

Dear Author,

Please, note that changes made to the HTML content will be added to the article before publication, but are not reflected in this PDF.

Note also that this file should not be used for submitting corrections.



ELSEVIER

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Spatial Statistics

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

A changepoint analysis of spatio-temporal point processes

Q1 Linda Altieri^{a,*}, E. Marian Scott^b, Daniela Cocchi^a,
Janine B. Illian^{c,d}

^a University of Bologna, Department of Statistical Sciences, Italy

^b University of Glasgow, School of Mathematics and Statistics, United Kingdom

^c University of St. Andrews, CREEM, School of Mathematics and Statistics, United Kingdom

^d NTNU Trondheim, Norway

ARTICLE INFO

Article history:

Received 8 January 2015

Accepted 18 May 2015

Available online xxxx

Keywords:

Spatio-temporal point processes

Changepoint analysis

INLA

Radioactive particle data

ABSTRACT

This work introduces a Bayesian approach to detecting multiple unknown changepoints over time in the inhomogeneous intensity of a spatio-temporal point process with spatial and temporal dependence within segments. We propose a new method for detecting changes by fitting a spatio-temporal log-Gaussian Cox process model using the computational efficiency and flexibility of integrated nested Laplace approximation, and by studying the posterior distribution of the potential changepoint positions. In this paper, the context of the problem and the research questions are introduced, then the methodology is presented and discussed in detail. A simulation study assesses the validity and properties of the proposed methods. Lastly, questions are addressed concerning potential unknown changepoints in the intensity of radioactive particles found on Sandside beach, Dounreay, Scotland.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

In this work, we propose a method for carrying out a changepoint analysis in the complex context of spatio-temporal point processes. Increasingly, spatio-temporal point process data are becoming

* Corresponding author.

E-mail address: linda.altieri@unibo.it (L. Altieri).

<http://dx.doi.org/10.1016/j.spasta.2015.05.005>

2211-6753/© 2015 Elsevier B.V. All rights reserved.

1 routinely available allowing questions concerning changes in the intensity of the process to be
2 addressed, such as in earthquake studies, where locations of earthquake epi-centres and strength
3 are mapped over time and where there is an interest assessing changes in the intensity and spatial
4 distribution of seismic events over recent years (Shearer and Stark, 2011; Touati et al., 2014), or
5 in occurrence of conflict data (Mangion et al., 2012). Other case studies derive from the field of
6 environmental monitoring, such as the dataset presented here.

7 Our study is motivated by questions on the monitoring and recovery of radioactive particles from
8 Sandside beach, North of Scotland, close to the former Dounreay nuclear facilities (Tyler et al., 2010).
9 Minute fragments of irradiated nuclear fuel particles, generally similar to grains of sand, have been
10 generated by historic practices at UKAEA Dounreay (<http://www.dounreay.com/particle-cleanup>).
11 UKAEA Dounreay historically released particles primarily into the marine environment. Since 1984,
12 particles have been found on the publicly accessible beach at Sandside Bay. A variety of different mon-
13 itoring systems have been used and the frequency of monitoring has also varied over the period since
14 routine monitoring first began. The beach monitoring campaign has several purposes, primarily driven
15 by the recovery of particles, and thus reduction of risk to the public from encountering such particles
16 and hence public reassurance. At the same, it provides improved understanding and knowledge of the
17 particle population abundance and its change as a result of continued monitoring and retrieval of par-
18 ticles offshore and historic site practices. Monitoring of the beaches (and in particular Sandside beach)
19 has been ongoing for a number of years with the first particles being detected in the 1980s. Over the
20 past 15 years, two major changes in the equipment used to detect the onshore particles have taken
21 place, representing known potential changepoints. In addition, offshore particle retrieval campaigns
22 are believed to have reduced the particle intensity for particles moved onshore with tides and cur-
23 rents with an unknown temporal lag, potentially generating multiple unknown changepoints in the
24 intensity function of the particle distribution. Questions on how to construct a method able to detect
25 unknown changepoints in such a complex dataset are raised; the proposed method has to deal with
26 the issues of spatial inhomogeneity, spatial dependence among points and temporal dependence of
27 the process.

28 1.1. Background and tools

29 For our work, we use a class of point process models called log-Gaussian Cox Processes (LGCPs).
30 Cox processes assume that the point distribution over space is due to stochastic environmental
31 heterogeneity, modelled as a random intensity function; given a realization of the intensity surface,
32 the distribution of points follows an inhomogeneous Poisson process. In LGCPs, the logarithm of the
33 intensity surface over an observation window W is assumed to be a (latent) Gaussian random field.
34 For a review on point process models and LCGPs we refer to Illian et al. (2008) and Diggle (2014).
35 LGCPs constitute a very flexible class of models that can potentially be extended to spatio-temporal
36 data (Diggle et al., 2013); tractability issues that have impeded the use of these models up to very
37 recent years can now be overcome using Integrated Nested Laplace Approximation (INLA, Rue et al.,
38 2009).

39 INLA is an alternative option to MCMC methods for approximating the posterior distribution of
40 the parameters of interest; it is simulation free, which is the key to its speed, and it exploits two
41 approximations. Firstly, a Laplace approximation is employed to represent the posterior distributions
42 with a Gaussian shape; secondly, the Gaussian Field is substituted by a Gaussian Markov Random Field
43 with a sparse precision matrix, which makes calculations very efficient. INLA returns the posterior
44 probability of every time point of being a changepoint, allowing the changepoint positions to be
45 inferred *a posteriori*.

46 Changepoint analysis is a well-established area of statistical research, frequently applied in a
47 temporal context, and less frequently over space (Raftery, 1993). The basic assumption is that data
48 are ordered and split into time segments following the same model but under different parameter
49 specifications (Chen and Gupta, 2012); the other common assumption is that observations are *i.i.d.*.
50 The interest lies in detecting the time and magnitude of the change(s). For a review of changepoint
51 techniques for temporal data we refer to Eckley et al. (2011). We aim at understanding what happens

when the usual assumptions of a changepoint analysis (simply temporal *i.i.d.* data) do not hold, which raises a few challenging issues especially when applied to the context of point processes.

1.2. Theoretical issues

While some of the existing changepoint methods can potentially be extended to the general spatio-temporal context, for spatio-temporal point processes this branch of analysis appears to be as yet relatively unexplored. Some substantial differences with regard to the standard changepoint analysis in time or in space have to be taken into account: firstly, at every time point the datum is not a single point but an irregular pattern of points, distributed over a possibly irregular observation window. Secondly, in many real situations, spatial dependence among points and temporal dependence within time segments have to be taken into account, and analytically obtaining mathematical quantities of interest, such as likelihood values and posterior distributions, is not trivial; modelling dependence within data segments in the context of unknown multiple changepoints is currently a challenge even for simple temporal series. Thirdly, frequently point process data are collected over space, and it is not common to have repeated measurements in the same observation window over time, in a sequence large enough to allow changepoint analysis. Most of the studies on point processes aim at describing the behaviour of the intensity function, therefore its changes over time are certainly of interest, and the provision of tools for changepoint analysis on spatio-temporal processes would enlarge the number of questions that can be answered, especially when spatial dependence among points and temporal dependence within time segments are included.

1.3. Research questions

The aim of our work is to propose a method to find multiple unknown changepoints over time in the inhomogeneous intensity of a spatio-temporal point process, allowing for spatial and temporal dependence within segments.

When dependence is allowed, the segment marginal likelihood usually becomes intractable, hence there is a need for fast computational methods, such as INLA, providing an accurate and tractable approximation of the likelihood. The computational speed and flexibility of INLA has not yet been exploited for a spatio-temporal changepoint analysis.

When applied to spatio-temporal point processes, a changepoint analysis of the behaviour of the intensity function over time can address different aspects: firstly, a change in scale, when the overall number of points increases or decreases significantly after a certain time point. Another option is a change in spatial structure, when the expected number of points remains constant, but their distribution over space changes after a certain time point. Lastly, the change can occur in both scale and spatial distribution. In the special case of a change in scale only, and when spatial homogeneity can be assumed throughout the whole time series, our method reduces to a traditional changepoint analysis on the time series of point counts. Our method allows changepoint detection to be extended to any point process where the information concerning the spatial distribution is of interest (as in the examples at the beginning of Section 1 state).

We aim at developing a method that is able to detect any of these changes over time, and that can therefore provide answers to a wider variety of cases and carry much more information than a traditional changepoint analysis that ignores spatial structure.

1.4. Motivation for the approach

In this study, we take a Bayesian approach to changepoint analysis for two main reasons. First of all, Bayesian inference allows knowledge brought by data (the likelihood) to be enriched by including extra information in the prior distributions of the parameters; in many real situations some changepoints might be considered more likely than others. Secondly, a Bayesian approach allows dependence to be dealt with, while there are currently no satisfactory frequentist solutions to the problem.

We use INLA to fit latent Gaussian models such as LGCPs as it brings substantial advantages when it comes to detecting multiple changepoints in a spatio-temporal point process context; first of all, very complex models can be fitted using INLA; the extension to spatio-temporal models is computationally challenging but feasible, and accurate and tractable approximations of the segment marginal likelihoods can be produced. Secondly, we can explore all the time series and compare the likelihood values resulting from different changepoint positions to choose the best position *a posteriori*. This is more efficient than traditional changepoint algorithms (Eckley et al., 2011); such a complex exploration in such a complex dataset would not be possible in reasonable time without INLA.

2. Methodology

2.1. Models

We consider a changepoint under four increasingly complex point process models, and consider the case of both a single changepoint and multiple changepoints at unknown locations; we discretize the observation window into a fine grid, and define $y_{ts} \sim \text{Poi}(|C|\lambda_{ts})$ as the number of points at time $t = 1, \dots, T$ in cell $s = 1, \dots, S$, where $|C|$ is the cell area. As is traditional in changepoint analysis [ref], we present the changepoint search as a test of two alternative hypotheses. H_0 means no changepoint; H_1 only concerns the number (1 or more) of changes and is therefore a complex hypothesis that may be decomposed in different sub-hypotheses for different changepoint positions.

We initially consider a model (Model 1) with an intercept, which assumes a spatially homogeneous intensity λ_t ; under each hypothesis (for the alternative, the case of a single changepoint is displayed for simplicity) we model the logarithm of the intensity function λ_t as:

$$\begin{aligned} H_0 : \log(\lambda_t) &= \mu + \epsilon_t \quad \text{for } t = 1, \dots, T \\ H_1 : \log(\lambda_t) &= \mu_1 + \epsilon_t \quad \text{for } t \leq \theta^* \\ &= \mu_2 + \epsilon_t \quad \text{for } t > \theta^* \end{aligned} \quad (1)$$

where μ is the intercept and $\epsilon_t \sim \text{iid}N(0, \tau_\epsilon)$ is an unstructured error term. Under H_0 all values over both space and time depend on a single value for μ , while under H_1 μ_t is constant over space but allowed to vary over time: a single changepoint in location θ^* splits the dataset into two time segments with a different value for the intensity function. In the more general case of $M \geq 2$ changepoints, the equation under H_1 is split into $M + 1$ segments defined by the ordered changepoint locations $\theta_1, \theta_2, \dots, \theta_M$.

Extensions to Model 1 (Eq. (1)) include adding a temporal effect (Model 2) and an extension to inhomogeneous processes (allowing for a spatially inhomogeneous intensity function λ_{ts}) by including a spatial effect (Model 3). All effects are included in Model 4 (given in Eq. (2)).

$$\begin{aligned} H_0 : \log(\lambda_{ts}) &= \mu + \phi + \psi_s + \epsilon_{ts} \quad \text{for } t = 1, \dots, T \text{ and } s = 1, \dots, S \\ H_1 : \log(\lambda_{ts}) &= \mu + \phi_1 + \psi_{1s} + \epsilon_{ts} \quad \text{for } t \leq \theta^* \text{ and } s = 1, \dots, S \\ &= \mu + \phi_2 + \psi_{2s} + \epsilon_{ts} \quad \text{for } t > \theta^* \text{ and } s = 1, \dots, S \end{aligned} \quad (2)$$

μ is a common intercept and, within each time segment, ϕ is a random effect modelled as an AR(1): $\phi_t = \phi_{t-1} + u_t$ where $u_t \sim N(0, \tau_\phi^{-1})$. Priors are needed for the precision $\tau_\phi \sim \text{Gamma}(\alpha_\phi, \beta_\phi)$. The spatial effect is ψ_s where the basic space unit s is the grid cell. This approximation is needed for tractability reasons, but INLA allows extremely fine grids while still being computationally feasible. ψ_s is modelled as an intrinsic CAR, i.e. as a Random Walk in two dimensions on a lattice, with a smooth neighbourhood structure that gives non-zero weights to the first 12 neighbours in the lattice (Illian et al., 2012). Again, the precision hyperparameter can be defined as $\tau_\psi \sim \text{Gamma}(\alpha_\psi, \beta_\psi)$.

The current implementation of INLA in the R-INLA software (www.r-inla.org) is not restricted to the spatial and temporal random fields chosen here.

2.2. Single changepoint detection

The single changepoint detection procedure starts by comparing the sub-hypotheses under H_1 , obtaining the marginal likelihoods conditional on different changepoint locations. Afterwards, the highest likelihood is either compared to the likelihood under H_0 , or compared to a chosen threshold. To do this, we now present two different Bayesian techniques.

2.2.1. Bayes factor method

We propose a modified version of the logarithm of the Bayes Factor, with only one term for θ^* instead of all possible θ s:

$$\gamma'_{\theta^*} = \log(\pi(\theta^*)) + q_1(\theta^*) + q_2(\theta^*) - l_0 = \log(\pi(\theta^*)) + l_1(\theta^*) - l_0 \quad (3)$$

where θ^* is the changepoint position returning the highest likelihood value under H_1 , $\pi(\theta^*)$ is the value of the prior distribution at the changepoint, and $l_1(\theta^*) = q_1(\theta^*) + q_2(\theta^*)$ is the corresponding maximum log-likelihood under the alternative hypothesis, obtained as a sum of two segment log-likelihoods. For the model with no changepoints, the maximum log-likelihood value under H_0 is greater than the maximum log-likelihood value under H_1 , as Bayes factors naturally incorporate penalization for model complexity. If $\gamma'_{\theta^*} > 0$, we reject the null model of no changepoint, and the changepoint is estimated to occur at θ^* . In conclusion, this method first compares the options under H_1 and then tests the best one against H_0 and is routinely used in Bayesian temporal changepoint analysis (Eckley et al., 2011).

2.2.2. Posterior threshold method

An alternative option is to fix a posterior probability threshold to identify changepoints. We consider the posterior distribution of the changepoint positions coming from the likelihood values conditional on different options under H_1 . Instead of testing them against H_0 , we use a decision threshold: if there are peaks in the posterior distribution above the threshold, the highest peak corresponds to the accepted changepoint. Suggestions for the choice of the threshold are given in Section 5. This method has the advantage of being visually immediate and easy to explain to non-statisticians; moreover, it is very flexible as the threshold choice can be adapted to the model fitting the data and to the analysis context.

2.3. Multiple changepoint detection

We now extend the method to an unknown number of changepoints; two approaches can be taken: a binary segmentation algorithm aimed at finding one changepoint in each step, or a simultaneous search aiming at finding all changepoints in one step.

2.3.1. Binary segmentation algorithms

For a general introduction to these methods we refer to Eckley et al. (2011), and in particular for point processes to Park et al. (2012). The idea of a binary segmentation procedure, and the key to its simplicity, is to split the multiple search into a series of subsequent searches; in each step, a single changepoint search is carried out, and either method (BF or PT) can be used. When running such an algorithm, number and positions of changepoints are estimated sequentially at the same time: in each step, if a changepoint is found, its position is immediately chosen before moving on to the next step, as we need to know where to split the data into further segments.

The analysis can become computationally very demanding as T and M become large, and methods are available for reducing time and memory storage requirements (Eckley et al., 2011). The computational efficiency of INLA makes this algorithm feasible even for complex spatio-temporal data.

2.3.2. Simultaneous changepoint search

The procedure we build follows a two level prior setting as in Wyse et al. (2011) and Fearnhead (2006), where a prior distribution is given to the number of changepoints and then a prior conditional on the number is given to the changepoint locations. As a consequence, we first estimate the number of changepoints, then conditional on that we identify the most likely positions. We then follow Wyse et al. (2011) with an extension to spatio-temporal models. $M + 1$ conditional likelihoods are computed using recursive equations, which give evidence for the model with m changepoints, $m = 0, \dots, M$. If $\pi(M)$ is non informative, the highest likelihood value $L(Y|m)$ corresponds to the chosen number of changepoints (see Section 2.4).

2.4. Posterior distribution

Irrespective of the detection method, the final goal of a Bayesian changepoint analysis is to obtain a posterior distribution of the number and positions of changepoints.

In a single changepoint search, the algorithm produces a posterior distribution assigning a probability to every potential changepoint position. For each model scenario, we run the model T times under the alternative hypothesis. In each run, we condition on the changepoint occurring at a different specific location $\theta \in \{1, \dots, T\}$ and fit one of the models; we obtain a conditional log-likelihood value $l_1(\theta) = q_1(\theta) + q_2(\theta)$ (see Section 2.2.1). The T -dimensional vector $l_1 = \{l_1(\theta), \theta \in \{1, \dots, T\}\}$ is then transformed following the usual Bayes Rule to obtain the posterior distribution: in the absence of prior knowledge, rescaling the likelihood vector to integrate to 1 gives the posterior distribution of interest. The time point corresponding to the maximum posterior value is taken as the most likely changepoint position *a posteriori* θ^* . The decision on the significance of the detected potential changepoint with respect to the null hypothesis can be based on either the Bayes Factor or the Posterior Threshold method. The computational expense of refitting each model many times for different potential changepoints is not prohibitive when using INLA.

In a multiple changepoint search, as for the binary segmentation algorithm, a criterion for decision making must be chosen first in order to proceed with the iterations. Any of the methods proposed in 2.2 can be used; once chosen, the posterior distribution for each potential changepoint position is obtained for each step and time segment just as for the single search. Since time is discrete, a final posterior distribution for the whole time series can be obtained by averaging values pointwise and then rescaling in order to integrate to 1 and to deal with a proper distribution. If a simultaneous search is carried out, we obtain $M + 1$ conditional data likelihoods $L(Y|m = 0), L(Y|m = 1), \dots, L(Y|m = M)$, where Y is the whole dataset. Following the Bayes Rule $P(m|Y) \propto L(Y|m)\pi(m)$, if the prior is uniform $\hat{M} = \arg \max_m \{L(Y|m), m = 1, \dots, M\}$. Conditional on \hat{M} , the posterior positions for the changepoints are then estimated given the previous change (with the convention $\theta_0 = 0$), the data and \hat{M} . Assuming the changepoint process is a Markov process we find them iteratively:

$$P(\theta_1, \dots, \theta_{\hat{M}}|Y, \hat{M}) = P(\theta_1|Y, \hat{M}) \times P(\theta_2|\theta_1, Y, \hat{M}) \times \dots \times P(\theta_{\hat{M}}|\theta_{\hat{M}-1}, Y, \hat{M}).$$

3. Simulation study

In order to assess and compare the performance of the methods proposed in Section 2.1, we carry out a simulation study covering different situations.

3.1. Simulation design

We fix a time series of $T = 50$ time points, and a grid of $S = 20 \times 20 = 400$ cells. The observation window W is a square of area 100. We also ran simulations in irregular polygonal windows, which did not lead to substantially different conclusions; we therefore only consider the square window here.

We simulate point pattern data that follow an *i.i.d.* and an AR(1) process in time, respectively, for both single and multiple changepoint detection under H_0 (no changepoint, $\lambda = 1$ over the whole series) and under different options for H_1 :




Fig. 1. Examples of simulated data, with the corresponding underlying intensity function: initial intensity (a), large scale change (b), spatial change (c), spatial and scale change (d).

- one large changepoint in scale (from $\lambda_1 = 1$ to $\lambda_2 = 2$)
- one small changepoint in scale (from $\lambda_1 = 1$ to $\lambda_2 = 1.2$)
- three changepoints in scale ($\lambda_1 = 1, \lambda_2 = 1.4, \lambda_3 = 2.3, \lambda_4 = 2$)
- one changepoint in spatial structure
- one changepoint in both scale and spatial structure.

Fig. 1 shows examples of the above listed changes. Panel (a) displays the initial value for the intensity function; panel (b) shows a large change in scale (small change and multiple changes not displayed here); panel (c) a change in spatial structure; panel (d) a change in both space and scale. Data with a change in scale are generated under both a homogeneous (a single value for λ over the window) and an inhomogeneous process (different values for λ_s over the window, with a mean value over space equal to the homogeneous λ). 100 replicates are generated for each case.

In order to find a sensible and not too arbitrary threshold for the PT method, it is possible to use simulated data under the null hypothesis for assessing the significance level α based on different threshold values. Once we find a value such that the significance level does not exceed a certain limit (usually $\alpha \leq \{0.01, 0.05, 0.1\}$), we use that threshold on data generated under the alternative hypothesis in order to evaluate its power level, the ability to detect the correct changepoint locations and the accuracy of the produced estimates.

3.2. Simulation results

After we fit Model 1 to 4 (Section 2.1) to all simulated patterns, we assess the performance of the proposed detection techniques: for a single search, the BF and PT methods (2.2); for a multiple search, the simultaneous approach and the binary segmentation algorithm combined with both BF and PT methods (2.3). Methods are evaluated based on type I and type II errors (see Table 1), number and position of detected changepoints and accuracy of the intensity estimates (tables not reported here). Table 1 does not report results where the data generation process and the nature of the model do not match: Model 1 and 2 are appropriate for spatially homogeneous data, while Model 3 and 4 are only relevant to inhomogeneous data. As for the ability to detect changepoints (*i.e.* to reject H_0), Table 1 shows that when a change in scale is considered all methods perform quite well over the first two models, but all detection techniques that use the BF method and the simultaneous search become too conservative as soon as spatial dependence is included (too many 0s in the table); the PT method performs better over all models, also due to the possibility to tune the threshold value according to the model. We also covered different types of change in the intensity function which involve the spatial distribution (labelled as 'spatial' and 'spat+scale' in Table 1). When a single change in spatial structure is considered, the performance of all methods is excellent: changes in the spatial structure are detected with both methods when using Models 3 and 4, which include a spatial effect, and changes in both scale and space are detected for all models.

Once the ability to reject H_0 is assessed, we focus on the number of detected changepoints. Results are correct in all H_0 data: even in situations where some changepoints were found, as in AR(1) data, they were correctly identified as spurious changepoints (in a different location for every replicate). As regards the detection in H_1 data, when a changepoint is detected with any method, its location is either correctly identified or in a position so close to the true changepoint that the slight mislocation would be irrelevant in practice. Again, spurious changes in time dependent data do not affect the conclusions,

Table 1

'Significance levels' (type I errors) in data generated under H_0 and 'power levels' ($1 -$ type II errors) in data generated under H_1 .

		Homogeneous process				Inhomogeneous process				
		Model 1		Model 2		Model 3		Model 4		
		BF	PT	BF	PT	BF	PT	BF	PT	
IID	H_0	$\lambda = 1$	0	≤ 0.05	0	≤ 0.05	0	≤ 0.1	0	≤ 0.1
		$\lambda_2 = 2$	1	1	1	1	1	1	0	1
		$\lambda_2 = 1.2$	1	1	1	0.98	0	0.34	0	0.3
		spatial	0.38	0.44	0.42	0.26	1.00	1.00	1.00	1.00
		spat +scale	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		mult BinSeg mult Simult	1	1	0.99	1	0	0.93	0	0.26
AR(1)	H_0	$\lambda = 1$	0.96	0.66	0.38	0.24	0.15	0.26	0	0.18
		$\lambda_2 = 2$	1	0.97	0.81	0.75	0.53	0.81	0	0.52
		$\lambda_2 = 1.2$	1	0.73	0.43	0.36	0.19	0.54	0	0.37
		mult BinSeg	1	0.98	0.94	0.91	0.55	0.84	0	0.67
		mult Simult	1	1	1	1	0	0	0	0




Fig. 2. Simulated intensity functions for four time segments, with three changepoints (top panels). Estimated segment intensities under the detection of two changepoints (bottom panels).

i.e. our methods are still able to detect the correct changepoint locations despite the variability in the time series.

Lastly, given the detected changepoints, estimates are very accurate over all the simulated scenarios; see Fig. 2 for an example. When a changepoint is not detected (e.g. third/fourth column in Fig. 2), estimated values are located in between the two segments' true values, and when a changepoint was only detected in a proportion of the replicates (as happens with very small changes), the true magnitude of the change is reduced. When a change in scale is considered, the correct (increasing or decreasing) trend is always captured; as can be seen from Fig. 2, results are very accurate in reproducing both scale and spatial trend.

4. Case study

Since the 1950s, Dounreay has been the site of several nuclear research establishments, most of which are now being decommissioned. Radioactive particles have been found on local beaches in the North of Scotland since the 1990s as a result of historic practices during nuclear fuel reprocessing at the Dounreay plant (<http://www.dounreay.com>). The dataset used gives the particles' locations on one of the local beaches, Sandside beach, during each of the years of monitoring (Dounreay-Particle-Advisory-Group, 2012) (see Fig. 3). The temporal data series consists of yearly point pattern




Fig. 3. Radioactive particles detected on the Sandside beach between 1999 and 2013.

Table 2

Results for a changepoint search on Sandside beach data – four log-Gaussian Cox process models, two methods for single changepoint detection (BF and PT), three methods for multiple changepoint detection (BF or PT with binary segmentation algorithm and a simultaneous approach).

(a) Single changepoint detection			
Model	BF	PT	
Intercept	2006	2006	
Temporal	2003	2003	
Spatial	–	2006	
Sp-temp	–	–	
(b) Multiple changepoint detection			
Model	BF BinSeg	PT BinSeg	Simultaneous
Intercept	2003, 2006, 2012	2003, 2006, 2012	2006
Temporal	2003	2003, 2006, 2012	2003
Spatial	–	2006, 2012	–
Sp-temp	–	–	–

realizations, and additional information on the retrieval and radioactivity level for each of the particles. The underlying intensity and its spatial structure are of interest, along with potential changes. For the questions presented in Section 1, this motivating dataset represents an adequate example of an inhomogeneous spatio-temporal point process with changepoints over time.

MCMC goodness of fit tests (Illian et al., 2008) show that Cox processes fit data very well; in particular, the flexible class of log-Gaussian Cox processes is suitable for the problem as the distribution of particles could be due to an underlying driver (tides and winds). Moreover, it is very straightforward to modify these models by adding an intercept, random or smooth effects to the structured predictor; the estimation with INLA is very fast (and precise) even for complex models and this allows us to fit several different models without becoming computationally prohibitive.

4.1. Results on particle data

Table 2 displays a summary of the number and positions of detected changepoints in the data series for both a single and a multiple search. Results must be interpreted carefully since the time series is very short. The first two detected changepoints in 2003 and 2006 correspond to the periods of equipment changes and produce an increase in the point intensity; this supports that the changes in equipment have significantly improved the probability of detecting particles. The third changepoint in 2012 is very close to the end of the series, therefore conclusions must be drawn with some caution here; the results show some evidence of a decreasing intensity, which might be related to the offshore retrieval campaign, suggesting a reduction of the arrival of particles on Sandside beach. The results highlight the flexibility and reliability of our method with the detection of the same changepoints using several techniques, as happens with 2003 and 2006, which gives us more confidence in the decision. Further discussion can be found in Section 5.




Fig. 4. Estimated intensities for model 2 with Binary segmentation and PT method. (a) 1999–2005; (b) 2006–2011; (c) 2012–2013.

1 An example of the analysis output is given in Fig. 4: this shows the estimated intensity functions
 2 for multiple changepoint detection using the model including spatial dependence and the Posterior
 3 Threshold method; they are derived by the INLA estimates of the latent field. It can be seen in the
 4 Figure that the spatial structure of the intensity function is inhomogeneous with a low density value
 5 in most of the window and a hot spot in the bottom-right area. The scale of values in the intensity
 6 plots shows that there is an increase in the intensity after 2006, and then a decrease in the last two
 7 years of the series.

8 5. Discussion

9 The proposed new method is able to find unknown multiple changepoints in the intensity of
 10 a spatio-temporal point process, including dependence within time segments. A flexible class of
 11 models, Log-Gaussian Cox Processes, is used, that allows unobserved environmental heterogeneity
 12 in the spatial patterns to be accounted for. Changes in both intensity and structure of the patterns
 13 can therefore be detected. The models have been conveniently fitted using INLA, thus providing a
 14 practically relevant case study of INLA in the new context of spatio-temporal changepoint analysis.

15 We conclude with some comments and remarks on our method and results.

16 Firstly, the choice of threshold for the PT method tunes the conservativeness of conclusions:
 17 larger values (closer to 1) lead to more conservative results, and smaller values (closer to 0) detect
 18 changepoints more easily. This choice can therefore be informed by prior knowledge, if information
 19 on the location of changepoints in the data series is available. In our simulation study, we propose an
 20 objective way of obtaining a threshold.

21 Secondly, as a final note on the multiple changepoint detection techniques, we would like to
 22 highlight that models with several changepoints will not necessarily be preferred. Indeed, in the
 23 binary segmentation algorithm up to 4 changepoints were detectable, but a fourth one was hardly
 24 ever found; moreover, in many scenarios less than three changes (*i.e.* two, or even zero) were detected,
 25 meaning that a model allowing for more data segments does not always describe data better. The
 26 simultaneous approach too proves to be conservative on our data.

27 Thirdly, as regards results, the proposed method is not only satisfactory in detecting where the
 28 changepoint(s) occur(s), but also in producing accurate estimates. The ability to detect changepoints
 29 does not depend directly on the INLA approach, but depends on the choice of the detection method.
 30 Given the detection of changepoints, estimates are very accurate in reproducing both spatial trend
 31 and scale of values of the intensity function over all the simulation study.

32 Lastly, our approach provides flexible methods for obtaining fast results and reducing arbitrariness
 33 in taking decisions. Indeed, in the simulation study (Section 3.2) we compare the performances of the
 34 different methods and we highlight which ones are more conservative and which less so. If a single
 35 method must be chosen for a specific case study, when there is prior knowledge about the occurrence
 36 of many changes a less conservative technique can be used; when strong evidence is needed, a more
 37 conservative one ought to be preferred. In addition, the different methods often identify the same
 38 changepoints, therefore in absence of prior knowledge we suggest using all detection methods to
 39 ensure robustness: the changepoints found in all cases will be the most likely to be true, while the
 40 other ones must be dealt with carefully and this is where prior knowledge can play a role. Thus, the
 41 speed and flexibility of our approach reduces subjectivity and improves reliability of the results. The
 42 methodology discussed in this paper aims at exploiting different types of information derived from
 43 fitting a number of models. This is very different from a setting where one wishes to compare several
 44 models; when this is of interest, the DIC is routinely used in Bayesian inference.

In addition, the required computational time is acceptable: in the simulation study running the models took a few minutes for each replicate, and on real data all results were obtained in less than 30 min in total.

The set of presented models constitutes the basis for extensively analysing spatial inhomogeneity, spatial dependence and temporal dependence. Many promising extensions are possible. It is worth mentioning the role of the prior distributions: they can concern the number of changepoints, their locations and all the effects hyperparameters, and it would certainly be of interest to test different prior settings than the default ones in R-INLA (Simpson et al., 2014), and check how strongly they can affect conclusions. Moreover, once the ability to include dependence in the model is assessed, the focus may be on adding covariates and marks and running some Bayesian model selection. Furthermore, the analysis can be further extended to gradual changes in the intensity function, in order to compare its performance with a study of the spatio-temporal trend of the series.

Acknowledgement

As regards author Linda Altieri, the research work underlying this paper was partially funded by a FIRB 2012 grant (project no. RBF12URQJ; title: Statistical modeling of environmental phenomena: pollution, meteorology, health and their interactions) for research projects by the Italian Ministry of Education, Universities and Research.

References

- Chen, J., Gupta, A.K., 2012. Parametric Statistical Changepoint Analysis, second ed. Birkhauser, Boston.
- Diggle, P., 2014. Statistical Analysis of Spatial and Spatio-Temporal Point Patterns, third ed. Taylor & Francis Group, Boca Raton.
- Diggle, P., Moraga, P., Rowlingson, B., Taylor, B.M., 2013. Spatial and spatio-temporal log-Gaussian Cox processes: extending the geostatistical paradigm. *Statist. Sci.* 28, 542–563.
- Dounreay-Particle-Advisory-Group, 2012. Dounreay Site Restoration Limited. Annual Report, Tech. Rep. Scottish Environment Protection Agency.
- Eckley, I.A., Fearnhead, P., Killick, R., 2011. Analysis of changepoint models. In: Barber, D, C.A.T., Chiappa, S., C.U.P. (Eds.), *Bayesian Time Series Models* (Chapter 10).
- Fearnhead, P., 2006. Exact and efficient Bayesian inference for multiple changepoint problems. *Stat. Comput.* 16, 203–213.
- Illian, J., Penttinen, A., Stoyan, H., Stoyan, D., 2008. *Statistical Analysis and Modelling of Spatial Point Patterns*. Wiley, Chichester.
- Illian, J.B., Sørbye, H.S., Rue, H., 2012. A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA). *Ann. Appl. Stat.* 6, 1499–1530.
- Mangion, A.Z., Dewar, M., Kadiramanathan, V., Sanguinetti, G., 2012. Point process modelling of the Afghan War Diary. *Proc. Natl. Acad. Sci. USA* 109, 12414–12419.
- Park, T., Krafty, R.T., Sánchez, A.L., 2012. Bayesian semi-parametric analysis of Poisson change-point regression models: application to policy making in Cali, Colombia. *J. Appl. Stat.* 39 (10), 2285–2298.
- Raftery, A.E., 1993. *Change Point and Change Curve Modeling in Stochastic Processes and Spatial Statistics*, Tech. Rep. University of Washington.
- Rue, H., Martino, S., Chopin, N., 2009. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B* 71, 319–392. Part 2.
- Shearer, P.M., Stark, P.B., 2011. Global risk of big earthquakes has not recently increased. *Proc. Natl. Acad. Sci. USA* 109, 717–721.
- Simpson, D.P., Martins, T.G., Riebler, A., Fuglstad, G.A., Rue, H., 2014. Penalising model component complexity: A principled, practical approach to constructing priors, arXiv preprint [arXiv:1403.4630](https://arxiv.org/abs/1403.4630).
- Touati, S., Naylor, M., Main, I.G., 2014. Statistical modeling of the 1997–1998 Colfiorito earthquake sequence: Locating a stationary solution within parameter uncertainty. *Bull. Seismol. Soc. Am.* 104, 885–897.
- Tyler, A.N., Scott, E.M., Dale, P., Elliott, A.T., Wilkins, B.T., Boddy, K., Toole, J., Cartwright, P., 2010. Reconstructing the abundance of Dounreay hot particles on an adjacent public beach in Northern Scotland. *Sci. Total. Environ.* 408, 4495–4503.
- Wyse, J., Friel, N., Rue, H., 2011. Approximate simulation-free Bayesian inference for multiple changepoint models with dependence within segments. *Bayesian Anal.* 6, 501–528.