

R. Kluszczyński · M. N. M. van Lieshout
T. Schreiber

Image segmentation by polygonal Markov fields

Received: 29 August 2005 / Revised: 12 December 2005 /
Published online: 19 August 2006
© The Institute of Statistical Mathematics, Tokyo 2006

Abstract This paper advocates the use of multi-coloured polygonal Markov fields for model-based image segmentation. The formal construction of consistent multi-coloured polygonal Markov fields by Arak–Clifford–Surgailis and its dynamic representation are specialised and adapted to our context. We then formulate image segmentation as a statistical estimation problem for a Gibbsian modification of an underlying polygonal Markov field, and discuss the choice of Hamiltonian. Monte Carlo techniques, including novel Gibbs updates for the Arak model, to estimate the model parameters and find an optimal partition of the image are developed. The approach is applied to image data, the first published application of polygonal Markov fields to segmentation problems in the mathematical literature.

Keywords Arak process · Dynamic representation · Image segmentation · Gibbsian modification · Gibbs sampler · L_p misclassification rate · Multi-coloured polygonal Markov field

Work carried out under project PNA4.3 ‘Stochastic Geometry’. This research was supported by the EC 6th Framework Programme Priority 2 Information Society Technology Network of Excellence MUSCLE (Multimedia Understanding through Semantics, Computation and Learning; FP6-507752), and partially by the Foundation for Polish Science (FNP) and by the Polish Minister of Scientific Research and Information Technology, grant 1 P03A 018 28 (2005-2007) [the third author].

R. Kluszczyński · T. Schreiber
Faculty of Mathematics and Computer Science, Nicolaus Copernicus University, ul. Chopina
12-18, 87-100 Toruń, Poland

M. N. M. van Lieshout (✉)
CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
E-mail: Marie-Colette.van.Lieshout@cwi.nl

1 Introduction

One of the fundamental image analysis tasks is that of *segmentation*, i.e. to partition the image in relatively homogeneous regions (Rosenfeld and Kak 1982). Indeed, segmenting the data is often the first step in image interpretation problems. The partition may be achieved at several conceptual levels. At the lowest level, that of individual pixels, segmentation amounts to classification of pixel values. At the other extreme, the focus of attention are the objects that form a given image and the goal is to extract them from the image.

A myriad of segmentation methods has been proposed, from elementary thresholding through contour extraction methods to scene modelling. For a comprehensive overview, the reader is referred to Chap. 10 in Rosenfeld and Kak (1982), or to Chap. 4 in the state of the art report for single modalities published within the framework of the MUSCLE Network of Excellence. Recent material can also be found in Vincent (1999) which treats morphological methods, in Chap. 10 in Green, Richardson, and Hjort (2003) on Bayesian models, or in the volume edited by Osher and Paragios (2003) which discusses a.o. level set approaches, active contour models and variational methods. We should also mention the textbooks by Gimmel'farb (1999) and Winkler (2003), the former of which focuses on random field texture models, the latter on computational aspects of the segmentation problem.

In this paper, we propose to use polygonal field models. Thus, we place ourselves at the intermediate conceptual level that regards a segmentation as a coloured tessellation Hurn, Husby, and Rue (2003). In contrast to pixel based Markov random field (MRF) priors with small neighbourhoods, our model is able to capture global aspects of the image. To do so in an MRF, a larger neighbourhood system should be used, at large computational cost. Moreover, when there are long range interactions, boundary effects tend to become apparent. Since the polygonal fields we propose exhibit zero-range Markovianity, such long-range artefacts are prevented. Conceptually too, reality is continuous rather than discrete. At the same time, compared to a high-level scene interpretation approach, there is no need to model all objects in the image, which is feasible in restricted application domains only.

The idea to use polygonal field models for image segmentation purposes can be traced back to Clifford and Middleton (1989); from a computational point of view, a Metropolis–Hastings style sampler was developed by Clifford and Nicholls (1994) and applied to an image reconstruction problem within a Bayesian framework. To the best of our knowledge, the original implementation has not been put to the test for real image data, though. This is also noted by Paskin and Thrun (2005), who propose a modification of the sampler of Clifford and Nicholls (1994) and illustrate its usefulness in practice for robotic mapping.

Other kinds of tessellations have also been proposed. Heikkinen and Arjas (1998) combine a Voronoi tessellation with a Markov random field to build a prior distribution for smoothly varying intensity surfaces; Green (1995) and Møller and Skare (2001) propose Voronoi based models for image segmentation, Nicholls (1998) advocates triangulations, and Gimmel'farb (1999), Møller and Waagepetersen (1998), Tjelmeland and Besag (1998) suggest Markov random field models with longer range, higher order, or region based interaction structures.

It should be emphasised that a major reason for using Voronoi tessellations or triangulations rather than polygonal Markov random fields is that simulation of

the latter by a variety of Metropolis–Hastings updates (Borkar and Mitter 1993; Clifford and Nicholls 1994) is rather cumbersome. The algorithm in Schreiber (2005) by contrast is both conceptually and computationally easier. The goal of the present paper is to exploit this fact and describe a flexible, general framework for segmentation, to present the computational tools needed, and to indicate the feasibility of the method. In a follow-up paper, we will focus on benchmarking in much greater detail than space allows here.

The plan of this paper is as follows. In Sect. 2, we recall the family of consistent polygonal random field models introduced by Arak and Surgailis (1989), and adapt their dynamic representation to our purposes. Section 3 presents a coherent framework for image interpretation as a statistical inference problem for a Gibbsian modification of a suitably chosen Arak process; the Markov chain Monte Carlo machinery required to select model parameters and find an optimal segmentation is developed in Sect. 4. The next section describes the implementation, and presents results on toy examples. The paper is concluded with a critical discussion and plans for further research.

2 Multi-coloured polygonal Markov fields

In the present section we briefly recall the formal construction of consistent multi-coloured polygonal Markov fields introduced by Arak and Surgailis (1989). We do not use the full generality of this construction; rather, we specialise and adapt to suit our particular purposes. The choices we make are aimed at obtaining an isotropic polygonal field with uniformly distributed colourings. This leaves us with two principal degrees of freedom, one related to the number of admissible colours, the other to the density of so-called V -shaped nodes. Note that the fields discussed below will be used in our image segmentation framework only upon further Gibbsian modification and not in the raw form presented in this section.

2.1 Basic Arak–Surgailis construction

Consider a bounded convex and open domain $D \subseteq \mathbb{R}^2$, fix $k \geq 2$, and let $J := \{1, \dots, k\}$ be the set of (labels of) admissible colours. We define the family $\hat{\Gamma}_D$ of admissible coloured polygonal configurations in D by taking all possible planar graphs γ in $D \cup \partial D$, with non-intersecting straight line segments as edges and colouring $\hat{\gamma}$ the resulting partition of D by one of the members of J , so that the following conditions are satisfied:

- (CPC1) all the interior vertices of γ (lying in D) are of order 2 (we shall refer to them as the V -shaped nodes in the sequel), of order 3 with two out of the three emanating edges co-linear (T -shaped nodes), or of order 4 with two co-linear pairs of emanating edges (X -shaped nodes);
- (CPC2) all the boundary vertices are of degree 1;
- (CPC3) for each straight line $l \subseteq \mathbb{R}^2$ the intersection $l \cap \gamma$ consists of at most one interval of non-zero length and possibly of some isolated points;
- (CPC4) no two sub-regions of D assigned the same colour by $\hat{\gamma}$ are adjacent to each other (i.e. share a common boundary edge).

In the sequel we shall use the notation $\hat{\cdot}$ for coloured polygonal configurations, while the lack of $\hat{\cdot}$ will indicate the corresponding graph separating regions of different colours. Write $\ell(\gamma)$ for the total edge length of γ , $\hat{\gamma} \in \hat{\Gamma}_D$, let $E(\gamma)$ be the set of all edges of γ , and let $N_V(\gamma)$, $N_T(\gamma)$ and $N_X(\gamma)$ stand for the respective numbers of V -, T - and X -shaped nodes. For a family $\{l\} = \{l_i\}_{i=1}^n$ of straight lines intersecting D , denote by $\hat{\Gamma}_D(l)$ the set of those $\hat{\gamma} \in \hat{\Gamma}_D$ for which $\gamma \subseteq \bigcup_{i=1}^n l_i$, and, for each line l_i , the intersection $l_i \cap \gamma$ consists of exactly one non-zero length interval, possibly with some isolated points added.

With the admissible number of colours, k , fixed, the family of polygonal fields discussed below will be parametrised by one further parameter $\alpha_V \in [0, 1]$ that controls the density of V -shaped nodes, see the discussion at the end of this section for further details. Put

$$\alpha_X := 1 - \alpha_V, \quad \alpha_T := \frac{1}{2} \left(1 - \frac{k-2}{k-1} \alpha_X \right). \tag{1}$$

The parameters α_X and α_T control the density of X - and T -shaped nodes, respectively (in a sense that will become clear in the equivalent dynamic representation of the polygonal field model discussed in the next subsection), yet in view of (1) they are uniquely determined by α_V . Further, set

$$\varepsilon := \frac{\alpha_V}{k-1} + \frac{(k-2)\alpha_T}{k-1}.$$

For each coloured polygonal configuration $\hat{\gamma} \in \hat{\Gamma}_D$, define its energy $\Phi(\hat{\gamma}) = \Phi_D^{\alpha_V; k}(\hat{\gamma})$ by

$$\begin{aligned} \Phi(\hat{\gamma}) := & -N_V(\gamma) \log \alpha_V - N_T(\gamma) \log [(k-1)\alpha_T] - N_X(\gamma) \log [(k-1)\alpha_X] \\ & - \text{card}[E(\gamma)] \log(k-1) + 2\varepsilon \ell(\gamma). \end{aligned} \tag{2}$$

It should be noted for reference purposes that this energy expression corresponds to that given in (5.4) in Arak and Surgailis (1989) with the following choice of constants:

- $p_{ij} := \begin{cases} \frac{1}{k-1}, & i \neq j \\ 0, & i = j; \end{cases}$
- $b(i, j) := \alpha_V$, $c(i; j, k) := (k-1)\alpha_T$, $d(i, j, k, m) := (k-1)\alpha_X$;
- $e(i, j) := \varepsilon$, $f(i) = \frac{\pi\alpha_V}{k-1}$;
- $\gamma(dv) := dv/(1+v^2)^{3/2}$.

To proceed, we let Λ_D be the restriction to D of a homogeneous Poisson line process Λ with intensity measure given by the standard isometry-invariant Lebesgue measure μ on the space of straight lines in \mathbb{R}^2 . To construct μ , identify a straight line l with the pair $(\phi, \rho) \in [0, \pi) \times \mathbb{R}$, where $(\rho \sin(\phi), \rho \cos(\phi))$ is the vector orthogonal to l , and join it to the origin; then endow the parameter space $[0, \pi) \times \mathbb{R}$ with the usual Lebesgue measure. With the above notation, define the polygonal Markov field $\hat{\Lambda}_D := \hat{\mathcal{A}}_D^{\alpha_V}$ as the Gibbsian modification of the process induced on $\hat{\Gamma}_D$ by Λ_D with the Hamiltonian (2). In other words, writing $\hat{\Gamma}_D(\Lambda_D)$

for the random set of coloured polygonal configurations induced by the almost surely finite Poisson line process Λ_D hitting D as before,

$$\mathbb{P}(\hat{\mathcal{A}}_D \in \mathcal{E}) = \frac{\mathbb{E} \sum_{\hat{\gamma} \in \hat{\Gamma}_D(\Lambda_D) \cap \mathcal{E}} \exp[-\Phi(\hat{\gamma})]}{\mathbb{E} \sum_{\hat{\gamma} \in \hat{\Gamma}_D(\Lambda_D)} \exp[-\Phi(\hat{\gamma})]} \quad (3)$$

for each $\mathcal{E} \subseteq \hat{\Gamma}_D$, Borel measurable with respect to the topology introduced by Arak and Surgailis (1989, Sect. 2). The expectation is with respect to the distribution of Λ_D . We refer the reader to Arak and Surgailis (1989, Sect. 5) for further details.

The resulting process $\hat{\mathcal{A}}_D$ has a number of remarkable properties. It is isotropic (as easily seen from (3)), exactly solvable in the sense that an explicit formula for the partition function is available, consistent in that $\hat{\mathcal{A}}_D$ coincides in distribution with the restriction of $\hat{\mathcal{A}}_{D'}$ to D for $D' \supseteq D$, and it enjoys a two-dimensional Markov property stating that the conditional behaviour of the process in an open bounded domain depends on the exterior configuration only through arbitrarily close neighbourhoods of the boundary, see Arak and Surgailis (1989). Moreover, for a given straight line l , the one-dimensional section of the field with l is a homogeneous Poisson point process on l of intensity 2, with the induced colouring of the segments between its points distributed uniformly under the condition that no two adjacent segments share the same colour, see Arak and Surgailis (1989, Theorem 5.1, Sect. 5). These nice features are shared by a much broader class of processes, known as consistent polygonal Markov fields, introduced and investigated in detail by Arak and Surgailis (1989, 1991). For interesting alternative point-rather than line-based representations of such models, see Arak, Surgailis and Clifford (1993).

2.2 Dynamic representation

An equivalent description of polygonal Markov fields is available in terms of the equilibrium evolution of one-dimensional particle systems that trace the polygonal realisations of the process in two-dimensional space-time. This description, introduced already in the original Arak paper (Arak 1982) and usually referred to as the *dynamic representation*, turned out to be very useful in establishing the essential properties of these models.

Below, we specialise the general dynamic representation of Arak and Surgailis (1989, Sect. 6) for our particular process $\hat{\mathcal{A}}_D$. We interpret the open convex domain D as a set of *space-time* points $(t, y) \in D$, with t referred to as the *time* coordinate and with y standing for the *spatial* coordinate of a particle at time t . For language convenience, we orient the planar time-space representation such that the time axis extends from the left (past) to the right (future) and the spatial coordinate grows when moving upwards. In these terms, a straight line segment in D stands for a piece of the time-space trajectory of a freely moving particle, and it separates two regions of two different colours, in the sequel referred to as the upper (above the trajectory graph) and lower (below the trajectory) colour of the given particle. For any straight line l that is not parallel to the spatial axis and that crosses the domain D , we define in the obvious way its entry point in D , $\text{in}(l, D) \in \partial D$ and its exit point out(l, D) $\in \partial D$.

We choose the space-time birth coordinates for new particles according to the superposition of a homogeneous Poisson point process with rate $\pi\alpha_V/(k-1)$ in D (interior birth sites) and a Poisson point process on the boundary (boundary birth sites) with intensity measure

$$\kappa(B) = \mathbb{E} \text{card}\{l \in \Lambda, \text{in}(l, D) \in B\}, \quad B \subseteq \partial D. \tag{4}$$

Each interior birth site emits two particles that move with initial velocities $v' < v''$ chosen according to the joint distribution

$$\theta(dv', dv'') := \pi^{-1} |v' - v''| \left(1 + v'^2\right)^{-3/2} \left(1 + v''^2\right)^{-3/2} dv' dv''. \tag{5}$$

This can be shown to be equivalent to choosing the directions of the straight lines representing the space-time trajectories of the emitted particles according to the distribution of the *typical angle* between two lines of Λ , see Sects. 3 and 4 in Arak and Surgailis (1989) and the references therein. The colour in the interior of the angle created by the trajectories of the two particles thus born is chosen uniformly at random from the set of all admissible colours in J except for the one assigned to the exterior of the angle (in other words, except for the colour extending to the left of the birth site).

Each boundary birth site $x \in \partial D$ yields one particle with initial speed v chosen according to the distribution $\theta_x(dv)$ identified by requiring that the direction of the line entering D at x and representing the space-time trajectory of the emitted particle be chosen according to the distribution of a straight line $l \in \Lambda$ conditioned on the event $\{x = \text{in}(l, D)\}$. As before, the colour for the new region (to the right of the new particle trajectory as time progresses from left to right) is chosen uniformly from the set of all admissible colours in J except for the one extending to the left of this trajectory.

In case the region D has exactly one left-extreme point, the choice of the initial colour at this point is done uniformly from J . Otherwise, if D admits a non-trivial left-extreme boundary segment I parallel to the spatial axis, we produce the initial condition for the evolution by constructing the collection of particles born on the segment I , in the way discussed above, and we choose the colours for the resulting subsegments of I by conditioning the uniform choice of colour from J independently for each subsegment on there being no two adjacent subsegments of the same colour.

All particles evolve independently in time according to the following rules:

- (E1) between the critical moments listed below the particles move with constant velocity v so that $dy = vdt$;
- (E2) the time evolution v_t of the velocity of an individual particle is given by a pure-jump Markov process so that

$$\mathbb{P}(v_{t+dt} \in du \mid v_t = v) = \frac{\alpha_V}{k-1} q(v, du)dt$$

for the transition kernel

$$q(v, du) := |u - v|(1 + u^2)^{-3/2} du;$$

(E3) an individual particle moving with velocity v splits with probability

$$\frac{(k-2)\alpha_T}{k-1}q(v, du)dt$$

into two particles, one of which preserves the velocity v while the second moves with velocity $u \in du$. The colour of the angle between the trajectories of the two particles is chosen at random uniformly from the set J excluding the upper and lower colour of the original particle;

(E4) upon a collision of two particles (equal spatial coordinates y at some moment t with $(t, y) \in D$)

(E4a) if the colours $i, j \in J$, respectively, above and below the angle formed by the trajectories of the particles before their collision do coincide then

- With probability α_V both colliding particles die,
- With probability $\alpha_X = 1 - \alpha_V$ [see (1)] both particles survive and their trajectories continue after intersecting at the collision point. The colour for the interior of the resulting angle to the right of the collision point is chosen at random uniformly from the set $J \setminus \{i\} = J \setminus \{j\}$;

(E4b) if the colours $i, j \in J$, respectively, above and below the angle formed by the trajectories of the particles before their collision do not coincide then

- With probability α_T the upper particle survives and the lower one dies,
- With probability α_T the lower particle survives and the upper one dies,
- With probability $((k-2)\alpha_X)/k-1 = 1 - 2\alpha_T$ [see (1)] both particles survive, their trajectories continue past the collision site, and the resulting new angle to the right of the collision site is assigned a random colour uniformly from $J \setminus \{i, j\}$;

(E5) when a particle touches the boundary ∂D , it dies.

The following result is a corollary of Arak and Surgailis (1989, lemma 6.1).

Proposition 1 *The multi-coloured polygonal field \hat{A}_D defined by (3) coincides in distribution with the field resulting from the dynamic construction above.*

2.3 Interpretation of parameters

Below, we briefly discuss the meaning of the parameters α_V , α_T and α_X that play a role in the dynamic construction above. As follows from the dynamic rules (E2-4),

- α_V controls the density of V-shaped nodes of the field,
- α_T controls the density of T-shaped nodes,
- α_X controls the density of X-shaped nodes.

It should be emphasised, however, that these parameters are not independent and, given the number of colours k , they are uniquely determined from α_V by the relation (1). To illustrate the nature of this dependency, we explore two extreme possibilities for choosing α_V .

- Observe first that setting α_V very small results in large survival probabilities for colliding particles, cf. (E4), and thus strongly favours X-shaped nodes at the expense of V-shaped ones. To avoid density explosion of such long-living

particles, the system responds by decreasing the birth rates. For large values of k the parameter α_T also gets small in view of (1). Consequently, in this case, our process resembles the tessellation of D generated by the Poisson line process Λ .

- Setting α_V close to 1 results in small survival probabilities upon particle collisions so that the density of X -shaped nodes decreases. The value of α_T gets close to $1/2$. The behaviour of the system strongly depends on the number of admissible colours k . If $k = 2$, we obtain a field with predominantly V -shaped nodes, close in behaviour to the original Arak process [for which $k = 2$ and $\alpha_V = 0$, see Arak and Surgailis (1989, Sect. 4)]. As k grows, the relative density of T -shaped nodes grows as well.

3 Image segmentation as a statistical inference problem

Here we present a framework for the interpretation of images in terms of multi-coloured polygonal Markov fields.

3.1 The model

The data consist of a discretised image. Write S for the set of sites (pixels in 2D), and L for the set of labels at each site. Thus, the data can be formalised as a vector $\vec{y} = (y_s)_{s \in S}$ with $y_s \in L$. The pixel values y_s may be binary, real or \mathbb{R}^d valued, i.e. $L = \{0, 1\}, \mathbb{R}$ or \mathbb{R}^d . The goal is to interpret \vec{y} in terms of a coloured polygonal configuration $\hat{\gamma} \in \hat{\Gamma}_D$. As before, assume that D is a bounded, convex and open subset of \mathbb{R}^2 . We will impose the further constraint that D contains S .

The model to be used for inference will be a Gibbsian modification of the polygonal random field \mathcal{A}_D by means of a Hamiltonian (sometimes referred to as energy function)

$$\mathcal{H}(\hat{\gamma}) + \mathcal{H}(\vec{y}; \hat{\gamma})$$

that is the sum of two terms. The first one, $\mathcal{H}(\hat{\gamma})$, is a prior or regularisation energy that constrains the geometry so as to favour large polygons with smooth boundaries. In the examples discussed in Sect. 6 below, we took

$$\mathcal{H}(\hat{\gamma}) = \beta \ell(\gamma)$$

proportional to the total edge length. A realisation is shown in Fig. 1. The second term $\mathcal{H}(\vec{y}; \hat{\gamma})$ describes the goodness-of-fit between the coloured configuration $\hat{\gamma}$ and the data \vec{y} . In our examples, we considered Hamiltonians of the form

$$\mathcal{H}(\vec{y}; \hat{\gamma}) = \beta_p \sum_{\hat{p} \in \hat{\gamma}} \mathcal{H}_p(\vec{y}; \hat{p})$$

for $\beta_p > 0$. The sum is over the coloured polygons \hat{p} induced by $\hat{\gamma}$. For computational reasons, $\mathcal{H}_p(\vec{y}; \hat{p})$ depends only on those y_s for which $s \in p$. The idea is



Fig. 1 Sample of Arak process

to penalise inhomogeneous polygons of the wrong colour. For instance, if $\vartheta(\hat{p})$ is the colour of polygon p , $\mathcal{H}_p(\vec{y}; \hat{p})$ could be defined as

$$\sum_{s \in \hat{p}} |y_s - \vartheta(\hat{p})|^q$$

($q = 1, 2$), the choice we make in the example section.

3.2 Probabilistic interpretation

Suppose the observed data \vec{y} can be seen as a realisation of an independent noise model (Baddeley and Van Lieshout 1992). More precisely, the log likelihood $-\mathcal{H}(\vec{y}; \hat{\gamma})$ of \vec{y} given a coloured polygonal configuration $\hat{\gamma}$ is a sum of pixel error terms

$$-\mathcal{H}(\vec{y}; \hat{\gamma}) = \sum_{s \in S} \log h(y_s | \vartheta_s(\hat{\gamma})) \tag{6}$$

where $\{h(\cdot | \vartheta) : \vartheta \in \Theta\}$ is a family of probability densities indexed by a parameter space Θ , and $\vartheta(\hat{\gamma}) = (\vartheta(\hat{\gamma})_s)_{s \in S}$ is a Θ -valued image determined by $\hat{\gamma}$ which we call the signal. Typical examples include the following.

Gauss pixel values are normally distributed with fixed variance σ^2 and mean $\vartheta(\hat{\gamma}) \in \Theta = \mathbb{R}$, so $h(y | \vartheta) = (2\pi\sigma^2)^{-1/2} \exp[-(y - \vartheta)^2 / (2\sigma^2)]$, hence

$$\mathcal{H}(\vec{y}; \hat{\gamma}) = \frac{1}{2} |S| \log(2\pi\sigma^2) + \frac{1}{2\sigma^2} \sum_{s \in S} (y_s - \vartheta_s(\hat{\gamma}))^2.$$

Laplace pixel values are double exponentially distributed with fixed dispersion parameter λ and mean $\vartheta(\hat{\gamma}) \in \Theta = \mathbb{R}$, so $h(y|\vartheta) = \frac{\lambda}{2} \exp[-\lambda|y - \vartheta|]$, hence

$$\mathcal{H}(\vec{y}; \hat{\gamma}) = -|S| \log(\lambda/2) + \lambda \sum_{s \in S} |y_s - \vartheta_s(\hat{\gamma})|$$

Binary pixel values fall into $\{0, 1\}$ and are Bernoulli distributed with $\mathbb{P}(y = \vartheta|\vartheta) = p$ and $\mathbb{P}(y = 1 - \vartheta|\vartheta) = 1 - p$, so that

$$\mathcal{H}(\vec{y}; \hat{\gamma}) = -|S| \log(1 - p) - |\vec{y} \Delta \hat{\gamma}| \log\left(\frac{p}{1 - p}\right),$$

where $\vec{y} \Delta \hat{\gamma} = \{s \in S : y_s \neq \vartheta_s(\hat{\gamma})\}$ is the set of sites at which the observed colour does not match that of the polygon which contains the site.

We conclude that for binary images an L_1 or misclassification criterion is a natural choice that corresponds to a Bernoulli distribution for each pixel. For real valued observations, a white noise model gives rise to an L_2 choice for the Hamiltonian; for other noise distributions using the L_1 distance may be preferred (Baddeley and Van Lieshout 1992).

For the models discussed above, the maximum likelihood solution $\hat{\gamma}^*$ minimising $\mathcal{H}(\vec{y}; \hat{\gamma})$ with respect to $\hat{\gamma} \in \hat{\Gamma}_D$ is readily described. Indeed, for *Gauss*, it is the solution of a least squares regression of \vec{y} on the class of functions $\{s \mapsto \vartheta_s(\hat{\gamma}) : \hat{\gamma} \in \hat{\Gamma}_D\}$, for *Laplace* of a least absolute deviation regression, and for the binary model, provided $p < 1/2$, $\hat{\gamma}^*$ minimises the misclassification rate $|\vec{y} \Delta \hat{\gamma}|/|S|$. In general, the maximum likelihood estimator is not unique. Moreover, it tends to result in an over-segmentation, in the sense that a polygon of the closest colour around each cluster of like coloured pixels would be a solution to the maximum likelihood equation.

Clearly, more complicated probability densities $f(\vec{y}|\hat{\gamma})$ can be devised to deal with textures, or unequal variance and dispersion parameters; in this case the Hamiltonian reads $\mathcal{H}(\vec{y}; \hat{\gamma}) = -\log f(\vec{y}|\hat{\gamma})$. One could additionally include a term $\beta e \sum_{e \in E(\gamma)} \mathcal{H}_e(\vec{y}; e)$ in $\mathcal{H}(\vec{y}; \hat{\gamma})$ to favour clear edges. For instance $-\mathcal{H}_e(\vec{y}; e)$ may be the (squared) difference in average label in two regions on either side of the edge (since this term is bounded, by adjusting the prior edge term, finiteness of the corresponding partition functions may be ensured), or be based on the gradient along profiles perpendicular to the edge (Ortner 2004).

3.3 Model parameters

In general, the Hamiltonian will depend on parameters, i.e. $\mathcal{H}(\hat{\gamma}) = \mathcal{H}(\hat{\gamma}; \omega)$ and $\mathcal{H}(\vec{y}; \hat{\gamma}) = \mathcal{H}(\vec{y}; \hat{\gamma}; \tilde{\omega})$. Thus, statistical inference must be based on the Gibbsian modification with Hamiltonian

$$\mathcal{H}(\hat{\gamma}; \omega) + \mathcal{H}(\vec{y}; \hat{\gamma}; \tilde{\omega}). \tag{7}$$

Formula (6) is constructed out of four basic ingredients. The coloured configuration $\hat{\gamma}$ provides a concise description of the image, the parameter $\tilde{\omega}$ is crucial in linking the description to the observed image, while the vector ω regulates the

strength and type of geometric constraints, or ‘smoothness’ of the description. Finally, the parameters α_V and k of the dominating random field $\hat{\mathcal{A}}_D$ govern the relative frequency of V-shaped joins and the number of colour labels into which to segment.

More specifically, write

$$\ell(\vec{\mathbf{y}}|\hat{\gamma}; \tilde{\omega}) = \frac{\exp[-\mathcal{H}(\vec{\mathbf{y}}; \hat{\gamma}; \tilde{\omega})]}{Z(\hat{\gamma}; \tilde{\omega})} \quad (8)$$

for the normalised likelihood with respect to the product of counting or Lebesgue measure on L^S , depending on whether L is finite or Euclidean. Since the partition function $Z(\hat{\gamma}; \tilde{\omega}) := \sum_{\vec{\mathbf{z}}} \exp[-\mathcal{H}(\vec{\mathbf{z}}; \hat{\gamma}; \tilde{\omega})]$ for discrete L (or $Z(\hat{\gamma}; \tilde{\omega}) := \int \exp[-\mathcal{H}(\vec{\mathbf{z}}; \hat{\gamma}; \tilde{\omega})] d\vec{\mathbf{z}}$ for continuous L , in the sequel we shall always use \sum in this context to avoid doubling equations) is explicitly known (see the discussion in Subsect. 3.2), we can incorporate it in the Hamiltonian and assume without loss of generality that

$$Z(\hat{\gamma}; \tilde{\omega}) \equiv 1. \quad (9)$$

We assume in addition that the regularisation parameter vector ω is fixed and not subject to statistical inference, and that a reference measure $d\pi(\tilde{\omega})$ (possibly improper) is given on the space of admissible values of the goodness-of-fit parameter vector $\tilde{\omega}$. Write $\mathcal{L}(\hat{\mathcal{A}}_D)$ for the probability distribution of the coloured polygonal random field $\hat{\mathcal{A}}_D$, and define a joint Radon–Nikodym derivative

$$f(\vec{\mathbf{y}}; \hat{\gamma}; \tilde{\omega}) = \frac{\exp[-\mathcal{H}(\vec{\mathbf{y}}; \hat{\gamma}; \tilde{\omega}) - \mathcal{H}(\hat{\gamma}; \omega)]}{Z(\omega)},$$

where

$$Z(\omega) := \sum_{\vec{\mathbf{z}}} \int \int \exp[-\mathcal{H}(\vec{\mathbf{z}}; \hat{\gamma}; \tilde{\omega}) - \mathcal{H}(\hat{\gamma}; \omega)] d[\mathcal{L}(\hat{\mathcal{A}}_D)](\hat{\gamma}) d\pi(\tilde{\omega}),$$

with respect to the product measure of $d\mathcal{L}(\hat{\mathcal{A}}_D)$, $d\pi(\tilde{\omega})$, and the reference distribution for $\vec{\mathbf{y}}$, provided $Z(\omega) < \infty$. This is of course in agreement with (8) and can be interpreted probabilistically by imposing the joint density

$$f(\hat{\gamma}; \tilde{\omega}) := \frac{\exp[-\mathcal{H}(\hat{\gamma}; \omega)] \sum_{\vec{\mathbf{y}}} \exp[-\mathcal{H}(\vec{\mathbf{y}}; \hat{\gamma}; \tilde{\omega})]}{Z(\omega)} = \frac{\exp[-\mathcal{H}(\hat{\gamma}; \omega)]}{Z(\omega)}$$

with respect to $d\mathcal{L}(\hat{\mathcal{A}}_D) \times d\pi(\tilde{\omega})$ on $(\hat{\gamma}, \tilde{\omega})$, cf. (9). For a given data image $\vec{\mathbf{y}}$, we thus arrive at the Gibbsian modification corresponding to the Hamiltonian in (7), with normalising constant

$$Z(\omega; \vec{\mathbf{y}}) = \int \int \exp[-\mathcal{H}(\vec{\mathbf{y}}; \hat{\gamma}; \tilde{\omega}) - \mathcal{H}(\hat{\gamma}; \omega)] d[\mathcal{L}(\hat{\mathcal{A}}_D)](\hat{\gamma}) d\pi(\tilde{\omega}).$$

Since $\pi(\cdot)$ may be improper, we must impose the condition that the integral is finite.

Our goal is to find a vector $(\hat{\gamma}^*, \tilde{\omega}^*)$ for which the value of the Hamiltonian in (7) is a sufficiently good approximation of the infimum. To this end, we use a simulated annealing algorithm that will be discussed in detail in the next section.

Estimation of α_V and k is complicated by the fact that the algorithms developed in this paper can only simulate the polygonal fields for fixed values of α_V and k ; in particular, the state space $\hat{\Gamma}_D = \hat{\Gamma}_D^k$ depends on the choice of k . Furthermore, we were unable to construct random dynamics for updating k and α_V that do not lead to a considerable deterioration of the quality of the configuration currently reached by the algorithm. For instance the transition $k \mapsto k - 1$ would require removing one colour in the present configuration. If $k \leq 4$, this may well be impossible without completely rebuilding the current configuration (recall that colours assigned to adjacent regions have to differ in our model), thus leading to a waste of computational effort. For $k > 4$ recolouring the current configuration in $k - 1$ colours, though theoretically possible, is a computationally demanding (NP-hard) task.

To overcome such problems, we propose to restrict the admissible values of α_V and k to some finite subsets of $[0, 1]$ and \mathbb{N} , respectively, to perform the optimisation of other parameters separately for each (α_V, k) within this restricted domain, and then to compare the resulting solutions outputting the optimal one. In spite of its obvious general disadvantages, such an exhaustive search technique over the parameter space seems nevertheless preferable if the range of admissible values for α_V and k is relatively small. Recast as a model selection problem, we choose that admissible value for which the Bayesian information criterion

$$\text{BIC}(\alpha_V, k) = -2\mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}_{\alpha_V, k}^*; \tilde{\omega}_{\alpha_V, k}^*) - \dim(\tilde{\omega}_{\alpha_V, k}^*) \log(|S|) \quad (10)$$

achieves its maximal value. Recall that $(\hat{\gamma}_{\alpha_V, k}^*, \tilde{\omega}_{\alpha_V, k}^*)$ is the (approximate) mode of the model (7) with parameters α_V and k for the reference distribution. For Hamiltonians of the form (6), (10) reduces to

$$\text{BIC}(\alpha_V, k) = 2 \sum_{s \in S} \log h(y_s | \hat{\gamma}_{\alpha_V, k}^*, \tilde{\omega}_{\alpha_V, k}^*) - \dim(\tilde{\omega}_{\alpha_V, k}^*) \log(|S|).$$

4 Monte Carlo inference

The purpose of the current section is to develop a simulated annealing algorithm for image segmentation by multi-colour polygonal Markov fields. To that end, we begin by developing, much along the lines of Schreiber (2005, sect. 2.1), random dynamics on the space $\hat{\Gamma}_D$ of admissible coloured polygonal configurations which are reversible and leave the law of $\hat{\mathcal{A}}_D$ invariant. This will allow us further to provide modifications suitable for Metropolis–Hastings simulation of general Gibbsian modifications of $\hat{\mathcal{A}}_D$, which in turn will be incorporated in a Gibbs sampler for (7).

4.1 Basic dynamics

In the sequel, particular care is needed to distinguish between the notion of time considered in the dynamic representation of the polygonal field $\hat{\mathcal{A}}_D$ given in Sect. 2,

and the notion of time to be introduced for the continuous time random dynamics on $\hat{\Gamma}_D$ constructed below. To make this distinction clear, we shall refer to the former as the *representation time* (r-time for short) and shall keep for it the notation t , while the latter will be called the *simulation time* (s-time for short) and denoted by s in the sequel.

For the exposition below, it is convenient to imagine that each individual interior or boundary particle birth site in the dynamic representation comes with an associated infinite sequence of i.i.d. standard exponential random variables, a *Poisson clock*, to determine the subsequent moments of critical events **(E2)**, **(E3)** (velocity updates and splits) as well as with an associated independent infinite sequence of i.i.d. uniform random variables on $[0, 1]$ used to determine the new velocities upon critical events, to produce the death/survival events of **(E4)** with the required probabilities, and to choose with the specified probabilities the colours for new regions arising in the course of the evolution. In other words, each birth site is assumed to carry a package enclosing all randomness the resulting particles may encounter during their evolution, and the above is just one technical possibility of how this can be achieved. Note that such a package can easily be implemented in a computer simulation by storing random number generator seeds. We shall use the name *birth package* for a birth site carrying such randomness-determining information. In these terms, it is now easily seen that the coloured polygonal configuration obtained in the course of the dynamic construction depends *deterministically* on the underlying collection of birth packages.

Consider a coloured polygonal configuration $\hat{\gamma}$ and a new birth site $x_0 \in D$ not yet present in $\hat{\gamma}$, extended to a birth package in the standard way discussed above. Adding this birth package to the collection of birth packages that determine $\hat{\gamma}$ and keeping the evolution rules **(E1–5)** of the dynamic representation results in a new configuration to be denoted by $\hat{\gamma} \oplus x_0$. Likewise, removing a birth site x_1 from a coloured configuration $\hat{\gamma}$ in which it was present yields a new configuration $\hat{\gamma} \ominus x_1$.

Recall that the collection of birth sites for \hat{A}_D is chosen according to a Poisson point process with intensity measure $\pi\alpha_V(k-1)^{-1}dx + \kappa(dx)$ specified in (4). Hence, the law of \hat{A}_D is invariant with respect to the following pure-jump Markovian birth site birth and death dynamics on $\hat{\Gamma}_D$, denoted by **(BS)** in the sequel, with $\hat{\gamma}_s$ standing for the state at time s :

- (BS:birth)** with intensity $[\pi\alpha_V(k-1)^{-1}dx + \kappa(dx)]ds$ set $\hat{\gamma}_{s+ds} := \hat{\gamma}_s \oplus x$;
- (BS:death)** for each birth site in $\hat{\gamma}_s$ with intensity $1 \cdot ds$ set $\hat{\gamma}_{s+ds} := \hat{\gamma}_s \ominus x$;
- (BS:recolour)** with intensity τ , where $\tau > 0$ is a fixed parameter of the dynamics, recolour the configuration $\hat{\gamma}_s$ by keeping the configuration γ_s and assigning new colours to the corresponding polygonal regions uniformly at random on the event that no two adjacent regions share the same colour.

In fact, more can be stated, see also Schreiber (2005, Prop. 1).

Theorem 1 *The distribution of the polygonal field \hat{A}_D is the unique invariant law of the dynamics given by **(BS:birth)**, **(BS:death)** and **(BS:recolour)**. The resulting stationary process is reversible. Moreover, for any initial distribution of $\hat{\gamma}_0$, the laws of the random polygonal fields $\hat{\gamma}_s$ converge in total variation to the law of \hat{A}_D as $s \rightarrow \infty$.*

Indeed, with the set of moves restricted to **(BS:birth)** and **(BS:death)** the detailed balance, implying the invariance and reversibility, follows just by basic properties of the Poisson point process determining the birth sites in the dynamic construction as discussed in subsection 2.2 (recall that the ‘contents’ of the birth packages are independent). Now, it is easily seen that the **(BS:recolour)** step is self-reversible which yields the required reversibility for the complete **(BS:birth, death, recolour)** dynamics. A short justification is still required for the uniqueness and convergence statements in the above theorem. They both follow from the observation that, regardless of the initial state, the process $\hat{\gamma}_s$ spends a non-null fraction of time in the state with no polygonal contours and the whole domain D coloured in one fixed colour, say $1 \in J$. Indeed, this observation allows us to conclude the required uniqueness and convergence by a standard coupling argument as in Theorem 1.2 in Liggett (1985).

4.2 Dynamics for Gibbsian modifications

Assume that a Hamiltonian \mathcal{H} is defined on the space $\hat{\Gamma}_D$ of admissible coloured polygonal configurations. In our applications \mathcal{H} will be of the form

$$\mathcal{H}(\hat{\gamma}) := \mathcal{H}(\vec{y}; \hat{\gamma}, \tilde{\omega}) + \mathcal{H}(\hat{\gamma}; \omega),$$

see the framework of Sect. 3, yet the setting below is more general and we only require Ruelle stability, i.e. that

$$\mathcal{H}(\hat{\gamma}) \geq -B$$

for some positive constant B , which clearly guarantees that the partition function

$$Z_D[\mathcal{H}] := \mathbb{E} \exp \left[-\mathcal{H}(\hat{A}_D) \right]$$

is finite. Consequently, the corresponding Gibbsian modification $\hat{\mathcal{A}}_D^{\mathcal{H}}$ can be defined by

$$\frac{d\mathcal{L}(\hat{\mathcal{A}}_D^{\mathcal{H}})}{d\mathcal{L}(\hat{A}_D)}[\hat{\gamma}] = \frac{\exp \left[-\mathcal{H}(\hat{\gamma}) \right]}{Z_D[\mathcal{H}]}. \tag{11}$$

Consider the following modification of the basic **(BS)** dynamics:

(BS[H]:birth) with intensity $[\pi \alpha_V (k - 1)^{-1} dx + \kappa(dx)] ds$ set $\hat{\delta} := \hat{\gamma}_s \oplus x$. Then, with probability $\min \left(1, \exp \left[\mathcal{H}(\hat{\gamma}_s) - \mathcal{H}(\hat{\delta}) \right] \right)$ put $\hat{\gamma}_{s+ds} := \hat{\delta}$, otherwise keep $\hat{\gamma}_{s+ds} := \hat{\gamma}_s$;

(BS[H]:death) for each birth site x in $\hat{\gamma}_s$, with intensity $1 \cdot ds$ set $\hat{\delta} := \hat{\gamma}_s \ominus x$. Then, with probability $\min \left(1, \exp \left[\mathcal{H}(\hat{\gamma}_s) - \mathcal{H}(\hat{\delta}) \right] \right)$ put $\hat{\gamma}_{s+ds} := \hat{\delta}$, otherwise keep $\hat{\gamma}_{s+ds} := \hat{\gamma}_s$;

(BS[H]:recolour) with intensity τ , where $\tau > 0$ is a fixed parameter of the dynamics, produce a recoloured version $\hat{\delta}$ of $\hat{\gamma}_s$ by keeping the configuration γ_s and assigning new colours to the corresponding polygonal regions uniformly at random on the event that no two adjacent regions share the same colour. Then, with probability $\min\left(1, \exp\left[\mathcal{H}(\hat{\gamma}_s) - \mathcal{H}(\hat{\delta})\right]\right)$ put $\hat{\gamma}_{s+ds} := \hat{\delta}$, otherwise keep $\hat{\gamma}_{s+ds} := \hat{\gamma}_s$.

In other words, the original dynamics **(BS)** are used to propose a new configuration $\hat{\delta}$, which is then accepted with probability $\min\left(1, \exp\left[\mathcal{H}(\hat{\gamma}_s) - \mathcal{H}(\hat{\delta})\right]\right)$ and rejected otherwise. By a straightforward verification of the detailed balance conditions and an appeal to theorem 1, we obtain the following result.

Theorem 2 *The distribution of the polygonal field $\hat{\mathcal{A}}_D^{\mathcal{H}}$ given by (11) is the unique invariant law of the dynamics described by **(BS[H]:birth)**, **(BS[H]:death)**, and **(BS[H]: recolour)**. The resulting stationary process is reversible. Moreover, for any initial distribution of $\hat{\gamma}_0$, the laws of the random polygonal fields $\hat{\gamma}_s$ converge in total variation to the law of $\hat{\mathcal{A}}_D^{\mathcal{H}}$ as $s \rightarrow \infty$.*

4.3 Gibbs sampler

To simulate from (7), we use Gibbs sampling, i.e. we sample iteratively from the conditional distribution of each component in turn given that the values of all others are fixed. Thus, if $\tilde{\omega}_n$ is the current value of the likelihood parameters (at the first step assigned perhaps by some crude histogram based method), the iterative procedure is as follows:

(G:configuration) draw a realisation $\hat{\gamma}_n$ of $\hat{\Gamma}$ from the set $\hat{\Gamma}_D$ of coloured polygonal configurations distributed (approximately) as

$$\frac{\exp\left[-\mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}_n; \tilde{\omega}_n) - \mathcal{H}(\hat{\gamma}_n; \omega)\right] d[\mathcal{L}(\hat{\mathcal{A}}_D)](\hat{\gamma}_n)}{\int \exp\left[-\mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}; \tilde{\omega}_n) - \mathcal{H}(\hat{\gamma}; \omega)\right] d[\mathcal{L}(\hat{\mathcal{A}}_D)](\hat{\gamma})}$$

by the method described in the previous section;

(G:parameters) sample a new parameter value $\tilde{\omega}_{n+1}$ from the distribution

$$\frac{\exp\left[-\mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}_n; \tilde{\omega}_{n+1})\right] d\pi(\tilde{\omega}_{n+1})}{\int \exp\left[-\mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}_n; \tilde{\omega})\right] d\pi(\tilde{\omega})}.$$

Clearly, each step preserves the Gibbs measure given by (7).

4.3.1 Simulated annealing

Let $\beta > 0$, and define β -modified probability measures by their Radon–Nikodym derivatives

$$\frac{d\mu^\beta(\hat{\gamma}; \tilde{\omega})}{d[\mathcal{L}(\hat{\mathcal{A}}_D)](\hat{\gamma}) \times \pi(\tilde{\omega})} = \frac{\exp\left[-\beta(\mathcal{H}(\hat{\gamma}; \omega) + \mathcal{H}(\tilde{\mathbf{y}}; \hat{\gamma}; \tilde{\omega}))\right]}{Z^\beta(\omega, \tilde{\mathbf{y}})}, \quad (12)$$

where $Z^\beta(\omega, \vec{y}) = Z[\beta(\mathcal{H}(\cdot; \omega)) + \mathcal{H}(\vec{y}; \cdot; \cdot)]$. Note that multiplication by a constant does not affect Ruelle stability, hence μ^β is well-defined. Intuitively, for small β , the dominating measure is approximated, while for large values (12) tends to put its mass on the set M of modes. Since M might well have mass 0 in our context, the series of measures μ^β will not converge in total variation (Haario and Saksman 1991, Thm. 3.3) as $\beta \rightarrow \infty$, although weak convergence to a measure concentrated on M could be proven rigorously.

5 The implementation

The purpose of the present section is to briefly discuss our implementation of the simulation algorithm as introduced in Sect. 4. Note that in its current version our implementation has been developed only for the bi-coloured case ($k=2$) which excludes T- and X-shaped nodes and leaves us with V-shaped nodes only, whence the simulated polygonal field becomes a collection of non-intersecting contours, possibly chopped off by the boundary, as indeed it was in the original Arak (1982) construction. Consequently, in the discussion below, we specialise to the bi-coloured case; the implementation of the simulation algorithm for the general multi-coloured setting is the subject of our current work in progress, beyond the scope of the present paper.

5.1 Representing polygonal configurations

A configuration of a bi-coloured polygonal Markov field is represented as a list of labelled vertices. The full description of a vertex is provided by

- The Cartesian coordinates of the vertex;
- A pointer to the *parent vertex*, i.e. the vertex initiating the segment which terminates at the current vertex; this field is set to NULL for vertices at birth sites;
- The angle the segment terminating at the current vertex forms with the horizontal axis;
- The *virtual length* of the segment initiated by the parent of the current vertex; this is the length this segment would have if the corresponding particle were the only one present in the system and evolved in an empty environment. The actual length of this segment is usually smaller due to collisions.

The list of vertices is sorted by increasing x -coordinate (r-time coordinate).

5.2 Generating the initial configuration

The initial configuration for our simulation procedure is generated according to the basic non-interacting Arak process, see Arak (1982) or Arak and Surgailis (1989, Sect. 4). This is in fact an arbitrary choice in view of the mixing property of our birth site birth and death dynamics stated in Theorems 1 and 2, yet we find it convenient to describe briefly our implementation of the original Arak dynamic

construction first, and then to use it as a basis for the discussion of the proposal steps in **(BS[H])**.

The initial Arak-distributed configuration is generated in a single left-to-right sweep through the image domain, according to the dynamic construction of Sect. 2.2. This is done by successively updating in r -time two priority queues that store, respectively,

- The birth sites with r -time coordinate exceeding the current r -time, and *virtual end points* of segments generated so far, for which the r -time coordinate exceeds the current r -time. Note that by a *virtual end point* of a segment we understand the point where it would end should the corresponding particle move in empty environment not subject to collisions with other particles; by definition the distance between the initial point of a segment and its virtual endpoint coincides with its virtual length,
- *Virtual collision points* which are all possible pairwise intersection points of currently existing *virtual segments* (i.e. segments joining an initial point to its corresponding virtual end point) lying forward in r -time.

At each step of the algorithm, the next vertex to arise in course of the r -time evolution is determined by choosing that vertex that minimises the r -time coordinate in both queues. Consequently, the contents of these queues can be regarded as a current collection of ‘candidates’ for the next vertex.

5.3 Modifying polygonal configurations

As discussed in Sects. 4.1 and 4.2, the crucial update operations of our Monte Carlo algorithm are those of adding **(BS[·] : birth)** and removing **(BS[·] : death)** a birth site in a polygonal configuration, the third step **(BS[·] : recolour)** is greatly simplified by restricting the number of colours to two.

To add a new birth site, we first choose its position uniformly at random within the image domain, and then we let the new-born particles evolve and interact with the existing ones according to the rules **(E1–5)**. Note that it often results in updating a large portion of the original configuration. Likewise, removing a given birth site while obeying the evolution rules **(E1–5)** for all remaining particles may also lead to a non-local configuration update. Both these updates are implemented using the same data structures as discussed in Sect. 5.2.

5.4 Evaluating the Hamiltonian

Due to the bi-coloured nature of our current implementation, we have tested our algorithm mainly for the L_1 -type Hamiltonians discussed in subsect. 3.2. This required the evaluation of the number of misclassified pixels upon each update proposal. In the current version of our software this is done by using a line sweep-and-scan technique.

Various modifications can be used to improve the efficiency of the basic algorithm described above. For example, due to isotropy, one could choose a random direction for time before each birth or death step; the line-scanning computation of the areas can be avoided by an application of the divergence theorem.

6 Examples

In this section we present sample segmentation results produced by our algorithm. The current version of our implementation is designed predominantly to segment binary images. However, we provide as well an example showing that our software is also able to capture main objects from grayscale images with high contrast between the object and the background regions. Below, Figs. 2 and 3 present segmentation results of binary images (flower, stickman) in the presence of heavy correlated noise, with their boundaries strongly blurred and spread. Further, Fig. 4 presents a segmentation of a grayscale scene (mushroom). In all cases we used a linear cooling schedule, which turned out to be very effective. The results for the flower, the stickman and the mushroom were achieved after, respectively, 2.1, 2.6 and 3.5 million steps, with the average execution time being close to 0.015 s per iteration for 2 CPUs architecture with Intel Xeon 2.4GHz processors and 2 GB RAM memory. It should be emphasised that updating the polygonal configurations contributes just a minor fraction of the computation time, its major part being spent on evaluating the amount of misclassified pixels as required for determining the Hamiltonian – further optimisation is envisioned here, which is a subject of our current research in progress.

The images in Figs. 2a and 4a are accompanied by the corresponding graphs (Figs. 2b and 4b, respectively) visualising the decay of the pixel misclassification rate during the execution of our algorithm. We see that rates close to 10% are usually achieved after about 100, 000 iterations, thanks to the global character of initial configuration updates. After that, the PMR exhibits much slower decay, eventually achieving in all examples 3%. Note that the fact that the misclassification rates do not fall below that value is a natural consequence of the noisy character (Figs. 2, 3) and 256 gray-levels of pixel values (Fig. 4) of the images processed. In Fig. 3b, a graph of the acceptance rate is shown as a function of the number of iterations. It can be seen that initially, almost all moves are accepted; when a reasonable fit is achieved, the rate drops to below 10%.

It is seen from the sample segmentations above that our technique easily deals with the noise and extracts the major objects present in the scenery. The price for this noise immunity and robustness of our approach is that its performance at the level of fine details is much worse. These features stand in strong opposition to the usual lattice-based Markov field methods, where fine details are usually well processed but essential problems arise at the level of the global structure of the image. We aim at solving this problem in our future work, by using multi-resolution techniques and various kinds of post-processing and smoothing.

7 Conclusion

We envision to continue and extend the work reported in the present paper in a number of natural directions. One clear major task is to push our implementation into the general multi-coloured setting of the paper, getting rid of the bi-coloured restriction. Further, as discussed above, the Hamiltonians imposed in our model were mainly of integral nature, based on global information, the prominent example being the number of misclassified pixels. This is reflected by the behaviour of our algorithm which, as remarked in Sect. 6, well recognises large- and

medium-sized components of the image but is much less accurate at the level of fine details. Consequently, stable with respect to local image perturbations as the global Hamiltonian terms are, they may nevertheless require complementary terms, based on local information such as gradient-based terms that penalise discrepancy between segmenting polygonal edges and natural edges present in the image. Optimising the choice of a particular form of such local contributions and incorporating them into our implementation is a subject of our current work in progress. Another

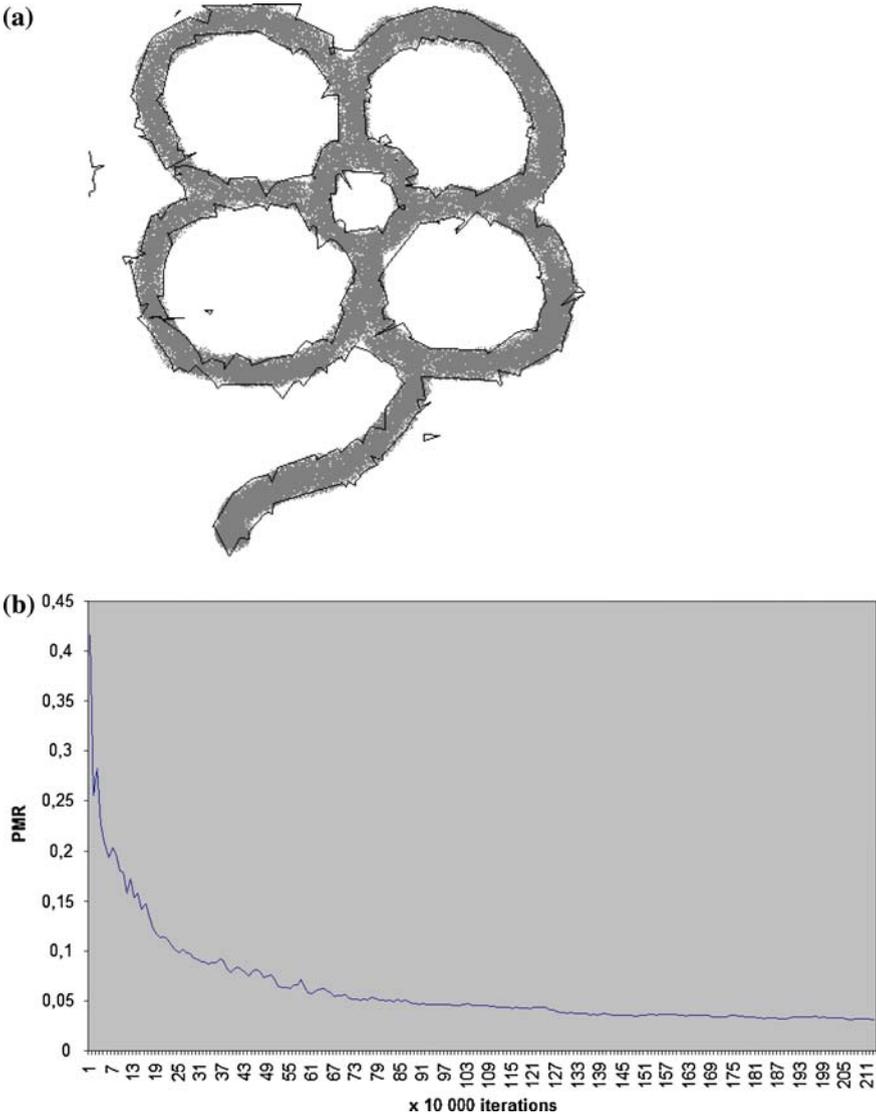


Fig. 2 Segmented flower

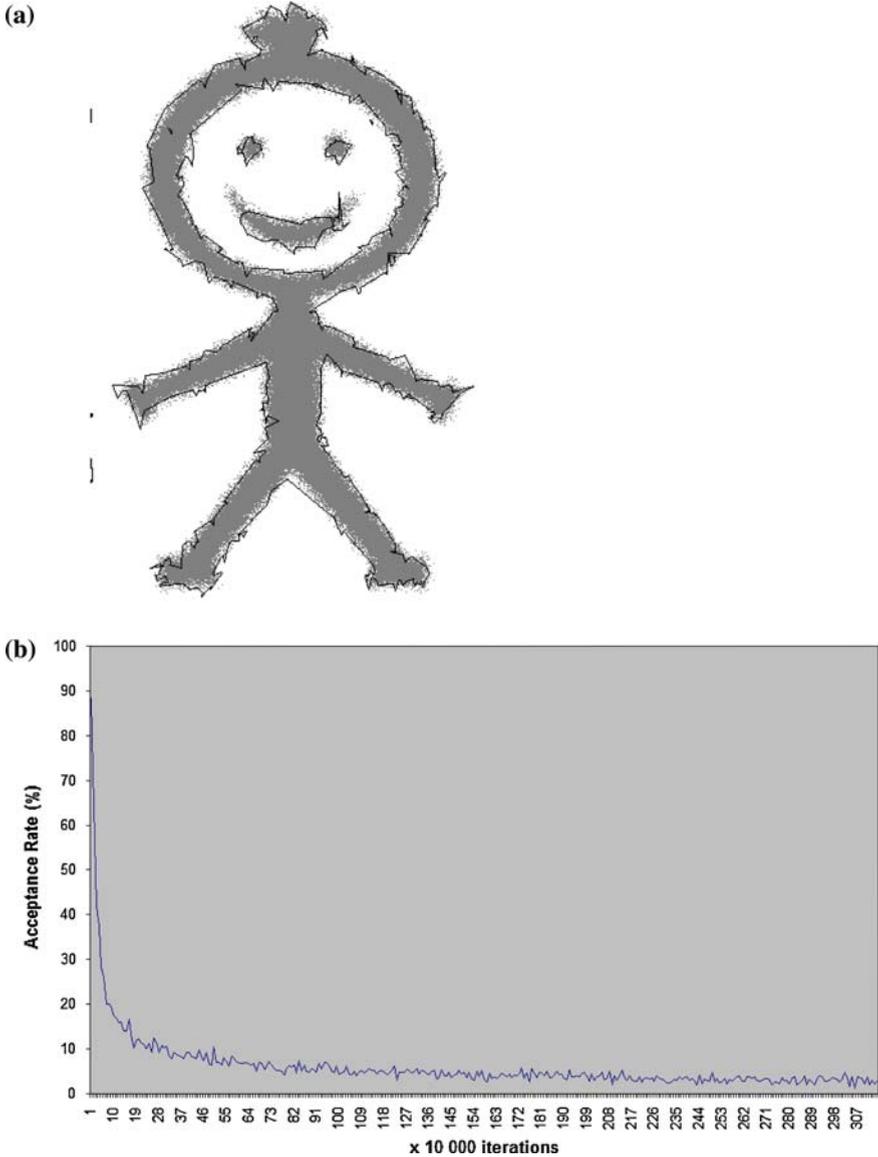


Fig. 3 Segmented stickman

important aspect of our further research is to enable our algorithm to deal with textured regions. It should be noted that this is covered by our theoretical discussion in Sects. 3.3 and 4.3, since the texture information can be interpreted as a component of the likelihood parameter vector $\tilde{\omega}$. A further possible method of improving the behaviour of our algorithm on local details is to introduce some local moves to our (BS[-]) dynamics, close in spirit to those proposed by Clifford and Nicholls (1989).

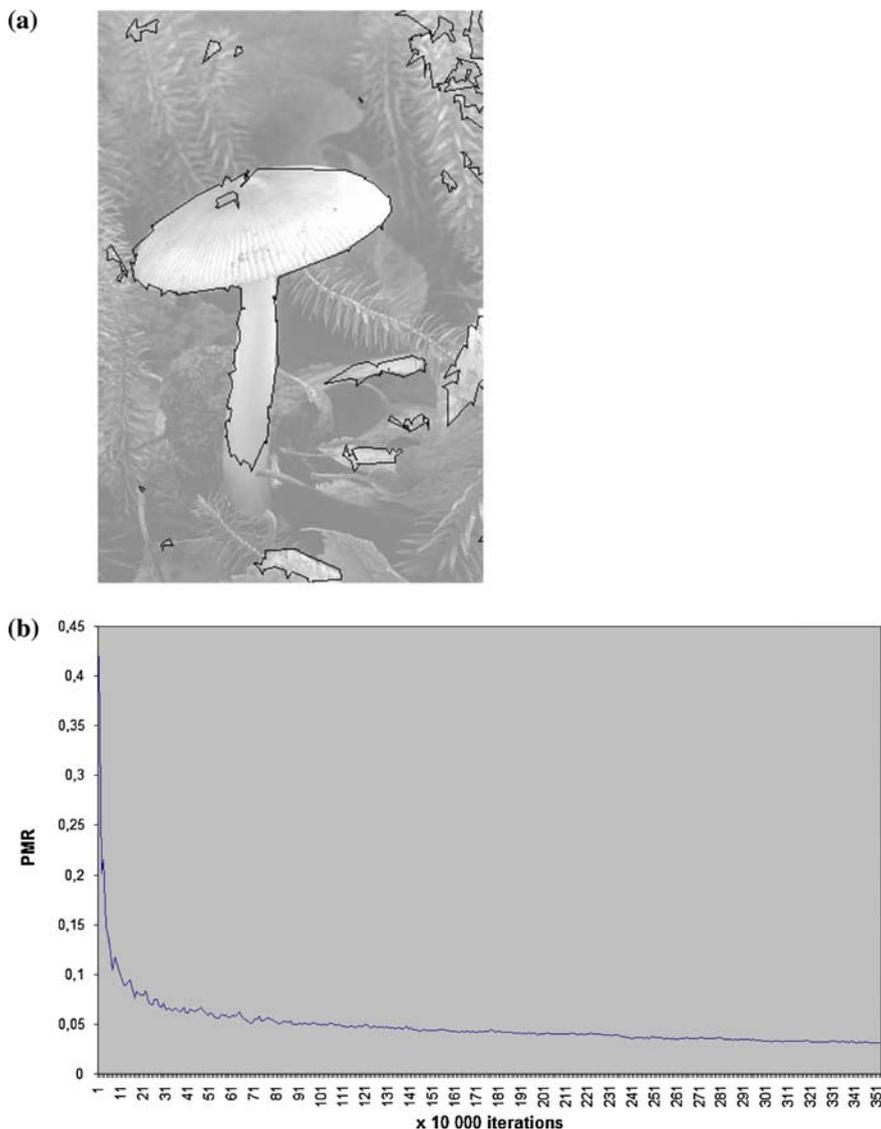


Fig. 4 Segmentation of a grayscale image

Acknowledgements We gratefully acknowledge support from the EC 6th Framework Programme Priority 2 Information Society Technology Network of Excellence MUSCLE (Multimedia Understanding through Semantics, Computation and Learning; FP6-507752). Tomasz Schreiber also wishes to express his gratitude for support from the Foundation for Polish Science (FNP) as well as from the Polish Minister of Scientific Research and Information Technology, grant 1 P03A 018 28 (2005-2007), and to warmly thank for hospitality of the Centrum voor Wiskunde en Informatica (CWI), Amsterdam, The Netherlands, where this research was carried out.

References

- Arak, T. (1982). On Markovian random fields with finite number of values. In: 4th USSR-Japan symposium on probability theory and mathematical statistics, Abstracts of Communications. Tbilisi
- Arak, T., Surgailis, D. (1989). Markov fields with polygonal realisations. *Probabability Theory and Related Fields*, 80, 543–579
- Arak, T., Surgailis, D. (1991). Consistent polygonal fields. *Probabability Theory and Related Fields*, 89, 319–346
- Arak, T., Clifford, P., Surgailis, D. (1993). Point-based polygonal models for random graphs. *Advances in Applied Probability*, 25, 348–372
- Baddeley, A. J., van Lieshout, M. N. M. (1992). ICM for object recognition. In: Y. Dodge, J. Whittaker (Eds.) *Computational Statistics*, Vol. 2 pp. 271–286. Heidelberg: Physica/Springer
- Borkar, V., Mitter, S. K. (1993). Stochastic processes that generate polygonal and related random fields, Research Report LIDS-P2216, MIT
- Clifford, P., Middleton, R.D. (1989). Reconstruction of polygonal images. *Journal of Applied Statistics*, 16, 409–422
- Clifford, P., Nicholls, G. K. (1994). A Metropolis sampler for polygonal image reconstruction, Electronic version available at: http://www.stats.ox.ac.uk/~clifford/papers/met_poly.html
- Gimel'farb, G. L. (1999). Image textures and Gibbs random fields. Dordrecht: Kluwer
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82, 711–732
- Green, P. J., Richardson, S., Hjort N. L. (Eds.) (2003). *Highly structured stochastic systems*, *Oxford Statistical Science Series*, Vol. 27. Oxford: Oxford University Press
- Haario, H., Saksman, E. (1991). Simulated annealing process in general state space. *Advances in Applied Probability*, 23, 886–893
- Heikkinen, J., Arjas, E. (1998). Non-parametric Bayesian estimation of a spatial Poisson intensity. *Scandinavian Journal of Statistics*, 25, 435–450
- Hurn, M. A., Husby, O., Rue, H. (2003). Advances in Bayesian image analysis. In: P.J. Green, S. Richardson, N. L. Hjort (Eds.) *Highly structured stochastic systems*, *Oxford Statistical Science Series* (Vol. 27, pp. 323–325). Oxford: Oxford University Press
- Liggett, T. (1985). *Interacting particle systems*. New York: Springer
- Møller, J., Skare, Ø. (2001). Bayesian image analysis with coloured Voronoi tessellations and a view to applications in reservoir modelling. *Statistical Modelling*, 1, 213–232
- Møller, J., Waagepetersen, R. P. (1998). Markov connected component fields. *Advances in Applied Probability (SGSA)*, 30, 1–35
- Nicholls, G.K. (1998). Bayesian image analysis with Markov chain Monte Carlo and coloured continuum triangulation models. *Journal of the Royal Statistical Society, Series B*, 60, 643–659
- Ortner, M. (2004). Processus ponctuels marqués pour l'extraction automatique de caricatures de bâtiments à partir de modèles numériques d'élévation. PhD Thesis, University of Nice–Sophia Antipolis.
- Osher, S., Paragios, N. (Eds.) (2003). *Geometric level set methods in imaging, vision, and graphics*. New York: Springer
- Paskin, M. A., Thrun, S. (2005). Robotic mapping with polygonal random fields. In: *Proceedings in Artificial Intelligence UAI-05*
- Rosenfeld, A., Kak, A. C. (1982). *Digital picture processing*, 2nd edn, Vol. 2. Orlando: Academic Press
- Schreiber, T. (2005). Random dynamics and thermodynamic limits for polygonal Markov fields in the plane. *Advances in Applied Probability*, 37 (4) (in press).
- Tjelmeland, H., Besag, J. (1998). Markov random fields with higher order interactions. *Scandinavian Journal of Statistics*, 25, 415–433
- Vincent, L. (1999). Mathematical morphology. In: O. Barndorff-Nielsen, W. S. Kendall, M. N. M. van Lieshout (Eds.) *Stochastic geometry, likelihood and computation*. Boca Raton: CRC Press/Chapman and Hall
- Winkler, G. (2003). *Image analysis, random fields and Markov chain Monte Carlo methods*, A mathematical introduction, 2nd edn, *Applications of Mathematics, Stochastic Modelling and Applied Probability* vol. 27. Berlin Heidelberg New York: Springer