

Thinning spatial point processes into Poisson processes

Jesper Møller¹ and Frederic Paik Schoenberg²

¹ Department of Mathematical Sciences, Aalborg University, Fredrik Bajers Vej 7G, DK-9220 Aalborg, Denmark.

² Department of Statistics, 8125 Math-Science Building, University of California, Los Angeles, CA 90095–1554, USA.

Abstract

This paper describes methods for randomly thinning certain classes of spatial point processes. In the case of a Markov point process, the proposed method involves a dependent thinning of a spatial birth-and-death process, where clans of ancestors associated with the original points are identified, and where one simulates backwards and forwards in order to obtain the thinned process. In the case of a Cox process, a simple independent thinning technique is proposed. In both cases, the thinning results in a Poisson process if and only if the true Papangelou conditional intensity is used, and thus can be used as a diagnostic for assessing the goodness-of-fit of a spatial point process model. Several examples, including clustered and inhibitive point processes, are considered.

Keywords: area-interaction point process, clans of ancestors, coupling, Cox process, dependent and independent thinning, Markov point process, Papangelou conditional intensity, Poisson process, spatial birth-and-death process, Thomas process.

AMS Classifications: 60G55 (point processes), 62M30 (spatial processes).

1 Introduction

A useful method for evaluating a fitted model $\hat{\lambda}$ for the conditional intensity λ of a temporal or space-time point process, where the conditioning is based on all observations at all previous times, is via random thinning. This technique, described in Schoenberg (2003), involves keeping each point τ_i independently with a probability proportional to $\hat{\lambda}(\tau_i)$. If the true conditional intensity λ is used in place of $\hat{\lambda}$, then the resulting process is a homogeneous Poisson process (Schoenberg, 2003). The residual points, obtained after thinning using the estimate $\hat{\lambda}$ in place of λ , may readily be inspected for being a homogeneity homogeneous Poisson process using standard methods.

The question of how to extend this technique to the case of a purely spatial point process remains open. For example, Markov (or Gibbs) point processes are characterized by their Papangelou conditional intensities, where the conditioning is based on the observations at all other locations, see e.g. Ripley and Kelly (1977) and Baddeley and Møller (1989). While Schoenberg (2005) conjectured that a thinning method similar to that used for space-time point processes should be valid for purely spatial point processes as well, this conjecture was shown by Baddeley et al. (2005, page 664) to be false. Schoenberg and Zhuang (2008) introduced a method for thinning spatial point processes based on considering all possible subsets of points and selecting among these subsets with the appropriate probability, but this method relies on certain rather restrictive assumptions and also, since it requires $O(2^n)$ computations for thinning a point pattern consisting of n points, is excessively computationally intensive and impractical for all but very small values of n . Furthermore, the method of Schoenberg and Zhuang (2008) only applies to spatial point processes with known Pa-

pangelou intensities, but not to many important classes of point processes, including Cox processes (Cox, 1955), whose Papangelou intensities may be intractable.

Here, we introduce alternative methods for thinning a spatial point process X . For simplicity and specificity, we assume that X is finite, simple (i.e. has no multiple points), and defined on a bounded Borel set $S \subset \mathbb{R}^k$ ($k \in \{1, 2, \dots\}$) so that X can be considered as a random subset of S . This setting covers most cases of practical interest, but our methods can easily be extended both to non-simple Markov point processes defined on a general state space and using an exponential state space setting (Carter and Prenter, 1972; Preston, 1977; Ripley and Kelly, 1977), and to non-simple Cox processes which may be defined on \mathbb{R}^k or more complicated state spaces. For background material on spatial point processes, particularly Markov point processes and Cox processes, see Møller and Waagepetersen (2004, 2007) and the references therein.

Section 2 considers the case where X is a Markov point process and $\lambda(x, \xi)$ denotes its Papangelou conditional intensity, which is assumed to be bounded from below by a non-negative deterministic function $\rho(\xi)$ for all finite point configurations $x \subset S$ and points $\xi \in S \setminus x$. We consider $X = X_0$ as an equilibrium state of a spatial birth-and-death process X_t at time 0, and we consider coupling this with another spatial birth-and-death process W_t with equilibrium distribution given by a Poisson process with intensity function ρ such that $W_t \subseteq X_t$ for all times $t \in \mathbb{R}$. The proposed method involves first simulating backwards X_t for $t \leq 0$. At the i th jump backwards in time, with respect to the neighbour relation used for defining the Markov point process, we identify the i th generation ancestors associated with the points in X , and the simulation is stopped the first time T before time 0 such a generation is empty. Second, $W_T = \emptyset$ is the empty point configuration, and we use a dependent thinning

technique to obtain W_t for $T \leq t \leq 0$. This procedure has some similarity to perfect simulation algorithms for spatial point processes (Kendall, 1998; Kendall and Møller, 2000; Fernández, Ferrari, and Garcia, 2002) but in contrast these algorithms assume an upper bound $V(\xi) \geq \lambda(x, \xi)$ (the so-called local stability condition) and uses this to couple X_t with a birth-and-death process D_t such that $X_t \subseteq D_t$, $t \in \mathbb{R}$, and each D_t is a Poisson process with intensity function V . In fact, our procedure needs only first to simulate the jump chain of births and deaths for X_t with t running from time 0 to time T , second what happens at these jump times in W_t with t running from T to 0, and finally return W_0 . We demonstrate how this works for an area-interaction point process (Baddeley and Van Lieshout, 1995).

Section 3 considers the case where instead $\lambda(\xi)$ is a non-negative random function, X is a Cox process driven by λ , and λ is bounded from below by a non-negative deterministic function ρ . This case is much simpler, since X can be viewed as a superposition of a Cox process driven by $\lambda - \rho$ and an independent Poisson process with intensity function ρ . We show how an independent thinning technique applies to obtain a Poisson process with intensity function ρ . This is exemplified in the case of a modified Thomas process (Thomas, 1949).

In both cases of a Markov point process and a Cox process model, the thinning results in a Poisson process with intensity function ρ if and only if the true λ is used, and thus can be used as a diagnostic for assessing the goodness-of-fit of a spatial point process model. Sections 2.4 and 3.3 show specific examples, using the L -function to check if the thinned process is Poisson. The L -function is a standardized version of Ripley's K -function (Ripley, 1976, 1977) and both can be extended to the inhomogeneous case (Baddeley, Møller and Waagepetersen, 2000), though many other techniques for checking whether a point process

is Poisson have been developed; see e.g. Cressie (1993) and Stoyan, Kendall and Mecke (1995). In practical applications, in both Bayesian and frequentist settings, λ is typically replaced by an estimate $\hat{\lambda}$ (for a discussion on how such estimates can be obtained, see Møller and Waagpetersen (2004, 2007) and the references therein), but for simplicity and specificity, in all examples we consider simulated data where the true λ is known and hence is not estimated.

2 Thinning Markov point processes

Suppose that X has a density f with respect to the homogeneous Poisson process on S with intensity equal to unity. We assume that f is hereditary, that is, $f(x \cup \{\xi\}) > 0$ whenever $f(x) > 0$ for a finite point configuration $x \subset S$ and a point $\xi \in S \setminus x$. The Papangelou conditional intensity (Kallenberg, 1984) is defined by

$$\lambda(x, \xi) = f(x \cup \{\xi\})/f(x)$$

where the hereditary condition ensures that f and λ are in a one-to-one correspondence. (Here and throughout this paper, we use the convention that $0/0 = 0$.) Heuristically, $\lambda(x, \xi) d\xi$ can be interpreted as the conditional probability of X having a point in an infinitesimal small region containing ξ and of size $d\xi$ given the rest of X is x .

Let \sim denote a symmetric relation on S , and $N_\xi = \{\eta \in S : \xi \sim \eta\}$ the neighbours to $\xi \in S$. If for any finite $x \subset S$ and $\xi \in S \setminus x$, $\lambda(x, \xi)$ depends only on ξ and N_ξ , then X is said to be a Markov point process (with respect to \sim), cf. Ripley and Kelly (1977) and Van Lieshout (2000). Obviously, any hereditary density f defines a Markov point process if

we let all pairs of points be neighbours ($\xi \sim \eta$ for all $\xi, \eta \in S$), but as noticed in Sections 2.3-2.4 we have much more restricted relations in mind.

2.1 Lower bound on the Papangelou conditional intensity

Assume that for any finite $x \subset S$ and $\xi \in S \setminus x$,

$$\lambda(x, \xi) \geq \rho(\xi) \geq 0 \tag{1}$$

where ρ is a (deterministic) Borel function. Since the Poisson process with intensity function ρ is considered below, in order to avoid the trivial case where this process is almost surely empty, we also assume that the Lebesgue integral $\int_S \rho(\xi) \, d\xi$ is positive. Thus (1) is a rather strong condition on λ . For example, pairwise interaction processes are in general excluded, cf. the Strauss process, the hard core Gibbs point process, and many other examples in Baddeley and Møller (1989), Van Lieshout (2000), and Møller and Waagepetersen (2004). However, condition (1) is satisfied for a saturated Strauss process (Geyer, 1999) and for the area-interaction point process studied in Section 2.4.

2.2 Coupling

Consider a spatial birth-and-death process $(X_t; t \in \mathbb{R})$ with birth rate λ and death rate equal to unity, whereby $(X_t; t \in \mathbb{R})$ is reversible with respect to the equilibrium distribution specified by f (Preston, 1977; Ripley, 1977). Let X_{-t} denote the state just before time t . The spatial birth-and-death process is a jump process, i.e. it is a continuous-time Markov process where a jump at time t is either a birth $X_t = X_{-t} \cup \{\xi\}$ of a new point ξ or a death $X_t = X_{-t} \setminus \{\eta\}$ of an old point $\eta \in X_{-t}$. If the process is in a state x (a finite subset of

S) after a jump, the waiting time to the next jump is exponentially distributed with mean $1/A(x)$, where $A(x) = B(x) + n(x)$, $B(x) = \int_S \lambda(x, \xi) d\xi$, and $n(x)$ is the number of points in x . Conditional on that t is a jump time and $X_{-t} = x$, a birth $X_t = x \cup \{\xi\}$ happens with probability $B(x)/A(x)$, in which case the newborn point ξ has density $\lambda(x, \xi)/B(x)$, and otherwise a death $X_t = x \setminus \{\eta\}$ happens, where η is a uniformly picked point from x . For the present paper, as described in Section 2.3, we need only to generate a finite number of jumps of X_t for $t \leq 0$.

Assume that \emptyset (the empty point configuration) is an ergodic state of $(X_t; t \in \mathbb{R})$. This condition implies that f specifies the unique equilibrium distribution, cf. Preston (1977) where also conditions ensuring the condition is discussed. For instance, ergodicity of \emptyset is implied by the local stability condition,

$$\lambda(x, \xi) \leq V(\xi) \tag{2}$$

where V is an integrable function. In turn (2) is satisfied for most models used in practice, see Møller and Waagepetersen (2004).

Suppose that $X = X_0$ follows f , and imagine that we have generated X_t forwards for $t \geq 0$ and backwards for $t \leq 0$ (by reversibility the same generation can be used forwards and backwards). We can then couple $(X_t; t \in \mathbb{R})$ with a lower spatial birth-and-death process $(W_t; t \in \mathbb{R})$ obtained as follows. Each time t where $X_t = \emptyset$ is a regeneration time, and then we set $W_t = \emptyset$. These regeneration times split $((X_t, W_t); t \in \mathbb{R})$ into independent and identically distributed cycles, and so it suffices to consider the generation of $(W_t; t \in \mathbb{R})$ within each cycle. If a birth $X_t = X_{-t} \cup \{\xi\}$ happens, then with probability $\rho(\xi)/\lambda(x, \xi)$ (independently of what else has happened at previous jump times) we let $W_t = W_{-t} \cup \{\xi\}$,

and otherwise $W_t = W_{-t}$ is unchanged. If a death $X_t = X_{-t} \setminus \{\eta\}$ happens, then we let $W_t = W_{-t} \setminus \{\eta\}$ (meaning that $W_t = W_{-t}$ is unchanged if $\eta \notin W_{-t}$). Thereby $(W_t; t \in \mathbb{R})$ is a spatial birth-and death process with birth rate ρ and death rate equal to unity, and W_0 follows f (we omit the proof, since it follows along similar lines as in Appendix G in Møller and Waagepetersen, 2004). Clearly, $W_t \subseteq X_t$ for all $t \in \mathbb{R}$.

2.3 Dependent thinning procedure

As above suppose that (1) is satisfied and \emptyset is an ergodic state of the spatial birth-and-death processes $(X_t; t \in \mathbb{R})$, where $X = X_0$ follows f . Assume also that X is a Markov point process (Section 2.1). We can then generate the Poisson process W_0 with intensity function ρ within a random but finite number of steps, as described below.

Denote Z_0, Z_1, \dots the jump chain of $(X_t; t \leq 0)$ considered backwards in time, meaning that $Z_0 = X_0$, Z_1 is the state just before the first jump time before time 0 occurs, and so on. Similarly, define Y_0, Y_1, \dots , again in reverse chronological order, as the jump chain of $(W_t; t \leq 0)$. Let $G_0 = Z_0$, and for $i = 1, 2, \dots$, define recursively the i th generation ancestors of X by

$$G_i = \{\xi \in Z_i : N_\xi \cap G_{i-1} \neq \emptyset\}.$$

Since \emptyset is an ergodic state, we can define a discrete non-negative random variable I by $I = 0$ if $X = \emptyset$, and else $G_0 \neq \emptyset, \dots, G_{I-1} \neq \emptyset, G_I = \emptyset$. We also define a discrete non-negative random variable J by $J = 0$ if $X = \emptyset$, and else $Z_0 \neq \emptyset, \dots, Z_{J-1} \neq \emptyset, Z_J = \emptyset$.

Now, the dependent thinning procedure works as follows. If $X_0 = \emptyset$, then simply $W_0 = \emptyset$. In most practical applications we expect $X_0 = \emptyset$ to be a very unlikely event. Below we assume

that $X_0 \neq \emptyset$.

First, for each $i = 1, \dots, I$, simulate Z_i and determine G_i . Here we use that conditional on $Z_{i-1} = x$, a (backwards) birth $Z_i = x \cup \{\xi\}$ happens with probability $B(x)/A(x)$, in which case the newborn point ξ has density $\lambda(x, \xi)/B(x)$, and otherwise a (backwards) death $Z_i = x \setminus \{\eta\}$ happens, where η is a uniformly picked point from x . This requires the evaluation of the function $B(x)$, where numerical integration may be needed.

Second, considering the jump chain for the lower birth-and-death process forwards in time, set $Y_I = \emptyset$, and for $i = I - 1, \dots, 0$, generate Y_i in the same way as in the coupling construction in Section 2.2, but where of course we only need to consider the ancestors of X , since all other points will be irrelevant for the output W_0 . Specifically, for $i = I - 1, \dots, 1$, let U_i be a uniform random variable on $[0, 1]$ which is independent of what so far been generated forwards from step I to step i ; that is, (Z_I, Z_{I-1}, Y_I) if $i = I - 1$, and $(Z_I, \dots, Z_i, Y_I, \dots, Y_{i+1}, U_{I-1}, \dots, U_{i+1})$ if $i > I - 1$. Then

- if $Z_i = Z_{i+1} \cup \{\xi_i\}$ is a (forwards) birth and $U_i \leq \rho(\xi_i)/\lambda(Z_{i+1}, \xi_i)$, then $Y_i = Y_{i+1} \cup \{\xi_i\}$, and else $Y_i = Y_{i+1}$;
- if $Z_i = Z_{i+1} \setminus \{\eta_i\}$ is a (forwards) death, then $Y_i = Y_{i+1} \setminus \{\eta_i\}$.

Thereby we can return $W_0 = Y_0$ within the $2I$ steps given above. In many applications, the mean value of J might be extremely large (Berthelsen and Møller, 2002). In the extreme case where all points are neighbours, $I = J$. However, in most application examples, $I \ll J$. See Fernández *et al.* (2002) and Berthelsen and Møller (2002).

2.4 Example: area-interaction point process

For S a bounded planar region, the area-interaction point process X has Papangelou conditional intensity

$$\lambda(x, \xi) = \beta \gamma^{-|b(\xi, r) \setminus \bigcup_{\eta \in x} b(\eta, r)|} \quad (3)$$

where β, γ, r are positive parameters, $|\cdot|$ denotes area, and $b(\xi, r)$ is the disc with radius r centered at ξ (Baddeley and Van Lieshout, 1995). For $\gamma > 1$, $\lambda(x, \xi)$ is increasing in x (the attractive case, originally studied by Widom and Rowlinson, 1970). For $\gamma < 1$, $\lambda(x, \xi)$ is decreasing in x (the repulsive case). For $\gamma = 1$, X is simply a homogeneous Poisson process on S with intensity β . Consequently, both (1) and (2) are satisfied, with

$$\rho(\xi) = \beta \gamma^{-\pi r^2}, \quad V(\xi) = \beta, \quad \text{if } \gamma \geq 1$$

and

$$\rho(\xi) = \beta, \quad V(\xi) = \beta \gamma^{-\pi r^2}, \quad \text{if } 0 < \gamma \leq 1.$$

An inhomogeneous version of the area-interaction point process and satisfying (1)-(2) is obtained by replacing β by a non-negative Borel function $\beta(\xi)$ all places above, assuming $\int_S \beta(\xi) d\xi < \infty$ which ensures the existence of the process. Clearly, no matter which version is used, (3) implies that X is Markov with respect to the relation given by that $\xi \sim \eta$ if and only if $b(\xi, r) \cap b(\eta, r) \neq \emptyset$, i.e. when the distance between ξ and η is $\leq 2r$.

Figures 1a and 2a show examples of simulated homogeneous area-interaction processes with $\gamma > 1$ and $\gamma < 1$, respectively. The aggregation in the area-interaction process for $\gamma > 1$ can be seen in Figure 1a and is confirmed in Figure 1b which shows the estimated centered L -function $\hat{L}(d) - d$ corresponding to the realization in Figure 1b, along with pointwise 95%-

confidence bounds based on simulations of 1000 homogenous Poisson processes whose rates are equivalent to that observed by the process in Figure 1a. Specifically, in terms of Ripley's K -function (Ripley, 1976, 1977), $L(d) = \sqrt{K(d)/\pi}$, where $d > 0$ denotes distance and we use Ripley's non-parametric estimate \hat{K} to obtain the estimate $\hat{L}(d) = \sqrt{\hat{K}(d)/\pi}$. Note that in the special case of a stationary Poisson process, the centered L -function $L(d) - d$ is equal to zero. Figure 1c shows W , the random thinning of the process shown in Figure 1a, using the method described in Section 2.3, and Figure 1d shows the estimated L -function for the thinned process in Figure 1c. Figure 1 demonstrates that the clustering in the area-interaction process with $\gamma > 1$ is removed by the thinning procedure, resulting in a homogeneous Poisson process.

Similarly, Figure 2 shows the random thinning of an area-interaction process with $\gamma < 1$, and the estimated centered L -functions for both the realization of the area-interaction process and its corresponding thinned process. As with Figure 1, one sees in Figure 2b that the inhibition in the original process is statistically significant, compared to the homogeneous Poisson process, for distances of 0.3 to 0.6, and that this inhibition is removed by the random thinning procedure, as confirmed by Figure 2d.

3 Thinning Cox processes

In the sequel, let $\lambda(\xi)$ be a random non-negative function defined for all $\xi \in S$ such that conditional on λ , X is a Poisson process with intensity function λ . In other words, X is a Cox process driven by λ . We assume that almost surely $\int_S \lambda(\xi) d\xi$ is finite, meaning that X is a finite point process.

3.1 Lower bound on the random intensity

Assume that, with probability one,

$$\lambda(\xi) \geq \rho(\xi) \geq 0, \tag{4}$$

where as in Section 2, ρ is a (deterministic) Borel function with $\int_S \rho(\xi) d\xi > 0$. Condition (4) implies X can be viewed as the superposition $Q \cup R$ of a Cox process Q driven by $\gamma = \lambda - \rho$ and an independent Poisson process R with intensity function ρ . For example, γ may be log Gaussian and Q then a log Gaussian Cox process (Møller, Syversveen and Waagepetersen, 1998), or a shot noise process and Q then a shot noise Cox process (Møller, 2003).

3.2 Independent thinning procedure

If λ were known, then we could directly obtain a Poisson process W on S with intensity function ρ as an independent thinning of X with retention probabilities $\rho(\xi)/\lambda(\xi)$. Indeed, since the independent thinning of the Poisson process $X|\lambda$ is a Poisson process with intensity function ρ , W is a Poisson process independent of λ and with intensity function ρ .

In practice we usually only observe a realization of $X = x$, and hence we need first to generate λ conditional on $X = x$. Hence the independent thinning procedure works by:

- first, generating a realization of λ conditional on $X = x$
- second, generating W as an independent thinning of X with retention probabilities $\rho(\xi)/\lambda(\xi)$, $\xi \in x$.

How to simulate λ conditional on $X = x$ depends on the particular model. For example, if γ is log Gaussian, a Langevin-Hastings algorithm can be used (Møller *et al.*, 1998; Møller and

Waagepetersen, 2004), and if γ is a shot-noise process, a birth-death Metropolis-Hastings algorithm applies (Møller, 2003; Møller and Waagepetersen, 2004). We run one of these Metropolis-Hastings algorithms until it is effectively in equilibrium, and then return an (approximate) simulation of λ conditional on $X = x$.

3.3 Example: Thomas process

Let S be a bounded planar region and Q a modified Thomas process (Thomas, 1949), i.e. a Cox process driven by

$$\gamma(\xi) = \omega \sum_{\eta \in \Phi} \varphi(\xi - \eta), \quad \xi \in S$$

where φ is the bivariate normal density with mean 0 and covariance matrix $\sigma^2 I$, Φ is a homogeneous Poisson process with intensity κ defined on a bounded Borel set $S_{\text{ext}} \supseteq S$, and ω, σ, κ are positive parameters. Here S_{ext} is chosen sufficient large so that edge effects can effectively be ignored, see Møller (2003) and Møller and Waagepetersen (2004). In the following examples, $S = [0, 10] \times [0, 10]$ and $S_{\text{ext}} = [-10, 20] \times [-10, 20]$.

Figure 3 shows the thinning of the superposition $X = Q \cup R$ when the Poisson process R has constant rate $\rho = 6$ and the Thomas process has parameters $(\kappa, \omega, \sigma) = (2, 3, 0.2)$. From Figure 3b one sees that the process is highly clustered, which is a result of the fact that points tend to be clustered around the cluster centers Φ . Figure 3c shows the corresponding thinned process using the method described in Section 3.2, and the centered L -function estimate in Figure 3d verifies that the resulting thinned process is homogeneous Poisson.

One may also consider an inhomogeneous version where $\rho(\xi)$ depends on ξ . For instance, Figure 4a shows a simulation of such a process with $\rho(x, y) = \exp(ax + by)$, where (x, y) are

the Cartesian coordinates of ξ . Unlike the previous examples, in this case the thinned process is an inhomogeneous Poisson process with intensity function ρ , rather than a homogeneous Poisson process. Figure 4b shows the resulting thinned process, and Figure 4c displays a centered version of the estimated inhomogeneous L -function (Baddeley *et al.*, 2000; Veen and Schoenberg, 2005), confirming that the process in Figure 4c is an inhomogeneous Poisson process.

Acknowledgments

Supported by the Danish Natural Science Research Council, grant 272-06-0442, "Point process modelling and statistical inference". We made use of several functions in the R package *spatstat* (Baddeley and Turner, 2005).

References

- Baddeley, A. and Møller, J. (1989). Nearest-neighbour Markov point processes and random sets. *International Statistical Review*, 2, 89–121.
- Baddeley, A., Møller, J., and Waagepetersen, R. (2000). Non and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statistica Neerlandica*, 54, 329–350.
- Baddeley, A. and Turner, R. (2005). Spatstat: an R package for analyzing spatial point patterns. *Journal of Statistical Software*, 1–42.
- Baddeley, A., Turner, R., Møller, J. and Hazelton, M. (2005). Residual analysis for spatial point processes (with discussion). *Journal of Royal Statistical Society Series B*, 67, 617–666.
- Baddeley, A.J., and Lieshout, M.N.M. van (1995). Area-interaction point processes. *Annals of the Institute of Statistical Mathematics*, 46, 601–619.
- Berthelsen, K.K. and Møller, J. (2002). A primer on perfect simulation for spatial point processes. *Bulletin of the Brazilian Mathematical Society*, 33, 351–367.
- Carter, D.S. and Prenter, P.M. (1972). Exponential spaces and counting processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 21, 1–19.
- Cox, D.R. (1955). Some statistical models related with series of events. *Journal of the Royal Statistical Society Series B*, 17, 129–164.
- Cressie, N.A. (1993). *Statistics for Spatial Data, revised ed.* Wiley, New York.

- Fernández, R., and Ferrari, P.A., and Garcia, N.L. (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stochastic Processes and their Applications*, 102, 63–88.
- Geyer, C.J. (1999). Likelihood inference for spatial point processes. In *Stochastic Geometry: Likelihood and Computation*, Barndorff-Nielsen, O.E., Kendall, W.S., and Lieshout, M.N.M. van, eds. Chapman & Hall / CRC, Boca Raton, Florida, pp. 79–140.
- Kallenberg, O. (1984). An informal guide to the theory of conditioning in point processes. *International Statistical Review*, 52, 151–164.
- Kendall, W.S. (1998). Perfect simulation for the area-interaction point process. in *Probability Towards 2000*, L. Accardi and C.C. Heyde, eds., Springer Lecture Notes in Statistics 128, Springer Verlag, New York, pp. 218–234.
- Kendall, W.S. and Møller, J. (2000). Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Advances in Applied Probability*, 32, 844–865.
- Lieshout, M.N.M. van (2000). *Markov Point Processes and Their Applications*. Imperial College Press, London.
- Møller, J. (2003). Shot noise Cox processes. *Advances in Applied Probability*, 35, 4–26.
- Møller, J., Syversveen, A.R., and Waagepetersen, R.P. (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25, 451–482.

- Møller, J. and Waagepetersen, R.P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall / CRC, Boca Raton.
- Møller, J. and Waagepetersen, R.P. (2007). Modern spatial point process modelling and inference (with discussion). *Scandinavian Journal of Statistics*, 34, 643–711.
- Preston, C.J. (1977). Spatial birth-and-death processes. *Bulletin of the International Statistical Institute*, 46, 371–391.
- Ripley, B. (1976). The second-order analysis of stationary point processes. *Journal of Applied Probability*, 13, 255–266.
- Ripley, B.D. (1977). Modelling spatial patterns (with discussion). *Journal of the Royal Statistical Society Series B*, 39, 172–212.
- Ripley, B.D. and Kelly, F.P. (1977). Markov point processes. *J. London Math. Soc.*, 15, 188–192.
- Schoenberg, F.P. (2003). Multidimensional residual analysis of point process models for earthquake occurrences. *Journal of the American Statistical Association*, 98, 789–795.
- Schoenberg, F.P. (2005). Comment on "Residual analysis for spatial point processes" by Baddeley, A., Turner, R., Møller, J. and Hazelton, M. *Journal of the Royal Statistical Society Series B*, 67, 661.
- Schoenberg, F.P. and Zhuang, J. (2008). Residual analysis via thinning for spatial point processes. *Annals of the Institute of Statistical Mathematics*, in review.

- Stoyan, D., Kendall, W.S., and Mecke, J. (1995). *Stochastic Geometry and Its Applications*, 2nd ed. Wiley, Chichester.
- Thomas, M. (1949). A generalization of Poisson's binomial limit for use in ecology. *Biometrika*, 36, 18–25.
- Veen, A. and Schoenberg, F.P. (2005). Assessing spatial point process models for California earthquakes using weighted K-functions: analysis of California earthquakes. In *Case Studies in Spatial Point Process Models*, Baddeley, A., Gregori, P., Mateu, J., Stoica, R., and Stoyan, D. (eds.), Springer, New York, pp. 293–306.
- Widom, B., and Rowlinson, J.S. (1970). A new model for the study of liquid-vapor phase transitions. *Journal of Chemical Physics*, 52, 1670–1684.

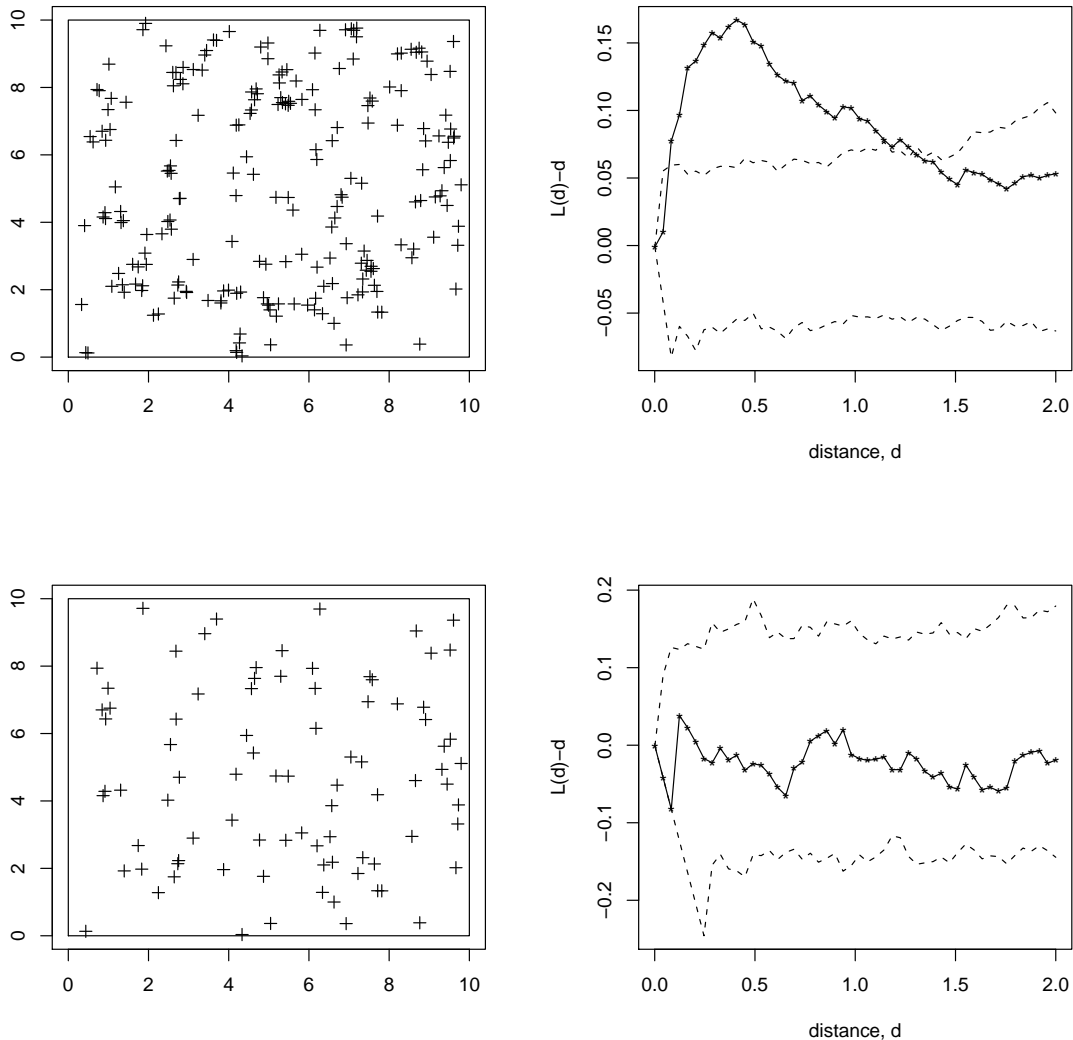


Figure 1: Thinning of a simulated attractive area-interaction process on $S = [0, 10] \times [0, 10]$, with $\gamma = 10^3, \beta = 7, r = 0.3$. Top-left panel (a): simulated area-interaction process, X . Top-right panel (b): estimated centered L -function corresponding to X , along with empirical pointwise 95%-confidence bounds obtained by simulating 1000 homogeneous Poisson processes on $[0, 10] \times [0, 10]$ each with expected number of points equal to the number observed in X . Bottom-left panel (c): thinned process W corresponding to X . Bottom-right panel (d): estimated centered L -function corresponding to W , along with empirical pointwise 95%-confidence bounds from 1000 simulated homogeneous Poisson processes each with expected number of points equal to the number observed in W .

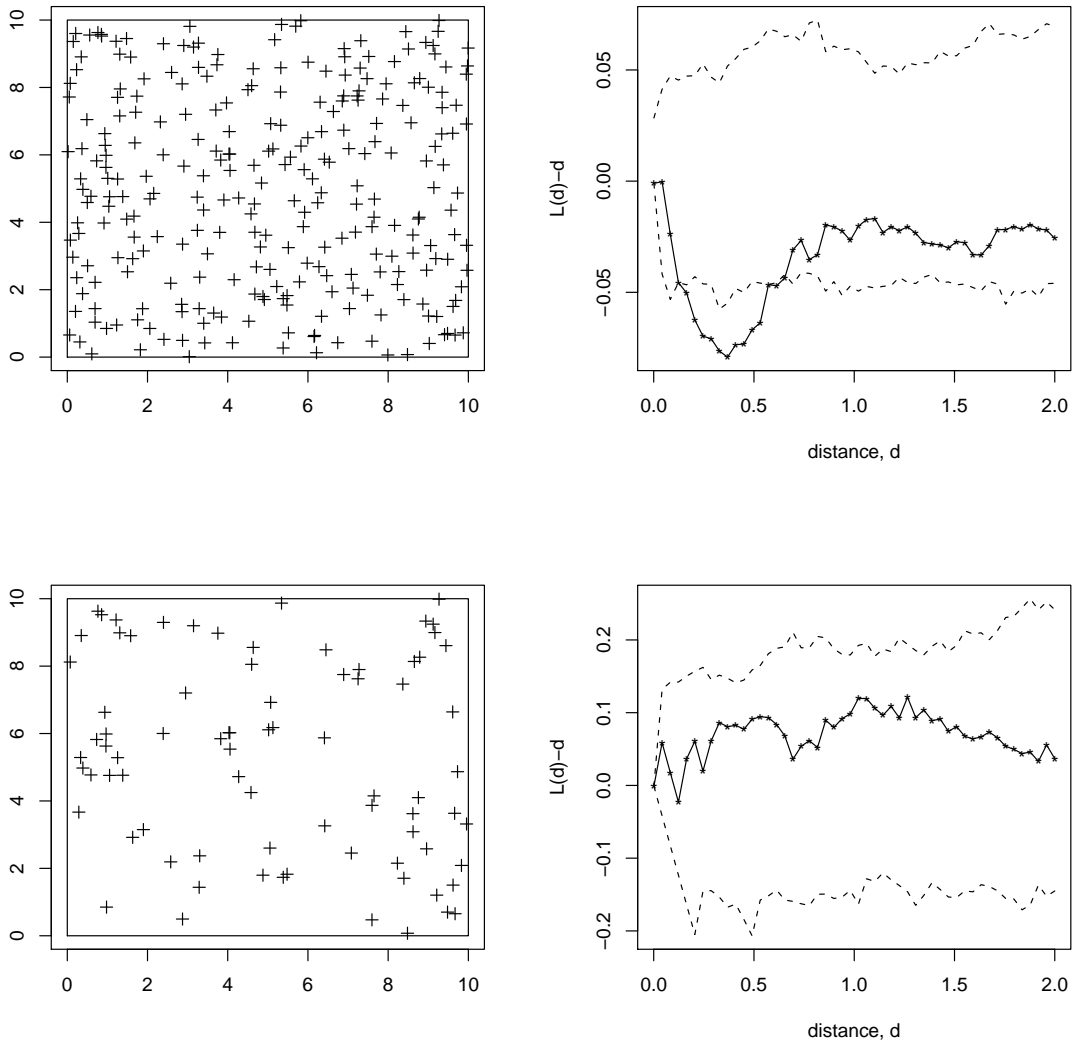


Figure 2: The same four plots as in Figure 1, but for the inhibitive case where $\gamma = 10^{-3}$, $\beta = 1$, $r = 0.3$. (a): simulated area-interaction process, X . Top-right panel (b): estimated centered L -function for X with empirical pointwise 95%-confidence bounds obtained by simulating 1000 homogeneous Poisson processes on $[0, 10] \times [0, 10]$ each with expected number of points equal to the number observed in the process X . Bottom-left panel (c): thinned process W corresponding to X . Bottom-right panel (d): estimated centered L -function corresponding to W with empirical pointwise 95%-confidence bounds from 1000 simulated homogeneous Poisson processes on $[0, 10] \times [0, 10]$ each with expected number of points equal to the number observed in the process W .

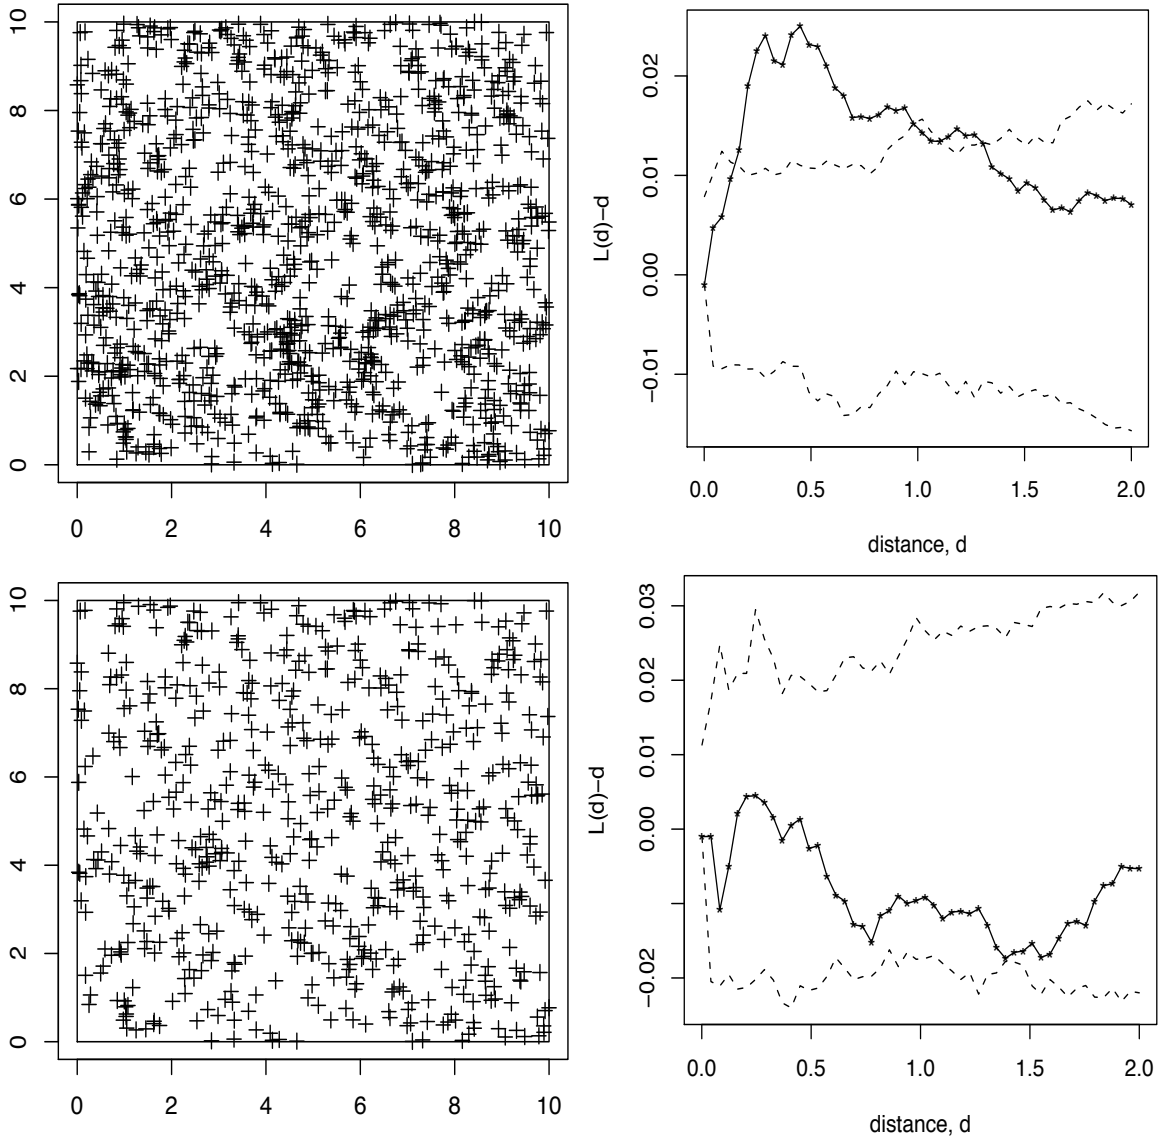


Figure 3: Thinning of a simulated superposition of a homogeneous Poisson process with rate $\rho = 6$ and a Thomas process on $S = [0, 10] \times [0, 10]$ with $\kappa = 2$, $\omega = 3$, and $\sigma = 0.2$. Top-left panel (a): simulated superposition, X . Top-right panel (b): estimated centered L -function for X with empirical pointwise 95%-confidence bounds obtained via 1000 simulated homogeneous Poisson processes each with expected number of points equal to the number observed in X . Bottom-left panel (c): thinned process W corresponding to X . Bottom-right panel (d): estimated centered L -function for W with empirical pointwise 95%-confidence bounds obtained via 1000 simulated homogeneous Poisson processes each with expected number of points equal to the number observed in W .

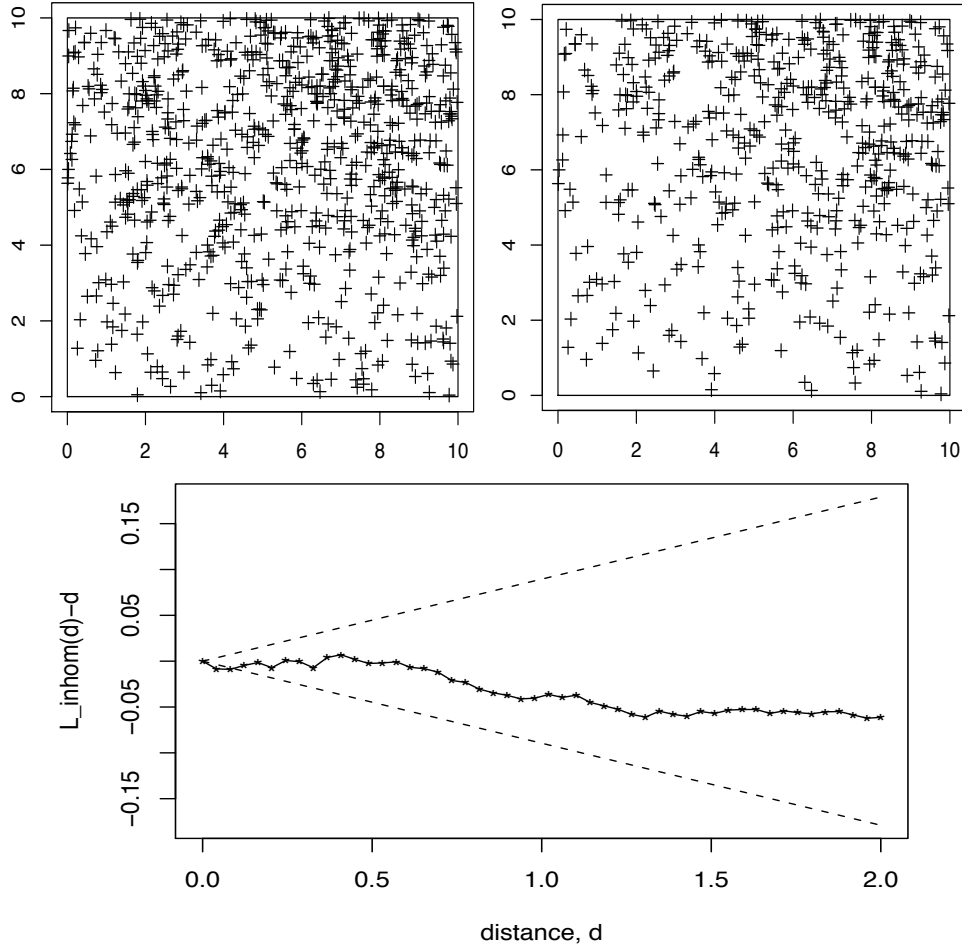


Figure 4: Thinning of a simulated superposition of an inhomogeneous Poisson process and Thomas process on $S = [0, 10] \times [0, 10]$, where the Poisson process has intensity function $\rho(x, y) = \exp(ax + by)$ with $(a, b) = (0.1, 0.2)$, and the Thomas process has parameter $(\kappa, \omega, \sigma) = (1, 5, 0.5)$. Top-left panel (a): simulated superposition, X . Top-right panel (b): thinned process W corresponding to X . Bottom panel (c): centered estimate of the inhomogeneous $L(d)$ corresponding to W , using intensity ρ , with pointwise 95%-confidence bounds based on the normal approximation in Veen and Schoenberg (2005), i.e. $\pm 1.96\sqrt{2\pi 10^2 d}/E(\#R)$, where $E(\#R) = \int_S \rho(x, y) dx dy = [\exp(10a) - 1][\exp(10b) - 1]/(ab)$ is the expected number of points in the Poisson process R .