Scandinavian Journal of Statistics, Vol. 37: 264–285, 2010
doi: 10.1111/j.1467-9469.2009.00678.x

Disagreement Loop and Path Creation/ Annihilation Algorithms for Binary Planar Markov Fields with Applications to Image Segmentation

TOMASZ SCHREIBER

Faculty of Mathematics & Computer Science, Nicolaus Copernicus University

MARIE-COLETTE VAN LIESHOUT

CWI Amsterdam/Eindhoven University of Technology

ABSTRACT. We introduce a class of Gibbs–Markov random fields built on regular tessellations that can be understood as discrete counterparts of Arak–Surgailis polygonal fields. We focus first on consistent polygonal fields, for which we show consistency, Markovianity and solvability by means of dynamic representations. Next, we develop disagreement loop as well as path creation and annihilation dynamics for their general Gibbsian modifications, which cover most lattice-based Gibbs–Markov random fields subject to certain mild conditions. Applications to foreground–background image segmentation problems are discussed.

Key words: consistent polygonal field, disagreement loop, dynamic representation, image segmentation, Markov random field, path creation and annihilation

1. Introduction

A classical problem in image analysis is that of segmenting the data into relatively homogeneous areas (see, e.g. Rosenfeld & Kak, 1982). It is often the first step in further analysis. Given the centrality of the problem, it is not surprising that a myriad of segmentation methods have been proposed, both deterministic and stochastic in nature. Indeed, the first seminal papers in statistical image analysis (Geman & Geman, 1984; Besag, 1986) concerned this problem. The authors proposed using *Markov random fields* – sometimes also referred to as *Gibbs distributions or fields* – to favour spatially coherent image partitions over more noisy ones. More precisely, a graph is formed by taking the pixel lattice as vertices, joining nearby pixels by an edge between them, and assigning high probability to images in which pixels sharing an edge have similar values. A more recent account of this approach can be found in the volume edited by Chellapa & Jain (1993), or the textbooks by Gimel'farb (1999) and Winkler (2003).

The approach described here uses models that operate on the pixel level. Alternative methods focus on the partition of the image that is the outcome of a segmentation. Green (1995) and Møller & Skare (2001) propose Voronoi-based models, and Nicholls (1998) suggests triangulations. In a previous work, Kluszczyński *et al.* (2005, 2007) advocated the use of *polygonal field models* (introduced in a range of papers by Arak *et al.*, 1989, 1991, 1993), an idea entertained by Clifford & Middleton (1989). In contrast to Markov random fields, the coloured Arak & Surgailis fields take as starting point not the pixel lattice but the Poisson line process. Any realization of this process induces a tessellation, which is then coloured. The boundaries between different-coloured regions form the edges of a graph. Arak & Surgailis (1989) showed that a careful choice of Hamiltonian for a Gibbs field on the set of

(admissible) coloured graphs gives rise to consistent polygonal fields that enjoy remarkable properties which make them eminently suitable for simulation. Building on the concept of *disagreement loops* (Schreiber, 2005), Monte Carlo algorithms were developed that allowed for more global updates than the local ones proposed by Clifford & Nicholls (1994). It was found that the method was robust with respect to noise and able to capture the topology and large- or medium-sized image components well and rapidly; fine details may be lost, though, see Kluszczyński *et al.* (2005, 2007).

The purpose of the present article is to introduce a class of Gibbs–Markov random fields that can be understood as discrete counterparts of polygonal fields. Our construction is two-staged: first a collection of lines inducing a tessellation of the image is fixed, then a polygonal field is constructed on this tessellation, that is to say, the edges of the field consist of (possibly multiple) segments of the tessellation and the vertices of the field are a subset of the collection of nodes of the tessellation. The generic examples of the tessellation-generating collection of lines are realizations of the Poisson line process as well as the line sets corresponding to the regular planar lattice, but it is important to stress that these are not the only valid choices. One may for instance pre-process the image to extract the lines across which the (absolute) gradient flux is particularly high. Not every line is equally desired to turn up in the final segmentation. This is taken into account by ascribing a likelihood parameter to each line, reflecting the potential utility of the line in segmentation. For example, lines in dense regions may be down-weighted to give sparse regions a fair chance; alternatively, after pre-processing, the image gradient may guide the choice of weight. As in our previous work, Kluszczyński *et al.* (2005, 2007), the regions of the tessellation are coloured so that the polygonal boundaries of the field coincide with the interfaces separating different colours. Thus, the edges of the field are the maximal linear segments that form the boundaries between different coloured regions, whereas the field vertices are the intersection points between these boundary segments. Note that each edge of the field may contain many tessellation nodes, that is, intersection points of the tessellation lines.

The analogy with continuum polygonal Markov fields is exploited to define Hamiltonians that are such that desirable properties of these processes (consistency, Markovianity, explicit expressions for the partition function) can be carried over to the discrete context. Moreover, the analogy gives rise to new attractive sampling schemes complementing the usual local Gibbs and Metropolis methods employed for Gibbs fields on finite graphs.

The plan of this article is as follows. In section 2, we construct a family of admissible coloured polygonal configurations built on regular tessellations, and introduce the concept of a discrete polygonal field. The special class of consistent polygonal fields is treated in detail in section 3 with emphasis on its dynamic representation. Invariant birth and death process dynamics for such consistent polygonal fields are derived in section 4 exploiting the notion of a disagreement loop. A simple modification for general polygonal fields is the topic of section 5. More general dynamics are introduced in section 6 which lead naturally to the path creation and annihilation dynamics that form the topic of section 7. In section 8, the image analysis task of foreground–background separation is recast as a statistical inference problem for a discrete polygonal field model. Examples are presented in section 9. We conclude with a critical discussion of our approach and indicate some topics that merit further research.

2. Polygonal fields on regular linear tessellations

By a *regular linear tessellation* of the plane we shall understand a countable family $\mathcal{T}$ of straight lines in $\mathbb{R}^2$ such that no three lines of $\mathcal{T}$ intersect at one point and such that a bounded subset of $\mathbb{R}^2$ is hit by at most a finite number of lines from $\mathcal{T}$, that is to say, $\mathcal{T}$ is

locally finite. Even though we will admit $\mathcal{T}$ random in the sequel, it is assumed to be deterministic in this section, an assumption that does not lead to any loss of generality because in case of random $\mathcal{T}$ the construction below can be verbatim repeated realization-wise.

For a bounded open convex set $D \subseteq \mathbb{R}^2$ the tessellation $\mathcal{T}$ induces a partition of D into a finite collection $D_{\mathcal{T}}$ of regions of polygonal shapes, possibly chopped off by the boundary. Next, we shall always assume that ∂D contains no nodes of $\mathcal{T}$, defined here as intersection points of lines from $\mathcal{T}$. We shall also require that the intersection of each $l \in \mathcal{T}$ with ∂D consists of exactly two points, that is, there are no segments of ∂D along lines of $\mathcal{T}$. Consider the set $\hat{\Gamma}_D(\mathcal{T})$ of all possible colourings of the regions in $D_{\mathcal{T}}$ into *black* and *white* enjoying the additional property that two regions of the same colour can share a node only if either they share a segment or there is a third region of the same colour sharing segments with both of them. In other words, there are no regions of the same colour connected only by corners. The family $\Gamma_D(\mathcal{T})$ of *admissible polygonal configurations* in D built on $\mathcal{T}$ is defined to consist of all planar graphs γ in $\bar{D} := D \cup \partial D$ arising as interfaces between black and white regions of colourings $\hat{\gamma} \in \hat{\Gamma}_D(\mathcal{T})$. Note that in the sequel we shall consistently use the $\hat{}$ notation for coloured elements of $\hat{\Gamma}_D(\mathcal{T})$ whereas omitting the $\hat{}$ will stand for the corresponding *colourless contour configuration* in $\Gamma_D(\mathcal{T})$. Observe that the family $\Gamma_D(\mathcal{T})$ could equivalently be defined to consist of all planar graphs γ in $D \cup \partial D$ such that

- all edges of γ lie on the lines of $\mathcal{T}$,
- all interior vertices of γ, that is, those lying in D, are of degree 2,
- all boundary vertices of γ, that is, those lying on ∂D, are of degree 1,

see also Arak & Surgailis (1989). In other words, the elements of $\Gamma_D(\mathcal{T})$ are collections of disjoint polygonal contours built on $\mathcal{T}$, possibly nested and possibly chopped off by the boundary. Note that by an edge of γ we mean a *maximal* union of connected co-linear tessellation segments in γ, likewise by a vertex of γ we mean a point where two edges of γ meet (interior vertex) or where an edge of γ meets the boundary ∂D (boundary vertex). Thus, the nodes of $\mathcal{T}$ lying in the interior of edges of γ are not considered as vertices of γ, likewise the segments of $\mathcal{T}$ which are not maximal in γ are not edges of γ. To avoid possible ambiguities in the sequel, we shall always use the notions of *vertices* and *edges* in the context of the polygonal configurations built on $\mathcal{T}$, whereas the respective terms *nodes* and *segments* will be reserved for $\mathcal{T}$. When discussing the relations between polygonal fields and general Gibbs fields next, we will also need a notation for the state space of the latter, which is $\hat{\Omega}_D(\mathcal{T})$, standing for the collection of all possible black–white colourings of $D_{\mathcal{T}}$ without the additional requirement of having no *corner-only connections*. The corresponding family $\Omega_D(\mathcal{T})$ of *colour-blind* contour collections is a superset of $\Gamma_D(\mathcal{T})$, as it also admits interior vertices of degree 4.

A reader aware of the original Arak & Surgailis (1989) continuum formalism might wonder at this point why we admit here co-linear edges of graphs $\gamma \in \Gamma_D(\mathcal{T})$, which corresponds to situations where a linear section of $\gamma \in \Gamma_D(\mathcal{T})$ along $l \in \mathcal{T}$ is a union of disjoint edges, forbidden in the Arak and Surgailis' set-up. Paradoxically, the answer is that this is aimed at ensuring the Markov property of the field, which is the same reason for which such situations were excluded in the continuum setting. A solution to this apparent paradox lies in crucial differences between the continuum and discrete set-ups.

- In the continuum Arak & Surgailis (1989) set-up, forbidding multiple disjoint edges along the same line does not introduce any long-range dependencies because even though the edges far away should *avoid the lines already used*, the measure of these forbidden lines is anyway zero and thus the exclusion is not effectively felt by the field. In contrast, any co-linear edge would have to arise as a result of a mechanism introducing

long-range dependencies because the local field-generating rules proposed by Arak & Surgailis yield co-linearity with probability zero.

- In contrast to this, in our discrete setting exclusion of co-linear edges would lead to a situation where the *forbidden lines* would be felt in arbitrary distance from the corresponding edges, thus obviating the Markovianity of the field in any reasonable sense. However, in the discrete set-up there do exist local mechanisms producing co-linear edges with positive probabilities.

In particular, it should be emphasized in this context that the tessellation-based polygonal fields constructed in this article, although analogous to the Arak–Surgailis fields, are not versions thereof conditioned on having their edges on the lines of $\mathcal{T}$ – in fact, such a conditioning would yield a non-Markovian field, owing to the presence of *forbidden lines* as discussed before. Note that these remarks are only meant to compare our setting with the original Arak and Surgailis continuum framework, and will not be referred to in the sequel of this article.

Assume now that fixed activity/probability parameters $\pi_l \in (0,1)$ are ascribed to the straight lines $l \in \mathcal{T}$. The term *probability parameter* is used here to reflect the interpretation of π_ls in terms of probabilities of creating an edge along l, as made precise next, thus in particular no requirement of π_ls summing up to one is imposed. For a function $\mathcal{H}_D : \hat{\Gamma}_D(\mathcal{T}) \mapsto \mathbb{R} \cup \{+\infty\}$ the (*discrete*) *polygonal field* $\hat{\mathcal{A}}_{\mathcal{H}_D}$ with Hamiltonian $\mathcal{H}_D$ is defined to be the random element in $\hat{\Gamma}_D(\mathcal{T})$ such that

$$\mathbb{P}(\hat{\mathcal{A}}_{\mathcal{H}_D} = \hat{\gamma}) = \frac{\exp(-\mathcal{H}_D(\hat{\gamma})) \prod_{e \in E(\gamma)} \pi_{l[e]}}{\mathcal{Z}[\mathcal{H}_D]}, \tag{1}$$

where $E(\gamma)$ stands for the collection of edges of γ considered here to be open, that is, not to contain their vertices, for formal convenience next, whereas $l[e] \in \mathcal{T}$ is the straight line containing e, and

$$\mathcal{Z}[\mathcal{H}_D] := \sum_{\hat{\theta} \in \hat{\Gamma}_D(\mathcal{T})} \exp(-\mathcal{H}_D(\hat{\theta})) \prod_{e \in E(\theta)} \pi_{l[e]} \tag{2}$$

is the corresponding partition function. In other words, the probability of seeing some $\hat{\gamma} \in \hat{\Gamma}_D(\mathcal{T})$ as the realization of $\hat{\mathcal{A}}_{\mathcal{H}_D}$ is proportional to the Boltzmann factor $\exp(-\mathcal{H}_D(\hat{\gamma}))$ times the product of probabilistic costs $\pi_{l[e]}$ of creating its edges. In particular, under fixed Hamiltonian, putting π_ls low for one family of lines, say $\mathcal{L}_1$, and setting π_ls high over some other family $\mathcal{L}_2$ of lines, amounts to promoting field edges along the lines of $\mathcal{L}_2$ while penalizing edge creation along lines from $\mathcal{L}_1$.

Recalling that the Gibbs field $\hat{\mathcal{G}}_{\Psi_D}$ on $\hat{\Omega}_D(\mathcal{T})$ with Hamiltonian $\Psi_D : \hat{\Omega}_D(\mathcal{T}) \mapsto \mathbb{R} \cup \{+\infty\}$ is given by

$$\mathbb{P}\left(\hat{\mathcal{G}}_{\Psi_D} = \hat{\gamma}\right) = \frac{\exp(-\Psi_D(\hat{\gamma}))}{\sum_{\hat{\theta} \in \hat{\Omega}_D(\mathcal{T})} \exp(-\Psi_D(\hat{\theta}))}, \tag{3}$$

we easily see that the polygonal field $\hat{\mathcal{A}}_{\mathcal{H}_D}$ coincides in law with $\hat{\mathcal{G}}_{\Psi_D}$ regarded as a $\hat{\Gamma}_D(\mathcal{T})$-valued random element for

$$\Psi_D(\hat{\gamma}) := \begin{cases} \mathcal{H}_D(\hat{\gamma}) - \sum_{e \in E(\gamma)} \log \pi_{l[e]}, & \hat{\gamma} \in \hat{\Gamma}_D(\mathcal{T}), \\ +\infty, & \text{otherwise.} \end{cases} \tag{4}$$

Thus, all polygonal fields on $\mathcal{T}$ are Gibbs fields and a vast class of Gibbs fields, namely those concentrated on $\hat{\Gamma}_D(\mathcal{T})$, admit a representation as polygonal fields. In spite of this apparent redundancy in definitions there are good reasons for considering the notion of a discrete polygonal field though, one of them being that, unlike (3), the definition (1) admits a natural continuum version and in fact it is the continuum set-up where it has originally arisen, see

Arak (1982), Arak & Surgailis (1989) and Arak *et al.* (1993). It should also be emphasized that for suitable natural choices of Hamiltonian in (1) the resulting field exhibits striking properties, as will be discussed in section 3. There is another important reason for introducing (1), however, which is crucial for the purposes of the present article: there are simulation techniques available for continuum polygonal fields, see Schreiber (2005) and Kluszczyński *et al.* (2007), whose discrete adaptations can be used to provide new attractive simulation algorithms for discrete Gibbs fields. In the context of image segmentation, the Hamiltonian will include terms that quantify how well a coloured contour collection describes the data (see section 8). To proceed with the presentation of these ideas, we shall begin with a discussion of the so-called *consistent polygonal fields* and their *dynamic representations* in section 3.

To complete the present section, we remark that formula (1) can be regarded as the discrete equivalent of the *line-based* representation for polygonal fields as developed in Arak & Surgailis (1989). It is natural to ask whether an alternative *point-based* representation in the spirit of Arak *et al.* (1993) is available as well. This is indeed the case, but, unlike in the continuum set-up, this alternative representation arises by a simple re-arrangement of the line-indexed product $\prod_{e\in E(\gamma)} \pi_{l[e]}$ as $(\prod_{v(l_1,l_2)\in V(\gamma)\cap D} \sqrt{\pi_{l_1}\pi_{l_2}})(\prod_{v(l)\in V(\gamma)\cap \partial D} \sqrt{\pi_l})$ with $v(l_1,l_2)$ standing for vertices of γ falling into D and arising at the intersection of lines $l_1, l_2 \in \mathcal{T}$ and with $v(l)$ ranging through boundary vertices of γ in ∂D lying on $l \in \mathcal{T}$. Thus, in sharp contrast to the continuum case, the point-based representation here does not seem to offer a relevant alternative to the line-based setting.

3. Consistent polygonal fields on regular tessellations

In the seminal papers Arak (1982) and Arak & Surgailis (1989), it was observed that for some particular choices of the Hamiltonian the corresponding polygonal fields enjoy remarkable properties which make them very well suited for simulation – these special processes are the so-called *consistent polygonal fields*. In the present section, we shall adopt to the discrete case of polygonal fields on regular tessellations the argument of Arak and Surgailis originally developed mainly in the continuum setting (with some exceptions though, see, e.g. model D in Surgailis, 1991).

To proceed, consider the Hamiltonian

$$\Phi_D(\hat{\gamma}) := -\sum_{e\in E(\gamma)} \sum_{l\in\mathcal{T}, l\not\sim e} \log(1-\pi_l) + \sum_{n(l_1,l_2)\in\gamma} \log(1-\pi_{l_1}\pi_{l_2}), \quad \hat{\gamma}\in\hat{\Gamma}_D(\mathcal{T}), \tag{5}$$

with $n(l_1,l_2)\in\gamma$ ranging through all nodes of the tessellation $\mathcal{T}$ arising as intersection points of $l_1, l_2 \in \mathcal{T}$ and lying on γ, that is to say, either lying on the edges of γ or coinciding with one of its vertices. Here and in the sequel, $l \not\sim e$ means that the line l intersects e but is not co-linear with it. The polygonal field $\hat{\mathcal{A}}_{\Phi_D}$ is a *consistent polygonal field* in the sense made precise by theorem 1.

Theorem 1
The polygonal field $\hat{\mathcal{A}}_{\Phi_D}$ with Hamiltonian Φ_D enjoys the following properties:

Consistency: *For bounded open convex $D' \subseteq D \subseteq \mathbb{R}^2$, the field $\hat{\mathcal{A}}_{\Phi_D} \cap D'$ coincides in law with $\hat{\mathcal{A}}_{\Phi_{D'}}$. By increasing $D \uparrow \mathbb{R}^2$ this allows us to construct the whole plane extension of the process $\hat{\mathcal{A}}_\Phi$ such that $\hat{\mathcal{A}}_{\Phi_D}$ coincides in law with $\hat{\mathcal{A}}_\Phi \cap D$ for all bounded open convex $D \subseteq \mathbb{R}^2$.*

Linear sections: *For a straight line l containing no nodes of $\mathcal{T}$, the intersection points and intersection directions of l with the edges of the polygonal field $\hat{\mathcal{A}}_\Phi$ coincide in distribution with*

the intersection points and directions of l with the line field $\Lambda_{\mathcal{T}}$ defined to be the random sub-collection of $\mathcal{T}$ where each straight line $l^* \in \mathcal{T}$ is chosen to belong to $\Lambda_{\mathcal{T}}$ with probability $\pi_{l^*}/(1+\pi_{l^*})$ and rejected otherwise, and all these choices are made independently.

Solvability: *An explicit formula is available for the partition function:*

$$\mathcal{Z}[\Phi_D] = 2\left(\prod_{n(l_1,l_2)\in D}(1-\pi_{l_1}\pi_{l_2})\right)^{-1}\left(\prod_{l\in\mathcal{T},\, l\cap D\neq\emptyset}\frac{1}{(1+\pi_l)}\right)^{-1}. \tag{6}$$

Markov property: *For a smooth closed curve $\theta \subset \mathbb{R}^2$ containing no nodes of $\mathcal{T}$, the conditional distribution of $\hat{\mathcal{A}}_\Phi$ in the interior of θ depends on the configuration outside θ only through the intersection points and intersection directions of θ with the edges of the polygonal field and through the colouring of the field along θ.*

Proof. The proof of theorem 1 is based on the so-called *dynamic representation* for consistent polygonal fields, being a discrete version of and constructed in full analogy with the corresponding representation in sections 4 and 5 in Arak & Surgailis (1989). The idea underlying this construction is to represent the considered polygonal field in terms of the equilibrium evolution of a one-dimensional particle system tracing the polygonal boundaries of the field in two-dimensional time-space. To this end, we interpret the open convex domain D as a set of *time-space* points $(t, y) \in D$, with $t \in \mathbb{R}$ referred to as the *time* coordinate and with $y \in \mathbb{R}$ standing for the *spatial* coordinate of a particle at the time t. In this language, a straight line segment in D stands for a piece of the time-space trajectory of a freely moving particle. For a straight line l non-parallel to the spatial axis and crossing the domain D we define in the obvious way its entry point to D, $\mathrm{in}(l, D) \in \partial D$, and its exit point, $\mathrm{out}(l, D) \in \partial D$. Without loss of generality we assume that no line of $\mathcal{T}$ is parallel to the spatial axis, possibly rotating the coordinate system if this is not the case.

We choose the time-space birth sites for the new particles by independently placing a birth site

- at each node $n(l_1, l_2)$ of the tessellation $\mathcal{T}$ falling into D, with probability $\pi_{l_1}\pi_{l_2}$ (interior birth site),
- at each entry point $\mathrm{in}(l, D)$ of lines $l \in \mathcal{T}$ into D, with probability $\pi_l/(1+\pi_l)$ (boundary birth site).

Each interior birth site $n(l_1, l_2)$ emits two particles moving with initial velocities such that the initial segments of their trajectories lie on the lines l_1 and l_2 of the tessellation going out from the birth site, unless another particle previously born hits the site in which case the birth does not occur. Note that this prevents creation of degree 3 and degree 4 vertices in the resulting graph. Each boundary birth site $\mathrm{in}(l, D)$ emits a single particle moving with the initial velocity such that the initial segment of its trajectory lies on l (no precaution similar to the one for interior birth sites before is present because boundary birth sites cannot be hit by previously born particles). All the particles evolve independently in time according to the following rules:

(E1) Between the critical moments listed next, each particle moves with constant velocity so that $\mathrm{d}y = v\,\mathrm{d}t$ with v standing for the actual velocity.
(E2) When a particle touches the boundary ∂D, it dies.
(E3) In case of a collision of two particles (equal spatial coordinates y at some moment t with $(t, y) \in D$), both of them die.
(E4) Whenever a particle moving in time-space along $l_1 \in \mathcal{T}$ reaches a node $n(l_1, l_2)$, it changes its velocity so as to move along l_2 with probability π_{l_2}, and keeps moving along l_1 otherwise.

The claim constituting the core of the proof is that the union of the time-space trajectories traced by the particles of the aforegiven system coincides in distribution with the contour ensemble $\mathcal{A}_{\Phi_D}$ of the polygonal field $\hat{\mathcal{A}}_{\Phi_D}$, whereas the law of the field itself is recovered by picking one of the two possible black–white colourings at random with probability 1/2. To verify this statement we choose some $\hat{\gamma} \in \hat{\Gamma}_D(\mathcal{T})$ and compute the probability that the colour-blind contour ensemble γ is traced by the previous particle system. To this end, we observe the following.

- Each edge $e \in E(\gamma)$ whose initial (lower time coordinate) vertex lies on ∂D yields a factor $\pi_{l[e]}/(1+\pi_{l[e]})$ (boundary birth site) times $\prod_{l \in \mathcal{T}, l \not\sim e}(1-\pi_l)$ (no velocity updates along e).
- Each of the two edges $e_1, e_2 \in E(\gamma)$ stemming from a common interior birth site $n(l_1, l_2)$ yields a factor $\pi_{l_i}, i = 1, 2$ (coming from the birth probability) times $\prod_i \prod_{l \in \mathcal{T}, l \not\sim e_i}(1-\pi_l)$ (no velocity updates along e_i).
- Each of the edges $e \in E(\gamma)$ arising owing to a velocity update of a particle yields a factor $\pi_{l[e]}$ (velocity update probability) times $\prod_{l \in \mathcal{T}, l \not\sim e}(1-\pi_l)$ (no velocity updates along e).
- The absence of birth sites in nodes $n(l_1, l_2)$ of $\mathcal{T}$ in D not belonging to γ yields the factor $\prod_{n(l_1,l_2) \in D \setminus \gamma}(1-\pi_{l_1}\pi_{l_2})$ (note that birth sites are allowed in all points of γ – either they give rise to particles tracing γ or are discarded if hit by a previously born particle).
- The absence of boundary birth sites at those entry points to D of lines of $\mathcal{T}$ which do not give rise to an edge of γ yields the factor $\prod_{l \in \mathcal{T}, l \cap D \neq \emptyset, \text{in}(l,D) \notin \gamma}(1+\pi_l)^{-1}$.

Putting these factors together allows us to evaluate the probability of γ being traced by the particle system to

$$\left(\prod_{e \in E(\gamma)} \pi_{l[e]}\right)\left(\prod_{e \in E(\gamma)} \prod_{l \in \mathcal{T}, l \not\sim e}(1-\pi_l)\right)\left(\prod_{n(l_1,l_2) \in \gamma}(1-\pi_{l_1}\pi_{l_2})\right)^{-1}\left(\prod_{n(l_1,l_2) \in D}(1-\pi_{l_1}\pi_{l_2})\right) \times \left(\prod_{l \in \mathcal{T}, l \cap D \neq \emptyset} \frac{1}{1+\pi_l}\right) = \frac{2\exp(-\Phi_D(\hat{\gamma}))}{\mathcal{Z}[\Phi_D]} \prod_{e \in E(\gamma)} \pi_{l[e]} \tag{7}$$

with $\mathcal{Z}[\Phi_D]$ given by (6). Taking into account that the choice between the two possible colourings of the field is made with probability 1/2, independently of γ, we see that the probability of obtaining $\hat{\gamma}$ as the outcome of the particle system evolution is exactly

$$\frac{\exp(-\Phi_D(\hat{\gamma}))}{\mathcal{Z}[\Phi_D]} \prod_{e \in E(\gamma)} \pi_{l[e]},$$

and hence the resulting polygonal field coincides in law with $\hat{\mathcal{A}}_{\Phi_D}$ as required – this fact will be referred to as the *dynamic representation for the polygonal field* $\hat{\mathcal{A}}_{\Phi_D}$ in the sequel. The **Solvability** statement (6) follows from this as well. The remaining properties follow in full analogy with the corresponding argument in Arak & Surgailis (1989), whence we only provide a brief discussion in the following. First, the **Markov property** stated before is a direct consequence of the Gibbsian definition of $\hat{\mathcal{A}}_{\Phi_D}$, whereas the **Linear sections** statement will follow from the form of the boundary birth mechanism described before as soon as we establish the remaining **Consistency** property. To this end, choose a bounded open convex set $D \subseteq \mathbb{R}^2$ and a straight line l intersecting D, and define D' to be the set of points of D lying to the left of l (lower time coordinates). Clearly then, from the dynamic representation we conclude the **Consistency** statement for so chosen D and D'. Noting that the dynamic representation is equally available upon rotating the time-space coordinate system, we see that the **Consistency** holds as well upon cutting off the part of the set D lying to the left of l. This means however that the consistency holds upon cutting off pieces of the original set with arbitrary straight lines – a repetitive use of this procedure and a possible passage to

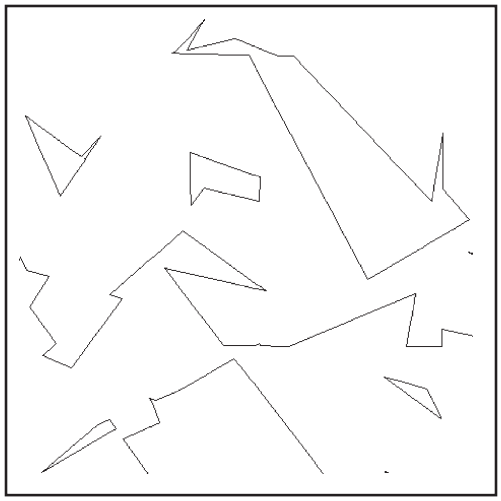

Fig. 1. A typical realization of the consistent field on a Poisson line tessellation.

the limit allows us to *carve* from D any arbitrary convex subset. This proves the **Consistency** claim and completes the proof of theorem 1.

We conclude this section by remarking that, as seen from the **Linear sections** statement of theorem 1, the probability that a given edge e is present in the field $\mathcal{A}_{\Phi_D}$ is $\pi_{l[e]}/(1+\pi_{l[e]})$, which does not exceed 1/2. This may seem bizarre at the first look, as increasing $\pi_{l[e]}$ towards 1 would suggest maximizing the probability of edge occurrence at e. Upon a second look, though, it is readily seen that this fact is because of the exclusion of X-shaped nodes in the definition of $\Gamma_D(\mathcal{T})$, which prevents the maximum obtainable edge density from exceeding 1/2.

A typical realization of the consistent polygonal field on a homogeneous Poisson line tessellation is shown in Fig. 1.

4. Disagreement loop birth and death dynamics for consistent polygonal fields

A crucial concept next will be that of a *disagreement loop*, borrowed from Schreiber (2005), section 2.1. This arises from the dynamic construction of the polygonal fields as provided by the evolution rules **(E1–4)** and the corresponding birth rules specified in the proof of theorem 1 before.

Suppose that we observe a particular realization $\gamma \in \Gamma_D(\mathcal{T})$ of the colourless contour ensemble $\mathcal{A}_{\Phi_D}$ and that we modify the configuration by adding an extra birth site x_0 to the existing collection of birth sites for γ arising in the dynamic representation, while keeping the evolution rules **(E1–4)** for all the particles, including the two newly added ones if x_0 is an interior birth site and the single newly added one if x_0 is a boundary birth site. Denote the resulting new random (colourless) polygonal configuration by $\gamma \oplus x_0$. A crucial observation is that, under appropriate coupling of the dynamics of newly added particles with that of the previously existing ones along their sub-trajectories annihilated owing to the creation of new particles, for an interior birth site x_0 the symmetric difference $\gamma \triangle [\gamma \oplus x_0]$ is almost surely a single loop (a closed polygonal curve), possibly self-intersecting and possibly chopped off by the boundary. Likewise, under the same coupling, the symmetric difference $\gamma \triangle [\gamma \oplus x_0]$ for boundary x_0 is almost surely a single polygonal path with no self-intersections. We describe

this coupling next. It should be noted that for formal convenience we assume that a polygonal configuration γ comes together with the knowledge of all attempted birth sites discarded during its creation.

If x_0 happens to lie on γ, the new birth attempt gets immediately discarded, yielding $\gamma \triangle [\gamma \oplus x_0] = \emptyset$. Apart from this degenerate case, the leftmost point of the loop $\gamma \triangle [\gamma \oplus x_0]$ is of course x_0. Each of the two *new* particles p_1, p_2 emitted from x_0 move independently, according to **(E1–4)**, each giving rise to a *disagreement path*. The initial segments of such a disagreement path correspond to the movement of a particle, say p_1, before its annihilation in the first collision. If this is a collision with the boundary, the disagreement path gets chopped off and terminates there. If this is a collision with a segment of the original configuration γ corresponding to a certain *old* particle p_3, the *new* particle p_1 dies but the disagreement path continues along the part of the trajectory of p_3 which is contained in γ but not in $\gamma \oplus x_0$. A particular case may happen when the collision occurs with a birth site in γ, where one of the two particles born (denote it by p_3) gets annihilated, whereas the evolution of the second one is re-used under our coupling to yield the extension of the evolution of p_1 – we do so instead of annihilating both particles emitted from the birth site and building the extension of the evolution of p_1 anew. The choice of the annihilated and surviving particles upon such a collision at a node $n(l_1, l_2)$ with the current direction of p_1 being l_1 is made as follows: with probability $(1 - \pi_{l_2})$ we annihilate the particle moving along l_2 (no velocity update for p_1) and with probability π_{l_2} we annihilate the particle moving along l_1 (p_1 undergoes a velocity update). Clearly, this is consistent with the dynamic representation. Thus, we note that under the so-constructed coupling p_1 always moves according to the usual evolution rules **(E1–4)** because the particles emitted by the hit birth site did so. At some further moment p_3 dies itself in γ, touching the boundary or killing another particle p_4 in γ. In the second case, however, this collision only happens for γ and not for $\gamma \oplus x_0$, so the particle p_4 survives (for some time) in $\gamma \oplus x_0$ yielding a further connected portion of the disagreement path for p_1, which is contained in $\gamma \oplus x_0$ but not in γ. Likewise, it may happen that p_3 reaches a node $n(l_1, l_2)$ where a birth attempt was made in γ but was discarded because of the presence of p_3. Now that p_3 is absent in $\gamma \oplus x_0$, this birth does occur for $\gamma \oplus x_0$. We require in our coupling that one of the emitted particles follow the remaining trajectory of p_3 in γ, thus re-using for the new particle the random choices made before for p_3. Clearly, the second particle emitted by the birth site denoted by p_4 as before, adds one further sub-portion of the disagreement loop in $[\gamma \oplus x_0] \setminus \gamma$. Again, the particle evolving according to the previous trajectory of p_3 moves in accordance with all the rules of the dynamic representation because so did p_3.

A recursive continuation of this construction shows that the disagreement path initiated by p_1 consists alternately of connected polygonal sub-paths contained in $[\gamma \oplus x_0] \setminus \gamma$ (call these *positive* parts) and in $\gamma \setminus [\gamma \oplus x_0]$ (call these *negative* parts). Note that this disagreement path is self-avoiding and, in fact, it can be represented as the graph of some piecewise linear function $t \mapsto y(t)$. Clearly, the same applies for the disagreement path initiated by the second initial particle p_2. An important observation is that whenever two *positive* or two *negative* parts of the two disagreement paths hit each other, both disagreement paths may die at this point and the disagreement loop may close (as opposed to intersections of segments of distinct signs which do not have this effect). Obviously, if the disagreement loop does not close in this way, it gets eventually chopped off by the boundary. Note that upon the intersection of two *positive* or *negative* sub-paths at a node v, instead of getting killed, a disagreement loop may also continue owing to the alteration of the status of a birth site, should it occur at v. Indeed, for negative sub-paths a birth site previously discarded may be reactivated, whereas for positive sub-paths a birth site may be inactivated.

We shall write $\Delta^{\oplus}[x_0;\gamma]=\gamma\triangle[\gamma\oplus x_0]$ to denote the (random) disagreement loop constructed before. It remains to consider the case of x_0 being a boundary birth site, which is much simpler because there is only one particle emitted and so, under our coupling, $\Delta^{\oplus}[x_0;\gamma]=\gamma\triangle[\gamma\oplus x_0]$ is easily seen to be a single self-avoiding polygonal path eventually chopped off by the boundary. We abuse the language and call such $\Delta^{\oplus}[x_0;\gamma]$ a (degenerate) disagreement loop as well.

Likewise, a disagreement loop arises if we *remove* one birth site x_0 from the collection of birth sites of an admissible polygonal configuration $\gamma\in\Gamma_D(\mathcal{T})$, while keeping the evolution rules for all the remaining particles. We write $\gamma\ominus x_0$ for the configuration obtained from γ by removing x_0 from the list of the birth sites, while the resulting random disagreement loop is denoted by $\Delta^{\ominus}[x_0;\gamma]$, so that $\Delta^{\ominus}[x_0;\gamma]=\gamma\triangle[\gamma\ominus x_0]$. Note that again an empty disagreement loop may occur, should we annihilate a birth site where the birth attempt was discarded owing to the presence of previously created particles. We refer the reader to Schreiber (2005) for further formal details of the disagreement loop concept.

With this terminology we are in a position to describe random dynamics on the coloured configuration space $\hat{\Gamma}_D(\mathcal{T})$ which leave invariant the law of the polygonal field $\hat{\mathcal{A}}_{\Phi_D}$. Particular care is needed, however, to distinguish between the notion of time considered in the dynamic representation of the polygonal field as well as throughout the construction of the disagreement loops above, and the notion of time to be introduced for the random dynamics on $\hat{\Gamma}_D(\mathcal{T})$ 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, whereas the latter will be called the *simulation time* (s-time for short) and will be denoted by s in the sequel.

Consider the following pure jump birth and death type continuous time Markovian **(DL)** dynamics on $\hat{\Gamma}_D(\mathcal{T})$.

> **DL:birth**: At each $x:=n(l_1,l_2)\in D, l_1,l_2\in\mathcal{T}$, with intensity $\pi_{l_1}\pi_{l_2}\,\mathrm{d}s$ set $\gamma_{s+\mathrm{d}s}:=\gamma_s\oplus x$ (interior births), then construct $\hat{\gamma}_{s+\mathrm{d}s}$ by randomly choosing, with probability 1/2, either of the two possible colourings for $\gamma_{s+\mathrm{d}s}$. Proceed likewise at each $x:=l\cap\partial D, l\in\mathcal{T}$, with intensity $\pi_l/(1+\pi_l)\,\mathrm{d}s$ (boundary births).
>
> **DL:death**: For each interior and boundary birth site x in γ_s, with intensity 1 set $\gamma_{s+\mathrm{d}s}:=\gamma_s\ominus x$, then construct $\hat{\gamma}_{s+\mathrm{d}s}$ by randomly choosing, with probability 1/2, either of the two possible colourings for $\gamma_{s+\mathrm{d}s}$.

If none of these updates occurs, we keep $\hat{\gamma}_{s+\mathrm{d}s}=\hat{\gamma}_s$. It is convenient to perceive these dynamics in terms of generating random disagreement loops λ and setting $\gamma_{s+\mathrm{d}s}:=\gamma_s\triangle\lambda$, with the loops of the type $\Delta^{\oplus}[\cdot;\cdot]$ corresponding to the rule **DL:birth** and $\Delta^{\ominus}[\cdot;\cdot]$ to the rule **DL:death**.

As a consequence of the dynamic representation developed in the proof of theorem 1 we obtain proposition 1.

Proposition 1

The distribution of the consistent polygonal field $\hat{\mathcal{A}}_{\Phi_D}$ is the unique invariant law of the dynamics given by **DL:birth** *and* **DL:death**. *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 variational distance to the law of $\hat{\mathcal{A}}_{\Phi_D}$ as $s\to\infty$.*

Proof. To establish the invariance, we note first that the **DL** dynamics is easily seen to preserve the Bernoulli law imposed by the dynamic representation on the collection of birth sites of the process (including the birth attempt sites discarded owing to the incidence with previously born particles). Moreover, the **DL** dynamics is explicitly constructed so as to

ensure that if the original configuration is traced by particles evolving according to the rules of the dynamic construction and emitted from the given collection of birth sites, then so is the updated configuration, with the accordingly updated collection of birth sites. This ensures the required invariance. Alternatively, a direct detailed balance check can also be made, for instance in analogy to section 6 in Schreiber (2008), see especially subsection 6.4 there, where direct balance calculations are made in a very general continuum set-up.

The reversibility of the dynamics follows from the obvious reversibility of the birth site birth and death process. To see this, assign to each birth site (including the discarded birth attempts) the (random) sequence of velocity update choices to be made in the course of evolution of the particles it emits (would emit), which corresponds to the full knowledge of the (unlimited) particle evolution. This makes birth sites into *birth packages*, algorithmically representable, for example, by assigning to each birth site the corresponding seed for random number generation. Thus, we end up with a random collection of birth packages containing all the randomness of the field, that is to say, fully determining the resulting field now obtainable by a deterministic procedure. For definiteness, we assume that if a particle passes on its way a birth site, which thus becomes inactivated, starting from the next evolution step after this event the particle evolves further according to the birth package of the just inactivated birth site rather than according to its own. Now, in this context it is easy to check that the reverse move to adding a birth package at x_0 is removing the same birth package at x_0 and vice versa, that is to say $[\gamma \oplus x_0] \ominus x_0 \equiv \gamma$ and $[\gamma \ominus x_0] \oplus x_0 \equiv \gamma$ with x_0 carrying always the same birth package. This ensures the reversibility under the aforegiven *birth package* interpretation. To get the required reversibility for polygonal configurations not containing the full knowledge of entire birth packages (unnecessarily determining the unlimited extension of the trajectories of annihilated particles as if they were to survive forever) it is now enough to integrate out the spurious components of birth packages.

The uniqueness and convergence statements in proposition 1 require a short justification as well. They both follow from the observation that, in finite volume, regardless of the initial state, the process $\hat{\gamma}_s$ spends a non-null fraction of time in the state 'black' (no contours, the whole domain D coloured black). Indeed, this observation allows us to conclude the required uniqueness and convergence by a standard coupling argument, for example, along the lines of the proof of theorem 1.2 in Liggett (1985).

5. Disagreement loop birth and death dynamics for general polygonal fields

Take now a general polygonal field $\hat{\mathcal{A}}_{\mathcal{H}_D+\Phi_D}$ with a Hamiltonian $\mathcal{H}_D+\Phi_D : \hat{\Gamma}_D(\mathcal{T}) \to \mathbb{R}$ as in (1). Consider the following modification of the basic **(DL)** dynamics constructed previously in section 4:

> **DL[H]:birth**: At each $x := n(l_1, l_2) \in D, l_1, l_2 \in \mathcal{T}$, with intensity $\pi_{l_1}\pi_{l_2}\,\mathrm{d}s$, set $\delta_{s+\mathrm{d}s} := \gamma_s \oplus x$, whereupon construct $\hat{\delta}$ by choosing with probability 1/2 one of the two possible colourings for δ. Then, with probability $\min(1, \exp[\mathcal{H}_D(\hat{\gamma}_s) - \mathcal{H}_D(\hat{\delta})])$, put $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\delta}$, otherwise keep $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\gamma}_s$. Proceed likewise at each $x := l \cap \partial D, l \in \mathcal{T}$, with intensity $\pi_l/(1+\pi_l)\,\mathrm{d}s$.
> **DL[H]:death**: for each interior and boundary birth site x in γ_s, with intensity $1 \cdot \mathrm{d}s$, set $\delta := \gamma_s \ominus x$, whereupon construct $\hat{\delta}$ by choosing with probability 1/2 one of the two possible colourings for δ. Then, with probability $\min(1, \exp[\mathcal{H}_D(\hat{\gamma}_s) - \mathcal{H}_D(\hat{\delta})])$, put $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\delta}$, otherwise keep $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\gamma}_s$.

In other words, the original dynamics **DL** are used to propose a new configuration $\hat{\delta}$, which is then accepted with probability $\min(1, \exp[\mathcal{H}_D(\hat{\gamma}_s) - \mathcal{H}_D(\hat{\delta})])$, and rejected otherwise. By a

straightforward verification of the detailed balance conditions and an appeal to proposition 1, we obtain the following result, theorem 2.

Theorem 2
The distribution of the polygonal field $\hat{\mathcal{A}}_{\mathcal{H}_D+\Phi_D}$ *given by* (1) *is the unique invariant law of the dynamics described by* **DL[H]:birth** *and* **DL[H]:death**. *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 variational distance to the law of* $\hat{\mathcal{A}}_{\mathcal{H}_D+\Phi_D}$ *as* $s \to \infty$.

6. Generalized dynamic representation for consistent polygonal fields

The dynamic construction of consistent polygonal fields borrowed from Arak & Surgailis (1989) and adapted for tessellation-based fields in the proof of theorem 1 can be regarded as *revealing* increasing portions of the polygonal field in the course of time flow. Under this interpretation, with probability 1 the portion of a polygonal field in a bounded open convex domain D *uncovered* by time t is, upon closure, precisely the closure of its intersection with $D_t = \bar{D} \cap (-\infty, t] \times \mathbb{R}$. The idea underlying our generalized dynamic representation developed in the present section in the following, and constituting a discrete counterpart of the continuum generalized representation in section 4 of Schreiber (2008), is to replace the family D_t by some other time-increasing family of subsets of D, also denoted by D_t in the sequel, eventually covering the whole D, and to try to provide a natural construction of the polygonal field being gradually uncovered on the growing domain D_t in the course of time flow. We shall always assume that D_t is convex, for otherwise we would have to deal with situations where two or more disconnected parts of an edge of the field have been revealed, which leads to unwanted dependencies along the segments connecting these parts. Taking this into account, and having formal convenience in mind, we impose the following natural assumptions on $D_t, t \in [0, 1]$,

- $(D_t)_{t\in[0,1]}$ is an increasing family of compact convex subsets of $\bar{D} = D \cup \partial D$,
- D_0 is a single point x in $\bar{D} = D \cup \partial D$,
- D_1 coincides with $\bar{D}$,
- D_t is continuous in the usual Hausdorff metric,
- For each $l \in \mathcal{T}$ the intersection $l \cap D_{\tau_l}$ consists of exactly one point $\mathbb{A}(l)$, where $\tau_l := \inf\{t \in [0, 1], D_t \cap l \neq \emptyset\}$. Moreover, $\mathbb{A}(l)$ is not a node of $\mathcal{T}$.

The condition requiring that D_0 be a singleton can be easily weakened, in fact it is enough if D_0 be a linear segment, yet this requires certain technical changes in the dynamics below without providing essential generalizations and hence we do not discuss this option here. The point $\mathbb{A}(l)$ will be referred to as the *anchor point* for l; this induces the *anchor* mapping $\mathbb{A} : \mathcal{T} \to \bar{D}$. Note that the so-defined *growing window* D_t introduces a natural notion of time over the field domain $\bar{D}$, where the time mark of a point $x \in \bar{D}$ is given as the first moment where $\bar{D}_t$ hits x. Unlike in the standard dynamic representation discussed in the proof of theorem 1 though, in general there can be no definite and fixed *time flow direction* here as the time flow geometry can change over time. Consider now the following dynamics in continuous time $t \in [0, 1]$, with all updates given by the following rules performed independently of each other and with $\mathrm{d}t$ standing for an infinitesimal time increment.

(GE:Initialise) Begin with an empty field at time 0.
(GE1) Between critical moments listed next, during the time interval $[t, t + \mathrm{d}t]$, the field lines in D_t hitting ∂D_t extend straight to $D_{t+\mathrm{d}t} \setminus D_t$.

(GE2) When a field line hits the boundary ∂D, it stops growing in this direction (recall that ∂D contains no segments along lines of $\mathcal{T}$ and so the intersection of a field line with ∂D consists of at most two points).

(GE3) When two field lines intersect in $D_{t+\mathrm{d}t} \setminus D_t$, they are not extended any further beyond the intersection point (stop growing in the direction marked by the intersection point).

(GE4) Whenever a node $n(l_1, l_2), l_1, l_2 \in \mathcal{T}$, falls into $D_{t+\mathrm{d}t} \setminus D_t$, and l_1 is a current field line hitting ∂D_t, with probability π_{l_2} we update the direction of the line to l_2, extending away from the anchor point $\mathbb{A}(l_2)$. We keep the previous direction along l_1 otherwise. Should several such vertices belong to $D_{t+\mathrm{d}t} \setminus D_t$, the direction updates are performed independently.

(GE:LineBirth) Whenever an anchor point $\mathbb{A}(l), l \in \mathcal{T}$, falls into $D_{t+\mathrm{d}t} \setminus D_t$, with probability $\pi_l/(1+\pi_l)$ a new field line l is born at $\mathbb{A}(l)$, extending in both directions if possible.

(GE:VertexBirth) Whenever a node $n(l_1, l_2), l_1, l_2 \in \mathcal{T}$, falls into $D_{t+\mathrm{d}t} \setminus D_t$, with probability $\pi_{l_1}\pi_{l_2}$ two new field lines l_1 and l_2 are born, each extending in the direction away from its anchor point, unless another field line present at time t hits $n(l_1, l_2)$ in which case the birth does not occur.

It is worth noting that if we choose the family D_t so that $D_t := \bar{D} \cap (-\infty, (1-t)x_{\min} + tx_{\max}] \times \mathbb{R}$, where $x_{\min}$ and $x_{\max}$ are the minimal and maximal x-coordinates of points in D (assume that $\bar{D}$ contains exactly one point with x-coordinate $x_{\min}$ for formal convenience), the generalized dynamic representation **(GE)** coincides with the original Arak and Surgailis' one determined by rules **(E1–4)** and we have $\mathbb{A}(l) = \mathrm{in}(l, D)$ where $\mathrm{in}(l, D)$ is the entry point of l into D as considered in the proof of theorem 1 before; that is to say the point of $l \cap \bar{D}$ with the minimal x-coordinate.

It should be emphasized at this point that the geometry of the *growing window* D_t determines the field-generating dynamics only through the induced order in which the tessellation nodes and segments are revealed, whereas the probabilities of the corresponding critical events (directional updates and line and vertex births) are not affected. This is because of the fact that these critical events can only occur at the deterministic moments where D_t hits a new node or segment, which stands in contrast to the continuum case discussed in section 4 of Schreiber (2008), where the critical events could occur at any time moment with rates essentially depending on the evolution of D_t.

In analogy with the corresponding result for the usual dynamic construction, as established in the proof of theorem 1, we show that the field resulting from the aforegiven **(GE)** construction does coincide in law with the contour ensemble $\mathcal{A}_{\Phi_D}$ and, consequently, with $\hat{\mathcal{A}}_{\Phi_D}$ upon choosing with probability 1/2 one of the two possible colourings.

Theorem 3
The random contour ensemble resulting from the aforementioned construction **(GE)** *coincides in law with* $\mathcal{A}_{\Phi_D}$.

Proof. To verify the statement of the theorem, we choose some $\gamma \in \Gamma_D(\mathcal{T})$ and compute the probability that the colour-blind contour ensemble γ is obtained as a result of the aforementioned construction. To this end, we observe the following.

- Each edge $e \in E(\gamma)$ containing the anchor point $\mathbb{A}(l[e])$ in its interior yields a factor $\pi_{l[e]}/(1+\pi_{l[e]})$ (line birth) times $\prod_{l \in \mathcal{T}, l \not\sim e}(1-\pi_l)$ (no direction updates along e).

- Each of the two edges $e_1, e_2 \in E(\gamma)$ stemming from a common interior birth site $n(l_1, l_2)$ yields a factor $\pi_{l_i}, i = 1, 2$ (coming from the birth probability) times $\prod_i \prod_{l \in \mathcal{T}, l \not\sim e} (1 - \pi_l)$ (no direction updates along e_i).
- Each of the edges $e \in E(\gamma)$ arising as a result of a direction update yields a factor $\pi_{l[e]}$ (direction update probability) times $\prod_{l \in \mathcal{T}, l \not\sim e} (1 - \pi_l)$ (no direction updates along e).
- The absence of birth sites at nodes $n(l_1, l_2)$ of $\mathcal{T}$ in D not belonging to γ yields the factor $\prod_{n(l_1,l_2) \in D \setminus \gamma} (1 - \pi_{l_1} \pi_{l_2})$ (note that birth sites are allowed in all points of γ – either they give rise to lines tracing γ or are discarded if hit by a previously born line).
- The absence of line births at those anchor points $\mathbb{A}(l), l \in \mathcal{T}$, which do not give rise to an edge of γ yields the factor $\prod_{l \in \mathcal{T}, l \cap D \neq \emptyset, \mathbb{A}(l) \notin \gamma} (1 + \pi_l)^{-1}$.

Putting these factors together allows us to evaluate the probability of γ arising as a result of the **(GE)** construction as:

$$\left(\prod_{e \in E(\gamma)} \pi_{l[e]} \right) \left(\prod_{e \in E(\gamma)} \prod_{l \in \mathcal{T}, l \not\sim e} (1 - \pi_l) \right) \left(\prod_{n(l_1,l_2) \in \gamma} (1 - \pi_{l_1} \pi_{l_2}) \right)^{-1} \left(\prod_{n(l_1,l_2) \in D} (1 - \pi_{l_1} \pi_{l_2}) \right) \times \left(\prod_{l \in \mathcal{T}, l \cap D \neq \emptyset} \frac{1}{1 + \pi_l} \right) = \frac{2 \exp(-\Phi_D(\hat{\gamma}))}{\mathcal{Z}[\Phi_D]} \prod_{e \in E(\gamma)} \pi_{l[e]} \tag{8}$$

in full analogy to (7) with $\mathcal{Z}[\Phi_D]$ given by (6). This completes the proof as in the argument establishing the dynamic representation in theorem 1.

7. Path creation and annihilation dynamics for general polygonal fields

The purpose of the present section is to use the generalized dynamic representation to construct dynamics on the space of polygonal configurations in a spirit similar to the derivation of the disagreement loop dynamics from the basic Arak and Surgailis' dynamic representation discussed previously in section 4. This means in particular that the path creation and annihilation procedure is a direct generalization of the disagreement loop dynamics; there are good reasons, though, to treat the disagreement loop procedure separately, as it is usually less complex from a computational viewpoint owing to the particularly simple nature of the anchor mapping $\mathbb{A}(l) = \mathrm{in}(l, D)$. Also, as we shall see next, the main advantage of choosing the family D_t different from the one corresponding to the Arak and Surgailis' construction relies on the particular form of line birth events, whereas the vertex birth Monte-Carlo moves are equally well generated by the simpler disagreement loop dynamics. It should be noted here that further interest in the generalized dynamic representation stems from its theoretical implications, see Schreiber (2008), giving an insight into the higher-order correlation structure of the field, which falls beyond the scope of the present application-oriented article.

To proceed towards the so-posed objective, we note first that, in full analogy with the **DL** dynamics given before, adding a new line birth site for $l \in \mathcal{T}$ at its anchor point l as specified in **(GE:LineBirth)** and letting it evolve thereupon according to the rules of the **(GE)** dynamics under the same coupling as for **DL**, yields a disagreement loop/path as the symmetric difference between the previous configuration γ and the new one $\gamma \oplus_{D_t} l$, where $\oplus_{D_t}$ is the line creation operator corresponding to the choice of $(D_t)_{t \in [0,1]}$. Likewise, annihilating a line birth site, with the corresponding operator denoted by $\ominus_{D_t}$, yields under the usual coupling a disagreement loop/path as well. Clearly, disagreement paths or loops are also obtained upon adding or removing a vertex, with the corresponding operators denoted again by $\oplus_{D_t}$ and $\ominus_{D_t}$ with no ambiguity arising between the line creation and annihilation operators which can be distinguished by the different nature of their right argument.

It should be noted at this point that the so-defined operations may result in a disagreement loop wrapping around the growing window D_t, with its two branches meeting and annihilating on the other side – a careful reader might wonder at this point whether such a geometry of a disagreement loop does not imply that its parts evolve *backwards in time*. The answer is always *no*, which formally follows by the construction of the dynamics, yet the more intuitive explanation stems from the fact that, unlike in the original space-time dynamics discussed in the proof of theorem 1, here we in general do not have globally defined *time axis* and *spatial axis* – on the contrary, the time flow direction may vary from point to point and is determined by the local normal direction to the boundary of D_t.

Consider the following *path creation and annihilation* Markovian (**PCA**) dynamics in continuous time:

PCA:CreateLine For each $l \in \mathcal{T}$, with intensity $\pi_l/(1+\pi_l)\,\mathrm{d}s$, set $\gamma_{s+\mathrm{d}s} := \gamma_s \oplus_{D_t} l$.
PCA:AnnihilateLine For each $l \in \mathcal{T}$ such that $\mathbb{A}(l) \in \gamma$, with intensity 1, set $\gamma_{s+\mathrm{d}s} := \gamma_s \ominus_{D_t} l$.
PCA:CreateVertex For each $x := n(l_1, l_2) \in D, l_1, l_2 \in \mathcal{T}$, with intensity $\pi_{l_1}\pi_{l_2}\,\mathrm{d}s$, set $\gamma_{s+\mathrm{d}s} := \gamma_s \oplus_{D_t} x$.
PCA:AnnihilateVertex For each interior birth site x in γ_s, with intensity 1, set $\gamma_{s+\mathrm{d}s} := \gamma_s \ominus_{D_t} x$.

In full analogy with proposition 1 we get proposition 2.

Proposition 2
The distribution of the consistent colour-blind polygonal field $\mathcal{A}_{\Phi_D}$ is the unique invariant law of the **PCA** *dynamics. The resulting stationary process is reversible. Moreover, for any initial distribution of γ_0 the laws of the random polygonal fields γ_s converge in variational distance to the law of $\mathcal{A}_{\Phi_D}$ as $s \to \infty$.*

Clearly, a version of this dynamics can be constructed for general polygonal fields with Hamiltonian $\Phi_D + \mathcal{H}_D$ for some $\mathcal{H}_D : \hat{\Gamma}_D(\mathcal{T}) \to \mathbb{R}$. To this end, we regard the **PCA** dynamics as a generator of update proposals, which are then accepted with probability 1 if $\mathcal{H}_D$ decreases and with probability $\exp(-\Delta\mathcal{H}_D)$ otherwise, where $\Delta\mathcal{H}_D$ is the increase in $\mathcal{H}_D$. To put it in formal terms, proceed as follows:

PCA[H] Given $\hat{\gamma}_s$ generate an update proposal δ for $\gamma_{s+\mathrm{d}s}$, colour it randomly in one of two possible ways with probability 1/2 each, thus obtaining the coloured update proposal $\hat{\delta}$. Now, if $\mathcal{H}_D(\hat{\delta}) < \mathcal{H}_D(\hat{\gamma}_s)$, set $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\delta}$, otherwise set $\hat{\gamma}_{s+\mathrm{d}s} := \hat{\delta}$ with probability $\exp(-[\mathcal{H}_D(\hat{\delta}) - \mathcal{H}_D(\hat{\gamma}_s)])$ and keep $\hat{\gamma}_{s+\mathrm{d}s} = \hat{\gamma}_s$ with the complementary probability.

Using proposition 2 and verifying the usual detailed balance conditions we come to theorem 4.

Theorem 4
The distribution of the polygonal field $\hat{\mathcal{A}}_{\mathcal{H}_D+\Phi_D}$ is the unique invariant law of the dynamics **PCA[H]**. *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 variational distance to the law of $\hat{\mathcal{A}}_{\Phi_D+\mathcal{H}_D}$ as $s \to \infty$.*

8. Image segmentation as a statistical inference problem

Here we recall the framework for Bayesian image interpretation using planar random field priors (see, e.g. Chellapa & Jain, 1993).

Let the data consist of some discretized grey-level image. Write S for the set of sites and $L=\{0,\ldots,255\}$ for the label set. Formally, an image is a vector $y=(y_\varsigma)_{\varsigma\in S}$ with $y_\varsigma\in L$. 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.

In the Bayesian setting, the task of segmenting y into foreground and background regions is formulated as a statistical parameter estimation problem. As the aim is to interpret y in terms of a black-and-white-coloured polygonal configuration, we seek to find a $\hat{\gamma}^*\in\hat{\Gamma}_D(\mathcal{T})$ that explains the data 'best' in the sense of having a small *misclassification rate* as well as exhibiting desirable 'smoothness' properties. The former aspect is captured in the *likelihood*, the latter in the *prior*. Inference is based on the *posterior polygonal field distribution*

$$\frac{\exp\left(-Y_{D,S}^{y,\beta}(\hat{\gamma})\right)}{\mathcal{Z}[Y_{D,S}^{y,\beta}]}\prod_{e\in E(\gamma)}\pi_{l[e]} \tag{9}$$

on $\hat{\Gamma}_D(\mathcal{T})$ given y with

$$Y_{D,S}^{y,\beta}(\hat{\gamma})=\Phi_D(\hat{\gamma})+\beta X_{D,S}(\hat{\gamma};y),\quad \hat{\gamma}\in\hat{\Gamma}_D(\mathcal{T}), \tag{10}$$

for inverse temperature $\beta>0$, and partition function

$$\mathcal{Z}[Y_{D,S}^{y,\beta}]=\sum_{\hat{\theta}\in\Gamma_D(\mathcal{T})}\exp(-Y_{D,S}^{y,\beta}(\hat{\theta}))\prod_{e\in E(\theta)}\pi_{l[e]}.$$

Thus, (9) is a law of the form (1) with a Hamiltonian that is a weighted sum of two terms: a prior $\Phi_D(\hat{\gamma})$ given by (5) balanced by a term $X_{D,S}(\hat{\gamma};y)$ that describes the goodness-of-fit between the coloured configuration $\hat{\gamma}$ and the data y. Observe that (9) can also be regarded as the (posterior) distribution of a classical Gibbs field on $\hat{\Omega}_D(\mathcal{T})$, the space consisting of all binary colourings of $D_{\mathcal{T}}$, with a Hamiltonian of the form (4),

$$\Psi_D(\hat{\gamma})=\begin{cases}Y_{D,S}^{y,\beta}(\hat{\gamma})-\sum_{e\in E(\gamma)}\log\pi_{l[e]}, & \hat{\gamma}\in\hat{\Gamma}_D(\mathcal{T}),\\ +\infty, & \text{otherwise.}\end{cases} \tag{11}$$

In the examples to be presented in section 9, we shall use the 'misclassification rate'

$$X_{D,S}(\hat{\gamma};y)=\frac{1}{n_S}\sum_{\varsigma\in S}\left|\frac{y_\varsigma}{255}-\hat{\gamma}[\varsigma]\right|, \tag{12}$$

where $\hat{\gamma}[\varsigma]\in\{0,1\}$ denotes the (normalized) colour of $\hat{\gamma}$ at pixel ς, and n_S is the number of pixels. Note that as D contains S, $\hat{\gamma}[\varsigma]$ is well defined.

Upon observation of y, our goal is to estimate the underlying coloured polygonal configuration by Hamiltonian optimization. More specifically, we use the simulated annealing algorithm (Geman & Geman, 1984; Haario & Saksman, 1991) and let $\beta\uparrow\infty$ to find the polygonal configurations $\hat{\gamma}^*$ having minimal misclassification rate while keeping the hyperparameters $\pi_{l[e]}$ for line desirability fixed.

9. Examples

In this section, we present results produced by our C++ implementation of the aforementioned algorithms on toy and real-life examples. The model (9) was used with components

(5) and (12). We used a combination of the **DL[H]:birth–death** algorithm of section 5 and the **PCA[H]** dynamics of section 7 (with the window D_l initiating its growth in the middle of a randomly chosen edge and growing as its parallel rectangle thereupon) combined with the usual local basic polygon recolourings constituting the core of the standard simulation algorithms for lattice-indexed Gibbs–Markov fields (see, e.g. Winkler, 2003). Path creation and annihilation, being global in nature, is useful in the beginning of the simulated annealing procedure, as are the disagreement loop updates. In later stages of the simulation they are complemented with local polygon updates to take care of fine details.

We applied simulated annealing, treating (5) as the Hamiltonian of the reference distribution, and a linear cooling schedule that amounts to setting β equal to the number of iterates divided by 60. For the family $\mathcal{T}$ we chose a random non-homogeneous *gas* of lines with intensity determined by the gradient field of the processed image and interacting by hard-core exclusions within a fixed distance. More precisely, new line proposals were generated by

- choosing a random point in the image domain with density proportional to gradient length,
- choosing the new line direction perpendicular to the gradient vector,
- accepting the line whenever it is far enough from the ones already present, with the distance between lines identified with the distance between their points closest to the origin.

Moreover, lines were also removed with constant intensity and the system was allowed to relax towards equilibrium (for a short time). We decided to apply this computationally inexpensive relaxation scheme because theoretically the equilibrium gas should ensure better packing effectiveness and regularity than the simpler random sequential packing process. In practice, the choice between equilibrium packing and random sequential packing seems to have no effect on the overall performance of our software. However, an important improvement of the segmentation performance was achieved by making the line activity parameters $\pi_l, l \in \mathcal{T}$, proportional to the absolute gradient flux through l; in fact this has been pushed further in our implementation and local activities ascribed to edges rather than lines were tested with a good effect.

The general conclusion from the simulations discussed next is that the right choice of tessellation is crucial for a good performance of our algorithm. The present random tessellation-generating mechanism has the property of producing quite reasonable tessellations, among which very good and very bad ones do happen sometimes. This leads to striking differences of segmentation performances between subsequent simulation runs with the same parameters and the same picture processed, see the following examples. Consequently, our present active research effort is concentrated on developing a good picture-driven tessellation-generating mechanism well integrating with the polygonal field set-up and replacing the *ad hoc* gradient-driven equilibrium line packing as applied so far. The envisioned results will be presented in a separate paper.

Our present software was first tested on synthetic images. At each iteration, the disagreement loop dynamics were chosen with probability 0.85, path creation and annihilation with probability 0.05, leaving probability 0.1 for local polygon updates. Approximately 12,500 Monte-Carlo updates were carried out per second of simulation time on an Intel Pentium M 2 GHz with 2 GB RAM memory. Each resulting segmentation is presented together with the underlying tessellation and the number of updates as well as the pixel misclassification rate (PMR) reached are given. The target PMR was set to 0.03 and the maximal number of iterates was 750,000 corresponding to about 1 minute simulation time.

Figure 2 (left) shows a blurry 'A' overlaid on the segmentation result obtained after 478,000 steps (92,000 updates accepted) with the target misclassification rate of 0.03 reached. The corresponding tessellation is shown on the right-hand side of the figure.

Figure 3 (left) is the result after 260,000 iterations (45,000 accepted updates) on a 'B' with the target of 0.03 PMR achieved. The underlying tessellation is shown in Fig. 3 (right).

Further, in Fig. 4 (left) we show the result of another, less successful run of the software, on the same picture and with the same parameters – the source of trouble was that the underlying random tessellation, as seen in the right of the figure, did not contain lines needed to mark the horizontal inner bar of the letter 'B'. The achieved misclassification rate was 0.038 after 750,000 steps (132,000 accepted updates).

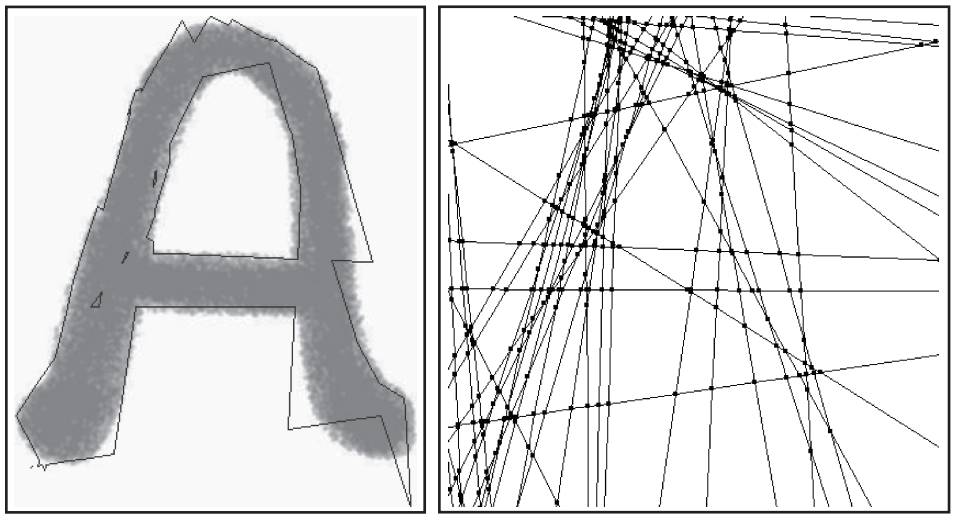

Fig. 2. Segmented letter A (left) and the underlying tessellation (right).

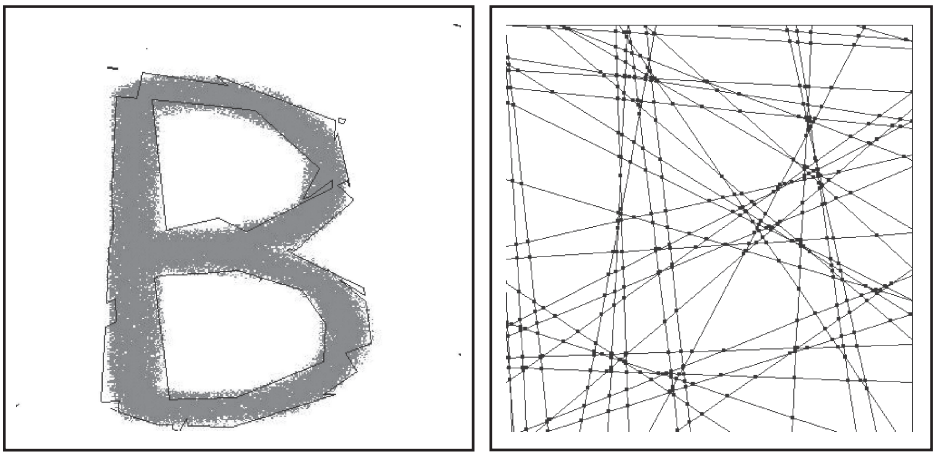

Fig. 3. Segmented letter B (left) and the underlying tessellation (right).

Also, we used grey-level images from the Berkeley Segmentation Dataset and Benchmark web-site (http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/) to evaluate our approach, as well as images from the PASCAL Network of Excellence challenge 2006 (http://www.pascal-network.org/challenges/VOC/thumbs/VOC2006/index1.html).

Figure 5 (left) depicts a cat. The target 0.03 PMR was obtained after 194,000 steps (40,000 accepted updates). The corresponding tessellation is shown in the right-hand side of the figure.

Finally, Fig. 6 shows a picture of a mushroom against a bushy background, with the segmentation result obtained after 750,000 steps (100,000 accepted updates). This picture turned out to be significantly more difficult and the PMR achieved was 0.067.

The discrete nature of polygonal fields defined using tessellations implies that, unless the tessellation is properly chosen, it is not possible in principle to follow exactly the outline of the foreground object. This is why the tessellation choice is so crucial. However, the state

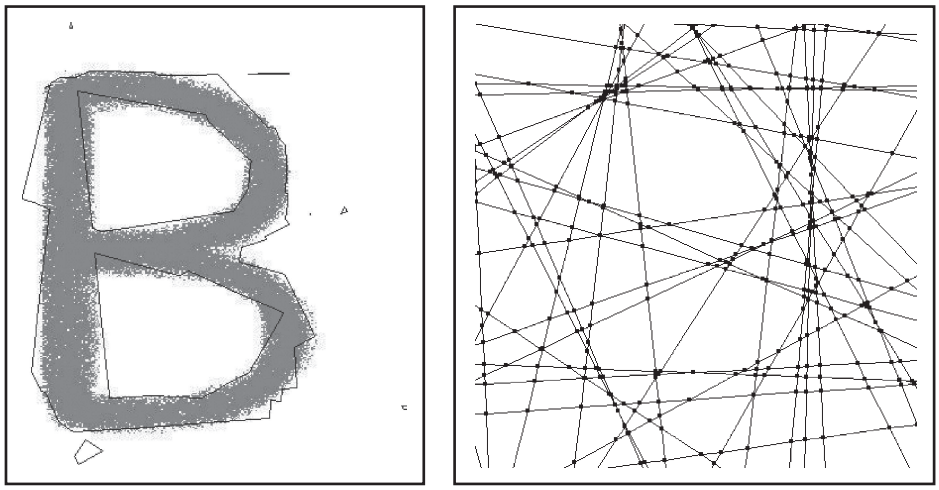

Fig. 4. Another segmentation of B (left) and the underlying tessellation (right).

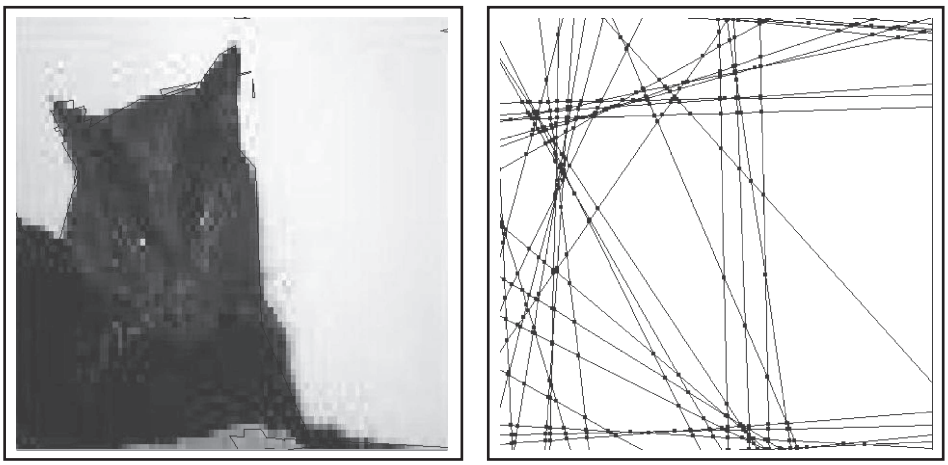

Fig. 5. Segmented image of a cat (left) and the tessellation (right).

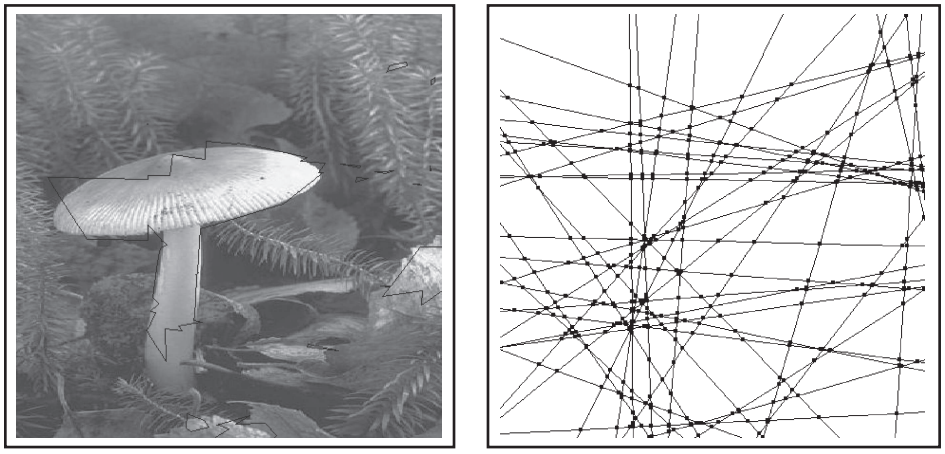

Fig. 6. Segmented image of a mushroom (left) and the tessellation (right).

space of the annealing algorithm is much reduced, resulting in a faster procedure than that described for the continuum models by Kluszczyński *et al.* (2005, 2007).

10. Discussion and conclusion

In this article, we presented a class of Gibbs fields on finite graphs generated by regular tessellations, which can be interpreted as discretizations of the consistent polygonal Markov fields introduced by Arak & Surgailis (1989). We showed that these fields enjoy a number of striking properties (including consistency, solvability and Markovianity) similar to those of the continuous polygonal Markov fields by which their definition was inspired. In particular, a dynamic representation can be constructed, which may be used to design simulation algorithms based on disagreement loops as well as PCA. By addition of a Hamiltonian that measures the goodness-of-fit by an L_p criterion, the model and techniques were applied to the image analysis problem of foreground–background separation.

It should be noted that, in contrast to the definition of a usual Markov random field, we have excluded X-shaped nodes, as doing so allows us to obtain a neat theory and avoids ambiguity in distinguishing the contours that comprise a polygonal configuration. The condition can easily be lifted, though, with only a few differences to the theory.

- In the dynamic construction developed in section 3 we admit births of new particles also at sites where a prior collision occurred.
- The Hamiltonian $\Phi_D(\hat{\gamma})$ gets an extra term $-\log(1-\pi_{l_1}\pi_{l_2})$ for each particle collision (death) site $n(l_1,l_2)\in D$ of γ at which no new particle is born – this extra term takes into account the probability of not having an X-shaped node (not observing a new birth) at the death site $v(l_1,l_2)$, which is $1-\pi_{l_1}\pi_{l_2}$. Unfortunately, the so-modified Hamiltonian becomes dependent on the choice of the coordinate system (spatial and time axes) unlike the original Hamiltonian in (5); this is because the notion of the birth site does depend on the direction of the time flow.

The algorithms proposed in this article can be modified accordingly. Clearly, if a rich enough collection $\mathcal{T}$ is used, the choice to include or exclude X-shaped nodes hardly matters for segmentation purposes.

Exploration of the full potential of the class of consistent polygonal fields built on finite tessellations is beyond the scope of the present article and part of our work in progress. In particular, the choice of the collection $\mathcal{T}$ and the assignment of probabilities π_l to each $l \in \mathcal{T}$ remains as subject of further study. From the perspective of image segmentation, it would be of particular interest to base the choice of tessellation on gradient information. Indeed, as edges lie at the boundaries between more or less homogeneous regions, they are characterized by abrupt changes in intensity values, and contain valuable information about the structures and objects present in the image. Hence, computation of the edge map of an image is useful, as the amount of data that needs to be stored may be greatly reduced, while preserving most of the information of interest (Rosenfeld & Kak, 1982). The Canny (1986) filter is widely regarded as the best general purpose edge detector, designed to combine high signal-to-noise ratio and precise localization and composed of Gaussian smoothing, Sobel filter, non-maximum suppression and hysteresis steps.

We believe that foreground–background separation is not the only application for which the models and techniques discussed in this article can be applied. As the Arak (1982) process can be extended to allow for more than two colour labels (Arak & Surgailis, 1991; Kluszczyński *et al.*, 2007), so is it of interest to explore their discrete counterparts for multi-class segmentation (current work in progress).

Another application area is the detection of linear features such as road networks or edges from image data. Several marked point process models have been proposed as prior distributions for this task (Stoica, 2001; van Lieshout & Stoica, 2003; Lacoste, 2004). However, being defined by a density with respect to a Poisson process, such line segment models are not capable of reproducing such characteristics of real-life images as an abundance of parallel lines, joined endpoints or preferred angles of crossing. Consistent discrete polygonal field models, in contrast, are. Related, somewhat easier, conditional simulation problems occur in cognitive experiments aimed at understanding the human visual system.

Acknowledgements

The authors 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) which made possible mutual visits. Van Lieshout also wishes to express gratitude for the warm hospitality and additional financial support received from the Nicolaus Copernicus University, Toruń, Poland. The authors acknowledge R. Kluszczyński for technical assistance. Schreiber acknowledges support from the Polish Minister of Science and Higher Education grant N N201 385234 (2008–10). Special thanks are also due to the anonymous referees whose comments were very helpful in improving this paper.

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. *Probab. Theory Related Fields* **80**, 543–579.

Arak, T. & Surgailis, D. (1991). Consistent polygonal fields. *Probab. Theory Related Fields* **89**, 319–346.

Arak, T., Clifford, P. & Surgailis, D. (1993). Point-based polygonal models for random graphs. *Adv. Appl. Probab.* **25**, 348–372.

Besag, J. (1986). On the statistical analysis of dirty pictures (with discussion). *J. Roy. Statist. Soc. Ser. B* **48**, 259–302.

Canny, J. (1986). A computational approach to edge detection. *IEEE Trans. PAMI* **8**, 679–698.

Chellapa, R. & Jain, A. (eds) (1993). *Markov random fields: Theory and applications*. Academic Press, Boston.

Clifford, P. & Middleton, R. D. (1989). Reconstruction of polygonal images. *J. Appl. Statist.* **16**, 409–422.

Clifford, P. & Nicholls, G. (1994). A Metropolis sampler for polygonal image reconstruction. Available at: http://www.stats.ox.ac.uk/clifford/papers/met_poly.html

Geman, S. & Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Trans. PAMI* **6**, 721–741.

Gimel'farb, G. L. (1999). *Image textures and Gibbs random fields*. Kluwer, Dordrecht.

Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–732.

Haario, H. & Saksman, E. (1991). Simulated annealing process in general state space. *Adv. Appl. Probab.* **23**, 886–893.

Kluszczyński, R., van Lieshout, M. N. M. & Schreiber, T. (2005). An algorithm for binary image segmentation using polygonal Markov fields. In *Image Analysis and Processing, Proceedings of the 13th International Conference on Image Analysis and Processing* (eds F. Roli & S. Vitulano), 383–390. *Lecture Notes in Computational Science* 3615. Springer Berlin, Heidelberg.

Kluszczyński, R., van Lieshout, M. N. M. & Schreiber, T. (2007). Image segmentation by polygonal Markov fields. *Ann. Inst. Statist. Math.* **59**, 465–486. Springer, Berlin-Heidelberg.

Lacoste, C. (2004). *Extraction de réseaux linéiques à partir d'images satellitaires et aériennes par processus ponctuels marqués*. PhD Thesis, University of Nice-Sophia, Antipolis.

van Lieshout, M. N. M. & Stoica, R. S. (2003). The Candy model: properties and inference. *Statist. Neerlandica* **57**, 177–206.

Liggett, T. (1985). *Interacting particle systems*. Springer-Verlag, New York.

Møller, J. & Skare, Ø. (2001). Bayesian image analysis with coloured Voronoi tessellations and a view to applications in reservoir modelling. *Stat. Modelling* **1**, 213–232.

Nicholls, G. K. (1998). Bayesian image analysis with Markov chain Monte Carlo and coloured continuum triangulation models. *J. Roy. Statist. Soc. Ser. B Statist. Methodol.* **60**, 643–659.

Rosenfeld, A. & Kak, A. C. (1982). *Digital picture processing*, Vol. II, 2nd edn. Academic Press, New York.

Schreiber, T. (2005). Random dynamics and thermodynamic limits for polygonal Markov fields in the plane. *Adv. Appl. Probab.* **37**, 884–907.

Schreiber, T. (2008). Non-homogeneous polygonal Markov fields in the plane: graphical constructions and geometry of higher-order correlations. *J. Statist. Phys.* **132**, 669–705.

Stoica, R. (2001). *Processus ponctuels pour l'extraction des réseaux linéiques dans les images satellitaires et aériennes*. PhD Thesis, University of Nice-Sophia, Antipolis.

Surgailis, D. (1991). The thermodynamic limit of polygonal models. *Acta Appl. Math.* **22**, 77–102.

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 27. Springer-Verlag, Berlin.

Received May 2008, in final form August 2009

T. Schreiber, Faculty of Mathematics & Computer Science, Nicolaus Copernicus University, ul. Chopina 12-18, 87-100 Toruń, Poland.
E-mail: tomeks@mat.uni.torun.pl