Spatial Statistics 17 (2016) 199–219

Contents lists available at ScienceDirect

Spatial Statistics

journal homepage: www.elsevier.com/locate/spasta

Statistical analysis of spatially homogeneous dynamic agent-based processes using functional time series analysis

Jack D. Hywood [a,*], Mark N. Read [b], Gregory Rice [c]

[a] Sydney Medical School, The University of Sydney, Sydney, Australia

[b] School of Life and Environmental Sciences & Charles Perkins Centre, University of Sydney, Sydney, Australia

[c] Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Canada

ARTICLE INFO

Article history:
Received 8 January 2016
Accepted 7 June 2016
Available online 29 June 2016

Keywords:
Agent-based processes
Agent-based models
Functional time series
Moment dynamics
Spatio-temporal point process
Spatial statistics

ABSTRACT

Dynamic systems consisting of multiple interacting autonomous individuals are of particular interest in a number of scientific fields, including ecology, biology, and swarm robotics. Such systems are commonly referred to as agent-based processes. Detection and characterisation of agent–agent interactions is an important step in the analysis of agent-based processes, however existing statistical methods are relatively limited. This paper presents a novel framework for investigating spatial interactions between agents combining techniques from spatial statistics and functional time series analysis. Assuming second order spatial equilibrium of the agent-based process, we develop a test for identifying the specific nature of interactions between agents. We also consider methodology for validating the assumption of spatial equilibrium for a given realisation of the agent-based process. The efficacy of this methodology is demonstrated via Monte Carlo simulation studies and an application to experimental data obtained by observing a species of flightless locust.

* Corresponding author.
E-mail address: jhyw1620@uni.sydney.edu.au (J.D. Hywood).

http://dx.doi.org/10.1016/j.spasta.2016.06.002
2211-6753/

1. Introduction

Dynamic systems of interacting autonomous individuals are encountered in a diverse range of sciences, including ecology, biology, and swarm robotics (Balch, 2000; Plotkin et al., 2000; Cavagna et al., 2008; Buhl et al., 2012; Brambilla et al., 2013; Dale and Fortin, 2014; Johnston et al., 2014; Michalec et al., 2015b; Russell et al., 2015; Binder and Simpson, 2015). Technological developments increasingly allow for the tracking or imaging of individuals in processes arising in these fields, from GPS-tagging of animals (Cagnacci et al., 2010; Dell et al., 2014) and video recordings of animal movements (Cavagna et al., 2008; Buhl et al., 2011, 2012; Michalec et al., 2015a,b) to *in situ* imaging of cells within live animals (Chtanova et al., 2008; Ng et al., 2011; Meijering et al., 2012; Tong et al., 2015). Such systems are typically referred to as agent-based processes and models of these are similarly referred to as agent-based models, or individual-based models.

Observing agent-based processes over time results in sequential sets of positional point data, or point patterns, representing agent positions. For example, an animation of wingless locust nymph movement data, from which two frames are presented in Fig. 1, is presented in the supplementary material. These data are analysed in a later section of this paper. A primary interest in the analysis of such dynamic positional data is to determine the existence and nature of interactions between agents.

While many commonly observed agent-based processes evolve via agent movements, they can also evolve via births and deaths of agents. For example, a population of cells could increase in number via mitosis. Moreover, over longer time scales processes that are not necessarily considered dynamic can in fact be thought of as agent-based processes. For example, the evolution of tree locations in a forest could be considered as a birth–death agent-based process with the probabilities of births and deaths being influenced by the relative position of other trees.

A novel approach used to analyse agent-based models is moment dynamics (Bolker and Pacala, 1997; Law and Dieckmann, 2000; Plank and Law, 2015). Dynamic spatial moments are analogous to the intensity terms in spatial statistics (Diggle, 2003), but describe dynamic agent-based models rather than static point processes. The dynamic first and second spatial moments of an agent-based process can be combined to produce a dynamic pair correlation function and related K function. These functions indicate ranges and types of spatial interaction between agents. Moment dynamic analysis is becoming increasingly important and widespread in the analysis of agent-based models (Raghib et al., 2011; Plank and Law, 2015; Binny et al., 2015).

The study of agent-based processes from a statistical perspective has been limited. Although approaches for the analysis of agent-based processes with respect to pair correlation functions have recently been developed (Cavagna et al., 2008; Binder and Simpson, 2013) these methods have several limitations. Primary among these is the lack of a statistical test of significance for detecting agent–agent interactions from experimental data. Furthermore, existing approaches lack a means of determining when it is reasonable to pool data across observations of an agent-based process to produce a single estimate for the pair correlation function or K function. If the expected spatial arrangement of agents is time invariant, then such data can be pooled, making estimation more accurate. However, if the expected spatial arrangement is not invariant with time, as might occur if the spatial arrangement at observed times is dependent on initial conditions, or if agent–agent interactions are dependent on an environmental factor that varies with time, then the data should not be pooled.

In this paper, we present a general framework for the analysis of agent-based processes that addresses these issues. The methodology that we propose is based on analysing the functional time series that is generated by calculating empirical K functions for each set of agent positional data. The resulting series of curves often exhibit significant autocorrelation due to the dependence between successive spatial arrangements of the agents. The discipline of functional time series analysis has grown considerably in the last decade to provide methodology for the analysis of autocorrelated functions (Ferraty and Vieu, 2006; Hörmann et al., 2010; Horváth and Kokoszka, 2012).

By using techniques in functional time series analysis, we develop a test for stationarity of the empirical K functions. As we will show, the stationarity of empirical K functions suggests that the underlying process is in a type of spatial equilibrium, which we refer to as reduced second moment temporal stationarity. Importantly, if it is determined that an agent-based process exhibits such

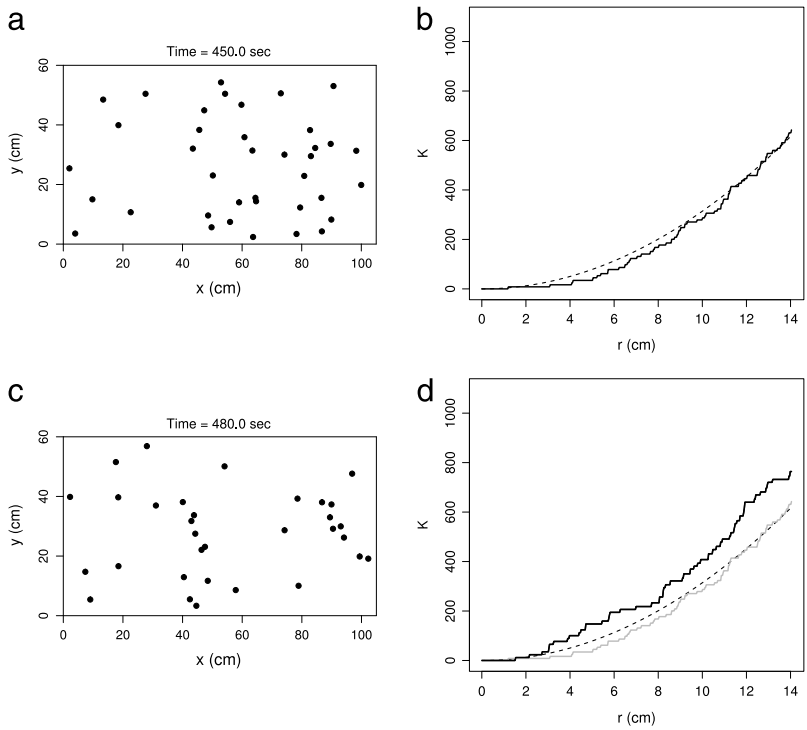

Fig. 1. Positional data from a field experiment recording wingless locust nymph movements, presented in Buhl et al. (2012), and the associated empirical K functions associated with each set of positional data. (a) and (c) give the positional data from two consecutive time points, and (b) and (d) give the empirical K functions (———) associated with each respective set of positional data, and πr^2 (– – –), the theoretical K function for agent-based processes with no interaction. (d) also indicates the previous empirical K function from the previous time point from (b) (———).

stationarity, one can pool data from the entire realisation of the process to produce a single accurate estimate for a time invariant K function that may be used for inference regarding agent–agent interactions. We develop a hypothesis test based on this estimated K function for the hypothesis that agent–agent interactions are of a specific nature. For example, our test can be used to determine if agent positional data is consistent with completely random motion.

The rest of the paper is organised as follows. In Section 2 we provide some background on spatial moments of agent-based processes, and introduce the main objects of study in this paper. In Section 3 we introduce a statistical test for second moment stationarity for agent-based processes. In Section 4 we present a statistical test for detecting statistically significant clustering and regularity. In Section 5 we introduce an example of a continuous time agent-based model, and perform analysis for this model to demonstrate the performance on the presented statistical tests. We also consider in Section 6 an application to dependent sequences of spatial point patterns that arise from the simulation of Gibbs processes via the Metropolis–Hastings algorithm. In Section 7 we use our methods to analyse field experiment data for wingless locust nymphs.

2. Agent-based processes, spatial moments, and the K function

2.1. Notation and definitions

Agent-based processes may be modelled by spatiotemporal stochastic processes consisting of agents, or points, in $\mathbb{R}^d$ evolving through time via a set of rules. Such rules can allow for interaction between agents. We envision that the agents in such a process are evolving in either of, or a mixture

of, two modes: (1) that agents undergo continuous movement within $\mathbb{R}^d$, or (2) new agents enter via births and are removed via deaths, and that these actions are based on their current location and the relative position of other agents. In Section 4 of Ripley (1977), both types of such evolving point processes are formally defined, focusing on the spatial birth–death process and interacting particles performing diffusions. Further information on the mathematical foundations of such processes may be found in Preston (1975) and Schweitzer (2003).

In practice, dynamic agent-based processes can only be observed in a bounded region at discrete time points. Let us define an *observation region*, $A \subseteq \mathbb{R}^d$, and let the time points at which observations occur be given as $\mathbf{T} = \{t_1, t_2, \ldots, t_N\}$, with N giving the total number of time points observed. For simplicity, we assume that the time between observations is constant; i.e. $t_{i+1} - t_i = a$, for $i = 1, \ldots, N-1$, with a being constant. For times $\mathbf{T}$, we only observe agents with positional data $x \in A$. As such, for any given time $t \in \mathbf{T}$, the realisation of the dynamic agent-based process results in a spatial point pattern within A.

Let X_i give the spatial co-ordinate data for **observed** agents recorded for time t_i, such that $X_i = \{x_1, x_2, \ldots, x_{n_i}\}$, with x_j giving the spatial co-ordinates of a single agent within the observation region A, and n_i giving the total number of agents occurring within the observation region A at time t_i. Let $\boldsymbol{X} = \{X_1, X_2, \ldots, X_N\}$ give the complete set of spatial data for all N time points.

Raghib et al. (2011), Plank and Law (2015) and Binny et al. (2015) define the first spatial moment as

$$\lambda(u; t) = \lim_{|du| \to 0} \frac{\mathbb{E}[Y(du|t)]}{|du|},$$

with du defining a small region around the point u, $|du|$ giving the volume of du, $Y(du|t)$ giving the number of agents within a region du given time t, and $\mathbb{E}[\cdot]$ denoting expectation.

The second spatial moment is defined as

$$\lambda_2(u, u'; t) = \lim_{|du|, |du'| \to 0} \frac{\mathbb{E}[Y(du|t)Y(du'|t) - Y(du \cap du'|t)]}{|du||du'|},$$

with the second term in the expectation designed to remove the Dirac delta peak that would otherwise arise for $u = u'$ (Plank and Law, 2015). The first and second spatial moments are analogous to the first and second-order intensities common to spatial and spatiotemporal statistics (Diggle, 2003; Gabriel and Diggle, 2009; Cressie and Wikle, 2011). Higher order spatial moments can be similarly defined, although the study of agent-based processes typically focuses on the first and second moments.

As in Plank and Law (2015), we refer to processes for which statistical properties with respect to spatial arrangements of agents are invariant with translation as being "spatially homogeneous", i.e. such processes are spatially stationary (Diggle, 2003). Spatial homogeneity implies that the first spatial moment is not dependent on position and that higher order spatial moments are only dependent on the distances between positions. Thus, for a spatially homogeneous agent-based process, for all $u, u' \in A$, we can express the first spatial moment $\lambda(u; t)$ as $\lambda(t)$ and we can express the second spatial moment $\lambda_2(u, u'; t)$ as $\lambda_2(r; t)$, with $r = |u - u'|$.

The pair correlation function, and the reduced second-order moment measure, the K function, are well studied in spatial statistics as they are typically more interpretable and easier to estimate than the second spatial moment itself (i.e. the second order intensity in spatial statistics). For a spatially homogeneous agent-based process, the pair correlation function and K function are defined as

$$g(r; t) = \frac{\lambda_2(r; t)}{\lambda(t)^2}, \quad \text{and} \quad K(r; t) = 2\pi \int_0^r g(u; t) u du. \tag{1}$$

We can express $K(r; t)$ as an expectation:

$$K(r; t) = \frac{1}{\lambda(t)} \mathbb{E}[\text{Number of additional agents within } r \text{ of a randomly chosen agent at time } t].$$

It is well known that the pair correlation function for a non-interacting spatially homogeneous agent-based process is $g(r; t) = 1$ (Raghib et al., 2011; Plank and Law, 2015; Binny et al., 2015). Thus

the K function for such processes is $K(r;t) = \pi r^2$. As such, pair correlation functions, and K functions, for interacting spatially homogeneous agent-based processes can be compared against the respective functions associated with non-interacting agent-based processes to indicate types, magnitudes, and ranges of agent–agent interactions. All concepts covered in this paper are equally applicable to both g and K functions. Because calculation of empirical estimates for the K function is substantially easier than for g, we focus our analysis within this paper on the K function.

For each spatial point pattern, X_i, we can produce an empirical $\hat{K}(r;t)$ function as an estimate of $K(r;t)$. Due to the discrete time nature of observations, we instead express $\hat{K}(r;t_i)$ as $\hat{K}_i(r)$, with $\hat{K}_i(r)$ corresponding to the estimate for time t_i. We define $\hat{K}_i(r)$ similarly to the K function of spatial statistics,

$$\begin{aligned}\hat{K}_i(r) &= \frac{1}{\hat{\lambda}_i n_i}\sum_{k=1}^{n_i}\sum_{l\neq k} w_{k,l}\mathbb{1}(\delta(k,l)\leq r),\\ &= \frac{|A|}{n_i^2}\sum_{k=1}^{n_i}\sum_{l\neq k} w_{k,l}\mathbb{1}(\delta(k,l)\leq r),\end{aligned} \tag{2}$$

where n_i is the number of agents within A at time t_i, $|A|$ is the volume of A, $\hat{\lambda}_i = n_i/|A|$ is the unbiased estimate for $\lambda(t)$ for time t_i, $\delta(k,l)$ is the distance between the kth and lth agents within A at time t_i, $\mathbb{1}(\cdot)$ is the indicator function, and $w_{k,l}$ is an appropriate edge correction weight associated with the kth and lth agents (Diggle, 2013). The inclusion of an edge correction term reduces the bias that would otherwise occur due to the fact that the observation region A represents only a subset of the entire process, with agents occurring external to the boundaries of A. An example of two time points from an agent-based process, movement data for wingless locust nymphs from a field experiment, and the associated $\hat{K}_i(r)$ functions are given in Fig. 1.

The estimate $\hat{K}_i(r)$ in (2) is an approximately unbiased estimator of $K(r;t_i)$ for small enough r (Ripley, 1977; Diggle, 2003, 2013). On account of this we restrict r values to $r \in [0,R]$, for some arbitrary R such that the bias of $\hat{K}_i(r)$ is relatively small for $r \leq R$. Appropriate values for R generally depend on the geometry of the observation region. For example, for a 2D rectangular observation region a recommended value for R is $1/4$ the length of the smallest side length (Baddeley et al., 2015).

We note that it is typical to use $\hat{\lambda}_i = (n_i - 1)/|A|$ in (2) instead of $n_i/|A|$ for reasons pertaining to bias, though the difference between estimates is clearly small for large enough n_i. We use $\hat{\lambda}_i = (n_i - 1)/|A|$ in all numerical analysis presented in this paper. For a single longitudinal set of spatial point patterns, $\boldsymbol{X} = \{X_1, X_2, \ldots, X_N\}$, arising from a single realisation of an agent-based process, we can produce a series of functions, $\hat{\boldsymbol{K}} = \{\hat{K}_1(r), \hat{K}_2(r), \ldots, \hat{K}_N(r)\}$, ordered sequentially in time.

Before proceeding to specific applications, we briefly address some of the benefits and limitations to using K and pair correlation functions in order to study point patterns generated by observing agent-based processes. One benefit is that they do not rely on a specific model to quantify agent–agent interactions. Though model-based analysis of agent-based processes requires a thorough prior knowledge of the underlying process to facilitate construction of a realistic model, this approach has the potential to elucidate detailed characteristics about the interaction of agents. For example, Buhl et al. (2012) attempt to fit an agent-based model to experimental data by calculating the model parameters to minimise the difference between moment measure computed from the data, and those computed from simulations of the model. Other examples of model-based approaches include Gautrais et al. (2012), Cavagna and Giardina (2014), Plank and Law (2015), Russell et al. (2015) and Cavagna et al. (2015). The proposed methodology may supplement model based approaches by assessing the homogeneity of second moment measures estimates used to fit such models (see Appendix C). Additionally, K functions do not require that the agents are each individually tracked, and all agents within an observation region are utilised in analysis regardless of whether they entered the region after the beginning of the observation period, or whether they leave it prior to the end of the observation period. On the other hand, any inference about the process at hand based on second moment measures is restricted to the information available in them. This could be misleading since, for

example, one may construct point patterns for which $K(r) = \pi r^2$ that are not Poisson point processes (Baddeley and Silverman, 1984).

2.2. Temporal stationarity for spatially homogeneous agent-based processes

Determining whether certain properties of agent-based processes with respect to spatial arrangements are invariant with time has major implications for the types of analysis that can be performed. Here we define an important type of temporal stationarity.

Definition 2.1. For a spatially homogeneous agent-based process, the process is reduced second-moment (RSM) temporally stationary if the K function is time invariant, such that $K(r; t) = K_{\mathbb{1}}(r)$.

We note that RSM temporal stationarity is independent of whether $\lambda(t)$ is also invariant with time. From Eq. (1) it is clear that if $K(r, t)$ is time invariant then $g(r; t) = g_{\mathbb{1}}(r)$ is also time invariant.

Examples of agent-based processes that would be expected to be RSM temporally stationary would be processes for which the ranges, types, and strengths of agent–agent interactions are invariant with respect to time. In contrast, processes in which agent–agent interactions can change with time would be unlikely to be RSM temporally stationary. Such processes might include processes for which agent–agent interactions are influenced by a time dependent environmental factor.

3. Testing for RSM temporal stationarity

For an RSM temporally stationary processes, estimates for $g_{\mathbb{1}}(r)$ or $K_{\mathbb{1}}(r)$ can provide information regarding agent–agent interactions similar to the analogous functions in spatial statistics. This motivates developing a test that measures the validity of the assumption that a given realisation **K** has been generated by an agent-based process that is RSM temporally stationary. Our procedure is based on the fact that, under RSM temporal stationarity, $\hat{K}_i(r)$ is an approximately unbiased estimator for $K_{\mathbb{1}}(r)$. This suggests testing the hypothesis

$$H_0 : \hat{K}_i(r) = K_{\mathbb{1}}(r) + \varepsilon_i(r), \quad \text{for } r \in [0, R],\ 1 \le i \le N,$$

where the error functions in the estimation $\varepsilon_i(r)$ are mean zero, stationary, and temporally dependent. Hypothesis testing for H_0 is considered in Horváth et al. (2014a) in the context of testing for stationarity with functional time series data. In order to adapt their test to this setting, we must assume a weak dependence condition on the error functions under H_0. Let

$$S_{N,\varepsilon}(x, r) = \frac{1}{\sqrt{N}} \sum_{i=1}^{\lfloor Nx \rfloor} \varepsilon_i(r), \quad \text{for } r \in [0, R],\ 0 \le x \le 1.$$

Assumption 3.1. There exists a sequence of bivariate Gaussian processes $\Gamma_N(r, x)$, $0 \le r \le R$, $0 \le x \le 1$, and $N \ge 1$ satisfying that

$$\sup_{0 \le x \le 1} \int_0^R \left(S_{N,\varepsilon}(x, r) - \Gamma_N(x, r)\right)^2 dr = o_P(1),$$

where $\mathbb{E}[\Gamma_N(x, r)] = 0$, $Cov(\Gamma_N(x, r), \Gamma_N(y, s)) = \min(x, y)c(r, s)$, with $c(r, s) = \sum_{\ell=-\infty}^{\infty} \mathbb{E}[\varepsilon_0(r)\, \varepsilon_\ell(s)]$.

Assumption 3.1 implies that the partial sum process of the error functions under H_0 admits a Gaussian approximation. We note that the limiting sequence of Gaussian processes Γ_N does not depend on N, and so this assumption can be readily used to approximate the distribution of test statistics based on the partial sum process. The function $c(r, s)$ is typically referred to as the long-run covariance function, and incorporates the information on the autocorrelation of the series of K functions.

Assumption 3.1 holds for weakly dependent, stationary function space valued time series in great generality. Specifically, Assumption 3.1 has been shown to hold if we assume a nonlinear dynamical dependence structure as in Wu (2005) and further second order moment conditions (Berkes et al., 2013; Jirak, 2013). Similar results have been shown for mixing sequences of random functions, see Dehling (1983) and Politis and Romano (1994). In the context of empirical K functions computed from agent-based processes, the dynamic and stochastic elements of agent movements or births/deaths imply that the spatial arrangement of agents within an observation region, A, for a particular time point will be dependent on spatial arrangements occurring at previous time points, but that the dependence will dissipate as temporal separation increases. Since $\hat{K}_i(r)$ functions are dependent only on the spatial arrangement of agents, it is a reasonable assumption that the error functions will also be weakly dependent, and hence satisfy Assumption 3.1.

Let

$$S_{N,\hat{K}}(x,r) = \frac{1}{\sqrt{N}} \sum_{i=1}^{\lfloor Nx \rfloor} \hat{K}_i(r), \quad \text{and} \quad S^0_{N,\hat{K}}(x,r) = S_{N,\hat{K}}(x,r) - xS_{N,\hat{K}}(1,r).$$

The bivariate process $S^0_{N,\hat{K}}$ is a function space version of the cumulative sum (CUSUM) process. In order to test H_0, we employ the test statistic

$$P_N = \int_0^R \int_0^1 \left(S^0_{N,\hat{K}}(x,r) \right)^2 dxdr. \tag{3}$$

By utilising the Karhunen–Loéve expansion, it follows under H_0 and Assumption 3.1 that

$$P_N \xrightarrow{\mathcal{D}} \sum_{i=1}^{\infty} \lambda_i \int_0^1 B_i^2(x)dx, \tag{4}$$

where the $B_i, i \geq 1$ are independent and identically distributed Brownian Bridges on the unit interval, and the $\lambda_i,\ i \geq 1$ are the eigenvalues of the Hilbert–Schmidt integral operator with kernel c; namely

$$\lambda_i \phi_i(r) = \int_0^R c(r,s)\phi_i(s)ds, \quad 1 \leq i \leq \infty,$$

with the ϕ_i being orthonormal eigenfunctions. An approximate test of size α of H_0 is obtained by rejecting if $P_N > q_\alpha$, where q_α is the $1-\alpha$ quantile of the distribution on the right hand side of (4).

The limiting distribution is non-pivotal, since it depends on the eigenvalues of the long run covariance operator, which depend on $c(r,s)$. To get around this, we use a Monte Carlo simulation to estimate the limiting quantile q_α from the empirical eigenvalues $\hat{\lambda}_i$ based on an estimator $\hat{c}$ of c. The details are outlined in Appendix A.

In the case when the agent-based process is not RSM temporally stationary, we have $K(r,t)$ varying with t for $t \in [t_1, t_N]$, and

$$H_A:\ \hat{K}_i(r) = K(r,t_i) + \varepsilon_i(r), \quad \text{for } r \in [0,R],\ 1 \leq i \leq N,$$

holds. In the case when $K(r,t_1) = K(r,t_2) = \cdots = K(r,t_i^*) \neq K(r,t_{i+1}^*) = \cdots = K(r,t_N)$ for some time point t_i^*, i.e. there exists a change point in the sequence $K(r,t_i)$, then $P_N \to \infty$ in probability. As shown in Horváth et al. (2014a), $P_N \to \infty$ in probability under H_A for even mild conditions on the nature of the change.

4. Analysis of $\bar{K}(r)$ for a single realisation of an agent-based process

Working under the assumption of RSM temporal stationarity, an estimate for $K_{\mathcal{I}}(r)$ may be obtained by averaging over all functions in the functional time series:

$$\bar{K}(r) = \frac{1}{N} \sum_{i=0}^{N} \hat{K}_i(r).$$

A problem of interest is to determine the nature, or even existence, of interaction between agents. To study this, the estimate $\bar{K}(r)$ can be compared to a given K function of interest. For example, a sensible K function to compare to to check for completely random motion of the agents in a two dimensional space is $K(r) = \pi r^2$. The function for comparison could also be produced by a separate simulation, for example by estimating the K function from a large number of independent simulations of a Strauss process, since the theoretical K function for such a process has no known analytic form. Making such comparisons over different ranges of the radius r allows for detailed inference about the scales at which the agent-based process tends to cluster or disperse.

Specifically, if the functional time series $\hat{\mathbf{K}}$ is stationary, and satisfies 3.1, then we can test the hypotheses

$$H_0 : K_{\mathbb{1}}(r) = f(r) \quad \text{vs.} \quad H_A : K_{\mathbb{1}}(r) \neq f(r),$$

for $r_1 \leq r \leq r_2$, for an arbitrary function $f(r)$.

We propose to test H_0 using the test statistic

$$Q_N(r_1, r_2) = \int_{r_1}^{r_2} \left(S_{N,\hat{K}}(1, r) - \frac{1}{\sqrt{N}} f(r) \right)^2 dr.$$

$Q_N(r_1, r_2)$ defines an analog of the one sample "T-test" statistic for functional time series over the range r_1 to r_2. Under Assumption 3.1, one can show that

$$Q_N(r_1, r_2) \xrightarrow{\mathcal{D}} \sum_{i=1}^{\infty} \lambda_i(r_1, r_2) Z_i^2, \tag{5}$$

where Z_i are i.i.d. standard normal variables, and the eigenvalues $\lambda_i(r_1, r_2)$ satisfy

$$\lambda_i(r_1, r_2)\phi_{r_1,r_2}(r) = \int_{r_1}^{r_2} c(r, s)\phi_{r_1,r_2}(s)ds,$$

with orthonormal eigenfunctions ϕ_{r_1,r_2} defined on $[r_1, r_2]$. Again, since the limiting distribution of $Q_N(r_1, r_2)$ is non-pivotal, we perform Monte Carlo simulations to estimate the asymptotic distribution from estimates of the eigenvalues $\hat{\lambda}_i(r_1, r_2)$, with these estimated in the manner described in Appendix A. Under H_A, $Q_N(r_1, r_2) \to \infty$ in probability at the rate of $\sqrt{N}$.

Due to the fact that agent–agent interactions at different scales result in deviations in $K_{\mathbb{1}}(r)$ from πr^2 at different ranges of r it is of interest to test for equality of $\bar{K}(r)$ and πr^2 for specific ranges of r. Similar considerations are important in spatial statistics (Loosmore and Ford, 2006; Baddeley et al., 2014).

5. Simulation study of the 2D jump agent-based model

In the following section we use the approaches outlined above to analyse simulated data from a specific dynamic agent-based model. For our purposes, we adapt the one-dimensional agent-based simulation model presented in Binny et al. (2015) to two dimensional space. We will refer to the agent-based process presented here as the *2D jump model*. This simulation model is of interest with respect to modelling cell movements.

The model consists of M individuals positioned on a two-dimensional continuous rectangular domain with periodic boundary conditions. Let x_l, and x_r give the left and right x boundaries respectively, and y_L, and y_U give the lower and upper y boundaries respectively. We refer to this region as the *simulation region*. We then define a rectangular observation region, A, with left, right, lower, and upper boundaries given by x_l^*, x_r^*, y_L^*, and y_U^*, respectively, such that we only observe the agents positioned such that their x and y values satisfy $x_l^* \leq x \leq x_r^*$ and $y_L^* \leq x \leq y_U^*$.

The model is a continuous time Markov process, with the state of the system at time t given by $\mathbf{x}(t) = \{(\mathbf{x}_1(t), \mathbf{x}_2(t), \ldots, \mathbf{x}_M(t))\}$, with $\mathbf{x}_i(t) = (x_i(t), y_i(t))$ giving the position of the ith agent in the simulation region.

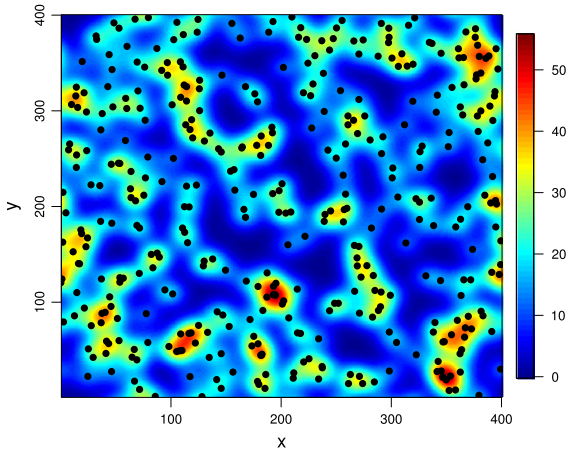

Fig. 2. Plot of $\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})$, with $\beta = 10$ and $\sigma^2 = 10$, for a $[0, 400] \times [0, 400]$ simulation region with periodic boundary conditions. The vertical colour scale to the right of the plot indicates which colours correspond to which values of $\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})$. The black dots indicate positions of agents. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The movement event $\mathbf{x}_i \mapsto \mathbf{x}_i + \mathbf{s}$, for agent i, occurs with the rate $m\mu(\mathbf{x}_i, \mathbf{x}_i + \mathbf{s})$, with m giving the movement rate, and $\mu(\mathbf{x}_i, \mathbf{y})$ giving the probability density function for agent i moving from $\mathbf{x}_i$ to position $\mathbf{y}$. The process is simulated using the Gillespie algorithm (Gillespie, 1976).

To define how neighbouring agents bias movements of other agents, we define a bias function $\mathbf{b}(\mathbf{x})$ by

$$\mathbf{b}(\mathbf{x}) = \nabla\left(\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})\right),$$

with $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})$, and $v(\mathbf{z}) = \beta \ \exp\left(-\frac{|\mathbf{z}|^2}{2\sigma^2}\right)$. As such, $\mathbf{b}(\mathbf{x})$ is defined as the gradient of the bivariate function $\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})$. The function $\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})$ produces an undulating 2D surface, with the gradient of the surface being determined by the locations of agents, and the sign of parameter β. See Fig. 2 for an example of $\sum_{j=1}^{n} v(\mathbf{x}_j - \mathbf{x})$. Agents are biased to move down the gradient of this surface. For $\beta > 0$ agents are biased to move away from one another, and for $\beta < 0$ agents are biased to move towards one another. Parameter σ^2 indicates the range at which this interaction is strong, with range of interaction increasing with σ^2.

Using this bias function, we define the probability density function for an agent at $\mathbf{x}$ moving in direction θ as

$$p(\theta, \mathbf{x}) = \frac{1}{\pi\left(1 + e^{-\mathbf{b}(\mathbf{x})\cdot(\cos(\theta), \sin(\theta))}\right)}, \quad 0 \le \theta < 2\pi,$$

with $\mathbf{b}(\mathbf{x}) \cdot (\cos(\theta), \sin(\theta))$ being the dot product between these vectors.

Similarly to Binny et al. (2015), we assume that the distance moved by an agent is independent of the direction it moves. We utilise the same distribution for the distance moved as in Binny et al. (2015), an exponential distribution with mean $1/\rho$.

Consequently, we can express the movement density function as $\mu(\mathbf{x}, d, \theta)$, giving the density function for an agent at position $\mathbf{x}$ moving distance d at an angle of θ from its current position:

$$\mu(\mathbf{x}, d, \theta) = \rho\, e^{-\rho d} p(\theta, \mathbf{x}), \quad 0 \le d, \text{ and } 0 \le \theta < 2\pi.$$

We applied both of the tests outlined in Sections 3 and 4 to positional data obtained from simulations of the 2D jump model. For all simulations we set $x_l = 0, x_r = 400, y_L = 0$, and $y_U = 400$,

Table 1
The empirical rejection rates for nominal levels 10%, 5%, and 1% for testing stationarity for the 2D jump process with no interaction. R gives the maximum r used for each test. M gives the number of agents. T defines the time period of simulations.

Parameters			Nominal levels		
R	M	T	0.1	0.05	0.01
25	100	50	0.114	0.038	0
		100	0.123	0.051	0.003
		250	0.113	0.049	0.006
		500	0.11	0.054	0.008
	200	50	0.122	0.045	0.003
		100	0.1	0.043	0.004
		250	0.118	0.045	0.008
		500	0.115	0.05	0.009
	400	50	0.097	0.039	0
		100	0.114	0.041	0.002
		250	0.096	0.036	0.002
		500	0.1	0.056	0.006
50	100	50	0.113	0.043	0.002
		100	0.12	0.048	0.003
		250	0.109	0.053	0.005
		500	0.108	0.049	0.004
	200	50	0.13	0.056	0
		100	0.097	0.035	0.001
		250	0.107	0.052	0.005
		500	0.114	0.056	0.008
	400	50	0.131	0.038	0
		100	0.122	0.04	0.002
		250	0.1	0.04	0.006
		500	0.111	0.059	0.005

with the parameters of the observation region given by $x_l^* = 100$, $x_r^* = 300$, $y_L^* = 100$, and $y_U^* = 300$, such that the observation region is given by $[100, 300] \times [100, 300]$, 1/4 of the total simulation region.

We considered the cases for no interaction, $\beta = 0$, attraction, $\beta < 0$, and repulsion, $\beta > 0$. We set $\rho = 1/10$, and $m = 10$ for all simulations. For interacting processes we use $\sigma^2 = 10$. This selection of parameters is designed to allow for an effective illustration of the approaches outlined in this paper.

Each simulation is repeated 1000 times. The empirical sizes of each test are reported in Tables 1–4 for nominal sizes of 10%, 5%, and 1%.

5.0.1. Analysis for $\beta = 0$

We begin analysis with the case $\beta = 0$, i.e. no interaction. To demonstrate the effect of the total number M of agents in the simulation we compare the statistical tests for $M = 100, 200, 400$. We let these agents be distributed randomly within the simulation region, $[0, 400] \times [0, 400]$, with uniform distribution. Note that since the observation region is $[100, 300] \times [100, 300]$ the expected number of observed agents for any time point will be $M/4$.

The lack of interaction between agents implies that $K(r; t) = \pi r^2$ for all t. As such, the non-interacting 2D jump process is RSM temporally stationary for all $t \geq 0$.

For the case $\beta = 0$, we wish to determine whether the statistical tests for stationarity, and equality of the mean K function to πr^2, behave as expected. Let the observation periods considered be $[0, T]$, with $T = 50, 100, 250, 500$. We let $t_{i+1} - t_i = 0.5$ be the length of a time step for all observations. Consequently, the number of time points, and hence functions, in the time series for each T value is given by $N = 2T + 1$, since we take $t_1 = 0$. We produce 1000 independent simulations for each combination of M and T.

Both tests must be performed over a specified range of r values. We perform the tests over the ranges $[0, R]$, with $R = 25, 50$.

Table 2
The empirical rejection rates for nominal levels 10%, 5%, and 1% for testing $\bar{K}(r) = \pi r^2$ for the 2D jump process with no agent–agent interactions. R gives the maximum r used for each test. M gives the number of agents. T defines the time period of simulations.

Parameters			Nominal levels		
R	M	T	0.1	0.05	0.01
25	100	50	0.138	0.081	0.025
		100	0.128	0.066	0.011
		250	0.119	0.06	0.011
		500	0.105	0.054	0.005
	200	50	0.159	0.085	0.031
		100	0.139	0.088	0.028
		250	0.106	0.052	0.007
		500	0.103	0.057	0.014
	400	50	0.129	0.075	0.034
		100	0.122	0.065	0.021
		250	0.102	0.056	0.012
		500	0.103	0.05	0.012
50	100	50	0.176	0.107	0.044
		100	0.124	0.074	0.021
		250	0.124	0.068	0.022
		500	0.108	0.057	0.014
	200	50	0.166	0.097	0.044
		100	0.139	0.085	0.031
		250	0.114	0.068	0.015
		500	0.118	0.062	0.016
	400	50	0.172	0.105	0.036
		100	0.136	0.079	0.025
		250	0.12	0.056	0.013
		500	0.12	0.063	0.017

Testing for stationarity: The empirical rejection rates when the nominal levels are 10%, 5%, and 1% are given in Table 1. We can see from these results that the statistical test for stationarity performs well for all combinations of R, M, and T.

Testing for interaction: To test for agent–agent interactions we test each of the simulations produced for equality between $\bar{K}(r)$ and πr^2. The empirical rejection rates when the nominal levels are 10%, 5%, and 1% are given in Table 2. From these results we can see that the test for equality between $\bar{K}(r)$ and πr^2 performs relatively well for the different combinations of R, M, and T, with the test performing better for larger T values, and better for the smaller R value.

We summarise the above results of the simulations as follows: When analysing relatively small time series, e.g. $N = 50, 100$, it is advisable to either (a) use a smaller R value than might otherwise be used when testing for $\bar{K}(r) = \pi r^2$, or (b) appreciate that this test, for small N, is biased towards rejection of the null hypothesis, and thus to reduce any nominal level used for rejection.

5.0.2. Analysis for cases $\beta \neq 0$

We continue with analysis of simulations of the 2D jump model for cases $\beta \neq 0$. While the non-interacting 2D jump process was known to be RSM stationary, with $K_{\mathbb{1}}(r) = \pi r^2$, it is not known whether this will be the case for any particular non-zero β. However, we expect it to be the case for range of β values selected for analysis, which typically result in weak to moderate dispersion or clustering.

We compare several different β values, $\beta = 10, 5, -2.5, -3$, with $\sigma^2 = 10$ for all simulations. The observation period for all simulations is [0, 125], since, as discussed below, we remove a burn-in period from this total observation period prior to analysis. As for the non-interacting case we populate the simulation regions with the M agents distributed randomly with a uniform distribution. We compare results for $M = 100, 200, 400$.

Table 3
The empirical rejection rates for nominal levels 10%, 5%, and 1% for testing stationarity for the 2D jump process with agent–agent interactions for the observation period of [25, 125]. R gives the maximum r used for each test. β defines the strength and type of agents. M gives the number of agents.

Parameters			Nominal levels		
R	β	M	0.1	0.05	0.01
25	10	100	0.112	0.035	0.003
		200	0.12	0.043	0
		400	0.116	0.04	0.001
	5	100	0.114	0.045	0.001
		200	0.105	0.046	0.004
		400	0.111	0.044	0.005
	−2.5	100	0.115	0.046	0.004
		200	0.114	0.043	0.003
		400	0.105	0.036	0.003
	−3	100	0.123	0.049	0.004
		200	0.115	0.048	0.002
		400	0.11	0.044	0.004
50	10	100	0.099	0.04	0.001
		200	0.108	0.044	0.001
		400	0.105	0.032	0
	5	100	0.118	0.053	0.004
		200	0.102	0.039	0.001
		400	0.118	0.05	0.004
	−2.5	100	0.105	0.033	0.001
		200	0.119	0.041	0.001
		400	0.11	0.047	0
	−3	100	0.128	0.061	0.002
		200	0.11	0.039	0.002
		400	0.126	0.048	0.003

Now, if a time invariant $K_{\mathcal{I}}(r)$ function exists for a given $\beta \neq 0$ and $K_{\mathcal{I}}(r) \neq \pi r^2$, then the agent–agent interactions are expected to drive the agents away from the initial uniform arrangement towards an arrangement corresponding with the associated $K_{\mathcal{I}}(r)$ function, i.e. towards equilibrium. If this is the case then there may be an initial period of non-RSM temporal stationarity before the process becomes RSM temporally stationary. We are interested in analysing the RSM temporally stationary portion. To this end we remove an initial segment, or burn-in period, of the observations from the sample. Specifically, the burn-in period of $[0, 25)$ was removed from sample over all simulations, since the vast majority of simulations for each β appear to have reached equilibrium after $t = 25$. This leaves the final observation period of $[25, 125]$.

Testing for stationarity: The empirical rejection rates when the nominal levels are 10%, 5%, and 1% when testing for stationarity for the different β values are given in Table 3. We can see from these results that for the observation period of $[25, 125]$, the empirical rejection rates are relatively comparable to those for the non-interacting process.

We expected that the 2D jump processes for these β values would be RSM temporally stationary, and these rejection rates suggest this is the case, confirming that the test for RSM temporal stationarity performs well for interacting agent-based processes as well as non-interacting processes as observed above.

Testing for interaction: The empirical rejection rates when the nominal levels are 10%, 5%, and 1% when testing for $\bar{K}(r) = \pi r^2$ for the different β values are given in Table 4. These values demonstrate that increasing M, and thus the number of agents observed at each time point, for a given β value significantly increases the rejection rates at all levels. We also note that rejection rates are higher for smaller R values, reflecting the fact that agent–agent interactions are strongest over small distances.

Table 4
The empirical rejection rates for nominal levels 10%, 5%, and 1% for testing $\bar{K}(r) = \pi r^2$ for the 2D jump process with agent–agent interactions for the observation period of $[25, 125]$. R gives the maximum r used for each test. β defines the strength and type of agents. M gives the number of agents.

Parameters			Nominal levels		
R	β	M	0.1	0.05	0.01
25	10	100	1	1	0.999
		200	1	1	1
		400	1	1	1
	5	100	0.946	0.905	0.773
		200	1	1	1
		400	1	1	1
	−2.5	100	0.536	0.402	0.198
		200	0.982	0.949	0.831
		400	1	1	1
	−3	100	0.674	0.552	0.299
		200	0.998	0.986	0.93
		400	1	1	1
50	10	100	0.955	0.913	0.762
		200	1	1	1
		400	1	1	1
	5	100	0.55	0.437	0.259
		200	0.97	0.942	0.795
		400	1	1	0.999
	−2.5	100	0.216	0.132	0.039
		200	0.591	0.428	0.203
		400	0.998	0.969	0.834
	−3	100	0.309	0.172	0.068
		200	0.732	0.584	0.321
		400	0.999	0.996	0.948

An estimate, $\bar{g}(r)$, for the pair correlation function can be produced via smoothing and transforming the estimate $\bar{K}(r)$ (Illian et al., 2008; Chiu et al., 2013; Baddeley et al., 2015). Examples of estimated $\bar{K}(r)$ and $\bar{g}(r)$ functions for single realisations of 2D jump processes for $\beta \neq 0, M = 400$, are presented in Figs. 1 and 2 in the online supplementary material.

6. Simulation study of point patterns generated by the Metropolis–Hastings algorithm

The 2D jump model considered in the previous section evolves in continuous time via agents performing movements via jumps. Agent positions are observed at discrete times. As mentioned in Section 2, agent-based processes can also evolve via births and deaths of agents. We here demonstrate the applicability of the statistical methods presented above to birth–death processes. In particular, we consider an application of the proposed methodology to dependent sequences of point patterns that arise as iterations of the Metropolis–Hastings (M–H) algorithm when used to simulate Gibbs processes. The M–H algorithm is explained in detail in Chapter 13 of Baddeley et al. (2015), and, in brief, may be used to simulate spatial point processes from a target distribution by describing and simulating a discrete-time Markov Chain whose states are spatial point patterns, and whose transitions involve adding, removing, or shifting points. The target point process is described by a density function, which is used to assign a probability of whether a proposed change to the point process should be accepted or rejected. Depending on the clustering behaviour and intensity of the target process, it can take anywhere from several thousand to hundreds of thousands of iterations to achieve acceptable convergence to the target distribution.

Our methodology may be used to analyse the point patterns that result from this procedure. Formally, let $\mathbf{X}_1, \ldots, \mathbf{X}_N$ denote the sequence of point patterns that result from applying the M–H

Table 5
The empirical rejection rates for nominal levels 10%, 5%, and 1% for testing stationarity and $\bar{K}(r) = \pi r^2$ for M–H Strauss model with $\beta = 2$, $r_{int} = 0.7$, $p = 0$, and $q = 0.5$.

γ	Test	Nominal levels		
		10%	5%	1%
1	Stationarity	0.101	0.049	0.006
	$\bar{K}(r) = \pi r^2$	0.099	0.049	0.007
0.7	Stationarity	0.098	0.04	0.002
	$\bar{K}(r) = \pi r^2$	1	1	1
0.2	Stationarity	0.091	0.047	0.006
	$\bar{K}(r) = \pi r^2$	1	1	1

algorithm to simulate a specified point process with density $f(\mathbf{x})$, starting from a given initial point pattern (perhaps a uniform distribution of a certain number of points). In order to control the dependence between K functions computed from these point patterns, we only consider every wth pattern, which results in a sample of point patterns of length N/w. The specific form of the M–H algorithm that we used is described on page 533 of Baddeley et al. (2015), which involves the choice of parameters p and $q/1-q$, defining the probabilities of proposing a point shift and, given a shift is not proposed, a death/birth, respectively.

6.1. Application to point patterns simulating a Strauss process

We first consider point patterns that arise during the simulation of the well known Strauss point process (Strauss, 1975; Kelly and Ripley, 1976):

$$f(\mathbf{x}) = \alpha \beta^{|\mathbf{x}|} \gamma^{s(\mathbf{x})}, \tag{6}$$

where $s(\mathbf{x}) = \sum_{x,y \in \mathbf{x}} \mathbb{1}(\|x - y\| \leq r_{int})$, α is a normalising constant, $\beta > 0$ is the 'abundance' parameter associated with intensity of agents per unit area, $0 \leq \gamma \leq 1$ is the interaction parameter, r_{int} is the range of interaction, and $|\mathbf{x}|$ gives the number of agents in $\mathbf{x}$. The Strauss point process models inhibition, with $\gamma < 1$ resulting in "repulsion" between points (for $\gamma = 1$ there is no interaction. $\gamma = 0$ results in a hard core process) (Baddeley et al., 2015). We are interested in studying whether the proposed methodology can detect dispersion, and if the sequence of K functions computed from successive point patterns appears stationary.

We set the number of agents in the initial point pattern, $\mathbf{X}_1$, to be $M_1 = 100$, and the observation area to be $[0, 10] \times [0, 10]$. Setting $\beta = 2$, $r_{int} = 0.7$ in (6), we perform three sets of 1000 independent simulations with $\gamma = 1, 0.7$, and 0.2 for each set respectively. We note that for $\gamma = 1$ (6) reduces to the density of a homogeneous Poisson point process, i.e. there are no interactions between agents (Baddeley et al., 2015). We set the total number of iterations N for each simulation to be 100,000, with $w = 100$, so that each 100th iteration of the Metropolis–Hastings algorithm was retained to produce a sample of size 1000. We set the simulation probabilities to $p = 0$ and $q = 0.5$. We refer to this agent-based model as the M–H Strauss model.

For similar reasons as for the 2D jump model with interaction we removed a burn-in period from each simulation before performing the statistical tests on the remaining data. We removed the first 500 point patterns to constitute a burn-in period, leaving 500 patterns for analysis for each simulation.

The empirical rejection rates when the nominal levels are 10%, 5%, and 1% are given in Table 5 for running both the stationarity test and the test of significance for the mean K function being equal to πr^2. Since we removed a significant burn-in period from the simulated data before performing the statistical tests, we expected that data produced for the M–H Strauss model would indicate RSM stationarity. We can see from the results in Table 5 that this appears to be the case for each value of γ considered, and that the statistical test for stationarity performs well for this model. The rejection rates with respect to the nominal levels of 10%, 5%, and 1% for testing if $\bar{K}(r)$ is significantly different from πr^2 for no interaction, $\gamma = 1$, behave as expected. When the process exhibited dispersion, i.e.

for $\gamma = 0.7, 0.2$, a significant difference between the mean K function and πr^2 was found in every associated simulation and significance level considered.

6.2. Investigating convergence of the M–H algorithm for simulating the Geyer saturation process

Additionally we consider an application of our methodology to investigate the number of iterations required to achieve convergence when using the M–H algorithm. A criterion that seems reasonable for assessing if convergence has been achieved when simulating a Gibbs process is to monitor whether a corresponding sequence of second moment measures appears to have stabilised. This may be achieved by applying the proposed K-function stationarity test to rolling windows of a fixed size to the sequence of spatial point patterns, which we illustrate here with a short application to the Geyer saturation processes.

For these simulations, we used the probability density associated with the Geyer saturation process (Geyer, 1999):

$$f(\mathbf{x}) = \alpha\beta^{|\mathbf{x}|}\prod_{i=1}^{|\mathbf{x}|}\gamma^{\min(s,t(x_i,r_{int},\mathbf{x}))}, \tag{7}$$

where α is a normalising constant, β, γ, r_{int}, and s are parameters, and $t(x_i, r_{int}, \mathbf{x})$ is the number of other data points x_j from $\mathbf{x}$ lying within a distance r_{int} of the point x_i. The effect of these parameters for the Geyer saturation point process is explained elsewhere (Geyer, 1999; Baddeley et al., 2015).

Similarly to our application to the Strauss process, we set the initial number of agents to $M_1 = 100$, and let the observation area be $[0, 10] \times [0, 10]$, with the initial agents uniformly distributed in the observation area. We compared the results from changing the value of p while keeping the other parameters constant, performing 40 independent simulations for each set of parameters. We set $\beta = 0.4$, $\gamma = 1.25$, $r_{int} = 1$, $s = 10$, and $q = 0.9$, and compared the results for $p = 0.99, 0.975, 0.95, 0.9$. We set the total number of iterations N for each simulation to be 1,000,000. We again set $w = 100$, producing a sample size of 10,000. We did not remove a burn-in period. We then broke up each sample into "windows" of length 500, giving 20 total windows per simulation, and computed p-values of the stationarity test applied to each window. We refer to this agent-based model as the M–H Geyer saturation model.

We expected that since the agents were initially uniformly distributed the M–H algorithm would, over the initial windows, tend to change the point patterns fairly dramatically, and this would manifest in the p-values computed from the initial windows being quite small. As convergence is achieved, the changes in the point processes become much more gradual, and we expected the p-values to become essentially uniformly distributed on (0, 1) for later windows.

The results of the tests for each set of simulations are presented in Fig. 3 in terms of heat maps of the resulting p-values. We can observe in Fig. 3 that for each set of simulations there is an initial band of small p-values occurring for at least the first window, and that p-values for the remaining windows demonstrate a relatively uniform distribution. This suggests that for each set of simulations there is an initial period of non-RSM temporal stationarity, after which the process reaches RSM temporal stationarity. We observed in Fig. 3 that the width of the bands of small p-values decreased with decreasing p, i.e. decreasing the probability of proposing a shift appeared to decrease the number of iterations taken for the process to reach RSM temporal stationarity. Other data, not presented here, suggests that there may be a complicated relationship between parameter values and the length of the non-RSM temporal stationarity for this process.

7. Data application: wingless locust nymph movement data

Prior to developing into winged adults and forming large flying swarms, wingless locust nymphs (juveniles) aggregate into large migratory locust marching bands. The Australian locust *Chortoicetes terminifera* forms large crescent shaped bands that can spread over kilometres and contain millions of locusts (Buhl et al., 2011). Such marching bands maintain their global structure, such that minimal

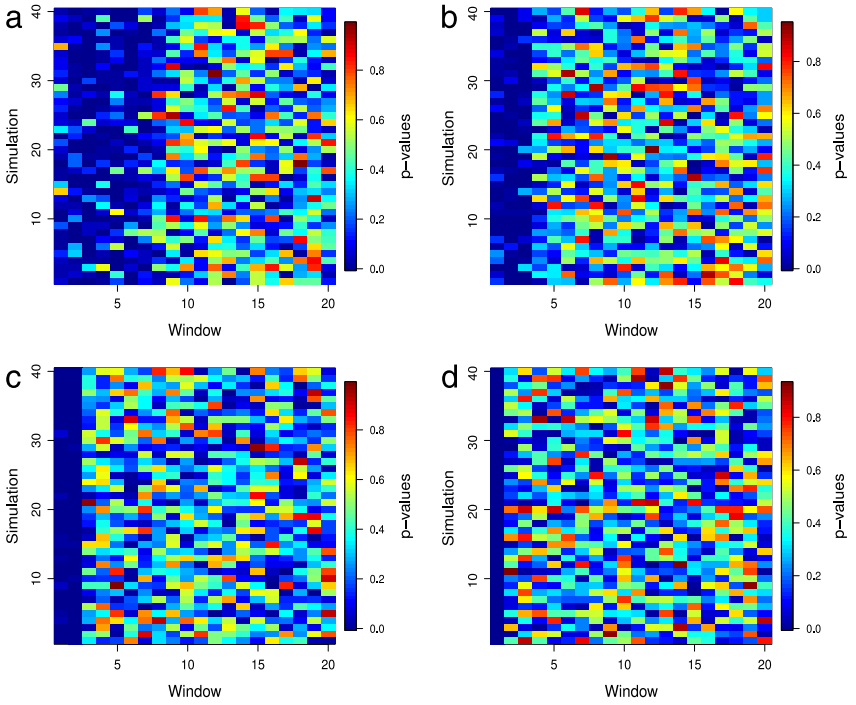

Fig. 3. Plot of p-values for each window associated with each simulation of the M–H Geyer saturation model. The vertical colour scale to the right of the plot indicates which colours correspond to which p-value. The following parameters are identical for each figure: $\beta = 0.4$, $\gamma = 1.25$, $r_{int} = 1$, $s = 10$, and $q = 0.5$. The value of p changes for each figure: (a) $p = 0.99$; (b) $p = 0.975$; (c) $p = 0.95$; (d) $p = 0.9$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

dispersion of the locusts occurs during migration. The interactions between individual locust nymphs are of considerable interest to researchers, since these local interactions lead to emergent cohesive global behaviour.

Here we apply the approaches presented above to a single set of experimental data recording the movements of individual *C. terminifera* locusts in a locust marching band. The data analysed here represents a single experiment of the collection presented in Buhl et al. (2011, 2012).

Data was collected during field experiments described in detail in Buhl et al. (2011). In brief, footage of migrating bands of juvenile *C. terminifera* were recorded via a tripod-mounted 1080p camera filming vertically overhead. One frame from every 30 s was then analysed, with the middle of each locust being recorded as the locust's position (locust nymphs are approximately 7 mm in length and a few mm in width). As discussed in Buhl et al. (2011), the recording area was flat, smooth, without vegetation and raked to ensure homogeneity, with the recording area being of significant distance from obstacles and landmarks that might otherwise influence the homogeneity. Buhl et al. (2012) found that the process was isotropic. While there are no existing tests for spatial homogeneity of agent-based processes, visual inspection of the data suggests that these attempts were successful and that the observed process was spatially homogeneous.

Buhl et al. (2011) describe how the density of locusts, analogous to the first spatial moment, is higher at the front of migrating bands and lower at the rear. As such, the global structure of locusts is spatially inhomogeneous. However, the local homogeneity observed in the observation region during the observation period ensures that our approach for analysis is valid, i.e. we can consider the observed process within the fixed observation region to be a spatially homogeneous process with time dependent $\lambda(t)$.

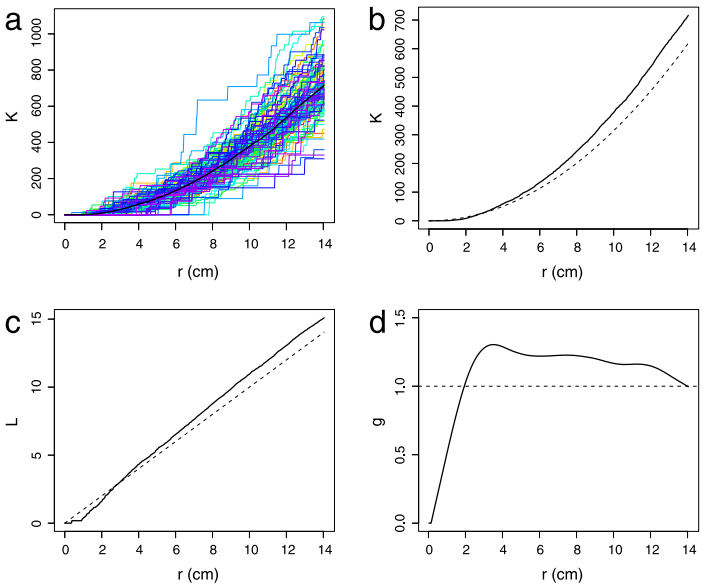

Fig. 4. Empirical functions produced from a field experiment recording wingless locust nymph movements presented in Buhl et al. (2012) (a) The rainbow plot for the functional time series $\hat{\mathbf{K}}$, with $\bar{K}(r)$ (———); (b) $\bar{K}(r)$ (———), $K(r) = \pi r^2$ (– – –); (c) $\bar{L}(r)$ (———), $L(r) = r$ (– – –); (d) $\bar{g}(r)$ (———), $g(r) = 1$ (– – –). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Furthermore, we note that the movements of the locusts are biased in the direction of the movement of the overall locust marching band, though this has no impact on our methods for analysis due to the maintenance of spatial homogeneity throughout the observation period.

For the specific data set being analysed here the observation region was a 0.583 m^2 rectangle (103.8 cm $\times$ 56.2 cm), with 114 time points, each 30 s apart. The minimum and maximum number of locusts observed at a particular time point was 11 and 52 respectively, with the mean number of locusts observed being 29. We note that we removed the first time step from our analysis as this appeared to be the time point at which the first locusts entered the visualisation area and thus was deemed to be inhomogeneous.

An animation indicating the positions of locusts at each time point and the corresponding $\hat{K}_i(r)$ functions for each time point is provided in the online supplementary material (see Fig. 1 for the associated details). Fig. 4 presents the estimates for $\bar{K}(r)$ and $\bar{g}(r)$, as well as the rainbow plot for this functional data and the estimated L function. A rainbow plot of functional data is a useful means of visualising a functional time series, with earlier functions shown in red, followed by orange, yellow, green, blue and indigo with the last few functions plotted in violet (Hyndman and Shang, 2010; Shang and Hyndman, 2015). The L function is a useful transformation of the K function commonly used in spatial statistics that makes visualisation easier (Diggle, 2003; Illian et al., 2008). Specifically, $L(r) = \sqrt{K(r)/\pi}$, so that the theoretical $L(r)$ function associated with a non-interacting process is $L(r) = r$.

We continued by performing the statistical tests for stationarity. We set R, the maximum value of r analysed, to be 1/4 of the smallest edge of the observation region; $R = 14.04$ cm. The p-value for the stationarity for $r \in [0, 14.04]$ was 0.50. The large p-value suggests that the data is stationary, and therefore we continue under the assumption that the underlying process is RSM temporally stationary. As such, we are justified in continuing to test for equality between $\bar{K}(r)$ and πr^2.

Analysis in Buhl et al. (2011, 2012) suggests that locusts interact differently at two different inter-locust distance scales. First, locusts appear to be have repulsive interactions for inter-locust distances up to some distance likely between 2 and 3 cm. Secondly, locusts appear to have attractive interactions

at inter-locust distances between approximately 3 cm and 13.5 cm, though the maximum distance at which interactions occur is unresolved (Buhl et al., 2012). Based on this prior knowledge of these potential ranges of interaction, and the cumulative nature of the K function, we would expect that $\bar{K}(r) < \pi r^2$ for approximately $r \in [0, 3]$ and $\bar{K}(r) > \pi r^2$ for approximately $r \in [3, 13.5]$. As such we decided, prior to observing the estimate for $\bar{K}(r)$, to test for $\bar{K}(r) = \pi r^2$ over $r \in [0, 3]$ and $r \in [3, 13.5]$. The p-values obtained when testing for $\bar{K}(r) = \pi r^2$ over these ranges were 0.009 for $r \in [0, 3]$, and 0 for $r \in [3, 13.5]$, indicating that the null hypothesis of $\bar{K}(r) = \pi r^2$ could be strongly rejected over these ranges of r.

As seen in images (b) and (c) in Fig. 4, we have $\bar{K}(r) < \pi r^2$ and $\bar{L}(r) < r$ for $r \in [0, 2.77]$, and $\bar{K}(r) > \pi r^2$ and $\bar{L}(r) > r$ for $r \in [2.77, 14.04]$. This matches our expectations for $\bar{K}(r)$ given our prior knowledge regarding locust behaviour, and combined with the above p-values strongly indicates that locusts are dispersed over the range of $r \in [0, 2.77]$ and clustered over the range of (at least) $[2.77, 13.5]$.

We can utilise the estimated pair correlation function in image (d) in Fig. 4 to further characterise the agent–agent interactions present between locusts. Interpretation of pair correlation functions has been extensively studied for spatial statistics (Illian et al., 2008) and is an active area of research for agent-based processes (Buhl et al., 2012; Agnew et al., 2014). Interpretation of pair correlation function functions typically requires some understanding of the underlying process, and can be aided by the use of simulations (Buhl et al., 2012; Agnew et al., 2014). We utilise the work of Buhl et al. (2012) in aiding our interpretation. Since the centres of two locusts cannot occur closer than the width of a locust, i.e. a few millimetres, the process can be described as a *hard core* process (Illian et al., 2008). As such we expect $\bar{g}(r) = 0$ for very small r values, which is evident in Fig. 4. The peak of $\bar{g}(r)$ occurs around $r \approx 3.5$ cm suggesting that, in addition to the space occupying restrictions on locust positions, agent–agent interactions for locusts are repulsive for inter-agent distances below 3.5 cm with attractive interactions occurring at larger distances. As r approaches 14.0 cm, $\bar{g}(r)$ approaches 1, suggesting that the maximum range of attraction may occur around this distance. As mentioned previously, such descriptions of interactions match well with existing evidence regarding locust interactions (Buhl et al., 2012), though we note the presence of quantitative differences arising from methodological differences in how the correlation function is calculated.

8. Discussion

In this paper we have presented a framework for the statistical analysis of spatially homogeneous dynamic agent-based processes. Our approach utilises methods from both spatial statistics and functional time series. The combination of methods from these areas for the analysis of dynamic agent-based processes has not been explored previously and represents a novel method for analysis. The significant contributions of our approach are the addition of statistical rigour to the analysis of agent-based processes and the added efficiency for the analysis of such processes.

By transforming spatial data arising from agent-based processes into longitudinal series of functional data, we are able to perform important statistical tests using newly developed statistical methods from the functional time series literature. By doing so, we can make important inferences about the underlying agent-based processes. With appropriate assumptions being met, testing for stationarity of the functional series supplies evidence of reduced second moment (RSM) temporal stationarity of the underlying agent-based process.

If an observed agent-based process is RSM temporally stationary we can then pool the functional data from an observed realisation to efficiently produce an accurate estimate for the time invariant $K_I(r)$ function, $\bar{K}(r)$. In contrast, this cannot be done for a process that is not RSM temporally stationary. For an RSM temporally stationary agent-based process we can perform a test for equality between $\bar{K}(r)$ and the theoretical K function associated with a non-interacting agent-based process, πr^2 for $\mathbb{R}^2$. This provides statistical evidence for whether or not interaction exists over different spatial ranges. Finally, both the estimated $\bar{K}(r)$ and $\bar{g}(r)$ functions can be utilised to characterise interactions for a given process.

We employed these methods for several Monte Carlo simulation studies. Specifically, we introduced the 2D jump model, and agent-based models derived from the Metropolis–Hastings

algorithm. As such, different types of processes, strengths and types of interactions, and time lengths were analysed and the results matched the asymptotic theory of the tests, validating our approach.

To demonstrate the application of our techniques to real world data we examined experimental wingless locust nymph movement data. Statistical testing suggests that the process is RSM temporally stationary and that agent–agent interactions were present. The interpretation of locust interactions from the estimates $\bar{K}(r)$ and $\bar{g}(r)$ matches existing descriptions of locust nymph behaviour.

Observations made during this work point to several avenues of further work. Specifically, we noted that the dependence structure of the spatial arrangement of agents and, thus, the dependence structure of the empirical K functions influenced the performance of the statistical tests employed. The relationship between such dependence and the appropriate bandwidth selection to use for estimating the long run covariance for functional time series is an open area of research and warrants further examination.

Finally, several other extensions of the work presented here exist. Our approaches could be developed for the analysis of agent-based processes with multiple types of agents (i.e. multi-species systems) which often are of interest in biological and ecological research. Also, adapting this approach for inhomogeneous environments would open up analysis to a much wider range of real world scenarios. Adapting the approaches presented here for the analysis of multiple realisations of dynamic agent-based processes would be another important extension that would allow for these methods to be employed in areas in which conducting multiple experiments or observations is feasible.

Acknowledgements

We are thankful to J. Buhl, S. Simpson, and G. Sword for providing the wingless locust nymph field experiment data. We also thank T. Chtanova for helpful discussions regarding cellular agent-based processes that motivated this work, and F-G Michalec for helpful discussions regarding copepod swimming data. We are grateful for the comments of the associate editor and reviewer, which led to substantial improvements to the paper.

Appendix A. Estimation of empirical eigenvalues for use in Monte Carlo simulations

To produce estimates $\hat{\lambda}_i$ for the eigenvalues λ_i used in the statistical tests presented in this paper, we use the sample autocovariance functions, defined as

$$\hat{\gamma}_i(r,s) = \frac{1}{N}\sum_{j=i+1}^{N}\Big(\hat{K}_j(r) - \bar{K}(r)\Big)\Big(\hat{K}_{j-i}(s) - \bar{K}(s)\Big), \quad 0 \le i \le N-1,$$

to produce an estimate for the long run covariance function $c(r,s)$. We refer to this as the sample long-run covariance (a kernel estimator), and define it as

$$\hat{c}_N(r,s) = \hat{\gamma}_0(r,s) + \sum_{i=1}^{N-1} U\Big(\frac{i}{h}\Big)\big(\hat{\gamma}_i(r,s) + \hat{\gamma}_i(s,r)\big).$$

An appropriate kernel function U and bandwidth h are required for accurate estimation of $c(r,s)$.

Suggestions for U and h are found in Horváth and Kokoszka (2012) and Horváth et al. (2013a,b, 2014a,b). We used a kernel, lag-window estimator as defined in Horváth et al. (2013a), with the flat-top kernel

$$U(s) = \begin{cases} 1, & 0 \le |s| < 0.5 \\ 2(1-s), & 0.5 \le |s| < 1 \\ 0, & |s| \ge 1 \end{cases}$$

and bandwidth parameter $h = N^{1/2}$. While we found this bandwidth to be appropriate for all analyses conducted in this paper, we noted in the analysis of highly dependent functional data, not presented here, that using a larger bandwidth than $h = N^{1/2}$ appeared to produce more accurate results.

If for each i, $\hat{K}_i(r)$ is defined for a vector of evenly spaced r values, then the estimated $\hat{c}_N(r,s)$ is likewise defined for an even grid of r and s values. By treating the resultant grid as a matrix, we can produce the empirical eigenvalues by finding the eigenvalues for this matrix. We performed Monte Carlo simulation to estimate the limiting quantiles for the limiting distributions (4) and (5). We performed 5000 simulations for each statistical test.

Appendix B. Numerical methods for spatial statistics

All numerical methods employed to produce the relevant output for this paper were produced using the software R (R Core Team, 2015), with the major package used being *spatstat* (Baddeley et al., 2015).

All $\hat{K}_i(r)$ functions were estimated using the *Kest spatstat* function, with the "Ripley" edge correction method. The $\bar{g}(r)$ estimates were produced using an approach similar to the *pcf.fv spatstat* function, with a smoothing parameter of 0.9.

Appendix C. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.spasta.2016.06.002.

References

Agnew, D., Green, J., Brown, T., Simpson, M., Binder, B., 2014. Distinguishing between mechanisms of cell aggregation using pair-correlation functions. J. Theoret. Biol. 352, 16–23.

Baddeley, A., Diggle, P.J., Hardegen, A., Lawrence, T., Milne, R.K., Nair, G., 2014. On tests of spatial pattern based on simulation envelopes. Ecol. Monograph 84 (3), 477–489.

Baddeley, A., Rubak, E., Turner, R., 2015. Spatial Point Patterns: Methodology and Applications with R. CRC Press.

Baddeley, A.J., Silverman, B.W., 1984. A cautionary example on the use of second-order methods for analyzing point patterns. Biometrics 1089–1093.

Balch, T., 2000. Hierarchic social entropy: An information theoretic measure of robot group diversity. Auton. Robots 8 (3), 209–238.

Berkes, I., Horváth, L., Rice, G., 2013. Weak invariance principles for sums of dependent random functions. Stochastic Process. Appl. 123 (2), 385–403.

Binder, B.J., Simpson, M.J., 2013. Quantifying spatial structure in experimental observations and agent-based simulations using pair-correlation functions. Phys. Rev. E 88 (2), 022705.

Binder, B.J., Simpson, M.J., 2015. Spectral analysis of pair-correlation bandwidth: application to cell biology images. Roy. Soc. Open Sci. 2 (2), 140494.

Binny, R.N., Plank, M.J., James, A., 2015. Spatial moment dynamics for collective cell movement incorporating a neighbour-dependent directional bias. J. R. Soc. Interface 12 (106), 20150228.

Bolker, B., Pacala, S.W., 1997. Using moment equations to understand stochastically driven spatial pattern formation in ecological systems. Theor. Popul. Biol. 52 (3), 179–197.

Brambilla, M., Ferrante, E., Birattari, M., Dorigo, M., 2013. Swarm robotics: a review from the swarm engineering perspective. Swarm Intell. 7 (1), 1–41.

Buhl, J., Sword, G.A., Clissold, F.J., Simpson, S.J., 2011. Group structure in locust migratory bands. Behav. Ecol. Sociobiol. 65 (2), 265–273.

Buhl, J., Sword, G.A., Simpson, S.J., 2012. Using field data to test locust migratory band collective movement models. Interface Focus 20120024.

Cagnacci, F., Boitani, L., Powell, R.A., Boyce, M.S., 2010. Animal ecology meets gps-based radiotelemetry: a perfect storm of opportunities and challenges. Philos. Trans. R. Soc. B 365 (1550), 2157–2162.

Cavagna, A., Cimarelli, A., Giardina, I., Orlandi, A., Parisi, G., Procaccini, A., Santagati, R., Stefanini, F., 2008. New statistical tools for analyzing the structure of animal groups. Math. Biosci. 214 (1), 32–37.

Cavagna, A., Del Castello, L., Giardina, I., Grigera, T., Jelic, A., Melillo, S., Mora, T., Parisi, L., Silvestri, E., Viale, M., et al., 2015. Flocking and turning: a new model for self-organized collective motion. J. Stat. Phys. 158 (3), 601–627.

Cavagna, A., Giardina, I., 2014. Bird flocks as condensed matter. Annu. Rev. Condens. Matter Phys. 5 (1), 183–207.

Chiu, S.N., Stoyan, D., Kendall, W.S., Mecke, J., 2013. Stochastic Geometry and its Applications. John Wiley & Sons.

Chtanova, T., Schaeffer, M., Han, S.-J., van Dooren, G.G., Nollmann, M., Herzmark, P., Chan, S.W., Satija, H., Camfield, K., Aaron, H., et al., 2008. Dynamics of neutrophil migration in lymph nodes during infection. Immunity 29 (3), 487–496.

Cressie, N., Wikle, C.K., 2011. Statistics for Spatio-Temporal Data. John Wiley & Sons.

Dale, M.R., Fortin, M.-J., 2014. Spatial Analysis: A Guide for Ecologists. Cambridge University Press.

Dehling, H., 1983. Limit theorems for sums of weakly dependent banach space valued random variables. Z. Wahrscheinlichkeitstheor. Verwandte Geb. 63 (3), 393–432.

Dell, A.I., Bender, J.A., Branson, K., Couzin, I.D., de Polavieja, G.G., Noldus, L.P., Pérez-Escudero, A., Perona, P., Straw, A.D., Wikelski, M., et al., 2014. Automated image-based tracking and its application in ecology. Trends Ecol. Evol. 29 (7), 417–428.

Diggle, P.J., 2003. Statistical Analysis of Spatial Point Patterns. Hodder Arnold.

Diggle, P.J., 2013. Statistical Analysis of Spatial and Spatio-Temporal Point Patterns. CRC Press.
Ferraty, F., Vieu, P., 2006. Nonparametric Functional Data Analysis: Theory and Practice. Springer Science & Business Media.
Gabriel, E., Diggle, P.J., 2009. Second-order analysis of inhomogeneous spatio-temporal point process data. Stat. Neerl. 63 (1), 43–51.
Gautrais, J., Ginelli, F., Fournier, R., Blanco, S., Soria, M., Chaté, H., Theraulaz, G., 2012. Deciphering interactions in moving animal groups. PLoS Comput. Biol. 8 (9), e1002678.
Geyer, C., 1999. Likelihood inference for spatial point processes. In: Stochastic Geometry: Likelihood and Computation, Vol. 80. pp. 79–140.
Gillespie, D.T., 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. J. Comput. Phys. 22 (4), 403–434.
Hörmann, S., Kokoszka, P., et al., 2010. Weakly dependent functional data. Ann. Statist. 38 (3), 1845–1884.
Horváth, L., Hušková, M., Rice, G., 2013a. Test of independence for functional data. J. Multivariate Anal. 117, 100–119.
Horváth, L., Kokoszka, P., 2012. Inference for Functional Data with Applications, Vol. 200. Springer Science & Business Media.
Horváth, L., Kokoszka, P., Reeder, R., 2013b. Estimation of the mean of functional time series and a two-sample problem. J. R. Stat. Soc. Ser. B Stat. Methodol. 75 (1), 103–122.
Horváth, L., Kokoszka, P., Rice, G., 2014a. Testing stationarity of functional time series. J. Econometrics 179 (1), 66–82.
Horváth, L., Rice, G., Whipple, S., 2014b. Adaptive bandwidth selection in the long run covariance estimator of functional time series. Comput. Statist. Data Anal..
Hyndman, R.J., Shang, H.L., 2010. Rainbow plots, bagplots, and boxplots for functional data. J. Comput. Graph. Statist. 19 (1).
Illian, J., Penttinen, A., Stoyan, H., Stoyan, D., 2008. Statistical Analysis and Modelling of Spatial Point Patterns, Vol. 70. John Wiley & Sons.
Jirak, M., 2013. On weak invariance principles for sums of dependent random functionals. Statist. Probab. Lett. 83 (10), 2291–2296.
Johnston, S.T., Simpson, M.J., McElwain, D.S., Binder, B.J., Ross, J.V., 2014. Interpreting scratch assays using pair density dynamics and approximate Bayesian computation. Open Biol. 4 (9), 140097.
Kelly, F.P., Ripley, B.D., 1976. A note on strauss's model for clustering. Biometrika 357–360.
Law, R., Dieckmann, U., 2000. A dynamical system for neighborhoods in plant communities. Ecology 81 (8), 2137–2148.
Loosmore, N.B., Ford, E.D., 2006. Statistical inference using the g or k point pattern spatial statistics. Ecology 87 (8), 1925–1931.
Meijering, E., Dzyubachyk, O., Smal, I., et al., 2012. Methods for cell and particle tracking. Methods Enzymol. 504 (9), 183–200.
Michalec, F.-G., Holzner, M., Souissi, A., Stancheva, S., Barras, A., Boukherroub, R., Souissi, S., 2015a. Lipid nanocapsules for behavioural testing in aquatic toxicology: Time–response of eurytemora affinis to environmental concentrations of pahs and pcb. Aquat. Toxicol..
Michalec, F.-G., Souissi, S., Holzner, M., 2015b. Turbulence triggers vigorous swimming but hinders motion strategy in planktonic copepods. J. R. Soc. Interface 12 (106), 20150158.
Ng, L.G., Qin, J.S., Roediger, B., Wang, Y., Jain, R., Cavanagh, L.L., Smith, A.L., Jones, C.A., De Veer, M., Grimbaldeston, M.A., et al., 2011. Visualizing the neutrophil response to sterile tissue injury in mouse dermis reveals a three-phase cascade of events. J. Invest. Dermatol. 131 (10), 2058–2068.
Plank, M.J., Law, R., 2015. Spatial point processes and moment dynamics in the life sciences: A parsimonious derivation and some extensions. Bull. Math. Biol. 77 (4), 586–613.
Plotkin, J.B., Potts, M.D., Leslie, N., Manokaran, N., LaFrankie, J., Ashton, P.S., 2000. Species–area curves, spatial aggregation, and habitat specialization in tropical forests. J. Theoret. Biol. 207 (1), 81–99.
Politis, D.N., Romano, J.P., 1994. Limit theorems for weakly dependent hilbert space valued random variables with applications to the stationary bootstrap. Statist. Sinica 4 (2), 461–476.
Preston, C.J., 1975. Spatial birth-and-death processes. Bull. Int. Statist. Inst. 46, 371–391. 405–408.
Raghib, M., Hill, N.A., Dieckmann, U., 2011. A multiscale maximum entropy moment closure for locally regulated space–time point process models of population dynamics. J. Math. Biol. 62 (5), 605–653.
R Core Team, 2015. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org.
Ripley, B.D., 1977. Modelling spatial patterns. J. R. Stat. Soc. Ser. B Stat. Methodol. 39 (2), 172–212.
Russell, J.C., Hanks, E.M., Haran, M., 2015. Dynamic models of animal movement with spatial point process interactions. J. Agric. Biol. Environ. Stat. 1–19.
Schweitzer, F., 2003. Brownian Agents and Active Particles. On the Emergence of Complex Behavior in the Natural and Social Sciences. In: Springer Series in Synergetics, Springer, Berlin, Germany.
Shang, H.L., Hyndman, R.J., 2015. rainbow: Rainbow Plots, Bagplots and Boxplots for Functional Data. R package version 3.3. URL http://CRAN.R-project.org/package=rainbow.
Strauss, D.J., 1975. A model for clustering. Biometrika 62 (2), 467–475.
Tong, P.L., Roediger, B., Kolesnikoff, N., Biro, M., Tay, S.S., Jain, R., Shaw, L.E., Grimbaldeston, M.A., Weninger, W., 2015. The skin immune atlas: three-dimensional analysis of cutaneous leukocyte subsets by multiphoton microscopy. J. Invest. Dermatol. 135 (1), 84–93.
Wu, W.B., 2005. Nonlinear system theory: Another look at dependence. Proc. Natl. Acad. Sci. USA 102 (40), 14150–14154.