

Bayesian inference for multivariate point processes observed at sparsely distributed times

Rasmussen, Jakob Gulddahl; Møller, Jesper; Aukema, B.H.; Raffa, K.F.; Zhu, J.

Publication date:
2006

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Rasmussen, J. G., Møller, J., Aukema, B. H., Raffa, K. F., & Zhu, J. (2006). *Bayesian inference for multivariate point processes observed at sparsely distributed times*. Department of Mathematical Sciences, Aalborg University. Research Report Series No. R-2006-24

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

**Bayesian inference for
multivariate point processes
observed at sparsely distributed times**

by

J.G. Rasmussen, J. Møller, B.H. Aukema, K.F. Raffa
and J. Zhu

R-2006-24

July 2006

DEPARTMENT OF MATHEMATICAL SCIENCES
AALBORG UNIVERSITY

Fredrik Bajers Vej 7 G ■ DK-9220 Aalborg Øst ■ Denmark

Phone: +45 96 35 80 80 ■ Telefax: +45 98 15 81 29

URL: <http://www.math.aau.dk>



Bayesian inference for multivariate point processes observed at sparsely distributed times

J. G. Rasmussen and J. Møller

Department of Mathematical Sciences, Aalborg University
jgr@math.aau.dk and jm@math.aau.dk

B. H. Aukema

Canadian Forest Service and University of Northern British Columbia
baukema@nrcan.gc.ca

K. F. Raffa

Department of Entomology, University of Wisconsin - Madison
raffa@entomology.wisc.edu

J. Zhu

Department of Statistics, University of Wisconsin - Madison
jzhu@stat.wisc.edu

Abstract

We consider statistical and computational aspects of simulation-based Bayesian inference for a multivariate point process which is only observed at sparsely distributed times. For specificity we consider a particular data set which has earlier been analyzed by a discrete time model involving unknown normalizing constants. We discuss the advantages and disadvantages of using continuous time processes compared to discrete time processes in the setting of the present paper as well as other spatial-temporal situations.

Keywords: Bark beetle, conditional intensity, forest entomology, Markov chain Monte Carlo, missing data, prediction, spatial-temporal process.

1 Introduction

This paper concerns statistical inference of spatial-temporal processes that are observed on a spatial lattice but only at sparsely distributed time points. Zhu *et al.* (2006) analyzed

such a data set using a spatial-temporal autoregressive type of model, which assumes that time is discrete and coincides with the observation times. Here we propose an alternative continuous time model, based on multivariate point processes (Daley and Vere-Jones, 2003), and develop Bayesian inference for estimating the model parameters and the times of events. The proposed methodology is illustrated by a subset of the data set featured in Zhu *et al.* (2006), but we do not attempt to address all the scientific questions in that paper. Instead the focus is on the methodology and our conclusion is that a multivariate point process model may have several advantages compared to a spatial-temporal autoregressive type of model.

The data set in Zhu *et al.* (2006) is from a study of a plantation of red pines located near Spring Green, Wisconsin, USA. Each tree was examined annually, and three types of data were recorded. The tree’s condition (alive or dead) was recorded from 1986 to 1992. The number of *Dendroctonus valens* (LeConte), a bark beetle hereafter called “turpentine beetle” that attacks the base of the tree, was recorded from 1987 to 1992. The presence or absence of *Ips* spp. (predominantly *Ips pini* (Say) and to a lesser extent *Ips grandicollis* (Eichhoff)), another bark beetle that mass attacks the main stem, was also recorded from 1987 to 1992. The turpentine beetle has one generation per year, with new attacks occurring from late April through June, and each beetle attacking only one tree (Furniss and Carolin, 1980). *Ips* spp. have two to three generations per year, depending on temperature. They become active in early May and dormant in September. Again, each beetle attacks only one tree. The primary objectives in Zhu *et al.* (2006) were quantifying the relations between the two types of beetles and the condition of the trees, as well as capturing the spatial and temporal structure of colonization by the two beetle types and the condition of the trees. Of special interest was the fact that a large gap of dead trees appeared.

In the present paper, we model *Ips* spp., using the turpentine beetles as a covariate. Furthermore, we only consider the area immediately around the gap, where Figure 1 shows the *Ips* spp. data. Focusing on a subset of the trees leads of course to loss of important biological information, but since we want to illustrate the methodological aspects, we do not aim at an overly complex model. We specify a multivariate point process model for *Ips* spp., using a Bayesian setting where we regard what happens between the observation times as missing data. For short we refer to the multivariate point process model as a continuous time model and the spatial-temporal autoregressive type of model in Zhu *et al.* (2006) as a discrete time model.

The paper is organized as follows. Section 2 introduces the needed notation. Section 3 specifies the multivariate point process model and prior assumptions, and Section 4 discusses simulation based Bayesian inference. Section 5 concludes with a comparison between continuous time and discrete time models for the specific models considered in Zhu *et al.* (2006) and the present paper as well as for more general models.



Figure 1: White squares are locations without trees or trees that were already dead in 1986, and gray squares are trees that were alive in 1986. The numbers indicate which year a tree has been attacked by *Ips* spp.; trees without numbers have not been attacked during the period of observation.

2 Notation

The data were collected during autumn, when *Ips* spp. were dormant until the beginning of the next attack period the following spring. For convenience we can therefore assume that the data were observed at the times $k = 0, 1, \dots, 5$ which correspond to the end of the years 1987, 1988, \dots , 1992. Further, we let $k = -1$ correspond to the end of year 1986, and we say that time t is in year k if $k - 1 < t \leq k$.

We let $i = 1, \dots, 807$ index the sites, i.e. the locations with living trees at time -1 . Let $y_{t,i} = 0$ if site i has not been previously attacked by *Ips* spp. at time t , $y_{t,i} = 1$ if it has been attacked earlier in the same year, and $y_{t,i} = 2$ if it has been attacked in a previous year. By an “event” at the tree i we mean a transition of the zero-one process $v_{t,i} = \mathbf{1}[y_{t,i} \geq 1]$, where $\mathbf{1}[\cdot]$ denotes the indicator function. We denote the time of this transition t_i . We do not consider a transition $1 \rightarrow 2$ for $y_{t,i}$ as an event, since it is certain that this transition happens at the end of the year at which the event took place. Note that there is a one-to-one correspondence between the process $v_{s,i}$ for $s < t$ (or $s \leq t$) and the process $y_{s,i}$ for $s < t$ (or $s \leq t$).

The process $v_s = (v_{s,1}, \dots, v_{s,807})$ is a particular kind of multivariate point process (or counting process), where each $v_{t,i}$ is restricted to be either zero or one. Such a process is specified by the conditional intensity function (Daley and Vere-Jones, 2003): For each tree i , given the history of the process v_s for times $s < t$, let

$$\lambda_{t,i} = \mathbf{E}[dv_{t,i} | (v_s)_{s < t}] / dt$$

denote the conditional intensity of the tree being attacked by *Ips* spp. In the next section we specify models for the conditional intensity, allowing $\lambda_{t,i}$ to depend on covariate information $(x_{s,i})_{s < t}$, where $x_{t,i}$ denotes the number of turpentine beetles at time t and site i . It will also depend on external information related to *Ips* spp. activity and on “neighboring information”. For a site i , consider its first-, second-, \dots order neighbors, which are the (up to) four nearest, four second nearest, \dots sites to i , and let N_i denote the set of first to fifth order neighbors of i . Finally, let $u_{t,i} = \mathbf{1}[y_{t,i} = 1]$ and let $n_{t,i} = \sum_{j \in N_i} u_{t,j}$ be the number of neighbors of i that are infested with *Ips* spp. by time t in the same year.

3 Model

As usual, $t-$ means the time just before time t . Assume that for t in year k ,

$$\lambda_{t,i} = \mathbf{1}[v_{t-,i} = 0] \rho(t) \left\{ \psi_0 + \psi_1 n_{t-,i}^{\alpha_1} + \psi_2 n_{(k-1)-,i}^{\alpha_2} + \psi_3 x_{k-,i} \right\} \quad (1)$$

where ρ is a non-negative function, $\psi_0, \psi_1, \psi_2, \psi_3$ are non-negative parameters, and $\alpha_1 = \alpha_2 = 2$ (this choice and alternative models are discussed at the end of this section). The

term $\mathbf{1}[v_{t-,i} = 0]$ is included, since *Ips* spp. do not attack the same tree twice. The meaning of the other terms in (1) is given below.

The function ρ incorporates external information about *Ips* spp. activity due to seasonal variation. Figure 2 shows ρ and reflects the fact that *Ips* spp. has a window of activity about four-five months and normally peak around July. Specifically, for t in year k , $\rho(t) = \varphi((t - \mu_k)/\sigma_k)$ where φ is the standard normal density function and the parameters μ_k and σ_k are determined as follows. Aukema *et al.* (2005) modeled *Ips* spp. activity as the number of *Ips* spp. caught in traps every week during the flight period in 2001–2002, using a linear regression model with various explanatory variables, but only the temperature is available in our study. Therefore we refit the model with temperature as the only explanatory variable, and estimate μ_k and σ_k by the empirical mean and standard deviation obtained from predicting the number of *Ips* spp. that would have been caught during each week of the k th year. Since $\mu_k - k$ and σ_k do not depend greatly on k , the five normal densities in Figure 2 look rather similar relative to the years.




Figure 2: The function ρ .

The function ρ is the same for all sites, and is a rough description of *Ips* spp. activity depending only on the temperature. The term $\{\dots\}$ in (1) adjusts for this. The individual terms in $\{\dots\}$ play the following roles. The term ψ_0 is included since the other terms can be zero and we need then to scale the function ρ . The term $\psi_1 n_{t-,i}^{\alpha_1}$ appears, since *Ips* spp. may emerge from a tree attacked earlier in the year to attack nearby trees. The term $\psi_2 n_{(k-1)-,i}^{\alpha_2}$ appears, since *Ips* spp. overwinter in the ground close to a previously attacked tree and emerge to attack nearby trees in the following year. The term $\psi_3 x_{k-,i}$ reflects that *Ips* spp. tend to attack trees that previously have been attacked by turpentine beetles. Since we have only observed which trees the turpentine beetles have attacked at the end of the year, we should ideally also model the turpentine beetles as a continuous time process. However, to keep the model simple we refrain from doing that, and instead assume that the turpentine beetles contribute to the conditional intensity throughout the year. Since turpentine beetles usually attack earlier in the year than *Ips* spp. and $\rho(t)$ is close to being zero early in the year, this is probably not so unrealistic.

For each year k , the processes $(v_t)_{t < k}$ and $(v_t)_{t > k}$ are conditionally independent given v_k (however, v_t is not a continuous time Markov process). We therefore condition on v_0 and consider the likelihood function based on the remaining data $d = (v_t)_{0 < t \leq 5}$ (i.e. including the missing data on the time interval $[0, 5]$). Letting $\psi = (\psi_0, \psi_1, \psi_2, \psi_3)$, the likelihood is

$$L(\psi; d|v_0) = \prod_{k=1}^5 \left[\prod_{i: k-1 < t_i \leq k} \lambda_{t_i, i} \right] \exp \left(- \int_{k-1}^k \sum_i \lambda_{t, i} dt \right). \quad (2)$$

Furthermore, we specify an improper uniform prior for ψ on $[0, \infty)^4$. Thus the conditional density of ψ given d (and v_0) is proportional to the likelihood (2). A rigorous proof that this conditional distribution of ψ is proper seems difficult, but from a practical point of view we would expect the MCMC runs described in Section 4.1 to diverge if the distribution was improper. This is not the case, and alternatively we could replace the range of ψ by a very large but bounded region.

Actually, before considering the model (1), we analyzed the model with

$$\lambda_{t, i} = \rho(t) \exp \left(\psi_0 + \psi_1 n_{t-, i} + \psi_2 n_{(k-1)-, i} + \psi_3 x_{k-, i} \right),$$

where $\psi_0, \psi_1, \psi_2, \psi_3$ are real parameters. This model is of a somewhat similar form as the model in Zhu *et al.* (2006), but a model check along similar lines as in Section 4.2 showed that the model did not fit the data well. Partly inspired by the form of the Hawkes process (Hawkes, 1971; Daley and Vere-Jones, 2003), we then turned to the model (1) but with $\alpha_1 = \alpha_2 = 1$. We observed a misfit and considered therefore alternative values $\alpha_1, \alpha_2 \in \{0, 1, 2\}$. We finally concluded that the model (1) with $\alpha_1 = \alpha_2 = 2$ fit best.

4 Inference

4.1 Posterior simulation and estimation

The posterior distribution is the conditional joint distribution of ψ and the missing data on $[0, 5]$ given the observations y_0, y_1, \dots, y_5 . We use a Metropolis-within-Gibbs algorithm to simulate from the posterior distribution — we assume that the reader is familiar with Markov chain Monte Carlo (MCMC) methods; see e.g. Robert and Casella (2004).

Specifically, we propose to update each of the four parameters $\psi_0, \psi_1, \psi_2, \psi_3$ one at a time using Metropolis random walk steps with normal proposal distributions, where the proposal variances are chosen to reach an average acceptance ratio of approximately 0.25 (Roberts *et al.*, 1997). Moreover, within a given year k at a site i either one event $t_i \in (k-1, k]$ has happened, where we do not know the exact value of t_i , or nothing has happened. If an event has happened at site i , we need to simulate the unknown t_i from its conditional distribution given “everything else”, i.e. from the density proportional

to $\lambda_{t,i}$, with $t \in (k - 1, k]$. We do this by visiting all the sites with events in some pre-determined order, updating t_i by an independent Metropolis sampler with proposal density $\rho(t)$, $t \in (k - 1, k]$.

Figure 3 shows plots of the posterior distributions of $\psi_0, \psi_1, \psi_2, \psi_3$ based on an MCMC run length of 100,000 with a burn-in length of 1,000. Note that all of the parameters are clearly bounded away from zero. Obviously, since ψ_1 respective ψ_2 is significant, the probability that a tree is going to be attacked increases when more trees in the neighborhood have been attacked earlier in the year respective in the previous year. This leads to several potential biological mechanisms. One is that once a beetle enters a tree and emits aggregation pheromones (Wood, 1982), large numbers of beetles from previously attacked neighboring trees are available to respond and hence rapidly exhaust the tree's defenses (Raffa and Berryman, 1983). Second, factors that predispose trees to being susceptible to attack may be distributed in a highly clustered fashion, and the resident population of beetles again respond to the first entries thereby generating the observed pattern (Erbilgin and Raffa, 2003). Obviously these are not mutually exclusive. Finally, the significance of ψ_3 in (1) implies that trees are more susceptible to being attacked by *Ips* spp. if they have been attacked by turpentine beetles previously.




Figure 3: Posterior distributions of $\psi_0, \psi_1, \psi_2, \psi_3$.

Figure 4 shows estimated results for the missing event times t_i . The results are based on an MCMC run length of 100,000 with a burn-in length of 1,000 and sampling every 25th missing data set. The first row in the figure shows for each year 1–5 the estimated mean value of the t_i , using a gray scale where white means that no event happened at the site and darker values correspond to early events. During the five years the missing data are located further and further away from the initial gap of dead trees, and within each year the mean values tend to be larger at locations further away from this gap. The second row in Figure 4 shows the standard deviations of the event times, where dark values correspond to small standard deviations. For each year, the standard deviations are about twice as large for isolated attacks than for those occurring in clumps. Since the degree of clumping varies from year to year, the standard deviations in e.g. year 3 are

higher than in the other years. The histograms in the last row in Figure 4 show for each year 1–5 the empirical distribution of all times of events during the year. The histograms are rather similar except for the difference in years, and they look much like the normal densities given by $\rho(t)$ except for a shift about half a month to the right.




Figure 4: Columns: years 1, 2, 3, 4, 5. First row: gray scale plots of estimated mean event times. Second row: gray scale plots of estimated standard deviations of event times. Last row: histograms of all event times and $\rho(t)$ within each year.

4.2 Model checking

Following the idea of posterior predictive model checking (Gelman *et al.*, 1996), we consider posterior predictions: for each $l = 1, 2, \dots$, we simulate a realization from the posterior distribution and use this when simulating “new data” $u_{t,i}^{(l)}$ and $v_{t,i}^{(l)}$ from the observation model; see Appendix A. For convenience, let $u_{t,i}^{(0)} = u_{t,i}$ and $v_{t,i}^{(0)} = v_{t,i}$ denote the data.

To check how well the model fits the data, we first consider the number of attacks in each year,

$$w_1^{(l)}(t) = \sum_i u_{t,i}^{(l)},$$

where $t = 1, \dots, 5$. The left plot in Figure 5 shows the 2.5, 50, 97.5 percentiles estimated from $w_1^{(l)}(t)$ for $l = 0, \dots, 199$. The values of $w_1^{(0)}(t)$ is also shown in the figure. In all years except year 3 where the number of *Ips* spp. attacks was particularly low, $w_1^{(0)}(t)$ is located in the central interval.

Second, let $d(i, j)$ denote Euclidean distance between site i and j . For $\delta \geq 1$,

$$w_2^{(l)}(\delta) = \sum_{i < j: d(i, j) \in (\delta-1, \delta]} v_{5,i}^{(l)} v_{5,j}^{(l)},$$

is the number of pairs of sites between $\delta - 1$ and δ apart and attacked at some time during the observation period. Thus $w_2^{(l)}(\delta)$ quantifies the degree of spatial clustering. The center plot in Figure 5 shows $w_2^{(l)}(\delta)$ for $\delta = 1, \dots, 5$ in the same way as in the first plot. Again there are no discrepancies between the model and the data.

Third, in order to combine temporal and spatial information, consider the number of neighboring sites that are attacked in years time t apart,

$$w_3^{(l)}(t) = \sum_{i, j, k: j \in N_i, k=t, \dots, 5} u_{k,i}^{(l)} u_{k-t,j}^{(l)}$$

for $t = 1, \dots, 4$, while

$$w_3^{(l)}(0) = \sum_{i, j, k: j \in N_i, k=0, \dots, 5} u_{k,i}^{(l)} u_{k,j}^{(l)} / 2$$

where we divide by 2 to avoid counting all the pairs twice. The right plot in Figure 5 shows $w_3^{(l)}(t)$ for $t = 0, \dots, 4$ in the same way as the plot for $w_1^{(l)}(t)$, except that we have taken the logarithm of $w_3^{(l)}(t)$ to be able to see what happens at times $t = 3, 4$. For $t = 0, \dots, 2$, $w_3^{(0)}(t)$ is located in the 95% central interval, but for $t = 3, 4$ the model underestimates the number of pairs.




Figure 5: Plots of $w_1(t)$, $w_2(\delta)$, and $\log(w_3(t))$. The crosses indicate the data and the bars indicate the 95% central interval and the median.

Finally, the upper row in Figure 6 shows plots of the data for each year $t = 1, \dots, 5$ where a site i at time t has been colored black if $v_{t,i}^{(0)} = 1$ and gray otherwise. Simulating 1000 posterior predictions, we let $\tilde{v}_{t,i} = 1$ if $v_{t,i}^{(l)} = 1$ in more than 50% of the simulations and $\tilde{v}_{t,i} = 0$ otherwise. The lower row in Figure 6 shows plots of $\tilde{v}_{t,i}$. Comparing the two rows in the figure, we see that both the data and the posterior predictive simulations show a clear formation of a large cluster of infested trees in the middle. Furthermore, the clusters seem to be roughly of the same size in the data and the simulations. On the other hand, there are some discrepancies between the shape of the cluster in the data and in the simulations: the shape of the cluster is more circular in the data than in the simulations. However, in spite of this deviation, the general behavior still seems to have been captured adequately well by the model.




Figure 6: Upper row: A site is colored black if $v_{t,i} = 1$ and gray if $v_{t,i} = 0$ for times $t = 1, \dots, 5$ (left to right). Lower row: same as upper row, but for $\tilde{v}_{t,i}$.

5 Comparison between continuous and discrete time models

We have illustrated how discrete time observations of a multivariate point process can be analyzed using a Bayesian missing data approach. We conclude with a discussion of the advantages and disadvantages of using continuous time processes compared to discrete time processes in the setting of the present paper as well as other spatial-temporal situations.

Computation: The most important advantage is the ease of computation. The likelihood function for a spatial-temporal process with discrete time often involves one or more unknown normalizing constants which need to be estimated using MCMC methods, see

e.g. the Markov random field models in Besag and Tantrum (2003) and Zhu *et al.* (2006). In contrast the likelihood function for a multivariate point process is completely specified by modeling the conditional intensity, cf. equation (2). In Zhu *et al.* (2006), we modeled the data set consisting of the three types of data given by turpentine beetles, *Ips* spp. and tree conditions. The likelihood function for this model factorized into three terms, one for each type of data, where we can compare the term corresponding to *Ips* spp. and the model used in the present paper, when we include the *Ips* spp. at all sites (rather than restricting the data to the subset of sites considered so far in the present paper). For this particular comparison, the computation time of the MCMC algorithm for posterior distributions for the discrete time model is roughly two hundred times longer than the corresponding computation time for the continuous time model. This improvement in speed means that we have been able to investigate several variations of the model (1) (see the end of Section 3), something we could not do within practical time limits in the case of the discrete time process for *Ips* spp. in Zhu *et al.* (2006). On the other hand, in Zhu *et al.* (2006) the likelihood terms corresponding turpentine beetle colonization and tree conditions were easy to specify, and calculations for these terms were very fast. Extending the continuous time model to include turpentine beetles and tree conditions may be much more involved, and a comparison between such a model and the full model in Zhu *et al.* (2006) may well turn out differently.

External information: For a discrete time process as compared to a continuous time process, it may be difficult to incorporate external (or, in a broad sense, covariate) information in the form of another stochastic process observed at a different time scale than the discrete time process. For example, the ρ term in (1) incorporates external information about *Ips* spp. activity due to seasonal variation; such information can only be incorporated into the model in Zhu *et al.* (2006) after some form of aggregation. On the other hand, we require the external information that was collected in the field across a full season of *Ips* spp. activity modeled into ρ to obtain a realistic continuous time model for *Ips* spp. attacks. If we had no such information or only unreliable information available, the continuous time model approach would be problematic. It would also be a problem if the modeling form of ρ is misspecified. Specifying ρ is time consuming, and this partially offsets the advantage of the shorter computation time for the continuous time model.

Time scale comparability: While the parameters of two continuous time processes may be compared even if their time scales for observations (i.e. time lengths between observation times) are different, it may not be meaningful to compare the parameters of two discrete time processes with different time scales. The choice of observation times in Zhu *et al.* (2006) is biologically meaningful, since we have annual observation after cessation of insect activity; however, suppose that we had another data set with observations every second year. Further, imagine that we wish to use a model of the same form as in Zhu *et al.* (2006) for this other data set. Then, because of the different time scales, we cannot directly compare the parameters governing e.g. *Ips* spp. activity.

Consistency: Although observed only at discrete times, the types of data considered in the present paper come from an underlying continuous time process, and the existence of an underlying continuous time process is not ensured by specifying a discrete time process. Obviously, this is not an issue when we have specified a continuous time model, but our continuous time model only approximates the complexity of the system under study. Indeed, *Ips* spp. colonization does not occur exactly at one time point; it is a complicated process involving hundreds or thousands of beetles attacking over a short period of time. On the other hand, the discrete time model reflects the cumulative *Ips* spp. attacks throughout the past season, and in this sense it is a perfectly sensible model for the system under study.

Estimation of missing data: A continuous time process allows us to model what has happened between observation times, where in our case the event times of within-season *Ips* spp. attacks can be readily estimated by the MCMC algorithm in Section 4.1. In contrast it is not possible to do this kind of estimation for missing data based on the model in Zhu *et al.* (2006), although it was not a part of the scientific objectives there. In the present paper, the estimates provide some trustful qualitative results regarding how the times of events depend on the distance to the gap of dead trees and the similar behavior over the years. However, quantitative results depend much on a careful modeling of external information, particularly the function ρ .

Predictions: It is straightforward to predict what may happen at any time after the final observation time, applying Ogata's modified thinning algorithm for the continuous time process in the present paper (see Appendix A). The discrete time model in Zhu *et al.* (2006) allows us only to predict what may happen annually.

In conclusion, spatial-temporal processes often suffer from being computationally intensive, and we have demonstrated that in the present case, using a continuous time process is indeed a viable alternative to a discrete time process.

Appendix A

For model checking in Section 4.2, we need to simulate new data from the observation model. For this we use Ogata's modified thinning algorithm extended to marked processes (Daley and Vere-Jones, 2003; Ogata, 1981).

Factorize the intensity, $\lambda_{t,i} = \lambda_t \times p(i|t)$, where $\lambda_t = \sum_i \lambda_{t,i}$ is the intensity for the temporal process of events disregarding the location, and $p(i|t) \propto \lambda_{t,i}$ is the probability function for the location given the time of the event. Note that the dependence on the past events (locations and times) of both functions are suppressed in the notation. Furthermore, define two functions $l(t)$ and $m(t)$ by $l(t) = t + 0.1$ for $k - 1 < t \leq \mu_k$ and $l(t) = k$ for $\mu_k < t \leq k$, and $m(t) = \max_{s \in [t, l(t)]} \lambda_s$. Ogata's modified thinning algorithm is then started at time $t = 0$ and the following steps are repeated until $t > 5$:

1. Compute $l(t)$ and $m(t)$.

2. Generate an exponentially distributed variable T with inverse mean $m(t)$ and a uniform variable U on $[0, 1]$ independently of each other.
3. If $T > l(t)$, set $t = t + l(t)$.
4. Else if $t + T > 5$ or $U > \lambda_{t+T}/m(t)$, set $t = t + T$.
5. Otherwise, let the next event be $t_i = t + T$, where i is generated using the probability function $p(i|t)$, and set $t = t + T$.

The output is the set of all t_i obtained in step 5.

Combining this algorithm with a sample of the posterior distributions in Section 4.1, we get posterior predictions. In practice we get a sample of the posterior distributions by taking values from the Markov chains at regular intervals chosen such that the sample is effectively independent.

Acknowledgments

This work was supported by the Danish Natural Science Research Council, USDA NRI (2003-3502-13528), NSF (DEB0314215), and the UW-Madison College of Agricultural and Life Sciences. We are grateful to the Wisconsin Department of Natural Resources for providing the study site.

References

- Aukema, B. H., Clayton, M. K., and Raffa, K. F. (2005). Modeling flight activity and population dynamics of the pine engraver, *Ips pini*, in the Great Lakes region: Effects of weather and predators at short time scales. *Population Ecology*, **47**, 61–69.
- Besag, J. and Tantrum, J. (2003). Likelihood analysis of binary data in space and time. In P. J. Green, N. L. Hjort, and S. Richardson, editors, *Highly Structured Stochastic Systems*, pages 289–295. Oxford University Press, Oxford.
- Daley, D. J. and Vere-Jones, D. (2003). *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. Springer, New York, 2nd edition.
- Erbilgin, N. and Raffa, K. F. (2003). Spatial analysis of forest gaps resulting from bark beetle colonization of red pines experiencing belowground herbivory and infection. *For. Ecol. & Manag.*, **177**, 145–153.
- Furniss, R. L. and Carolin, V. M. (1980). *Western Forest Insects*. USDA FS Misc. Publ. 1339. Washington DC. 654 pp.

- Gelman, A., Meng, X. L., and Stern, H. S. (1996). Posterior predictive assessment of model fitness via realized discrepancies (with discussion). *Statistica Sinica*, **6**, 733–807.
- Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, **58**, 83–90.
- Ogata, Y. (1981). On Lewis’ simulation method for point processes. *IEEE Transactions on Information Theory*, **IT-27**, 23–31.
- Raffa, K. F. and Berryman, A. A. (1983). The role of host plant resistance in the colonization behavior and ecology of bark beetles. *Ecol. Monogr.*, **53**, 27–49.
- Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer, New York, 2nd edition.
- Roberts, G. O., Gelman, A., and Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Annals of Applied Probability*, **7**, 110–120.
- Wood, D. L. (1982). The role of pheromones, kairomones, and allomones in the host selection and colonization behavior of bark beetles. *Ann. Rev. of Entomol.*, **27**, 411–446.
- Zhu, J., Rasmussen, J. G., Møller, J., Aukema, B. H., and Raffa, K. (2006). Spatial-temporal modeling of forest gaps generated by colonization from below- and above-ground beetle species. Research report R-2006-04, Department of Mathematical Sciences, Aalborg University. Available at <http://www.math.aau.dk/~jm>. (Submitted for publication).