

A changepoint analysis on spatio-temporal point processes

Un'analisi changepoint per processi di punto spazio-temporali

Linda Altieri, E. Marian Scott, Janine B. Illian

Abstract This work introduces a likelihood ratio based test statistic for detecting change points over time in the inhomogeneous intensity of a spatio-temporal point process. We propose a new method for detecting changes by fitting a spatio-temporal log-Gaussian Cox process model using INLA. A simulation study assessing the validity and properties of the test is presented.

Abstract *Questo lavoro presenta una statistica test basata sul rapporto di verosimiglianza per identificare nel tempo change points nell'intensità non-stazionaria di un processo di punto spazio-temporale. I cambiamenti sono identificati utilizzando un modello di Cox log-Gaussiano spazio-temporale che impiega INLA. Uno studio di simulazione valuta la validità e le proprietà del test.*

Key words: changepoint, spatio-temporal point processes, LGCP, INLA

1 Motivation

Changepoint analysis is a well-established area of statistical research; while some of the existing changepoint methods can be extended to the spatio-temporal context, for spatio-temporal point processes this appears to be as yet relatively unexplored. Some differences with regard to the standard changepoint analysis have to be taken into account: firstly, an individual datum is a pattern of points, which makes any analysis more complex; secondly, the measured response variable is the point location,

Linda Altieri
University of Bologna, Dept of Statistics, e-mail: linda.altieri@unibo.it

E. Marian Scott
University of Glasgow, School of Mathematics and Statistics e-mail: Marian.Scott@glasgow.ac.uk

Janine B. Illian
University of St Andrews, CREEM, School of Mathematics and Statistics e-mail: janine@mcs.st-and.ac.uk

different from a common quantitative/qualitative variable. Moreover, the issues of spatial dependence among points and of temporal dependence within time segments have to be addressed. We therefore believe a statistical analysis of changepoint detection methods in the context of spatio-temporal point processes is a challenging and interesting study area.

1.1 Background

The basic assumption in a changepoint analysis is that data are ordered and split into segments, following the same model but under different parameter specifications [6]. The other common assumption is that observations are *i.i.d.*. Modelling dependence within data segments in the context of unknown multiple change points is currently a challenge: when dependence is allowed, the segments' marginal likelihood usually becomes intractable, and the computational complexity of the problem increases; there is a need for fast methods providing an accurate and tractable approximation of the likelihood. Recent work by Wyse et al. [6] uses Integrated Nested Laplace Approximation (INLA) to make a model allowing for dependence within segments feasible. INLA is an alternative, computationally efficient approach to MCMC methods [4] to obtain the posterior distribution of model parameters; Wyse et al. employ INLA to produce approximations for the posterior of both the number of change points and their positions.

Our work extends these new techniques to the context of spatio-temporal point processes, in the case of log-Gaussian Cox point processes (LGCPs). Cox processes assume that the point distribution over space is due to stochastic underlying heterogeneity, modelled as a random intensity function $\Lambda(s)$ [2]; given $\Lambda(s)$, the distribution of points follows an inhomogeneous Poisson process. In LGCPs the logarithm of the intensity surface in an observation window W is assumed to be a (latent) Gaussian field $Z(s)$, *i.e.* $\Lambda(s) = \int_W \lambda(s) ds = \exp(Z(s))$. LGCPs constitute a very flexible class of models that can be adapted to the spatio-temporal case and fitted using INLA [3]. The INLA approach has several fundamental advantages: first of all, it is an effective computational tool for model implementation; its efficiency allows an extension from the temporal to the spatio-temporal context; the likelihood values resulting from different changepoint positions can be evaluated, to choose the best position a posteriori. We present a simulation study of this approach in the spatio-temporal point process context; unlike traditional changepoint detection algorithms (see [1]), with this method the 3 dimensions of the problem (two spatial and one temporal) are maintained. We define a test statistic allowing decisions on whether, and how many, temporal change points are present.

Our study is motivated by questions on the monitoring and recovery of radioactive particles from Sandside beach, North of Scotland [5]; over approximately 30 years, there have been three major changes in the equipment used to detect the particles, representing known potential change points. In addition, offshore particle re-

trieval campaigns are believed to have reduced the particle intensity onshore with an unknown temporal lag, generating multiple unknown change points in the intensity.

2 Methodology

We define what is meant by a change point under several common, but increasingly complex, point process models. We consider the case of one change point at an unknown location; we focus on discrete models, and we assume $y_{ts} \sim \text{Poi}(|C|\lambda_{ts})$ where $|C|$ is the cell area, t indexes time and s space. We initially consider a model with fixed effect which assumes a spatially homogeneous intensity λ_t

$$\log(\lambda) = \mu + \varepsilon \quad (1)$$

which is expressed under each hypothesis as :

$$\begin{cases} H0 : \log(\lambda_t) = \mu + \varepsilon_t \\ H1 : \log(\lambda_t) = \mu_1 + \varepsilon_t \text{ for } t < \tau^* \\ \quad \log(\lambda_t) = \mu_2 + \varepsilon_t \text{ for } t \geq \tau^* \end{cases}$$

Under H0 all values over both space and time depend on a single value for μ , while under H1 μ_t is constant over space but allowed to vary over time. For a single change point, we first have to detect where the change occurs, and then we estimate two values for μ_t .

We separately add a random effect describing temporal dependence and a random effect describing spatial inhomogeneity and dependence; finally, for the most complicated model we consider both effects:

$$\log(\lambda) = \alpha + \phi_1 + \phi_2 + \varepsilon \quad (2)$$

where α is an intercept, ϕ_1 is modelled as an intrinsic CAR, i.e. as a Random Walk in two dimensions on a lattice, and ϕ_2 is modelled as an AR(1). The temporal dependence is only assumed to be within, not across, segments. The precision parameter for both effects has a $\text{Gamma}(\alpha, \beta)$ prior. Specifically,

$$\begin{cases} H0 : \log(\lambda_{ts}) = \alpha + \phi_{1s} + \phi_{2s} + \varepsilon_{ts} \\ H1 : \log(\lambda_{ts}) = \alpha + \phi_{1ts} + \phi_{2ts} + \varepsilon_{ts} \end{cases}$$

For the spatially inhomogeneous case, we build the model assuming that the spatial structure is the same over time up to a scale parameter, and the changepoint detection identifies the time point that describes the change of scale in our data, taking spatial and temporal dependence into account.

The test statistic we choose is

$$\gamma_{\tau^*} = \frac{\pi(\tau^*)L_1(\tau^*)}{L_0} \quad (3)$$

where τ^* is the changepoint position returning the maximum likelihood value under H1, $\pi(\tau^*)$ is the prior distribution on its position, and $L_1(\tau^*)$ and L_0 are the maximum likelihoods under H1 and H0 respectively. This is a more conservative version of the Bayes Factor, where the numerator would be the sum over all τ s, used in place of the posterior ratio (when the prior ratio is 1) to decide if there is a change point [1]. Equivalently, we can take the logarithm of (3)

$$\gamma'_{\tau^*} = \log(\pi(\tau^*)) + l_1^* - l_0. \quad (4)$$

If $\gamma'_{\tau^*} > 0$, we reject the null model of no change point, and the change point is estimated to occur at τ^* .

3 Simulation study

3.1 Simulation design

To evaluate the proposed changepoint test procedure, a simulation study is implemented on a square observation window with a 10×10 unit area. We consider the case of detecting a single change point, with a time series length $T = 50$.

A prior distribution must be set on both number (0 or 1) of change points and its position; we give the same probability mass to both values: $\pi(0) = \pi(1) = 0.5$. As for its position (excluding the endpoints of the range):

$$\pi(\tau) = \frac{1}{T-9} \text{ for } \tau \in \{5, 6, \dots, 44, 45\}.$$

We build four model scenarios: a simple one assuming a spatially homogeneous process and introducing a fixed effect, a model adding a temporal random effect, a third model with intercept and spatial dependence and heterogeneity, and finally a spatially inhomogeneous model with two smooth effects allowing for both spatial and temporal dependence within segments. The discretised response vector Y is $[T \times S]$ long, where $S = 400$ is the number of grid cells, and is made up of blocks Y_t , counting, for each time point, the number of points in each cell s .

In the simplest model, data are generated from a stationary Poisson process; for the more general case, they are generated from a spatially inhomogeneous process, whose spatial structure is constant over time up to a scale factor, using an inhomogeneous Poisson process (analogous to generating from a LGCP, but fixing the latent field). For both cases, we produce datasets under the null hypothesis of no change point and under the alternative hypothesis of one change point; we choose to try different change magnitudes: given an initial function λ_s , the changes are $10\lambda_s$, $2\lambda_s$ and $1.2\lambda_s$. The number of simulation replicates, for each scenario and under both hypotheses, is 100. An example of an inhomogeneous pattern is shown in Fig 1.

For every model scenario, under the alternative, the model is run $T - 9$ times for each simulation replicate; in each run, we condition on the change point occurring at location $\tau = (5, 6, \dots, 44, 45)$ and return $l_1(\tau)$ given that changepoint location. We

take as the most likely a posteriori changepoint position τ^* the one returning $l_1^* = l_1(\tau^*) = \max_{\tau}(l_1(\tau))$. The model is then run once under H_0 . Finally, the test statistic (4) is evaluated to choose between the two competing models. The computational efficiency of INLA and the ability to return the (approximate) likelihood value for any model makes this changepoint search algorithm feasible.

3.2 Simulation results

Fig 2 shows the performance of the method under the first model scenario with a single change point and a fixed effect; the test statistic (4) leads to the correct conclusion in all cases. Under H_0 no clear peak can be found in the pointwise averaged posterior probability function; under H_1 , with the two greater intensity change magnitudes, the correct change point is detected in all replicates. Even with 1.2λ , despite the magnitude of the change being very small compared to the data variance, a change point is detected, and in 70 % of the replicates the correct location is identified, while in all other cases it is detected in adjacent locations.

When adding temporal dependence, results remain substantially identical for H_0 data and for the two greater changes; as for 1.2λ , a change point is detected only in 40% of cases, but a greater difficulty in detecting such a small change is to be expected when data are assumed to be dependent over time. According to DIC, the more complicated model is preferred to the first one in all cases.

Even when accounting for both temporal and spatial dependence and heterogeneity, the INLA detection algorithm maintains a very good performance and the graphs returned highlight change points in the correct position or extremely close to it.

4 Conclusions and discussion

We developed a likelihood based changepoint detection technique with the two important innovations of being applicable to spatio-temporal point processes and allowing for both spatial and temporal dependence within segments. The choice of a conservative version of the test statistic should protect against type 1 errors; there is still a need to improve the statistic (4) to obtain better results for the spatially inhomogeneous models under the alternative hypothesis. The simulation study with the simple scenarios considered shows that the proposed method performs well. Further work will develop more realistic scenarios to further check the chosen statistic's significance and power.

There are several extensions to the current work under consideration. In the first, application to a more general model allowing for any type of spatial inhomogeneity, where change could occur in both space and time. In the second, generalising to multiple change points. even in these relatively simple cases, computational time is relatively long; it may be reduced by using the recent stochastic partial differential

equation approach. Another option for large datasets is the reduced filtering recursion (RFR) presented by [6], based on selecting a representative subset of points from the data and looking for change points inside the sampled series. The reported study can nevertheless be considered as a first step toward the analysis of intensity change points in a general spatio-temporal context.




Fig. 1 Simulated data example, different intensity levels




Fig. 2 Examples of change point detection and posterior probabilities

References

1. Eckley, I.A., Fearnhead, P., Killick, R.: Analysis of Change point Models. In: Barber, D., Cengil, A.T., Chiappa, S. (eds.) Bayesian Time Series Models, Chapter 10. Cambridge University Press (2011)
2. Illian, J., Penttinen, A., Stoyan, H., Stoyan, D.: Statistical Analysis and Modelling of Spatial Point Patterns. Wiley, Chichester (2008)
3. Illian, J., Sørbye, S., Rue, H.: A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA). *Annals of Applied Statistics* **6**, 1499-1530 (2012)
4. Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of The Royal Statistical Society Series B*, **71**, 319-392 (2009)
5. Tyler, A.N., Scott, E.M., Dale, P., Elliott, A.T., Wilkins, B., Boddy, K., Toole, J., Cartwright, P.: Reconstructing the abundance of Dounreay hot particles on an adjacent public beach in Northern Scotland. *Science of the Total Environment* **408**, 4495-4503 (2010)
6. Wyse, J., Friel, N., Rue, H.: Approximate simulation-free Bayesian inference for multiple change point models with dependence within segments. *Bayesian Analysis* **6**, 501-528 (2011)