Lecture 10
Spatio-Temporal Point Processes

Dennis Sun
Stats 253

July 23, 2014

Outline of Lecture

Where are we?

Intensity Estimation in Poisson Processes

$\lambda(\boldsymbol{s})$ can be estimated nonparametrically...

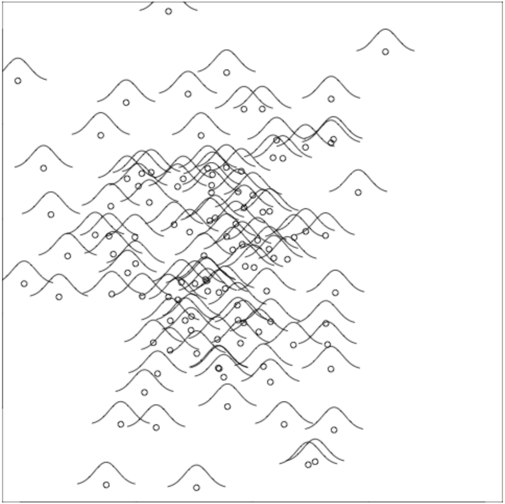

Intensity Estimation in Poisson Processes

or parametrically, such as...

$$\log \lambda(\boldsymbol{s}) = \boldsymbol{x}(\boldsymbol{s})^T \boldsymbol{\beta}$$

To do this, note that the likelihood of $(\boldsymbol{s}_1, ..., \boldsymbol{s}_{N(D)}, N(D))$ is:

$$L(\lambda(\cdot)) = e^{-\int_D \lambda(\boldsymbol{s})\, d\boldsymbol{s}} \frac{\left(\int_D \lambda(\boldsymbol{s})\, d\boldsymbol{s}\right)^{N(D)}}{N(D)!} \prod_{i=1}^{N(D)} \frac{\lambda(\boldsymbol{s}_i)}{\int_D \lambda(\boldsymbol{s})\, d\boldsymbol{s}}$$

so the log-likelihood is

$$\ell(\lambda(\cdot)) = -\int_D \lambda(\boldsymbol{s})\, d\boldsymbol{s} + \sum_{i=1}^{N(D)} \log \lambda(\boldsymbol{s}_i) + \text{constants},$$

or in terms of $\boldsymbol{\beta}$...

$$\ell(\boldsymbol{\beta}) = -\int_D \exp\{\boldsymbol{x}(\boldsymbol{s})^T \boldsymbol{\beta}\}\, d\boldsymbol{s} + \sum_{i=1}^{N(D)} \boldsymbol{x}(\boldsymbol{s}_i)^T \boldsymbol{\beta} + \text{constants}$$

Second-Order Properties

Processes may still exhibit **clustering** or **inhibition**.

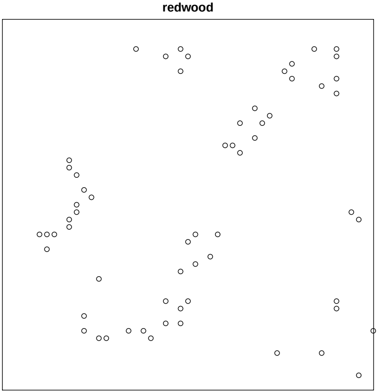

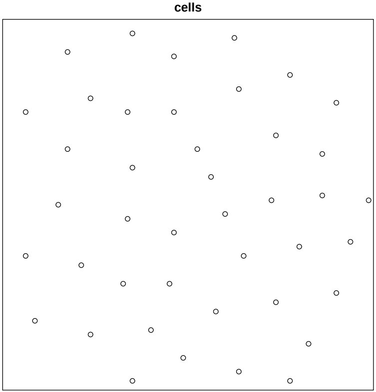

Second-Order Intensity

$$\lambda(\boldsymbol{s}) = \lim_{|d\boldsymbol{s}|\to 0} \frac{\mathrm{E}(N(d\boldsymbol{s}))}{|d\boldsymbol{s}|} \qquad \mu(\boldsymbol{s}) = \mathrm{E}(y(\boldsymbol{s}))$$

$$\lambda_2(\boldsymbol{s}, \boldsymbol{s}') = \lim_{|d\boldsymbol{s}|\to 0} \frac{\mathrm{E}(N(d\boldsymbol{s})N(d\boldsymbol{s}'))}{|d\boldsymbol{s}||d\boldsymbol{s}'|} \qquad \Sigma(\boldsymbol{s}, \boldsymbol{s}') = \mathrm{E}(y(\boldsymbol{s})y(\boldsymbol{s}')) - \mu(\boldsymbol{s})\mu(\boldsymbol{s}')$$

λ_2 is called the **second-order intensity**.

A process is stationary if $\lambda(\boldsymbol{s}) \equiv \lambda$ and $\lambda_2(\boldsymbol{s}, \boldsymbol{s}') = \lambda_2(\boldsymbol{s} - \boldsymbol{s}')$.

Ripley's K-function

$$K(r) = \frac{1}{\lambda}\mathrm{E}\#\{\text{events within distance } r \text{ of a randomly chosen event}\}$$
$$= \frac{1}{\lambda}\mathrm{E}\left[\frac{1}{N(D)}\sum_{i=1}^{N(D)}\sum_{j\neq i} 1\{d(\boldsymbol{s}_i, \boldsymbol{s}_j) \leq r\}\right]$$

This has a natural estimator:

$$\hat{K}(r) = \frac{1}{\hat{\lambda}}\frac{\#\{(i,j) : d(\boldsymbol{s}_i, \boldsymbol{s}_j) \leq r, i \neq j\}}{N(D)}$$

Provides us with a strategy for fitting models:

$$\underset{\theta}{\text{minimize}} \int_0^{r_0} w(r)(\hat{K}(r) - K_\theta(r))^2\,dr$$

Relationship between K and λ_2

If the point process is stationary and isotropic,

$$K(r) = \frac{2\pi}{\lambda^2} \int_0^r \lambda_2(r')r'\,dr'$$

Handling Inhomogeneity

- What if $\lambda(\boldsymbol{s}) \not\equiv \lambda$ but is known?
- Then we might hope $\dfrac{\lambda_2(\boldsymbol{s}, \boldsymbol{s}')}{\lambda(\boldsymbol{s})\lambda(\boldsymbol{s}')} = \rho(||\boldsymbol{s} - \boldsymbol{s}'||)$ is stationary and isotropic.
- New definitions:

$$K_I(r) = \mathrm{E}\left[\frac{1}{|D|}\sum_{i=1}^{N(D)}\sum_{j\neq i}\frac{1\{d(\boldsymbol{s}_i, \boldsymbol{s}_j) \leq r\}}{\lambda(\boldsymbol{s}_i)\lambda(\boldsymbol{s}_j)}\right] = 2\pi\int_0^r \rho(r')r'\,dr'$$

$$\hat{K}_I(r) = \frac{1}{|D|}\sum_{i=1}^{N(D)}\sum_{j\neq i}\frac{1\{d(\boldsymbol{s}_i, \boldsymbol{s}_j) \leq r\}}{\lambda(\boldsymbol{s}_i)\lambda(\boldsymbol{s}_j)}$$

- All computations proceed with K_I and $\hat{K}_I$ instead of K and $\hat{K}$.

Where are we?

What is a Spatio-temporal Point Process?

We now observe locations and times $(\boldsymbol{s}_i, t_i)$.

```
library(stpp)
data(fmd, northcumbria)
fmd <- as.3dpoints(fmd)
plot(fmd, s.region=northcumbria)
```

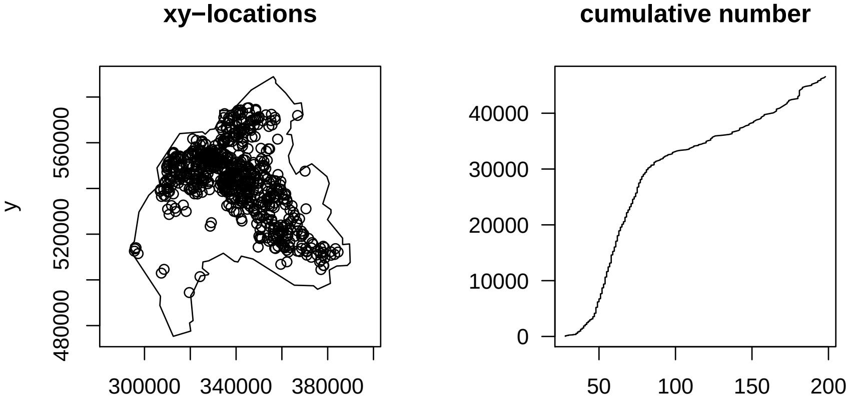

Summarizing spatial and temporal information jointly

We now observe locations and times $(\boldsymbol{s}_i, t_i)$.

```
plot(fmd, s.region=northcumbria, pch=19, mark=T)
```

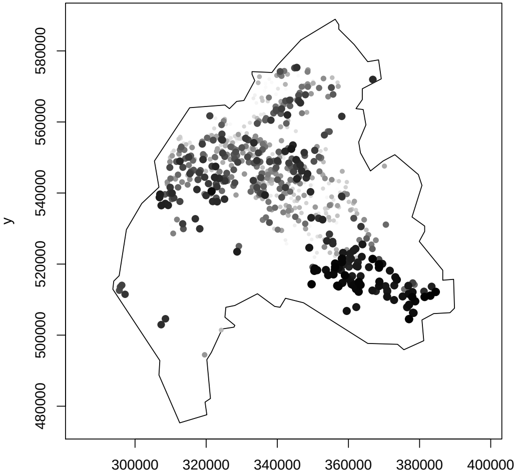

Animations tell the best story

We now observe locations and times $(\boldsymbol{s}_i, t_i)$.

```
animation(fmd, s.region=northcumbria)
```

Where are we?

Spatio-temporal Poisson Process

- Model events as occurring in space time with intensity $\lambda(\boldsymbol{s}, t)$.
- If $A \subset D \times [0, T]$, then $N(A) \sim \text{Pois}\left(\int_A \lambda(\boldsymbol{s}, t)\, d\boldsymbol{s}\, dt\right)$.
- How do we estimate $\lambda(\cdot, \cdot)$?

Estimating the Intensity Function

Exactly the same as before!

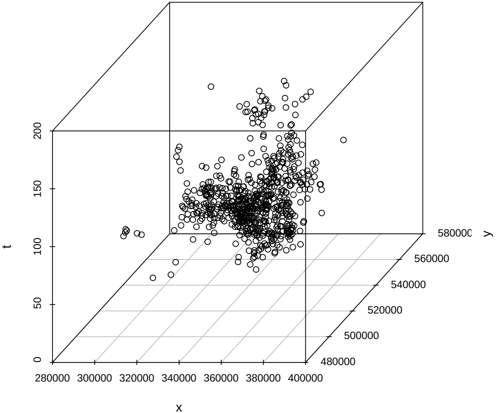

How can differences between space and time be captured when estimating $\lambda(\cdot,\cdot)$?

What if $\lambda(\cdot,\cdot)$ is **separable**?

$$\lambda(\boldsymbol{s},t) = \lambda_{\boldsymbol{s}}(\boldsymbol{s})\lambda_t(t)$$

The need for modeling interactions between observations becomes even more acute in the space-time setting.

Where are we?

Approach 1: K-functions

The second-order spatio-temporal intensity function is:

$$\lambda_2((\boldsymbol{s},t),(\boldsymbol{s}',t')) \stackrel{def}{=} \lim_{|d\boldsymbol{s}|,|d\boldsymbol{s}'|,dt,dt'\to 0} \frac{\mathrm{E}(N(d\boldsymbol{s}\times dt)N(d\boldsymbol{s}'\times dt'))}{|d\boldsymbol{s}||d\boldsymbol{s}'|\,dt\,dt'}$$

If the point process is stationary and isotropic in space and in time:

$$\begin{aligned} K(r,h) &= \frac{1}{\lambda}\mathrm{E}\begin{pmatrix}\#\text{ events within radius } r \text{ and time } h \\ \text{of randomly chosen event}\end{pmatrix} \\ &= \frac{1}{\lambda}\mathrm{E}\left[\frac{1}{N}\sum_{i=1}^{N}\sum_{j>i} 1\{d(\boldsymbol{s}_i,\boldsymbol{s}_j)\le r\}1\{t_j - t_i \le h\}\right] \end{aligned}$$

which can be estimated in the usual way. We also have the relation:

$$K(r,h) = \frac{2\pi}{\lambda^2}\int_0^h\int_0^r \lambda_2(r',h')r'\,dr'\,dh'$$

Approach 1: K-functions

- Suppose we have a process that depends on some parameters θ.
- e.g., Parents generate $S \sim \mathrm{Pois}(\mu)$ offspring. Each of the S offspring appear at times T_i according to a Poisson process with rate r.
- If we can calculate the theoretical K-function K_θ, then we can estimate θ by solving

$$\underset{\theta}{\text{minimize}} \int_0^{r_0} \int_0^{h_0} w(r,h)(\hat{K}(r,h) - K_\theta(r,h))^2$$

Approach 2: Conditional Intensity Function

- In space-time, it's natural to model the conditional intensity given the past $\mathcal{H}_t = \{(\boldsymbol{s}_i, t_i) : t_i < t\}$.

$$\lambda_c(\boldsymbol{s}, t|\mathcal{H}_t) = \lim_{|d\boldsymbol{s}|\to 0, dt\to 0} \frac{\mathrm{E}(N(d\boldsymbol{s} \times dt)|\mathcal{H}_t)}{|d\boldsymbol{s}|dt}$$

- For a Poisson process, $\lambda_c(\boldsymbol{s}, t|\mathcal{H}_t) = \lambda(\boldsymbol{s}, t)$.
- e.g., pairwise interaction model:

$$\lambda_c(\boldsymbol{s}, t|\mathcal{H}_t) = \alpha(t) \prod_{i=1}^{|\mathcal{H}_t|} h_\theta(\boldsymbol{s}, \boldsymbol{s}_i)$$

1. $h_\theta(\boldsymbol{s}, \boldsymbol{s}_i) = 1 - e^{-\theta||\boldsymbol{s}-\boldsymbol{s}_i||}$
2. $h_\theta(\boldsymbol{s}, \boldsymbol{s}_i) = e^{-\theta||\boldsymbol{s}-\boldsymbol{s}_i||}$

Approach 2: Conditional Intensity Function

We typically estimate the intensity function using maximum likelihood. The log-likelihood of $(\boldsymbol{s}_i, t_i)$ is

$$\ell(\lambda_c(\cdot,\cdot|\mathcal{H}_\cdot)) = -\int_0^T \int_D \lambda_c(\boldsymbol{s}, t|\mathcal{H}_t)\, d\boldsymbol{s}\, dt + \sum_{i=1}^n \log \lambda_c(\boldsymbol{s}_i, t_i|\mathcal{H}_{t_i}) + \text{constants}.$$

NB. The integral will probably have to be approximated.

Where are we?

Summary

- Spatio-temporal point processes are a messy and emerging field.
- Poisson processes are not viable models for spatio-temporal processes; must take into account interactions across space-time.
- There are two main approaches for modeling and fitting interaction models: K-functions and conditional intensity.
- The `stpp` package contains many useful routines for visualizing, simulating, and (less so) fitting spatio-temporal point process models.

Homeworks

- Homework 3 due Friday. Don't forget to upload your files to the website (should be in .csv format).
- Homework 4a will be posted tonight.