

Independent component models for replicated
point processes

Daniel Gervini
Department of Mathematical Sciences
University of Wisconsin–Milwaukee

August 9, 2016

Abstract

We propose a semiparametric independent-component model for the intensity functions of replicated point processes. We show that the maximum likelihood estimators of the model parameters are consistent and asymptotically normal when the number of replications goes to infinity. The finite-sample behavior of the estimators is studied by simulation. As an example of application, we analyze the temporal variation in the spatial distribution of street robberies in Chicago.

Key words: Doubly-stochastic process; functional data analysis; latent-variable model; Poisson process; spline smoothing.

1 Introduction

Point processes in time and space have a broad range of applications, in diverse areas such as neuroscience, ecology, finance, astronomy, seismology, and many others. Examples are given in classic textbooks like Cox and Isham (1980), Diggle (2013), Møller and Waagepetersen (2004), Streit (2010), and Snyder and Miller (1991), and in the papers cited below. However, the point-process literature has mostly focused on single-realization cases, such as the distribution of trees in a single forest (Jalilian et al., 2013) or the distribution of cells in a single tissue sample (Diggle et al., 2006). Situations where several replications of a process are available are increasingly common. Among the few papers proposing statistical methods for replicated point processes we can cite Diggle et al. (1991), Baddeley et al. (1993), Diggle et al. (2000), Bell and Grunwald (2004), Landau et al. (2004), Wager et al. (2004), and Pawlas (2011). However, these papers propose estimators for summary statistics of the processes (the so-called F , G and K statistics) rather than the intensity functions, which would be more informative.

When several replications of a process are available, it is possible to “borrow strength” across replications to estimate the intensity functions. This is the basic idea behind many functional data methods that are applied to stochastic processes which, individually, are only sparsely sampled (James et al., 2000; Yao et al., 2005; Gervini, 2009). Functional Data Analysis has mostly focused on continuous temporal processes; little work has been done on discrete or spatial point processes. We can mention Bouzas et al. (2006, 2007) and Fernández-Alcalá et al. (2012), which have rather limited scopes since they only estimate the mean of temporal Cox processes, and Wu et al. (2013), who estimate the mean and the principal components of independent and identically distributed realizations of a temporal Cox process, but their method is based on kernel estimators of covariance functions that cannot be easily extended beyond i.i.d. replications of temporal processes. The semiparametric model we propose in this paper, on the other hand, can be applied just as easily to temporal or spatial processes, and it can accommodate covariates and be extended to non-i.i.d. situations like ANOVA models or more complex data structures like marked point processes and multivariate processes, although we will not consider those extensions in this paper.

As an example of application we will analyze the spatial distribution of street

robberies in Chicago in the year 2014, where each day is seen as a replication of the spatial process. Our model is presented in Section 2 and the estimation process in Section 3. We establish the consistency and asymptotic normality of the component estimators in Section 4, study their finite-sample behavior by simulation in Section 5. The Chicago crime data is analyzed in Section 6.

2 The model

A point process X is a random countable set in a space \mathcal{S} , where \mathcal{S} is usually \mathbb{R} for temporal processes and \mathbb{R}^2 or \mathbb{R}^3 for spatial processes (Møller and Waagepetersen, 2004, ch. 2; Streit, 2010, ch. 2). Locally finite processes are those for which $\#(X \cap B) < \infty$ with probability one for any bounded $B \subseteq \mathcal{S}$. For such processes we can define the count function $N(B) = \#(X \cap B)$. Given a locally integrable function $\lambda : \mathcal{S} \rightarrow [0, \infty)$, i.e. a function such that $\int_B \lambda(t) dt < \infty$ for any bounded $B \subseteq \mathcal{S}$, the process X is a Poisson process with intensity function $\lambda(t)$ if (i) $N(B)$ follows a Poisson distribution with rate $\int_B \lambda(t) dt$ for any bounded $B \subseteq \mathcal{S}$, and (ii) conditionally on $N(B) = m$, the m points in $X \cap B$ are independent and identically distributed with density $\tilde{\lambda}(t) = \lambda(t) / \int_B \lambda$ for any bounded $B \subseteq \mathcal{S}$.

Let X be a Poisson process with intensity function $\lambda(t)$ and $X_B = X \cap B$ for a given B . Then the density function of X_B at $x_B = \{t_1, \dots, t_m\}$ is

$$\begin{aligned} f(x_B) &= f(m)f(t_1, \dots, t_m|m) \\ &= \exp\left\{-\int_B \lambda(t) dt\right\} \frac{\left\{\int_B \lambda(t) dt\right\}^m}{m!} \times \prod_{j=1}^m \tilde{\lambda}(t_j). \end{aligned} \tag{1}$$

Since the realizations of X_B are sets, not vectors, what we mean by ‘density’ is the following: if \mathcal{N} is the family of locally finite subsets of \mathcal{S} , i.e. $\mathcal{N} = \{A \subseteq \mathcal{S} : \#(A \cap B) < \infty \text{ for all bounded } B \subseteq \mathcal{S}\}$, then for any $F \subseteq \mathcal{N}$,

$$\begin{aligned} P(X_B \in F) &= \sum_{m=0}^{\infty} \int_B \cdots \int_B \mathbb{I}(\{t_1, \dots, t_m\} \in F) f(\{t_1, \dots, t_m\}) dt_1 \cdots dt_m \\ &= \sum_{m=0}^{\infty} \frac{\exp\left\{-\int_B \lambda(t) dt\right\}}{m!} \int_B \cdots \int_B \mathbb{I}(\{t_1, \dots, t_m\} \in F) \left\{\prod_{j=1}^m \lambda(t_j)\right\} dt_1 \cdots dt_m, \end{aligned}$$

and, more generally, for any function $h : \mathcal{N} \rightarrow [0, \infty)$

$$E\{h(X_B)\} = \sum_{m=0}^{\infty} \int_B \cdots \int_B h(\{t_1, \dots, t_m\}) f(\{t_1, \dots, t_m\}) dt_1 \cdots dt_m. \quad (2)$$

A function h on \mathcal{N} is, essentially, a function well defined on \mathcal{S}^m for any integer m which is invariant under permutation of the coordinates; for example, $h(\{t_1, \dots, t_m\}) = \sum_{j=1}^m t_j/m$. We note that an alternative definition of density function, given e.g. in Møller and Waagepetersen (2004, ch. 6), is the density function of X_B with respect to the standard (unit rate) Poisson process on B at the point pattern $x_B = \{t_1, \dots, t_m\}$, which is

$$f(x_B) = \exp\left\{|B| - \int_B \lambda(t) dt\right\} \prod_{j=1}^m \lambda(t_j),$$

but the difference with (1) is only a constant factor $\exp(|B|)/m!$ that does not affect the likelihood-based inference that follows, so both definitions of $f(x_B)$ are equivalent for our purposes; we will use (1).

Single realizations of point processes are often modeled as Poisson processes with fixed λ s. But for replicated point processes a single intensity function λ rarely provides an adequate fit for all replications. It is more reasonable to assume that the λ s are subject-specific and treat them as latent random effects. These processes are called doubly stochastic or Cox processes (Møller and Waagepetersen, 2004, ch. 5; Streit, 2010, ch. 8). We can think of a doubly stochastic process as a pair (X, Λ) where $X|\Lambda = \lambda$ is a Poisson process with intensity λ , and Λ is a random function that takes values on the space \mathcal{F} of non-negative locally integrable functions on \mathcal{S} . This provides an appropriate framework for the type of applications we have in mind: the n replications can be seen as i.i.d. realizations $(X_1, \Lambda_1), \dots, (X_n, \Lambda_n)$ of a doubly stochastic process (X, Λ) , where X is observable but Λ is not. In this paper we will assume that all X_i s are observed on a common region B of \mathcal{S} ; the method could be extended to X_i s observed on non-conformal regions B_i at the expense of a higher computational complexity.

For many applications we can think of the intensity process Λ as a finite combination of independent components,

$$\Lambda(t) = \sum_{k=1}^p U_k \phi_k(t), \quad (3)$$

where ϕ_1, \dots, ϕ_p are functions in \mathcal{F} that will be estimated from the data, and U_1, \dots, U_p are independent nonnegative random variables. This is the functional equivalent of a multivariate Independent Component decomposition (Hyvärinen et al., 2001). For identifiability we assume that $\int_B \phi_k(t) dt = 1$ for all k , i.e. that the ϕ_k s are density functions. Then $\int_B \Lambda(t) dt = \sum_{k=1}^p U_k$ and, if we define $S = \sum_{k=1}^p U_k$ and $W_k = U_k/S$, we have

$$\Lambda(t) = S\tilde{\Lambda}(t), \text{ with } \tilde{\Lambda}(t) = \sum_{k=1}^p W_k \phi_k(t). \quad (4)$$

The intensity process $\Lambda(t)$, then, can be seen as the product of an intensity factor S and a scatter factor $\tilde{\Lambda}(t)$ which is a mixture of the ϕ_k s. Conditionally on $\Lambda(t) = \lambda(t)$ the count function $N(B)$ has a Poisson distribution with rate $S = s$, and conditionally on $N(B) = m$ the m points in X_B are independent identically distributed realizations with density $\tilde{\lambda}(t)$ which is determined by the W_k s. In general S and \mathbf{W} are not independent, but if the U_k s are assumed to follow independent Gamma distributions, say $U_k \sim \text{Gamma}(\alpha_k, \beta)$ with a common β , then S and \mathbf{W} are independent with $S \sim \text{Gamma}(\sum \alpha_k, \beta)$ and $\mathbf{W} \sim \text{Dirichlet}(\alpha_1, \dots, \alpha_p)$ (Bilodeau and Brenner, 1999, ch. 3).

As pointed out by a referee, the motivation for model (3) is heuristic, not mathematical. Ordinary principal components can be justified by Mercer's theorem (Ash and Gardner, 1975, ch. 1.4) for any stochastic process with finite variance, but not every stochastic process $\Lambda(t)$ in \mathcal{F} will follow model (3), not even with $p = \infty$. The justification for model (3) is the same as for multivariate independent components (Hyvärinen et al., 2001, ch. 7): for many, but not all, cases of interest, it provides a practical, parsimonious and interpretable approximation to the true process $\Lambda(t)$.

3 Estimation

The marginal density of X_B , with a slight abuse of notation in writing $x_B = (m, \mathbf{t})$, can be expressed as

$$\begin{aligned} f(x_B) &= \int \int f(m, \mathbf{t}, s, \mathbf{w}) ds d\mathbf{w} \\ &= \int \int f(\mathbf{t} | m, \mathbf{w}) f(m | s) f(s, \mathbf{w}) ds d\mathbf{w} \end{aligned} \quad (5)$$




Figure 1: Hierarchical diagram of model variables. Observed quantities are in boxes.

with

$$f(\mathbf{t} \mid m, \mathbf{w}) = \prod_{j=1}^m \sum_{k=1}^p w_k \phi_k(t_j).$$

Since $\tilde{\lambda}(t)$ is a mixture of the ϕ_k s, as is customary in the implementation of the EM algorithm for mixtures (Dempster et al., 1977) we introduce additional p -dimensional indicator vectors $\mathbf{y}_1, \dots, \mathbf{y}_m$ that, in words, indicate from which ϕ_k each point t_j comes from. More rigorously, given m and \mathbf{w} , for $j = 1, \dots, m$ we define $\mathbf{y}_j = (y_{j1}, \dots, y_{jp})$ with $y_{jk} \in \{0, 1\}$ and $\sum_{k=1}^p y_{jk} = 1$, with $P(y_{jk} = 1 \mid m, \mathbf{w}) = w_k$. So, conditionally on m and \mathbf{w} , the \mathbf{y}_j s are independent Multinomial(1, \mathbf{w}). Collecting all \mathbf{y}_j s into a single \mathbf{y} for notational simplicity, and denoting by \mathcal{Y}_m the set of all possible \mathbf{y} s, by marginalizing over the \mathbf{y} s we obtain

$$\begin{aligned}
f(x_B) &= \int \int \sum_{\mathbf{y} \in \mathcal{Y}_m} f(m, \mathbf{t}, \mathbf{y}, s, \mathbf{w}) ds d\mathbf{w} \\
&= \int \int \sum_{\mathbf{y} \in \mathcal{Y}_m} f(\mathbf{t} \mid m, \mathbf{y}) f(\mathbf{y} \mid m, \mathbf{w}) f(m \mid s) f(s, \mathbf{w}) ds d\mathbf{w}
\end{aligned}$$

with

$$f(\mathbf{t} \mid m, \mathbf{y}) = \prod_{j=1}^m \prod_{k=1}^p \phi_k(t_j)^{y_{jk}}$$

and

$$f(\mathbf{y} \mid m, \mathbf{w}) = \prod_{j=1}^m \prod_{k=1}^p w_k^{y_{jk}},$$

which are more amenable for implementation of the EM algorithm.

Figure 1 helps clarify the dependence between the variables, as well as which variables are observed and which ones are latent.

To estimate the ϕ_k s we use a semiparametric approach: we assume they belong to a family of functions \mathcal{B} with non-negative basis functions $\beta_1(t), \dots, \beta_q(t)$ defined on the region of interest B ; for example, we can use B-splines for temporal processes and tensor-product splines or radial basis functions for spatial processes.

Then $\phi_k(t) = \mathbf{c}_k^T \boldsymbol{\beta}(t)$ with $\mathbf{c}_k \geq 0$ to guarantee that $\phi_k(x) \geq 0$. The strategy, as usual in semiparametric modelling, is to use a rich family \mathcal{B} , with a relatively large q , and control the smoothness of the ϕ_k s via a smoothing parameter, as explained below. The specific nature of \mathcal{S} , i.e. whether the process is temporal or spatial, will only be reflected in the choice of \mathcal{B} ; for all other purposes, including the theory developed below, we need not make any distinctions between temporal and spatial processes.

The model parameters are estimated by penalized maximum likelihood. In addition to the component coefficients $\mathbf{c}_1, \dots, \mathbf{c}_p$, the distribution of S and \mathbf{W} will depend on certain parameters that we collect in a vector $\boldsymbol{\eta}$. The whole model, then, is parameterized by $\boldsymbol{\theta} = (\mathbf{c}_1, \dots, \mathbf{c}_p, \boldsymbol{\eta})$ and the parameter space is $\Theta = \mathcal{C} \times \mathcal{H}$, where

$$\mathcal{C} = \{(\mathbf{c}_1, \dots, \mathbf{c}_p) : \mathbf{a}^T \mathbf{c}_k = 1, \mathbf{c}_k \geq 0 \text{ for all } k\},$$

with $\mathbf{a} = \int_B \boldsymbol{\beta}(t) dt$, and \mathcal{H} is the space of the $\boldsymbol{\eta}$ s. Then, given n independent realizations x_{B1}, \dots, x_{Bn} with $x_{Bi} = \{t_{i1}, \dots, t_{im_i}\}$, the likelihood function is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^n f(x_{Bi}; \boldsymbol{\theta}) \quad (6)$$

with $f(x_B; \boldsymbol{\theta})$ given by (5). Since the dimension of \mathcal{B} is typically large, a roughness penalty has to be added to (6) to obtain smooth $\hat{\phi}_k$ s. We use a penalty of the form $-\zeta P(\boldsymbol{\theta})$ with $\zeta \geq 0$ and $P(\boldsymbol{\theta}) = \sum_{k=1}^p g(\phi_k)$, where

$$g(\phi) = \int_B |\mathbf{H}\phi(t)|_F^2 dt,$$

\mathbf{H} denotes the Hessian and $|\cdot|_F$ the Frobenius matrix norm. So for a temporal process $g(\phi) = \int (\phi'')^2$ and for a spatial process $g(\phi) = \int \{(\frac{\partial^2 \phi}{\partial t_1^2})^2 + 2(\frac{\partial^2 \phi}{\partial t_1 \partial t_2})^2 + (\frac{\partial^2 \phi}{\partial t_2^2})^2\}$, both of which are quadratic functions in the coefficients \mathbf{c}_k . Then $\hat{\boldsymbol{\theta}}$ maximizes

$$\rho_n(\boldsymbol{\theta}) = n^{-1} \log L(\boldsymbol{\theta}) - \zeta P(\boldsymbol{\theta}) \quad (7)$$

among $\boldsymbol{\theta} \in \Theta$ for a given smoothing parameter ζ . The smoothing parameter ζ as

well as the number of components p can be chosen by cross-validation, maximizing

$$\text{CV}(\zeta, p) = \sum_{i=1}^n \log f(x_{Bi}; \hat{\boldsymbol{\theta}}_{(-i)}), \quad (8)$$

where $\hat{\boldsymbol{\theta}}_{(-i)}$ denotes the penalized maximum likelihood estimator for the reduced sample with x_{Bi} deleted. An EM algorithm for the computation of $\hat{\boldsymbol{\theta}}$ is described in detail in the Supplementary Material. As by-products of the EM algorithm, at no extra computational cost, we obtain individual predictors of the random effects in (4), $\hat{s}_i = \hat{E}(S | X_i)$ and $\hat{w}_{ik} = \hat{E}(W_k | X_i)$, and then as individual predictors of the random effects in (3) we can use $\hat{u}_{ik} = \hat{s}_i \hat{w}_{ik}$, obtaining estimates of the individual intensity functions $\hat{\lambda}_i(t) = \sum_{k=1}^p \hat{u}_{ik} \hat{\phi}_k(t)$.

4 Asymptotics

The asymptotic behavior of $\hat{\boldsymbol{\theta}}$ as $n \rightarrow \infty$ can be studied via standard empirical-process techniques (Pollard, 1984; Van der Vaart, 2000), since (7) is the average of independent identically distributed functions plus a non-random roughness penalty, as in e.g. Knight and Fu (2000).

In principle two types of asymptotics may be of interest: the “nonparametric” asymptotics, where the dimension of the basis space \mathcal{B} grows with n , and the simpler “parametric” asymptotics, where the dimension of \mathcal{B} is large but fixed, and the ϕ_k s are assumed to belong to \mathcal{B} . In effect, the second approach ignores smoothing bias, but this is not a serious problem in practice as long as the dimension q of \mathcal{B} is reasonably large. We will then follow the second approach, as others have in similar semiparametric contexts (e.g. Yu and Ruppert, 2002, and Xun et al., 2013).

The constraints of the space \mathcal{C} introduce some complications in an otherwise standard asymptotic analysis. We will follow the approach of Geyer (1994). Proofs of the results presented here can be found in the Supplementary Material. The first result of this section, consistency of the estimator, is essentially a consequence of model identifiability. It is not obvious that model (3) is identifiable, so this is proved first in the Supplementary Material.

Theorem 1 *If the smoothing parameter $\zeta = \zeta_n$ goes to zero as $n \rightarrow \infty$, either deterministically or in probability, then $\hat{\boldsymbol{\theta}}_n \xrightarrow{P} \boldsymbol{\theta}_0$.*

As explained in Section 3, the parameter space Θ is of the form $\mathcal{C} \times \mathcal{H}$ with an \mathcal{H} that we will assume convex and open. For example, for the model with independent Gamma scores, $\boldsymbol{\eta} = (\alpha_1, \dots, \alpha_p, \beta)$ and $\mathcal{H} = (0, +\infty)^{p+1}$. So \mathcal{H} is not a problem for the asymptotics because it does not contribute active constraints. The problem is the space \mathcal{C} , which is a convex but closed set. In particular, the non-negativity constraints create some unusual asymptotics; for example, if a given $c_{0,kj}$ is zero, it is clear that $\sqrt{n}(\hat{c}_{n,kj} - c_{0,kj})$ cannot be asymptotically normal since it is a nonnegative quantity for all n .

To handle these constraints we need to introduce the notion of tangent cone. Let r and d be such that $\mathcal{H} \subseteq \mathbb{R}^r$ and $\Theta \subseteq \mathbb{R}^d$. The tangent cone of Θ at $\boldsymbol{\theta}_0$ is the set of all directions and limits of directions from which $\boldsymbol{\theta}_0$ can be approached with sequences that stay inside Θ . If $\boldsymbol{\theta}_0$ is in the interior of Θ it can be approached from any direction, so the tangent cone is the whole \mathbb{R}^d in that case. But if $\boldsymbol{\theta}_0$ is on the border of Θ it can only be approached from certain directions; for example, if $c_{0,11} = 0$ then the approaching sequences must necessarily satisfy $c_{11} \geq 0$ to stay inside Θ . An in-depth treatment of tangent cones can be found in Hiriart-Urruty and Lemaréchal (2001, ch. A.5). The tangent cone of the product of convex sets is the product of the respective tangent cones, so the tangent cone of Θ at $\boldsymbol{\theta}_0$ is $\mathcal{D}_0 = \mathcal{T} \times \mathbb{R}^r$, with

$$\mathcal{T} = \{(\mathbf{v}_1, \dots, \mathbf{v}_p) : \mathbf{a}^T \mathbf{v}_k = 0, v_{kj} \geq 0 \text{ for } j \in I_k, k = 1, \dots, p\}, \quad (9)$$

and $I_k = \{j : c_{0,kj} = 0\}$.

The following theorem gives the asymptotic distribution of $\hat{\boldsymbol{\theta}}_n$ in a form as explicit as can be given for a general $\boldsymbol{\theta}_0$. Let \mathbf{F}_0 be Fisher's Information Matrix,

$$\begin{aligned} \mathbf{F}_0 &= E_{\boldsymbol{\theta}_0} \{ \nabla \log f(X_B; \boldsymbol{\theta}_0) \nabla \log f(X_B; \boldsymbol{\theta}_0)^T \} \\ &= -E_{\boldsymbol{\theta}_0} \{ \nabla^2 \log f(X_B; \boldsymbol{\theta}_0) \}, \end{aligned}$$

where the derivatives of $\log f(x_B; \boldsymbol{\theta})$ are taken with respect to the parameter $\boldsymbol{\theta}$. Explicit expressions for these derivatives can be found in the Supplementary Material.

Theorem 2 *If $\sqrt{n}\zeta_n \rightarrow \kappa$ as $n \rightarrow \infty$, either deterministically or in probability, then*

$\sqrt{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \xrightarrow{D} \boldsymbol{\delta}(\mathbf{Z})$, where $\boldsymbol{\delta}(\mathbf{Z})$ is the maximizer of

$$W(\boldsymbol{\delta}) = \{\mathbf{Z} - \kappa \nabla P(\boldsymbol{\theta}_0)\}^T \boldsymbol{\delta} - \frac{1}{2} \boldsymbol{\delta}^T \mathbf{F}_0 \boldsymbol{\delta}$$

over $\boldsymbol{\delta} \in \mathcal{D}_0$, and $\mathbf{Z} \sim \mathbf{N}(\mathbf{0}, \mathbf{F}_0)$.

Although a closed expression for $\boldsymbol{\delta}(\mathbf{Z})$ in Theorem 2 is not available in general, $\boldsymbol{\delta}(\mathbf{Z})$ is easy to simulate: one generates M Monte Carlo samples $\mathbf{Z}_1, \dots, \mathbf{Z}_M$ from a $\mathbf{N}(\mathbf{0}, \mathbf{F}_0)$ distribution and then computes $\boldsymbol{\delta}(\mathbf{Z}_i)$ for each \mathbf{Z}_i , which involves a simple quadratic optimization problem with linear constraints. But in the particular case where all component coefficients $c_{0,kj}$ are strictly positive, we can be more explicit: write a $\boldsymbol{\delta} \in \mathcal{D}_0$ as $\boldsymbol{\delta} = (\boldsymbol{\delta}_1, \boldsymbol{\delta}_2)$ with $\boldsymbol{\delta}_1 \in \mathcal{S}$ and $\boldsymbol{\delta}_2 \in \mathbb{R}^r$; since there are no inequality constraints in (9), the only restriction for $\boldsymbol{\delta}_1$ is that $(I_p \otimes \mathbf{a}^T) \boldsymbol{\delta}_1 = \mathbf{0}_p$. The matrix $I_p \otimes \mathbf{a}^T$ is $p \times pq$ of rank p , so there exists an orthogonal $pq \times (pq - p)$ matrix $\boldsymbol{\Gamma}$ such that $(I_p \otimes \mathbf{a}^T) \boldsymbol{\Gamma} = \mathbf{0}$ and $\boldsymbol{\delta}_1 = \boldsymbol{\Gamma} \boldsymbol{\xi}$ for some unconstrained $\boldsymbol{\xi} \in \mathbb{R}^{pq-p}$. Let

$$\mathbf{A} = \begin{pmatrix} \boldsymbol{\Gamma}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix}.$$

Then

$$W(\boldsymbol{\delta}) = \tilde{\mathbf{Z}}^T \begin{pmatrix} \boldsymbol{\xi} \\ \boldsymbol{\delta}_2 \end{pmatrix} - \frac{1}{2} (\boldsymbol{\xi}^T, \boldsymbol{\delta}_2^T) \tilde{\mathbf{F}}_0 \begin{pmatrix} \boldsymbol{\xi} \\ \boldsymbol{\delta}_2 \end{pmatrix}$$

with $\tilde{\mathbf{Z}} = \mathbf{A} \{\mathbf{Z} - \kappa \nabla P(\boldsymbol{\theta}_0)\}$ and $\tilde{\mathbf{F}}_0 = \mathbf{A} \mathbf{F}_0 \mathbf{A}^T$, so

$$\boldsymbol{\delta}(\mathbf{Z}) = \mathbf{A} \begin{pmatrix} \boldsymbol{\xi}(\tilde{\mathbf{Z}}) \\ \boldsymbol{\delta}_2(\tilde{\mathbf{Z}}) \end{pmatrix} \quad \text{with} \quad \begin{pmatrix} \boldsymbol{\xi}(\tilde{\mathbf{Z}}) \\ \boldsymbol{\delta}_2(\tilde{\mathbf{Z}}) \end{pmatrix} = \tilde{\mathbf{F}}_0^{-1} \tilde{\mathbf{Z}}.$$

Note that $\tilde{\mathbf{Z}} \sim \mathbf{N}(\boldsymbol{\mu}, \tilde{\mathbf{F}}_0)$ with $\boldsymbol{\mu} = -\kappa \mathbf{A} \nabla P(\boldsymbol{\theta}_0)$, so $(\boldsymbol{\xi}(\tilde{\mathbf{Z}}), \boldsymbol{\delta}_2(\tilde{\mathbf{Z}}))$ has a non-singular $\mathbf{N}(\tilde{\mathbf{F}}_0^{-1} \boldsymbol{\mu}, \tilde{\mathbf{F}}_0^{-1})$ distribution but $\boldsymbol{\delta}(\mathbf{Z})$ itself has a singular $\mathbf{N}(\mathbf{A} \tilde{\mathbf{F}}_0^{-1} \boldsymbol{\mu}, \mathbf{A} \tilde{\mathbf{F}}_0^{-1} \mathbf{A}^T)$ distribution. Nevertheless, the diagonal elements of $\mathbf{A} \tilde{\mathbf{F}}_0^{-1} \mathbf{A}^T / n$ can be used as variance estimators to construct, for example, confidence intervals for the elements of $\boldsymbol{\theta}$.

5 Simulations

We ran some simulations to study the finite-sample performance of the estimators under various temporal and spatial models, and to compare the performance of

the new method with the method of Wu et al. (2013) for estimation of individual intensity functions.

The first pair of models we considered were temporal. We considered two models like (3), both with two components: Model 1 had $\phi_1(t) = \exp\{-100(t - .3)^2\}/.177$ and $\phi_2(t) = \exp\{-100(t - .7)^2\}/.177$, $t \in [0, 1]$, which are two peaks with practically no overlap; Model 2 had $\phi_1(t) = \exp\{-20(t - .3)^2\}/.385$ and $\phi_2(t) = \exp\{-20(t - .7)^2\}/.385$, $t \in [0, 1]$, which are also unimodal but flatter functions with more overlap. For both models the component scores U_1 and U_2 had lognormal distributions with $E(U_1) = 30$, $E(U_2) = 20$, $V(U_1) = 10$ and $V(U_2) = 1$. Two sample sizes were considered, $n = 50$ and $n = 150$.

For estimation of the ϕ_k s we used cubic B -splines with $K = 5$ and $K = 10$ equispaced knots, so $q = 9$ and $q = 14$ respectively, and Gamma-distributed component scores as working model. The estimating models are then somewhat different from the data-generating models, allowing us to assess the robustness of the estimators to model misspecification. The smoothing parameter ζ was chosen by five-fold cross-validation; for comparison we also computed estimators with the optimal ζ_{opt} that minimizes the mean squared error, which gives us a lower bound that cannot be attained in practice but serves as a benchmark to judge the adequacy of five-fold cross-validation as a method for choosing ζ .

Table 1 summarizes the results. For each estimator we report $\text{bias} = (\int [E\{\hat{\phi}(t)} - \phi(t)]^2 dt)^{1/2}$, $\text{std} = \{\int E([\hat{\phi}(t) - E\{\hat{\phi}(t)}])^2 dt\}^{1/2}$ and $\text{rmse} = (\int E[\{\hat{\phi}(t) - \phi(t)\}^2] dt)^{1/2}$, where the expectations have been approximated by 500 Monte Carlo replications. We do not report Monte Carlo standard errors on the table, but, for example, the first mean squared error reported, $.732^2 = .536$, has a Monte Carlo standard deviation of .001, which gives a 95% confidence interval (.534, .538) for the mean squared error and (.730, .733) for the root mean squared error, so the quantities in Table 1 are accurate up to the two reported decimals. We see in Table 1 that the comparative behavior of the estimators with respect to sample size and number of knots is similar for both models and both components, although ϕ_1 is easier to estimate than ϕ_2 because it is the dominant component, having both larger expectation and variance in the true model. We see that ten knots produce much better estimators than five knots, even for a sample size as small as 50. For a fixed sample size we see that increasing the number of knots reduces the bias and increases the variance, as expected, but the gains in bias amply compensate for the increases in variance, so the overall

Model 1												
$n = 50$												
$n = 150$												
Estimator	$q = 9$			$q = 14$			$q = 9$			$q = 14$		
	bias	std	rmse	bias	std	rmse	bias	std	rmse	bias	std	rmse
ϕ_1 cv	.73	.07	.73	.30	.14	.33	.73	.09	.73	.17	.12	.20
ϕ_2 cv	.81	.08	.82	.35	.17	.39	.82	.09	.82	.19	.14	.24
ϕ_1 opt	.73	.04	.73	.18	.07	.19	.73	.03	.73	.14	.07	.16
ϕ_2 opt	.81	.04	.81	.20	.09	.22	.82	.03	.82	.15	.08	.17
Model 2												
ϕ_1 cv	.31	.07	.32	.19	.08	.20	.31	.05	.32	.19	.04	.19
ϕ_2 cv	.33	.08	.34	.22	.10	.24	.33	.07	.34	.22	.05	.23
ϕ_1 opt	.30	.08	.31	.09	.07	.11	.30	.05	.30	.08	.06	.10
ϕ_2 opt	.32	.07	.33	.07	.08	.11	.32	.04	.32	.07	.07	.10

Table 1: Simulation Results. Bias, standard deviation and root mean squared error of independent component estimators for temporal models.

mean squared errors are much smaller. The estimators, then, successfully “borrow strength” across replications, and it is best to err on the side of using too many basis functions rather than too few. Overall, the performance of the cross-validated estimators is not much worse than that of the optimal estimators, so five-fold cross-validation is an adequate method for choosing the smoothing parameter, although there is some room for improvement.

To corroborate these findings in a spatial scenario, we simulated a spatial model (3) with components $\phi_1(\mathbf{t}) \propto \exp(-.5 \|\mathbf{t} - (.3, .3)\|^2 / .1^2)$ and $\phi_2(\mathbf{t}) \propto \exp(-.5 \|\mathbf{t} - (.7, .7)\|^2 / .1^2)$, for $\mathbf{t} \in [0, 1]^2$, and component scores $U_1 \sim \text{Gamma}(30, 1)$ and $U_2 \sim \text{Gamma}(20, 1)$. For estimation we used renormalized Gaussian radial kernels as basis family \mathcal{B} . Explicitly, $\beta_k(t) = \exp\{-\|\mathbf{t} - \boldsymbol{\tau}_k\|^2 / 2\sigma_k^2\} / \sum_{j=1}^q \exp\{-\|\mathbf{t} - \boldsymbol{\tau}_j\|^2 / 2\sigma_j^2\}$, where the $\boldsymbol{\tau}_k$ s are points in $[0, 1]^2$ and each σ_k is taken as half the distance between $\boldsymbol{\tau}_k$ and the closest $\boldsymbol{\tau}_j$. As $\boldsymbol{\tau}_k$ s we took q uniformly spaced points in $[0, 1]^2$, and we considered two cases: $q = 25$ and $q = 64$. We considered two sample sizes, $n = 50$ and $n = 100$. The smoothing parameter ζ was chosen by five-fold cross-validation on a few preliminary trial samples. Biases, standard deviations and root mean squared errors (defined above) based on 300 Monte Carlo replications are shown in Table 2. The conclusions are in line with the temporal models: the dominant component of the

Component	$n = 50$					
	$q = 25$			$q = 64$		
	bias	std	rmse	bias	std	rmse
ϕ_1	1.34	.020	1.34	.38	.13	.40
ϕ_2	1.38	.052	1.38	.50	.17	.53
	$n = 100$					
ϕ_1	1.34	.012	1.34	.37	.09	.38
ϕ_2	1.38	.035	1.38	.48	.12	.50

Table 2: Simulation Results. Bias, standard deviation and root mean squared error of independent component estimators for spatial models.

mean squared error is the bias, not the variance, so increasing the sample size has little effect on the rmse whereas increasing the number of basis functions does reduce the rmse significantly. Therefore, once again the recommendation is to use a family \mathcal{B} as large as feasible (given the computational time) in order to reduce bias, since the availability of replications (and the smoothing parameter ζ) will take care of the variance.

To conclude this section, we will go back to the temporal models and compare the behavior of the new method with the method of Wu et al. (2013). Direct comparison of the component estimators is not possible because the method of Wu et al. (2013) is not based on a non-negative decomposition for the process $\Lambda(t)$ like model (3). Briefly, the method of Wu et al. (2013) works as follows: first, $\mu(t) = E\{\tilde{\Lambda}(t)\}$ and $\rho(s, t) = \text{Cov}\{\tilde{\Lambda}(s), \tilde{\Lambda}(t)\}$ are estimated by kernel smoothers, and the eigenfunctions of ρ , which can be negative, are computed; then the individual densities $\tilde{\lambda}_i$ are estimated using the eigenfunctions of $\rho(s, t)$ as basis, truncating and renormalizing the $\tilde{\lambda}_i$ s if necessary to make them non-negative; finally, the intensities λ_i are estimated as $\lambda_i = m_i \tilde{\lambda}_i$, where m_i is the number of observations for replication i . Since we cannot compare the component estimators directly, we will compare the intensity estimators and the density estimators. The method of Wu et al. (2013) has three main tuning parameters: the bandwidths of the kernel smoothers and the number of components to include in the expansions of the $\tilde{\lambda}_i$ s. Wu et al. (2013) discuss a number of strategies to choose these parameters. Here we use the optimal Gaussian bandwidths for the kernel smoothers (Silverman, 1986, ch. 3.4 and 4.3), which produce reasonably smooth estimators of $\mu(t)$ and $\rho(s, t)$ in our simulations, and instead of choosing the number of components p using Wu et al.'s suggestions we

Model 1				
Estimator	$n = 50$		$n = 150$	
	intensity	density	intensity	density
IC-based, cv	14	.26	11	.18
IC-based, opt	10	.17	10	.15
PC-based, opt	13	.18	12	.14

Model 2				
IC-based, cv	7.0	.10	6.5	.09
IC-based, opt	6.8	.10	6.7	.10
PC-based, opt	9.6	.12	9.3	.11

Table 3: Simulation Results. Root mean squared errors of intensity and density function estimators using the independent component model (IC) and Wu et al.’s principal component model (PC).

simply report the estimation errors for the optimal p , which happens to correspond to $p = 1$. Note that Wu et al.’s model has a mean while (3) does not, so a one-principal-component model has the same dimension as a two-independent-component model.

Table 3 shows the estimated root mean squared errors $\{E(\sum_{i=1}^n \|\hat{\lambda}_i - \lambda_i\|^2/n)\}^{1/2}$ and $\{E(\sum_{i=1}^n \|\hat{\tilde{\lambda}}_i - \tilde{\lambda}_i\|^2/n)\}^{1/2}$, where $\|\cdot\|$ is the L^2 norm, based on 500 Monte Carlo replications. In fact, the same simulated data and the same independent-component estimators as in Table 1 were used. We only report the results for the ten-knot estimators, since they were uniformly better than the five-knot estimators. Again, we report the errors of both the cross-validated and the optimal independent-component-based estimators. Note that for Wu et al.’s principal-component-based estimators we are choosing the optimal p , so the errors observed in practice will be somewhat higher than those reported in Table 3. Unlike the results in Table 1, the results in Table 3 do change with the model. For Model 1 the results are mixed, depending on the sample size and on whether we are estimating the density function or the intensity function. But for Model 2 the independent-component-based estimators, even the cross-validated ones, clearly outperform Wu et al.’s principal-component-based estimators.

6 Application: Street Theft in Chicago

In this section we analyze the spatial distribution of street robberies in Chicago during the year 2014. The data was downloaded from the City of Chicago Data Portal, a very extensive data repository that provides, among other things, detailed information about every reported criminal incident in the city. The information provided includes type, date, time, and coordinates (latitude and longitude) of the incident. Here we will focus on crimes typified as of primary type “theft” and location “street”. There were 16,278 reported incidents of this type between January 1, 2014 and December 31, 2014. Their locations cover most of the city, as shown in Figure 2(a). These data can be seen as a single realization of a spatiotemporal process (as Li and Guan, 2014, do for a similar data set), or as $n = 365$ realizations of a spatial process, where the variable ‘time’ (the day) is just the index of replication. For the purposes of this paper, we take the second view.

Investigating the spatial and temporal distribution of crime is important because, as noted by Ratcliffe (2010), crime opportunities are not uniformly distributed in space and time, and the discovery and analysis of spatial patterns may help better understand the role of geography and opportunity in the incidence of crime. The geographical analysis of crime is not new and goes back at least to Guerry and Quetelet in the early 1800s (Friendly, 2007), but technological and data limitations hampered earlier attempts at crime mapping. Only in the last few decades have software and hardware developed to the point of allowing efficient crime mapping, and only in the last few years have extensive crime data repositories like the City of Chicago Data Portal become available.

Awareness of the geographical and temporal distribution of crime is necessary, for example, for efficient allocation of crime-prevention resources. Sometimes even the most basic spatial analysis reveals unexpected results. Figure 2(b) shows a kernel-density estimator for the points in Figure 2(a). Those familiar with the Chicago area will immediately notice that the mode occurs at the neighborhoods of Wicker Park and Pulaski, an entertainment area with high pedestrian traffic that is not typically associated with crime and violence in people’s mind. But it is well known to criminologists that street robbery tends to concentrate not on deprived neighborhoods per se but on places that they denominate “crime attractors”, areas that “bring together, often in large numbers, people who carry cash, some of whom are distracted

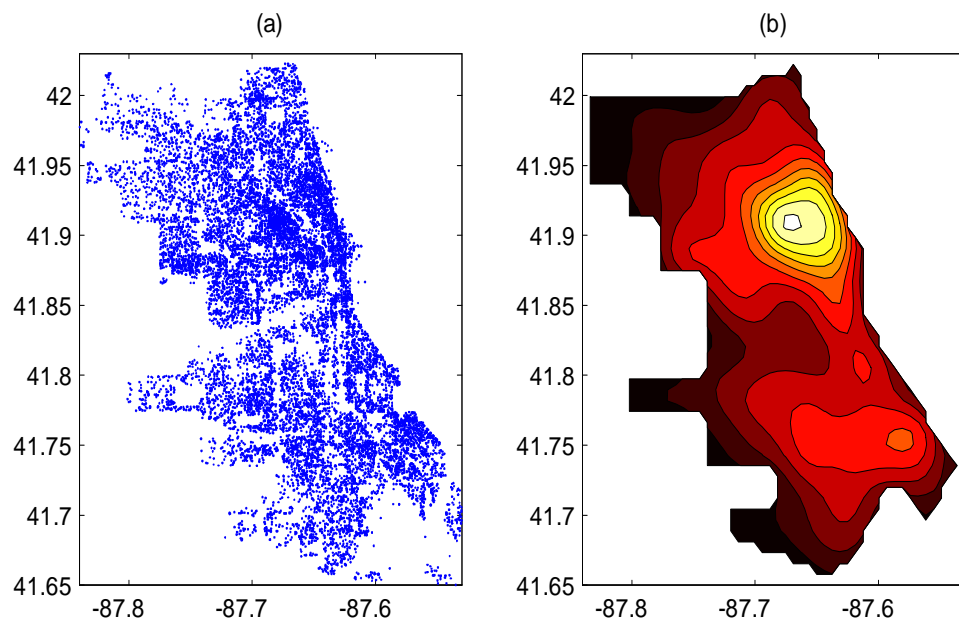


Figure 2: Chicago Street Theft. (a) Location of reported incidents in the year 2014. (b) Kernel density estimator of the data in (a).

and vulnerable” (Bernasco and Block, 2011).

We fitted independent component models for these data, using the 365 days as replications. We considered p s ranging from 2 to 5 and component scores with Gamma distribution. As basis family \mathcal{B} for the components we used renormalized Gaussian radial kernels $\beta_k(\mathbf{t}) = \exp\{-\|\mathbf{t} - \boldsymbol{\tau}_k\|^2 / 2\sigma_k^2\} / \sum_{j=1}^q \exp\{-\|\mathbf{t} - \boldsymbol{\tau}_j\|^2 / 2\sigma_j^2\}$ with $q = 49$ uniformly spaced centers $\boldsymbol{\tau}_k$ in the rectangle $[-87.84, -87.53] \times [41.65, 42.03]$, which is the smallest rectangle that includes the domain of interest B (the city of Chicago), and σ_k is taken as half the distance between $\boldsymbol{\tau}_k$ and the closest $\boldsymbol{\tau}_j$. The smoothing parameter ζ was chosen by five-fold cross-validation for each p . The cross-validated criteria for $p = 2, \dots, 5$ were 179.37, 177.73, 186.04 and 185.05, respectively, so we chose the model with four components. For this model the $\hat{\alpha}_k$ s were 4.95, 4.38, 4.04 and 3.51, and $\hat{\beta}$ was 2.70, so the respective expected values of the U_k s were 13.37, 11.85, 10.93 and 9.50. The overall expected number of robberies per day in the city is then 45.64.

Contour plots of the four components are shown in Figure 3. We computed the asymptotic variance estimators derived in Section 4 for the component coefficients and it turned out that many \hat{c}_{kj} s were indeed not significantly different from zero, but the plots obtained after filtering out these components were virtually identical to those in Figure 3, so they are not shown here.

The components in Figure 3 are well-localized and easily interpretable. The first component has a mode on the border of Near North and Lincoln Park neighborhoods, and extends northwards towards Lakeview and Uptown. These are the most affluent neighborhoods in Chicago, normally not associated with crime in people’s perceptions, but the large number of affluent people on the streets make them ideal “attractors” for street theft. The second component has a mode on Washington Park and extends eastwards towards the University of Chicago campus, neighborhoods in the South side of the city that are well-known to be problematic; the proportion of households below poverty level in Washington Park is 39%, and the unemployment level is 23%.

The third component is bimodal, with the main mode at West Town, roughly where the mode of the overall mean was (Figure 3(b)), and a second mode at Woodlawn and South Shore neighborhoods. These are also neighborhoods with generally high levels of crime and poverty, although not as high as Washington Park. The fourth component is also bimodal. The main mode concentrates on the West side

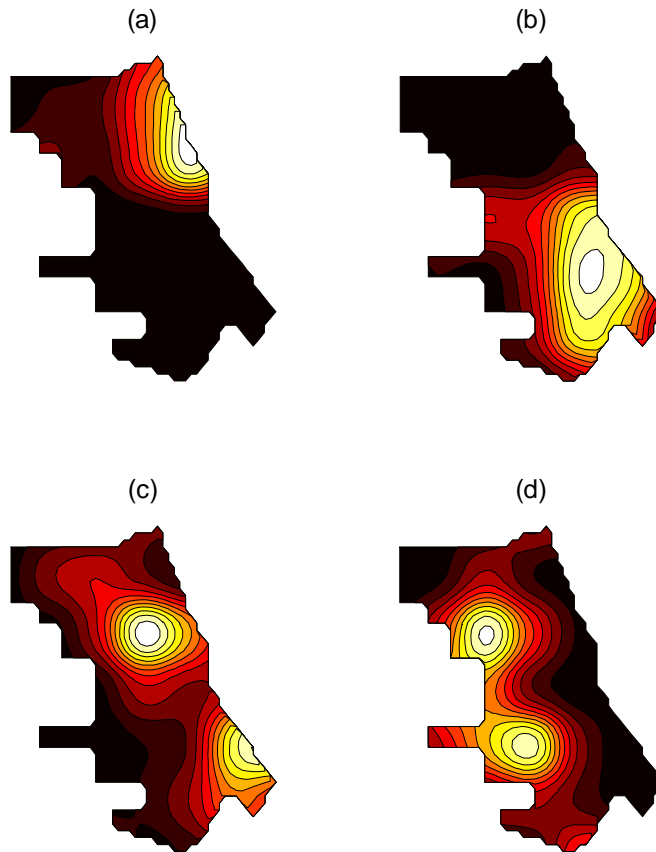


Figure 3: Chicago Street Theft. Contour plot of (a) first, (b) second, (c) third, and (d) fourth component.

neighborhoods of Humboldt Park and West Garfield, and the second mode on the South side neighborhoods of West Englewood and Gage Park. These are the worst neighborhoods in the city in terms of overall crime and poverty; for example, the proportion of households below poverty level is 33% in Humboldt Park and 32% in West Englewood, and the percentages of people with no high school diploma are 37% and 30%, respectively. These similarities in demographics suggest that the bimodality of the components is not an artifact of the estimators but a real feature of the data.

Model (3) also provides estimators for the individual intensity functions, $\hat{\lambda}_i(t) = \sum_{k=1}^p \hat{u}_{ik} \hat{\phi}_k(t)$. Two examples are shown in Figure 4. Figure 4(a) corresponds to November 8, the day with the largest first component score, and Figure 4(c) to July 25, the day with the largest fourth component score. There was a total of 43 thefts in November 8, 35 of which were associated with the first component, as determined by the \hat{y}_{ijk} s; specifically, for each incident t_{ij} we found $\hat{k}_{ij} = \operatorname{argmax}_k \hat{y}_{ijk}$, which indicates what component t_{ij} is most strongly associated with. The corresponding estimator of the intensity function is shown in Figure 4(b), which accurately reflects the distribution of the points in Figure 4(a). In July 25 there was a total of 72 street robberies, 17 of which were associated with the first component, 17 with the second, 0 with the third and 38 with the fourth. The corresponding intensity function is shown in Figure 4(d), and it is again an accurate representation of the distribution of the points in Figure 4(c). If we were using individual kernel smoothers for the intensity functions, with only 43 or even 72 data points we would not be able to obtain estimators that are smooth and fine-detailed at the same time, like those in Figure 4; a small bandwidth would produce irregular estimators and a large bandwidth would produce featureless estimators.

An analysis of the component scores also reveals interesting patterns. It is known that some types of crime tend to follow seasonal patterns; violent crimes in particular tend to increase during the summer (Field, 1992). Figure 5 shows that this is indeed the case for street theft in Chicago. Even though there is a large variability, the seasonal trends are plain to see, in particular for the second component which shows a very clear peak in July. The first and third components also reach their maxima in July, although their curves are flatter. The fourth component, on the other hand, has two local peaks in April and September. Whether these peaks are systematic across the years may be established by analyzing data from previous years, since the

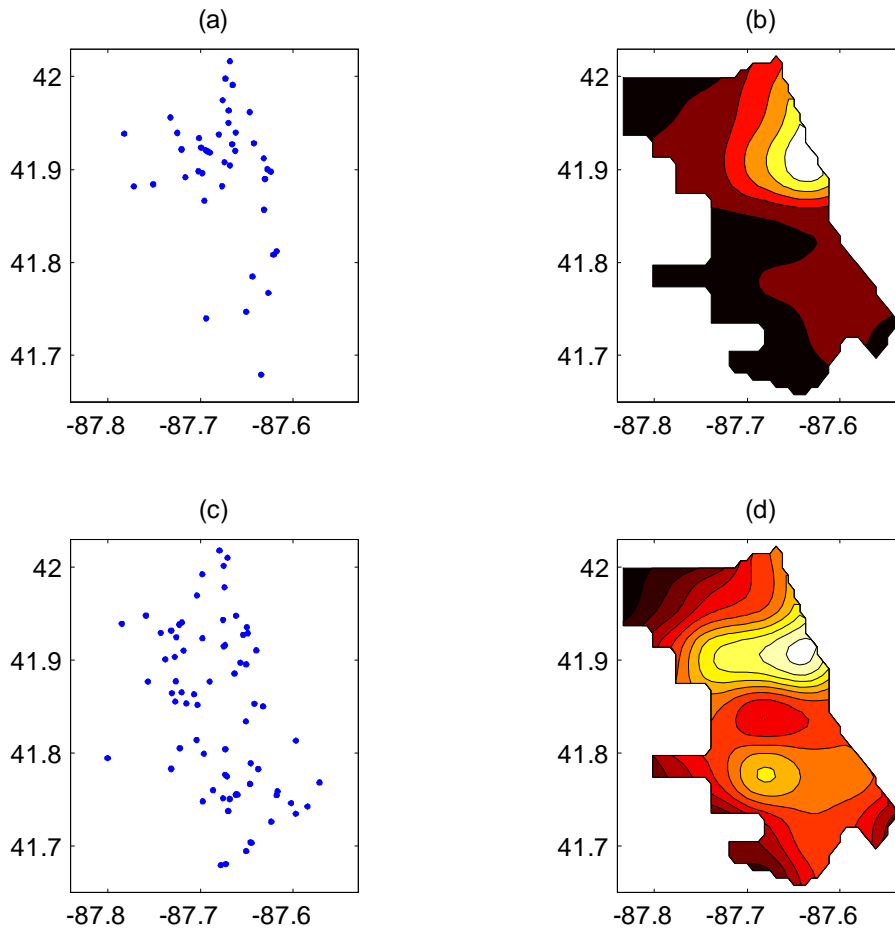


Figure 4: Chicago Street Theft. Incidents reported in (a) November 8 and (b) July 25 of 2014. Corresponding intensity estimators based on a four-component model, for (b) November 8 and (d) July 25.




Figure 5: Chicago Street Theft. Daily component scores and lowess smoother (solid line) for (a) first, (b) second, (c) third, and (d) fourth component.




Figure 6: Chicago Street Theft. Kernel density estimators of the distribution of time of reported incidents, for incidents associated with (a) first, (b) second, (c) third, and (d) fourth component.

City of Chicago Data Portal has daily crime data going back to 2001. These trends seem to indicate that the replications, although independent, may not be identically distributed. This temporal pattern can be accommodated with a slight modification of model, by assuming that the component scores u_{1k}, \dots, u_{nk} are independent with $u_{ik} \sim \text{Gamma}(\alpha_k(i), \beta)$ for functions $\alpha_k(i)$ that can be modeled, parametrically or semiparametrically, to fit the seasonal patterns in Figure 5.

Finally, we investigated the distribution of the time of the incidents. Figure 6 shows kernel density estimators of the time of robbery for each component (using Matlab's function `ksdensity` with its default options). Since the minimum always occurs at 5am, for better visualization we shifted the points between 0 and 4am to

the other end by adding 24, so e.g. 3am became 27. We see that the largest number of incidents tend to occur around 8 pm for all regions, but while the distribution is strongly unimodal for the North side neighborhoods (Figure 6(a)), for the other regions there are clear plateaus earlier in the day (Figures 6(b–d)). These plots, however, may hide seasonal variations like those seen in Figure 5. A more thorough analysis of these variations would need models that incorporate covariates, but this is beyond the scope of this paper.

Acknowledgement

This research was partly supported by US National Science Foundation grant DMS 1505780.

References

- Ash, R.B. and Gardner, M.F. (1975). *Topics in stochastic processes*. Academic Press, New York.
- Baddeley, A.J., Møller, R.A., Howard, C.V., and Boyde, A. (1993). Analysis of a three-dimensional point pattern with replication. *Applied Statistics* **42** 641–668.
- Bell, M.L., and Grunwald, G.K. (2004). Mixed models for the analysis of replicated spatial point patterns. *Biostatistics* **5** 633–648.
- Bernasco, W., and Block, R. (2011). Robberies in Chicago: a block-level analysis of the influence of crime generators, crime attractors, and offender anchor points. *Journal of Research in Crime and Delinquency* **48** 33–57.
- Bilodeau, M., and Brenner, D. (1999). *Theory of Multivariate Statistics*. Springer, New York.
- Bouzas, P.R., Valderrama, M., Aguilera, A.M., and Ruiz-Fuentes, N. (2006). Modelling the mean of a doubly stochastic Poisson process by functional data analysis. *Computational Statistics and Data Analysis* **50** 2655–2667.

- Bouzas, P.R., Ruiz-Fuentes, N., and Ocaña, F.M. (2007). Functional approach to the random mean of a compound Cox process. *Computational Statistics* **22** 467–479.
- Cox, D.R., and Isham, V. (1980). *Point Processes*. Chapman and Hall/CRC, Boca Raton.
- Dempster, A.P., Laird, N.M., and Rubin, D.B. (1977). Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society Series B* **39** 1–38.
- Diggle, P.J. (2013). *Statistical Analysis of Spatial and Spatio-Temporal Point Patterns, Third Edition*. Chapman and Hall/CRC, Boca Raton.
- Diggle, P.J., Lange, N., and Beneš, F.M. (1991). Analysis of variance for replicated spatial point patterns in clinical neuroanatomy. *Journal of the American Statistical Association* **86** 618–625.
- Diggle, P.J., Mateau, J., and Clough, H.E. (2000). A comparison between parametric and nonparametric approaches to the analysis of replicated spatial point patterns. *Advances in Applied Probability* **32** 331–343.
- Diggle, P.J., Eglen, S.J., and Troy, J.B. (2006). Modeling the bivariate spatial distribution of amacrine cells. In *Case Studies in Spatial Point Process Modeling*, eds. A. Baddeley et al., New York: Springer, pp. 215–233.
- Fernández-Alcalá, R.M., Navarro-Moreno, J., and Ruiz-Molina, J.C. (2012). On the estimation problem for the intensity of a DSMPP. *Methodology and Computing in Applied Probability* **14** 5–16.
- Field, S. (1992). The effect of temperature on crime. *British Journal of Criminology* **32** 340–351.
- Friendly, M. (2007). A.-M. Guerry’s moral statistics of France: challenges for multivariable spatial analysis. *Statistical Science* **22** 368–399.
- Gervini, D. (2009). Detecting and handling outlying trajectories in irregularly sampled functional datasets. *The Annals of Applied Statistics* **3** 1758–1775.

- Geyer, C.J. (1994). On the asymptotics of constrained M-estimation. *The Annals of Statistics* **22** 1993–2010.
- Hiriart-Urruty, J.-B., and Lemaréchal, C. (2001). *Fundamentals of Convex Analysis*. Springer, NY:
- Hyvärinen, A., Karhunen, J., and Oja, E. (2001). *Independent Component Analysis*. Wiley, New York.
- Jalilian, A., Guan, Y., and Waagpetersen, R. (2013). Decomposition of variance for spatial Cox processes. *Scandinavian Journal of Statistics* **40** 119–137.
- James, G., Hastie, T., and Sugar, C. (2000). Principal component models for sparse functional data. *Biometrika* **87** 587–602.
- Knight, K., and Fu, W. (2000). Asymptotics for lasso-type estimators. *The Annals of Statistics* **28** 1356–1378.
- Landau, S., Rabe-Hesketh, S., and Everall, I.P. (2004). Nonparametric one-way analysis of variance of replicated bivariate spatial point patterns. *Biometrical Journal* **46** 19–34.
- Li, Y., and Guan, Y. (2014). Functional principal component analysis of spatiotemporal point processes with applications in disease surveillance. *Journal of the American Statistical Association* **109** 1205–1215.
- Møller, J., and Waagepetersen, R.P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall/CRC, Boca Raton.
- Pawlas, Z. (2011). Estimation of summary characteristics from replicated spatial point processes. *Kybernetika* **47** 880–892.
- Pollard, D. (1984). *Convergence of Stochastic Processes*. New York: Springer.
- Ratcliffe, J. (2010). Crime mapping: spatial and temporal challenges. In *Handbook of Quantitative Criminology*, A.R. Piquero and D. Weisburd (eds.), pp. 5–24. New York: Springer.
- Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall.

- Snyder, D.L., and Miller, M.I. (1991). *Random Point Processes in Time and Space*. Springer, New York.
- Streit, R.L. (2010). *Poisson Point Processes: Imaging, Tracking, and Sensing*. Springer, New York.
- Van der Vaart, A. (2000). *Asymptotic Statistics*. Cambridge University Press, Cambridge, UK.
- Wager, C.G., Coull, B.A., and Lange, N. (2004). Modelling spatial intensity for replicated inhomogeneous point patterns in brain imaging. *Journal of the Royal Statistical Society Series B* **66** 429–446.
- Wu, S., Müller, H.-G., and Zhang, Z. (2013). Functional data analysis for point processes with rare events. *Statistica Sinica* **23** 1–23.
- Xun, X., Cao, J., Mallick, B., Maity, A., and Carroll, R.J. (2013). Parameter estimation of partial differential equations. *Journal of the American Statistical Association* **108** 1009–1020.
- Yao, F., Müller, H.-G., and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association* **100** 577–590.
- Yu, Y., and Ruppert, D. (2002). Penalized spline estimation for partially linear single-index models. *Journal of the American Statistical Association* **97** 1042–1054.