

LIKELIHOOD-BASED INFERENCE FOR MATÉRN TYPE-III REPULSIVE POINT PROCESSES

MARK L. HUBER * AND

ROBERT L. WOLPERT,** *Duke University*

Abstract

In a repulsive point process, points act as if they are repelling one another, leading to underdispersed configurations when compared to a standard Poisson point process. Such models are useful when competition for resources exists, as in the locations of towns and trees. Bertil Matérn introduced three models for repulsive point processes, referred to as types I, II, and III. Matérn used types I and II, and regarded type III as intractable. In this paper an algorithm is developed that allows for arbitrarily accurate approximation of the likelihood for data modeled by the Matérn type-III process. This method relies on a perfect simulation method that is shown to be fast in practice, generating samples in time that grows nearly linearly in the intensity parameter of the model, while the running times for more naive methods grow exponentially.

Keywords: Coupling from the past; point process; thinning; underdispersed

2000 Mathematics Subject Classification: Primary 65C60

Secondary 62M30; 68U20

1. Introduction

Spatial data are often more regularly spaced than would be expected under a simple Poisson point process model because of competition for resources among the entities represented by the points of the model. These types of point process are called *repulsive* because the points act as if they are repelling one another. See [11] and [30] for examples of this type of data, which usually arises when entities are in competition for finite resources.

Various methods have been introduced to model repulsive or underdispersed point processes, including simple inhibition processes [17, Chapter 3], thinned processes [3], [27], the Strauss process [31], and Markov point processes [23], which model the configuration by assigning a density (to discourage ‘closeness’ of the points) with respect to a homogeneous Poisson point process.

Matérn [17, p. 48] offered a different approach to modeling repulsion. Rather than use a density to describe the model, he gave three algorithms for generating configurations that implicitly (rather than explicitly) build in repulsion.

In Section 2 we describe this approach in more detail, and present the central questions that need to be addressed in order to pursue a Bayesian (or other likelihood-based) approach to inference for these problems. Section 3 gives a simple acceptance–rejection approach to solve

Received 7 November 2008; revision received 4 June 2009.

* Current address: Department of Mathematics and Computer Science, Claremont McKenna College, 850 Columbia Avenue, Claremont, CA 91711, USA. Email address: mhuber@cmc.edu

** Postal address: Department of Statistical Science and Nicholas School of the Environment, Duke University, Durham, NC 27708-0251, USA. Email address: wolpert@stat.duke.edu


this problem that is suitable for small data sets, but which scales poorly with problem size. In Section 4 we present a new perfect simulation method that can handle much larger sets, and develop an arbitrarily accurate product estimator for the likelihood function. In Section 5 this new method is applied to simulated data to test the effectiveness and running time.

2. The Matérn point processes

Matérn [17, pp. 47–48] introduced what are now called his *type-I* and *type-II* thinned point processes (see Figure 1). Begin with an intensity $\lambda > 0$, a σ -finite Borel reference measure $\mu(\cdot)$ on \mathbb{R}^d , and a ‘primary’ Poisson point process with mean $\lambda\mu(\cdot)$ that will assign to any region \mathcal{S} of finite measure a Poisson-distributed number of points with expected value $\lambda\mu(\mathcal{S})$.

Given a fixed radius $R > 0$, Matérn’s type-I process consists of the ‘secondary’ points remaining upon removing all primary points that lie within distance R of any other primary point (so, if a pair lie within distance R , both are removed).


For the type-II process, Matérn assigns independent and identically distributed (i.i.d.) marks $t_j \stackrel{\text{i.i.d.}}{\sim} \text{Un}(0, 1]$ (called ‘times’) to each of the primary points. The secondary points are those whose time mark is smaller than all neighbors within distance R . In other words, Matérn’s type-II process only removes the later of two points which lie within distance R ; for fixed λ and R , this will have a higher density of points than the type-I process.




Primary Poisson process
 $R = 0.4$, $\text{dist}(a, b) = 0.30$, $\text{dist}(b, c) = 0.23$,
 $\text{dist}(a, c) = 0.53$, $t_a = 0.27$, $t_b = 0.52$, $t_c = 0.78$.



Matérn I
 Pair a and b eliminated; pair b and c eliminated; no points remain.



Matérn II
 $t_a < t_b$: a eliminates b ; $t_b < t_c$: b eliminates c ; only point a remains.



Matérn III
 $t_a < t_b$: a eliminates b ; since b has been eliminated, it does not eliminate c . Points a and c remain.

FIGURE 1: Example of Matérn repulsive point processes.

Matérn [17, p. 48] briefly mentioned a third point process featuring ‘dynamic’ thinning in which points are removed only if they lie within distance R of an earlier *secondary* point; the remaining points, which we call the Matérn type-III process (or, more briefly, Matérn III), will have still higher density (for fixed λ and R) than types I or II. It is this third process which is considered in detail below.

Let μ be a finite measure on a Borel subset \mathcal{S} of \mathbb{R}^d . All three Matérn processes begin with a Poisson-distributed primary random finite subset $z \sim \text{Po}(\lambda\mu(dz))$ of \mathcal{S} , and yield random sets $x \subseteq z \subseteq \mathcal{S}$ of secondary points with the property that balls of radius $R/2$ centered at each point do not intersect. Let $\#(s)$ denote the number of points in any set $s \subseteq \mathcal{S}$. Call the $\#(x)$ secondary points $x \subseteq z$ that remain in the process *seen points*, and the $\#(y) = \#(z) - \#(x)$ points $y = z \setminus x$ that have been removed *hidden points*. Renumber if necessary so we can write $x = (x_1, \dots, x_{\#(x)})$ for the seen points and $t = (t_1, \dots, t_{\#(x)})$ for the associated (unobserved) time marks.

Note that the time marks could have been drawn from any continuous distribution on the real numbers [28]. Suppose that instead of uniform distributions, independent draws (\tilde{t}_i) are taken from some other distribution over the real numbers with continuous cumulative distribution function $F(\cdot)$. Then $\{t_i := F(\tilde{t}_i)\}$ will be independent and uniformly distributed on $(0, 1]$. Moreover, ties are impossible under both (continuous) distributions and $\tilde{t}_i < \tilde{t}_j$ if and only if $t_i < t_j$, so no generality is lost in assigning uniform distributions on $(0, 1]$ to the time marks.

Let $\theta \in \Theta$ index possible parameter vectors $(\lambda_\theta, R_\theta)$. The Matérn process may also be constructed from a Poisson point process on $\mathcal{S} \times (0, 1]$ with product intensity measure $\Lambda_\theta(dz dt) := \lambda_\theta \mu(dz)m(dt)$, where m denotes the Lebesgue (or uniform) measure on $(0, 1]$. The first coordinate is the location in \mathcal{S} , while the second coordinate is the time mark. The *seen* points are those for which no other seen point with a lower time mark lies within distance R . Note that $\Lambda_\theta(\mathcal{S} \times (0, 1]) = \lambda_\theta \mu(\mathcal{S})$.

2.1. Connection with RSA

Physicists and chemists have long studied the ‘random sequential adsorption’ or continuum RSA model for the irreversible binding of proteins to surfaces. In its most common form this model constructs a hard-core process as a sequence of points drawn from the uniform distribution on the complements (within some region $\mathcal{S} \subset \mathbb{R}^2$) of the unions of disks of radius R centered at each of the previously drawn points. When the disks cover \mathcal{S} entirely, the ‘jamming limit’ has been reached, whereupon the process halts. Variations include replacing the disks with squares or other convex shapes, constructing the process in \mathbb{R}^d rather than the plane, employing independent random radii R_i , stopping after a specified number of points have been drawn or attempts have been made, etc. For more details and some of the historical development, see [5], [7], [8], [22], [24], [25], and [32].

Now consider stopping the RSA process before the jamming limit by drawing a primary Poisson process with intensity 1 on $\mathcal{S} \times (0, T]$ and thinning points sequentially in time by removing those born within distance R of retained points with smaller time marks. Now divide each of the time marks by T . The result is a Matérn III process where once again the time marks lie in $(0, 1]$, but the intensity of the process is now T . The limit as $T \rightarrow \infty$ of the stopped RSA process is just the jamming limit of the RSA process. Therefore, the limit of the Matérn III process as the intensity goes to infinity is also the jamming limit of the RSA process.

2.2. Likelihood function for Matérn III

Informal estimates of $(\lambda_\theta, R_\theta)$ are simple to find for any of the Matérn processes. The maximum likelihood estimate (MLE) \hat{R} is simply the smallest observed point spacing, $\hat{R} = \min_{i \neq j} \{d(x_i, x_j)\}$, for example, and method of moments estimates are available for λ_θ from

empirical estimates of their packing densities (which have simple closed-form relations to λ_θ and R_θ for the Matérn I and II processes, and a relation available from simulation for Matérn III). Our concern in this paper is quite different—we wish to construct the likelihood function for θ , to support fully coherent inference from the Bayesian or other perspectives using this model (see [1]).

For one-dimensional point sets, a modification of dynamic programming techniques can be used to evaluate the Matérn III likelihood function. The data sets considered in Section 5 are two-dimensional, so the dynamic programming approach fails; moreover, the techniques developed in this paper work not only for any window in \mathbb{R}^d with finite measure, but also for far more general sets. This includes any complete separable metric space $(\mathcal{S}, \text{dist})$ with finite Borel measure μ .

For any possible configuration $x = (x_1, \dots, x_{\#(x)}) \subseteq \mathcal{S}$ of seen points, the vector of time marks $t = (t_1, \dots, t_{\#(x)})$ could take any value in the cube $(0, 1]^{\#(x)}$. Any hidden point $y \in \mathcal{S}$ must lie within distance R_θ from some seen point $x_i \in x$ with time mark $t_y > t_i$, since otherwise (y, t_y) would not have been removed in thinning. For $s \in \mathcal{S}$ and $r > 0$, let $B_r(s) := \{y \in \mathcal{S} : \text{dist}(s, y) < r\}$ denote the open r -ball in \mathcal{S} centered at s . Then define

$$\begin{aligned}
 D(x, t) &:= \{(y, t_y) : (\text{there exists } x_i \in x)[\text{dist}(x_i, y) < R_\theta] \wedge [t_i < t_y]\} \\
 &= \bigcup_{i=1}^{\#(x)} B_{R_\theta}(x_i) \times (t_i, 1], \\
 A_\theta(x, t) &:= \Lambda_\theta(D(x, t)). \tag{2.1}
 \end{aligned}$$

Call region $D(x, t)$ the *shadow* of configuration x with times t , and call $A_\theta(x, t)$ the shadow’s *area*, the expected number of hidden points for seen configuration (x, t) and parameter θ . Figure 2 shows an example of the shadow region for three points in $\mathcal{S} = [0, 10]$. Note that $D(x, t)$ depends implicitly on R_θ , and $A_\theta(x, t)$ depends on both R_θ and λ_θ .

Van Lieshout [33] introduced a density for RSA (and, hence, for Matérn III), but the formulation there was given as the product of $\#(x)$ factors, each of which could be difficult to evaluate. Here a new result is presented that gives a closed-form expression for the density in terms of a simple expression of the measure of the shadow.




FIGURE 2: Example of a shadow region $D(x, t)$. Here $2R = 1.4$.

Theorem 2.1. *The density function of the seen points and their times with respect to the Poisson point process with mean $\mu \otimes m$ is*

$$f_{\text{seen points}}(x, t \mid \theta) = \mathbf{1}_{\{\rho(x) > R_\theta\}} \lambda_\theta^{\#(x)} \exp(\mu(\mathcal{S})\{1 - \lambda_\theta\}) \exp(A_\theta(x, t)), \quad (2.2)$$

where $\rho(x) := \min_{i \neq j} \text{dist}(x_i, x_j)$ denotes the minimum interpoint distance.

Proof. There are two ways to create a random draw of the seen points. Firstly, primary points (Z, S) are drawn from the $\text{Po}(\lambda_\theta(\mu \otimes m)(dz ds))$ distribution, and then the seen points are built by thinning, as described in Section 2. Secondly, use an acceptance–rejection algorithm [34]. Draw a candidate (X, T) from the $\text{Po}(\lambda_\theta(\mu \otimes m)(dx dt))$ distribution and then accept or reject it as a draw from the seen points. If rejection occurs, draw again, and continue until a draw is accepted.

To determine if (X, T) should be accepted as a draw from the seen points, use the following procedure. First, draw (Z, S) from $\text{Po}(\lambda_\theta(\mu \otimes m)(dz ds))$ conditioned on the points (X, T) being contained in (Z, S) . From a recursive application of Palm distributions (see [29, Section 4.4]), this can be accomplished by drawing $(Z \setminus X, S \setminus T)$ as an independent draw from the $\text{Po}(\lambda_\theta(\mu \otimes m)(dy ds))$ distribution. Next, accept (X, T) if it is exactly the set of seen points in (Z, S) , i.e. if no points of (Z, S) lie outside the shadow of (X, T) , an event with probability

$$\exp(-\lambda_\theta(\mu \otimes m)(\mathcal{S} \setminus D(X, T))) = \exp(A_\theta(X, T)) \exp(-\mu(\mathcal{S})).$$

Also, for (X, T) to be the seen points, the minimum distance $\rho(X)$ between any two points in X must exceed R_θ . So the probability of acceptance if $X = x$ and $T = t$ is

$$p(x, t) = \mathbf{1}_{\{\rho(x) > R_\theta\}} \exp(A_\theta(x, t)) \exp(-\lambda_\theta \mu(\mathcal{S})).$$

The theory of the acceptance–rejection algorithm states that the result has a density with respect to the proposal distribution that is proportional to $p(x, t)$:

$$f(x, t \mid \theta) = \frac{p(x, t)}{p_\theta},$$

where p_θ is the probability that a draw is accepted. So next determine p_θ .

Recall that (Z, S) is the union of two independent draws from $\text{Po}(\lambda_\theta(\mu \otimes m)(dz ds))$. The sum of Poisson random variables is Poisson with parameter equal to the sum of the individual parameters, and the points in (Z, S) are all i.i.d., so $(Z, S) \sim \text{Po}(2\lambda_\theta(\mu \otimes m)(dz ds))$. This has density $2^{\#(z)} \exp(-\lambda_\theta \mu(\mathcal{S}))$ with respect to the measure $\text{Po}(\lambda_\theta(\mu \otimes m)(dz ds))$.

Conditioned on the points (Z, S) , the points (X, T) are a subset of cardinality $\#(x)$ drawn uniformly at random from (Z, S) . This can be checked directly. The probability of drawing any number, n , of points for (Z, S) from its Poisson point process distribution is $\exp(-2\mu(\mathcal{S})) (2\mu(\mathcal{S}))^n / n!$. This must also be the probability of drawing i seen points (X, T) and $n - i$ hidden points $(Z \setminus X, S \setminus T)$,

$$\frac{\exp(-\mu(\mathcal{S}))\mu(\mathcal{S})^i}{i!} \frac{\exp(-\mu(\mathcal{S}))\mu(\mathcal{S})^{n-i}}{(n-i)!} = \frac{\exp(-2\mu(\mathcal{S}))\mu(\mathcal{S})^n}{i!(n-i)!}.$$

Dividing the right-hand side of the previous equation by the density of (Z, S) , which is $2^n \exp(-\lambda_\theta \mu(\mathcal{S}))$ from above, shows that the probability of drawing i points from (X, T) conditioned on drawing n points from (Z, S) is $2^{-n} n! / (i!(n-i)!)$, or simply the binomial distribution with parameters n and $\frac{1}{2}$.

This is the same as choosing (X, T) by flipping a fair coin for each point in (Z, S) , and putting the points with heads into (X, T) . For acceptance to occur, the points chosen to fall into (X, T) must exactly match the points selected from (Z, S) by the Matérn III process. This has probability $1/2^n$ if there are n points in (Z, S) . Canceling all the common factors shows that the total probability of acceptance is

$$p_\theta = \sum_{n=0}^{\infty} \exp(-2\mu(\mathcal{S})) \frac{\mu(\mathcal{S})^n}{n!} = \exp(-\mu(\mathcal{S})),$$

thereby showing that the density of (X, T) with respect to $\text{Po}(\lambda_\theta(\mu \otimes m)(dx dt))$ is just

$$f(x, t | \theta) = \frac{p(x, t)}{p_\theta} = \mathbf{1}_{\{\rho(x) > R_\theta\}} \exp(A_\theta(x, t)).$$

Finally, note that a draw (Z, S) from $\text{Po}(\lambda_\theta(\mu \otimes m)(dz ds))$ will have density function $\exp(\mu(\mathcal{S})(1 - \lambda_\theta)) \lambda_\theta^{\#(Z)}$ with respect to the $\text{Po}((\mu \otimes m)(dz ds))$ distribution, which completes the proof.

The marginal density function $g(x | \theta)$ for the seen points is then given by integrating the latent time marks:

$$g(x | \theta) = \mathbf{1}_{\{\rho(x) > R_\theta\}} \exp(\mu(\mathcal{S})\{1 - \lambda_\theta\}) \lambda_\theta^{\#(x)} \int_{(0, 1]^{\#(x)}} \exp\{A_\theta(x, t)\} dt, \tag{2.3}$$

where ‘ dt ’ denotes the Lebesgue measure differential $m(dt_1) \cdots m(dt_{\#(x)})$ on the unit cube $(0, 1]^{\#(x)}$. This is also the *likelihood function* for θ , upon observing $X = x$.

Note that $A_\theta(x, t)$ is an increasing function of R_θ , making $g(x | \theta)$ an increasing function of R_θ over $[0, \rho(x))$. In principle, for any prior distribution $\pi(d\theta)$, the posterior distribution for θ can be found as

$$\pi(d\theta | X = x) = Z_x^{-1} g(x | \theta) \pi(d\theta),$$

where $Z_x := \int_{\Theta} g(x | \theta') \pi(d\theta')$ is the appropriate normalizing constant. The awkward integral in (2.3) is an obstacle for this straightforward approach to inference. Instead, we will follow a data augmentation strategy in which random samples of hidden points Y and unobserved time marks T , for both seen and hidden points, are drawn from their conditional distributions (given x and θ); after this, inference about θ can be accomplished using a variety of methods. For use in Section 4.1, we first compute the conditional density function of the time mark vector $T \in (0, 1]^{\#(x)}$ for seen points x . From (2.3), this is simply

$$h_\theta(t | x) = Z_{\theta, x}^{-1} \exp(A_\theta(x, t)) \tag{2.4}$$

with normalizing constant $Z_{\theta, x} := \int_t \exp\{A_\theta(x, t)\} dt$.

3. Acceptance–rejection method

A simple approach for approximating $g(x | \theta)$ is the *acceptance–rejection* method of von Neumann [34]. This also provides a method to draw hidden points Y and time marks T from their conditional distributions that works well for small examples, but (as shown below) it scales badly in the number $\#(x)$ of seen points.

Algorithm 3.1. (*Acceptance–rejection.*) **Inputs:** parameter θ and seen points x .

Outputs: random draws of hidden points Y and time marks T .

1. repeat
2. **draw** hidden points $Y \subseteq \mathcal{S}$ from Poisson process with intensity $\lambda_\theta \mu(\cdot)$
3. **draw** time marks T for x and Y uniformly from $(0, 1]^{\#(x)+\#(Y)}$
4. **until** for all $a \in Y$, there exists $b \in x$ with $\text{dist}(a, b) \leq R_\theta$ and $t_b < t_a$
5. return Y and T

This is an example of a perfect simulation algorithm: the outputs come exactly from the desired distribution, but the running time is itself an unbounded random variable. The number of failure steps in the repeat loop has a geometric distribution with success parameter $p = \int_{(0,1]^{\#(x)}} \exp(A_\theta(x, t) - \lambda_\theta \mu(\mathcal{S})) dt$ and so the expected number of loop steps is $1/p$.

Unfortunately, for simple choices of reference measure such as Lebesgue, the acceptance probability p decreases (and the expected number of loop steps increases) geometrically in the number $\#(x)$ of seen points. Consider, for example, a rectangular region $\mathcal{S} \subseteq \mathbb{R}^d$ for some $d \in \mathbb{N}$, with sides of length at least R_θ and with periodic boundary conditions (to simplify the argument by eliminating edge effects). With Lebesgue reference measure $\mu(\cdot)$, the total volume of the region is $V = \mu(\mathcal{S})$, the product of the lengths of the sides.

Since balls of diameter R centered at the seen points x are disjoint, each with volume $v = v_d(R/2)^d$ in \mathbb{R}^d for $v_d := \pi^{d/2}/\Gamma(1+d/2)$, a lower bound on the *unshadowed* volume leads to an upper bound on the shadowed volume and, hence, on the acceptance probability:

$$\lambda_\theta V - A_\theta(x, t) \geq \lambda_\theta v \sum_{j=1}^{\#(x)} t_j$$

so

$$p := \int_{(0,1]^{\#(x)}} \exp(A_\theta(x, t) - \lambda_\theta V) dt \leq \int_{(0,1]^{\#(x)}} \exp\left(-\lambda_\theta v \sum_{j=1}^{\#(x)} t_j\right) dt = \rho^{\#(x)}, \quad (3.1)$$

where $\rho := [1 - \exp(-\lambda_\theta v)]/(\lambda_\theta v) < 1$. In the one-dimensional case ($d = 1$) we have $v_1 = 2$ and $p \leq \rho^{\#(x)}$ with $\rho = [1 - \exp(-\lambda_\theta R_\theta)]/(\lambda_\theta R_\theta)$, for example, while $v_2 = \pi$ and $\rho = [1 - \exp(-\lambda_\theta \pi R_\theta^2/4)]/(\lambda_\theta \pi R_\theta^2/4)$ in the two-dimensional case ($d = 2$).

A similar argument gives a lower bound for the shadow volume, giving a lower bound for p :

$$\begin{aligned} A_\theta(x, t) &\geq \lambda_\theta v \sum_{j=1}^{\#(x)} (1-t_j), \\ p &:= \int_{(0,1]^{\#(x)}} \exp(A_\theta(x, t) - \lambda_\theta V) dt \\ &\geq \int_{(0,1]^{\#(x)}} \exp\left(\lambda_\theta \left\{ v \sum_{j=1}^{\#(x)} (1-t_j) - V \right\}\right) dt \\ &= \rho^{\#(x)} \exp(-\lambda_\theta \{V - v\#(x)\}), \end{aligned}$$

so Algorithm 3.1 may be very reasonable for sufficiently small $\#(x)$ and $\lambda_\theta V$. The geometric decrease of (3.1) makes the basic form of acceptance–rejection untenable for large $\#(x)$; we now turn to an alternative.

4. Coupling from the past

Acceptance–rejection works well for small sets of seen points, but quickly becomes impractical as $\#(x)$ grows. In this section a new Markov chain for this problem is created that allows for perfect sampling when used with read-once coupling from the past (ROCFTP) [35].

Our coupling-from-the-past (CFTP) approach employs form 1 of the bounding chains of Huber [14], closely related to the antimonotonicity approach introduced by Kendall [16] and developed by Møller [19], and Häggström and Nelander [12]. The bounding chain view is as follows. Suppose that the state space for a Markov chain (X_t) is of the form $\Omega_X = \mathcal{T}^{\mathcal{I}}$ for a set of ‘labels’ \mathcal{T} and a set of ‘indices’ \mathcal{I} , that is, each state is an assignment of labels to the indices. The state space for the time marks is $(0, 1]^{1, \dots, \#(x)}$, so here $\mathcal{T} = (0, 1]$ and $\mathcal{I} = \{1, \dots, \#(x)\}$. Now consider a new state space $\Omega_B = (2^{\mathcal{T}})^{\mathcal{I}}$ so that each index $i \in \mathcal{I}$ is assigned a *subset* of labels rather than a single label. The bounding chain method runs a Markov chain (B_t) on this larger space that bounds the possible values for the original chain (X_t) .

Definition 4.1. Suppose that (A_t) and (B_t) are Markov chains, both adapted to the same filtration \mathcal{F}_t . Say that (A_t, B_t) is a *coupling* of the processes if, for all measurable A, B and states a, b ,

$$\begin{aligned} P(A_{t+1} \in A \mid A_t = a, \mathcal{F}_t) &= P(A_{t+1} \in A \mid A_t = a), \\ P(B_{t+1} \in B \mid B_t = b, \mathcal{F}_t) &= P(B_{t+1} \in B \mid B_t = b). \end{aligned}$$

A coupling is simply a construction of both Markov chains on the same probability space with the same filtration.

Definition 4.2. A Markov chain (B_t) with state space $(2^{\mathcal{T}})^{\mathcal{I}}$ is a *bounding chain* for the Markov chain (X_t) with state space $\mathcal{T}^{\mathcal{I}}$ if there exists a coupling (X_t, B_t) such that

$$(\text{for all } i \in \mathcal{I}) \quad (X_t(i) \in B_t(i)) \implies (\text{for all } i \in \mathcal{I}) \quad (X_{t+1}(i) \in B_{t+1}(i)).$$

For any $x \in \mathcal{T}^{\mathcal{I}}$ and $B \in (2^{\mathcal{T}})^{\mathcal{I}}$, write ‘ $x \in B$ ’ if $x(i) \in B(i)$ for all $i \in \mathcal{I}$. Now CFTP with bounding chains can be described. Suppose that the initial state of the chain is drawn from its stationary distribution, $X_0 \sim \pi$, and that the initial state of the bounding chain is $B_0 = \mathcal{T}^{\mathcal{I}}$, so $B_0(i) = \mathcal{T}$ for every $i \in \mathcal{I}$ and, hence, $X_0 \in B_0$. Then, after a block of any fixed number of steps \tilde{t} , by stationarity and boundedness we also have $X_{\tilde{t}} \sim \pi$ and $X_{\tilde{t}} \in B_{\tilde{t}}$.

After taking such a block of length \tilde{t} starting in states $X_0 \sim \pi$ and $B_0 = \mathcal{T}^{\mathcal{I}}$, there are two possibilities. Either $B_{\tilde{t}}$ contains exactly one state (namely $X_{\tilde{t}}$) or it contains more than one state. When $\#(B_{\tilde{t}}) = 1$, call the block a *success*, otherwise call it a *failure*. Each block succeeds or fails independently, and the probability p that any particular block is a success is the same for all blocks. The ROCFTP method of Wilson takes advantage of the following fact, first shown in a different form by Wilson [35]. For completeness, we present a proof similar in form to Wilson’s.

Lemma 4.1. Let $G \sim \text{Ge}(p)$ for some $0 < p \leq 1$ (so $P(G = j) = p(1-p)^j$ for integers $j \geq 0$). Let Y_0 have any distribution, and run the Markov chain forward where the first block is conditioned to be a success, and the next G blocks are conditioned to be failures. Then $Y_{\tilde{t}(1+G)} \sim \pi$.

Proof. Consider $X_0 \sim \pi$ and, for some large number $k \in \mathbb{N}$ to be specified later, take $k\tilde{t}$ steps in the Markov chain. Divide these steps into k blocks of length \tilde{t} . Since π is stationary, $X_{k\tilde{t}} \sim \pi$ as well.

Let G be the number of failure blocks following the *last* success block, if any, among the first k blocks, and otherwise set $G = k$. Then G is distributed as the minimum of k and a $\text{Ge}(p)$ random variable.

Let Y denote a random variable chosen by simulating a success block (which returns a single state in the bounding chain) followed by G failure blocks. Note that $Y_{\tilde{t}(1+G)} \sim Y$, regardless of the distribution of Y_0 . Let Z be a random variable found by starting with $Z_0 \sim \pi$, and then running k failure blocks. Now consider $X_{k\tilde{t}}$. Conditioned on $G < k$, we have $X_{k\tilde{t}} \sim [Y \mid G < k]$, while conditioned on $G = k$, $X_{k\tilde{t}} \sim Z$. In other words, for measurable A ,

$$P(X_{k\tilde{t}} \in A) = P(Z \in A)(1-p)^k + P(Y \in A \mid G < k)[1 - (1-p)^k].$$

Since π is stationary, $P(X_{k\tilde{t}} \in A) = \pi(A)$ for all k . By assumption, $p > 0$, so, as $k \rightarrow \infty$, the right-hand side converges to $P(Y \in A)$, completing the proof.

The coupling between chain and bounding chain is typically accomplished by use of update functions (see, e.g. [10]). Let (X_t) be a time homogeneous Markov chain with state space Ω . Let U be a uniform random number on $(0, 1]$. Call $\phi: \Omega \times (0, 1] \rightarrow \Omega$ an *update function* if, for all $x \in \Omega$ and measurable A , $P(\phi(x, U) \in A) = P(X_1 \in A \mid X_0 = x)$. The chain can be constructed recursively from a countable set $(U_t) \stackrel{\text{i.i.d.}}{\sim} \text{Un}(0, 1]$ by the recipe $X_0 = x$ and, for $t > 0$, $X_t = \phi(X_{t-1}, U_t)$.

A single uniform variate may be regarded as an infinite sequence of uniform $\{0, 1\}$ random bits, which in turn can be used to construct a countable number of new independent uniform random variates on $(0, 1]$. Therefore, the update function described here is very general, and includes any method for updating that uses at most a countable number of random draws. This generality will be useful when a Markov chain step requires a draw from a Poisson point process, which (as typically implemented) requires a Poisson-distributed number of uniform draws.

Any update function for a Markov chain will induce one for the bounding chain, and a coupling of the chains. Suppose that (B_t) is a bounding chain for process (X_t) with update function ϕ . Then, to update B_t given $U_{t+1} \sim \text{Un}(0, 1]$, use $B_{t+1} = \Phi(B_t, U_{t+1})$, with update function $\Phi(B_t, U) := \bigcup_{x \in B_t} \phi(x, U)$. The chains X_t and B_t are coupled, since both are adapted to $\mathcal{F}_t := \sigma\{U_s : s \leq t\}$.

Lemma 4.1 says that a sample may be generated from π by running a success block followed by a geometrically distributed number, G , of failure blocks. To determine G , run blocks successively. Once a success block has occurred, the state just prior to each subsequent success block will be a draw from the stationary distribution (see Figure 3). The following pseudocode




FIGURE 3: Read-once coupling from the past. Here S and F denote success and failure blocks, respectively.

will run the Markov chain through \tilde{t} steps to form a single block and indicate its success or failure.

Algorithm 4.1. (*Evaluate_block.*) **Inputs:** current state X , update function ϕ , block length \tilde{t} , and state space \mathcal{T}^I .

Outputs: X output of block and `success_block` \in {true, false}.

1. $B \leftarrow \mathcal{T}^I$
2. **for** t from 1 to \tilde{t} **do**
3. **draw** $U \leftarrow \text{Un}(0, 1]$
4. $X \leftarrow \phi(X, U)$
5. $B \leftarrow \Phi(B, U) := \bigcup_{x \in B} \phi(x, U)$
6. **end for**
7. `success_block` $\leftarrow (\#(B) = 1)$

With the block evaluation routine in hand, the following pseudocode runs ROCFTP. Begin by running blocks until a successful block is found. Since it is successful, the output of the block is a single state X . Save X in variable Y , and run forward through another block. If this block is a success then Y is the state immediately preceding the success block, and so Y is a draw from π . Otherwise, continue taking blocks until a success is found.

The time needed to generate n draws with this procedure will be the time needed to generate $n + 1$ successful blocks, with mean $(n + 1)/p$, where p denotes the success probability.

Algorithm 4.2. (*Read-once coupling from the past.*) **Inputs:** update function ϕ , number of samples n , block length \tilde{t} , and state space \mathcal{T}^I .

Outputs: $X_1, \dots, X_n \stackrel{\text{i.i.d.}}{\sim} \pi$.

1. **repeat**
2. $(X, \text{success_block}) \leftarrow \text{Evaluate_block}(\emptyset, \phi, \tilde{t}, \mathcal{T}^I)$
3. **until** `success_block`
4. **for** i from 1 to n **do**
5. **repeat**
6. $Y \leftarrow X$
7. $(X, \text{success_block}) \leftarrow \text{Evaluate_block}(X, \phi, \tilde{t}, \mathcal{T}^I)$
8. **until** `success_block`
9. $X_i \leftarrow Y$
10. **end for**

4.1. The update function for Matérn time marks

The key to success using ROCFTP is the construction of an update function $\phi(\cdot)$ for which it is easy to evaluate the induced update function $\Phi(\cdot)$ for the bounding chain.

The conditional density for the time marks, given the seen points x , was given in (2.4) for $t \in (0, 1]^{\#(x)}$ as $h_\theta(t \mid x) = Z_{\theta,x}^{-1} \exp(A_\theta(x, t))$, where (see (2.1)) $A_\theta(x, t) := \lambda_\theta \mu(D(x, t))$ is the expected number of hidden points in the shadow of configuration x with time marks t .

The shadow $D(x, t)$ is a union of overlapping cylinders, making exact computation of $A_\theta(x, t)$ very difficult even for Lebesgue measure $\mu(\cdot)$ in only two dimensions. A Metropolis–Hastings approach will be used to avoid the explicit computation of $A_\theta(x, t)$.

The basic Metropolis–Hastings procedure [13], [18] converts a proposal Markov chain into one whose stationary distribution is exactly the target distribution. Suppose that the target distribution has a density function $f(t)$ with respect to a reference measure $\eta(dt)$ and, when at point T , the proposal chain proposes a point S using density $q(T, s)$ with respect to $\eta(ds)$. Furthermore, suppose that $f(t)q(t, s) = 0$ if and only if $f(s)q(s, t) = 0$. Then the following procedure takes one step in the Metropolis–Hastings chain. Here $\text{Be}(r)$ denotes the Bernoulli distribution, which is equal to 1 with probability $r \wedge 1$ and 0 otherwise.

Algorithm 4.3. (*General_Metropolis_Hastings.*) **Input:** current state T .

Output: next state T .

1. **draw** S from distribution $q(T, s)\eta(ds)$
2. $r \leftarrow \{f(S)q(S, T)\}/\{f(T)q(T, S)\}$
3. **draw** $B \leftarrow \text{Be}(r)$
4. **if** $B = 1$ **then**
5. $T \leftarrow S$
6. **end if**

Our present goal in making inference about the Matérn type-III process is to draw samples of the time marks T , conditionally on the observed values of the seen points x . Note that the Metropolis–Hastings ratio, $r = \{f(S)q(S, T)\}/\{f(T)q(T, S)\}$, depends on the target density $f(\cdot)$ only up to a scale factor, so the conditional density $h_\theta(t \mid x)$ of the time marks (see (2.4)) may be replaced by the simpler unnormalized version $f(t) = \exp(A_\theta(x, t))$.

Consider a proposal chain in which an index $i \in \{1, \dots, \#(x)\}$ is chosen uniformly and then a new time mark t_{prop_i} is proposed for x_i from some distribution with symmetric density $q(s, t) = q(t, s)$. Call the resulting proposed set of time marks t_{prop} . By symmetry, the Metropolis–Hastings ratio becomes simply

$$r = \exp(A_\theta(x, t_{\text{prop}}) - A_\theta(x, t)). \tag{4.1}$$

Calculating the change in shadow area under the new time mark is a difficult task even in only two dimensions. Therefore, it is necessary to construct an approach that does not require direct computation of r in line 2 of Algorithm 4.3.

In the case that the proposed new time mark is smaller, $t_{\text{prop}} \leq t$, the shadow of x will grow ($A_\theta(x, t_{\text{prop}}) \geq A_\theta(x, t)$) and the Metropolis–Hastings ratio $r \geq 1$ will exceed unity, so the proposal will always be accepted and we could replace lines 2–3 of Algorithm 4.3 with ‘set $B = 1$ ’.

In the case that $t_{\text{prop}} > t$, the Hastings ratio r of (4.1) is *exactly* the probability that a Poisson point process with intensity Λ_θ will assign zero points to the newly unshadowed region $\Delta := D(x, t) \setminus D(x, t_{\text{prop}})$. Thus, in lieu of lines 2–3 of Algorithm 4.3, generate a

Poisson point process $W \sim \text{Po}(\Lambda_\theta)$ on Δ and set $B \leftarrow \mathbf{1}_{\{0=\#(W)\}}$. The set Δ never needs to be evaluated explicitly—generating a Poisson point process on Δ is a simple two-step process, as follows.

Each point in Δ lies within distance R_θ of x_i , and has a time mark between t_i and t_{prop_i} , so

$$\Delta := D(x, t) \setminus D(x, t_{\text{prop}}) \subseteq B_{R_\theta}(x_i) \times (t_i, t_{\text{prop}_i}].$$

Thus, the Poisson point process W can be generated by first generating such a process over $B_{R_\theta}(x_i) \times (t_i, t_{\text{prop}_i}]$ and then thinning it by retaining only points outside $D(x, t_{\text{prop}})$. If there are no such points then $B = 1$ and the proposal is accepted; otherwise $B = 0$ and the previous state is retained. This idea of using an auxiliary Poisson process to generate a Bernoulli distributed random variate without the need to calculate the parameter exactly also appears in [2], where it was used in the very different context of generating diffusion sample paths.

For later use in Section 4.2, we present as Algorithm 4.4, below, a method for drawing a Metropolis–Hastings update of the vector T of time marks for a configuration x of seen points.

Algorithm 4.4. (*Matern_III_Metropolis_Hastings.*) **Input:** seen points x with current state of time marks T .

Output: next state of time marks T .

1. **draw** i uniformly from $\{1, \dots, \#(x)\}$
2. $t \leftarrow T$
3. **draw** $t_i \leftarrow \text{Un}(0, 1]$
4. **if** $t_i \leq T_i$ **then**
5. $T_i \leftarrow t_i$
6. **else**
7. **draw** Poisson point process W on $B_{R_\theta}(x_i) \times (T_i, t_i]$ with intensity Λ_θ
8. **if** $\#(W \cap D(x, t)^c) = 0$ **then**
9. $T_i \leftarrow t_i$
10. **end if**
11. **end if**

Now consider the question of how to use this update function for a single state to update an entire bounding chain state. The bounding states for this chain will be the intervals $B_t = [T_t^{\text{lo}}, T_t^{\text{hi}}]$, where $B_t := [T_t^{\text{lo}}, T_t^{\text{hi}}] = \{T : (T_t^{\text{lo}}(i) \leq T(i) \leq T_t^{\text{hi}}(i) \text{ for all } i)\}$. Pseudocode for the bounding chain update is presented in Algorithm 4.5, below, followed by a brief explanation of why it works.

Algorithm 4.5. (*Matern_III_bounding_chain.*) **Inputs:** seen points x and current bounds $B = [T^{\text{lo}}, T^{\text{hi}}]$.

Output: next state of time mark bounds $B = [T^{\text{lo}}, T^{\text{hi}}]$.

1. **draw** i uniformly from $\{1, \dots, \#(x)\}$
2. $t^1 \leftarrow T^{\text{hi}}, t_i^1 \leftarrow T_i^{\text{lo}}$

3. $t^2 \leftarrow T^{\text{lo}}, t_i^2 \leftarrow T_i^{\text{hi}}$
4. **draw** $t_i \leftarrow \text{Un}(0, 1]$
5. $t^3 \leftarrow T^{\text{hi}}, t_i^3 \leftarrow t_i$
6. $t^4 \leftarrow T^{\text{lo}}, t_i^4 \leftarrow t_i$
7. **if** $t_i \leq T_i^{\text{lo}}$ **then** (case I)
8. $T_i^{\text{lo}} \leftarrow t_i, T_i^{\text{hi}} \leftarrow t_i$
9. **else if** $T_i^{\text{lo}} < t_i \leq T_i^{\text{hi}}$ **then** (case II)
10. $T_i^{\text{hi}} \leftarrow t_i$
11. **draw** Poisson point process W on $B_{R_\theta}(x_i) \times (T_i, t_i]$ with intensity Λ_θ
12. **if** $\#(W \cap \{D(x, t^1) \setminus D(x, t^3)\}) = 0$ **then**
13. $T_i^{\text{lo}} \leftarrow t_i$
14. **end if**
15. **else if** $T_i^{\text{hi}} < t_i$ **then** (case III)
16. **draw** Poisson point process W on $B_{R_\theta}(x_i) \times (T_i, t_i]$ with intensity Λ_θ
17. **if** $\#(W \cap \{D(x, t^1) \setminus D(x, t^3)\}) = 0$ **then**
18. $T_i^{\text{lo}} \leftarrow t_i$
19. **end if**
20. **if** $\#(W \cap \{D(x, t^2) \setminus D(x, t^4)\}) = 0$ **then**
21. $T_i^{\text{hi}} \leftarrow t_i$
22. **end if**
23. **end if**

To summarize, when the proposed time mark t_i is smaller than the existing one, always accept the proposal and collapse $[T_i^{\text{lo}}, T_i^{\text{hi}}]$ to (t_i) . When the proposed t_i is higher, the worst case for acceptance is when the original time marks are as small as possible (i.e. equal to T_i^{lo}), and, for all $j \neq i$, the time marks t_j are as high as possible (i.e. equal to T_j^{hi}). This is shown in the proof to the following theorem.

Theorem 4.1. *Algorithm 4.5 is a valid update for the bounding chain.*

Proof. Suppose that the call

$$(T'^{\text{lo}}, T'^{\text{hi}}) \leftarrow \text{Matern_III_bounding_chain}(x, T^{\text{lo}}, T^{\text{hi}})$$

is used for the bounding chain and that the call

$$T' \leftarrow \text{Matern_III_Metropolis_Hastings}(x, T)$$

is used for the time marks, where T satisfies $T^{\text{lo}} \leq T \leq T^{\text{hi}}$. Then with the same choices for i, t_i , and W as in Algorithm 4.5, we must show that $T'^{\text{lo}} \leq T' \leq T'^{\text{hi}}$. Three cases arise.

Case I: $t_i \leq T_i^{lo}$. Then $t_i \leq T_i^{lo} \leq T_i$ and, hence, $T'_i = t_i$, so setting $T_i^{lo} = T_i^{hi} = t_i$ still satisfies $T^{lo} \leq T' \leq T^{hi}$.

Case II: $T_i^{lo} < t_i \leq T_i^{hi}$. When $T_i \geq t_i$, $T'_i = t_i$, and when $T_i < t_i$, $T'_i \in \{t_i, T_i\} \leq t_i$, so $T_i^{hi} = t_i \geq T'_i$ is a valid move.

Now consider T^{lo} . When $t_i \leq T_i$, again $T'_i = t_i$, and so $T^{lo} = t_i \leq T_i$ is a valid move, so suppose that $t_i > T_i$.

The shadow change upon increasing T_i to t_i can also be written in the form

$$\Delta = \{D(x, T) \setminus D(x, t)\} = \{B_{R_\theta}(x_i) \times (T_i, t_i]\} \cap D(x, t)^c. \tag{4.2}$$

The set $B_{R_\theta}(x_i) \times (T_i, t_i]$ is as large as possible when T_i is as small as possible. The smallest it can be is T_i^{lo} .

The set $D(x, t)^c$ is as large as possible when $D(x, t)$ is as small as possible, which happens when, for all $j \neq i$, $t_j = T_j^{hi}$. Combining these, $\{D(x, T) \setminus D(x, t)\}$ is as large as possible when $T = t^1$ and $t = t^3$. In other words,

$$D(x, T) \setminus D(x, t) \subseteq D(x, t^1) \setminus D(x, t^3).$$

Now suppose that $\#(W \cap \{D(x, t^1) \setminus D(x, t^3)\}) = 0$. Then, necessarily,

$$\#(W \cap \{D(x, T) \setminus D(x, t)\}) = 0$$

as well for every $T \in [T^{lo}, T^{hi}]$, so the proposal $T' = t$ will be accepted and $T_i^{lo} = t_i \leq T'_i$ is a valid move in the bounding chain.

Case III: $T_i^{hi} < t_i$. The case for T^{lo} is the same as in case II. It remains to show that the move for T^{hi} is valid. Use representation (4.2) for Δ . This is as small as possible when $(T_i, t_i]$ is as small as possible, so $T_i = T_i^{hi}$. Similarly, $D(x, t)^c$ is as small as possible when $D(x, t)$ is as large as possible, which happens when $t_j = T_j^{lo}$ for all $j \neq i$. Therefore,

$$D(x, t^2) \setminus D(x, t^4) \subseteq D(x, T) \setminus D(x, t).$$

If $\#(W \cap \{D(x, t^2) \setminus D(x, t^4)\}) > 0$ then $\#(W \cap \{D(x, T) \setminus D(x, t)\}) > 0$ for all T satisfying $T^{lo} \leq T \leq T^{hi}$, so the proposal t will be rejected and $T_i^{hi} = T_i^{hi} \geq T'_i = T_i$ is fine. Otherwise, there might be some T that accepts the move, and so $T_i^{hi} = t_i \geq T'_i$, as in line 21 of Algorithm 4.5.

This completes the proof.

4.2. A product estimator of the likelihood

Given the ability to draw samples from π , there exist many Monte Carlo methods for estimating the likelihood $g(x \mid \theta)$. Here a version of the product estimator is presented (see [9, Section 5.31] for more details of this approach, and [26] for a more advanced variant). The product estimator algorithm has the advantage of giving approximations to the likelihood that are immune to the effects of multimodality.

Begin by fixing x and R , and consider a range of possible values of λ . Fix $0 < \gamma \leq 1/\mu(\mathcal{S})$ and, for integers $k \geq 0$, fix $\theta_k \in \Theta$ such that $\lambda_{\theta_k} = \tilde{\lambda}_k := \gamma k$ and $R_{\theta_k} = R$. Note that $A_{\theta_k} = kA_{\theta_1} = k\gamma(\mu \otimes m)\{D(x, t)\}$.

From (2.3) we have $g(x \mid \theta_0) = \mathbf{1}_{\{x=\emptyset\}}$ and, for $k \geq 1$,

$$g(x \mid \theta_k) = a_k \tilde{\lambda}_k^{\#(x)} \exp(\mu(S)[1 - \tilde{\lambda}_k]),$$

where

$$a_k := \int_{(0,1]^{\#(x)}} \exp(A_{\theta_k}(x, t)) dt = \int_{(0,1]^{\#(x)}} \exp(kA_{\theta_1}(x, t)) dt.$$

Evidently, $a_0 = 1$; for $k \geq 1$, estimate a_k by first estimating each ratio in the telescoping product

$$\frac{a_0 a_1 a_2 \dots a_{k-1}}{a_1 a_2 a_3 \dots a_k} = \frac{1}{\int_{(0,1]^{\#(x)}} \exp(A_{\theta_k}(x, t)) dt}.$$

Estimate each term (a_{i-1}/a_i) as follows. First, given the seen points x and radius R , fix n (to be chosen later) and, using Algorithm 4.4 (or 3.1), draw time marks $T_1, \dots, T_n \stackrel{\text{i.i.d.}}{\sim} h_{\theta_i}(t | x) dt$ (see (2.4)) for intensity $\tilde{\lambda}_i = i\gamma$. For each $1 \leq j \leq n$, draw a Poisson point process W_j with intensity $\gamma(\mu \otimes m)(ds dt)$ over the shadow $D(x, T_j)$. The independent Bernoulli random variables $B_j = \mathbf{1}_{\{\#(W_j)=0\}}$ have conditional expectation $E(B_j | T_j) = P(\#(W_j) = 0 | T_j) = \exp(-A_{\theta_1}(x, T_j))$, so

$$E(B_j) = \frac{\int_t \exp(-A_{\theta_1}(x, t) + iA_{\theta_1}(x, t)) dt}{\int_t \exp(iA_{\theta_1}(x, t)) dt} = \frac{a_{i-1}}{a_i}.$$

The probability $p_i = (a_{i-1}/a_i) = \exp(-\gamma\mu(\mathcal{S}))$ that W is empty is at least e^{-1} , since $\gamma \leq 1/\mu(\mathcal{S})$, so the Bernoulli variables (B_j) have means $E(B_j) = p_i > e^{-1}$ and coefficients of variation

$$\text{cov}(B_j) = \frac{\sqrt{p_i(1-p_i)}}{p_i} = \sqrt{\frac{1-p_i}{p_i}} \leq \sqrt{e-1}.$$

Set $\bar{B}_i := (B_1 + \dots + B_n)/n$. Then $E(\bar{B}_i) = (a_{i-1}/a_i)$ and $\text{cov}(\bar{B}_i)^2 \leq (e-1)/n$. Let $\hat{a}_i := 1/\prod_{j=1}^i \bar{B}_j$ be an estimate for a_i . Dyer and Frieze [6, pp. 135–136] showed that under these conditions, for any $\varepsilon > 0$, the condition $n \geq 16(e-1)i/\varepsilon^2$ is sufficient to guarantee that

$$P\left(\left|\frac{1}{\hat{a}_i} - \frac{1}{a_i}\right| > \frac{\varepsilon}{a_i}\right) \leq \frac{1}{4}.$$

The event $(1-\varepsilon)/a_i \leq 1/\hat{a}_i \leq (1+\varepsilon)/a_i$ is equivalent to $a_i(1+\varepsilon)^{-1} \leq \hat{a}_i \leq a_i(1-\varepsilon)^{-1}$. For $\varepsilon < 0.1$, $(1-\varepsilon/1.1)^{-1} \leq 1+\varepsilon$ and $(1-\varepsilon) \leq (1+\varepsilon)^{-1}$ always. Therefore, after $n \geq 16(e-1)i(1.1/\varepsilon)^2 = 19.36(e-1)i/\varepsilon^2$ steps, when $\varepsilon \leq 0.1$,

$$P(|\hat{a}_i - a_i| > \varepsilon a_i) \leq \frac{1}{4}.$$

The upper bound of $\frac{1}{4}$ can be made arbitrarily small by the standard method of repeating the process n' times and taking the median result, as follows. Let \hat{m} denote this median. For $\frac{1}{4} > \delta > 0$, Chernoff bounds [4] can be used to show that, for $n' \geq 10.4 \ln \delta^{-1}$,

$$P(|\hat{m} - a_i| > \varepsilon a_i) \leq \delta.$$

The resulting method is presented in pseudocode as Algorithm 4.6, below, where, as usual, $[x]$ denotes the greatest integer less than or equal to a real number $x \in \mathbb{R}$.

Algorithm 4.6. (*Matern_III_product_estimator.*) **Inputs:** relative accuracy $\varepsilon \in (0, 0.1)$, failure probability δ , steps k , and step size γ .

Output: estimate \hat{m} of a_k .

1. $n \leftarrow \lceil 19.36(e-1)k/\varepsilon^2 \rceil, n' \leftarrow \lceil 10.4 \ln \delta^{-1} \rceil$
2. **for** ℓ from 1 to n' **do**
3. **for** i from 1 to k **do**

4. **for** j from 1 to n **do**
5. **draw** time marks T_j using $\tilde{\lambda} = i\gamma$
6. **draw** Poisson point process W_j with intensity $\gamma(\mu \otimes m)$ over $\mathcal{S} \times (0, 1]$
7. **end for**
8. $\bar{B}_\ell \leftarrow (1/n) \sum_{j=1}^n \mathbf{1}_{\{\#(W_j \cap D(x, T_j))=0\}}$.
9. **end for**
10. $m_\ell \leftarrow \prod_{i=1}^k \bar{B}_\ell^{-1}$
11. **end for**
12. $\hat{m} \leftarrow \text{median of } \{m_1, \dots, m_{n'}\}$

Some remarks are in order.

- Upon replacing δ by $\delta' = \delta/k$, the Bonferroni bound ensures that the probability that *all* of the estimates a_j for j from $1, \dots, k$ are correct to within a factor of $1+\varepsilon$ will be at least $1-\delta$. Typically, the goal is to develop a profile of the posterior across values of λ , which is why all of these products are being estimated rather than using the more efficient method for a single estimate of [26].
- The user interested in a specific value of $\lambda > 0$ may either set $\gamma = \lambda/k$ for $k = \lceil \lambda\mu(\mathcal{S}) \rceil$ or take $\gamma = 1/\mu(\mathcal{S})$ and separately estimate a_λ/a_{k-1} by the same methods, still with expectation at least $\exp(-1)$.

5. Simulated data

The technique was tested on four data sets, drawn exactly from the Matérn III point process on the unit square \mathcal{S} . These data sets are shown in Figure 4. Sets (a) and (c) have $\lambda = 50$, while (b) and (d) have $\lambda = 100$. Sets (a) and (b) have $R = 0.025$ and sets (c) and (d) have $R = 0.05$. These parameter values were chosen to be comparable to those found in scaled data sets studied in [11] and [30]. Data sets (a), (b), (c), and (d) contain 40, 42, 82, and 80 points, respectively.

In Figure 5 the likelihood functions for the four data sets are presented. The parameter space is two-dimensional, and so the likelihood function is presented using three slices at different values of R . In each subfigure of Figure 5 the first value of R is the largest possible for the data set (the minimum interpoint distance). Unsurprisingly, the likelihood appears to be close to that of a normal distribution—however, the peak is shifted to the right by an amount that depends on R . Note that at $R = 0$ the Matérn III coincides with the standard Poisson process model, whose likelihood attains a maximum at a value of λ equal to the observed number of points.

Now consider the maximum likelihood estimator for these data sets. The value of R that maximizes the likelihood is always the minimum distance between pairs of points in the data set, since the area of the shadow is an increasing function of R . For the data sets with $R = 0.025$, the maximum likelihood occurred at values of λ only slightly higher than those for a basic Poisson point process. The data set with $(\lambda, R) = (50, 0.025)$ has 40 points and an MLE of $(42.5, 0.034)$, while that with $(\lambda, R) = (100, 0.025)$ has 82 data points and an MLE of $(91.3, 0.028)$.

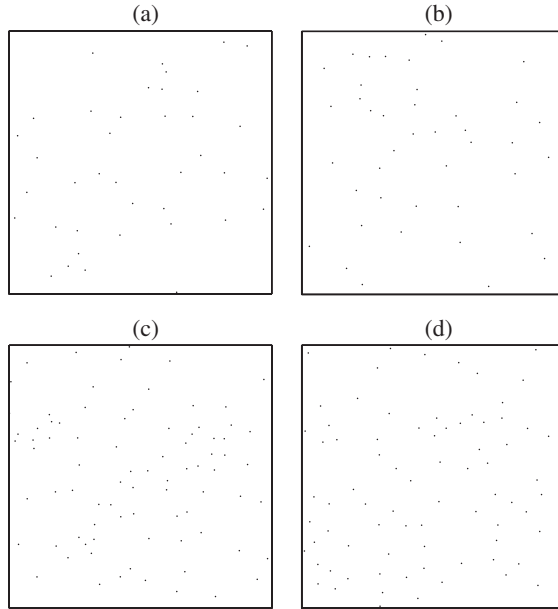


FIGURE 4: Simulated data sets.




FIGURE 5: Slices of likelihood functions for simulated data.




FIGURE 6: Running times for different intensities λ .

For the data sets with $R = 0.05$, the difference in the MLE between the basic Poisson process and Matérn III is substantial. The first data set had only 42 data points with $\lambda = 50$, but the MLE occurs at $(50.1, 0.051)$. The second data set with $(\lambda, R) = (100, 0.05)$ had its MLE at $(113.5, 0.051)$.

5.1. Running times

Figure 6 shows the average number of Markov chain steps evaluated per sample for the four data sets, based on 10 000 generated samples. For small values of λ , the number of points in the sample appears to determine the running time. The slopes of the time appear to be a function of R . This illustrates another advantage of modeling with Matérn III processes: the Markov chain appears to be rapidly mixing for all values of λ , in sharp contrast to methods based on pair repulsion models [15], which tend to have phase transitions that make sampling difficult for high values of λ_θ . This is in part a direct result of features of the model: for simulation-based inference with pair repulsion models, it is necessary to move point locations around. This is not needed with Matérn III, where the locations are fixed in the Markov chain and only the latent time marks move. Because the state space is of a completely different type, the issues associated with phase transitions and slow mixing of the Markov chain do not arise.

6. Discussion

Despite notable recent progress in the study of Gibbs random fields [20], [21], the problem of making likelihood-based inference for repulsive point processes remains challenging. In the absence of calculable likelihood functions, the common approach to inference is to simulate random samples from a specified class of processes and compute summary statistics such as the first-order packing density and second-order K , L , or g functions, and to compare these with similar statistics computed from data. Similar values of these statistics (or similar features of their plots) are regarded as evidence of model fit, while the inability to reconstruct features of the observed data is regarded as evidence against the suitability of a particular model family. Our goal has been to support quantitative, likelihood-based inference for some family of repulsive point processes to enable investigators to generate both Bayesian and maximum likelihood

estimators of model parameters, or to compute Bayesian predictive distributions of as yet unobserved quantities, within either of the present dominant paradigms of statistical inference.

In this work we have developed and presented an algorithm for the perfect simulation of latent time marks for data under the Matérn III model, and an algorithm for using these estimates to construct a product estimator for approximating the likelihood function to arbitrarily high accuracy. There remain several extensions of this method to be considered.

- *Soft-core processes.* Here R_θ is constant for all points, but the model becomes far more flexible when R_θ is allowed to vary from point to point.
- In a similar vein, λ_θ can be allowed to vary across the space of interest. This can be accomplished (for instance) by treating the intensity as a linear combination of basis functions with unknown coefficients.
- *Improved proposal distribution for time marks.* The Metropolis proposals for time marks used in Section 4.1 were uniformly distributed on the unit interval. Since the marginal distributions of time marks for seen points are skewed, typically smaller than those for the hidden points, it may be more efficient to sample using a proposal distribution that places more weight towards lower times. Better proposals might improve the mixing properties of the Markov chain, while having the deleterious effect of making the perfect simulation algorithm more complex.

Acknowledgements

The authors would like to thank Katja Ickstadt for early conversations about inhibitory processes, and Emery Brown, Rob Kass, Fabio Rigat, Jianfeng Feng, and Enrico Rossoni for helpful discussions about spike train data. This work was supported by US National Science Foundation grants DMS–0112069, DMS–0548153, DMS–0635449, and DMS–0757549. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

- [1] BERGER, J. O. AND WOLPERT, R. L. (1984). *The Likelihood Principle* (IMS Lecture Notes Monogr. Ser. **6**), 2nd edn. Institute of Mathematical Statistics, Hayward, CA.
- [2] BESKOS, A., PAPANASTASIIOPOULOS, O. AND ROBERTS, G. O. (2006). Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli* **12**, 1077–1098.
- [3] BROWN, T. (1979). Position dependent and stochastic thinning of point processes. *Stoch. Process. Appl.* **9**, 189–193.
- [4] CHERNOFF, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Statist.* **23**, 493–507.
- [5] DÖGE, G. (2001). Perfect simulation for random sequential adsorption of d -dimensional spheres with random radii. *J. Statist. Comput. Simul.* **69**, 141–156.
- [6] DYER, M. AND FRIEZE, A. (1991). Computing the volume of convex bodies: a case where randomness provably helps. In *Probabilistic Combinatorics and Its Applications* (Proc. Symp. Appl. Math. **44**), ed. B. Bollobás, American Mathematical Society, RI, pp. 123–169.
- [7] FEDER, J. (1980). Random sequential adsorption. *J. Theoret. Biol.* **87**, 273–254.
- [8] FINEGOLD, L. AND DONNELL, J. T. (1979). Maximum density of random placing of membrane particles. *Nature* **278**, 443–445.
- [9] FISHMAN, G. E. (1996). *Monte Carlo*. Springer, New York.
- [10] FOSS, S. G. AND TWEEDIE, R. L. (1998). Perfect simulation and backward coupling. *Commun. Statist. Stoch. Models* **14**, 187–203.
- [11] GLASS, L. AND TOBLER, W. R. (1971). Uniform distribution of objects in a homogeneous field: cities on a plain. *Nature* **233**, 67–68.

- [12] HÄGGSTRÖM, O. AND NELANDER, K. (1999). On exact simulation of Markov random fields using coupling from the past. *Scand. J. Statist.* **26**, 395–411.
- [13] HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109.
- [14] HUBER, M. (2004). Perfect sampling using bounding chains. *Ann. Appl. Prob.* **14**, 734–753.
- [15] KELLY, F. P. AND RIPLEY, B. D. (1976). A note on Strauss's model for clustering. *Biometrika*, **63**, 357–360.
- [16] KENDALL, W. S. (1998). Perfect simulation for the area-interaction point process. In *Probability Towards 2000* (New York, 1995; Lecture Notes Statist. **128**), eds L. Accardi and C. C. Heyde, Springer, New York, pp. 218–234.
- [17] MATÉRN, B. (1986). *Spatial Variation* (Lecture Notes Statist. **36**), 2nd edn. Springer, New York.
- [18] METROPOLIS, N. C. *et al.* (1953). Equations of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092.
- [19] MØLLER, J. (1999). Perfect simulation of conditionally specified models. *J. R. Statist. Soc. B* **61**, 251–264.
- [20] MØLLER, J., PETTITT, A. N., REEVES, R. AND BERTHELSEN, K. K. (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika*, **93**, 451–458.
- [21] MURRAY, I., GHAAHRAMANI, Z. AND MACKAY, D. J. C. (2006). MCMC for doubly-intractable distributions. In *Proc. 22nd Ann. Conf. on Uncertainty in Artificial Intelligence*, eds R. Dechter and T. Richardson, AUAI Press, Arlington, VA, pp. 359–366.
- [22] PÁLASTI, I. (1960). On some random space filling problems. *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 353–359.
- [23] RIPLEY, B. D. AND KELLY, F. P. (1977). Markov point processes. *J. London Math. Soc.* **15**, 188–192.
- [24] SCHAAF, P., VOEGEL, J.-C. AND SENGER, B. (1998). Irreversible deposition/adsorption processes on solid surfaces. *Annales de Physique* **23**, 1–89.
- [25] SOLOMON, H. (1967). Random packing density. In *Proc. 5th Berkeley Symp. Math. Statist. Prob.* Vol. 3, eds L. M. Le Cam and J. Neyman, University of California Press, Berkeley, CA, pp. 119–134.
- [26] ŠTEFANKOVIČ, D., VEMPALA, S. AND VIGODA, E. (2009). Adaptive simulated annealing: a near-optimal connection between sampling and counting. *J. Assoc. Comput. Mach.* **56**, 36 pp.
- [27] STOYAN, D. (1979). Interrupted point processes. *Biometrical J.* **21**, 607–610.
- [28] STOYAN, D. AND STOYAN, H. (1985). On one of Matérn's hard-core point process models. *Math. Nachr.* **122**, 205–214.
- [29] STOYAN, D., KENDALL, W. S. AND MECKE, J. (1995). *Stochastic Geometry and Its Applications*, 2nd edn. John Wiley, Chichester.
- [30] STRAND, L. (1972). A model for stand growth. In *IUFRO 3rd Conf. Advisory Group of Forest Statisticians*, Institut National de la Recherche Agronomique, Paris, pp. 207–216.
- [31] STRAUSS, D. J. (1975). A model for clustering. *Biometrika*, **62**, 467–475.
- [32] TANEMURA, M. (1979). On random complete packing by discs. *Ann. Inst. Statist. Math. B* **31**, 351–365.
- [33] VAN LIESHOUT, M. N. M. (2006). Maximum likelihood estimation for random sequential adsorption. *Adv. Appl. Prob.* **38**, 889–898.
- [34] VON NEUMANN, J. (1951). Various techniques used in connection with random digits. *NBS Appl. Math. Ser.* **12**, 36–38.
- [35] WILSON, D. B. (2000). How to couple from the past using a read-once source of randomness. *Random Structures Algorithms* **16**, 85–113.