Statistica Sinica **18**(2008), 1375-1393

NONPARAMETRIC BOOTSTRAP FOR K-FUNCTIONS ARISING FROM MIXED-EFFECTS MODELS WITH APPLICATIONS IN NEUROPATHOLOGY

Sabine Landau and Ian P. Everall

King's College London and University of California, San Diego

Abstract: Neuropathological studies frequently determine the positions of cells on multiple brain tissue sections taken from multiple individuals. Interest arises in group comparisons of the spatial dependencies between cells, in particular the spatial dependencies of a single cell type (clustering or regularity as measured by the univariate K-function), or the spatial interaction of two different cell types (attraction or repulsion as measured by the bivariate K-function). While the nonparametric statistical analysis of spatial dependencies in the one-way design is fairly well-established, investigations often employ more complex designs. In this paper we develop a residual bootstrapping approach for K-functions arising from a general repeated measures design by assuming an underlying linear mixed-effects model. We illustrate our methodology by re-analysing the spatial interaction between neurons and astrocytes (brain cells that are functionally related to neurons) in a study of HIV associated dementia.

Key words and phrases: Bivariate point process, bootstrap, K-function, mixed-effects model, neuropathology, nonstationarity, replicated spatial point pattern.

1. Introduction

Impaired functioning of brain cells has been hypothesized in brain disorders such as depression, schizophrenia, or dementia. Such abnormalities might manifest themselves as alterations of the cell densities and/or the spatial dependencies between brain cells. For example, the degree of spatial aggregation (clustering or regularity) of a brain cell type might be altered if non-random loss of this cell type were to occur in a degenerative disease. Similarly the degree of spatial interaction (attraction or repulsion) between two cell types might be affected. For example, brain cells called astrocytes normally support neurons and are typically found in their vicinity. So an impairment in the functional relationship between these two types of cells might manifest itself by an alteration of the degree of their spatial interaction. (We examine such a research question later.) To study brain cell function, neuropathological investigations therefore often determine the positions of cells on brain tissue sections from different groups of subjects

(e.g., patients and controls), and compare cell densities as well as spatial cell arrangement.

The resulting cell positions datasets can be considered as realisations of spatial point processes. Brain tissue sections are typically cut coronally, that is, reaching from the surface of the brain (the pia) to the grey-white matter junction. Coronal sections are thin relative to the cell sizes. In addition, coronal sectioning is known to produce inhomogeneous neuronal cell patterns since neurons are organised in the brain in (typically six) layers of differing densities that run parallel to the pial surface. Thus under standard cutting practice, brain cell positions have to be further considered as arising from planar inhomogeneous point processes.

For homogeneous replicated cell patterns, effects of experimental factors on densities can be analysed routinely in general purpose statistical packages using appropriate methods for count outcomes; and recently methods have been suggested that utilise replication to model spatial intensity for inhomogeneous cell patterns (Wager, Coull and Lange (2004)). In contrast, approaches for analysing spatial cell dependencies from replicated patterns are more restrictive. For some classes of homogeneous point process models, fixed or mixed-effects model formulations have been used to estimate process parameters from replicated point patterns (for a fixed-effects model for pairwise interaction point processes, see Diggle, Mateu and Clough (2000) and Mateu (2001); for a random-effects formulation for Gibbs point processes, see Bell and Grunwald (2004)). Alternatively, for the simple one-way design, nonparametric methods that make use of the replication and do not have to rely on stationarity have been suggested (see Diggle, Lange and Beneš (1991) for a residual bootstrap procedure to test overall group effects on the spatial aggregation of a univariate point process, and Landau, Rabe-Hesketh and Everall (2004) for extensions to spatial interaction of a bivariate process).

While nonparametric analysis of spatial cell dependencies has become more widespread in pathological investigations (e.g., Asare, Dunn, Glass, McArthur, Luthert, Lantos, and Everall (1996), Chana, Landau, Beasley, Everall, and Cotter (2003) and Schladitz, Särkkä, Pavenstädt, Haferkamp, and Mattfeldt (2003)), studies often employ more complex designs than can so far be accommodated by this approach. Several cell patterns per individual might be subject to analysis, and the effects of a number of continuous and/or categorical between- and/or within-subject variables might be under investigation. Repeated cell patterns per individual can arise due to multiple tissue sections, for example, from several brain regions. In addition, the information from a single multi-type cell pattern can be captured by several univariate or bivariate patterns when the research question is focused on comparing spatial arrangements of certain cell types. (We consider such a research question later.)

This paper aims to further develop the nonparametric approach which allows for inhomogeneous replicated point patterns to cover more general repeated measures designs. The existing approach operates first by summarizing the spatial dependencies of the cell patterns by functions of distance, typically by empirical K-functions, then modelling the relationship between expected K-function values and group by a one-way ANOVA type model for each distance, and finally applying a residual bootstrap of entire K-function curves to derive inferences. We propose to extend the approach by assuming an appropriate linear mixed-effects model for empirical K-function values at a given distance. This provides a framework that can account for correlation between cell patterns from the same individual, and can accommodate any type of between-subject or within-subject covariate.

We start by introducing a replicated spatial cell pattern dataset from a study of HIV associated dementia for illustrative purposes. Section 3 briefly reviews K-functions as measures of spatial aggregation and spatial interaction. Section 4 then specifies an appropriate mixed-effects model for repeated K-functions, and Section 5 provides details of a residual bootstrap procedure for analysing such functions. Section 6 concludes with a discussion.

2. Illustrative Example: HIV Associated Dementia Study

We re-analyse a replicated spatial point pattern dataset from a study of HIV associated dementia (HAD) to illustrate our proposed methodology; in particular, to show that the extended modelling framework enables inferences not previously possible under the simple one-way design (Landau, Rabe-Hesketh and Everall (2004)). All datasets and programs used are available from the Statistica Sinica website (http://www3.stat.sinica.edu.tw/statistica).

Some patients with HIV disease develop HAD. A hypothesized mechanism is that HIV infection in the brain may result in dysregulation of astrocytes which are normally crucial for supporting neurons and preventing excitotoxic damage to neurons. This dysregulation may lead to excitotoxic damage and neuronal death, particularly for large or pyramidal neurons which are most vulnerable to such damage. One way of assessing the experimental hypothesis of such a mechanism is to look for a specific alteration of the spatial interaction between large pyramidal neurons and astrocytes in HIV patients who have developed dementia relative to those who have not. (For more background about the HAD study see Roberts (2000).)

Brain tissue of 29 male patients who died of AIDS between 1985 and 1991 was available from the Johns Hopkins Hospital, Baltimore, USA (for demographic and clinical information, see Everall, Glass, McArthur, Spargo and Lantos (1994)). Prior to death, the patients had been assessed clinically for the presence and severity of HAD as rated on the Memorial Sloan-Kettering Scale (MSK, Price

and Brew (1988)). This was used to classify the patients into three dementia severity groups: non-demented group (7 patients, MSK stages 0 and 0.5), moderately demented group (13 patients, MSK stages 1 and 2), and severely demented group (9 patients, MSK stages 3 and 4). From each individual a brain block was supplied from the superior frontal gyrus, at the level of the genu or the corpus callosum. From each block a single thin coronal tissue section (5 microns thickness) was prepared and stained for neurons and astrocytes. Since the aim of the study was to examine the potentially different changes in HAD in the spatial relationships of astrocytes with different types of neurons, the latter were further classified into three subpopulations labelled "interneurons" (less than 4.9 microns in radius), "small pyramidal neurons" (radius from 4.9 to 6.9 microns) and "large pyramidal neurons" (radius more than 6.9 microns) on the basis of their size. (For more details of the data and discriminant analysis used to develop this rule see Roberts (2000)).

Each of the 29 patients therefore provided a single multivariate cell pattern containing cells of four different types. We captured the relevant information from this by considering three bivariate cell patterns per subject, each representing the spatial relationship between astrocytes and one of the neuronal cell types. This was reasonable in our context since we were focused on the spatial relationships between astrocytes and neuronal subtypes. However, it meant that we generated three potentially correlated repeated bivariate cell patterns. Some of the bivariate cell patterns are shown in Figure 1. The arrangement of neurons in layers is apparent, especially for interneurons. Also noticeable is the relative scarcity of large pyramidal neurons compared to other types of neurons. Astrocytes tended to cluster in the layer adjacent to the pial surface (layer I).

3. K-functions as Measures of Spatial Proximities between Cells

We briefly review empirical K-functions as summary measures for spatial pattern analysis.

3.1. The empirical K-function

For a stationary and isotropic process the theoretical univariate K-function, $K(t) \equiv (1/\lambda)E[\text{number of cells within distance } t \text{ of an arbitrary cell}]$ where λ is the intensity of the process, reflects spatial relationships between the cells. Complete spatial randomness (CSR) asserts that given n events of the process in planar region A, the spatial positions are an independent random sample from a uniform distribution over A. Under CSR $K(t) = \pi t^2$, and the value πt^2 further defines a benchmark for categorising the spatial aggregation of a univariate process: $K(t) > \pi t^2$ indicating clustering, and $K(t) < \pi t^2$ indicating regularity at distance t. Thus the ratio $I(t) = K(t)/(\pi t^2)$ provides a scale invariant measure of the extent to which the univariate K-function exceeds its expectation under CSR.

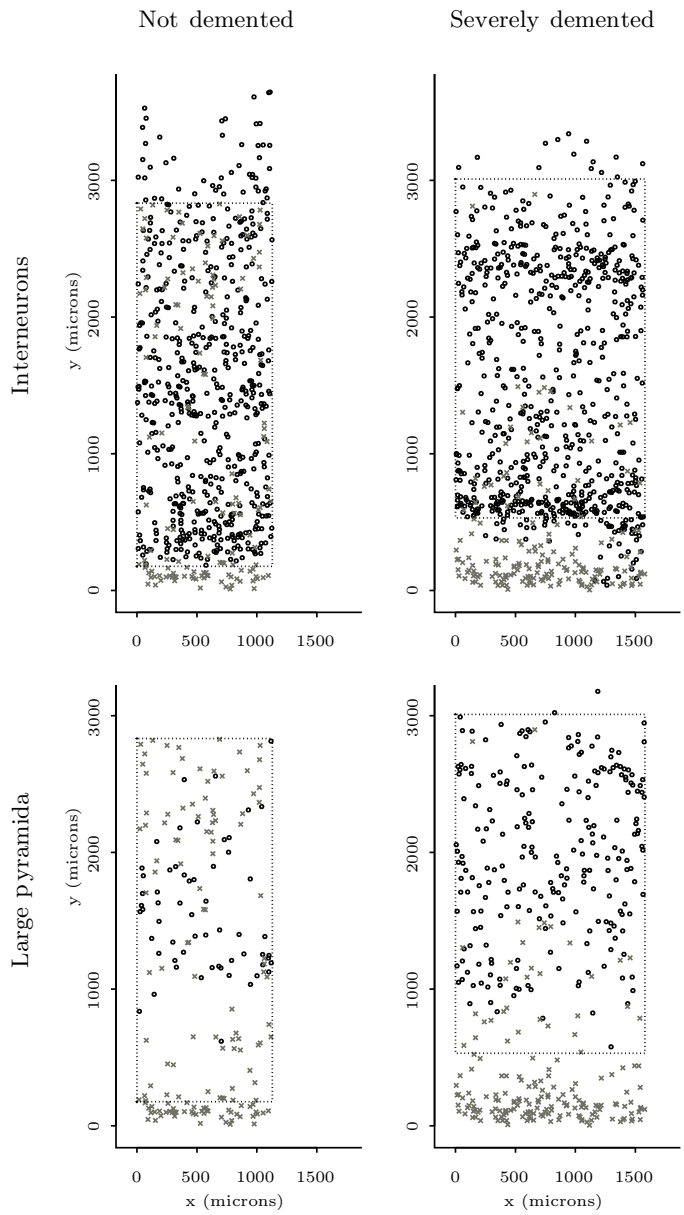

Figure 1. Spatial positions of neurons (black open symbols) and astrocytes (grey cross symbols) in grey matter for an individual from the non-demented and severely demented groups (columns) by two neuronal subpopulations (rows). The plots show cell positions and rectangular boundary boxes indicating the search areas employed. The boundary boxes reach from the top of layer I (the layer adjacent to the pial surface) to the grey-white matter junction.

Let u_{kj} denote the distance between the kth and jth cell position of the process in region A, $|A|$ the area of the region and define $I_t(u)$ to be 1 if $u \leq t$ and 0 otherwise. Then for a stationary and isotropic process, a commonly employed asymptotically unbiased estimator for $K(t)$ which takes account of edge effects (Ripley (1981)) is given $\hat{K}(t) \equiv (|A|/n^2)\{\sum_{k=1}^{n}\sum_{j=1;j\neq k}^{n} w_k^{-1} I_t(u_{kj})\}$, where n is the number of events of the cell process and w_k is the proportion of the circumference of the circle, with centre at the kth cell position and radius t, which lies within A.

For two stationary and isotropic processes the theoretical bivariate K-function $K_{12}(t) \equiv (1/\lambda_1)E$[number of type 1 cells within distance t of arbitrary type 2 cell] where λ_1 is the intensity of the type 1 cell process, reflects spatial interaction between the two different types of cells. The hypothesis of spatial independence between two stationary and isotropic univariate processes asserts that the positions taken by one cell type are independent of those taken by the other. Under spatial independence $K_{12}(t) = \pi t^2$, and the value πt^2 further defines a benchmark for categorising the spatial interaction between the two cell processes: $K_{12}(t) > \pi t^2$ indicating attraction, and $K_{12}(t) < \pi t^2$ indicating repulsion at distance t. Thus the ratio $I_{12}(t) = K_{12}(t)/\pi t^2$ provides a scale invariant measure of the extent to which the bivariate K-function exceeds its expectation under spatial independence.

Under stationarity and isotropy, two asymptotically unbiased estimators are given by

$$\hat{K}_{12}^{(1)}(t) \equiv \frac{|A|}{n_1 n_2}\left\{\sum_{k=1}^{n_1}\sum_{j=1}^{n_2} w_k^{-1} I_t(u_{kj})\right\} \text{ and } \hat{K}_{12}^{(2)}(t) \equiv \frac{|A|}{n_1 n_2}\left\{\sum_{j=1}^{n_2}\sum_{k=1}^{n_1} w_j^{-1} I_t(u_{kj})\right\},$$

where u_{kj} now denotes the distance between the kth cell position of process 1 and the jth cell position of process 2, and n_1 and n_2 are the number of cells of type 1 and 2, respectively. These may be optimally combined into the empirical bivariate K-function $\hat{K}_{12}(t) \equiv \{n_2/(n_1+n_2)\}\hat{K}_{12}^{(1)}(t) + \{n_1/(n_1+n_2)\}\hat{K}_{12}^{(2)}(t)$ (Lotwick and Silverman (1982)).

Under nonstationarity theoretical K-functions are not defined, and empirical K-functions can no longer be motivated as estimators with theoretical properties. However, as noted previously (Diggle, Mateu and Clough (2000)), they retain a tangible scientific interpretation as nonparametric summary measures of the degree of spatial aggregation (univariate cell pattern) or spatial interaction (bivariate pattern). Thus in the context of inherently nonstationary processes in neuropathological studies empirical K-functions remain outcomes of interest, but their interpretation needs to take account of the fact that they may partially reflect the nonstationarity mechanism (e.g., neuronal layering) itself.

3.2. Constructing empirical K-functions for the HAD study

For each bivariate cell pattern the spatial interaction between astrocytes and neuronal subtype was summarized by their empirical bivariate K-function. For this purpose rectangular grey matter search areas were chosen so that they reached as far as possible toward the potentially curved grey-white matter junction while keeping widths relatively stable across subjects. Since astrocytes tended to be prominent in layer I near the pial surface, and appeared homogeneously distributed across the remaining layers (see Figure 1), layer I was excluded from the search area in an attempt to reduce nonstationarity and ease interpretation. This led to one cell pattern from the moderate group being excluded, since all astrocytes were located in layer I. Bivariate K-functions were evaluated in steps of four microns up to a maximum search radius of $t_{\max} = 300$ microns (about a quarter of the shorter side of the boundary boxes). The outcome data for formal analysis therefore consisted of 84 empirical bivariate K-functions corresponding to 28 patients and three neuronal types.

We previously carried out simulations to assess the effects of neuronal layering on empirical bivariate K-functions (see Landau, Rabe-Hesketh and Everall (2004)). This showed that neuronal layering of a type similar to that observed in the HAD study had only a minor impact. (The interaction index was increased by 1.2% and 2.6% toward more attraction under moderate and severe layering.) Here we interpret substantial departures from the stationary independence model (interaction indices < 1 or > 1.03) as apparent cell interaction and not simply layering effects.

4. Mixed-effects Model for Repeated Empirical K-functions

4.1. Basic two-level mixed-effects model

We assume that the K-function values at a given distance $t \in \{t_1, \ldots, t_{\max}\}$ follow a two-level model. Specifically for subject $i \in \{1, \ldots, N\}$, let $\mathbf{K}_i(t)$ denote the m_i-dimensional vector of repeated empirical K-function values at distance t, $\mathbf{X}_i$ a known fixed-effects $m_i \times p$-design matrix, $\boldsymbol{\beta}(t)$ a p-dimensional vector of fixed effects, $\mathbf{Z}_i$ a known random-effects $m_i \times q$-design matrix whose columns are typically a subset of the columns of $\mathbf{X}_i$, $\mathbf{b}_i(t)$ a q-dimensional vector of subject random effects with expectation $\mathbf{0}$ and covariance matrix $\boldsymbol{\Psi}(t)$, and $\boldsymbol{\varepsilon}_i(t)$ a m_i-dimensional vector with expectation $\mathbf{0}$ and covariance matrix $\sigma^2(t)\boldsymbol{\Lambda}_i(t)$. Then a mixed-effects model for repeated empirical K-functions is given by

$$\mathbf{K}_i(t) = \mathbf{X}_i\boldsymbol{\beta}(t) + \mathbf{Z}_i\mathbf{b}_i(t) + \boldsymbol{\varepsilon}_i(t), \quad i = 1, \ldots, N. \tag{4.1}$$

The $\mathbf{b}_i(t)$ and $\boldsymbol{\varepsilon}_i(t)$ are assumed to be independent for different subjects and independent of each other. The level 1 variance-covariance matrices can be written as $\sigma^2(t)\boldsymbol{\Lambda}_i(t) = \sigma^2(t)\mathbf{V}_i(t)\mathbf{C}_i(t)\mathbf{V}_i(t)$, where $\mathbf{V}_i(t)$ is a diagonal matrix with jth diagonal element $v_{ij}(t) \equiv [\mathrm{Var}\,\{\boldsymbol{\varepsilon}_{ij}(t)\}/\sigma^2(t)]^{0.5}$, and $\mathbf{C}_i(t)$ is defined by $[\mathbf{C}_i(t)]_{jk} \equiv \mathrm{Corr}\{\boldsymbol{\varepsilon}_{ij}(t), \boldsymbol{\varepsilon}_{ik}(t)\}$. The matrix of within-subject correlations, $\mathbf{C}_i(t)$, is modelled by an (usually small) set of parameters that can vary with distance t.

4.2. Model for the sampling variances

The sampling variances of empirical K-functions depend on parameters of the underlying spatial point process. It has been shown theoretically for stationary and isotropic processes (Stoyan, Kendall and Mecke (1987)), and by simulation for layered non-stationary processes (Landau, Rabe-Hesketh and Everall (2004)), that the variance of the empirical K-function for subject i and univariate cell pattern j is approximately proportional to $1/(\lambda_{ij}n_{ij})$, where n_{ij} is the number of cells and λ_{ij} the expected density for that pattern. Similarly simulations of two independent stationary and isotropic processes, or involving one layered non-stationary process (Landau, Rabe-Hesketh and Everall (2004)) show that the variance of the empirical bivariate K-function is approximately proportional to $1/(\lambda_{1,ij}^{0.5}\lambda_{2,ij}^{0.5}n_{1,ij}^{0.5}n_{2,ij}^{0.5})$, where $n_{1,ij}$ and $n_{2,ij}$ are the numbers of cells of the two types, and $\lambda_{1,ij}$ and $\lambda_{2,ij}$ the respective expected densities. Therefore to account for process parameters we assume that the variance function at level 1 can be characterized by variance covariate values that do not vary with distance, i.e.,

$$\mathrm{Var}\,\{\boldsymbol{\varepsilon}_{ij}(t)\} = \sigma^2(t)v_{ij}^2 \tag{4.2}$$

with $v_{ij}^2 \equiv 1/(\hat{\lambda}_{ij}n_{ij})$ for modelling empirical univariate K-functions, or $v_{ij}^2 \equiv 1/(\hat{\lambda}_{1,ij}^{0.5}\hat{\lambda}_{2,ij}^{0.5}n_{1,ij}^{0.5}n_{2,ij}^{0.5})$ for empirical bivariate K-functions.

For simple one-way designs of empirical univariate K-functions, the proportionality factors have been approximated by assuming constant search area sizes (Baddeley, Moyeed, Howard and Boyde (1993), resulting in the choice $v_{ij}^2 = 1/n_{ij}^2$) or by assuming that the expected cell densities vary little between subjects (Diggle, Lange and Beneš (1991), with choice $v_{ij}^2 = 1/n_{ij}$). The appeal of the latter approach is that it avoids having to estimate further process parameters from the observed cell patterns. However, for most repeated measures applications, the assumption of constant expected densities across subjects and repeated measures is bound to be unrealistic, e.g., when the repeated measures represent coronal layers or when the repeated measures correspond to different types of cells. In addition, neuropathological studies often also hypothesize differences in expected densities across repeated measures and groups. We therefore suggest estimating the expected densities.

4.3. Repeated measures model for the HAD bivariate K-functions

We assumed that the 84 empirical bivariate K-functions followed a mixed model at each distance $t \in \{4, 8, \ldots, 300\}$. The fixed part $\mathbf{X}_i\boldsymbol{\beta}(t)$ of our mixed model was parameterised so that it contained linear and quadratic effects of (ordinal) dementia groups (fixed effects $\beta_2(t)$ and $\beta_3(t)$), differences between small or large pyramidal neurons and interneurons ($\beta_4(t)$ and $\beta_5(t)$), and respective interaction terms ($\beta_6(t)$ to $\beta_9(t)$). Sampling variances for processes similar to the one at hand had previously been shown to follow the stipulated proportionality relationship, and variance covariate values required to fully specify (4.2) were estimated from a saturated model for expected densities.

The neuronal types defined the three repeated measures per subject. To identify an appropriate model for the correlations between the repeated K-function values, simple fixed-effects models based on independent repeated measures were fitted, residuals constructed, scaled to constant variances using the variance covariate values, and correlations estimated. This suggested the need for a more complex correlation structure than could be accounted for by subject random intercepts. (The three correlations varied considerably; reaching coefficients of small positive size, small negative size and large negative size for the largest distances). We therefore opted for three (distance-varying) correlation parameters:

$$\mathbf{C}_i(t) \equiv \begin{array}{c} \begin{bmatrix} 1 & \rho_{12}(t) & \rho_{13}(t) \\ \rho_{12}(t) & 1 & \rho_{23}(t) \\ \rho_{13}(t) & \rho_{23}(t) & 1 \end{bmatrix} \end{array} \begin{array}{l} \text{interneurons} \\ \text{small pyramidal neurons} \\ \text{large pyramidal neurons} \end{array}.$$

5. Residual Bootstrap Procedure for Repeated K-functions

We now propose a residual bootstrap procedure for K-functions arising from the model at (4.1) and (4.2).

5.1. Review of existing approaches and outline of proposed extensions

Existing residual bootstrap procedures for empirical K-functions, arising from a one-way ANOVA type model, implement four basic steps: (i) fitting of a simple fixed-effects model to K-functions at each distance separately to generate residual functions; (ii) standardisation of residual functions to equalise their sampling variances across subjects; (iii) resampling of entire standardised residual functions; and (iv) generation of bootstrap replicates of individual and group mean K-functions. Generalisation of the modelling framework to the linear mixed-effects model here necessitates a number of modifications.

A. Resampling of random effects $\mathbf{b}_i(t)$ at subject level 2 as well as at level 1.
B. Preservation of the population covariance structure of the level-1 random effects.

C. Definition of a general estimator of parameter functions of interest.
D. Definition of a general test statistic.

We address A by implementing the residual bootstrap for multilevel models proposed by Carpenter, Goldstein and Rasbash (2003) to operate on entire K-function curves. Briefly, their procedure uses restricted maximum likelihood (REML) estimators of the multilevel model parameters to generate empirical best linear unbiased predictors (E-BLUPs) of the random effects at each level. These are known to provide unbiased predictions under normality (Kackar and Harville (1984)), and form the basis for independent resampling at each level. However, empirical sample moments of E-BLUPs differ from population moments of the random effects that they are predicting. While the random effects are assumed to have a population mean of zero, the mean of a sample of predictions will not necessarily be exactly zero, especially for higher level random effects and smaller sample sizes. Secondly, and more importantly, the variability of a sample of E-BLUPs is smaller than (the REML estimate of) the population variance of the random effects for all sample sizes, and in this sense the crude predictions are "shrunk" toward zero (Robinson (1991)). We therefore implement the Carpenter et al transformation for restoring population properties before resampling E-BLUPs.

Existing residual bootstrap methods transform level-1 residuals using only $\mathbf{V}_i(t)$ to equalise sampling variances before resampling. This practice would lead to resampled level-1 predictions that are independent between repeated measures as well as subjects, which is clearly not appropriate under a linear mixed model with $\mathbf{C}_i(t) \neq \mathbf{I}_{m_i}$. To preserve the population covariance structure of the level-1 random effects, $\sigma^2(t)\boldsymbol{\Lambda}_i(t)$, we suggest transforming level-1 predictions so that the modified values are E-BLUPs of a random sample from a common population with expectation 0 and variance $\sigma^2(t)$, resampling these, and then later re-instating the population covariance structure by backtransformation.

We suggest making use of the mixed model for K-functions to provide estimator functions for any linear combination of parameters (not only those involving group means), and to provide a general form of a test for addressing hypotheses of interest in spatial pattern investigations. The latter, in particular, requires some care in defining an appropriate scalar test statistic and to keep computational effort manageable.

5.2. Residual bootstrap procedure

Let $\hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t)$ denote the estimator of the fixed effects vector $\boldsymbol{\beta}(t)$ of the model at (4.1) and (4.2) derived by REML estimation under normality. Further let $\hat{\boldsymbol{\Psi}}_{\mathrm{REML}}(t)$, $\hat{\sigma}^2_{\mathrm{REML}}(t)$ and $\hat{\boldsymbol{\Lambda}}_{i,\mathrm{REML}}(t)$ denote the REML estimators of the variance-covariance parameters that provide $\hat{\boldsymbol{\Sigma}}_{i,\mathrm{REML}}(t) \equiv \widehat{\mathrm{Cov}}[\mathbf{K}_i(t)]$. The

E-BLUPS of the level-2 random effects, $\mathbf{b}_i(t)$, $i = 1, \ldots, N$, and the level-1 random effects $\boldsymbol{\varepsilon}_i(t)$, $i = 1, \ldots, N$, are derived by using expressions for the BLUPs $\hat{\mathbf{b}}_{i,\mathrm{BLUP}}(t) \equiv E\{\mathbf{b}_i(t)|\mathbf{K}_i(t)\}$ and $\hat{\boldsymbol{\varepsilon}}_{i,\mathrm{BLUP}}(t) \equiv \{\boldsymbol{\varepsilon}_i(t)|\mathbf{K}_i(t)\}$, but with unknown variance-covariance parameters replaced by their REML estimates.

Let $\mathbf{S}(t)$ denote the empirical $q \times q$-variance-covariance matrix at level 2, i.e., $\mathbf{S}(t) \equiv 1/N\{\hat{\mathbf{b}}_{\mathrm{E-BLUP}}(t)^T\hat{\mathbf{b}}_{\mathrm{E-BLUP}}(t)\}$ with $\hat{\mathbf{b}}_{\mathrm{E-BLUP}}(t) \equiv [\hat{\mathbf{b}}_{1,\mathrm{E-BLUP}}(t) \cdots \hat{\mathbf{b}}_{N,\mathrm{E-BLUP}}(t)]^T$ the matrix of (mean centred) level-2 E-BLUPS. For restoring population properties at level 2 we require a matrix, $\mathbf{A}(t)$, so that the linear transformation $\tilde{\mathbf{b}}(t) \equiv \hat{\mathbf{b}}_{\mathrm{E-BLUP}}(t)\mathbf{A}(t)$ has an empirical covariance matrix equal to the (estimated) population covariance matrix at level 2:

$$\tilde{\mathbf{S}}(t) \equiv \frac{1}{N}\{\tilde{\mathbf{b}}(t)^T\tilde{\mathbf{b}}(t)\} = \mathbf{A}(t)^T\mathbf{S}(t)\mathbf{A}(t) = \hat{\boldsymbol{\Psi}}_{\mathrm{REML}}(t).$$

A linear transformation that achieves this is (Carpenter, Goldstein and Rasbash (2003))

$$\mathbf{A}(t) \equiv \left[\hat{\boldsymbol{\Psi}}_{\mathrm{REML}}(t)^{0.5}\left\{\mathbf{S}(t)^{0.5}\right\}^{-1}\right]^T, \tag{5.1}$$

where $\mathbf{S}(t)^{0.5}$ denotes a $q \times q$-matrix that satisfies $\mathbf{S}(t)^{0.5}\{\mathbf{S}(t)^{0.5}\}^T = \mathbf{S}(t)$.

For preserving the population covariance structure of the level-1 random effects, $\sigma^2(t)\boldsymbol{\Lambda}_i(t)$, we require a set of matrices, $\mathbf{B}_i(t)$, $i = 1, \ldots, N$, so that the linear transformations $\mathbf{B}_i(t)\hat{\boldsymbol{\varepsilon}}_{i,\mathrm{E-BLUP}}(t)$, $i = 1, \ldots, N$, produce E-BLUPs of a random sample from a common population with expectation 0 and variance $\sigma^2(t)$. Linear transformations that achieve this are given by (Freedman and Peters (1984) and Solow (1985))

$$\mathbf{B}_i(t) \equiv \left\{\hat{\boldsymbol{\Lambda}}_{i,\mathrm{REML}}(t)^{0.5}\right\}^{-1}, \quad i = 1, \ldots, N. \tag{5.2}$$

We then suggest the following resampling procedure for E-BLUPs at level 1: (i) transform the level-1 predictions $\hat{\boldsymbol{\varepsilon}}_{i,\mathrm{E-BLUP}}(t)$ using the $\mathbf{B}_i(t)$, $i = 1, \ldots, N$; (ii) mean-centre and correct the empirical variance of the transformed values according to Carpenter et al's procedure; (iii) resample the (corrected) transformed level-1 predictions; (iv) backtransform these into level-1 bootstrap residual samples that exhibit heteroscedasticity and correlation by applying the inverse transformations $\mathbf{B}_i(t)^{-1}$, $i = 1, \ldots, N$. Resampling of the E- BLUPs at level 2 involves only two steps: (i) correct mean-centred level-2 predictions $\hat{\mathbf{b}}_{i,\mathrm{E-BLUP}}(t)$ using the transformation $\mathbf{A}(t)$, and (ii) resample the corrected level-2 predictions.

Our residual bootstrap procedure resamples entire curves, that is, the function values for the ith subject and jth repeated measure for different distances t are kept together at all times. This preserves any covariances between values at different distances (and avoids the need for modelling the latter explicitly).

Resampling of predictor curves is carried out independently at both levels. This breaks any correlation in line with the model assumption of independent random effects. At level 2 subjects are sampled with replacement to provide bootstrap replicates $\{\tilde{\mathbf{b}}_{i'}(t_1)^*, \ldots, \tilde{\mathbf{b}}_{i'}(t_{\max})^*\}$, $i' = 1, \ldots, N$. Note that the covariances of the level 2 random effects are preserved since the sampling of subjects ensures that multiple random effects predictions per subject are also kept together at all times. At level 1 subjects and repeated measures are sampled with replacement to (eventually) provide bootstrap replicates $\{\tilde{\boldsymbol{\varepsilon}}_{i'j'}(t_1)^*, \ldots, \tilde{\boldsymbol{\varepsilon}}_{i'j'}(t_{\max})^*\}$, $i' = 1, \ldots, N$; $j' = 1, \ldots, m_{i'}$ of the heteroscedastic and correlated level-1 errors.

K-functions can be rebuilt from the bootstrapped residual samples by defining

$$\tilde{\mathbf{K}}_{i'}(t)^* \equiv \mathbf{Z}_i \tilde{\mathbf{b}}_{i'}(t)^* + \tilde{\boldsymbol{\varepsilon}}_{i'}(t)^* + \mathbf{X}_i \hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t), \quad i = 1, \ldots, N, \tag{5.3}$$

where $\tilde{\boldsymbol{\varepsilon}}_{i'}(t)^* \equiv [\tilde{\varepsilon}_{i'1}(t)^*, \ldots, \tilde{\varepsilon}_{i'm_{i'}}(t)^*]^T$ and a bootstrap replicate of the estimator statistic $\mathbf{c}^T \hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t)$ for the linear combination $\mathbf{c}^T \boldsymbol{\beta}(t)$ constructed by $\mathbf{c}^T \tilde{\boldsymbol{\beta}}_{\mathrm{REML}}(t)^*$ where $\tilde{\boldsymbol{\beta}}_{\mathrm{REML}}(t)^*$ denotes the REML fixed effects estimator calculated from the new K-function replicates. Repeated bootstrap resampling then provides the sampling distribution of the estimator statistic. This simulated distribution can then be used to construct confidence bands for $\mathbf{c}^T \boldsymbol{\beta}(t)$, $t = t_1, \ldots, t_{\max}$ with pointwise confidence level $1-\alpha$ by using percentiles or other relevant methods (see Carpenter and Bithell (2000)). Since an analytic form of the standard error of the estimator is readily available in linear mixed models bootstrap t-confidence bands can be constructed (for details see Davison and Hinkley (1997)).

We now consider a residual bootstrap procedure for formally testing

$$H_0 : \mathbf{C}^T \boldsymbol{\beta}(t) = \mathbf{h}_0(t) \ \text{ for all } \ t \quad \text{against} \quad H_1 : \mathbf{C}^T \boldsymbol{\beta}(t) \neq \mathbf{h}_0(t) \ \text{ for any } \ t$$

where $\mathbf{C}^T = [\mathbf{c}_1 \cdots \mathbf{c}_q]^T$, $q \leq p$, and $\mathbf{h}_0(t) = [h_{01}(t) \cdots h_{0q}(t)]^T$, contain known values. Typically $\mathbf{h}_0(t) \equiv 0$ for all t but other choices might be of interest, for example $h_0(t) \equiv \pi t^2$ for testing the CSR (univariate K-functions) or spatial independence (bivariate K-functions) hypotheses. For confidence band construction the bootstrap replicates of the K-functions are not restricted in any way, and we simply add back the expected K-function values, $\mathbf{X}_i \hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t)$, to the resampled residual functions as shown at (5.3). In contrast, for hypothesis test derivation we need to generate the distribution of a test statistic under the null hypothesis. The resampling plan therefore needs to be modified to add back K-function values expected under the null hypothesis, $\mathbf{X}_i \hat{\boldsymbol{\beta}}_{\mathrm{REML,H_0}}(t)$. Computationally, this is achieved by reparametrisation of the model so that the linear combinations to be tested constitute parameters with explanatory variables that are orthogonal

to the explanatory variables of the remaining parameters. The K-function values expected under the null hypothesis can then be estimated without re-fitting the linear model, simply by using the linear predictor of the fitted models after restricting the effect(s) in question to the hypothesized values.

The estimator, $\mathbf{C}^T\hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t)$, as a measure of effect size, provides the basis of a test statistic. However, consideration needs to be given to how best to combine q-dimensional effect estimates at different distances into a scalar test statistic. In analogy with general linear models, we suggest using a quadratic test statistic

$$T \equiv \sum_{t=t_1}^{t_{\max}} \Big\{\mathbf{C}^T\hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t) - \mathbf{h}_0(t)\Big\}^T \Big[\mathbf{C}^T\Big\{\sum_{i=1}^{N} \mathbf{X}_i^T\hat{\boldsymbol{\Sigma}}_{i,\mathrm{REML}}(t)^{-1}\mathbf{X}_i\Big\}^{-1}\mathbf{C}\Big]^{-1} \times\Big\{\mathbf{C}^T\hat{\boldsymbol{\beta}}_{\mathrm{REML}}(t) - \mathbf{h}_0(t)\Big\}. \tag{5.4}$$

The test statistic standardizes the estimated effects at each distance by their estimated standard errors before combining effects across distances. Previous approaches have used variance stabilizing transformations (Diggle, Lange, and Beneš (1991) and Landau, Rabe-Hesketh and Everall (2004)) or distance weights (Diggle, Mateu and Clough (2000)) to account for the increased variability of empirical K-functions values at larger distances. We prefer (5.4) since it utilizes the replication in the data rather than approximate theoretical relationships to estimate distance weights, and also extends naturally to any q-dimensional linear hypothesis.

5.3. Bootstrapping the HAD bivariate K-functions

The empirical bivariate K-functions were assumed to follow the model described in Section 4.3. We now obtain inferences for the model parameters by using the modified residual bootstrap. All results were produced in `R` (R Development Core Team (2004)) using the `nlme` (Pinheiro and Bates (2000)) and `splancs` (Rowlingson and Diggle (1993)) packages; see http://www3.stat.sinica.edu.tw/statistica.

For descriptive purposes we started by generating unbiased estimates and confidence bands for the mean bivariate K-functions within dementia groups and neuronal types. K-functions were resampled as described in the previous section using 999 bootstrap simulations. Since the underlying model did not include any level 2 random effects, resampling operated at only one level. In this instance transformation (5.2) (which used moderate-sized correlation estimates: $\hat{\rho}_{23,\mathrm{REML}}(t)$ reaching 0.7 for most distances t; $\hat{\rho}_{12,\mathrm{REML}}(t)$ and $\hat{\rho}_{13,\mathrm{REML}}(t)$ consistently negative around -0.3 and -0.5, respectively) turned out to be more influential than transformation (5.1) (variance inflation factor after mean centring of the transformed residual: $\{84/(84-9)\}^{0.5} = 1.06$). Once K-function

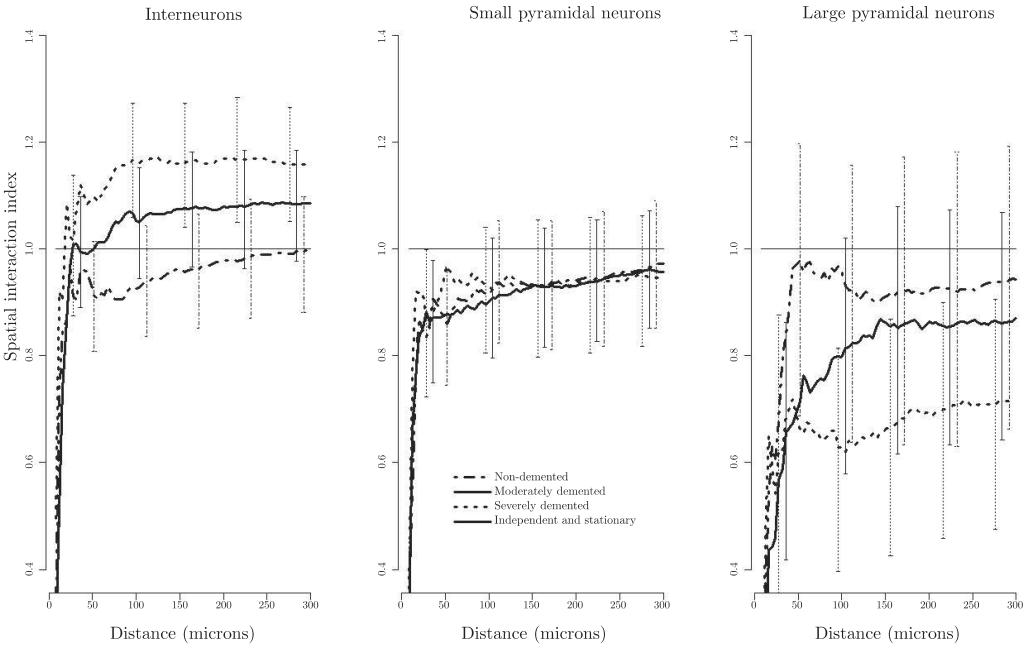

Figure 2. Group mean spatial interaction indices for neurons and astrocytes by neuronal subpopulation and dementia groups. Index function values above the threshold value 1 indicate apparent spatial attraction at the respective distance, while values below this threshold indicate apparent spatial repulsion. For each dementia group the relevant error bars indicate point-wise 95% confidence intervals constructed by bootstrapping with 999 runs.

replicates were rebuilt according to (5.3), bootstrap replicates of the respective estimator statistics and their estimated standard errors were also available and could be converted into bootstrap t-confidence bands.

To ease interpretation, Figure 2 displays the results in terms of spatial interaction indices. All estimated spatial interaction functions start at a value of (or almost) zero due to no (very few) astrocytes being found within the very small vicinities of the neurons (4 or 8 microns), and rapidly increase thereafter reflecting a hard core effect. For interneurons and astrocytes the interaction indices are continually increasing, eventually reaching attraction levels in the dementia groups. However, the confidence bands for the non-demented group and the moderately demented group (not shown completely in Figure 2 for graphical clarity) include the reference line at $y = 1.0$ (and also lines drawn at $y = 1.01$ or $y = 1.03$), indicating consistency with the spatial independence model. For small pyramidal neurons and astrocytes the three group-wise estimated spatial interaction indices were similar and consistent with the spatial interaction model.

Finally, for large pyramidal neurons, the spatial interaction indices in the non- and moderately demented groups reached levels consistent with the independence model, while the spatial relationship appeared repulsed in the severe group at all distances. (Note that the increased widths of the confidence bands for large pyramidal neurons relative to interneurons and small pyramidal neurons are a reflection of the relative scarcity of large pyramidal neurons due to variance function assumption (4.2), see Figure 1.)

The experimental hypothesis of dysfunctional astrocytes translated into an assessment of the interaction between the experimental factors dementia group (here coded as two separate contrasts) and neuronal cell type with regard to the spatial interaction between neurons and astrocytes (here measured by bivariate K-functions). We therefore tested the following:
(i) interaction between neuronal type and linear trend of dementia grading

$$H_0^{(I)} : \mathbf{C}_1^T \boldsymbol{\beta}(t) \equiv \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \boldsymbol{\beta}(t) = \begin{bmatrix} \beta_6(t) \\ \beta_7(t) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \text{for all} \quad t,$$

and (ii) interaction between neuronal type and quadratic trend of dementia grading

$$H_0^{(II)} : \mathbf{C}_2^T \boldsymbol{\beta}(t) \equiv \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \boldsymbol{\beta}(t) = \begin{bmatrix} \beta_8(t) \\ \beta_9(t) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \text{for all} \quad t,$$

against respective two-sided alternative hypotheses. Using quadratic test statistics of the form (5.4), the linear effect of dementia varied with neuronal type at the trend level ($p = 0.10$), while the observed K-functions were consistent with a constant quadratic effect for all three neuronal types ($p = 0.92$). We therefore dropped the latter interaction effect and also the main quadratic trend of dementia grading after further testing ($p = 0.80$). The nature of the detected interaction was such that, while spatial interaction between astrocytes and interneurons appeared more attracted for severely demented subjects than for non-demented subjects, the relationship between astrocytes and small pyramidal neurons was not affected in dementia, and the relationship between astrocytes and large pyramidal neurons became more repulsed (Figure 2).

The final model further served to quantify the size of the effects. Figure 3 shows the estimated differences between the bivariate K-functions in the severely demented and non-demented groups for each neuronal type, together with pointwise 95% bootstrap t-confidence bands. Figure 3 gives results in terms of index differences representing *absolute* differences in percentage change from the expected value under independence. For example, compared with the non-demented group, the estimated spatial interaction index for interneurons at distance 200 microns was 19.3% higher in the severely dementia group.

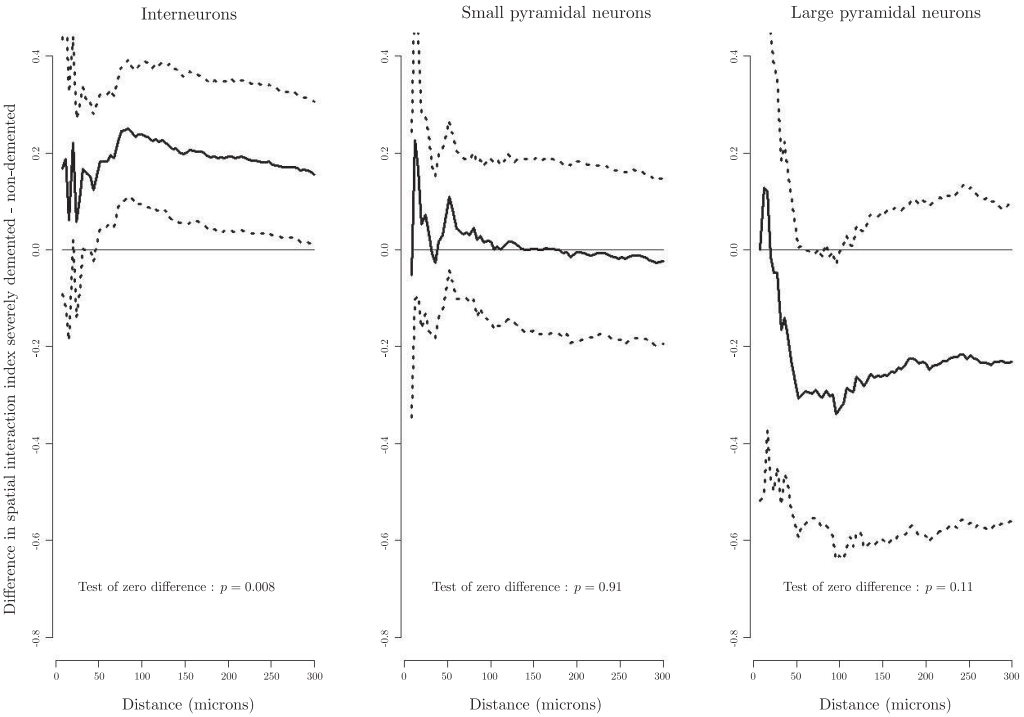

Figure 3. Difference in spatial interaction indices for neurons and astrocytes between severely demented and non-demented patients by neuronal subpopulation. Differences above the threshold value 0 indicate increased apparent attraction in the severely demented group relative to the non-demented group, and vice versa. The bands indicate pointwise 95% confidence bands constructed by bootstrapping with 999 runs.

We were able to directly address the experimental hypothesis that changes in spatial interaction with astrocytes in HAD depended on the type of neuron. The finding of an increased repulsion restricted to large pyramidal neurons is consistent with the excitotoxic damage hypothesis. In addition, the estimated correlations between the repeated measures showed that, within dementia groups, those subjects demonstrating more repulsion between astrocytes and large pyramidal neurons also tended to show more repulsion with small pyramidal neurons, consistent with a common function of pyramidal neurons. In contrast, increased repulsion with pyramidal neurons was associated with increased attraction with interneurons.

6. Discussion

We have described a nonparametric(perhaps one should say semi-parametric) residual bootstrap procedure for obtaining statistical inferences in linear mixed-

effects models for repeated empirical K-functions. The procedure extends existing approaches (Diggle, Lange and Beneš (1991) and Landau, Rabe-Hesketh, and Everall (2004)) to repeated measures designs, and provides a general modelling framework for the analysis of spatial cell dependencies, for example, to adjust for confounding variables or to investigate higher order effects.

While described here as an inference procedure for two-level mixed models for repeated K-functions, the bootstrap approach is in principle applicable to general multilevel models of K-functions. The residual resampling procedure described for the subject level (level 2) would simply have to be repeated at all higher levels. Thus an investigation assessing the spatial patterning of cells from multiple brain regions from several family members could be envisaged as a three-level model.

The suggested bootstrap approach for mixed-effects models for K-functions can yet be more generally considered a nonparametric procedure for obtaining inferences in multivariate mixed-effects models without specifying the correlation structure between the multiple dependent variables. The only aspect that is specific to K-functions is the modelling of the effect of process parameters on the precision of the function values. Thus the bootstrap approach could be used for spatial proximity functions other than K-functions, e.g., the F-, G- or J-functions (van Lieshout and Baddeley (1996)) and/or carried over to proximity functions for three-dimensional space (e.g., Baddeley, Moyeed, Howard and Boyde (1993)), provided that relationships between precisions and process parameters can be accommodated.

In practice, the use of summary measures such as the K-functions seems the best way forward. Fully parametric approaches for the analysis of replicated spatial point patterns have been proposed (Diggle, Mateu and Clough (2000), Mateu (2001) and Bell and Grunwald (2004)). These might provide more powerful inferences if a suitable bivariate point process model could be identified. However, this might prove difficult, especially when inhomogeneities have to be modelled and the gain in power might not be that great. In a simulation study, Diggle, Mateu and Clough (2000) showed that the nonparametric approach was reasonably powerful compared to the parametric approach even when the point patterns were generated from the correct point process model. (The study also confirmed that the parametric approach was not robust against mis-specifications of the point process model.)

An appealing feature of the bootstrap approach is that it is based purely on replication rather than stationarity, since histological structures, and coronal brain sections in particular, often display inhomogeneities. However, nonstationarity does affect the interpretation of spatial proximity measures in that detected effects can relate to the nonstationarity mechanism rather than cell relationships

per se. But at least the association between a combined measure of nonstationarity and cell interaction and explanatory variables can be formally investigated, and this we consider to be of value in itself. In this context a natural progression of our methodology that may provide for more clearcut interpretation, would be to adapt it to work with more refined summary functions. For example, an analogue of the K-function for measuring the interaction between points arising from an inhomogeneous process (Baddeley, Møller and Waagepetersen (2000)), or, if a suitable nonstationary point process model could be fitted to each cell pattern, their resulting K-function estimators, could be investigated.

A limitation of our proposed methodology is that variance covariate values, introduced to account for the dependence of the precision of K-functions on point process parameters, are assumed known. In practice they have to be estimated and the associated extra variability is currently ignored in our inferential procedures. Future work should look at ways of accounting for the imprecision of the variance covariate value estimates.

Acknowledgement

We are grateful to Eleanor Roberts of the Department of Psychiatry, University of California, San Diego, for providing the cell positions dataset. The authors gratefully acknowledge Dr. Lídice Galán for her useful comments during the revision of the manuscript.

References

Asare, E., Dunn, G., Glass, J., McArthur, J., Luthert, P., Lantos, P. and Everall, I. (1996). Neuronal pattern correlates with the severity of human immunodeficiency virus-associated dementia complex. *American Journal of Pathology* **148**, 1-8.

Baddeley, A. J., Moyeed, R. A., Howard, C. V. and Boyde, A. (1993). Analysis of three-dimensional point pattern with replication. *Appl. Statist.* **42**, 641-668.

Baddeley, A. J., Møller, J. and Waagepetersen, R. (2000). Non- and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statist. Neerlandica* **54**, 329-350.

Bell, M. L. and Grunwald, G. K. (2004). Mixed models for the analysis of replicated spatial point patterns. *Biostatistics* **5**, 633-648.

Carpenter, J. and Bithell, J. F. (2000). Bootstrap confidence intervals: when, which, what? A practical guide for medical statisticians. *Statist. Medicine* **19**, 1141-1164.

Carpenter, J. R., Goldstein, H. and Rasbash, J. (2003). A novel bootstrap procedure for assessing the relationship between class size and achievement. *Appl. Statist.* **52**, 431-443.

Chana, G., Landau, S., Beasley, C., Everall, I. P. and Cotter, D. (2003). Two-dimensional assessment of cytoarchitecture in the anterior cingulate cortex in major depressive disorder, bipolar disorder, and schizophrenia: Evidence for decreased neuronal somal size and increased neuronal density. *Biological Psychiatry* **35**, 1086-1098.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap Methods and their Applications*. Cambridge University Press, Cambridge.

Diggle, P. J., Lange, N. and Beneš, F. M. (1991). Analysis of variance for replicated spatial point patterns in clinical neuroanatomy. *J. Amer. Statist. Assoc.* **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. *Adv. Appl. Probab.* **32**, 331-343.

Everall, I. P., Glass, J. D., McArthur, J. C., Spargo, E. and Lantos, P. L. (1994). Neuronal density in the superior frontal and temporal gyri does not correlate with the degree of HIV associated dementia. *Acta Neuropathologica* **88**, 538-544.

Freedman, D. A. and Peters, S. C. (1984). Bootstrapping a regression equation: some empirical results. *J. Amer. Statist. Assoc.* **79**, 97-106.

Kackar, R. N. and Harville, D. A. (1984). Approximations for standard errors of estimators of fixed and random effects in mixed linear models. *J. Amer. Statist. Assoc.* **79**, 853-862.

Landau, S., Rabe-Hesketh, S. and Everall, I. P. (2004). Nonparametric one-way analysis of variance of replicated bivariate spatial point patterns, *Biometrical J.* **46**, 19-34.

Lotwick, H. W. and Silverman, B. W. (1982). Methods for analysing spatial processes of several types of points. *J. Roy. Statist. Soc. Ser. B* **44**, 406-413.

Mateu, J. (2001). Parametric procedures in the analysis of replicated pairwise interaction point patterns. *Biometrical J.* **43**, 375-394.

Pinheiro, J. C. and Bates, D. M. (2000). *Mixed-Effects Models in S and S-PLUS*. Springer, New York.

Price, R. W. and Brew, B. J. (1988). The AIDS dementia complex. *Journal of Infectious Diseases* **158**, 1079-1083.

R Development Core Team (2004). *R: A language and environment for statistical computing. R Foundation for Statistical Computing*. Vienna, Austria. URL http://www.R-project.org.

Ripley, B. D. (1981). *Spatial Statistics*. Wiley, New York.

Roberts, E. (2000). An investigation into the relationship between microglia, astrocytes and neuronal subpopulations in HIV affected cases. Unpublished PhD thesis, University of London, UK.

Robinson, G. K. (1991). That BLUP is a good thing: the estimation of random effects. *Statist. Sci.* **6**, 15-51.

Rowlingson, R. S. and Diggle, P. J. (1993). Spatial point pattern analysis code in S-Plus. *Computers in Geosciences* **19**, 627-655.

Schladitz, K., Särkkä, A., Pavenstädt, I., Haferkamp, O. and Mattfeldt, T. (2003). Statistical analysis of intramembranous particles using freeze fracture specimens. *J. Microsc.* **211**, 137-153.

Solow, A. R. (1985). Bootstrapping correlated data. *Math. Geol.* **17**, 769-775.

Stoyan, D., Kendall, W. S. and Mecke, J. (1987). *Stochastic Geometry and its Applications*. Wiley, New York.

van Lieshout, M. N. M. and Baddeley, A. J. (1996). A nonparametric measure of spatial interaction in point patterns. *Statist. Neerlandica* **50**, 344-361.

Wager, C. G., Coull, B. A. and Lange, N. (2004). Modelling spatial intensity for replicated inhomogeneous point patterns in brain imaging. *J. Roy. Statist. Soc. Ser. B* **66**, 429-446.

Department of Biostatistics, King's College London, Institute of Psychiatry, London SE5 8AF, UK.

E-mail: s.landau@iop.kcl.ac.uk

Department of Psychiatry, University of California, San Diego, CA 920-93-0603, U.S.A.

E-mail: ieverall@ucsd.edu

(Received April 2007; accepted April 2008)