New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

end p.426

13 Applications of Stochastic Geometry in Image Analysis

Marie-Colette N.M. van Lieshout

13.1 Introduction

The new millennium has opened with what can only be described as a data explosion following advances in digital technology, a development that has led to a strong demand for tools to analyse digital data such as still images, video streams, audio signals, and text.

As a field, statistical image analysis took off during the 1980s with the seminal work by Besag (1986) and Geman and Geman (1984) on the restoration of pictures degraded by noise. Much of the work in this period is focussed on 'low level' tasks, that is, it aims to de-noise, sharpen, segment, or classify the image. A good overview can be found in the supplement to *Journal of Applied Statistics* edited by Mardia and Kanji (1993) which includes reprints of the seminal papers mentioned above and a list of early references.

With the improvement in image quality, in the course of the 1990s the emphasis shifted towards the 'high level' task of describing an image by its content, for example in terms of the objects it contains and the spatio-temporal relations between them. Early work in this direction includes Baddeley and Van Lieshout (1992, 1993), Molina and Ripley (1989), Ripley and Sutherland (1990), as well as the work by Grenander and co-authors — although usually couched in pattern theoretic language (Grenander 1976, 1978, 1981).

Given the difference in goals, it is hardly surprising that different stochastic models are used in low and high level vision. For example, classification at the lowest conceptual level calls for a pixel based model. Even in this context, though, due to the high dimension of image data, fitting a model is a far from trivial task that cannot be done explicitly except for the simplest models. Instead, an algorithmic approach is often taken in which small changes — involving only a few pixels — are proposed iteratively. The crucial ingredient of many such algorithms is the difference in log likelihood between the new and old states. for

end p.427

computational reasons, it would be highly desirable if this ratio would be 'local', a concept that can be formalized by a Markov property. Further details can be found in Geman (1990) or Winkler (2003).

At the other extreme, the focus of attention are the objects in the image as described by their location, shape, and colour parameters. In this context, it is natural to use marked point processes (Daley and Vere-Jones 2003), in particular those satisfying a Markov property (Van Lieshout 2000a). Which mark to use depends on the application, ranging from simple geometric shapes (Baddeley and Van Lieshout 1992, Van Lieshout 1994, 1995) through deformable template models (Amit *et al.* 1991, Hansen *et al.* 2002, Hurn 1998, Mardia *et al.* 1997, Pievatolo and Green 1998, Rue and Hurn 1999, Rue and Husby 1998) via the complex ensembles of simple shapes studied by Lacoste *et al.* (2005), Stoica *et al.* (2002, 2004, 2007), and Ortner *et al.* (2007) to the spatio-temporal models of Van Lieshout (2007).

Intermediate level modelling tries to avoid a full scene description while preserving a global approach. It focusses therefore on image regions, either conceived as a conglomerate of pixels (Møller and Waagepetersen 1998, Tjelmeland and Holden 1993) or as a (continuous) tessellation of space (Clifford and Middleton 1989, Kluszczyński *et al.* 2007, Møller and Skare 2001, Nicholls 1998).

The present chapter is organized as follows. In Section 13.2, we present random fields, (sequential) object processes and polygonal field models. In Section 13.3, we discuss relations between the various Markov properties discussed in Section 13.2, indicate how a random field is affected by the choice of the number of class labels, and show that certain polygonal field models can be seen as non-overlapping object processes. In the last section, we present two applications. The first one concerns tracking of a variable number of moving objects through a video sequence, the other one a foreground/background separation problem. For the first example, a ground truth is known to which the outcome can be compared; in the second application, visual inspection must validate the results.

The emphasis in this chapter is on stochastic geometric modelling at the expense of a detailed discussion of the Monte Carlo algorithms employed for inference. The interested reader is referred to, e.g. Chapter 9 or Winkler (2003). We also refrain from discussing mathematical morphology, a branch of image analysis with similar roots to stochastic geometry (Matheron 1975).

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

Instead we refer to Serra (1982) or Dougherty (1992).

13.2 Stochastic geometric models: From random fields to object processes

13.2.1 Random fields

A digital image is simply a finite array of 'colour' or 'intensity' labels. The entries of the array S, typically a rectangle of raster points in $\mathbb{Z}^2$, are referred to as pixels. We shall use the notation $S = (s_1, \ldots, s_m)$, $m \in \mathbb{N}$. For the moment, assume the set of labels is finite as well, and denoted by Λ.

end p.428

Let $X = (X_1,\ldots,X_m)$ be a random field with values in Λ^S. Thus, X_i is the label in Λ assigned to pixel s_i. The labels can be categorical, for instance foreground/background, or be related to the intensity values of the image. The distribution of X is given by the joint probability mass function $\mathbf{P}\{X_1 = x_1,\ldots, X_m = x_m\}$ for $x = (x_1,\ldots, x_m) \in \Lambda^S$, which we assume to be strictly positive so that conditional distributions are well-defined. The condition may be relaxed a little (Besag 1974, Clifford 1990) to include models such as the lattice gas discussed below.

Let $\sim$ be a symmetric, reflexive relation on S. For instance, on $\mathbb{Z}^2$ one may define $s_i \sim s_j$ to be neighbours if and only if they are directly or diagonally adjacent, that is, if $\left|\,\left|s_i - s_j\right|\,\right| \leq \sqrt{2}$ In graph theoretical terms, the pixels are the vertices and an edge is drawn between $s_i \neq s_j$ if and only if $s_i \sim s_j$. The random field X is said to be Markov with respect to $\sim$ if for all $i = 1,\ldots, m$ the conditional distribution

$$\mathrm{P}\left\{X_i = x_i \,\middle|\, X_j = x_j,\ j \neq i\right\} = \mathrm{P}\left\{X_i = x_i \,\middle|\, X_j = x_j,\ s_i \sim s_j,\ j \neq i\right\} \tag{13.1}$$

depends only on x_i and the labels at those pixels s_j that share an edge with s_i (see e.g. Besag 1974).

Example: Lattice gas (Lebowitz and Gallavotti 1971) Let S be a finite square of raster points $\{s_1,\ldots,s_m\} \subseteq \mathbb{Z}^2$ equipped with the neighbourhood relation $s_i \sim s_j$ if and only if $\left|\,\left|s_i - s_j\right|\,\right| \leq \sqrt{2}$ Let Λ be the set {0,1, 2}, fix the parameter $\alpha > 0$, and consider the random field X on Λ^S whose joint probability mass function

$$\mathrm{P}\left\{X_1 = x_1,\ \ldots,\ X_m = x_m\right\} \propto \prod_{i=1}^{m} \alpha^{1_{x_i \neq 0}} \prod_{s_i \sim s_j, i<j} 1_{x_i = x_j \text{ or } x_i x_j = 0}$$

is proportional to the right hand side product. In words, labels are assigned independently with probabilities $1/(1 + 2\alpha)$ for 0 and $\alpha/(1 + 2\alpha)$ otherwise, conditional on the event that for every pair of neighbour pixels $s_i \sim s_j$, $i < j$, either one of the labels is zero ($x_i x_j = 0$) or they are identical ($x_i = x_j$). Clearly, some elements of Λ^S occur with probability zero. Nevertheless, whenever the conditioning event on the left hand side of (13.1) has strictly positive probability, the equation (13.1) holds, and we say that $\mathbf{P}$ is Markov with respect to $\sim$, with labels 1 and 2 'repelling' each other. A realization with $\alpha = 0.75$ is shown in Fig. 13.1.

From a computational point of view, the Markov property implies that iter- atively sampling X_i from the right hand side of (13.1) leaves $\mathbf{P}$ invariant, an observation that is crucial for Monte Carlo inference (see e.g. Geman 1990).

Assume that X is a random field with strictly positive joint probability mass function $\mathbf{P}$. Then, by the Hammersley—Clifford theorem (see the historical account by Clifford 1990), X is a Markov random field if and only if $\mathbf{P}$ can be written as

$$\mathrm{P}\left\{X_1 = x_1,\ \ldots,\ X_m = x_m\right\} = \prod_{C \in \mathcal{C}} \phi_C(x_c,\ c \in C) \tag{13.2}$$

end p.429

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

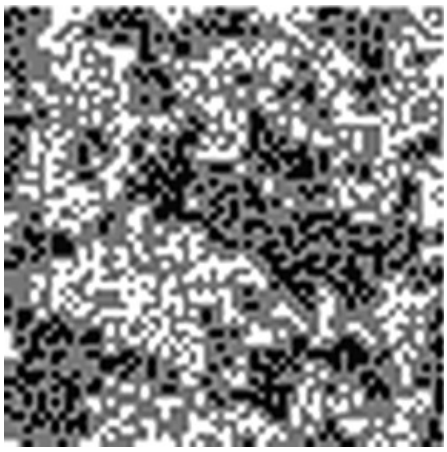

Fig. 13.1. Realization of a lattice gas model with α = 0.75. Black represents colour label '1', grey '0', and white '2'.

for some interaction functions $\varphi C : \Lambda^{C} \to \mathbb{R}^{+}$ defined for $C \in \mathbf{C}$, the family of sets consisting of pairwise neighbours. By convention, singletons and the empty set are included in $\mathbf{C}$. Indeed, φ_0 can be regarded as a normalizing constant.

Pairwise interaction models for which $\varphi_C = 1$ whenever the cardinality of C exceeds two are the most convenient, and widely used. The lattice gas model falls in this class. Indeed, $\phi_{\{s_i\}}(x_i) = \alpha^{1_{x_i \neq 0}}$ while $\phi_{\{s_i,s_j\}}(x_i, x_j) = 0$ if and only if $s_i \sim s_j$ and $x_i = 1$, $x_j = 2$ or vice versa, one otherwise. Such discouragement for neighbouring pixels to have different labels is typical for image analysis. Another famous example is the Potts interaction function $\phi_{\{s_i,s_j\}}(x_i, x_j) = \exp\left\{-\beta \ 1_{x_i \neq x_j}\right\}$ for some $\beta > 0$. The reader may find it useful to think of φ_C as a regularization term for the C-local behaviour of desirable realizations x of the random field X.

The global behaviour of a pairwise interaction Markov random field cannot be expected to reflect the global appearance of complex pictures, which is the reason why greedy ascent type algorithms tend to give better results than methods aimed at a global optimum (Winkler 2003). If global characteristics are the object of interest, it is advisable to use larger neighbourhoods and non-trivial interaction functions for a rich subfamily of $\mathbf{C}$ (Tjelmeland and Besag 1998). Note that (13.2) can be augmented with interactions on the dual graph (Geman *et al.* 1990, Geman and Reynolds 1992) when preservation of the natural edges in an image is important.

A stronger form of conditional local dependence is that of the Markov mesh models proposed by Abend *et al.* (1965). Let S be a finite rectangle of raster points in Z^2 and call pixel s_j a predecessor of s_i if it either lies above or to the left of s_i. Then X is a Markov mesh model if the conditional distribution of X_i,

end p.430

$i = 1,\ldots, m$, given the values x_j at all predecessors s_j of s_i depends only on a limited set of s_j, say the three nearest ones. These conditional distributions are often assumed to be given in closed form and easy to simulate from. In this case, the joint probability mass function factorizes as the product over such conditional distributions, and – in contrast to most Markov random field models – is analytically tractable and amenable to sequential procedures. However, due to their causal non-symmetric nature, Markov mesh models tend to produce realizations with striping effects not apparent in many natural images (Cressie and Davidson 1998, Lacroix 1987, Qian and Titterington 1991).

To conclude this subsection, Gaussian Markov random fields deserve special mention. So far, for convenience, we have assumed that Λ is finite. However, the theory discussed here remains valid in the context of real-valued labels if we replace probability mass functions by densities. When X is normally distributed, the precision matrix Q captures the spatial dependence. Typically in applications, Q will be sparse, so that Monte Carlo methods can often be avoided and replaced by numerical ones (Rue and Held 2005). Note that the class of Gaussian fields includes classic conditional autoregression models for smoothly varying images (Ripley 1988, see also Besag and Kooperberg 1995, Besag *et al.* 1991, Künsch 1987).

13.2.2 Intermediate level modelling

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

Intermediate level image analysis aims to describe an image in terms of its homogeneous regions. There are two main strands. In the first, one stays within the random field framework but considers a more global interaction structure; the second approach is based on the concept of a spatial tessellation (chapter 5 in this volume).

Random fields Pairwise interaction Markov random fields like the Potts model with moderate neighbourhood sizes can be used successfully as regular- ization terms for low level tasks such as the restoration of noisy pictures (Section 13.2.1). However, for large values ofβ, realizations of the Potts model tend to be dominated by a single colour.

In order to construct models that tend to produce realizations containing compact regions, Tjelmeland and Besag (1998) introduced higher order interaction functions φ_C that express the relative likelihood of various types of labellings on C. For example, if $\Lambda = \{0,1\}$, the C could be monochrome, contain convex/concave corners or edges between the two labels. See also Gimel'farb (1999).

With the same objective in mind, Møller and Waagepetersen (1998) focussed on image regions directly and proposed Markov connected component fields defined by a joint probability mass function of the form

$$\mathrm{P}\left\{X_1 = x_1,\ \ldots,\ X_m = x_m\right\} \propto \prod_{K \in \boldsymbol{K}(x)} \Psi_K(l(x_K)), \tag{13.3}$$

end p.431

where the product ranges over the maximal ~-connected components of identically labelled pixels in $x = (x_1,\ldots, x_m)$, $l(x_k)$ denotes the common label in the component $x_k = (x_k, k \in K)$, and $\Psi_k : \Lambda \to \mathrm{R}^+$ is the interaction function. The functions Ψ_k may be based on fundamental geometric quantities such as the area, perimeter, and Euler–Poincaré characteristic of K, as well as on corner features similar to the ones considered by Tjelmeland and Besag (1998). Inter- component interaction can be taken into account, leading to a joint probability mass function proportional to

$$\prod_{K \in \boldsymbol{K}(x)} \Psi_K(l(x_K)) \prod_{\{K,L\} \subseteq \boldsymbol{K}(x)} \Phi_{K,L}(l(x_K), l(x_L)), \tag{13.4}$$

the semi-Markov random fields introduced by Tjelmeland and Holden (1993). Here, $\Phi_{k,l} \geq 0$ are the symmetric inter-component interaction functions between components K and L in $\mathbf{K}(x)$. Note that the Potts model of Section 13.2.1 permits a factorization of the form (13.3)–(13.4). In general, though, Markov random fields and connected components fields or semi-Markov random fields with respect to the same neighbourhood relation ~ are not comparable in the sense that neither class is contained in the other.

Spatial tessellations A coloured spatial tessellation is a partition of the image window $D \subseteq \mathrm{R}^2$, for example the convex hull of the pixels, into a finite number of disjoint, usually polygonal, regions whose united closures fill the window (Stoyan *et al.* 1995) and where each polygon carries a label. The labels may be assigned uniformly under the condition that no adjacent regions share the same label, or, if the polygons are seen as building blocks for larger regions, by a conditional Markov random field model with respect to the adjacency relation that encourages similar labels between polygons separated by a common boundary segment (for example a Potts or Gaussian model).

A random tessellation can be obtained in various ways. We shall focus here on models derived from the Poisson line process L. Note that a line $\ell_{\theta,p}$ in the plane may be parametrized by the pair (θ,p), where p is the signed distance of the line to the origin and $\theta \in [0, \pi)$ is the angle between the normal to the line and the x-axis as illustrated in Fig. 13.2. The measure $\mathrm{d}\theta\, \mathrm{d}p$ can be shown to be invariant under translations and rotations, so that the random set of lines L parametrized by a Poisson process on $[0, \pi) \times \mathbb{R}$ with intensity measure $\mathrm{d}\theta\, \mathrm{d}p$ is stationary and isotropic. Clearly, the lines of L that intersect D form a partition of the image window. Moreover, the angle between two 'randomly chosen' lines of L is distributed according to the probability density $\sin(\phi)/2$ on $[0, \pi)$.

More refined tessellations can be obtained by using the lines of L as a skeleton (Arak 1982). Suppose that D is a bounded, non-empty, open convex set with a piecewise smooth border ∂D, for example a rectangle. Define the family $\mathcal{T}_D$ of admissible tessellations of D as the set of all planar graphs γ in $D \cup \partial D$ with non- intersecting straight line segments as edges such that no two edges are co-linear, all vertices in D have degree two, and $\gamma \cap \partial D$ is either empty or consists of

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

end p.432

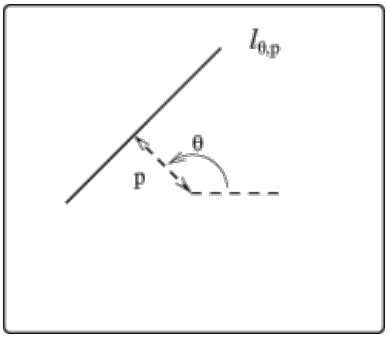

Fig. 13.2. Parametrization of lines in the plane.

vertices of degree 1. In statistical physics, the latter is referred to as an 'empty' or 'free' boundary condition, respectively. For a realization $l_D = \{l_i\}_{i=1}^n$ of the Poisson line process L_d restricted to D, let $\mathcal{X}_d(l_d)$ be the family of $\gamma \in \mathcal{X}_d$ such that the graph $\gamma \subseteq \cup_{i=1}^n l_i$ is constructed on lines of l_d and $\gamma \cap l_i$, $i = 1,\ldots, n$, contains a single interval of strictly positive length. Now, the polygonal Arak field A_d on D is defined by

$$\lim_{k\to\infty} \frac{\log \mathrm{P}\{W_1 \geq k\}}{\log(1/k)} = \tau_{\mathrm{W}} - 1. \tag{13.5}$$

where $l(\gamma)$ is the total edge length of γ and the expectation is with respect to the distribution of L_D. A realization of (13.5) with free boundary condition is shown in Fig. 13.3. The Arak field is Markovian in the following sense: the conditional distribution of A_D within a piecewise smooth closed curve C depends only on the intersection points and directions of A_d with C.

Fig. 13.3. Realization of an Arak field on [0,4] × [0,4].

end p.433

Several variations have been considered in the literature. Firstly, the probability distribution (13.5) may be used as a dominating measure to define further random tessellations. For instance, for $\beta > 0$, a length-interacting polygonal Markov field has probability density $f(\gamma) \propto \exp\{-\beta l(\gamma)\}$ with respect to the law of A_d (Arak 1982, Van Lieshout and Schreiber 2007, Schreiber 2005), thus favouring tessellations with small total edge lengths. Secondly, the class of admissable tessellations may be modified. In this vein, Nicholls (1998) restricts himself to triangles, Arak and Surgailis (1989) allow vertices of degrees 3 and 4, and Arak *et al.* (1993) consider point rather than line based constructions. Finally, Voronoi tessellations (Okabe *et al.* 2000) have been used as alternatives for polygonal fields (Blackwell and Møller 2003, Green 1995, Heikkinen and Arjas 1998, Møller and Skare 2001, Skare *et al.* 2007).

13.2.3 Object processes

As in the previous section, let D be the image window. Let Q be a Polish space that captures object attributes such as size,

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

shape and colour. We shall define an object process as a marked point process (see Chapter 1 of this volume) specified by its probability density f with respect to the distribution of a unit rate Poisson process on $D^{\circ} = D \cup \partial D$ marked in an i.i.d. fashion according to some mark distribution Q on Q. Realizations are denoted by $\mathbf{x} = \{x_1,\ldots, x_n\}$, where $n \in \mathrm{N0}$, and $x_i \in D^{\circ} \times Q$, $i = 1,\ldots,n$.

Example: Penetrable spheres (Widom and Rowlinson 1970) Let $Q = \{1, 2\}$. In analogy to the lattice gas model discussed in Section 13.2.1, the penetrable spheres mixture model is defined as a homogeneous Poisson process of rate $2\beta > 0$ in which points are independently assigned each of the two types with probability 1/2, conditional on the event that points of different type keep at least a distance $R > 0$ away from each other. A realization with $R = 0.1$ and $\beta = 100$ is shown in Fig. 13.4. Write X_1 for the ensemble of points with mark 1, X_2 for those marked 2. Then the model has joint probability density

$$f(\mathrm{x}_1, \mathrm{x}_2) \propto \beta^{n(\mathrm{x}_1)+n(\mathrm{x}_2)} \; \mathbf{1}_{d(\mathrm{x}_1,\mathrm{x}_2)>R}$$

with respect to the product measure of two independent unit rate Poisson processes on D°. Here $n(\mathbf{x}_i)$ denotes the cardinality of $\mathbf{x}_i$, and $d(\mathbf{x}_1,\mathbf{x}_2)$ is the shortest distance from a point in $\mathbf{x}_1$ to one in $\mathbf{x}_2$. The marginal distributions of X_i, $i \in \{1,2\}$, are called area-interaction processes (Baddeley and Van Lieshout 1995, Häggstrom *et al.* 1999) and have probability density proportional to

$$\beta^{n(\mathrm{x}_i)} \; \exp\left\{ -\beta \; \left| U_{\mathrm{x}_i} \cap \bar{D} \right| \right\}$$

with respect to a unit rate Poisson process on D°, where $|\cdot|$ denotes area, and U_{x_i} is the union of closed balls of radius R centred at the points of $\mathbf{x}_i$, $i = 1,2$. Note that points of the same type prefer to be close.

end p.434

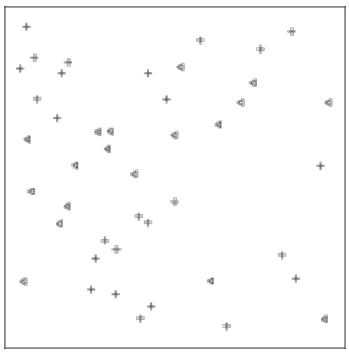

Fig. 13.4. Realization of a penetrable spheres model on [0,1] × [0,1] with $\beta = 100$ and $R = 1/10$. Points of type 1 are represented by triangles, points of type 2 by crosses.

Let ~ be a reflexive, symmetric neighbourhood relation on $D^{\circ} \times Q$. For example, the objects parametrised by x and y may be neighbours whenever they overlap (allowing for blur and shading where necessary). An object process X with probability density f is said to be Markov with respect to ~ if f is hereditary in the sense that whenever $f(\mathbf{x}) > 0$ for some configuration of objects, also $f(\mathbf{y}) > 0$ for all $\mathbf{y} \subseteq \mathbf{x}$, and, for all $u \in (D^{\circ} \times Q) \setminus \mathbf{x}$, the ratio

(13.6)

known as the conditional intensity (Papangelou 1974), depends only on u and $\{x_i \in \mathbf{x} : u \sim x_i\}$ (Ripley and Kelly 1977). A factorization is provided by the Ripley—Kelly theorem stating that a marked point process with probability density f is a Markov object process if and only if its probability density can be written as

$$f(\mathrm{x}) = \prod_{\text{cliques } \mathrm{y} \subseteq \mathrm{x}} \phi(\mathrm{y}) \qquad (13.7)$$

for some non-negative interaction function ϕ. The product ranges over object configurations $\mathbf{y} \subseteq \mathbf{x}$ that consist of pairwise neighbours (including singletons and the empty configuration). The resemblance to (13.1)is obvious. To see that the area-

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

interaction process is Markov, define $x \sim y$ if and only if $\|x - y\| \leq 2R$, and observe that the conditional intensity $\lambda(u \mid \mathbf{x}) = \beta\exp\{-\beta|(B(u, R) \setminus U_x) \cap D^-|\}$ depends only on u and those $x_i \in \mathbf{x}$ that are neighbours of u. The interaction function is given by

for $k \geq 1$.

end p.435

If $\phi(\mathbf{x}) = 1$ whenever $\mathbf{x}$ contains more than two objects, f defines a pairwise interaction process. Note that some care has to be taken to ensure that a model defined by some ϕ is integrable with respect to the dominating measure and hence can be normalized to unity. A sufficient condition is that $\phi(\mathbf{x}) \leq 1$ for all cliques $\mathbf{x} \neq 0$. Fortunately such a condition is not restrictive in the context of image analysis, as it is usually undesirable for a scene to contain many overlapping objects. Forbidding overlap altogether, that is, setting $\phi(\mathbf{x}) = 0$ whenever $x_1 \sim x_2$ for some $x_1 \neq x_2 \in \mathbf{x}$, results in a hard core model. For further details as well as a wide range of examples, the reader is referred to Van Lieshout (2000a).

Continuous analogues of the Markov mesh models and connected component fields discussed in Section 13.2.1 exist (Baddeley and Møller 1989, Baddeley *et al.* 1996, Cressie *et al.* 2000, Haggstrom *et al.* 1999) but to date seem not to have found many applications in high level image analysis.

Markov marked point processes are useful modelling tools for scenes composed of objects that do not overlap or are of similar appearance. However, they do not take into account the relative depth — distance to the camera — of objects in the image (Mardia *et al.* 1997) due to the invariance under permutations inherent in their definition, nor can they cope with non-symmetric neighbour relations (Ortner *et al.* 2007). In such cases, finite sequential spatial processes (Van Lieshout 2006a, 2006b) can be used. The realizations of such a process, denoted by $\vec{\mathbf{x}} = (x_1, \ldots, x_n)$ for $n \in$ No and $x_i \in D^- \times Q$, $i = 1, \ldots, n$, consist of vectors of arbitrary length – in contrast to the sets that arise as realizations of marked point processes. An example is given in Fig. 13.6, where for $i < j$, the coloured square x_i lies in the foreground compared to square x_j. The distribution of a sequential spatial process on D^- with marks in Q may be specified by its probability density with respect to the distribution of a random sequence of Poisson length with independent components that are uniformly distributed over D^- and marked in an i.i.d. fashion according to Q. See also Daley and Vere-Jones (2003).

Let $\sim$ be a reflexive relation on $D^- \times Q$, not necessarily symmetric. If $y \sim z$, the object parametrized by the marked point z is said to be a directed neighbour of y. Now, a sequential spatial process Y with probability density f is said to be Markov with respect to $\sim$ if $f(\vec{\mathbf{y}}) > 0$ implies $f(\vec{\mathbf{z}}) > 0$ for all subsequences $\vec{\mathbf{z}}$ of $\vec{\mathbf{y}}$, and, for $u \notin \vec{\mathbf{y}}$, the ratio $f((\vec{\mathbf{y}}, u))/f(\vec{\mathbf{y}})$ depends only on u and its directed neighbours $\{y_i \in \vec{\mathbf{y}} : u \sim y_i\}$ in $\vec{\mathbf{y}}$. Usually the set of neighbours is small compared to the global configuration, a fact that can be employed to advantage in the design of Monte Carlo algorithms. The ratio is related to the sequential conditional intensity

(13.8)

for inserting $u \notin \vec{\mathbf{y}}$ at position $i \in \{1, \ldots, n + 1\}$ of $\vec{\mathbf{y}} = (y_1, \ldots, y_n)$. Here $s_i(\vec{\mathbf{y}}, u) = (y_1, \ldots, y_{i-1}, u, y_i, \ldots, y_n)$. The overall conditional intensity for

end p.436

finding a marked point at u in any position in the vector given that the remain- der of the sequence equals $\vec{\mathbf{y}}$ is given by $\sum_{i=1}^{n+1} \lambda_1(u \mid \vec{\mathbf{y}})$. The expression should be compared to its classic counterpart (13.6).

An analogue of (13.7) holds in the sequential setting. Indeed, for $u \in D^- \times Q$, the sequence $\vec{\mathbf{y}}$ is said to be a u-clique with respect to $\sim$ if $\vec{\mathbf{y}}$ either has length zero or all its components y_i satisfy $u \sim y_i$. The definition is u-directed but otherwise permutation invariant, so we may map $\vec{\mathbf{y}}$ onto the set $\mathbf{y}$ by ignoring the permutation. Similarly, write $\mathbf{y} < i = \{y_1, \ldots, y_{i-1}\}$. Then, a sequential spatial process with probability density f is Markov with respect to $\sim$ if and only if it can be factorized as

(13.9)

for some non-negative interaction function ϕ such that $\phi(u, z) = 1$ if $\mathbf{z}$ is no u-clique with respect to $\sim$. As for marked point

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

processes, when defining a Markov density by its interaction function, integrability must be verified.

Clearly, any sequential spatial process Y immediately defines a classic object process by ignoring the permutation. Alternatively, provided it is integrable, one may define a marked point process X by geometric averaging, that is, by defining a probability density f_X with respect to a unit rate Poisson process on D^{∞} marked in an i.i.d. fashion according to Q by $f_X(\mathrm{x}) \propto \left(\prod_{\overrightarrow{x}} f\left(\overrightarrow{\mathrm{x}}\right)\right)^{1/n(x)!}$ where the product ranges over all permutations $\overrightarrow{\mathbf{x}}$ of object configuration $\mathbf{x}$, and f is a probability density of Y in the sequential setting. In this case, if f admits a Hammersley–Clifford factorisation of the form (13.9), the probability density of X factorizes as

$$f_X(\mathrm{x}) \propto \prod_{\emptyset \neq \mathrm{y} \subseteq \mathrm{x}} \tilde{\phi}(\mathrm{y})$$

where $\tilde{\phi}(y = \{\prod_{y \in \mathrm{y}} \phi(y, \mathrm{y} \setminus \{y\})\}^{1/n(\mathrm{y})}$. It follows that if Y is a pairwise interaction Markov sequential spatial process with respect to some relation $\sim$ having probability density $f(\emptyset)\prod_{i=1}^{n}\prod_{j<i}\phi(y_i, y_j)$ X is a pairwise interaction marked point process with respect to the symmetric relation $\approx$ for which $x \approx y$ if and only if $x \sim y$ or $y \sim x$. When interactions of higher order occur, $\tilde{\phi}(\mathbf{y}) \neq 1$ implies that some $y \in \mathbf{y}$ can be found such that for all $z \in \mathbf{y} \setminus \{y\}$, $y \sim z$. The interesting dual property that any finite sequential spatial process can be derived as the time-ordered vector of points in a classic spatio-temporal marked point process can be shown to hold as well. For further details, see Van Lieshout (2006a, 2006b).

13.3 Properties and connections

The success of Markov random fields in low level vision has given a boost to the development of Markov marked point processes. Indeed, as we saw, mesh or pairwise interaction models have their analogues in point process theory. Similar

end p.437

remarks hold for the intermediate level. For example, Markov connected component fields (13.3) are close in spirit to quermass interaction processes (Kendall *et al.* 1999, Møller and Helisova 2009). Reversely, methods designed for continuous polygonal fields have inspired algorithms for Markov random fields (Schreiber and Van Lieshout 2007).

In this section we consider the intermediate level models in closer detail and study some interesting properties.

13.3.1 Merging of labels

In image segmentation, one has to face the model selection problem of choosing the number of categorical labels, that is, the cardinality of the set Λ. It would be desirable if the local dependence structure of the model would not change with the number of labels. Indeed, the class of Markov connected component fields is closed under merging of labels, but the class of Markov random fields is not (Van Lieshout and Stoica 2009).

To be specific, let Λ be the set $\{1,\ldots, k\}$, $k \geq 2$, and X a k-label Markov connected component field (13.3) with respect to some reflexive, symmetric relation $\sim$ on S. Define the random field Y with labels in $\{0,\ldots,k-2\}$ by $Y_i = X_i \; 1_{X_i \leq k-2}$. Then Y is a $(k-1)$-label Markov connected component field with respect to $\sim$. To see this, fix $y = (y_1,\ldots,y_m) \in \{0,\ldots,k-2\}^S$. For $x \in \{z \in \Lambda^S : y_i = z_i \, 1_{z_i \leq k-2}\}$ and $j \in \Lambda$, write $\mathbf{K}_j(x)$ for the set of maximal connected components in x labelled j. Note that the maximal j-labelled connected components in x and y are identical for $j = 1,\ldots, k-2$ and that each $(k-1)$- or k-component is part of a single 0-component in y. Hence, $\mathbf{P}\{Y_1 = y_1,\ldots, Y_m = y_m\}$ is proportional to

which proves the claim and gives an expression for $\Psi_K(0)$. It is interesting to note that the above result can be seen as the discrete analogue of the fact that the class of Markov connected component point processes (Baddeley and Møller 1989) is closed under independent superposition (Chin and Baddeley 1999).

For Markov random fields, the situation is more complicated, as could be expected from the behaviour of Markov point processes with respect to superposition (Van Lieshout 2000b). By the above remarks, for Markov random fields such as the Potts model that can also be factorized as in (13.3), merging labels $k-1$ and k leads to a Markov connected component field

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

Y. However, Y is not necessarily a Markov random field with respect to the given relation. A counterexample is the Potts model with three colours, $\beta = 1$, and the horizontal and vertical neighbours relation: in configurations such as $\begin{pmatrix} 0 & ? \\ 0/1 & 0 \end{pmatrix}$ the conditional distribution of the top right pixel depends on whether Y takes value 0 or 1 at the bottom left pixel. Furthermore, since the class of Markov random fields is

end p.438

not contained in that of Markov connected component fields with respect to the same relation, merging labels does not always lead to a Markov connected component field. A counterexample is the 6-label Geman and Reynolds (1992) field X on $S = \{s_1, s_2\}$ with $s_1 \sim s_2$ defined by its joint probability mass function $\mathbf{P}\{X_1 = x_1, X_2 = x_2\} \propto \exp\{1/(|x_1 - x_2| + 1)\}$. Then clearly X is a pairwise interaction Markov random field, but, by arguments similar to those in Example A2 in Møller and Waagepetersen (1998), the joint probability mass function of Y does not factorize over maximal connected components.

13.3.2 The length-weighted Arak field is a hard core object process

In this section, we will consider the length-weighted Arak field with empty boundary condition introduced in Section 13.2.2. One can show that for β large enough the model is a Markov object process in the sense of Section 13.2.3 and give a convenient algorithmic description for the mark distribution.

To do so, let C_n be the set of all closed polygons in $[-n,n]^2$ which do not touch the boundary and set $\mathbf{C} = \bigcup_n C_n$. To each polygon θ, attach a unique $r(\theta)$, say the extreme lower left point, to serve as location parameter. Each polygon lying in the image window D as in the discussion surrounding (13.5), including the empty one, can than be described as $x + \theta$ with $x \in D$ and $\theta \in Q$ defined as $Q = \{\theta \in \mathbf{C} \mid r(\theta) = 0\} \cup \{\varnothing\}$.

For fixed $\beta \geq 2$, define a mark distribution Q by the continuous time random walk representation of Schreiber (2006) as follows. Fix Z_0 at 0, pick an initial direction uniformly in $(0, 2\pi)$, and, between direction updates, move in a constant direction with unit speed. Directions are updated with rate 4; the angle between the old and new direction is chosen according to the probability density $p(\phi) = |\sin(\phi)|/4$ on $(0, 2\pi)$, mirroring properties of the Poisson line process L discussed in Section 13.2.2. The resulting random walk $Z_t, t \geq 0$, is killed with rate $\beta - 2$ and whenever it hits its past trajectory to obtain $\tilde{Z}_t$, $t \geq 0$. Also draw a loop- closing half-line l^* through 0 whose angle with the initial segment of $(Z_t)_{t \geq 0}$ is distributed according to $p(\phi)$. Finally, if $\tilde{Z}_t$ hits l^* before being killed and the contour θ^* formed by l^* and the trajectory of $\tilde{Z}_t$ up to the moment of hitting has lower left extreme point 0, let e^* be the segment along l^* from the origin up to its intersection with $\tilde{Z}_t$, and generate a contour as follows:

- with probability $\exp\{-(\beta + 2)\, l(e^*)\}$, output θ^*;
- otherwise, output $\varnothing$

In all remaining cases, the empty polygon is outputted. The procedure is illustrated in Fig. 13.5.

The above algorithmic definition of Q is easy to implement. Moreover, the union of polygonal contours of the marked point process defined by its conditional intensity

end p.439

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

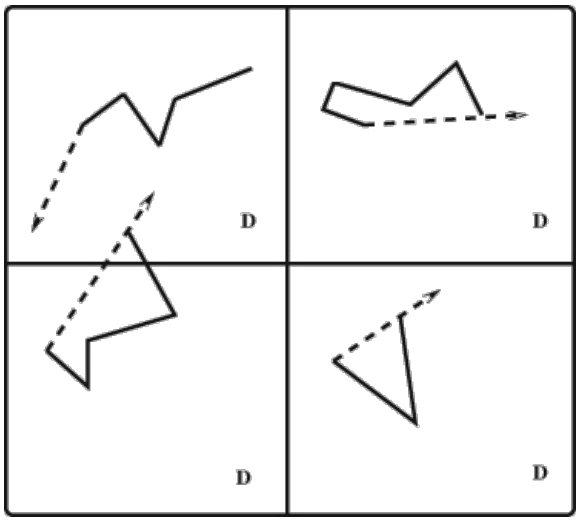

Fig. 13.5. Random walk construction. The dashed line is the loop-closing half— line, the solid one $\tilde{Z}_t$. The bottom pictures may output a non-empty polygon. The conditional intensity for the bottom left polygon is zero regardless of the remainder of the configuration since it does not lie entirely in *D*.

with respect to the product of Lebesgue measure on D and Q on Q coincides in distribution with the length-weighted Arak field (Van Lieshout and Schreiber 2007). Hence, one may regard it as a hard core object process of intensity 4π.

13.4 Examples

To illustrate the theory reviewed above, two image analysis problems are presented. The first example concerns the task of tracking a variable number of moving objects across a video sequence, the second one is devoted to foreground/background separation. In both cases, a regularization framework is chosen in which a goodness of fit term between the data and its semantic description is balanced by terms that favour smooth, sparse descriptors.

13.4.1 Motion tracking

Object tracking is important, as motion is a major source of semantic information in domains such as surveillance, robotics, and depth estimation. The aim is to find the objects that appear in a video sequence, and to follow their movements over time. Clearly, this is a high level task, hence we place ourselves in the framework of Section 13.2.3.

Consider the synthetic data given in Fig. 13.6. The objects are squares of various sizes and colours. Hence $Q = [s_{\min}, s_{\max}] \times \{0,\ldots, 255\}$ where the side

end p.440

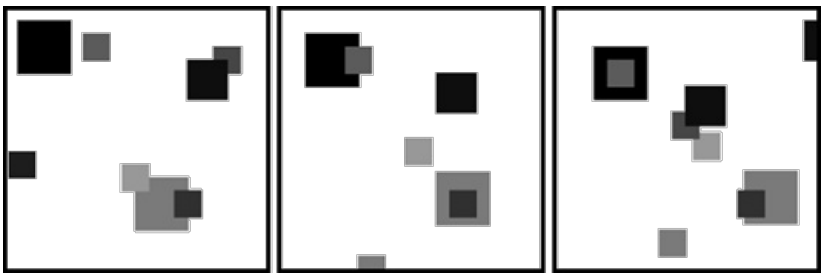

Fig. 13.6. Video sequence of moving squares.

length lies in some interval with $0 < s_{\min} < s_{\max} < \infty$, and the colour space is the set of 8-bit grey levels. We use the top left corner as our location parameter. Each square x occupies a set of pixels $R(x)$ in the digital image. As the object colour is not constant and there is significant overlap, we take vectors of squares to describe a scene. The components with a low index are those close to the camera. Doing so, we may define the signal $\theta_t(\vec{\mathbf{x}}^i)$ of configuration $\vec{\mathbf{x}}^i$ in video frame i at

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

pixel t as the colour of the object closest to the camera if $t \in \cup R(x_j^i)$ and as the background colour otherwise.

In order to formulate object tracking as a statistical inference problem, we shall define a probability density $f(\mathbf{x}) \propto \exp\{-U(\mathbf{x})\}$ at the sequence of object configurations $\mathbf{x} = (x^{\rightarrow 1}, \mathbf{x}^{\rightarrow 2}, \mathbf{x}^{\rightarrow 3})$ with respect to the product distribution of independent unit rate sequential Poisson processes marked by a colour chosen according to the data histogram and a uniformly chosen side length. For notational convenience, we suppress the dependence on the data. The unknown configuration $\mathbf{x}$ is now treated as a parameter to be estimated, in other words, the aim is to find $\mathbf{x}$ that minimise the energy $U(\mathbf{x})$.

Upon observation of the video sequence $\mathrm{y} = (\mathrm{y}^1, \mathrm{y}^2, \mathrm{y}^3) = (y_t^i)$ indexed by pixels $t \in T$ and image frames $i = 1,2,3$, the goodness of fit is taken to be proportional to

(13.10)

Note that the probability density function f defined by energy function U given by (13.10) has independent frame marginals that are Markov sequential spatial processes with respect to the overlapping objects relation $u \sim v$ if and only if $R(u) \cap R(v) \neq 0$. Optimisation of (13.10) over $\mathbf{x}$ amounts to a least absolute deviation regression. However, a minimum is not guaranteed to exist, nor, if it does exist, to be unique. Indeed, small spurious squares behind the signal of those closer to the camera (having a lower index) do not affect the goodness of fit.

end p.441

To overcome this problem, we add terms to the energy function that penalize overlap, favour temporal cohesion between squares in subsequent frames, and include matchings to keep track of a square as it moves across the video frames. To prevent spurious overlap, we add a positive penalty for each object and each pair of overlapping squares (known as the Strauss potential). Note that doing so results in a Markov overlapping objects process. Furthermore, we say that x_k^i and x_l^{i+1} are matched if they are instances of the same square in adjacent frames. The quality of a match is described by a weighted sum of the absolute difference in side lengths and colour and the squared distance between the top left points of the two squares involved. In order to quantify coherence between consecutive frames, we add a positive penalty term for each missing match, offset by the quality of matches that are present. Finally, in order to propagate relative depth information gathered when two squares overlap over time, a positive penalty is added to U for each overlapping pair of matched squares such that the relative depth is not preserved in the next or previous frame. Again, the dependence structure is local in the sense that only pairs of overlapping objects and consecutive video frames are involved.

To find an optimal configuration sequence with associated matchings, we use simulated annealing (Haario and Saksman 1991) within the Metropolis—Hastings framework (see Chapter 9). Inspired by Lund *et al.* (1999), we allow addition and deletion of matched or unmatched squares, modification of the permutation order, addition and deletion of a match, changes in location, colour, and side length, splitting a square in two, and merging two close squares into a single one. The synthetic data presented in Fig. 13.6 were chosen to include squares entering and leaving the image window or passing each other, a square leaving one connected component to join another, complete occlusion, as well as varying contrast. The signal of the near optimal sequence of configurations coincides with that of the data. Not visible in the signal but present in the near optimal configuration is a square hidden behind the dark one in the top right quadrant of the middle frame, see the listed sequence of object configurations with associated matchings in Table 13.1. The matches are correctly reproduced as well.

In order to quantify depth, we estimate the probabilities p_{kl}^i of square k in frame i lying closer to the camera than square $l > k$ in the same frame. Table 13.2 lists the result for $i = 2$. By standard combinatorial arguments, the correctness of these empirical values is verified. Other examples and further details are given in Van Lieshout (2008).

13.4.2 Foreground/background separation

Consider an image such as the one depicted in Fig. 13.7 and suppose the goal is to separate the foreground from the background. For this task, a large variety of models and methods is available, ranging from elementary thresholding through contour extraction methods to scene modelling. For a comprehensive overview, see for example Chapter 10 in Rosenfeld and Kak (1982). More recent material

end p.442

Table 13.1. Result after simulated annealing (see text) for Fig. 13.6. The columns contain the square parameters (*x* and *y* coordinates of top left corner, side length, and grey level) for frames *i* = 1, 2, 3. For

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

matched objects, the index of the associated square in the next frame is listed in the columns labelled '*m*'.

i=1				*m*	*i*=2				*m*	*i*=3			
89	119	20	180	7	119	49	30	30	1	99	59	30	30
129	139	20	80	6	49	29	20	127	4	89	79	20	105
1	109	20	55		59	189	20	155	9	190	9	30	30
99	129	40	155	8	19	19	40	0	8	39	39	21	127
139	39	30	30	1	125	57	20	105	2	139	139	20	80
59	19	20	127	2	129	139	20	80	5	144	124	40	155
9	9	40	0	4	95	99	20	180	7	105	95	20	180
159	29	20	105	5	119	125	40	155	6	29	29	40	0
										79	169	20	155

can be found in Vincent (1999), which treats morphological methods, or in the volume edited by Osher and Paragios (2003) which contains contributions on level set approaches, active contour models and variational methods, and more. In this chapter on stochastic geometric methods, we place ourselves at the intermediate conceptual level that regards the sought after segmentation as a coloured tessellation. The approach is a compromise between the full scene modelling feasible in restricted application domains only and the Markov random field models of Section 13.2.1.

As discussed in Section 13.2.2, at the intermediate level one may either choose sets of pixels or tessellations as the focus of interest. Since reality is continuous rather than discrete, we find the latter approach the most appealing and will use polygonal field models. The idea to use such a model as a prior distribution for image segmentation is due to Clifford and Middleton (1989). A Metropolis— Hastings style sampler was developed by Clifford and Nicholls (1994) and applied

Fig. 13.7. Segmented image of a cat.

end p.443

Table 13.2. Pairwise probabilities p_{kl}^2 of object *k* having a lower sequence index than object *l* after annealing (see text) for frame 2 and data as given in Fig. 13.6.

—	0.66	0.79	0.93	1.00	0.75	1.00	1.00
	—	0.67	1.00	0.62	0.59	0.82	0.95
		—	0.67	0.42	0.42	0.62	0.83
			—	0.22	0.24	0.43	0.71
				—	0.50	1.00	1.00
					—	0.75	1.00
						—	1.00
							—

to an image reconstruction problem. A modification was suggested by Paskin and Thrun (2005) for use in robotic mapping.

In order to formulate foreground extraction as a statistical inference problem, we shall define a probability density $f(\hat{\gamma}) \propto \exp\{-U(\hat{\gamma})\}$ on labelled tessellations $\hat{\gamma} \in \mathcal{A}_d$ with respect to the coloured Arak field $\hat{A}_d$ with free boundary conditions. The notation $\hat{A}_d$ is used to distinguish the dominating measure from its colourless counterpart A_d, cf. (13.5), with a similar remark for $\mathcal{A}_d$. The labels are binary: 1 indicates the foreground, 0 the background. Since adjacent regions must have different labels, each admissable γ can be coloured in two ways only. In the reference field $\hat{A}_d$, each possible colouring is given probability 1/2. Note that the energy may and will depend on the data image. Having formulated a probability density *f*, the unknown configuration $\hat{\gamma}$ is treated as a parameter to be estimated, in other words, the aim is to minimize the energy *U*

New Perspectives in Stochastic Geometry
Kendall, Wilfrid S. (Editor), Department of Statistics, University of Warwick, UK
Molchanov, Ilya (Editor), Department of Mathematical Statistics and Actuarial Science, University of Bern, Switzerland
Print publication date: 2009, Published to Oxford Scholarship Online: February 2010
Print ISBN-13: 978-0-19-923257-4, doi:10.1093/acprof:oso/9780199232574.001.0001

$(\hat{\gamma})$ over the family of admissible labelled tessellations.

Write $\mathbf{y} = (y_t)$, $t \in T$, for the data image. Here, each pixel value $y_t \in \{0,\ldots, 255\}$ is a grey level. Given $\mathbf{y}$, the energy function U is the sum of regression and regularization terms. In analogy to (13.10), the goodness of fit is described in terms of $L_1(\mathbf{y}, \hat{\gamma})$. Optimization of the goodness of fit criterion alone suffers from similar problems as those discussed in the context of motion tracking: In general, the minimum will not be unique and tend to result in an over-segmentation. To overcome these problems, we add a regularization term proportional to the total edge length $\ell(\gamma)$.

To find the optimal configuration, we again use simulated annealing within the Metropolis—Hastings framework. The update proposals are based on the dynamic representation of Arak (1982), an equivalent definition of A_d in terms of the evolution of a one-dimensional particle system. Briefly, the edges of the field that separate foreground and background components are interpreted as the traces of particles travelling in time and one-dimensional space. Particles are born in D in pairs according to a homogeneous Poisson point process, while at the boundary ∂D, single particles are generated according to a Poisson point process with an appropriate intensity measure. Particles change their velocity

end p.444

according to a pure jump process with a suitably chosen transition kernel, which show up as the vertices of γ in the trace. When particles collide or reach the boundary of D, they die. For details, see Arak and Surgailis (1989). Schreiber (2005)observed that when a new particle birth site $u \in D$ is added to configuration γ, the symmetric difference with the old configuration is a closed polygonal curve that may be self-intersecting or chopped off at the boundary. The addition of a birth site $u \in \partial D$ gives rise to a single self-avoiding polygonal curve that may be chopped off at the the boundary. Hence, a Metropolis—Hastings sampler can be set up that proposes to add or delete birth sites combined with swaps between foreground and background. The resulting moves are easy to implement and sufficiently flexible for an efficient exploration of the state space (Kluszczyński *et al.* 2007, Schreiber 2005).

The result for an image from the Pascal Network of Excellence challenge 2006 http://www.pascal-network.org/challenges/VOC/thumbs/VOC2006 is overlaid on the data in Fig. 13.7. The misclassification rate, calculated manually due to the lack of a ground truth, is about 3%. Further details can be found in Kluszczyński *et al.* (2006).

Acknowledgements

Much of the work presented here was done in collaboration with others. Special thanks are due to all co-authors, as well as to B. Lisser and A. Steenbeek for programming assistance.

References

Abend, K., Harley, T. J., and Kanal, L. N. (1965). Classification of binary random patterns. *IEEE Trans. Information Theory*, **11**, 538–544. Link

Amit, Y., Grenander, U., and Piccioni, M. (1991). Structural image restoration through deformable templates. *J. Amer. Statist. Assoc.*, **86**, 376–387. Link

Arak, T. (1982). On Markovian random fields with finite number of values. *4th USSR-Japan Symposium on Probability Theory and Mathematical Statistics, Abstracts of Communications*, Tbilisi.

Arak, T. and Surgailis, D. (1989). Markov fields with polygonal realizations. *Probab. Theory Related Fields*, **80**, 543–579. Link

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

Baddeley, A. J. and Lieshout, M. N. M. van (1992). Object recognition using Markov spatial processes. In *Proceedings 11th IAPR International Conference on Pattern Recognition*, B, 136–139. IEEE Computer Society Press, Los Alamitos.

Baddeley, A. J. and Lieshout, M. N. M. van (1993). Stochastic geometry models in high-level vision. In *Statistics and Images, Volume 1*, K. V. Mardia and G. K. Kanji (Eds.), *Advances in Applied Statistics*, a supplement to *Journal of Applied Statistics* , **20**, 231–256. Carfax, Abingdon.

end p.445

Baddeley, A. J. and Lieshout, M. N. M. van (1995). Area-interaction point processes. *Ann. Inst. Statist. Math.*, **46**, 601–619.

Baddeley, A. J. and Møller, J. (1989). Nearest-neighbour Markov point processes and random sets. *Internat. Statist. Rev.*, **57**, 89–121. Link

Baddeley, A. J., Lieshout, M. N. M. van, and Møller, J. (1996). Markov properties of cluster processes. *Adv. in Appl. Probab.*, **28**, 346–355. Link

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. Roy. Statist. Soc. Ser. B*, **36**, 192–236.

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

Besag, J. and Kooperberg, C. (1995). On conditional and intrinsic autoregressions. *Biometrika*, **82**, 733–746.

Besag, J., York, J., and Mollié, A. (1991). Bayesian image restoration, with two applications in spatial statistics (with discussion). *Ann. Inst. Statist. Math.*, **43**, 1–59. Link

Blackwell, P. G. and Møller, J. (2003). Bayesian analysis of deformed tessellation models. *Adv. in Appl. Probab.*, **35**, 4–26. Link

Chin, Y. C. and Baddeley, A. J. (1999). On connected component Markov point processes. *Adv. in Appl. Probab.*, **31**, 279–282. Link

Clifford, P. (1990). Markov random fields in statistics. In *Disorder in Physical Systems. A Volume in Honour of J. M. Hammersley*, G. R. Grimmett and D. J. A. Welsh (Eds.), pages 19–32. Oxford University Press, New York.

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

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

Cressie, N. and Davidson, J. L. (1998). Image analysis with partially ordered Markov models. *Comput. Statist. Data Anal.*, **29**, 1–26. Link

Cressie, N., Zhu, J., Baddeley, A. J., and Nair, M. G. (2000). Directed Markov point processes as limits of partially ordered Markov models. *Methodol. Comput. Appl. Probab.*, **2**, 5–21. Link

Daley, D. J. and Vere—Jones, D. (2003). *An Introduction to the Theory of Point Processes. Volume I. Elementary Theory and Methods* (2nd edn). Springer- Verlag, New York.

Dougherty, E. R. (1992). *An Introduction to Morphological Image Processing*. SPIE Press, Bellingham.

Geman, D. (1990). Random fields and inverse problems in imaging. *École d'été de Probabilités de Saint-Flour XVIII – 1988, Lecture Notes in Mathematics*, **1427**, 113–193. Springer-Verlag, Berlin. Link

Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721–741. Link

end p.446

Geman, D. and Reynolds, G. (1992). Constrained restoration and the recovery of discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**, 367–383. Link

Geman, D., Geman, S., Graffigne, C., and Dong, P. (1990). Boundary detection by constrained optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**, 609–628. Link

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. Link OUP Resource

Grenander, U. (1976). *Lectures on Pattern Theory, Vol. 1: Pattern Synthesis*. Applied Mathematical Sciences vol. 18. Springer-Verlag, New York.

Grenander, U. (1978). *Lectures on Pattern Theory, Vol. 2: Pattern Analysis*. Applied Mathematical Sciences vol. 24. Springer-Verlag, New York.

Grenander, U. (1981). *Lectures on Pattern Theory, Vol. 3: Regular Structures*. Applied Mathematical Sciences vol. 33. Springer-Verlag, New York.

Haario, H. and Saksman, E. (1991). Simulated annealing process in general state space. *Adv. in Appl. Probab.*, **23**, 866–893. Link

Häggström, O., Lieshout, M. N. M. van, and Møller, J. (1999). Characterization results and Markov chain Monte Carlo algorithms including exact simulation for some spatial point processes. *Bernoulli*, **5**, 641–658. Link

Hansen, M. B., Møller, J., and Tøgersen, F. A. (2002). Bayesian contour detection in a time series of ultrasound images through dynamic deformable template models. *Biostatistics*, **3**, 213–228. Link OUP Resource

Heikkinen, J. and Arjas, E. (1998). Non-parametric Bayesian estimation of a spatial Poisson intensity. *Scand. J. Statist.*, **25**, 435–450. Link

Hurn, M. A. (1998). Confocal fluorescence microscopy of leaf cells: an application of Bayesian image analysis. *J. Roy. Statist. Soc. Ser. C*, **47**, 361–377. Link

Kendall, W. S., Lieshout, M. N. M. van, and Baddeley, A. J. (1999). Quermass- interaction processes: Conditions for stability. *Adv. in Appl. Probab.*, **31**, 315–342. Link

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

Künsch, H. R. (1987). Intrinsic autoregressions and related models on the two- dimensional lattice. *Biometrika*, **74**, 517–524.

Lacoste, C., Descombes, X., and Zerubia, J. (2005). Point processes for unsu- pervised line network extraction in remote sensing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **27**, 1568–1579. Link

Lacroix, V. (1987). Pixel labelling in a second-order Markov mesh. *Signal Processing*, **12**, 59–82. Link

Lebowitz, J. L. and Gallavotti, G. (1971). Phase transitions in binary lattice gases. *J. Math. Phys.*, **12**, 1129–1133. Link

end p.447

Lieshout, M. N. M. van. (1994). Stochastic annealing for nearest-neighbour point processes with application to object recognition. *Adv. in Appl. Probab.*, **26**, 281–300. Link

Lieshout, M. N. M. van. (1995). *Stochastic Geometry Models in Image Analysis and Spatial Statistics, CWI Tract*, **108**. CWI, Amsterdam.

Lieshout, M. N. M. van. (2000a). *Markov Point Processes and their Applications*. Imperial College Press/World Scientific Publishing, London/ Singapore.

Lieshout, M. N. M. van. (2000b). Propagation of spatial interaction under superposition. In *Accuracy 2000, Proceedings of the 4th Inte rna tion al Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences*, G. B. M. Heuvelink and M. J. P. M. Lemmens (Eds.), pp. 687– 694. Delft University Press, Delft.

Lieshout, M. N. M. van. (2006a). Markovianity in space and time. In *Dynamics and Stochastics: Festschrift in Honour of Michael Keane*, D. Denteneer, F. den Hollander, and E. Verbitskiy (Eds.), *Lecture Notes Monograph Series*, 48, 154–167. Institute for Mathematical Statistics, Beachwood.

Lieshout, M. N. M. van. (2006b). Campbell and moment measures for finite sequential spatial processes. In *Proceedings Prague Stochastics 2006*, M. Hušková and M. Janžura (Eds.), pages 215–224. Matfyzpress, Prague.

Lieshout, M. N. M. van. (2008). Depth map calculation for a variable number of moving objects using Markov sequential object processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **30**, 1308– 1312. Link

Lieshout, M. N. M. van and Schreiber, T. (2007). Perfect simulation for length- interacting polygonal Markov fields in the plane. *Scand. J. Statist.*, **34**, 615–625. Link

Lieshout, M. N. M. van and Stoica, R. S. (2009). A note on pooling of labels in random fields. Research Report PNA-E0906, CWI, Amsterdam.

Lund, J., Penttinen, A., and Rudemo, M. (1999). Bayesian analysis of spatial point patterns from noisy observations. Research Report, Department of Mathematics and Physics, The Royal Veterinary and Agricultural University, Copenhagen.

Mardia, K. V. and Kanji, G. K. (Eds.) (1993). *Statistics and Images, Volume 1, Advances in Applied Statistics*, a supplement to *Journal of Applied Statistics* , **20**. Carfax, Abingdon.

Mardia, K. V., Qian, W., Shah, D., and De Souza, K. M. A. (1997). Deformable template recognition of multiple occluded objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**, 1036– 1042.

Matheron, G. (1975). *Random Sets and Integral Geometry* . John Wiley and Sons, New York.

end p.448

Molina, R. and Ripley, B. D. (1989). Using spatial models as priors in astronomical image analysis. *J. Appl. Stat.*, **16**, 193–206. Link

Møller, J. and Helisova, K. (2008). Power diagrams and interaction processes for unions of discs. *Adv. in Appl. Probab.*, **40**, 321 –347. Link

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

Møller, J. and Waagepetersen, R. P. (1998). Markov connected component fields. *Adv. in Appl. Probab.*, **30**, 1–35. Link

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

Okabe, A., Boots, B., Sugihara, K., and Chiu, S. N. (2000). *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams* (2nd edn). John Wiley and Sons, Chichester.

Ortner, M., Descombes, X., and Zerubia, J. (2007). Building outline extraction from digital elevation models using marked point processes. *International Journal of Computer Vision* , **72**, 107–132. Link

Osher, S. and Paragios, N. (Eds.) (2003). *Geometric Level Set Methods in Imaging, Vision, and Graphics* . Springer, New York.

Papangelou, F. (1974). The conditional intensity of general point processes and an application to line processes. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, **28**, 207–226. Link

Paskin, M. A. and Thrun, S. (2005). Robotic mapping with polygonal random fields. In *Proceedings in Artificial Intel ligence UAI-05*.

Pievatolo, A. and Green, P. J. (1998). Boundary detection through dynamic polygons. *J. Roy. Statist. Soc. Ser. B*, **60**, 609–626. Link

Qian, W. and Titterington, D. M. (1991). Multidimensional Markov chain models for image textures. *J. Roy. Statist. Soc. Ser. B*, **53**, 661–674.

Ripley, B. D. (1988). *Statistical Inference for Spatial Processes* . Cambridge University Press, Cambridge.

Ripley, B. D. and Kelly, F. P. (1977). Markov point processes. *J. London Math. Soc.*, **15**, 188–192. Link

Ripley, B. D. and Sutherland, A. I. (1990). Finding spiral structures in images of galaxies. *Philosophical Transactions of the Royal Society of London, Series A*, **332**, 477–485. Link

Rosenfeld, A. and Kak, A. C. (1982). *Digital Picture Processing* (2nd edn). Academic Press, Orlando.

Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications, Monographs on Statistics and Applied Probability*, **104**. Chapman & Hall/CRC, Boca Raton.

Rue, H. and Hurn, M. A. (1999). Bayesian object identification. *Biometrika*, **86**, 649–660. Link OUP Resource

end p.449

Rue, H. and Husby, O. K. (1998). Identification of partly destroyed objects using deformable templates. *Statistics and Computing*, **8**, 221–228. Link

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

Schreiber, T. (2006). Dobrushin—Kotecký—Schlosman theorem for polygonal Markov fields in the plane. *J. Statist. Phys.*, **123**, 631–684. Link

Schreiber, T. and Lieshout, M. N. M. van. (2007). Disagreement loop and path creation/annihilation algorithms for planar Markov fields with applications to image segmentation. EURANDOM Report 2007-045.

Serra, J. (1982). *Image Analysis and Mathematical Morphology*. Academic Press, London.

Skare, Ø., Møller, J., and Jensen, E. B. V. (2007). Bayesian analysis of spatial point processes in the neighbourhood of Voronoi networks. *Statistics and Computing*, **17**, 369–379. Link

Stoica, R., Descombes, X., and Zerubia, J. (2004). A Gibbs point process for road extraction in remotely sensed images. *International Journal of Computer Vision* , **57**, 121–136. Link

Stoica, R. S., Descombes, X., Lieshout, M. N. M. van, and Zerubia, J. (2002). An application of marked point processes to the extraction of linear networks from images. In *Spatial Statistics: Case Studies*, J. Mateu and F. Montes (Eds.), pages 287–312. WIT Press, Southampton.

Stoica, R. S., Martinez, V., and Saar, E. (2007). A three dimensional object point process for detection of cosmic filaments. *J. Roy. Statist. Soc. Ser. C*, **56**, 459–477. Link

Stoyan, D., Kendall, W. S., and Mecke, J. (1995). *Stochastic Geometry and its Applications* (2nd edn). John Wiley and Sons, Chichester.

Tjelmeland, H. and Besag, J. (1998). Markov random fields with higher-order interactions. *Scand. J. Statist.*, **25**, 415–433. Link

Tjelmeland, H. and Holden, L. (1993). Semi-Markov random fields. In *Geostatistics Troia '92*, A. Soares (Ed.). Kluwer, Amsterdam.

Vincent, L. (1999). Current topics in applied morphological image analysis. In *Stochastic Geometry, Likelihood and Computation*, O. Barndorff—Nielsen, W. S. Kendall and M. N. M. van Lieshout (Eds.), pages 199–283. CRC Press/Chapman and Hall, Boca Raton.

Widom, B. and Rowlinson, J. S. (1970). A new model for the study of liquid- vapor phase transitions. *J. Chem. Phys.*, **52**, 1670 –1684. Link

Winkler, G. (2003). *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods. A Mathematical Introduction* (2nd ed), *Applications of Mathematics, Stochastic Modelling and Applied Probability*, **27**. Springer- Verlag, Berlin.

end p.450

This page intentionally left blank.