A MARKOV RANDOM FIELD MODEL FOR EXTRACTING NEAR-CIRCULAR SHAPES

Tamas Blaskovics, Zoltan Kato

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

Ian Jermyn

ARIANA (joint INRIA/I3S research group), INRIA
B.P. 93, 06902 Sophia Antipolis, France

ABSTRACT

We propose a binary Markov Random Field (MRF) model that assigns high probability to regions in the image domain consisting of an unknown number of circles of a given radius. We construct the model by discretizing the 'gas of circles' phase field model in a principled way, thereby creating an 'equivalent' MRF. The behaviour of the resulting MRF model is analyzed, and the performance of the new model is demonstrated on various synthetic images as well as on the problem of tree crown detection in aerial images.

Index Terms— segmentation, Markov random field, shape prior

1. INTRODUCTION

When people interpret images, they take into account not only the image data, but also their prior knowledge about the phenomena that generated the image. If machine methods are ever to duplicate human performance, they must be able to include this knowledge in mathematical models. In particular, for the segmentation of entities in an image, the shape of the entity involved plays an essential role, and in consequence various approaches to the modelling of prior shape knowledge have been studied. Many incorporate shape knowledge by weighting different regions according to their 'proximity' to one or more template shapes [1, 2]. However, in many applications, there is an *a priori* unknown number of instances of an entity in the image, all of which must be segmented, a problem that is difficult to address using templates. On the other hand, Markov random fields are well suited to segmenting multiple instances, but little work has been done on including prior shape knowledge in such models [3].

Higher-order active contours (HOACs) [4] can incorporate sophisticated prior shape knowledge without using templates. They can therefore be used to segment multiple instances of an entity. In particular, in [5], a stability analysis was used to tune the parameters of the model in [4] so as to make regions consisting of an arbitrary number of circles of a given radius into local energy minima. This 'gas of circles' model is useful for several applications, notably to the segmentation of tree crowns from aerial images [5].

HOACs are described by energy functionals, and energy minimization is performed via gradient descent. A probabilistic version of these models, although implicit in the use of an energy, has not yet been formulated explicitly, but would open the door to learning shape model parameters and to using stochastic descent algorithms to make estimates. This means discretizing the space of regions so that everything is well-defined, and in a way that preserves the information content of the model: the discretized model should be 'equivalent' to the continuum model. This is complicated in terms of contours, but fortunately, HOACs can be reformulated as 'phase field' models [6]. Phase field models can naturally be viewed as the Gibbs energies of continuum Markov random fields, and can be discretized on a spatial grid, thereby greatly simplifying both the discretization and its subsequent implementation.

In this paper, we describe a Markov random field model that is approximately equivalent to the phase field HOAC model described in [7], including the 'gas of circles' parameter constraints. We test the performance of the model, minimized using simulated annealing with a standard Gibbs sampler, on synthetic images, and apply it to tree crown extraction from aerial images. The main contributions of the work are the construction of a probabilistic model of HOACs, the incorporation of prior shape knowledge into an MRF model, and the subsequent improvement in applications when compared to classical MRFs, including increased robustness to clutter and noise.

1.1. The 'gas of circles' phase field model

The phase field framework represents a region R by a function $\Phi \ni \phi : \mathcal{D} \to \mathbb{R}$ defined on the image domain $\mathcal{D} \subset \mathbb{R}^2$, and a threshold z: $R = \zeta_z(\phi) = \{x \in \mathcal{D} : \phi(x) \geq z\}$. The function is controlled by an energy, which imposes both a form on ϕ and energetic constraints on the corresponding region. We start from the phase field energy $E(\phi)$ of [4]:

$$E(\phi) = \int_{\mathcal{D}} \frac{D_f}{2}|\nabla\phi|^2 + \lambda_f\Big(\frac{\phi^4}{4} - \frac{\phi^2}{2}\Big) + \alpha_f\Big(\phi - \frac{\phi^3}{3}\Big) - \frac{\beta_f}{2}\int_{\mathcal{D}\times\mathcal{D}'} \nabla\phi\cdot\nabla'\phi'\,\Psi((x-x')/d)\;, \quad (1)$$

This research was partially supported by the Hungarian Scientific Research Fund (OTKA) – K75637. The authors would like to thank the Central Agricultural Office of Hungary for the aerial images.

where d controls the range of interaction, and

$$\Psi(z) = \begin{cases} \frac{1}{2}\big(2 - |z| + \frac{1}{\pi}\sin(\pi|z|)\big) & \text{if } ||z|-1| < 1, \\ 1 - H(|z|-1) & \text{otherwise.} \end{cases} \quad (2)$$

where H is the Heaviside step function. It is convenient to integrate the non-local term in E by integration by parts:

$$-\frac{\beta_f}{2}\int_{\mathcal{D}\times\mathcal{D}'} \nabla\phi\cdot\nabla'\phi'\,\Psi((x-x')/d) = \frac{\beta_f}{2}\int_{\mathcal{D}\times\mathcal{D}'} \phi\,\phi'\,\underbrace{\nabla^2\Psi((x-x')/d)}_{\mathcal{G}((x-x')/d)}\,.$$

The value ϕ_R that minimizes E for a fixed region R takes the values $+1$ inside R and -1 outside, away from the boundary ∂R, while changing smoothly from -1 to $+1$ in a narrow interface region around ∂R. In the 'gas of circles' model, the parameters of E are adjusted so that a circle of the desired radius is a local minimum and therefore stable [5, 7].

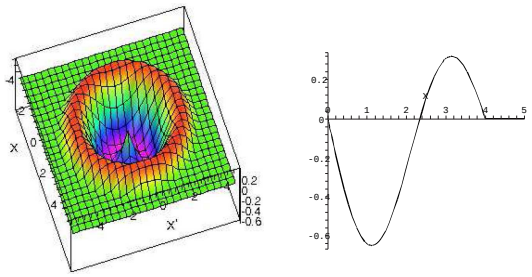

Fig. 1. The higher order interaction function $\mathcal{G}$ for $d = 2$.

2. DISCRETIZATION

We discretize $\mathcal{D}$ as a finite rectangular lattice $\mathcal{S} \subset \mathbb{Z}^2 \subset \mathbb{R}^2$. Each lattice site $s \in \mathcal{S}$ corresponds in the standard fashion to a rectangular cell $c_s \subset \mathcal{D}$. A region will be represented by a function $\Omega \ni \omega : \mathcal{S} \to \{\pm 1\}$. There is a map $W : \Phi \to \Omega$:

$$\omega_s = W(\phi)_s = 2H\Big(\int_{c_s}\phi\Big) - 1\,. \quad (3)$$

Using W, we want to construct an MRF that is 'equivalent' to the phase field model. Since we will compute MAP estimates, and since we wish to preserve the property that circles of a given radius have higher probability than neighbouring configurations, we define the MRF energy to be

$$U(\omega) = \min_{\phi:W(\phi)=\omega} E(\phi) = E(\phi_\omega)\,.$$

Note that ϕ_ω may not be unique due to the translation invariance of E. The nature of ϕ_ω can be deduced from the description of ϕ_R in the previous section. Were $D_f = \beta_f = 0$, we would have $\phi_\omega(x) = \omega_{\varsigma(x)}$, where ς is define by $x \in c_{\varsigma(x)}$. This value can be substituted into E. Writing $\int_{\mathcal{D}} = \sum_{s\in\mathcal{S}}\int_{c_s}$, and using the fact that $\varsigma(x)$ is constant on each cell, we then find that, up to an additive constant,

$$E(\phi_\omega) = \sum_s\Big\{\frac{2\alpha_f}{3}\omega_s + \frac{D_f}{2}\int_{c_s}|\nabla\omega_{\varsigma(x)}|^2\Big\} + \frac{\beta_f}{2}\sum_{s,s'}F_{ss'}\,\omega_s\,\omega_{s'}\,,$$

where $F_{ss'} = \int_{c_s\times c_{s'}}\mathcal{G}((x-x')/d)$. Note that $F_{ss'} = 0$ if $|s-s'| > 2d$, and that the derivative term is not well-defined.

For non-zero D_f and β_f, but far from ∂R, we will still have $\phi_\omega = \omega_{\varsigma(x)}$; indeed $D_f \neq 0$ reinforces this. Near ∂R, there will be a smooth transition between ± 1, controlled by the constraint (3). Due to the parameter settings used for the phase field model, and the resulting width of the boundary layer, the most significant deviations from $\omega_{\varsigma(x)}$, and hence the largest gradients, will be at the interfaces between pairs of nearest neighbour cells of opposite sign. These deviations will have little effect on the nonlocal term, since it is a sum over a large region, and so we can leave $\phi_\omega = \omega_{\varsigma(x)}$ in this term. Consider now the derivative term. It can be rewritten

$$\sum_s\int_{c_s}|\nabla\phi_\omega|^2 = \frac{1}{8}\sum_s\sum_{s'\sim s}\int_{c_s\cup c_{s'}}|\nabla\phi_\omega|^2\,.$$

where $\sim$ is the nearest neighbour relation. As just stated, the integrals will be small unless $(\omega_s - \omega_{s'})^2 \neq 0$. We assume that the integrals for each pair of such neighbours will be the same. The final MRF energy can therefore be written

$$U(\omega) = \alpha\sum_s\omega_s + \frac{D}{2}\sum_s\sum_{s'\sim s}(\omega_s-\omega_{s'})^2 + \frac{\beta}{2}\sum_{s,s'}F_{ss'}\,\omega_s\,\omega_{s'}\,. \quad (4)$$

where $\alpha = \frac{2\alpha_f}{3}$, $\beta = \beta_f$, and $D \propto D_f$ incorporates the integral over pairs of boundary cells.

3. MARKOV MODEL

The discrete energy functional in Eq. (4) defines an MRF with respect to an appropriate neighbourhood system. In Eq. (4), there are two type of interactions: the finite difference approximation of the gradient term, corresponding to classical nearest neighbours (*doubleton* interactions); and the interactions governed by kernel F_{sr}, corresponding to a neighbourhood of diameter $2d$ (*long-range* interactions). Since the size of the neighbourhood is dominated by the latter ($d \geq 2$ in practice), the neighbourhood of a site $s \in \mathcal{S}$ is the set $\{s' \in \mathcal{S} : |s-s'| < 2d\}$. In addition, we have the *singleton*

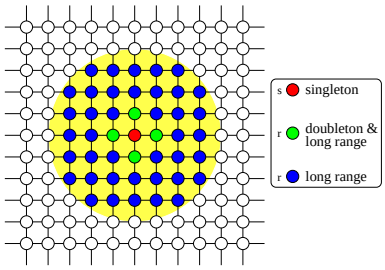

Fig. 2. MRF neighbourhood ($d = 2$, *i.e.* $|s - s'| < 4$).

terms. These structures are illustrated in Fig. 2. We now discuss the singleton, doubleton, and long-range clique potentials and their effects in more detail.

The first term in equation (4) has the form $\sum_s V_s$, where $V_s = \alpha\omega_s$ is thus the *singleton* potential. Its effect depends on the sign of α. Setting $\alpha > 0$ favours $\omega_s = -1$ everywhere. Thus as α is increased, typical configurations have less foreground pixels, yielding less circles. This is illustrated in Fig. 3), where samples from the MRF are shown for different α.

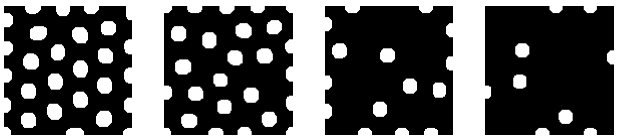

$\alpha = 0.1863$ $\alpha = 0.21$ $\alpha = 0.24$ $\alpha = 0.27$

Fig. 3. Typical samples from the MRF defined by U: the effect of altering α ($d = 8$, $\beta = 0.096$, $D = 0.1545$).

In addition, we will use a data likelihood that represents the background and foreground pixel classes by Gaussian distributions. This adds inhomogeneous terms to V_s. The result is that in the posterior probability for ω, V_s is given by

$$V_s = \alpha\omega_s + \gamma\left(\ln(\sqrt{(2\pi)}\sigma_{\omega_s}) + \frac{(I_s - \mu_{\omega_s})^2}{2\sigma^2_{\omega_s}}\right),$$

where $I : \mathcal{S} \to \mathbb{R}$ is the image data. The parameters of the Gaussian distributions $\mu_{\pm 1}$ and $\sigma_{\pm 1}$ are learned from representative samples provided by the user.

The *doubleton* potential $V_{\{s,s'\}}$ is defined on sets $\{s, s'\}$. Since the sum in U is over the nearest neighbours of each s, each doubleton is counted twice: this eliminates the factor $\frac{1}{2}$. To fix D, we use the fact that D_f is essentially equal to the contour length parameter in the original higher-order active contour model [6]. Since $\sum_{s\sim s'}(\omega_s - \omega_{s'})^2$ measures boundary length in the MRF, we should have $D = D_f$. However, the MRF length term is bigger than the boundary length of $\zeta(\phi_\omega)$ due to discretization effects. The ratio between these contour lengths is ≈ 0.82. Taking into account these corrections, the doubleton potential for $s \sim s'$ is

$$V_{\{s,s'\}} = D(\omega_s - \omega_{s'})^2 = \begin{cases} 0.82D_f & \text{if } \omega_s \neq \omega_{s'}, \\ 0 & \text{otherwise.} \end{cases}$$

The *long-range* potential $V'_{\{s,s'\}}$ introduces the prior shape knowledge. It is defined for $\{s, s'\}$ with $|s - s'| < 2d$:

$$V'_{\{s,s'\}} = \beta F_{ss'}\,\omega_s\,\omega_{s'} = \begin{cases} -\beta F_{ss'} & \text{if } \omega_s \neq \omega_{s'}, \\ +\beta F_{ss'} & \text{otherwise.} \end{cases}$$

From Fig. 1, it is clear that V' favours the same label when $|s - s'| < d'$ (*attractive* case) and different labels when $d' < |s - s'| < 2d$ (*repulsive* case), where $d' \simeq d$ is the zero of $\mathcal{G}$.

classical MRF		
Noise	FP (%)	FN (%)
0	0.3	0.3
-5	1.8	2.0
-10	2.5	7.2
-16	4.2	24.5
-20	9.8	39.6

gocMRF		
Noise	FP (%)	FN (%)
0	0.2	1.0
-5	1.2	1.7
-10	2.1	4.3
-16	3.9	8.5
-20	6.4	16.5

Table 1. Results on synthetic noisy images.

4. EXPERIMENTAL RESULTS

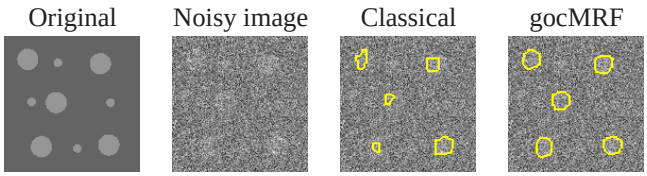

Fig. 4. Result on a synthetic image corrupted by additive noise of -16dB.

We use simulated annealing [8] with a standard Gibbs sampler [9] to minimize U. The initial temperature was set to 3 and we used an exponential annealing schedule $T_{n+1} = 0.97T_n$. The iterations were stopped when the temperature decreased below 0.01.

Table 1 shows the quantitative results obtained on a set of 160 synthetic noisy images. We compare the segmentation results to a classical MRF model [10], which doesn't include a shape prior. For a fair comparison, the false-positive (FP) and false-negative (FN) rates were computed while excluding the small circular regions. This is to avoid biasing the measure: the classical MRF should detect all regions having a particular intensity while our model will only detect the desired circles. Based on these numbers, it is clear that the proposed model is less sensitive to noise. Fig. 4 shows a sample synthetic image and demonstrates the results for a -16dB noise level.

Tree crown detection in aerial images can be used to compute a number of the tree plantation statistics used for efficient forestry management. The main challenge to successful detection is the cluttered background, which causes traditional segmentation methods to fail. Fig. 5 and Fig. 6 show some results. In Fig. 5, the trees are difficult to separate due to shadows, blur, and vegetation between neighbouring crowns. In Fig. 6, the classical MRF model fails to separate trees from background vegetation because they have similar intensity distributions. Obviously, the d parameter of our model, controlling the approximate radius of the detected trees, must be set correctly in order to achieve the best performance.

5. CONCLUSION

We have constructed a new binary MRF model that includes prior information about the shape of the region to be segmented, in this case, that it be composed of an arbitrary number of circles of a certain radius. In addition to classical homogeneity terms, the model includes terms that encourage inhomogeneity at long ranges, in order to capture nonlocal geometric properties. The model was derived in a principled fashion from a continuous 'gas of circles' phase field model. It is 'equivalent' to the continuous model in a well-defined way, thereby guaranteeing that stability conditions in the phase field model remain valid in the new model. Experimental tests on various synthetic and real images confirm the performance of the 'gas of circles' MRF model.

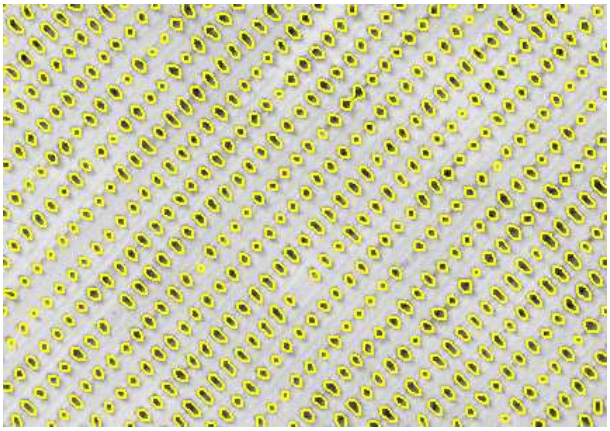

Fig. 5. Tree detection result on an aerial image of a regularly planted pine forest.

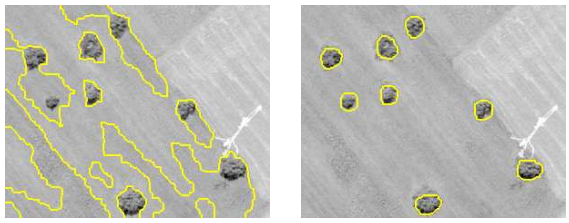

Classical MRF gocMRF

Fig. 6. Aerial image of sparsely planted trees. The classical MRF fails because the intensities of the foreground and background are similar making separation ambiguous.

6. REFERENCES

[1] A. Srivastava, S. Joshi, W. Mio, and X. Liu, "Statistical shape analysis: Clustering, learning, and testing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 4, pp. 590–602, 2003.

[2] D. Cremers, M. Rousson, and R. Deriche, "A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape," *International Journal of Computer Vision*, vol. 72, no. 2, pp. 195–215, 2007.

[3] R. Huang, V. Pavlovic, and D. N. Metaxas, "A graphical model framework for coupling MRFs and deformable models," in *Proc. of Int. Conf. on Computer Vision and Pattern Recognition*, 2004.

[4] M. Rochery, I. H. Jermyn, and J. Zerubia, "Higher order active contours," *International Journal of Computer Vision*, vol. 69, no. 1, pp. 27–42, August 2006.

[5] P. Horvath, 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," *Pattern Recognition*, vol. 42, no. 5, pp. 699–709, May 2009.

[6] M. Rochery, I. H. Jermyn, and J. Zerubia, "Phase field models and higher-order active contours," in *Proceedings of IEEE International Conference on Computer Vision*, Beijing, China, October 2005, vol. 2, pp. 970–976.

[7] 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, Sept. 2007.

[8] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, May 1983.

[9] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721–741, 1984.

[10] M. Berthod, Z. Kato, S. Yu, and J. Zerubia, "Bayesian image classification using Markov random fields," *Image and Vision Computing*, vol. 14, pp. 285–295, 1996.