

# Exploring patterns of demand in bike sharing systems via replicated point process models

Daniel Gervini and Manoj Khanal  
Department of Mathematical Sciences  
University of Wisconsin–Milwaukee

July 3, 2018

## Abstract

Understanding patterns of demand is fundamental for fleet management of bike sharing systems. In this paper we analyze data from the Divvy system of the city of Chicago. We show that the demand of bicycles can be modeled as a multivariate temporal point process, with each dimension corresponding to a bike station in the network. The availability of daily replications of the process allows nonparametric estimation of the intensity functions, even for stations with low daily counts, and straightforward estimation of pairwise correlations between stations. These correlations are then used for clustering, revealing different patterns of bike usage.

*Key words:* Functional canonical correlation; functional principal components; hierarchical clustering; Poisson process; spline smoothing.

# 1 Introduction

Bike sharing systems are becoming increasingly common in large cities around the world (Shaheen et al., 2010). These systems provide short-term bicycle rental services at unattended stations distributed throughout the city. A user checks out a bicycle at a station near the intended origin of the journey and returns it at a station near the intended destination. For the system to run smoothly, it is necessary that both bicycles and docks be available at every station. When no bicycles are available at the intended origin of a trip or no docks are available at the intended destination, users need to look for alternative nearby stations, which may dissuade them from using the system altogether. Since bike flow from one station to another is rarely matched by a similar flow in the reverse direction, imbalances in the spatial distribution of bikes inevitably arise (Nair and Miller-Hooks, 2011). There are different strategies to manage this problem. For example, bikes are manually relocated by trucks as part of the day-to-day operations of the system. From a longer-term perspective, careful planning of the location of new stations is important. In order to make good short- and long-term decisions, understanding the spatiotemporal patterns of bike demand is fundamental.

In this paper we show that bike demand at each station can be modeled as a temporal point process, where bike checkout times are the random events of interest. Bike return can be modeled in a similar way. We will analyze data from the Divvy system of the city of Chicago, publicly available at the Chicago Data Portal website (<https://data.cityofchicago.org>). Specifically, we will analyze bike trips that took place between April 1 and November 30 of 2016, since bike usage considerably decreases during the winter. There were a total of 3,068,211 bike trips and 458 active bike stations in that period. Demand varied a lot depending on the station location, from a lowest of 29 annual trips for station 386 in the South Side to a highest of 85,314 annual trips for station 35 at the Navy Pier. For stations with relatively large daily counts, the distribution of bike demand on any given day can be estimated by kernel smoothing or other density estimation methods (Silverman, 1986). But for stations with low daily counts this is not possible, at least not in a meaningful way. In this paper we propose a new method that overcomes this deficiency by “borrowing” data across replications, i.e. across different days. In the end, estimators of daily distributions of bike demand are obtained, even for stations with low daily counts,

but in an indirect way. These estimators are then used to study spatial correlations between stations and to derive clusters that correspond to different usage patterns.

To avoid confusion, let us clarify that in this paper we use the terms ‘realization’, ‘replication’ and ‘observation’ the way they are used in functional data analysis, which is somewhat different from the point-process literature. By ‘realization’ or ‘replication’, which for us are synonyms, we refer to a realization of the whole process, that is, the whole set of observations on any given day. By ‘observation’ we refer to an individual point in a realization of the process, that is, to a bike trip on any given day. Thus, our data set contains 244 replications or realizations of the process, each with a varying number of observations.

To put the problem in context, we note that different aspects of bike sharing systems have been studied in the specialized literature (e.g. Borgnat *et al.*, 2011; Vogel *et al.*, 2011; Nair *et al.*, 2013), but the problem of estimating and modeling daily demand distribution at every station in a network has not yet been addressed, to the best of our knowledge. From a statistical methodology perspective, we can mention early work on replicated point processes by Diggle *et al.* (1991), Baddeley *et al.* (1993), Diggle *et al.* (2000), Mateu (2001) and Landau *et al.* (2004), but these authors propose tests for various hypotheses using summary statistics of the process (see also Baddeley *et al.*, 2015, ch. 16; Diggle, 2013, ch. 5.4), rather than explicitly estimating the intensity functions of the processes, as we do here. More recent work that does address the intensity-function estimation problem was done by Wu *et al.* (2013), Bouzas and Ruiz-Fuentes (2015) and Gervini (2016), but only in the context of univariate processes, not multivariate ones as in this paper. Spatio-temporal processes have been widely studied in the literature, but mostly in the single-replication context (see e.g. Li and Guan, 2014; Shirota and Gelfand, 2017; Diggle, 2013, and references therein), not in the many-replication context of this paper. Finally, we mention that clustering methods for spatial functional data have been proposed by Delicado *et al.* (2010), Romano *et al.* (2010), Secchi *et al.* (2013), and Menafoglio and Secchi (2017), among others, but again in the context of a single datum per site, which does not allow direct estimation of spatial correlations and requires assumptions such as isotropy; in our application, the availability of many replications per site allows us to estimate spatial correlations directly and without isotropy assumptions, which, in fact, we show not to hold for the bike sharing network.

## 2 Modeling daily bike demand

### 2.1 Poisson point processes

Let  $X_{ij}$  be the set of checkout times for day  $i$  at bike station  $j$ . In our data set we have  $n = 244$  days and  $d = 458$  bike stations. Each  $X_{ij}$  is a finite but otherwise random set, so it is best modeled as a point process. The collection  $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$  can be seen as a realization of a multivariate point process. For an overview of point processes, see Møller and Waagepetersen (2004, ch. 2), Streit (2010, ch. 2) or Baddeley (2007).

A temporal point process  $X$  is a random countable set in  $[0, \infty)$ . A process is locally finite if  $\#(X \cap B) < \infty$  with probability one for any bounded interval  $B$ , in which case we can define the count function  $N(B) = \#(X \cap B)$ . A Poisson process is a locally finite process for which there exists a nonnegative locally integrable function  $\lambda(t)$  such that (i)  $N(B)$  follows a Poisson distribution with rate  $\int_B \lambda(t) dt$  for any bounded  $B$ , and (ii) for disjoint sets  $B_1, \dots, B_k$  the random variables  $N(B_1), \dots, N(B_k)$  are independent. A consequence of (i) and (ii) is that the conditional distribution of the points in  $X \cap B$  given  $N(B) = m$  is the distribution of  $m$  independent and identically distributed observations with density  $\lambda(t) / \int_B \lambda$ . The function  $\lambda$  is called the intensity function of the process.

### 2.2 The model

In our application we have  $nd$  processes  $X_{ij}$  with  $i = 1, \dots, n$  and  $j = 1, \dots, d$ , each with a corresponding intensity function  $\lambda_{ij}$  on the interval  $[a, b] = [0, 24]$ . Since the  $\lambda_{ij}$ s are nonnegative, for simplicity we will assume that they are positive everywhere, even if negligible in some regions, and model their logarithms using additive principal component models, similar to those used in functional data analysis (Ramsay and Silverman, 2005, ch. 8).

For each station  $j$  we assume

$$\log \lambda_{ij}(t) = \mu_j(t) + \sum_{k=1}^{p_j} u_{ikj} \phi_{kj}(t), \quad i = 1, \dots, n, \quad (1)$$

where  $\mu_j(t)$  is the annual mean function for station  $j$  and  $\{\phi_{kj}\}_{k=1}^{p_j}$  are orthonormal functions (across  $ks$ ) that account for various types of deviations from the mean.

We will refer to the  $\phi_{kj}$ s as components and the  $u_{ikj}$ s as component scores. The component scores are, in principle, random effects with  $E(u_{ikj}) = 0$ ,  $\sigma_{kj}^2 = V(u_{ikj})$  and  $\text{cov}(u_{ikj}, u_{ik'j}) = 0$  for  $k \neq k'$ . Without loss of generality we assume  $\sigma_{1j}^2 \geq \dots \geq \sigma_{p_j j}^2 > 0$ . However, for estimation purposes we will treat the  $u_{ikj}$ s as fixed effects, which does not require distributional assumptions on the  $u_{ikj}$ s that may be questionable.

Model (1) for  $\log \lambda_{ij}$  turns into a multiplicative model for  $\lambda_{ij}$ :

$$\lambda_{ij}(t) = \lambda_{0j}(t) \prod_{k=1}^{p_j} \psi_{kj}(t)^{u_{ijk}}, \quad (2)$$

where  $\lambda_{0j}(t) = \exp \mu_j(t)$  and  $\psi_{kj}(t) = \exp \phi_{kj}(t)$ . We will refer to  $\lambda_{0j}(t)$  as the baseline intensity function for station  $j$ .

Since the  $\lambda_{ij}$ s are not directly observable, the mean  $\mu_j(t)$  and the components  $\phi_{kj}(t)$  must be estimated from the data. To facilitate this, we use spline models (De Boor, 1978):

$$\mu_j(t) = \sum_{l=1}^q c_{l0j} \beta_l(t), \quad \phi_{kj}(t) = \sum_{l=1}^q c_{lkj} \beta_l(t), \quad (3)$$

where  $\{\beta_l\}_{l=1}^q$  is a spline basis. We use B-splines in this paper, but other bases can be used, even non-spline bases such as the Fourier basis. Modeling  $\mu_j$  and the  $\phi_{kj}$ s as spline functions turns the functional estimation problem into a simpler multivariate problem of estimating basis coefficients  $\mathbf{c}_{kj} = (c_{1kj}, \dots, c_{qkj})^T$ . It also simplifies the introduction of periodicity constraints: the intensity functions should satisfy  $\lambda_{ij}(0) = \lambda_{ij}(24)$  in this application, which is enforced by the simple linear constraints  $\mathbf{c}_{kj}^T \boldsymbol{\beta}(a) = \mathbf{c}_{kj}^T \boldsymbol{\beta}(b)$ , where  $\boldsymbol{\beta}(t) = (\beta_1(t), \dots, \beta_q(t))^T$ . Similarly, the orthonormality of the  $\phi_{kj}$ s is enforced by the constraints  $\mathbf{c}_{kj}^T \mathbf{J} \mathbf{c}_{k'j} = \delta_{k,k'}$ , where  $\mathbf{J} = \int_a^b \boldsymbol{\beta}(t) \boldsymbol{\beta}(t)^T dt$  and  $\delta_{k,k'}$  is Kronecker's delta.

## 2.3 Estimation

Fitting model (1), then, involves estimation of the parameters  $\mathbf{c}_{kj}$ s in (3) and of the component scores  $u_{ikj}$ , which, for estimation purposes, will be treated as fixed effects. We do this by maximum likelihood, using the Poisson model as working model. In view of the above-mentioned properties of the Poisson process, the density function

of  $X_{ij} = \{t_{ij1}, \dots, t_{ijm_{ij}}\}$  is

$$f_{ij}(m_{ij}, t_{ij1}, \dots, t_{ijm_{ij}}) = \exp \left\{ - \int_a^b \lambda_{ij}(t) dt \right\} \frac{1}{m_{ij}!} \prod_{l=1}^{m_{ij}} \lambda_{ij}(t_{ijl}), \quad (4)$$

where  $f_{ij}(0, \emptyset) = \exp \left( - \int_a^b \lambda_{ij} \right)$  if  $X_{ij} = \emptyset$  and  $m_{ij} = 0$ . Then the log-likelihood function for station  $j$ , ignoring the constant factor  $1/m_{ij}!$ , is

$$\ell_j = - \sum_{i=1}^n \int \lambda_{ij} + \sum_{i=1}^n \sum_{l=1}^{m_{ij}} \log \lambda_{ij}(t_{ijl}). \quad (5)$$

In principle, the estimators  $\hat{\mathbf{c}}_{kj}$ s and  $\hat{u}_{ikj}$ s would be the maximizers of  $\ell_j$ . However, maximizing  $\ell_j$  without any sort of roughness penalty will produce irregular estimators of  $\mu_j$  and the  $\phi_{kj}$ s if the spline basis dimension  $q$  is large. The roughness of a function  $g$  can be measured by the functional norm of its second derivative,  $\int_a^b (g'')^2$ . So we will define the  $\hat{\mathbf{c}}_{kj}$ s and preliminary estimators of the scores  $\tilde{u}_{ikj}$ s as the maximizers of the penalized log-likelihood function

$$\begin{aligned} P\ell_j &= \frac{1}{n} \ell_j - \xi_1 \int (\mu_j'')^2 - \xi_2 \sum_{k=1}^{p_j} \int (\phi_{kj}'')^2 \\ &= \frac{1}{n} \ell_j - \xi_1 \mathbf{c}_{0j}^T \mathbf{\Omega} \mathbf{c}_{0j} - \xi_2 \sum_{k=1}^{p_j} \mathbf{c}_{kj}^T \mathbf{\Omega} \mathbf{c}_{kj}, \end{aligned} \quad (6)$$

where  $\mathbf{\Omega} = \int_a^b \boldsymbol{\beta}''(t) \boldsymbol{\beta}''^T(t) dt$  and  $\xi_1$  and  $\xi_2$  are non-negative tuning parameters that regulate the degree of smoothness of  $\mu_j$  and the  $\phi_{kj}$ s, respectively. The maximization has to be carried out subject to the periodicity and orthonormality constraints

$$\begin{aligned} \mathbf{c}_{kj}^T \boldsymbol{\beta}(a) &= \mathbf{c}_{kj}^T \boldsymbol{\beta}(b), \quad k = 0, \dots, p_j, \\ \mathbf{c}_{kj}^T \mathbf{J} \mathbf{c}_{k'j} &= \delta_{k,k'}, \quad k, k' = 1, \dots, p_j, \end{aligned}$$

and, since the true random effects  $u_{ikj}$  are zero-mean uncorrelated variables (across

ks), we also impose the following constraints on the  $u_{ikj}$ s for estimation:

$$\begin{aligned}\frac{1}{n} \sum_{i=1}^n u_{ikj} &= 0, \quad k = 1, \dots, p_j, \\ \frac{1}{n} \sum_{i=1}^n u_{ikj} u_{ik'j} &= 0, \quad k, k' = 1, \dots, p_j, \quad k \neq k'.\end{aligned}$$

Our preliminary simulations showed that the sample variances of the  $\tilde{u}_{ikj}$ s obtained this way tend to overestimate the true variances of the  $u_{ikj}$ s. To ameliorate this problem we re-scale the component scores, letting  $\hat{u}_{ikj} = \tau_j \tilde{u}_{ikj}$  and finding the optimal  $\hat{\tau}_j$  by maximum likelihood based on the  $m_{ij}$ s. That is, since  $m_{ij} \sim \mathcal{P}(\int_a^b \lambda_{ij}(t) dt)$  for a Poisson process, the log-likelihood of the  $m_{ij}$ s is

$$\tilde{\ell}_j = - \sum_{i=1}^n I_{ij}(\tau) + \sum_{i=1}^n m_{ij} \log I_{ij}(\tau),$$

where  $I_{ij}(\tau) = \int_a^b \hat{\lambda}_{ij}^{(\tau)}(t) dt$  and  $\hat{\lambda}_{ij}^{(\tau)}(t)$  is as in (7) with  $u_{ikj}$  replaced by  $\tau \tilde{u}_{ikj}$ . Then  $\hat{\tau}_j$  is the maximizer of  $\tilde{\ell}_j$ , and  $\hat{u}_{ikj} = \hat{\tau}_j \tilde{u}_{ikj}$ .

Fully-detailed algorithms to compute these estimators are explained in the Supplementary Material, and Matlab programs are available on the first author's website.

Once the mean  $\mu_j$ , the components  $\phi_{kj}$ s and the scores  $u_{ikj}$ s have been estimated, the individual daily intensity functions can be estimated from model (1) as

$$\hat{\lambda}_{ij}(t) = \exp \left\{ \hat{\mu}_j(t) + \sum_{k=1}^{p_j} \hat{u}_{ikj} \hat{\phi}_{kj}(t) \right\}. \quad (7)$$

They can subsequently be used for spatial inference regarding, for instance, cross-correlations among bike stations, as we do in Section 3.

## 2.4 Choice of tuning parameters

The models introduced above have a number of tuning parameters that have to be chosen by the user: the number of components  $p_j$  in (1), the type and dimension  $q$  of basis functions in (3), and the smoothing parameters  $\xi_1$  and  $\xi_2$  in (6). The specific type of basis functions does not have much of an impact on the final estimator, provided the dimension is large enough; we simply take cubic  $B$ -splines with equally



spaced knots in our simulations and data analyses in this paper. The dimension  $q$  is more relevant and should be relatively large, since the regularity of the estimators will be taken care of by  $\xi_1$  and  $\xi_2$  (Eilers and Marx, 1996); but for the same reason, it is not necessary to agonize over an exact choice of  $q$ . As noted by Ruppert (2002, sec. 3), although  $q$  can be chosen systematically by cross-validation, there is little change in goodness of fit after a minimum dimension  $q$  has been reached, and the fit is essentially determined by the smoothing parameters thereafter.

The choice of  $\xi_1$  and  $\xi_2$ , then, is more important, and we do it by cross-validation (Hastie *et al.*, 2009, ch. 7). Leave-one-out cross-validation finds  $\hat{\xi}_{1j}$  and  $\hat{\xi}_{2j}$  that maximize

$$\text{CV}_j(\xi_1, \xi_2) = \sum_{i=1}^n \log \hat{f}_{ij}^{[-ij]}(m_{ij}, t_{ij1}, \dots, t_{ijm_{ij}}),$$

where  $\hat{f}_{ij}^{[-ij]}$  denotes the density (4) estimated without observation  $X_{ij}$ . A faster alternative is to use  $k$ -fold cross-validation, where the data is split into  $k$  subsets that are alternatively used as test data. We use five-fold cross-validation in our implementation of the method.

The choice of the number of components  $p_j$  can also be done by cross-validation or, more practically, by the usual ad-hoc methods for choosing the number of principal components (Jolliffe, 2002, ch. 6), which take into account the relative contribution of the estimated variances  $\hat{\sigma}_{kj}^2$  and stop at a  $p_j$  where further additions of components have a negligible impact on  $\hat{\sigma}_{1j}^2 + \dots + \hat{\sigma}_{p_j j}^2$ .

## 3 Spatial correlations and clustering

### 3.1 Measuring spatial correlation

In multivariate analysis, a measure of overall correlation between two random vectors  $\mathbf{U}$  and  $\mathbf{V}$  is the canonical correlation coefficient  $\rho = \max_{\mathbf{a}, \mathbf{b}} \text{corr}(\mathbf{a}^T \mathbf{U}, \mathbf{b}^T \mathbf{V})$ , the largest possible correlation between linear combinations of  $\mathbf{U}$  and  $\mathbf{V}$  (Izenman, 2008, ch. 7.3). This coefficient can be computed as follows: given  $\Sigma_{UU}$  the covariance matrix of  $\mathbf{U}$ ,  $\Sigma_{VV}$  the covariance matrix of  $\mathbf{V}$ , and  $\Sigma_{UV}$  the cross-covariance matrix of  $\mathbf{U}$  and  $\mathbf{V}$ , then  $\rho^2$  is the largest eigenvalue of  $\Sigma_{UU}^{-1/2} \Sigma_{UV} \Sigma_{VV}^{-1} \Sigma_{VU} \Sigma_{UU}^{-1/2}$ , or equivalently, of  $\Sigma_{VV}^{-1/2} \Sigma_{VU} \Sigma_{UU}^{-1} \Sigma_{UV} \Sigma_{VV}^{-1/2}$ . The sample canonical correlation coefficient is obtained by substituting sample covariance matrices.

In functional data analysis, where  $U(t)$  and  $V(t)$  are square-integrable random functions, an equivalent version is defined (Horváth and Kokoszka, 2012):  $\rho = \max_{\alpha, \beta} \text{corr}(\langle \alpha, U \rangle, \langle \beta, V \rangle)$ , where  $\alpha$  and  $\beta$  are square-integrable functions and  $\langle f, g \rangle = \int f(t)g(t)dt$ . As we show in the Appendix, the computation of  $\rho$  can ultimately be reduced to the multivariate case by using the principal component scores of  $U(t)$  and  $V(t)$ , and the sample version is obtained by substituting sample covariance functions and estimated component scores.

In our application, we are interested in the correlations of bike demand between different stations, say  $j$  and  $j'$ , so we compute the sample functional canonical correlation coefficient  $\hat{\rho}_{jj'}$  of their respective log-intensity functions,

$$\hat{\rho}_{jj'} = \max_{\alpha, \beta} \text{corr}(\langle \alpha, \log \hat{\lambda}_{ij} \rangle, \langle \beta, \log \hat{\lambda}_{ij'} \rangle). \quad (8)$$

As explained above,  $\hat{\rho}_{jj'}^2$  is the largest eigenvalue of  $\mathbf{S}_{jj}^{-1/2} \mathbf{S}_{jj'} \mathbf{S}_{jj'}^{-1} \mathbf{S}_{j'j} \mathbf{S}_{jj}^{-1/2}$ , or equivalently of  $\mathbf{S}_{j'j'}^{-1/2} \mathbf{S}_{j'j} \mathbf{S}_{jj}^{-1} \mathbf{S}_{jj'} \mathbf{S}_{j'j'}^{-1/2}$ , where

$$\mathbf{S}_{jj'} = \frac{1}{n} \sum_{i=1}^n (\hat{\mathbf{u}}_{ij} - \bar{\mathbf{u}}_{\cdot j})(\hat{\mathbf{u}}_{ij'} - \bar{\mathbf{u}}_{\cdot j'})^T,$$

$$\hat{\mathbf{u}}_{ij} = (\hat{u}_{i1j}, \dots, \hat{u}_{ip_j j})^T \text{ and } \bar{\mathbf{u}}_{\cdot j} = \sum_{i=1}^n \hat{\mathbf{u}}_{ij} / n.$$

### 3.2 Spatial clustering

Up to this point, we have treated the  $d = 458$  bike stations in our application as generic dimensions of a multivariate point process, fitting model (1) independently for each  $j = 1, \dots, d$ . However, when the  $d$  dimensions correspond to  $d$  locations in space, as in this case, there is a spatial aspect to the problem that is interesting to investigate.

We can think of the functional canonical correlation coefficients  $\rho_{jj'}$  defined above as discretizations of a spatial correlation function  $R$ ,  $\rho_{jj'} = R(\mathbf{s}_j, \mathbf{s}_{j'})$ , where  $\mathbf{s}_j$  and  $\mathbf{s}_{j'}$  are the spatial coordinates of bike stations  $j$  and  $j'$ . In applications of spatial functional data analysis where only one observation per site is available (e.g. Delicado et al., 2010; Menafoglio and Secchi, 2017), estimation of  $R$  requires assumptions such as isotropy, i.e. that  $R(\mathbf{s}_j, \mathbf{s}_{j'}) = g(\|\mathbf{s}_j - \mathbf{s}_{j'}\|)$  for some  $g$ , in order to pool data across neighboring sites. But in our case, the availability of  $n$  replications per site allow us

straightforward estimation of  $\rho_{jj'}$  by (8) without any assumptions on  $R$ . In fact, we will show in Section 5 that isotropy does not hold for our data.

The correlations  $\hat{\rho}_{jj'}$  can be used, for instance, to discover clusters among bike stations. They can be obtained by applying standard agglomerative techniques (Izenman 2008, ch. 12.3; Hastie *et al.*, 2009, ch. 14.3.12) to distances defined by  $d_{jj'} = 1 - \hat{\rho}_{jj'}$ . For our application we found that complete linkage generally produces better results than either single or average linkage.

When the dimension  $d$  is large, the number of different pairs  $(j, j')$  can be extremely large; for example,  $d(d-1)/2 = 104,196$  in our application. So it is advisable to trim non-significant  $\hat{\rho}_{jj'}$ s prior to clustering. A test for the hypothesis  $H_{0,jj'} : \rho_{jj'} = 0$  is the following (Seber 2004, ch. 5.7.3): let  $\{r_k^2\}$  be the  $p = \min(p_j, p_{j'})$  non-zero eigenvalues of  $\mathbf{S}_{jj}^{-1/2} \mathbf{S}_{jj'} \mathbf{S}_{j'j'}^{-1} \mathbf{S}_{j'j} \mathbf{S}_{jj}^{-1/2}$ , or equivalently of  $\mathbf{S}_{j'j'}^{-1/2} \mathbf{S}_{j'j} \mathbf{S}_{jj}^{-1} \mathbf{S}_{jj'} \mathbf{S}_{j'j'}^{-1/2}$ , and  $L = \prod_{k=1}^p (1 - r_k^2)$ ; then  $Q_{jj'} = -\{n - 1 - (p_j + p_{j'} + 1)/2\} \log L$  is asymptotically  $\chi_\nu^2$  with  $\nu = p_j p_{j'}$  under the null hypothesis. To determine non-significant  $\hat{\rho}_{jj'}$ s at a simultaneous level  $\alpha$  we use Benjamini and Hochberg (1995) procedure: let  $P_{jj'} = P(\chi_{p_j p_{j'}}^2 > Q_{jj'})$  be the  $p$ -value for  $H_{0,jj'}$ , and  $\{P_{(k)}\}$  the set of these  $p$ -values sorted in increasing order; then the correlations for which  $P_{(k)} \leq \alpha k / \{d(d-1)/2\}$  are considered significant. For the non-significant  $\hat{\rho}_{jj'}$ s, we set  $\hat{\rho}_{jj'} = 0$  and then proceed to apply the linkage algorithm. Clusters, if there are any, can be found from the dendrogram using standard techniques (see Izenman 2008, ch. 12.3; Hastie *et al.*, 2009, ch. 14.3.12). The consistency of the clusters can be evaluated using measures such as the Davies–Bouldin index (Davies and Bouldin, 1979) or the Dunn index (Dunn, 1974).

## 4 Simulations

In this section we study the consistency of the estimators by simulation. We simulated data from model (1) for  $d = 1$ , since estimation is done separately for each  $j$ . We considered three distributional situations that will arise in the Divvy data analysis of Section 5: component scores that (i) are independent and identically distributed, (ii) follow a trend, and (iii) are autocorrelated. We also studied the effect of the expected number of observations per replication, the baseline rate  $\int_a^b \lambda_0(t) dt$ , which is determined by  $\mu(t)$ .

To this end we considered model (1) with  $\mu(t) = \sin(\pi t) + c$ ,  $\phi_1(t) = \sqrt{2} \sin(\pi t)$

and  $\phi_2(t) = \sqrt{2} \sin(2\pi t)$ , for  $t \in [0, 1]$ . Since  $\int_0^1 \exp\{\sin(\pi t)\} dt = 1.98$ , we took  $c = \log 5$  and  $c = \log 15$ , which give approximate baseline rates 10 and 30, respectively. The  $u_{ik}$ s were generated as follows:

1. Independent:  $u_{1k}, \dots, u_{nk}$  were independent  $N(0, \sigma_k^2)$  with  $\sigma_1 = .3\sqrt{.6}$  and  $\sigma_2 = .3\sqrt{.4}$ , respectively, so that the overall variance was .09, with the first component accounting for 60% of the variability. (The  $u_{i1}$ s were independent of the  $u_{i2}$ s in all three scenarios, since the component scores are uncorrelated across  $k$ s by definition).
2. With quadratic trend: let  $s_i = -(i - n/2)^2$ , for  $i = 1, \dots, n$ . Then  $u_{i1} = \{(s_i - \bar{s})/\text{sd}(s_i)\}\sqrt{.75}\sigma_1 + z_i\sqrt{.25}\sigma_1$ , with  $z_i$ s independent and identically distributed  $N(0, 1)$ , and  $u_{i2}$ s independent and identically distributed  $N(0, \sigma_2^2)$ , with  $\sigma_1$  and  $\sigma_2$  as in Scenario 1. The variance of  $u_{i1}$  is still  $\sigma_1^2$ , but 75% of it now comes from the quadratic trend.
3. Autocorrelated: the  $u_{i1}$ s followed the autoregressive model  $u_{i1} = z_i\sigma_e$  for  $i = 1$  and  $u_{i1} = \rho u_{i-1,1} + z_i\sigma_e$  for  $i = 2, \dots, n$ , with  $z_i$ s independent and identically distributed  $N(0, 1)$ ,  $\rho = .8$  and  $\sigma_e = \sigma_1\sqrt{1 - \rho^2}$ , so the variance of the  $u_{i1}$ s was  $\sigma_1^2$  as in the previous scenarios. The  $u_{i2}$ s were independent and identically distributed  $N(0, \sigma_2^2)$ , and we took  $\sigma_1$  and  $\sigma_2$  as in Scenarios 1 and 2.

To get an idea of the  $m_i$ s produced by these models, we generated a sample of size 100 for each baseline rate, and observed  $m_i$ s between 4 and 20 for baseline rate 10 and between 22 and 44 for baseline rate 30. Four sample sizes were considered for each scenario:  $n = 50, 100, 200$  and 400.

For estimation of the functional parameters we used a cubic B-spline basis with five equally spaced knots in  $(0, 1)$ , which has dimension  $q = 9$ , large enough for the smooth functions we are estimating. We chose subjective but visually reasonable smoothing parameters  $\xi_1 = \xi_2 = 10^{-5}$ .

Tables 1 to 3 report the results. For  $\mu$  we defined bias =  $\|E(\hat{\mu}) - \mu\|$ , std =  $[E\{\|\hat{\mu} - E(\hat{\mu})\|^2\}]^{1/2}$  and rmse =  $\{E(\|\hat{\mu} - \mu\|^2)\}^{1/2}$ , where  $\|\cdot\|$  is the usual  $L^2[0, 1]$  norm. For the  $\phi_k$ s we could not use these quantities because of the sign indetermination (a priori, it is not possible to tell if  $\hat{\phi}_k$  is estimating  $\phi_k$  or  $-\phi_k$ ), so we considered the bivariate estimators  $\hat{\phi}_k(s)\hat{\phi}_k(t)$  of  $\phi_k(s)\phi_k(t)$  instead, which are sign-invariant,

$n$	Param	rate 10			rate 30		
		bias	std	rmse	bias	std	rmse
50	$\mu$	.45	1.33	1.40	.50	.88	1.01
	$\phi_1$	.39	.71	.81	.29	.62	.69
	$\phi_2$	.61	.83	1.04	.36	.69	.78
100	$\mu$	.53	.95	1.08	.55	.60	.82
	$\phi_1$	.30	.59	.66	.14	.44	.47
	$\phi_2$	.41	.72	.83	.22	.51	.56
200	$\mu$	.52	.64	.82	.53	.43	.68
	$\phi_1$	.20	.48	.52	.10	.30	.32
	$\phi_2$	.35	.62	.71	.17	.36	.40
400	$\mu$	.53	.45	.69	.53	.30	.61
	$\phi_1$	.21	.34	.40	.08	.21	.22
	$\phi_2$	.34	.57	.66	.16	.24	.29

Table 1: Simulation Results. Bias, standard deviation and root mean squared errors of parameter estimators for independent component scores (Scenario 1). Quantities for  $\mu$  were multiplied by 10.

$n$	Param	rate 10			rate 30		
		bias	std	rmse	bias	std	rmse
50	$\mu$	.46	1.29	1.37	.50	.82	.96
	$\phi_1$	.38	.70	.80	.25	.59	.64
	$\phi_2$	.60	.83	1.02	.35	.69	.78
100	$\mu$	.48	.91	1.02	.50	.59	.77
	$\phi_1$	.28	.60	.66	.17	.48	.51
	$\phi_2$	.45	.75	.87	.24	.54	.59
200	$\mu$	.50	.64	.82	.50	.40	.64
	$\phi_1$	.23	.46	.51	.10	.33	.35
	$\phi_2$	.34	.62	.71	.18	.38	.42
400	$\mu$	.49	.45	.66	.52	.30	.60
	$\phi_1$	.18	.34	.39	.08	.22	.23
	$\phi_2$	.34	.55	.65	.18	.31	.36

Table 2: Simulation Results. Bias, standard deviation and root mean squared errors of parameter estimators for component scores with a trend (Scenario 2). Quantities for  $\mu$  were multiplied by 10.

and defined bias, standard deviation and root mean squared error as before, except that  $\|\cdot\|$  was the bivariate  $L^2$  norm on  $[0, 1] \times [0, 1]$ . The expectations were approximated by Monte Carlo based on 200 replications of each scenario.

Table 1 shows that, for independent and identically distributed component scores, the estimators behave as expected: estimation errors decrease as  $n$  increases for each baseline rate, and they are lower for the higher baseline rate. The bias of  $\hat{\mu}$  does not decrease with  $n$ , but this is due to the suboptimal choice of smoothing parameter.

Table 2 shows the results for Scenario 2, where the first component score follows a quadratic trend, and we see that they are almost identical to those in Table 1, so the estimators work equally well in both situations. Table 3 shows the results for Scenario 3, the autoregressive first component scores. The mean squared errors of  $\hat{\phi}_1$  and  $\hat{\phi}_2$  are somewhat larger than in the previous scenarios, but only by 20% at most, and they still decrease as  $n$  increases, so the estimators are also consistent in this scenario.

In addition to consistency of the parameter estimators, it is also important to study the consistency of the component score estimators  $\hat{u}_{ik}$ s, since they are used

$n$	Param	rate 10			rate 30		
		bias	std	rmse	bias	std	rmse
50	$\mu$	.46	1.56	1.63	.46	1.21	1.29
	$\phi_1$	.45	.75	.88	.45	.71	.84
	$\phi_2$	.60	.83	1.03	.50	.77	.92
100	$\mu$	.45	1.12	1.21	.43	.90	.99
	$\phi_1$	.29	.63	.69	.21	.52	.56
	$\phi_2$	.48	.77	.91	.25	.57	.62
200	$\mu$	.56	.83	1.00	.49	.64	.81
	$\phi_1$	.21	.50	.54	.12	.36	.38
	$\phi_2$	.36	.65	.74	.17	.40	.44
400	$\mu$	.47	.56	.74	.52	.45	.69
	$\phi_1$	.20	.34	.39	.09	.25	.26
	$\phi_2$	.35	.54	.64	.18	.30	.35

Table 3: Simulation Results. Bias, standard deviation and root mean squared errors of parameter estimators for autoregressive component scores (Scenario 3). Quantities for  $\mu$  were multiplied by 10.

for inference (like clustering, in this paper). The distance between the  $\hat{u}_{ik}$ s and the true  $u_{ik}$ s cannot be measured directly, due to the sign indeterminacy, so we use the estimation error of the variations  $v_{ik}(t) = u_{ik}\phi_k(t)$  instead, which are sign-invariant. We define the expected average error  $eae = E(\sum_{i=1}^n \|\hat{v}_{ik} - v_{ik}\|/n)$ , where  $\|\cdot\|$  is the  $L^2[0, 1]$  norm. We also measure the association between the  $\hat{u}_{ik}$ s and the  $u_{ik}$ s by the expected absolute correlation,  $eac = E\{|\text{corr}(\hat{u}_{ik}, u_{ik})|\}$ , which is also sign-invariant.

Table 4 shows the results. Again we see consistency, an improvement in estimation as  $n$  and/or the baseline rate increase, but the latter has a bigger impact on the performance of the  $\hat{u}_{ik}$ s. This was expected, since the  $\hat{u}_{ik}$ s can only use the  $m_i$  observations available for replication  $i$ , whereas  $\hat{\mu}(t)$  and the  $\hat{\phi}_k(t)$ s pool data across replications. Regarding the three distributional scenarios, we see that there is almost no difference between the independent identically distributed case and the model with quadratic trend; the autoregressive model does show somewhat higher errors and lower correlations than the other two, especially for  $n = 50$ , but the difference tends to vanish as  $n$  increases. So we can say that the component score estimators are consistent under the three scenarios.

When the component scores reveal a trend or autocorrelation, model (1) can be modified to accommodate such relationships, and re-estimated. Other covariates on which the  $u_{ikj}$ s may depend can also be incorporated. However, a detailed elaboration of these possibilities goes beyond the scope of this paper.

## 5 Application: Chicago’s Divvy bike sharing system

As mentioned in the Introduction, we analyze in this section the checkout times of bike trips that took place between April 1 and November 31 of 2016 in Chicago’s Divvy system. First, we fitted model (1) for the 458 bike stations that were active during this period. As spline basis for the functional parameters we used cubic B-splines with ten equally spaced knots in  $(0, 24)$ . We fitted models with  $p = 6$  components, which were sufficient to capture the most important modes of variation in the data and can be estimated without inconvenient for most stations; only for station 386, the station with the lowest annual count (29 for the whole year), the model could not be fitted due to insufficient data.



$n$	Score	Independent				With trend				Autoregressive			
		rate 10		rate 30		rate 10		rate 30		rate 10		rate 30	
		eae	eac	eae	eac	eae	eac	eae	eac	eae	eac	eae	eac
50	pc 1	.28	.51	.18	.73	.28	.52	.17	.75	.29	.47	.20	.63
	pc 2	.26	.33	.17	.59	.26	.34	.17	.60	.26	.34	.19	.51
100	pc 1	.26	.54	.15	.77	.26	.54	.16	.77	.26	.52	.17	.73
	pc 2	.24	.40	.15	.67	.24	.39	.15	.65	.25	.37	.16	.63
200	pc 1	.24	.56	.14	.79	.24	.57	.14	.79	.24	.54	.15	.77
	pc 2	.23	.43	.14	.69	.23	.42	.14	.68	.23	.42	.14	.68
400	pc 1	.23	.58	.13	.80	.23	.57	.13	.80	.23	.57	.14	.79
	pc 2	.22	.43	.13	.70	.22	.43	.13	.68	.22	.43	.13	.69

Table 4: Simulation Results. Expected average error (eae) and expected absolute correlation (eac) of component score estimators under the three scenarios of Tables 1–3.

It is clearly infeasible to visually inspect the results for all stations, but as an illustration we will analyze in more detail the results for station 166, the station with median annual count. The estimated baseline intensity function  $\hat{\lambda}_{0,166}$  is shown in Figure 1. We see that  $\hat{\lambda}_{0,166}$  has three peaks: the first and largest one occurs at 7:30am, the second and smallest one at 1pm, and the third one at 5:30pm. The integral of  $\hat{\lambda}_{0,166}$  over  $[0, 24]$  is 17.66, very close to the mean daily count of 17.64, as expected.

To interpret the components  $\hat{\psi}_{k,j}$  it is instructive to plot the baseline function  $\hat{\lambda}_{0j}$  alongside  $\hat{\lambda}_{0j}^+ = \hat{\lambda}_{0j}\hat{\psi}_{kj}^c$  and  $\hat{\lambda}_{0j}^- = \hat{\lambda}_{0j}\hat{\psi}_{kj}^{-c}$ , for some positive constant  $c$  chosen for convenient visualization (here we take it as twice the standard deviation of the corresponding  $\hat{u}_{ikj}$ s). For the first component, this is shown in Figure 2(a). In Figure 2(b) we plotted the corresponding component scores  $\hat{u}_{i,1,166}$  as a time series on the index  $i$ . Figure 2(a) shows that a negative score corresponds to a sharpening of the morning peak and a positive score corresponds to a flattening of this peak. This corresponds to weekday versus weekend patterns of demand, respectively, as corroborated by Figure 2(b), which shows a steady weekly periodicity (the autocorrelation at lag 7 is .68), with peaks occurring almost always on Sundays and troughs mostly on Thursdays or Wednesdays. In Figure 3 we show the 244 estimated daily




Figure 1: Baseline intensity function of daily bike demand for Divvy station 166, located at the intersection of Wrightwood and Ashland avenues.

intensity functions, separating weekdays (Figure 3(a)) from weekends (Figure 3(b)); the absence of the morning peaks in Figure 3(b) is clear.

The second component (Figure 4(a)) explains overall count variation. Overall bike usage is strongly seasonal, as shown in Figure 4(b), with demand increasing from early Spring to Summer (the maximum occurs in June) and decreasing thereafter. The rest of the components explain finer-detailed aspects of bike demand.

After fitting model (2) for all bike stations, we computed the canonical correlations (8) for all pairs. The largest one turned out to be .98 and the smallest one .17. The largest correlation corresponds to bike stations 75 and 91, located at the main entrances of Union and Ogilvy train stations, respectively. Although these bike stations are relatively close to each other (556 m, four city blocks), they are not the closest. For example, station 73 is closer to station 75 (277 m, two city blocks) but their correlation is lower (.90), and station 169 is 452 m (three city blocks) away from station 75, closer than Ogilvy is but in the opposite direction and without any train stations nearby, so their correlation is only .72. It is clear, then, that correlations are not functions of distance alone but also of type of usage; the spatial correlations are not isotropic.

Then it is instructive to apply clustering methods to the correlations and try to associate the clusters with different patterns of usage. The clustering procedure of Section 3 gives the dendrogram shown in Figure 5. The vertical axis of the




Figure 2: First multiplicative component of daily bike demand for Divvy station 166. (a) Baseline (solid line) and baseline multiplied by a positive (dotted line) and negative (dashed line) exponent of the component. (b) Daily component scores as a time series.




Figure 3: Daily intensity functions of bike demand for Divvy station 166, (a) weekdays, (b) weekends.




Figure 4: Second multiplicative component of daily bike demand for Divvy station 166. (a) Baseline (solid line) and baseline multiplied by a positive (dotted line) and negative (dashed line) exponent of the component. (b) Daily component scores as a time series.




Figure 5: Dendrogram of complete-linkage clustering of bike stations in the Divvy system.

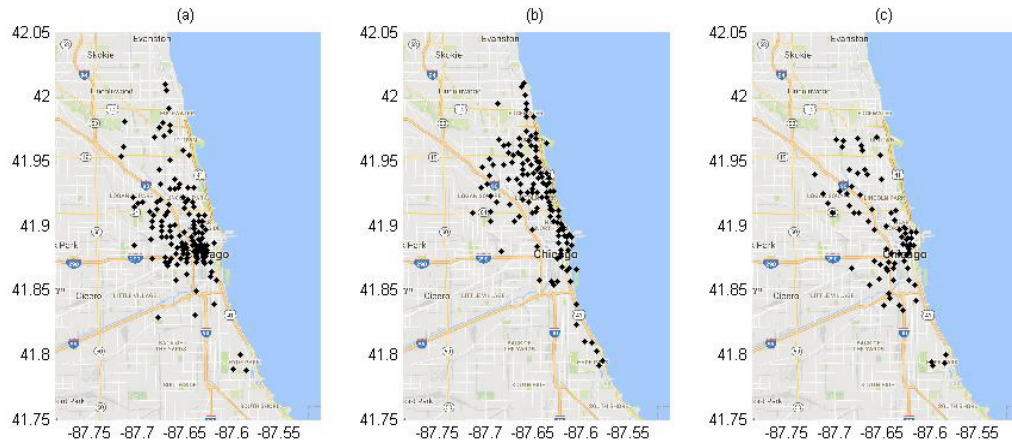


Figure 6: Clusters of bike stations in the Divvy system. (a) Largest cluster, 136 stations; (b) second largest cluster, 127 stations, (c) third largest cluster, 77 stations.

dendrogram indicates the distance of the objects being connected. Three big clusters are discernible in Figure 5, with a maximum distance of about .70, so the correlations of bike stations within the clusters are at least .30. These clusters include 136, 127 and 77 bike stations respectively, so they account for 340 of the 458 bike stations in the system (most of the others had non-significant correlations that were trimmed as explained in Section 3).

The locations of stations in each cluster are shown in Figure 6. We also show the baseline density functions for each station,  $\tilde{\lambda}_{0j} = \lambda_{0j} / \int \lambda_{0j}$ , in Figure 7. Although we are clustering by correlation and not by distance between baselines, the baseline densities do help interpret the type of usage given to the stations in each cluster. We see in Figure 7(a) that most densities in this cluster show a typical weekday-usage pattern (compare with Figure 3(a)); Figure 6(a) shows that most stations in downtown Chicago, and specifically in “the Loop”, belong to this cluster, so Cluster 1 consists of bike stations that are mostly used for commute. The densities in Figure 7(b) show a weekend-usage pattern (compare with Figure 3(b)), and Figure 6(b) shows that most stations along the lake shore belong to this cluster, so Cluster 2 consists of stations that are mostly used for leisure trips. The third cluster is somewhere in between.




Figure 7: Baseline density functions for the three clusters of bike stations in Figure 6.

## Acknowledgement

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

## 6 Appendix: Computation of functional canonical correlations

Let  $U(t)$  and  $V(t)$  be two stochastic processes admitting finite expansions  $U(t) = \mu_U(t) + \sum_{k=1}^p u_k \phi_k(t)$  and  $V(t) = \mu_V(t) + \sum_{k=1}^q v_k \psi_k(t)$ , where the  $\phi_k$ s and the  $\psi_k$ s are orthonormal. The canonical correlation coefficient is  $\rho = \max_{\alpha, \beta} \text{corr}(\langle \alpha, U \rangle, \langle \beta, V \rangle)$ , where  $\alpha$  and  $\beta$  are arbitrary square-integrable functions. Any  $\alpha$  and  $\beta$  can be decomposed as  $\alpha(t) = \sum_{k=1}^p a_k \phi_k(t) + \eta(t)$ , with  $\eta$  orthogonal to the  $\phi_k$ s, and  $\beta(t) = \sum_{k=1}^q b_k \psi_k(t) + \xi(t)$ , with  $\xi$  orthogonal to the  $\psi_k$ s. Then  $\langle \alpha, U - \mu_U \rangle = \sum_{k=1}^p a_k u_k$  and  $\langle \beta, V - \mu_V \rangle = \sum_{k=1}^q b_k v_k$ . Let  $\mathbf{a} = (a_1, \dots, a_p)$ ,  $\mathbf{U} = (U_1, \dots, U_p)$ ,  $\mathbf{b} = (b_1, \dots, b_q)$  and  $\mathbf{V} = (V_1, \dots, V_q)$ ; then  $\langle \alpha, U - \mu_U \rangle = \mathbf{a}^T \mathbf{U}$  and  $\langle \beta, V - \mu_V \rangle = \mathbf{b}^T \mathbf{V}$ . Since  $\text{corr}(\langle \alpha, U \rangle, \langle \beta, V \rangle) = \text{corr}(\langle \alpha, U - \mu_U \rangle, \langle \beta, V - \mu_V \rangle)$ , then  $\rho = \max_{\mathbf{a}, \mathbf{b}} \text{corr}(\mathbf{a}^T \mathbf{U}, \mathbf{b}^T \mathbf{V})$ , which is the standard multivariate canonical correlation coefficient for  $\mathbf{U}$  and  $\mathbf{V}$ . In particular, if  $U(t) = \log \Lambda_j(t)$  and  $V(t) = \log \Lambda_{j'}(t)$ ,

where  $\Lambda_j(t)$  is the random function that generates the  $\lambda_{ij}(t)$ s in model (1), we have that  $\mathbf{U}$  are the component scores for site  $j$  and  $\mathbf{V}$  are the component scores for site  $j'$ .

## References

- Baddeley, A. (2007). Spatial point processes and their applications. In *Stochastic Geometry*, Lecture Notes in Mathematics 1892, pp. 1–75. Springer, New York.
- Baddeley, B. A., Moyeed, R. A., Howard, C. V., and Boyde, A. (1993). Analysis of three-dimensional point pattern with replication. *Applied Statistics* **42** 641–668.
- Baddeley, A., Rubak, E., and Turner, R. (2015). *Spatial Point Patterns: Methodology and Applications with R*. CRC Press, Boca Raton, USA.
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B* **57** 289–300.
- Borgnat, P., Robardet, C., Rouquier, J., Abry, P., Flandrin, P., and Fleury, E. (2011). Shared bicycles in a city: A signal processing and data analysis perspective. *Advances in Complex Systems* **14** 1–24.
- Bouzas, P.R., and Ruiz-Fuentes, N. (2015). A review on functional data analysis for Cox processes. *Boletín de Estadística e Investigación Operativa* **31** 215–230.
- Davies, D., and Bouldin, D. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **1** 224–227.
- De Boor, C. (1978). *A Practical Guide to Splines*. Springer, New York.
- Delicado, P., Giraldo, R., Comas, C., and Mateu, J. (2010). Statistics for spatial functional data: Some recent contributions. *Environmetrics* **21** 224–239.
- Diggle, P. (2013). *Statistical Analysis of Spatial and Spatio-Temporal Point Patterns, Third Edition*. CRC Press, Boca Raton, USA.

- 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., Mateu, J. and Clough, H. E. (2000). A comparison between parametric and non-parametric approaches to the analysis of replicated spatial point pattern. *Advances in Applied Probability* **32** 331–343.
- Dunn, J. (1974). Well separated clusters and optimal fuzzy partitions. *Journal of Cybernetics* **4** 95–104.
- Eilers, P.H.C., and Marx, B.D. (1996). Flexible smoothing with B-splines and penalties (with discussion). *Statistical Science* **11** 89–121.
- Gervini, D. (2016). Independent component models for replicated point processes. *Spatial Statistics* **18** 474–488.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning. Data Mining, Inference, and Prediction. Second Edition.* Springer, New York.
- Horváth, L., and Kokoszka, P. (2012). *Inference for Functional Data with Applications.* Springer, New York.
- Izenman, A.J. (2008). *Modern Multivariate Statistical Techniques. Regression, Classification and Manifold Learning.* Springer, New York.
- Jolliffe, I.T. (2002). *Principal Component Analysis. Second Edition.* Springer, New York.
- 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.
- Mateu, J. (2001). Parametric procedures in the analysis of replicated pairwise interaction point patterns. *Biometrical Journal* **43** 375–394.



- Menafoglio, A., and Secchi, P. (2017). Statistical analysis of complex and spatially dependent data: A review of object oriented spatial statistics. *European Journal of Operational Research* **258** 401–410.
- Møller, J., and Waagepetersen, R.P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman and Hall/CRC, Boca Raton.
- Nair, R., and Miller-Hooks, E. (2011). Fleet management for vehicle sharing operations. *Transportation Science* **45** 524–540.
- Nair, R., Miller-Hooks, E., Hampshire, R.C., and Bušić, A. (2013). Large-scale vehicle sharing systems: Analysis of Vélib’. *International Journal of Sustainable Transportation* **7** 85–106.
- Ramsay, J. O., and Silverman, B. W. (2005). *Functional Data Analysis. Second Edition*. Springer, New York.
- Romano, E., Balzanella, A., and Verde, R. (2010). Clustering spatio-functional data: A model based approach. In *Classification as a tool for research. Studies in Classification, Data Analysis, and Knowledge Organization*, pp. 167–175. Springer, Berlin, Heidelberg.
- Ruppert, D. (2002). Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics* **11** 735–757.
- Seber, G.A.F. (2004). *Multivariate Observations*. Wiley, New York.
- Secchi, P. , Vantini, S. , and Vitelli, V. (2013). Bagging Voronoi classifiers for clustering spatial functional data. *International Journal of Applied Earth Observation and Geoinformation* **22** 53–64.
- Shaheen, S., Guzman, S., and Zhang, H. (2010). Bike sharing in Europe, the Americas and Asia: Past, present and future. *Transportation Research Record: Journal of the Transportation Research Board* **2143** 159–167.
- Shirota, S., and Gelfand, A.E. (2017). Space and circular time log Gaussian Cox processes with application to crime event data. *The Annals of Applied Statistics* **11** 481–503.

- Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall/CRC, Boca Raton.
- Streit, R.L. (2010). *Poisson Point Processes: Imaging, Tracking, and Sensing*. Springer, New York.
- Vogel, P., Greiser, T., and Mattfeld, D.C. (2011). Understanding bike-sharing systems using data mining: exploring activity patterns. *Procedia Social and Behavioral Sciences* **20** 514–523.
- Wu, S., Müller, H.-G., and Zhang, Z. (2013). Functional data analysis for point processes with rare events. *Statistica Sinica* **23** 1–23.