



HAL
open science

Spatio-temporal point process intensity estimation using zero-deflated subsampling applied to a lightning strikes dataset in France

Jean-François Coeurjolly, Thibault Espinasse, Anne-Laure Fougères, Mathieu Ribatet

► To cite this version:

Jean-François Coeurjolly, Thibault Espinasse, Anne-Laure Fougères, Mathieu Ribatet. Spatio-temporal point process intensity estimation using zero-deflated subsampling applied to a lightning strikes dataset in France. 2024. hal-04502501

HAL Id: hal-04502501

<https://hal.science/hal-04502501>

Preprint submitted on 13 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Spatio-temporal point process intensity estimation using zero-deflated subsampling applied to a lightning strikes dataset in France

Jean-François Coeurjolly*, Thibault Espinasse,
Anne-Laure Fougères †, Mathieu Ribatet ‡

Abstract

Cloud-to-ground lightning strikes observed in a specific geographical domain over time can be naturally modeled by a spatio-temporal point process. Our focus lies in the parametric estimation of its intensity function, incorporating both spatial factors (such as altitude) and spatio-temporal covariates (such as field temperature, precipitation, etc.). The events are observed in France over a span of three years. Spatio-temporal covariates are observed with resolution $0.1^\circ \times 0.1^\circ$ ($\approx 100\text{km}^2$) and six-hour periods. This results in an extensive dataset, further characterized by a significant excess of zeroes (i.e., spatio-temporal cells with no observed events). We reexamine composite likelihood methods commonly employed for spatial point processes, especially in situations where covariates are piecewise constant. Additionally, we extend these methods to account for zero-deflated subsampling, a strategy involving dependent subsampling, with a focus on selecting more cells in regions where events are observed. A simulation study is conducted to illustrate these novel methodologies, followed by their application to the dataset of lightning strikes.

*Univ. Grenoble Alpes, CNRS, LJK, 38000 Grenoble, France;
jean-francois@coeurjolly@univ-grenoble-alpes.fr

†Universite Claude Bernard Lyon 1, ICJ UMR5208, CNRS, Ecole Centrale de Lyon, INSA Lyon, Université Jean Monnet, 69622 Villeurbanne, France;
espinasse@math.univ-lyon1.fr, fougères@math.univ-lyon1.fr

‡Ecole Centrale de Nantes, Département Informatique et Mathématiques, Laboratoire de Mathématiques Jean Leray, 1, rue de la Noe 44321 Nantes Cedex, France;
mathieu.ribatet@ec-nantes.fr

Keywords and phrases: Spatio-temporal point process; Composite likelihood; Subsampling; High-dimensional data.

1 Introduction

Severe thunderstorms can be associated with personal injury or costly infrastructure damage. Among the major risks to be assessed are those associated with high lightning intensity, and they have been studied for about twenty years (see e.g. [Curran et al. \(2000\)](#), [Mona et al. \(2016\)](#), [Schulz et al. \(2005, 2016\)](#), [Simon et al. \(2017\)](#), [Hernandez-Magallanes and Genton \(2019\)](#), [Nampak et al. \(2021\)](#)). Additionally, the interest in evaluating the intensity of cloud-to-ground lightning strikes might increase in the context of an energy transition: wind farms are multiplying, and the risk of lightning is one of the five hazards typically identified by manufacturers. From a physical standpoint, a thunderstorm corresponds to the dissipation of energy involving various thermodynamic processes, such as convection. To describe these complex processes, several covariates are known to be useful: the CAPE (Convective Available Potential Energy), which acts as the 'fuel' in the development of the cloud; the θ'_w at 850 hPa (Wet-bulb potential temperature at 850 hPa), which is known to be a good synoptic predictor, giving important information on the position of the fronts; the temperature at two different altitudes (20m and 1500m), providing information on the vertical profile of the air mass (convection); humidity; and wind (zonal and meridional components). A few other purely spatial covariates like the altitude, the distance to the sea should also have an influence (see Section 2 for more details). Lightning strike impacts, although continuously studied by scientists in the fields of physics, climatology or statistics remains a phenomenon with an important part of randomness. The set of strikes collected over time and space constitute a spatio-temporal point pattern and the goal of this paper is to characterize its distribution in terms of climatological, topographical covariates.

Spatio-temporal point processes are the stochastic models generating such data. Point processes in (a Polish) space S ([Daley et al., 2003](#), [Møller and Waagepetersen, 2003](#)) model points, events, objects in interaction. The data consist in a set $\mathbf{y} = \{y_1, \dots, y_n\}$ with $y_i \in S$. When $S = \mathbb{R}^+, \mathbb{R}^d$ ($d \geq 1$) or $\mathbb{R}^d \times \mathbb{R}^+$ we speak of temporal, spatial or spatio-temporal point processes. Most of statistical methodologies and their implementation have been devoted to temporal or spatial point processes ([Møller and Waagepetersen,](#)

2003, Illian et al., 2008, Baddeley et al., 2015). Spatio-temporal are more challenging and modern due to the increase of data collection and have also received a lot of attention (Diggle, 2006, Cressie and Wikle, 2015, González et al., 2016).

The aim of this paper is to explore first-order structure of the lightning strikes dataset. The intensity function, say $\rho(y)$ for $y \in \mathbb{R}^d \times \mathbb{R}^+$, measures the mean local number of events in the vicinity of a space-time point y . A point pattern is said to be homogeneous if the intensity is constant and inhomogeneous otherwise. Our objective is to model parametrically the function ρ using spatial and spatio-temporal covariates, to estimate these parameters and to produce spatial prediction maps. This work can be viewed as a climatological risk assessment contribution. We characterize the distribution of lightning strikes using covariates. We do not intend to produce a short-term forecast of lightning strikes activity.

Before detailing further our approach and contributions, let us first highlight a few facts about the dataset (detailed in Section 2): we observe more than one million of impacts between 2013 and 2015, together with 6 spatial covariates and 7 spatio-temporal covariates. The spatio-temporal resolution is $v = 0.1^\circ \times 0.1^\circ \times 6\text{h}$. As outlined in Section 2, the point pattern is highly inhomogeneous in time with daily and seasonal effects, and in space. In addition, the point pattern exhibit some (natural) clustering effect not explained by inhomogeneity. Together these last two facts are the reasons why approximately 99% of cells of size v do not contain any data point.

We consider an exponential family model for the intensity function ρ , see Section 3 and in particular Equation (3.1). On the one hand, the clustering effect prevents us from modelling the data by a spatio-temporal Poisson point processes, the standard reference process which models events without any interaction (see e.g. Daley et al. (2003)). On the other hand, maximum likelihood estimation is also known (and in particular for spatio-temporal point processes) to be very computationally expensive (Møller and Waagepetersen, 2003). Composite likelihood methods offer excellent alternatives. They were not all developed for spatio-temporal point processes but the extension is in principle straightforward. Among methods, we have the Poisson likelihood (Rathbun and Cressie, 1994, Baddeley and Turner, 2000, Schoenberg, 2005, Waagepetersen and Guan, 2009), the logistic regression likelihood when one focuses not on the number of impacts per cell but on the presence/absence of impacts, a standard approach known as spatial pixel-based logistic regression and commonly used in the standard GIS community (see e.g. Baddeley et al.,

2010, and the many references therein on this topic), the conditional logistic regression likelihood estimation (Waagepetersen, 2008, Baddeley et al., 2014), the quasi-likelihood method (Guan et al., 2015) or the variational method (Coeurjolly and Møller, 2014). We review and compare these methods (except the last two ones which are less appropriate in the context of this application as explained in Section 4.5) in Section 4. More precisely, we particularize the different methods and discuss implementation issues, in the context of discretized covariates (ie piecewise constant covariates on cells of volume v). As noticed and studied by Baddeley et al. (2010), many papers in environmental sciences, ecology and, say, research areas using data based on Geographical Information Systems (GIS) analyse spatial point pattern data at a pixel level (Elliot et al., 2000, Wakefield, 2004, Waller and Gotway, 2004). Such a framework is often encountered in practice, see e.g. Raeisi et al. (2021). Our analysis in Section 4 yields interesting facts and links with generalized linear models. Some of these facts are known and can be found in Baddeley and Turner (2000), Baddeley et al. (2010, 2014). In particular, we demonstrate that the Poisson likelihood method, acknowledged for experiencing slight bias in the presence of continuous covariates due to an approximation in quasi-Poisson regression implementation, is precisely equivalent to a Poisson regression with a specific offset term when covariates are discretized.

The resolution of the covariates, combined with the fact that the observed phenomenon is highly inhomogeneous in time and space explain the overwhelming dominance of zero values of the indicators of presence/absence of an impact in a spatio-temporal cell. As detailed in Section 2, whatever the season or six-hour periods, we observed that approximately 99% of cells with volume v do not contain any data point. When one applies any method described above to the dataset or to any independent subsample (to reduce computational cost inherent to this large dataset), this causes numerical problems and instabilities. A practical strategy that is common in the GIS literature (Atkinson and Massari (1998), Gorsevski et al. (2006), see also Baddeley et al. (2015, Section 9.10.3)) when pixel-based logistic regression is considered is to take a sample of pixels where zero values are observed. Then logistic regression is applied to the data consisting of the ‘1’ pixels (pixels containing at least one data point) and the sampled ‘0’ pixels. The main contribution of the present paper is to revisit standard composite likelihood-methods for point patterns (temporal, spatial or spatio-temporal) when such ideas are considered. It is worth specifying that we do not aim at modelling

this excess of zeroes. Instead one aims at defining subsamples where the '0' (here) voxels can be deflated compared to the '1' voxels. Section 5 details the construction of such subsamples, called zero-deflated subsamples, and shows how these ideas can be applied to all composite likelihood methods presented in Section 4.

The rest of the paper is organised as follows. We detail the dataset and the problems it implies in Section 2. Section 3 provides a background on spatio-temporal point processes and specifies the parametric model considered in this paper. Section 4 reviews composite likelihood methods used to estimate the intensity function in the particular case where covariates are discretized. Section 5 defines the concept of zero-deflated subsamples in particular when they are applied to methods defined in Section 4. We end Section 5 with a brief simulation study showing the efficiency of the proposed methodologies when they are applied to zero-deflated subsamples even if we have a dataset with a large number of '0'. Section 6 proposes an application to the lightning strikes dataset. We first present criteria to measure quality of spatio-temporal predictions, one of them being very close to the Wasserstein distance largely used in the statistics community to compare two empirical measures. Second, we compare methods described in Section 5 in terms of computational cost and prediction quality. Our finding is that the subsampled version of the Poisson likelihood turns out to be the best one. This method is then used to illustrate temporal prediction curves and spatial prediction maps. Several perspectives can be addressed based on the present work. Some of them are described in Section 7. Finally, Appendix A gathers additional figures regarding the exploratory data analysis and simulation study.

2 Presentation of data and their complexity

In this section, our objective is to shed the light on the challenges the data of interest impose in terms of statistical methodology. Data are provided by Météorage and Météo-France. We observe impacts of cloud-to-ground lightning flashes over France from 2010 to 2015. Longitudes and latitudes are collected to a 100m accuracy and time events upto the nanosecond. Overall, we observe 1846533 events. Figure 1 is a brief illustration of this spatio-temporal dataset, from which we can readily see that this phenomenon turns out to be highly inhomogenous in time and space.

Even if a more rigorous background on point processes and intensity func-

Lightning strikes Sep.15–Sep.31, 2015




Figure 1: (Top) Map of France with altitudes and locations of impacts during the second half of September 2015. The dashed black box (resp. blue, red, orange) corresponds to the domain where impacts are observed (resp. to the subdomain called ‘Alps’, ‘Channel’ and ‘Pyrénées’ in Appendix A.1). The black triangle corresponds to the location denoted by x_0 in Figure 8. (Bottom) Spatio-temporal representation of lightning strikes impacts for different months in 2010. Only 20% of (randomly chosen) impacts are represented for August 2010.

tions is provided in Section 3, Figures 6 and 7 are quite easy to interpret. Figure 6 focuses on the temporal nature of the dataset (by aggregating data over space). A yearly seasonal effect is clearly seen with a much higher intensity during summer (June to August) than in spring (April to May) and autumn (September to November). The winter period (December to March) seems to be the most quiet period. As shown in Figure 6 (bottom), the temporal intensity is quite variable for different regions (Alps, Channel, Pyrénées for instance). Finally, it is worth pointing out that within a day the temporal intensity varies a lot. We notice a slow activity during the night and a higher activity in the afternoon/evening, a pattern which occurs quite similarly for the different considered regions and seasons. Figure 7 is more focused on the spatial nature of the dataset. We aggregate data by month and estimate non parametrically the spatial intensity (using a standard kernel intensity estimation method, e.g. Diggle (1985), Møller and Waagepetersen (2003)). We retrieve the previous comment that the activity is higher during warmer months. This figure illustrates more strenuously the spatial differences on the French territory.

Besides the obvious inhomogeneous nature of this dataset, Figure 2 explores the dependence between lightning strikes. This constitutes only a part of a larger exploratory analysis. Here, we have aggregated data by month and have estimated standard spatial summary statistics such as the inhomogeneous K and J functions (see Møller and Waagepetersen (2003), Illian et al. (2008), Van Lieshout (2011)). To estimate these functional summary statistics, we use kernel intensity estimates as plug-in estimates. We do not provide global envelopes tests (Myllymäki et al. (2017)) to not overload the graphs. The general conclusion is however crystal-clear: the point pattern cannot be simply modelled by a spatio-temporal inhomogeneous Poisson point process (otherwise monthly aggregated spatial point processes would be spatial Poisson point processes as well, which is not the conclusion from Figure 2). It is not the objective of the present paper to model second-order characteristics of the lightning strikes dataset. However, the previous comment means that if we estimate parametrically the intensity of the point pattern, we have to consider an estimation method which is robust to the model.

In addition to this spatio-temporal point pattern, we also have at our disposal spatial covariates and spatio-temporal covariates used to explain the distribution of lightning strikes. Spatial covariates are observed on a regular grid of $0.1^\circ \times 0.1^\circ$ ($\approx 100\text{km}^2$). They are constituted by the altitude map, the distance to the sea, as well as four additional variables computed from

the AURELHY method (Bénichou, 1994). Their interest has been shown recently, see e.g. Taillardat and Mestre (2020), for statistical post-processing of ensemble forecasts. Spatio-temporal covariates consist of model outputs (Météo-France ARPEGE¹ Model output) of climate covariates that are available only from 2013 to 2015. We have access to the following covariates observed on a regular grid of $0.1^\circ \times 0.1^\circ \times 6\text{h}$: CAPE (Convective Available Potential Energy); Humidity (at altitude 1500m); Temperatures (at altitudes 20m and 1500m); θ'_w at 850 hPa (Wet-bulb potential temperature at 850 hPa); zonal and meridional components of wind. The relevance of these latter covariates has been detailed in the Introduction. The altitude map is illustrated in Figure 1. Similarly to the point pattern, these covariates obviously vary a lot, as shown in Figure 8. Figure 9 explores the interest of some of these covariates to explain the presence/absence of a lightning strike at a specific location.

As noticed and studied by Baddeley et al. (2010), many papers in environmental sciences, ecology and, say, research areas using data based on Geographical Information Systems (GIS) analyse spatial point pattern data at a pixel level. Such a framework is often encountered in practice, see e.g. Raeisi et al. (2021) and in particular in the current paper, as spatio-temporal covariates are collected within a 0.1 degree squared and a 6-hour period. Given the observed phenomenon (remind that data are collected up to the nanosecond), this spatio-temporal grid appears quite coarse. The spatio-temporal observation domain $W \times \mathcal{T}$ (where W is the spatial domain and $\mathcal{T} = [2013, 2015]$) can therefore be viewed as a tessellation of spatio-temporal cells Δ_j ($j = 1, \dots, J$) over which covariates are constant. Given the resolution of spatio-temporal covariates, we have $J \approx 7 \times 10^7$. The following comment is crucial and constitutes the core of the present paper: despite the apparent high number of events (approximately 10^6 lightning strikes impacts for data restricted to the period 2013-2015 where all covariates and spatio-temporal covariates are observed), most of the cells Δ_j do not contain any data. This is illustrated by Figure 3: even if one splits data by season and 6-hour periods, more than 98% of cells Δ_j are empty.

To summarize this section, we observe 1050342 spatio-temporal events of cloud-to-ground lightning strikes in France over the period 2013-2015, 6 discretized spatial covariates and 7 discretized spatio-temporal covariates. The point pattern is highly inhomogeneous both in time and space and the

¹See e.g. <https://www.umn-cnrm.fr/spip.php?article121&lang=en>




Figure 2: Non parametric estimation of the spatial inhomogeneous J and Ripley's K functions. Data are aggregated by months for July and August from 2011 to 2015. The dashed line and curve correspond to the expected J and K functions under the Poisson assumption. An estimate of the J function (resp. K function) below 1 (resp. above πr^2) which corresponds to the situation under the Poisson case indicates clustering in data, that is strong positive spatial correlation between events (ie lightning strikes).




Figure 3: Number of impacts (per thousands) by season and 6-hour period and proportion of the spatio-temporal voxel-grid cells containing no lightning strike impact.

events also seem to exhibit some strong clustering (hence a departure to the independence assumption). Finally, at the voxel-grid level (i.e. at the resolution level), the dataset exhibits a huge excess of zeroes, in the sense that most of space-time observation domain is empty. All these different observations are studied in the next three sections before going back to the dataset in Section 6.

3 Notation and intensity modelling

A (planar) spatio-temporal point process \mathbf{Y} defined in $\mathbb{R}^2 \times \mathbb{R}^+$ can be viewed as a locally finite random measure on $\mathbb{R}^2 \times \mathbb{R}^+$ (see e.g. Daley et al. (2003), Møller and Waagepetersen (2003)), which means that for any bounded $W \subset \mathbb{R}^2$ and $\mathcal{T} \subset \mathbb{R}^+$, the number of (space-time) points is bounded. We also assume that \mathbf{Y} is simple, in the sense that two events do not occur at the same space-time point. In the following, we use the terminology point for an element of \mathbf{Y} , thus for a space-time point, and the notation $y = (x, t)$ where $x \in W$ and $t \in \mathcal{T}$ for an element of \mathbf{Y} .

Estimating first-order summary statistics is usually the first step to analyse a point process. The first-order intensity function ρ of \mathbf{Y} (assumed to exist) is a function of $y = (x, t)$ and is defined by the following integral characterization (also called Campbell theorem), see [Møller and Waagepetersen \(2003\)](#): for any non-negative function f ,

$$\mathbb{E} \sum_{y \in \mathbf{Y}} f(y) = \int_{\mathbb{R}^2 \times \mathbb{R}^+} f(y) \rho(y) dy$$

where we abuse notation by denoting $dy = dx dt$. In other words, using the notation $N(A)$ for the counting variable of any bounded borel set $A \subset \mathbb{R}^2 \times \mathbb{R}^+$

$$\rho(y) = \rho(x, t) = \lim_{dx \rightarrow 0, dt \rightarrow 0} \frac{\mathbb{E}\{N(dx \times dt)\}}{dx dt}$$

describes the local number of points in the vicinity of $y = (x, t)$. Higher-order intensity functions can be described in a similar way (using higher-order version of Campbell theorem). In this paper, we are interested in modelling ρ as a parametric form of spatial and spatio-temporal covariates (which themselves are realizations of spatial or spatio-temporal random fields) denoted here by $\mathbf{Z}(x) \in \mathbb{R}^{n_Z}$ and $\mathbf{C}(x, t) \in \mathbb{R}^{n_C}$. Hence, n_Z and n_C stand for the number of spatial and spatio-temporal covariates. Given $\mathbf{C}(x, t) = \mathbf{c}$ and $\mathbf{Z}(x) = \mathbf{z}$, we model ρ as a function of $x, t, \mathbf{c}, \mathbf{z}$, i.e. $\rho(y) = \rho((x, t), \mathbf{c}, \mathbf{z})$.

Given the explanatory study proposed in [Section 2](#), we propose to ‘aggregate’ data by season and 6-hour period of the day. Following this, we let $\mathcal{S} = \{\text{summer, fall, winter, spring}\}$ and $\mathcal{M} = \{1, \dots, 4\}$ denote the set of seasons (as described in the previous section) and the set of six-hour periods (0:00-06:00, 06:00-12:00, 12:00-18:00 and 18:00-0:00) of the day. And we view \mathbf{Y} as a union of marked point processes where the mark is an element of $\mathcal{S} \times \mathcal{M}$ (which actually depends on time), that is $\mathbf{Y} = \bigcup_{(s,m) \in \mathcal{S} \times \mathcal{M}} \mathbf{Y}_{s,m}$ where $\mathbf{Y}_{s,m}$ can be described as $\mathbf{Y}_{s,m} = \{(x, t) \in \mathbf{Y} : x \in W \text{ and } t \in \mathcal{T}_{s,m}\}$ where $\mathcal{T}_{s,m}$ is the union of time intervals restricted to the season s and the six-hour period m . Note that we do not assume that $\mathbf{Y}_{s,m}$ is independent of $\mathbf{Y}_{s',m'}$. However, we assume that the intensity function ρ of \mathbf{Y} has the form

$$\rho(y) = \sum_{s \in \mathcal{S}, m \in \mathcal{M}} \rho_{s,m}(y) \quad \text{with} \quad \rho_{s,m}(y) = \exp(\boldsymbol{\beta}_{s,m}^\top \{\mathbf{Z}(x), \mathbf{C}(x, t)\}) \quad (3.1)$$

for any $y = (x, t) \in W \times \mathcal{T}_{s,m}$ with $\boldsymbol{\beta}_{s,m} \in \mathbb{R}^{n_Z + n_C}$. We do not make any further assumption on the spatio-temporal point process. In particular, we

have to keep in mind that $\mathbf{Y}_{s,m}$ is probably not a Poisson point process, as revealed by Figure 2.

The estimation of ρ from (3.1) is tackled by estimating separately $\rho_{s,m}$ for any s, m . In the following two Sections 4-5, we focus on the specific problem of a single intensity. To make it general, we slightly simplify the notation by discarding the subscripts s, m to ease the reading and omitting (without loss of generality) the spatial covariates \mathbf{Z} . These sections are therefore written for general spatio-temporal point processes \mathbf{Y} with intensity modelled from an exponential family model with discretized spatio-temporal covariates.

4 Composite likelihoods methods with discretized covariates

Let \mathbf{Y} be a spatio-temporal point process on $W \times \mathcal{T}$. We consider the exponential model for its intensity given by $\rho(y) = \rho(x, t) = \exp(\boldsymbol{\beta}^\top \mathbf{C}(x, t))$ for any $y = (x, t) \in W \times \mathcal{T}$, where $\mathbf{C}(x, t) \in \mathbb{R}^{n_C}$. In particular, we assume that the covariates are deterministic (to ease the presentation) and are constant on Δ_j , that is $\mathbf{C}(x, t) = \mathbf{C}_j \in \mathbb{R}^{n_C}$ whenever $(x, t) \in \Delta_j$, where the cells Δ_j , with volume δ_j , form a tessellation of $W \times \mathcal{T}$, that is $W \times \mathcal{T} = \bigcup_{j=1}^J \Delta_j$. Note that the statements of Sections 4-5 are obviously valid for purely temporal or spatial point processes, or more generally for any point process in \mathbb{R}^d .

The methods we consider fall into the terminology of composite likelihoods methods and are quite standard in the literature, see e.g. Coeurjolly and Lavancier (2019) for a review. However we particularize them and discuss implementation issues, in the context of discretized covariates, which yields interesting facts. Some of these facts can be found in Baddeley and Turner (2000), Baddeley et al. (2010, 2014). A summary and comparison of methods described below is proposed in Table 1.

4.1 Poisson likelihood (method PL)

To estimate the vector $\boldsymbol{\beta}$, it is now well-known (see e.g. Waagepetersen and Guan (2009) or Coeurjolly and Lavancier (2019)) that the maximum of the Poisson likelihood exhibits interesting properties such as consistency, asymptotic normality even if the underlying point process is not a Poisson point process. The main argument comes from Campbell theorem which in particular states that the score of the Poisson likelihood is unbiased for general

point processes. For general intensity models, the Poisson likelihood writes

$$\text{PL}(\boldsymbol{\beta}) = \sum_{y \in \mathbf{Y}} \log \rho(y) - \int_{W \times \mathcal{T}} \rho(y) dy \quad (4.1)$$

which, for discretized covariates and exponential family models (3.1), reduces to

$$\begin{aligned} \text{PL}(\boldsymbol{\beta}) &= \sum_{j=1}^J \{N_j \boldsymbol{\beta}^\top \mathbf{C}_j - \delta_j \exp(\boldsymbol{\beta}^\top \mathbf{C}_j)\} \\ &= \sum_{j=1}^J \{N_j \log(\delta_j \exp(\boldsymbol{\beta}^\top \mathbf{C}_j)) - \delta_j \exp(\boldsymbol{\beta}^\top \mathbf{C}_j)\} - \sum_{j=1}^J N_j \log \delta_j. \end{aligned} \quad (4.2)$$

The first sum of (4.2) precisely corresponds to a Poisson regression likelihood with the canonical link and with offset term $(\log \delta_j)_j$. Therefore, when covariates are piecewise constant, the Poisson (or composite) likelihood estimator is strictly equivalent to a Poisson regression (with an offset term).

It is worth mentioning that when covariates are not piecewise constant, a numerical problem appears since one has to discretize the integral term of (4.1). [Baddeley and Turner \(2000\)](#) use Berman-Turner's approximation ([Berman and Turner, 1992](#)) and show that if one uses as quadrature points the union of a set of grid points and data points, Equation (4.1) can be implemented via a quasi-Poisson regression model. However, using data points in the integral discretization step leads to an approximation which can be problematic especially when the number of data points is high, as noticed by [Baddeley et al. \(2014\)](#).

4.2 Conditional logistic regression and its weighted version (methods CLRL and WCLRL)

When covariates are continuous, the implementation of the Poisson likelihood induces some bias. To overcome that numerical problem, [Waagepetersen \(2008\)](#), [Baddeley et al. \(2014\)](#) propose an alternative estimating function which can be implemented using a logistic procedure. Following [Waagepetersen \(2008\)](#), one can consider a dummy point process given by a stratified binomial point process with J points, one in each cell Δ_j , that is the point process, say \mathbf{D} , with J points and with intensity $\eta(x) = \sum_{j=1}^J \delta_j^{-1} 1(x \in \Delta_j)$. We then

leave the reader to check that the conditional logistic regression likelihood (for spatial point processes) is equal to

$$\text{CLRL}(\boldsymbol{\beta}) = \sum_{y \in \mathbf{Y}} \log \left(\frac{\rho(y)}{\eta(y) + \rho(y)} \right) + \sum_{y \in \mathbf{D}} \log \left(\frac{\eta(y)}{\eta(y) + \rho(y)} \right). \quad (4.3)$$

Equation (4.3) is indeed a logistic regression likelihood with a sample of $N + J$ Bernoulli random variables with probability $\pi(y) = \rho(y)/(\eta(y) + \rho(y))$ if $y \in \mathbf{Y}$ and $1 - \pi(y)$ if $y \in \mathbf{D}$. Let us see how (4.3) rewrites for piecewise constant covariates. Let $\mathbf{Y} = \{y_1, \dots, y_N\}$, $\mathbf{D} = \{d_1, \dots, d_J\}$ and $\rho_j = \exp(\boldsymbol{\beta}^\top \mathbf{C}_j)$, for $j = 1, \dots, J$. For any sequence a_1, \dots, a_N , we define the sequence $\check{a}_1, \dots, \check{a}_{N+J}$ by

$$\check{a}_k = \begin{cases} \sum_{l=1}^N a_l \mathbf{1}(y_k \in \Delta_l) & \text{if } k = 1, \dots, N \\ a_{k-N} & \text{if } k = N + 1, \dots, J \end{cases}$$

Using this notation, (4.3) can be rewritten as

$$\text{CLRL}(\boldsymbol{\beta}) = \sum_{k=1}^{N+J} \left\{ \check{i}_k \log \left(\frac{\check{\delta}_k \check{\rho}_k}{1 + \check{\delta}_k \check{\rho}_k} \right) + (1 - \check{i}_k) \log \left(\frac{1}{1 + \check{\delta}_k \check{\rho}_k} \right) \right\} \quad (4.4)$$

where $\check{i}_k = 1$ if $k = 1, \dots, N$ and 0 otherwise. Formally, the latter equation corresponds to a logistic regression with Bernoulli random variables \check{I}_k with mean $\text{P}(\check{I}_k = 1) = \check{\delta}_k \check{\rho}_k / (1 + \check{\delta}_k \check{\rho}_k)$. [Baddeley et al. \(2014\)](#) show that, for general spatial covariates and for any spatial point process, this procedure leads to an accurate estimation which is robust to the Poisson assumption. There is no integral discretization. It is also shown that if J is large, the conditional logistic regression likelihood tends almost surely to the Poisson likelihood. Finally, since covariates are constant on Δ_j , it is worth noting that (4.4), which is a logistic regression with response vector of length $N + J$, can be rewritten as

$$\text{WCLRL}(\boldsymbol{\beta}) = \sum_{j=1}^{\#J_1+J} \tilde{w}_j \left(\tilde{I}_j \log(\tilde{p}_j) + (1 - \tilde{I}_j) \log(1 - \tilde{p}_j) \right) \quad (4.5)$$

where $\#J_1$ is the number of elements of $J_1 = \{j : N_j > 0\}$ and where

$$(\tilde{w}_j, \tilde{I}_j, \tilde{p}_j) = \begin{cases} (N_{j_1(j)}, 1, p_{j_1(j)}) & \text{if } j \leq \#J_1 \\ (1, 0, p_{j-\#J_1+1}) & \text{if } J_1 < j \leq \#J_1 + J \end{cases}$$

where we use the notation $j_1(j)$ for the j th index of $\{1, \dots, J\}$ for which $N_j > 0$. Equation (4.5) can be simply implemented as a weighted logistic regression likelihood with canonical link. This implementation is faster since the response vector of the regression (4.5) is now $\#J_1 + J$ which can be much smaller than $N + J$ when a large number of cells Δ_j do not contain any point.

4.3 Bernoulli regression likelihood with logit link (method $\text{BRL}_{\text{logit}}$)

We continue with a method often considered in the literature (see e.g. [Baddeley et al. \(2010\)](#)), known as the pixel logistic regression. Let us start with Equation (4.4), which can be written using (4.3) as

$$\sum_{j=1}^J \{N_j \log(p_j) + \log(1 - p_j)\} \quad \text{with } p_j = \frac{\delta_j \rho_j}{1 + \delta_j \rho_j}. \quad (4.6)$$

The approximation suggested in the literature simply consists in replacing N_j by $I_j = 1(N_j > 0)$, i.e. the number of impacts in Δ_j is replaced by the presence/absence of an impact in Δ_j . This yields the criterion

$$\text{BRL}_{\text{logit}}(\boldsymbol{\beta}) = \sum_{j=1}^J \{I_j \log(p_j) + (1 - I_j) \log(1 - p_j)\}. \quad (4.7)$$

In other words, (4.7) is the likelihood of the Bernoulli random variables I_j if one (wrongly) assumes that $\text{logit}\mathbb{P}(I_j = 1) = \log(\delta_j) + \boldsymbol{\beta}^\top \mathbf{C}_j$. Replacing N_j by I_j in (4.3) is definitely a rough approximation which is difficult to quantify without any assumption on the underlying spatial point process \mathbf{Y} . This approach is also not robust to the Poisson assumption for \mathbf{Y} .

4.4 Bernoulli regression likelihood with complementary log-log link (method $\text{BRL}_{\text{cloglog}}$)

The previous approximation can be alleviated if we link correctly $\mu_j = \mathbb{P}(I_j = 1)$ to $\boldsymbol{\beta}$. A model is required for this. If one assumes that the variables N_j are independent and distributed as Poisson random variables, then standard

calculation shows that $g(\mu_j) = \boldsymbol{\beta}^\top \mathbf{C}_j + \log \delta_j$ where g is the log-log complementary function given by $g(t) = \log(-\log(1-t))$ for $t \in (0, 1)$. Then, the log-likelihood based on I_1, \dots, I_J is equal to

$$\begin{aligned} \text{BRL}_{\text{cloglog}}(\boldsymbol{\beta}) &= \sum_{j=1}^J \{I_j \log(\mu_j) + (1 - I_j) \log(1 - \mu_j)\} \\ &= \sum_{j=1}^J \left\{ I_j \log \left(\frac{\mu_j}{1 - \mu_j} \right) + \log(1 - \mu_j) \right\}. \end{aligned} \quad (4.8)$$

Therefore, if \mathbf{Y} is a (spatio-temporal) Poisson point process, $\boldsymbol{\beta}$ can be indeed (consistently) estimated using a logistic regression with log-log complementary link and offset term $(\log \delta_j)_j$. However, (4.8) is not a composite likelihood and so is not robust at all to the Poisson assumption. Indeed, the score of $\text{BRL}_{\text{cloglog}}(\boldsymbol{\beta})$ is equal to

$$\text{BRL}_{\text{cloglog}}^{(1)}(\boldsymbol{\beta}) = \sum_{j=1}^J C_j \log(1 - \mu_j) \left(\frac{I_j}{\mu_j} - 1 \right)$$

whereby we deduce, by denoting $\tilde{\mu}_j = \mathbb{P}(N_j > 0)$, that

$$\mathbb{E} \left(\text{BRL}_{\text{cloglog}}^{(1)}(\boldsymbol{\beta}) \right) = \sum_{j=1}^J \mathbf{C}_j \log(1 - \mu_j) \left(\frac{\tilde{\mu}_j}{\mu_j} - 1 \right).$$

For general point processes the assumption that $\tilde{\mu}_j = \mu_j, \forall j$ does not hold. Hence, $\text{BRL}_{\text{cloglog}}^{(1)}(\boldsymbol{\beta})$ is not, in general, an unbiased estimating function for non-Poisson point processes and thus this procedure should not be recommended if one suspects a significant departure of data from the Poisson assumption.

4.5 Other approaches

As we do not want to model second-order characteristics of the point process \mathbf{Y} , the maximum likelihood method is not considered. As far as we know, two other approaches could have been investigated. The first one is the quasi-likelihood method proposed by [Guan et al. \(2015\)](#). This method designed for spatial point processes (but which can easily be extended to spatio-temporal point processes) has the merit to be robust to the Poisson assumption and

to reduce the variance of the Poisson likelihood estimator. Its computation cost is however important in particular for large datasets. This approach also requires a preliminary estimation of the pair correlation function. Modelling the second-order structure of the spatio-temporal point process \mathbf{Y} is not the topic of this paper.

The second one, based on a variational type estimating equation, is proposed by [Coerjolly and Møller \(2014\)](#). The main interest of this method is that it does not require any optimization procedure and is very cheap from a computational point of view. However, it requires that covariates are observed at a very fine scale in particular in a neighborhood of the observed points. Such an assumption does not hold for the application considered in this paper.

5 Zero-deflated subsampled versions of composite likelihoods methods

5.1 Introduction

Given the size of regressions involved in the lightning strikes data ($J \approx 2.8 \times 10^6$ and the number of impacts can easily reach 300,000 for some seasons if one applies any method from Section 4 to $\mathbf{Y}_{s,m}$), the previous methodologies are computationally expensive and can lead to numerical instabilities. A way to alleviate these problems is to use subsampling strategies, which can also yield more robust estimates (see e.g. [Humbert et al. \(2022\)](#)). In the framework of point processes, subsampling corresponds to a thinning process (see [Daley et al., 2003](#), [Møller and Waagepetersen, 2003](#), [Cronie et al., 2023](#)). In particular, if π stands for the retaining probability (i.e. the thinning probability is equal to $1 - \pi$) assumed to be constant over space and time for the sake of simplicity, then it is well-known that the intensity ρ_{sub} of the thinned point process \mathbf{Y}_{sub} is equal to $\rho_{\text{sub}}(y) = \pi\rho(y)$. Hence if $\rho(y) = \exp(\boldsymbol{\beta}^\top \mathbf{C}(y))$, all methods described in Section 4 can be applied to estimate $\boldsymbol{\beta}$ based on \mathbf{Y}_{sub} up to an offset term. All remarks made on parametric methods summarized in Table 1 apply in the same way if one considers \mathbf{Y}_{sub} instead of \mathbf{Y} . The subsampling procedure could also be repeated and averaged over independent replications as investigated by [Cronie et al. \(2023\)](#) for cross-validation techniques.

However, if we see our dataset as counts of impacts per cell Δ_j (or pres-

ence/absence of an impact per cell), the dataset is highly unbalanced as shown in Figure 3. So if one samples cells independently, there is a high probability that the subsample contains only cells with no data point. For such subsamples, parametric methods presented in Section 4 will obviously fail or lead to numerical instabilities.

This practical problem is well-known in classification, machine learning tasks and regression problems, see e.g. Fernández et al. (2018), Spelmen and Porkodi (2018) and the references therein. In the context of spatial point processes, Baddeley et al. (2015, Section 9.10.3) have also considered this problem and describe a subsampled version of $\text{BRL}_{\text{logit}}$ (pixel logistic regression) where cells with impacts are more likely to appear in the subsample than empty cells. The goal of this section is to propose extensions of methods described in Section 4 for zero-deflated subsamples, that is for subsamples of N_j or I_j such that cells with no data point are less likely to appear (for instance) in the resulting subsample. A summary of these subsampled versions is proposed in Table 2.

5.2 Subsampled version of $\text{BRL}_{\text{logit}}$

We formalize and slightly extend Baddeley et al. (2015, Section 9.10.3). For $\ell = 0, 1$, let $J_\ell = \{j = 1, \dots, J : I_j = \ell\}$ be the set of empty and non-empty cells and let \mathcal{J}_ℓ be a sample from J_ℓ with constant inclusion probability $\pi_\ell = \mathbb{P}(j \in \mathcal{J}_\ell \mid I_j = \ell)$ and let $\mathcal{J} = \mathcal{J}_0 \cup \mathcal{J}_1$. Thus, \mathcal{J} is a discrete random set with on average $\pi_0 J \mathbb{P}(I_j = 0)$ elements corresponding to empty cells and $\pi_1 J \mathbb{P}(I_j = 1)$ indices corresponding to cells containing at least one data point. The variables $I_j \mid j \in \mathcal{J}$ are Bernoulli random variables with mean

$$\mathbb{P}(I_j = 1 \mid j \in \mathcal{J}) = \mathbb{P}(j \in \mathcal{J} \mid I_j = 1) \frac{\mathbb{P}(I_j = 1)}{\mathbb{P}(j \in \mathcal{J})} = \frac{\pi_1}{\mathbb{P}(j \in \mathcal{J})} \mathbb{P}(I_j = 1).$$

Proceeding similarly for $\mathbb{P}(I_j = 0 \mid j \in \mathcal{J})$ yields

$$\frac{\mathbb{P}(I_j = 1 \mid j \in \mathcal{J})}{\mathbb{P}(I_j = 0 \mid j \in \mathcal{J})} = \frac{\pi_1}{\pi_0} \frac{\mathbb{P}(I_j = 1)}{\mathbb{P}(I_j = 0)}. \quad (5.1)$$

Now, let us keep the spirit of the method $\text{BRL}_{\text{logit}}$ and (wrongly) assume that $\text{logit} \mathbb{P}(I_j = 1) = \log(\delta_j) + \boldsymbol{\beta}^\top \mathbf{C}_j$. Then, Equation (5.1) becomes the cornerstone of Section 5: the likelihood of the $I_j \mid j \in \mathcal{J}$ can *still* be used to estimate $\boldsymbol{\beta}$ using a standard logistic regression, but now with offset term $(\log(\delta_j) + \log(\pi_1/\pi_0))_{j \in \mathcal{J}}$.

5.3 Subsampled version of $\text{BRL}_{\text{cloglog}}$

We sample as in the previous section but we now assume that \mathbf{Y} is, as in Section 4.4, a spatio-temporal Poisson point process, which implies that $\mathbb{P}(I_j = 1) = 1 - \exp(-\delta_j \exp(\boldsymbol{\beta}^\top \mathbf{C}_j))$. Using this and (5.1), we can deduce that

$$g^{-1}(\mathbb{P}(I_j = 1 \mid j \in \mathcal{J})) = \log \delta_j + \boldsymbol{\beta}^\top \mathbf{C}_j \quad (5.2)$$

where for any $t \in (0, 1)$

$$g^{-1}(t) = \log \left(\log \left(1 + \frac{\pi_0}{\pi_1} \frac{t}{1-t} \right) \right). \quad (5.3)$$

In other words, given \mathcal{J} and under the Poisson assumption, the likelihood of the Bernoulli random variables $I_j, j \in \mathcal{J}$, corresponds to a logistic regression with link function (5.3) and offset term $(\log \delta_j)_{j \in \mathcal{J}}$. It is worth pointing out that when $\pi_0 = \pi_1$, the link function (5.3) reduces to the complementary log-log function. This is why we suggest the name of `clogclog.sub` for this link function (see Table 2). Similarly to the method $\text{BRL}_{\text{cloglog}}$, its subsampled version is in general not robust to the Poisson assumption.

5.4 Subsampled version of CLRL and WCLRL

Following Sections 4.2 and 5.2, we can propose the following subsampled version of CLRL (the subsampled version of WCLRL is similar). Let \mathcal{N} (resp. \mathcal{J}) be a sample of $\{j = 1, \dots, N + J : \tilde{I}_j = 1\}$ (resp. $\{j = 1, \dots, N + J : \tilde{I}_j = 0\}$) with inclusion probabilities $\pi_1 = \mathbb{P}(j \in \mathcal{N} \mid \tilde{I}_j = 1)$, $\pi_0 = \mathbb{P}(j \in \mathcal{J} \mid \tilde{I}_j = 0)$ and let $\mathcal{K} = \mathcal{N} \cup \mathcal{J}$. Then, it is clear from Section 5.2 that

$$\text{logit } \mathbb{P}(\tilde{I}_j = 1 \mid j \in \mathcal{K}) = \log \left(\frac{\pi_1}{\pi_0} \right) + \log(\delta_j) + \text{logit } \mathbb{P}(\tilde{I}_j = 1).$$

A subsampled version of CLRL is easily deduced: we estimate $\boldsymbol{\beta}$ by a logistic regression of $(\tilde{I}_j)_{j \in \mathcal{K}}$ in terms of $(\mathbf{C}_j)_{j \in \mathcal{K}}$ with logit link and offset term $(\log \pi_1 / \pi_0 + \log \delta_j)_{j \in \mathcal{K}}$. It is worth pointing out that the subsampled version of CLRL (and WCLRL) is fundamentally different from the other methods in the sense that we sample according to the values of the \tilde{I}_j , i.e. we sample from the data points with probability π_1 and independently from the cells with probability π_0 . This subsampling is less natural than the one used for the other methods where we sample from cells with or without impact.

5.5 Subsampled version of PL

As seen in Section 4.1, the method PL corresponds to a Poisson regression of the random variables N_j with the canonical log link. So a subsampled version of this method can be designed by subsampling these random variables. For any $\ell \in \mathbb{N}$, let \mathcal{J}_ℓ be a sample from $J_\ell = \{j = 1, \dots, J : N_j = \ell\}$ with inclusion probability $\tau_\ell = \mathbb{P}(j \in \mathcal{J}_\ell \mid N_j = \ell)$. Given $\mathcal{J} = \cup_{\ell \geq 0} \mathcal{J}_\ell$, the N_j 's are discrete random variables with (conditional) mean

$$\begin{aligned}
 \mathbb{E}(N_j \mid j \in \mathcal{J}) &= \sum_{\ell \geq 0} \ell \mathbb{P}(N_j = \ell \mid j \in \mathcal{J}_\ell) \\
 &= \sum_{\ell=0}^{\infty} \ell \mathbb{P}(j \in \mathcal{J}_\ell \mid N_j = \ell) \frac{\mathbb{P}(N_j = \ell)}{\mathbb{P}(j \in \mathcal{J}_\ell)} \\
 &= \sum_{\ell=0}^{\infty} \ell \mathbb{P}(j \in \mathcal{J}_\ell \mid N_j = \ell) \frac{\mathbb{P}(N_j = \ell)}{\mathbb{P}(j \in \mathcal{J})} \\
 &= \sum_{\ell > 0} \ell \mathbb{P}(N_j = \ell) \frac{\tau_\ell}{\mathbb{P}(j \in \mathcal{J})}. \tag{5.4}
 \end{aligned}$$

Assuming $\tau_0 = \pi_0$ and $\tau_\ell = \pi_1$ for any $\ell \geq 1$, which implies we sample according to the values of I_j as in Sections 5.2-5.3, leads to

$$\mathbb{E}(N_j \mid j \in \mathcal{J}) = \frac{\pi_1}{\mathbb{P}(j \in \mathcal{J})} \mathbb{E}(N_j)$$

and then to

$$\log \mathbb{E}(N_j \mid j \in \mathcal{J}) = \log \left(\frac{\pi_1}{\mathbb{P}(j \in \mathcal{J})} \right) + \log(\delta_j) + \boldsymbol{\beta}^\top \mathbf{C}_j.$$

The problem is that $\mathbb{P}(j \in \mathcal{J})$ is unknown and depends on the parameter vector $\boldsymbol{\beta}$. The likelihood of the $N_j \mid j \in \mathcal{J}$ is therefore no more a Poisson regression. We suggest to estimate $\mathbb{P}(j \in \mathcal{J})$ by $\pi_0 \mathbf{1}(j \in \mathcal{J}_0) + \pi_1 \mathbf{1}(j \in \mathcal{J}_1)$. This estimation leads to

$$\log \mathbb{E}(N_j \mid j \in \mathcal{J}) = o_j + \boldsymbol{\beta}^\top \mathbf{C}_j \quad \text{with } o_j = \log \left(\frac{\pi_1}{\pi_0} \mathbf{1}(j \in \mathcal{J}_0) \right) + \log(\delta_j)$$

which can now be implemented by a Poisson regression with log link and offset term $(o_j)_{j \in \mathcal{J}}$.

5.6 A brief simulation study, remarks and possible extensions

In this section, we propose a short simulation study to underline the interest and efficiency of zero-deflated subsampling. We consider a spatio-temporal Poisson point process with intensity model $\rho(y) = \exp(\boldsymbol{\beta}^\top \mathbf{C}(y))$ where $\boldsymbol{\beta} \in \mathbb{R}^3$ ($n_C = 3$) with $\beta_2 = \beta_3 = 1$. The first covariate is 1 (thus β_1 is an intercept term) and the covariate functions $C_2(y)$ and $C_3(y)$ are discretized such that C_{2j}, C_{3j} are realizations of standard normal random variables for $j = 1, \dots, J$. Thus, the N_j 's are Poisson random variables with mean $\mu_j = \exp(\boldsymbol{\beta}^\top \mathbf{C}_j)$ and we adjust β_1 such that on average on the sample with size J , we have 50%, 90%, 99% or 99.9% of 'empty cells'.

Figure 4 (top) constitutes the baseline. We generate $B = 500$ replications of the previous model and estimate $\boldsymbol{\beta}$ using methods described in Section 4 (no subsampling). We depict the logarithms of the empirical (average absolute) biases, standard deviations and root mean squared errors (RMSE) in term of $\log_2(J)$ for $J = 2^k$ and $k = 9, \dots, 16$. As mentioned in Section 4, the method $\text{BRL}_{\text{logit}}$ is the only method which is based on a biased estimating equation due to the misspecification of the link function. That is why this method leads to non negligible bias. All other methods lead to asymptotically unbiased estimates. We note first that the larger the number of empty cells, the larger the bias and second that the difference with the method $\text{BRL}_{\text{logit}}$ is less obvious when the proportion of empty cells reaches 99.9%. In terms of variance, we first observe the linear trend which means that the variance of estimates seems to decrease as a power J and this for any configuration of 'empty cells'. The best method seems to be the Poisson likelihood one (method PL). Again the variances are higher and the differences between all methods less important when the proportion of empty cells increases.

To evaluate methods described in Section 5 which are based on subsampling strategies, we start with a single simulation with length $J = 2^k$ ($k = 9, \dots, 16$) and estimate $\boldsymbol{\beta}$ based on $B = 500$ subsamples. We set $\pi_1 = 1$ and $\pi_0 = 2^{-2} = 25\%$, then $\pi_0 = 2^{-5} = 3.1\%$ (Figure 10) and $\pi_0 = 2^{-8} = 0.4\%$ (Figure 4 (bottom)). Thus, for these simulations, for each value of J and π_0 , estimates are based on data with length on average $J(p + \pi_0(1 - p))$ where p is the proportion of non-empty cells. The behaviors of empirical biases, standard deviations and thus root mean squared errors are more erratic than in the situation where $\pi_0 = 1$ (no subsampling). Also the values of the different scores are higher than the baseline situation but the comments on the

different methods still apply. In particular the root mean squared errors tend to 0 with J (except for the method $\text{BRL}_{\text{logit}}$) for any value of $1 - p$ (the proportion of empty cells). We have presented here results for a very small value of π_0 (more moderate situations are presented in Figure 10). This intrinsically shows the efficiency of zero-deflated procedures presented in Section 5 including the definition and estimations of offset terms.

This section convinces us of the interest of zero-deflated subsampling type procedures. Overall, when the proportion of empty cells is upto 90% or for moderate sample sizes and higher value of $1 - p$, the Poisson likelihood could be recommended. In extreme situations and for large sample sizes, it is harder to distinguish the different methods. Finally, as expected there are no differences between the methods CLRL and WCLRL. Only the last one is used in Section 6 as it is faster than the other one.

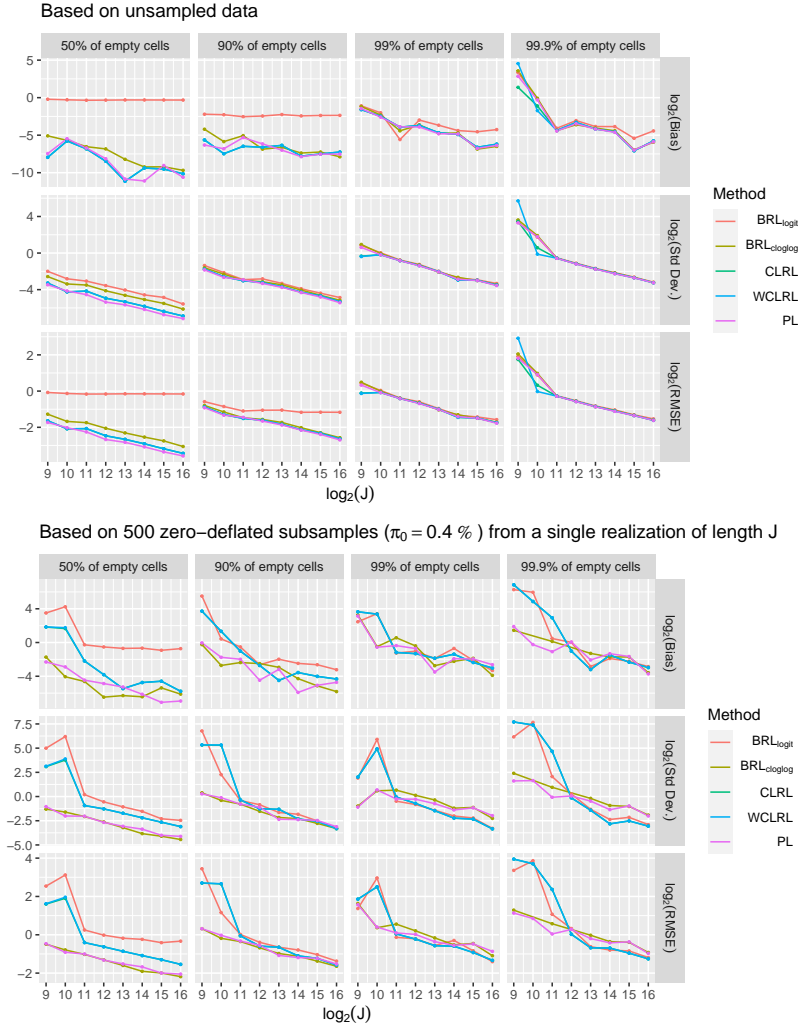


Figure 4: Logarithms (in base 2) of empirical (average absolute) biases, standard deviation and root-mean squared errors for estimates of $\beta = (\beta_1, 1, 1)$ in terms of $\log_2(J)$ based on $B = 500$ replications of the simulation model described in Section 5.6 yielding 50%, 90%, 99% or 99.9%. (Top) Data are unsampled (Bottom) From one simulation of length J , we subsample 500 replications with $\pi_1 = 1$ and $\pi_0 = 0.4\%$.

Method	BRL _{logit} (see Section 4.3)	BRL _{cloglog} (see Section 4.4)	CLRL (see Section 4.2)	WCLRL (see Section 4.2)	PL (see Section 4.1)
GLM family	binomial	binomial	binomial	binomial	poisson
Link	logit	cloglog	logit	logit	log
R formula	I~C	I~C	Icheck~Ccheck	Itilde~Ctilde w=wtilde	N~C
Offset term	(log δ_j) _j	(log δ_j) _j	(log δ_j) _j	(log δ_j) _j	(log δ_j) _j
Size of the response vector	J	J	$N + J$	$\#J_1 + J$	
Unbiased score	×	✓ (if Poisson)	✓	✓	✓
Robustness	×	×	✓	✓	✓
Variance	+/-	+/-	+	+	++

Table 1: Summary of subsampled versions of parametric methods used to estimate an exponential family model with piecewise constant covariates; I (resp. Icheck, Itilde, N) stands for the vector of the I_j (resp. $\tilde{I}_j, \tilde{I}_j, N_j$); C (resp. Ccheck, Ctilde) stands for the matrix of constant covariates; δ_j is the volume of the spatio-temporal cell Δ_j .

Subsampled version of	BRL _{logit}	BRL _{cloglog}	CLRL	WCLRL	PL
	(see Section 5.2)	(see Section 5.3)	(see Section 5.4)	(see Section 5.4)	(see Section 5.5)
GLM family	binomial	binomial	binomial	binomial	poisson
Link	logit	cloglog.sub	logit	logit	log
R formula	I.sub~C.sub	I.sub~C.sub	Icheck.sub~ Ccheck.sub	Itilde.sub~ Ctilde.sub, w=wtilde.sub	N.sub~C.sub
Offset term	$(\log \frac{\pi_1}{\pi_0} + \log \delta_j)_j$	$(\log \delta_j)_j$	$(\log \frac{\pi_1}{\pi_0} + \log \delta_j)_j$	$(\log \frac{\pi_1}{\pi_0} + \log \delta_j)_j$	$(\log \frac{\pi_1}{\pi_0}) \mathbf{1}(j \in \mathcal{J}_0) + \log(\delta_j)_j$
Size of the response vector	$\#\mathcal{J}_0 + \#\mathcal{J}_1$	$\#\mathcal{J}_0 + \#\mathcal{J}_1$	$\#\mathcal{N} + \#\mathcal{J}$	$\#\mathcal{N}_1 + \#\mathcal{J}$	$\#\mathcal{J}_0 + \#\mathcal{J}_1$
Unbiased score given subsample	×	✓ (if Poisson)	✓	✓	\approx ✓
Robustness	×	×	✓	✓	✓

Table 2: Summary of characteristics of parametric methods used to estimate an exponential family model with piecewise constant covariates; I.sub (resp. Itilde.sub, Icheck.sub, N.sub,), stands for a (dependent) subsample of the I_j (resp. \tilde{I}_j, \tilde{N}_j); C.sub (resp. Ccheck.sub, Ctilde.sub) stands for the subsampled matrix of constant covariates; δ_j is the volume of the spatio-temporal cell Δ_j ; (see Section 5 for a definition $\pi_0, \pi_1, \mathcal{J}_0, \mathcal{J}_1, \mathcal{N}, \mathcal{J}$ and \mathcal{N}_1 according to each method).

6 Back to the lightning strikes dataset

We return to the dataset described in Section 2, and model it as a spatio-temporal point process with intensity given by an exponential family with discretized spatial and spatio-temporal covariates (see details in Section 3). In the following, we decompose the spatio-temporal point process \mathbf{Y} into $\mathbf{Y}^{\text{learn}}$ and \mathbf{Y}^{test} as

$$\begin{aligned}\mathbf{Y}^{\text{learn}} &= \{(x, t) \in W \times \mathcal{T}^{\text{learn}}\} \text{ with } \mathcal{T}^{\text{learn}} = \{t \in \mathbb{R}^+ : \text{year} \in \{2013, 2014\}\} \\ \mathbf{Y}^{\text{test}} &= \{(x, t) \in W \times \mathcal{T}^{\text{test}}\} \text{ with } \mathcal{T}^{\text{test}} = \{t \in \mathbb{R}^+ : \text{year} \in \{2015\}\}.\end{aligned}$$

The goal of this section is to estimate the spatio-temporal parametric intensity function using $\mathbf{Y}^{\text{learn}}$ as specified in Section 3, and use the fitted model to predict the spatio-temporal intensity of \mathbf{Y}^{test} , so in 2015. More precisely, for any $s, m \in \mathcal{S} \times \mathcal{M}$ (ie for any season and 6-hour period of the day), we predict the spatio-temporal intensity of $\mathbf{Y}_{s,m}^{\text{test}}$ by learning the one of $\mathbf{Y}_{s,m}^{\text{learn}}$. Let $W \times \mathcal{T}_{s,m}^{\text{test}} = \bigcup_{j \in \mathcal{J}_{s,m}^{\text{test}}} \Delta_j$. Thus $\mathcal{J}_{s,m}^{\text{test}}$ is the set of indices of cells corresponding to s, m in the test dataset and we intend, in particular, to compare the variables $N_{j,s,m}$ and $\hat{N}_{j,s,m}$ for $j \in \mathcal{J}_{s,m}^{\text{test}}$, where

$$\begin{aligned}N_{j,s,m} &= N(\Delta_{j,s,m}) = \sum_{y \in \mathbf{Y}_{s,m}^{\text{test}}} \mathbf{1}(y \in \Delta_{j,s,m}) \\ \hat{N}_{j,s,m} &= \int_{\Delta_{j,s,m}} \exp(\hat{\beta}_{s,m}^\top \mathbf{C}(y)) dy = \exp(\hat{\beta}_{s,m}^\top \mathbf{C}_j)\end{aligned}$$

since the covariates are constant over $\Delta_{j,s,m}$ and since we assume without loss of generality that $\delta_{j,s,m} = |\Delta_{j,s,m}| = 1$. In what precedes, one has of course $\hat{\beta}_{s,m} = \hat{\beta}_{s,m}(\mathbf{Y}^{\text{learn}})$. We can also decide to compare $N_{j,s}$ and $\hat{N}_{j,s}$, ie observations and predictions averaged over the 6-hour periods, or N_j and \hat{N}_j , ie observations and predictions averaged over 2015.

Evaluating the quality of predictions is a tricky task for which several approaches exist. Section 6.1 discusses the criteria we focused on, while Section 6.2 provides empirical results on the lightning strikes dataset.

We end this section by emphasizing that we use the term 'prediction' in the machine learning way. Our goal is to characterize globally the intensity of lightning strikes by spatial and spatio-temporal covariates and we use training and test sets to compare the methodology proposed in this paper. We do not mean 'prediction' as a short-term forecast.

6.1 Criteria used to measure the quality of predictions

To evaluate the ability of different methods to predict the presence/absence of lightning strikes, we consider for any s, m and any $j \in \mathcal{J}_{s,m}^{\text{test}}$, $I_{j,s,m} = \mathbf{1}(N_{j,s,m} > 0)$ and $\hat{I}_{j,s,m} = \mathbf{1}(\hat{N}_{j,s,m} > 0)$ and compute standard metrics such as areas under ROC Curve and Precision-Recall Curve, denoted respectively by $\text{AUC}_{s,m}$ and $\text{PPV}_{s,m}$; see e.g. [Murphy \(2012, Section 5.7.2\)](#).

To evaluate the ability of predicting higher values of number of impacts, we suggest the metric $\text{WW}_{s,m}$ (for weighted Wasserstein). This metric, justified later, is defined as

$$\text{WW}_{s,m} = \omega \int_0^1 |L_{p,s,m} - R_{p,s,m}| dp \quad (6.1)$$

with $\omega = f(N_{s,m}/\hat{N}_{s,m})$ where $N_{s,m}$ and $\hat{N}_{s,m}$ denote the total number on $\mathcal{J}_{s,m}^{\text{test}}$ of observed and predicted lightning strikes in 2015's season s and 6-hour period of the day m , and $f(x) = x\mathbf{1}(x \geq 1) + x^{-1}\mathbf{1}(x < 1)$. Using the function $\tilde{j} : \mathbb{R}^d \times \mathbb{R}^+ \rightarrow \{1, \dots, J\}$ given by $\tilde{j}(y) = \sum_{s,m} \sum_{j \in \mathcal{J}_{s,m}^{\text{test}}} j \mathbf{1}(y \in \Delta_{j,s,m})$ which identifies the spatio-temporal cell an impact at location y belongs to, we define for $\bar{\nu}_{s,m}$, an upper-bound of $\max_{j \in \mathcal{J}_{s,m}^{\text{test}}} \hat{N}_{j,s,m}$

$$\begin{aligned} L_{p,s,m} &= \frac{1}{N_{s,m}} \sum_{j \in \mathcal{J}_{s,m}^{\text{test}}} N_{j,s,m} \mathbf{1}(\hat{N}_{j,s,m} \leq p\bar{\nu}_{s,m}) \\ &= \frac{1}{N_{s,m}} \sum_{y \in \mathbf{Y}_{s,m}^{\text{test}}} \mathbf{1}(\hat{N}_{\tilde{j}(y),s,m} \leq p\bar{\nu}_{s,m}) \\ R_{p,s,m} &= \frac{1}{\hat{N}_{s,m}} \sum_{j \in \mathcal{J}_{s,m}^{\text{test}}} \hat{N}_{j,s,m} \mathbf{1}(\hat{N}_{j,s,m} \leq p\bar{\nu}_{s,m}) \\ &= \frac{1}{\hat{N}_{s,m}} \int_{W \times \mathcal{T}_{s,m}^{\text{test}}} \mathbf{1}(\hat{N}_{\tilde{j},s,m} \leq p\bar{\nu}_{s,m}) \rho(y; \hat{\beta}_{s,m}) dy. \end{aligned} \quad (6.2)$$

As a first comment, let $\delta_p = N_{s,m}L_{p,s,m} - \hat{N}_{s,m}R_{p,s,m}$. It is worth pointing out that

$$\delta_p = \sum_{y \in \mathbf{Y}_{s,m}^{\text{test}}} h(y; \mathbf{Y}^{\text{test}}) - \int_{W \times \mathcal{T}_{s,m}^{\text{test}}} h(y; \mathbf{Y}^{\text{test}}) \rho(y; \hat{\beta}_{s,m}(\mathbf{Y}^{\text{learn}})) dy$$

where $h(y; \mathbf{Y}^{\text{test}}) = \mathbf{1}(\hat{N}_{\tilde{j}(y),s,m} \leq p\bar{\nu}_{s,m})$. Such a random variable is a bivariate residual functional for point processes. More precisely if $\rho(\cdot; \hat{\beta}) = \rho(\cdot; \beta)$

(ie the true intensity), δ_p would be a bivariate innovation functional as defined by [Cronie et al. \(2023\)](#), proved to be centered under the true model. The criterion (6.1) can therefore be viewed as a self-normalized weighted combination of generalized residuals.

But we can actually say more about the intuition of this criterion. By Campbell theorem, we see that $R_{p,s,m}$ is close to

$$\check{R}_{p,s,m} = \frac{1}{\hat{N}_{s,m}} \sum_{y \in \check{\mathbf{Y}}_{s,m}^{\text{test}}} \mathbf{1}(\hat{N}_{\check{j}(y),s,m} \leq p \bar{\nu}_{s,m}) \quad (6.3)$$

where $\check{\mathbf{Y}}_{s,m}^{\text{test}}$ is a spatio-temporal point process with $\hat{N}_{p,s,m}$ events and intensity function $\rho(\cdot; \hat{\beta}_{s,m}(\mathbf{Y}^{\text{learn}}))$. Now given $N_{s,m}$ (resp. given $\hat{N}_{s,m}$), we observe that (6.2) (resp. (6.3)) is the empirical cumulative distribution function evaluated at p based on the sample $\{\hat{N}_{\check{j}(y),s,m}/\bar{\nu}_{s,m}, y \in \mathbf{Y}_{s,m}^{\text{test}}\}$ (resp. $\{\hat{N}_{\check{j}(y),s,m}/\bar{\nu}_{s,m}, y \in \check{\mathbf{Y}}_{s,m}^{\text{test}}\}$). Let $\hat{F}_{s,m}^{\text{test}}$ and $\check{F}_{s,m}^{\text{test}}$ be these two empirical cdf, we deduce that

$$\int_0^1 |L_{p,s,m} - R_{p,s,m}| dp \approx \int_0^1 |\hat{F}_{s,m}^{\text{test}}(p) - \check{F}_{s,m}^{\text{test}}(p)| dp$$

which, given $N_{s,m}$ and $\hat{N}_{s,m}$ is the standard Wasserstein distance between the two empirical distributions. Finally, the coefficient $\omega = f(N_{s,m}/\hat{N}_{s,m})$ in (6.1) is aimed to take into account the total mass of $\mathbf{Y}_{s,m}^{\text{test}}$ and $\check{\mathbf{Y}}_{s,m}^{\text{test}}$. All these intuitions lead us to propose the notation $\text{WW}_{s,m}$ for weighted Wasserstein for the criterion defined by (6.1). For several years, the Wasserstein distance has been recognized as a robust optimal transport-based metric with numerous applications in classification and learning algorithms. See, for example, [Frogner et al. \(2015\)](#) [Frogner et al. \(2015\)](#) or [Arjovsky et al. \(2017\)](#) [Arjovsky et al. \(2017\)](#) for more insights.

To summarize, on the one hand the areas under ROC Curve and Precision-Recall Curve, $\text{AUC}_{s,m}$ and $\text{PPV}_{s,m}$ are used to evaluate the ability of different methods to predict the presence/absence of lightning strikes, and on the other hand, the $\text{WW}_{s,m}$ index is used to evaluate the ability of predicting higher values of number of impacts.

6.2 Results

As explained in the beginning of Section 6, we now estimate the spatio-temporal parametric intensity model from years 2013-2014 via $\mathbf{Y}^{\text{learn}}$, and

then predict the spatio-temporal intensity of \mathbf{Y}^{test} (for 2015). This is done by fixing a season s and a 6-hour period of the day m . For the sake of brevity, we show the results for two pairs of (s, m) only, namely (summer, 12:00-18:00) and (winter, 00:00-06:00), which respectively correspond to cases with the most and the least proportion of non empty cells among all cells. The competitors defined in Section 4 are combined with the subsampling strategies detailed in Section 5. All the combinations are tested with the whole sample ($\pi_1 = 100\%, \pi_0 = 100\%$), and for different subsampling rates $\pi_1 = 1, \pi_0 \in \{10\%, 1\%, 0.1\%\}$, each time aggregating either on 1, 3 or 10 bags (ie subsamples).

Tables 3-4 summarize the results obtained. We report the values of prediction criteria PPV, AUC and WW detailed in Section 6.1, as well as the calculation time per bag. From these two tables, a few comments can be formulated: first of all, whatever the criterion, it seems useless to choose more than one bag; second, note that AUC does not appear so informative, since all methods and subsampling rates achieve the same kind of performance. This is also the case for the PPV criterion in winter 00:00-06:00 (Table 4). Note also that PPV is quite small, in particular for winter (see Table 4). This confirms that the prediction of lightning storms is a very difficult task. However, this should be compared to the PPV achieved by the random classifier, which is the proportion of non-empty cells for these pairs of (season, moment) (0.0154 for (summer, 12:00-18:00) and 0.0004 for (winter, 00:00-06:00)). Two methods (PL and $\text{BRL}_{\text{cloglog}}$) stand out for their stability concerning the subsampling rate, as demonstrated by the behaviors of PPV and WW in Table 3. Additionally, in the absence of subsampling ($\pi_0 = 1$), PPV lacks discrimination. However, in terms of WW, PL and WCLRL slightly outperform the other methods. Notably, subsampling significantly impacts computational time reduction (by at least a factor of 40) without compromising results.

In a nutshell, the Poisson method PL with $\pi_0 = 0.1\%, N_{\text{bag}} = 1$ appears as the best trade off between prediction quality and computation cost. This method and parameters are used in Figure 5 to illustrate the quality of predictions over time and space in 2015. We report observed and predicted cumulative number of impacts and maps of intensities over a season of 2015 and over 2015. The cumulative number of impacts are quite well predicted even if some departures can be observed in the end of spring and summer. From a spatial point of view, it is interesting that the fitted model is able to reproduce most of activity phenomena close to the Channel (North West),

the French Alps (South East) and Pyrenees. There are however a few areas which are not well captured: the summer activity in the area of Bordeaux (South West), the Winter and summer activity north to Corsica island.

7 Conclusion

In this paper, we have proposed zero-deflated subsampled versions of standard composite likelihoods methods to estimate exponential family intensity models for spatio-temporal point processes. We have in particular focused on the situation where the covariates are piecewise constant taking the opportunity to review standard composite likelihood methods in this situation. We have theoretically justified these subsampled extensions. The methods perform well on simulations, at least for spatio-temporal Poisson point processes. When they are applied to the lightning strikes dataset, we have concluded that the subsampled Poisson likelihood provides the best compromise in terms of stability, prediction errors and computational time.

We believe this work leads to several perspectives. It would be first interesting to study more deeply these procedures. For instance, can we still be ensured that the estimator derived by these procedures is consistent and asymptotically normal? Extending these procedures for non piecewise constant covariates is a second perspective. Finally, investigating inhomogeneous subsampling, i.e. for instance letting π_0 and π_1 depend on the cell Δ_j is also a direction for further research.

Regarding the dataset application, modelling second-order characteristics of the process or even better modelling the distribution of \mathbf{Y} for instance as a mixture of a temporal Hawkes process and Neymann-Scott or log-Gaussian spatial Cox process for the spatial part could be an interesting first attempt. Such a modelling could be pertinent if one in addition intend to provide a short-term forecast of electrical activity.

8 Acknowledgements

The authors would like to take the opportunity to thank Météorage and Météo-France, and in particular Maxime Taillardat and Olivier Mestre for providing us with the data and for scientific exchanges about these data. The research of JF Coeurjolly is funded by Labex PERSYVAL-lab ANR-11-

	Method			
	BRL _{logit}	BRL _{cloglog}	WCLRL	PL
AUC_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	89.7	89.6	89.9	89.8
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	89.7	89.6	89.9	89.8
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	89.7	89.6-89.5	89.9	89.8
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	89.1-89.3	89.8	89.1	89.7-89.8
PPV_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	12.9	13	12.2	12.5
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	12.3	13.2	11.3	12.5
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	11.1-11	13.4-13.5	9.8-9.9	12.8-12.5
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	9.1-9.4	12.9-12.8	8.7	11.9-12.6
WW_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	8.1	6.4	3.5	3.2
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	9.4-9.3	5.8-5.9	4.9	3.2
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	11.9-12.1	12-12.2	9-8.6	3.2-3.1
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	13.2-12.8	5.2-5	19.7-18.5	4.5-3.1
Time (in sec.) per bag				
$\pi_0 = 100\%$	58	98	47	46
$\pi_0 = 10\%$	6	23	5	6
$\pi_0 = 1\%$	1	6	2	2
$\pi_0 = 0.1\%$	1	3	1	1

Table 3: Evaluation of prediction metrics (AUC_{s,m}, PPV_{s,m} and WW_{s,m}) and computational time for subsampled composite likelihood methods for different values of π_0 and N_{bags} ($\pi_1 = 1$). Predictions are done for the season s =summer and the time period m =12:00-18:00. Results for AUC_{s,m}, PPV_{s,m} and WW_{s,m} have been multiplied by 100. Most of criteria are similar for different values of N_{bags} . When these numbers are different, we provide the range of values (the first, resp. the second, value corresponding to $N_{\text{bags}} = 1$, resp. $N_{\text{bags}} = 10$).

	Method			
	BRL _{logit}	BRL _{cloglog}	WCLRL	PL
AUC_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	93.1	93.2	93.3	93.3
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	93.1	93.2	93.3	93.3
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	93	93.2-93.1	93.2	93.2-93.3
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	92-92.4	92.7-93.4	92.9	92.9-93.4
PPV_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	2	2.1	2.3	2.3
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	2	2.1	2.3	2.3
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	2	2.1	2.2	2.2-2.3
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	1.3-1.6	2-2.1	1.8-1.9	2-2.3
PPV_{s,m}				
$\pi_0 = 100\%$, $N_{\text{bags}} = 1$	58.8	58.8	42.4	42.4
$\pi_0 = 10\%$, $N_{\text{bags}} = 1, 3, 10$	57.5-58.9	57.4-58.4	43.5-42.7	42.7-41.4
$\pi_0 = 1\%$, $N_{\text{bags}} = 1, 3, 10$	61.8-65.6	51.3-53.3	48-45.2	39.7-42.6
$\pi_0 = 0.1\%$, $N_{\text{bags}} = 1, 3, 10$	40.3-53.9	38.8-47.2	40.2-45	25.8-41.9
Time (in sec.) per bag				
$\pi_0 = 100\%$	86	127	85	76
$\pi_0 = 10\%$	7	11	7	6
$\pi_0 = 1\%$	0.9	2	0.6	0.6
$\pi_0 = 0.1\%$	0.3	0.5	0.1	0.1

Table 4: Evaluation of prediction metrics (AUC_{s,m}, PPV_{s,m} and WW_{s,m}) and computational time for subsampled composite likelihood methods for different values of π_0 and N_{bags} ($\pi_1 = 1$). Predictions are done for the season s =winter and the time period m =0:00-6:00. Results for AUC_{s,m}, PPV_{s,m} and WW_{s,m} have been multiplied by 100. Most of criteria are similar for different values of N_{bags} . When these numbers are different, we provide the range of values (the first, resp. the second, value corresponding to $N_{\text{bags}} = 1$, resp. $N_{\text{bags}} = 10$).

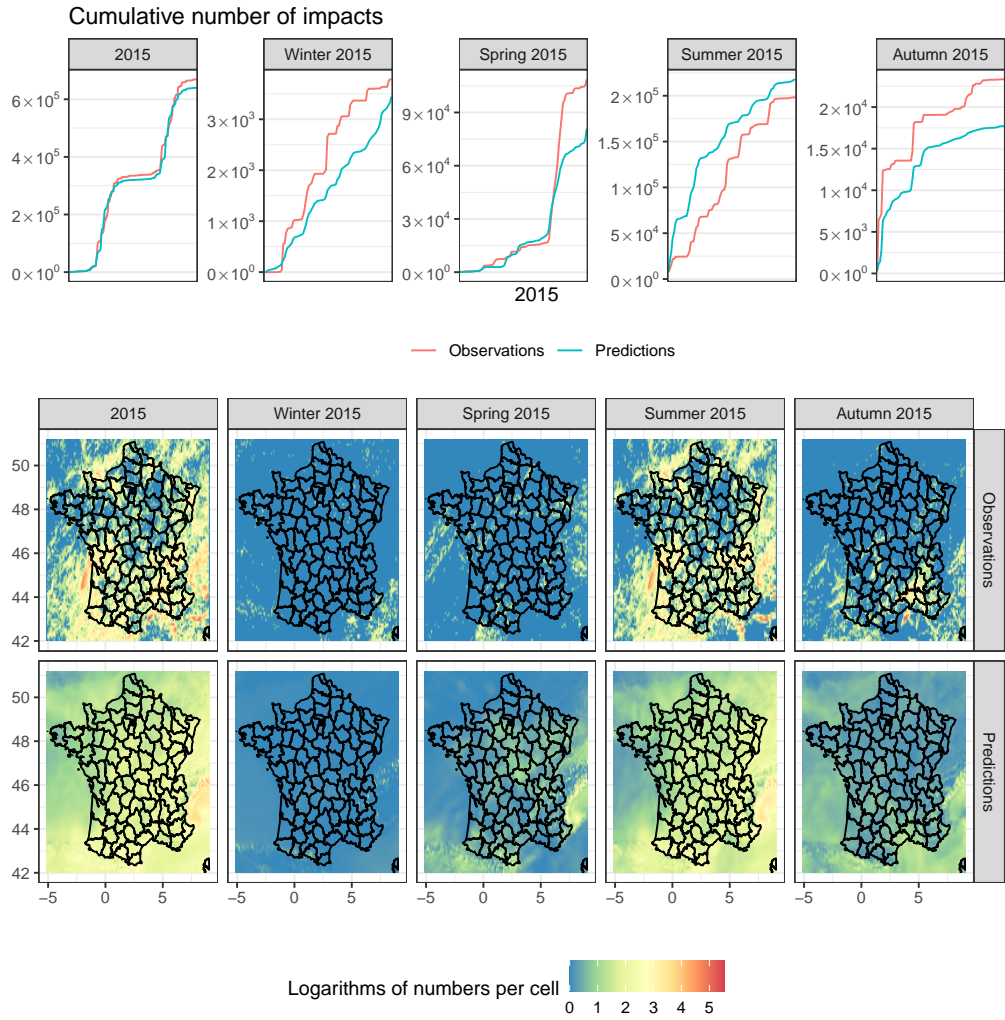


Figure 5: Observed and predicted cumulative number of impacts (top) and maps of logarithms of numbers per cell (bottom). Results are presented for each season in 2015 (thus aggregated over periods of days) and globally for 2015.

LABX-0025.

References

- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- P. M. Atkinson and R. Massari. Generalised linear modelling of susceptibility to landsliding in the central apennines, italy. *Computers & Geosciences*, 24(4):373–385, 1998.
- A. Baddeley and R. Turner. Practical maximum pseudolikelihood for spatial point patterns: (with discussion). *Australian & New Zealand Journal of Statistics*, 42(3):283–322, 2000.
- A. Baddeley, M. Berman, N. I. Fisher, A. Hardegen, R. K. Milne, D. Schumacher, R. Shah, and R. Turner. Spatial logistic regression and change-of-support in poisson point processes. *Electronic Journal of Statistics*, 4: 1151–1201, 2010.
- A. Baddeley, J.-F. Coeurjolly, E. Rubak, and R. Waagepetersen. Logistic regression for spatial gibbs point processes. *Biometrika*, 101(2):377–392, 2014.
- A. Baddeley, E. Rubak, and R. Turner. *Spatial point patterns: methodology and applications with R*. CRC press, 2015.
- M. Berman and T. R. Turner. Approximating point process likelihoods with glim. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 41(1):31–38, 1992.
- P. Bénichou. Cartography of statistical pluviometric fields with an automatic allowance for regional topography. In *Global Precipitations and Climate Change*, pages 187–199. Springer, Berlin and Heidelberg, 1994.
- J.-F. Coeurjolly and F. Lavancier. Understanding spatial point patterns through intensity and conditional intensities. In *Stochastic Geometry*, pages 45–85. Springer, 2019.

- J.-F. Coeurjolly and J. Møller. Variational approach for spatial point process intensity estimation. *Bernoulli*, 20(3):1097–1125, 2014.
- N. Cressie and C. K. Wikle. *Statistics for spatio-temporal data*. John Wiley & Sons, 2015.
- O. Cronie, M. Moradi, and C. A. Biscio. Statistical learning and cross-validation for point processes. *to appear in Biometrika*, 2023.
- E. B. Curran, R. L. Holle, and R. E. López. Lightning casualties and damages in the United States from 1959 to 1994. *Journal of Climate*, 13(19):3448–3464, 2000.
- D. J. Daley, D. Vere-Jones, et al. *An introduction to the theory of point processes: volume I: elementary theory and methods*. Springer, 2003.
- P. Diggle. A kernel method for smoothing point process data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 34(2):138–147, 1985.
- P. J. Diggle. Spatio-temporal point processes: methods and applications. *Monographs on Statistics and Applied Probability*, 107:1, 2006.
- P. Elliot, J. C. Wakefield, N. G. Best, D. J. Briggs, et al. *Spatial epidemiology: methods and applications*. Oxford University Press, 2000.
- A. Fernández, S. García, M. Galar, R. C. Prati, B. Krawczyk, and F. Herrera. *Learning from imbalanced data sets*, volume 10. Springer, 2018.
- C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio. Learning with a wasserstein loss. *Advances in neural information processing systems*, 28, 2015.
- J. A. González, F. J. Rodríguez-Cortés, O. Cronie, and J. Mateu. Spatio-temporal point process statistics: a review. *Spatial Statistics*, 18:505–544, 2016.
- P. V. Gorsevski, P. E. Gessler, R. B. Foltz, and W. J. Elliot. Spatial prediction of landslide hazard using logistic regression and roc analysis. *Transactions in GIS*, 10(3):395–415, 2006.

- Y. Guan, A. Jalilian, and R. Waagepetersen. Quasi-likelihood for spatial point processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(3):677–697, 2015.
- I. Hernandez-Magallanes and M. G. Genton. A point process analysis of cloud-to-ground lightning strikes in urban and rural Oklahoma areas. *Environmetrics*, 30(1):e2535, 2019.
- P. Humbert, B. Le Bars, and L. Minvielle. Robust kernel density estimation with median-of-means principle. In *International Conference on Machine Learning*, pages 9444–9465. PMLR, 2022.
- J. Illian, A. Penttinen, H. Stoyan, and D. Stoyan. *Statistical analysis and modelling of spatial point patterns*. John Wiley & Sons, 2008.
- T. Mona, Á. Horváth, and F. Ács. A thunderstorm cell-lightning activity analysis: The new concept of air mass catchment. *Atmospheric Research*, 169:340–344, 2016.
- K. P. Murphy. *Machine Learning: a Probabilistic Perspective*. MIT Press, 2012.
- M. Myllymäki, T. Mrkvička, P. Grabarnik, H. Seijo, and U. Hahn. Global envelope tests for spatial processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 381–404, 2017.
- J. Møller and R. P. Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC Press, 2003.
- H. Nampak, P. Love, P. Fox-Hughes, C. Watson, J. Aryal, and R. M. Harris. Characterizing spatial and temporal variability of lightning activity associated with wildfire over tasmania, australia. *Fire*, 4(1):10, 2021.
- M. Raeisi, F. Bonneu, and E. Gabriel. A spatio-temporal multi-scale model for Geyer saturation point process: application to forest fire occurrences. *Spatial Statistics*, 41:100492, 2021.
- S. L. Rathbun and N. Cressie. Asymptotic properties of estimators for the parameters of spatial inhomogeneous poisson point processes. *Advances in Applied Probability*, 26(1):122–154, 1994.

- F. P. Schoenberg. Consistent parametric estimation of the intensity of a spatial–temporal point process. *Journal of Statistical Planning and Inference*, 128(1):79–93, 2005.
- W. Schulz, K. Cummins, G. Diendorfer, and M. Dorninger. Cloud-to-ground lightning in austria: A 10-year study using data from a lightning location system. *Journal of Geophysical Research: Atmospheres*, 110, 2005.
- W. Schulz, G. Diendorfer, S. Pedeboy, and D. Poelman. The european lightning location system euclid – part 1: Performance validation. *Natural Hazards and Earth System Sciences*, 3:5325–5355, 2016.
- T. Simon, N. Umlauf, A. Zeileis, G. J. Mayr, W. Schulz, and G. Diendorfer. Spatio-temporal modelling of lightning climatologies for complex terrain. *Natural Hazards and Earth System Sciences*, 17:305–314, 2017.
- V. S. Spelman and R. Porkodi. A review on handling imbalanced data. In *2018 International Conference on Current Trends towards Converging Technologies (ICCTCT)*, pages 1–11. IEEE, 2018.
- M. Taillardat and O. Mestre. From research to applications – examples of operational ensemble post-processing in france using machine learning. *Nonlinear Processes in Geophysics*, 27(2):329–347, 2020.
- M. Van Lieshout. A J-function for inhomogeneous point processes. *Statistica Neerlandica*, 65(2):183–201, 2011.
- R. Waagepetersen. Estimating functions for inhomogeneous spatial point processes with incomplete covariate data. *Biometrika*, 95(2):351–363, 2008.
- R. Waagepetersen and Y. Guan. Two-step estimation for inhomogeneous spatial point processes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):685–702, 2009.
- J. Wakefield. A critique of statistical aspects of ecological studies in spatial epidemiology. *Environmental and Ecological Statistics*, 11:31–54, 2004.
- L. A. Waller and C. A. Gotway. *Applied spatial statistics for public health data*. John Wiley & Sons, 2004.

A Additional figures

A.1 Exploratory data analysis

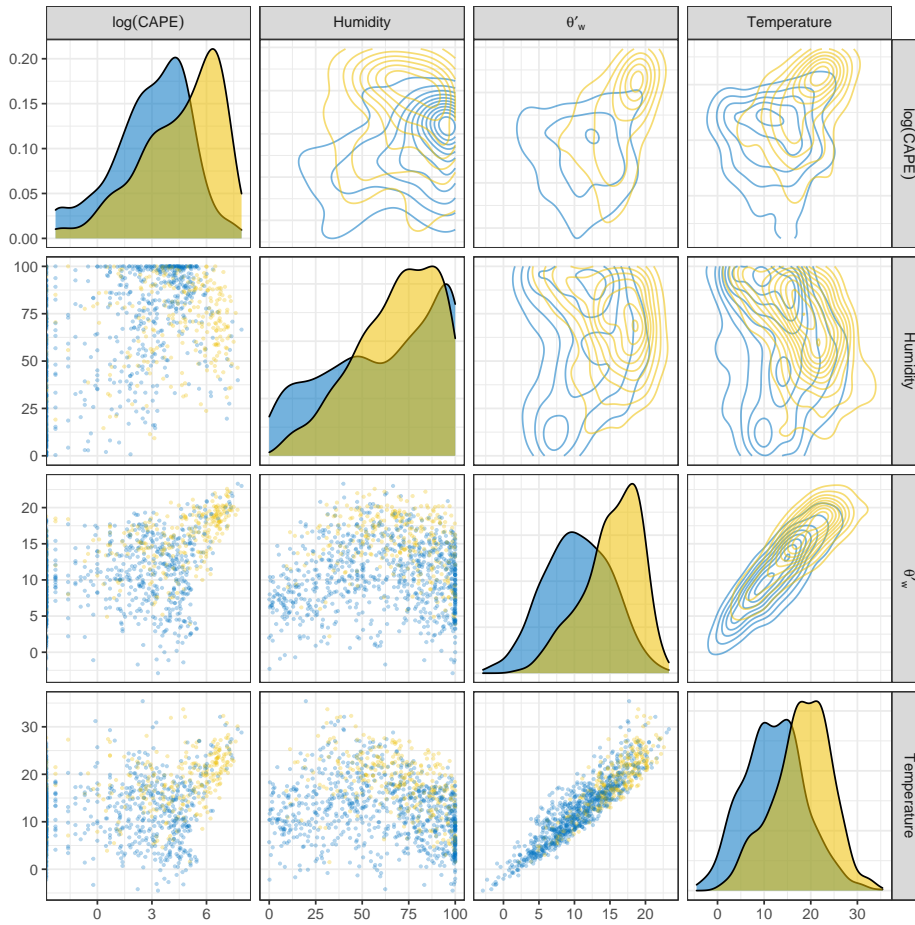


Figure 9: Pairplots (lower triangular part), marginal distribution (diagonal), and bivariate distributions (upper triangular part) of $\log(\text{CAPE})$ (0 values for CAPE are omitted), Humidity, θ'_w , Temperature at 20m around location x_0 at 12 am. Yellow (resp. blue) color corresponds to times for which at least one (resp. no) lightning strike is observed in the corresponding cell during the period 12:00-18:00. The time period spans from 2013 to 2015.

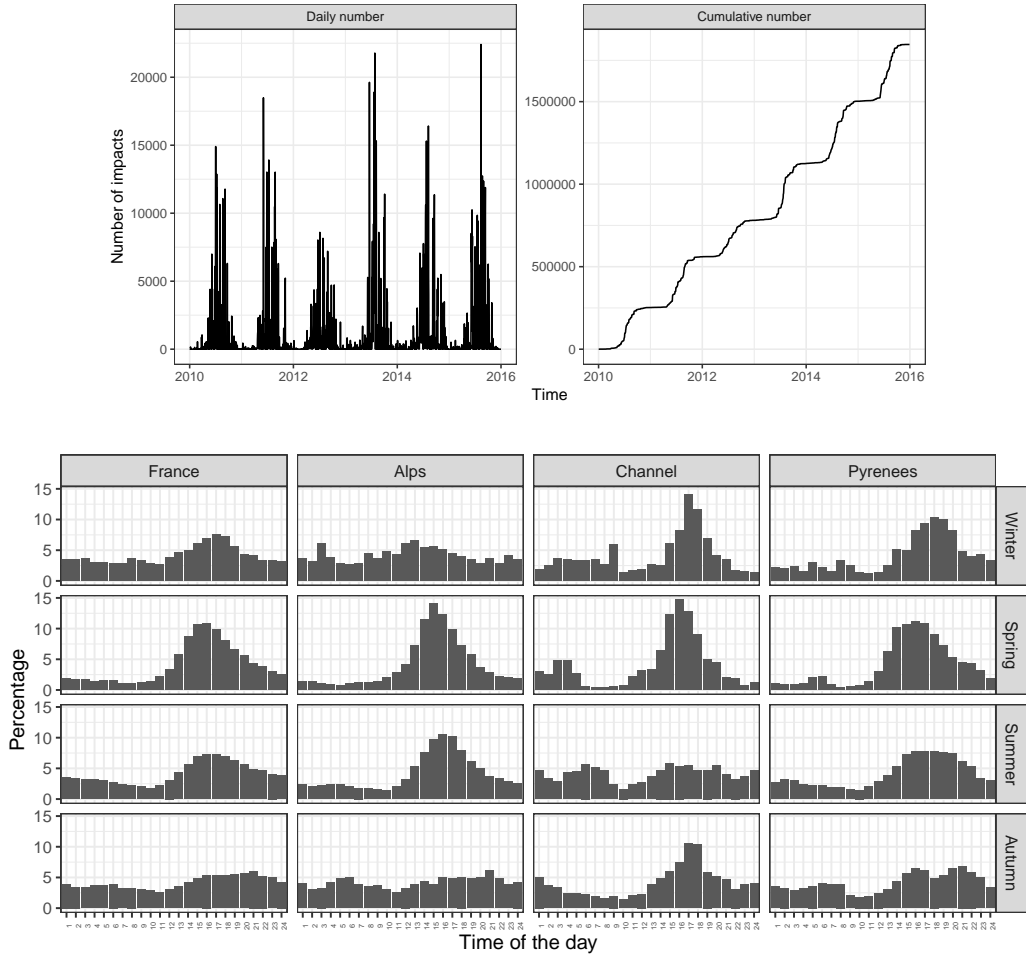


Figure 6: (Top) Daily counting process (left) and cumulative counting process (right) of impacts numbers over France from 2010 to 2015. (Bottom) Distribution of impacts in terms of the time of day for each season over France and Alps (first two columns on the left) and over the Channel and Pyrénées (last two columns on the right).




Figure 7: Logarithms of non-parametric kernel spatial intensity estimates of the intensity of impacts aggregated by month.

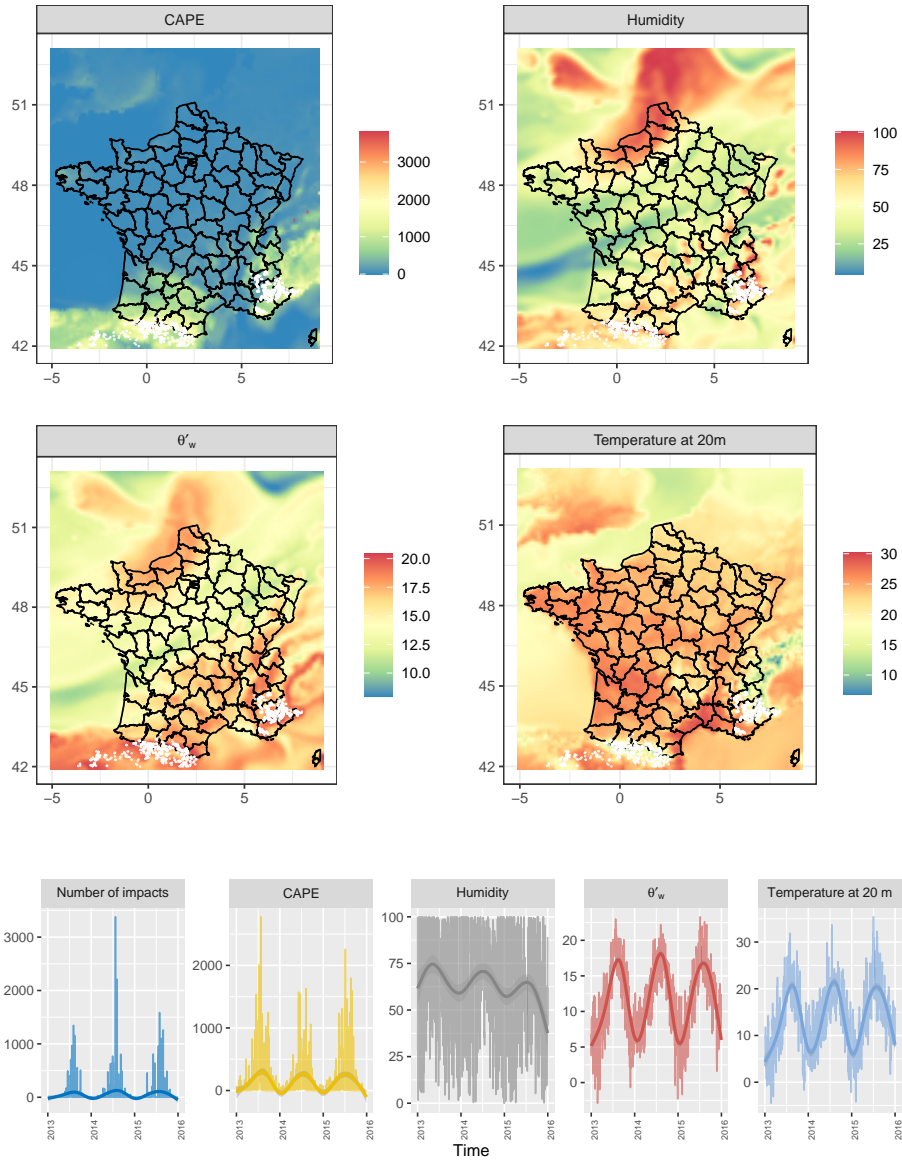


Figure 8: (Top) Maps of CAPE (top left); humidity (top right); θ'_w at 850 hPa (Wet-bulb potential temperature at 850 hPa) (bottom left); temperature at 20m (bottom right), overlaid by the locations of impacts on July 14, 2013 over the period 12:00-18:00. (Bottom, from left to right) Numbers of impacts over the period 12:00-18:00 for each day from 2013 to 2015, around location x_0 (defined in Figure 1); and covariates at 12am at location x_0 , in terms of time (CAPE; Humidity; θ'_w ; Temperature at 20m).

A.2 Simulation study

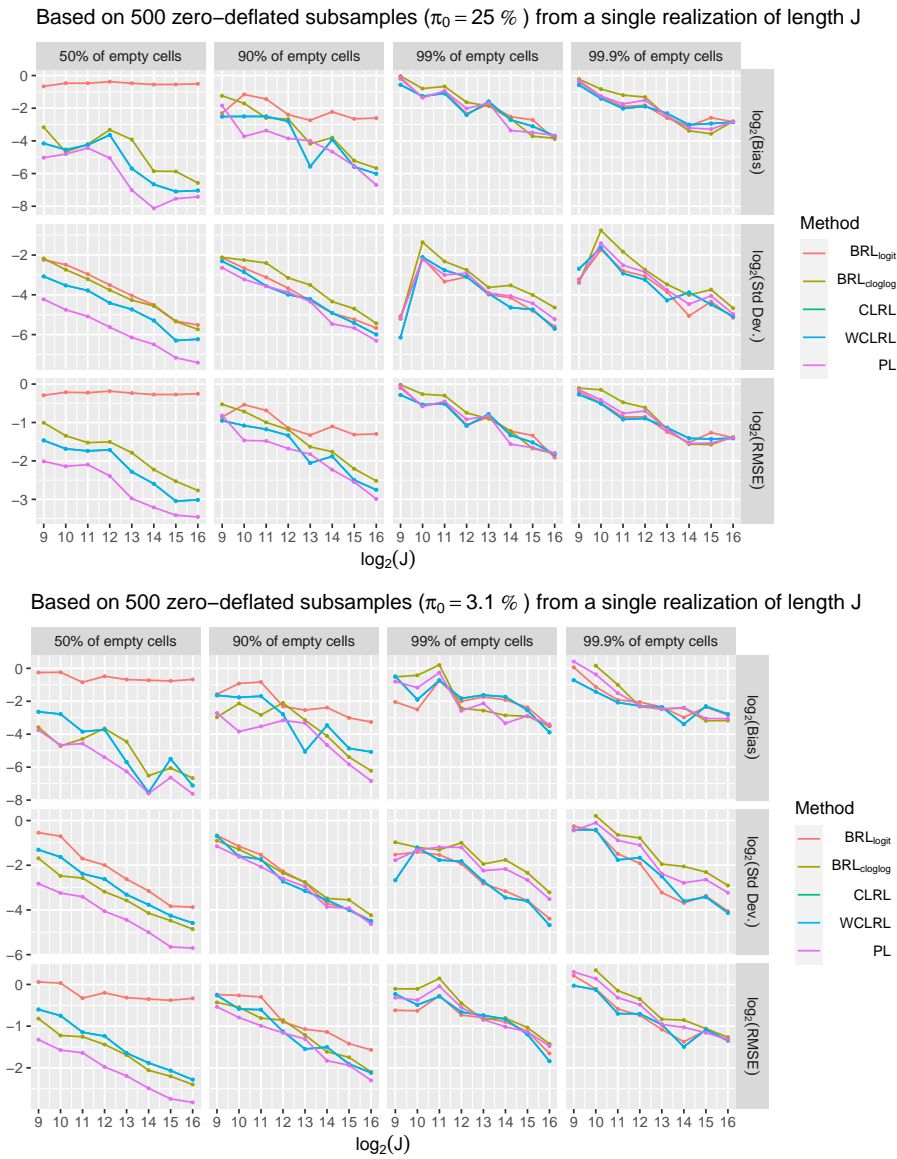


Figure 10: Description similar to Figure 4 with $\pi_0 = 25\%$ (top) and $\pi_0 = 3.1\%$ (bottom).