő paraméter korlátozásokat.
2. A szerző bemutatta a HOAC paraméterek egy új beállítását, mely az energiafüggvényen lokális minimum helyett inflexiós pontot eredményez, így azon körök, melyeket az adatkifejezés nem támogat eltűnnek, de már kicsiny, a kör létezését alátámasztó képi információ képes energia-minimumot képezni az adott sugárban. Ezenfelül, az új kritérium felére csökkentette a paraméterek számát.
3. A szerző megmutatta, hogyan lehetséges a már korábban meghatározott ‘gas of circles’ paramétereket áttranszformálni phase field paraméterekké úgy, hogy kívánt sugarú, stabil köröket képezzen, valamint meghatározta a paraméter átalakítás feltételeit.

Emellett bemutatta még az inflexiós pontot tartalmazó ‘gas of circles’ model paraméter transzformációjának lépéseit és megszorításait.

4. A szerző bevezetett két különböző adatkifejezést. Először a HOAC modellt egészítette ki egy a gradiensen és Gauss-eloszláson alapuló taggal, mely egycsatornás képekre használható, majd demonstrálta az intenzitáson alapuló tag kiterjesztését többcsatornás képekre többdimeziós Gauss-eloszlást használva.