Journal of the Royal Statistical Society
SERIES C Applied Statistics

Appl. Statist. (2013)
62, *Part* 2, *pp.* 181–199

Determining high-risk zones for unexploded World War II bombs by using point process methodology

Monia Mahling, Michael Höhle and Helmut Küchenhoff

Ludwig-Maximilians-Universität München, Germany

[Received May 2011. Final revision April 2012]

Summary. To prevent accidents caused by unexploded bombs from the Second World War, high-risk zones need to be determined. We introduce two statistical methods to determine such zones by considering patterns of exploded and unexploded bombs as realizations of inhomogeneous spatial Poisson point processes. The first method is based on the intensity of the point process; the second method on its nearest neighbour distance. Although the performance of the two methods is similar, the intensity-based method has the advantage of a more direct specification of risk and should hence be preferred. Risk that is associated with a zone is defined as the probability that not all unexploded bombs are in the zone. We propose and develop a procedure to calculate this probability. This procedure is then also used to investigate consequences of additional spatial clustering, in which case the intensity-based method is shown to be conservative. Our work is motivated and illustrated by World War II explosive ordnance disposal by German authorities on the basis of aerial pictures taken by the allies.

Keywords: Intensity; Kernel method; Risk assessment; Spatial point processes; Unexploded ordnance

1. Introduction

Even 65 years after the end of the Second World War, unexploded bombs still represent a serious problem in Germany. Often accidentally found during construction work, their clearance usually requires the evacuation of houses and the closing of roads and railway lines. Even worse, unintended detonations have resulted in severe accidents in several cases. Therefore, it is desirable to search high-risk areas for unexploded bombs before any construction work starts. As this causes high costs, the search must be restricted to carefully selected areas. So an important task to be addressed by statistical methods consists in finding as small regions as possible containing as many unexploded bombs as possible. A further goal is to perform a risk assessment of the investigated area, e.g. by estimating the probability that there are unexploded bombs outside the high-risk zone.

The data that are available to achieve this goal consist of aerial pictures taken by the allies showing the locations of exploded bombs. The georeferenced locations of the exploded bombs—derived from these aerial pictures—have been provided by *Oberfinanzdirektion Niedersachsen*, which supports the removal of unexploded ordnance in federal properties in Germany. As all information is derived from the aerial pictures, no data are available about the locations of unexploded bombs that have been found for the specific areas of interest.

Two examples of the data are presented in Fig. 1. Example A comprises 443 observations of exploded bombs in an area of approximately 400 ha. Example B consists of 104 observations in

Address for correspondence: Monia Mahling, Statistical Consulting Unit, Department of Statistics, Ludwig-Maximilians-Universität München, Ludwigstraße 33, 80539 München, Germany.
E-mail: monia.mahling@stat.uni-muenchen.de

 0035–9254/13/62181

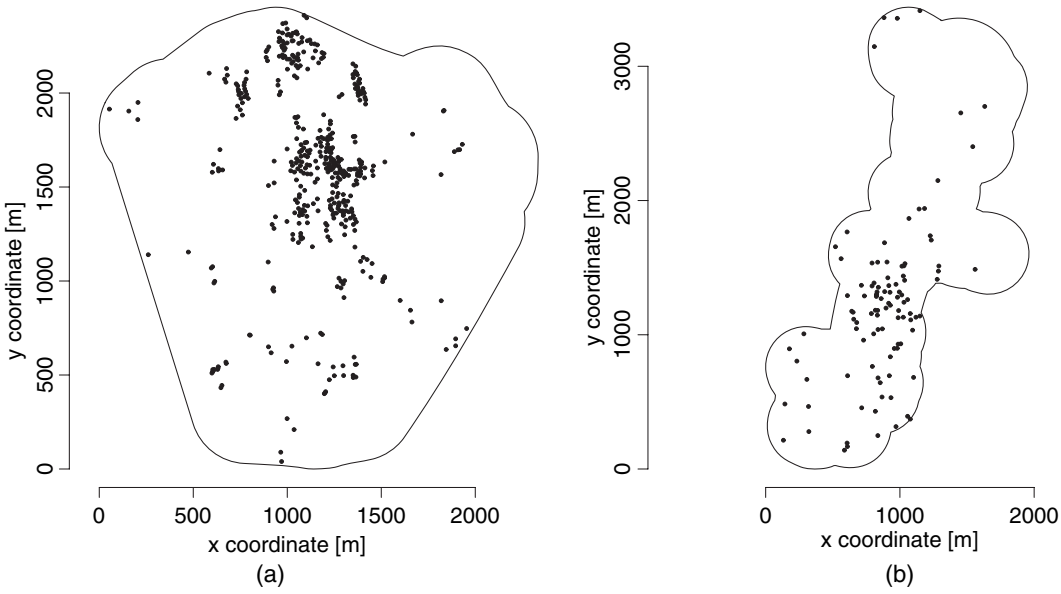

Fig. 1. Properties to be cleared for (a) example A and (b) example B: ———, observation window W (the areas of interest for which data are available); •, locations of exploded bombs

an area of approximately 350 ha. The properties to be cleared usually had military importance during the war (i.e. barracks, airfields and military training areas) or were important from a strategic or economic point of view (i.e. rivers, floodgates, roads and industrial plants). Since detailed information on specific type of premises and location could facilitate the identification of the property, no further information other than relative co-ordinates on the specific property was provided.

Up to now, high-risk zones have been defined in practice by taking the union of discs around all exploded bombs. The radius of these discs was fixed and determined by an expert in advance. This means that only very little information from the actual data is used, namely the co-ordinates of every single observation, but no characteristics of the pattern in general. A more sophisticated approach, which will be pursued in this paper, consists of interpreting the observed pattern of exploded bombs as a realization of a spatial point process. This point of view is widespread in the analysis of the locations of lightning strikes (Schabenberger and Gotway, 2005) or earthquakes (Choi and Hall, 1999). McDonald and Small (2009) used spatial point process methodology for analysing patterns of unexploded ordnance at former air force bombing ranges.

This paper is organized as follows. In Section 2, a brief introduction to spatial point processes is given. Then, two methods for determining high-risk zones are introduced. As the first method is based on the intensity of the spatial point process, the estimation of intensities is described in detail. The second method is based on the nearest neighbour distance and represents a heuristic development of the method which is currently used. An evaluation procedure for these methods is presented in Section 3. We find that the two new methods yield similarly good results and discuss reasons why the intensity-based method should be preferred nonetheless. Section 4 deals with a simulation procedure for assessing the risk that is associated with a high-risk zone which has been determined by using the intensity-based construction method. In Section 5, the consequences of additional spatial clustering are investigated and we show that the intensity-

based method can be applied in this case as it is conservative. In the final Section 6, the application of the presented methods in practice is discussed.

The data that are analysed in the paper can be obtained from

`http://www.blackwellpublishing.com/rss`

2. Method

2.1. Spatial point processes

The observed patterns of exploded bombs can be regarded as realizations of spatial point processes (Diggle, 2003; Illian *et al.*, 2008; Møller and Waagepetersen, 2003; Schabenberger and Gotway, 2005). Let W be the observation window of the spatial point process X, i.e. the border of the property of interest and of the area for which data are available. The process X represents the locations of all bombs, exploded as well as unexploded. Then the random variable $N_X(A)$ denotes the number of bombs in a region $A \subseteq W$. The intensity measure $\Lambda_X(A)$ equals the expected number of bombs in A, $E\{N_X(A)\}$, and the intensity function $\lambda_X(\mathbf{s})$, which is defined via

$$\Lambda_X(A) = \int_A \lambda_X(\mathbf{x})\,\mathrm{d}\mathbf{x},$$

represents the probability of an observation in an infinitesimal disc centred on a given location $\mathbf{s} \in W$.

In what follows, we shall assume that the bombing process X can be described by an inhomogeneous Poisson point process with intensity $\lambda_X(\mathbf{s})$. This means that the number of observations in a region $A \subseteq W$ follows a Poisson distribution, i.e. $N_X(A) \sim \mathrm{Po}\{\Lambda_X(A)\}$, and $N_X(A)$ and $N_X(B)$ are independent for disjoint regions $A \subseteq W$ and $B \subseteq W$. However, only a thinned version Y of the full process X, namely the exploded bombs, has been observed. It consists of the $N_Y(W) = n_Y$ observations whose co-ordinates can be derived from the aerial pictures. The process of unexploded bombs, $Z = X \backslash Y$, is unobserved, i.e. its locations are unknown. The probability q of non-explosion for every bomb is assumed to be homogeneous in W. Thus, Y and Z are inhomogeneous Poisson point processes with intensity functions $\lambda_Y(\mathbf{s}) = (1-q)\,\lambda_X(\mathbf{s})$ and $\lambda_Z(\mathbf{s}) = q\,\lambda_X(\mathbf{s})$ respectively.

2.2. Determination of a high-risk zone based on the underlying intensity

The aim is to construct a high-risk zone which comprises a large fraction of the unexploded bombs while covering a small area. We propose a method based on inhomogeneous Poisson point process methodology, called the *intensity-based construction method*: the high-risk zone consists of those locations in the observation window W for which the intensity is largest.

First, the intensity function $\lambda_Y(\mathbf{s})$ is estimated. As the probability q of non-explosion for every bomb is assumed to be homogeneous, we can use $\hat{\lambda}_Y(\mathbf{s})$ to estimate the intensity function of the process Z: $\hat{\lambda}_Z(\mathbf{s}) = q/(1-q) \times \hat{\lambda}_Y(\mathbf{s})$. The region within the contours defined by $\hat{\lambda}_Z(\mathbf{s}) = c$ forms the high-risk zone. In the following sections, we specify how $\lambda_Y(\mathbf{s})$ is estimated and explain how to determine the threshold $c > 0$.

2.2.1. Estimation of the intensity function by kernel methods

The intensity $\lambda_Y(\mathbf{s})$ is unknown and needs to be estimated from the pattern of exploded bombs. We use a kernel method (Diggle, 1985; Baddeley, 2010) for this:

$$\hat{\lambda}_Y(\mathbf{s}) = e(\mathbf{s}) \sum_{i=1}^{n_Y} K_H(\mathbf{s} - \mathbf{y}_i), \tag{1}$$

where $K_H(\cdot)$ is an anisotropic Gaussian kernel and $e(\mathbf{s})$ is an edge effect bias correction of the form $e(\mathbf{s})^{-1} = \int_W K_H(\mathbf{s} - \mathbf{v})\,\mathrm{d}\mathbf{v}$. The correction is necessary because observations lying outside the window, possibly near the boundary, are not taken into account, which would result in a negative bias around the boundary of the observation window. The variance–covariance matrix H of the Gaussian kernel determines the smoothing bandwidth. Wand and Jones (1993) recommended the use of unconstrained (i.e. not necessarily diagonal) matrices H, thus allowing the Gaussian kernel to have arbitrary orientation. We select the bandwidth matrix H by using smooth cross-validation (Duong, 2007), an approach which is traditionally used for kernel density estimators:

$$\hat{f}(\mathbf{s}, H) = \frac{1}{n_Y} \sum_{i=1}^{n_Y} K_H(\mathbf{s} - \mathbf{y}_i). \tag{2}$$

We note that edge correction is not part of this method. The normalizing constant $1/n_Y$ does not affect the optimal choice of the bandwidth because n_Y is fixed. Details about smooth cross-validation and related criteria can be found in Wand and Jones (1995) and Duong and Hazelton (2005). Fig. 1 in the on-line appendix to this paper illustrates the estimated intensity for examples A and B.

2.2.2. *Determination of the threshold c*

For a given threshold $c > 0$, a high-risk zone R_c can be defined as

$$R_c = \{\mathbf{s} \in W : \hat{\lambda}_Z(\mathbf{s}) \geqslant c\} = \left\{\mathbf{s} \in W : \frac{q}{1-q}\hat{\lambda}_Y(\mathbf{s}) \geqslant c\right\}. \tag{3}$$

The general idea is closely related to the principle of highest posterior density intervals known from Bayesian inference. However, the interpretation is different because—unlike the density—the intensity is not normalized. To find an appropriate value for c, the *failure probability* of high-risk zone R_c,

$$P\{N_Z(W \setminus R_c) > 0\}, \tag{4}$$

i.e. the probability that not all unexploded bombs are covered by the high-risk zone, is considered. We want to find a high-risk zone R_c for which this probability is low, say (less than or) equal to a fixed value $0 \leqslant \alpha \leqslant 1$. Since we aim at high-risk zones which are as small as possible, we consider only the case of equality.

As pointed out, the number of unexploded bombs $N_Z(W)$ is unknown. But Z is assumed to be an inhomogeneous Poisson point process, so

$$N_Z(A) \sim \mathrm{Po}\{\Lambda_Z(A)\}, \qquad \Lambda_Z(A) = q\,\Lambda_X(A) = \frac{q}{1-q}\Lambda_Y(A). \tag{5}$$

Given the probability of non-explosion q and an estimate of the intensity function, the estimated failure probability can be computed as

$$\begin{aligned}\hat{P}\{N_Z(W \setminus R_c) > 0\} &= 1 - \exp\{-\hat{\Lambda}_Z(W \setminus R_c)\}\{\hat{\Lambda}_Z(W \setminus R_c)\}^0 \frac{1}{0!} \\ &= 1 - \exp\left(-\left[\frac{q}{1-q}\left\{\int_{(W \setminus R_c)} \hat{\lambda}_Y(\mathbf{y})\,\mathrm{d}\mathbf{y}\right\}\right]\right).\end{aligned} \tag{6}$$

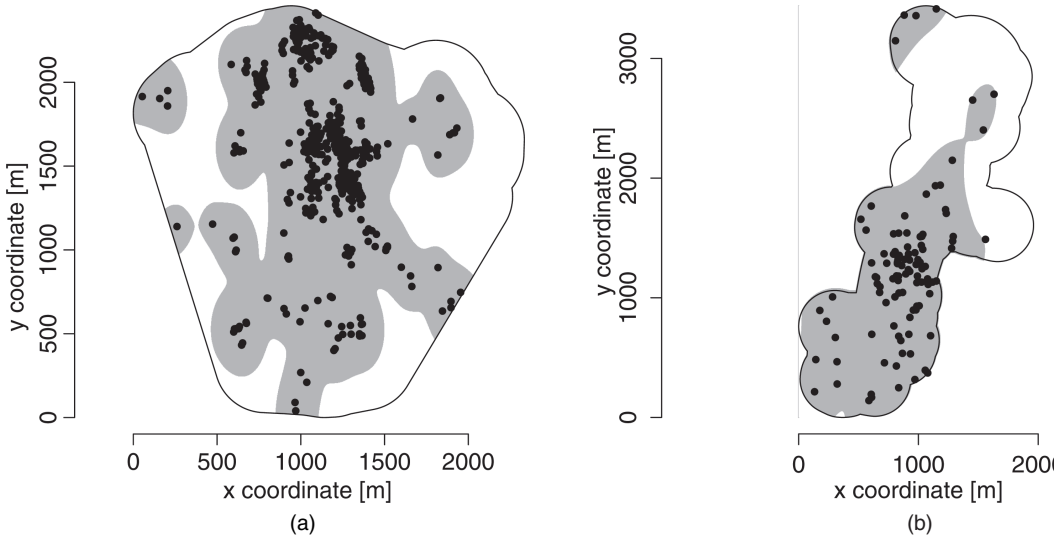

Fig. 2. Intensity-based high-risk zones (■) obtained for a probability of non-explosion of $q = 0.1$ and a maximal failure probability of $\alpha = 0.4$ for (a) example A and (b) example B

The threshold c for which $\hat{P}\{N_Z(W \setminus R_c) > 0\} = \alpha$ holds is determined by a numeric root finding procedure. Fig. 2 illustrates high-risk zones obtained with this construction method for the two example properties.

2.3. Implementation

The R package `spatstat` (Baddeley and Turner, 2005) for the analysis of spatial point patterns served as the main implementational toolbox for our work, e.g. point patterns were represented in its `ppp` data format and package functionality was used to estimate the intensity and to simulate point patterns. Bandwidth selection as described in Section 2.2.1 was taken from the R package `ks` (Duong, 2007).

2.4. Determination of a high-risk zone based on the nearest neighbour distance

The competing approach, called the *quantile-based construction method*, represents a heuristic development of the method that is currently used at the Oberfinanzdirektion Niedersachsen: discs of a fixed radius centred on each exploded bomb, whose union gives the high-risk zone. Instead of the arbitrary choice of the radius by an expert, a data-driven modification is applied. For every bomb $\mathbf{y}_i$ in Y, the distance to the nearest other bomb in Y, the so-called nearest neighbour distance

$$t_i = \min_{\mathbf{y}_j \in Y; j \neq i} \|\mathbf{y}_i - \mathbf{y}_j\|, \tag{7}$$

is computed and the empirical distribution function

$$G(r) = \frac{1}{n_Y} \sum_i \mathbf{1}\{t_i \leqslant r\} \tag{8}$$

of the nearest neighbour distances of the point pattern is determined. The radius of the discs is then obtained as the p-quantile $Q(p)$ of the distribution of this nearest neighbour distance,

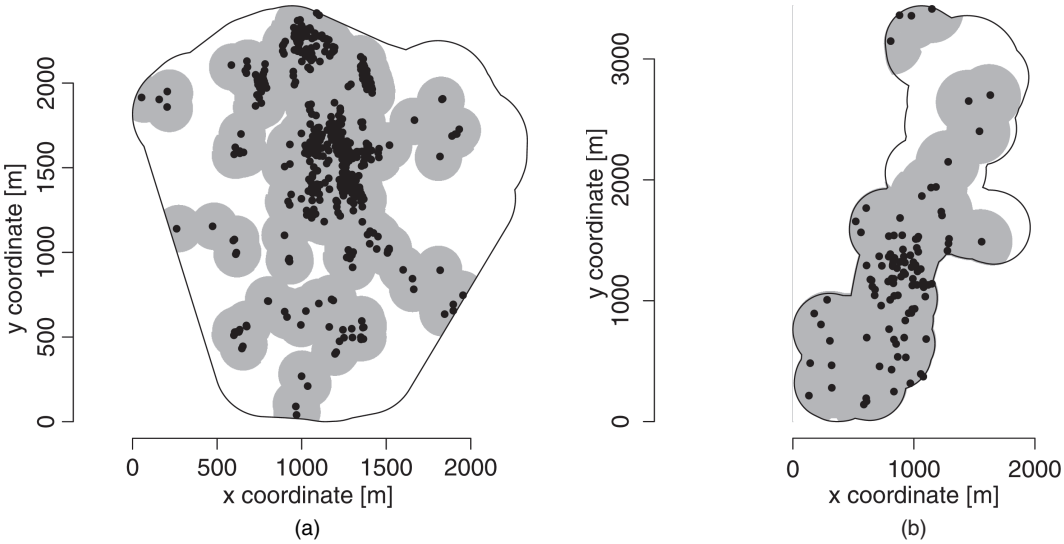

Fig. 3. Quantile-based high-risk zones (■) obtained for a probability of non-explosion of $q = 0.1$ by using the 99% quantile of the nearest neighbour distances as radius of the circles for (a) example A and (b) example B

where $0 \leqslant p \leqslant 1$ is a given value. For a given p, the high-risk zone consists of all locations $\mathbf{s}$ whose distance to the nearest observation does not exceed $Q(p)$, i.e. $R_p = \{\mathbf{s} \in W : \min_j \|\mathbf{s} - \mathbf{y}_j\| \leqslant Q(p)\}$. Quantile-based high-risk zones for examples A and B are shown in Figs 3(a) and 3(b) respectively.

3. Application and evaluation of the construction methods for high-risk zones

In this section, both construction methods from Sections 2.2 and 2.4 and the standard Oberfinanzdirektion Niedersachsen method are applied to examples A and B. First, the characteristics of these two patterns are examined by using Ripley's K-function. Then, we investigate whether the chosen inhomogeneous Poisson point process model with intensity function $\hat{\lambda}_Y(\mathbf{s})$ is appropriate. For this, a Monte Carlo test based on the (inhomogeneous) K-function is used. Finally, we present our evaluation procedure and discuss the results that we obtained for the three construction methods.

3.1. Model assumptions

A well-established tool for investigating homogeneous spatial point patterns is Ripley's K-function (Ripley, 1977)

$$K(r) = \frac{1}{\lambda} E[N\{b(o,r) \backslash \{o\}\}],$$

defined for $r \geqslant 0$. Note that the intensity function of homogeneous processes is $\lambda(\mathbf{s}) = \lambda$, a constant. Thus, multiplied with λ, the K-function corresponds to the expected number of observations in a circle $b(o,r)$ with radius r centred on a typical point o. Possible observations at the point o itself are not taken into account. The estimated K-function for an observed pattern is compared with the K-function of a pattern generated by complete spatial randomness, i.e. a homogeneous Poisson process, for which the K-function is $K(r) = \pi r^2$.

When estimating the K-function, we use Ripley's isotropic correction (Ripley, 1988) as implemented in `spatstat`. The values of the estimated K-function clearly exceed the theoretical

values for both example A and example B (Fig. 2 in the on-line appendix). This may suggest spatial clustering, but the deviation may also be due to inhomogeneity.

A generalization of the K-function is the inhomogeneous K-function (Baddeley *et al.*, 2000), which takes possible inhomogeneity into account by weighting every observation y_i with the inverse estimated intensity $1/\hat{\lambda}(y_i)$ at the respective location. Estimates for examples A and B are depicted in Fig. 3 and Fig. 4 in the on-line appendix.

For the intensity-based method, we assume that the observed patterns are realizations of an inhomogeneous Poisson point process with intensity function $\hat{\lambda}_Y(\mathbf{s})$. To check whether the model is appropriate, we performed a Monte Carlo test (Ripley, 1981) based on the (inhomogeneous) K-function with significance level 0.05, whose critical points are represented by a simultaneous envelope. The theoretical mean value of Ripley's K-function for patterns generated by our model was calculated as the average of 99 inhomogeneous Poisson point processes simulated from the estimated intensity function $\hat{\lambda}_Y(\mathbf{s})$. The symmetric envelopes are the result of adding and subtracting the fifth largest absolute difference between this average and the K-functions of 99 further simulated inhomogeneous Poisson point processes with intensity function $\hat{\lambda}_Y(\mathbf{s})$. Details can be found in Baddeley (2010). The same procedure was applied for the inhomogeneous K-function.

As seen from Fig. 4 and Fig. 5, all K-functions are entirely inside the envelopes, so the (inhomogeneous) K-functions of the observed patterns do not differ significantly from the (inhomogeneous) K-functions of an inhomogeneous Poisson point process with the estimated intensity function. As a consequence, there is no evidence against using such modelling to describe the observed bomb patterns. Note that, since the intensity of our Poisson process is estimated nonparametrically, it is not possible to consider residuals as proposed in Baddeley *et al.* (2005), as they are only defined for fully parametric models.

Furthermore, we performed a Monte Carlo test based on the pair correlation function (see Illian *et al.* (2008)). The results are presented in Fig. 5 in the on-line appendix. For example A, the inhomogeneous Poisson model is rejected. This is possibly due to clustering for $r < 50$ m. We discuss the consequences of such clustering in Section 5.

3.2. Simulation study

A simulation-based procedure is used to evaluate the construction methods for high-risk zones. To obtain patterns that are as realistic as possible, the following approach was chosen. The observed patterns Y were taken as 'full patterns', which means that we pretended to know the full process X. Denote by $\tilde{X} = Y$ this artificially defined full process. In the next step, each of the points of $\tilde{X}$ was thinned with probability q by drawing independent Bernoulli-distributed random variables, resulting in the process $\tilde{Y}$ of observed bombs and the process $\tilde{Z}$ of unobserved bombs. Of course, the intensities of $\tilde{X}$, $\tilde{Y}$ and $\tilde{Z}$ are smaller than those of the real processes X, Y and Z, but this procedure enables us to perform a simulation without making an assumption about the underlying point process model: it is not necessary to assume an inhomogeneous Poisson point process to obtain $\tilde{X}$, $\tilde{Y}$ and $\tilde{Z}$. Moreover, the parameter q never exceeds 0.15, so only a small fraction of the original observations is thinned and only little information is lost, so the scenario reflects the real problem in an appropriate way. The high-risk zone was then computed on the basis of only those observations belonging to $\tilde{Y}$, whereas the observations in $\tilde{Z}$ were used to evaluate the high-risk zone afterwards.

According to our co-operation partners from Oberfinanzdirektion Niedersachsen, a probability of non-explosion between 0.1 and 0.15 (depending on various factors such as the consistency of the subsoil and the year of the attack) can be regarded as a well-established value but, to our knowledge, no specific literature is available to document this further. As our work represents the

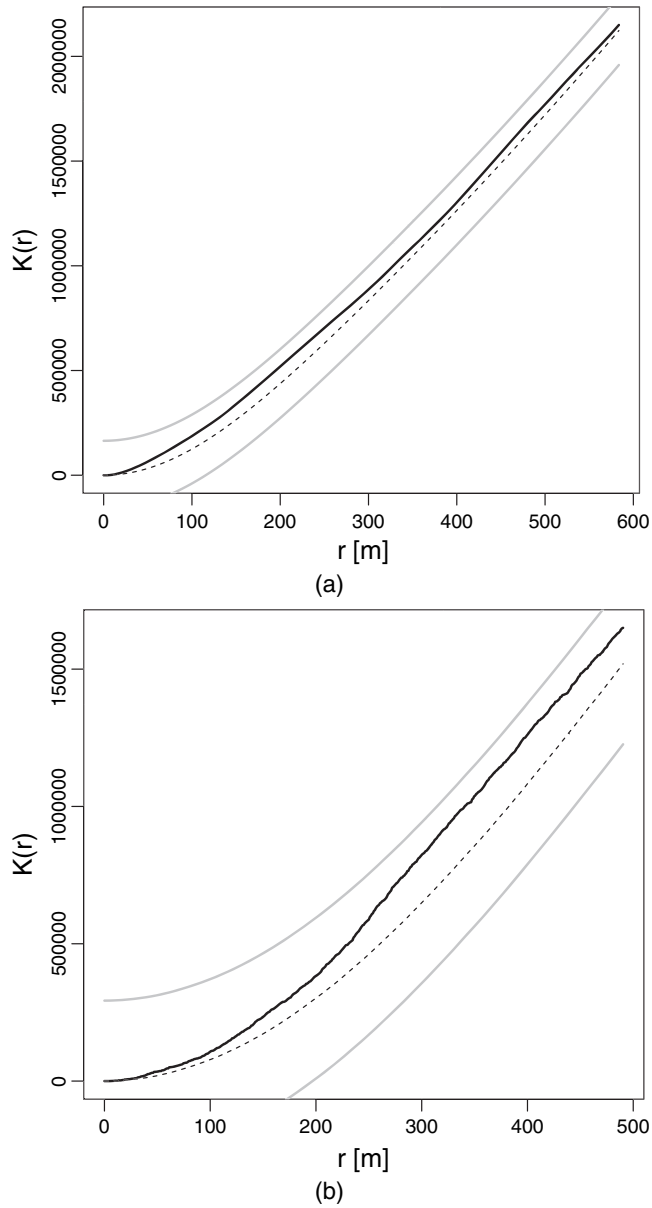

Fig. 4. Ripley's K-function and simultaneous envelopes generated by inhomogeneous Poisson point processes representing the critical points of a Monte Carlo test with significance level 0.05 for (a) example A and (b) example B: ———, observed K-function; - - - - - - -, estimated theoretical K-function; ———, simultaneous 95% envelope

first time that risk has been quantified in this field, no standard values for the failure probability α and the quantile p exist, so they were chosen on the basis of expert knowledge of our co-operation partners, and α was set to 0.4, 0.2 or 0.1 for both examples. These values may seem large. However, we need to take into account that an unexploded bomb outside the high-risk zone does not necessarily mean that somebody will die or be hurt. The bombs may be at depths of several metres where they might not be affected by construction work. Even if they are found,

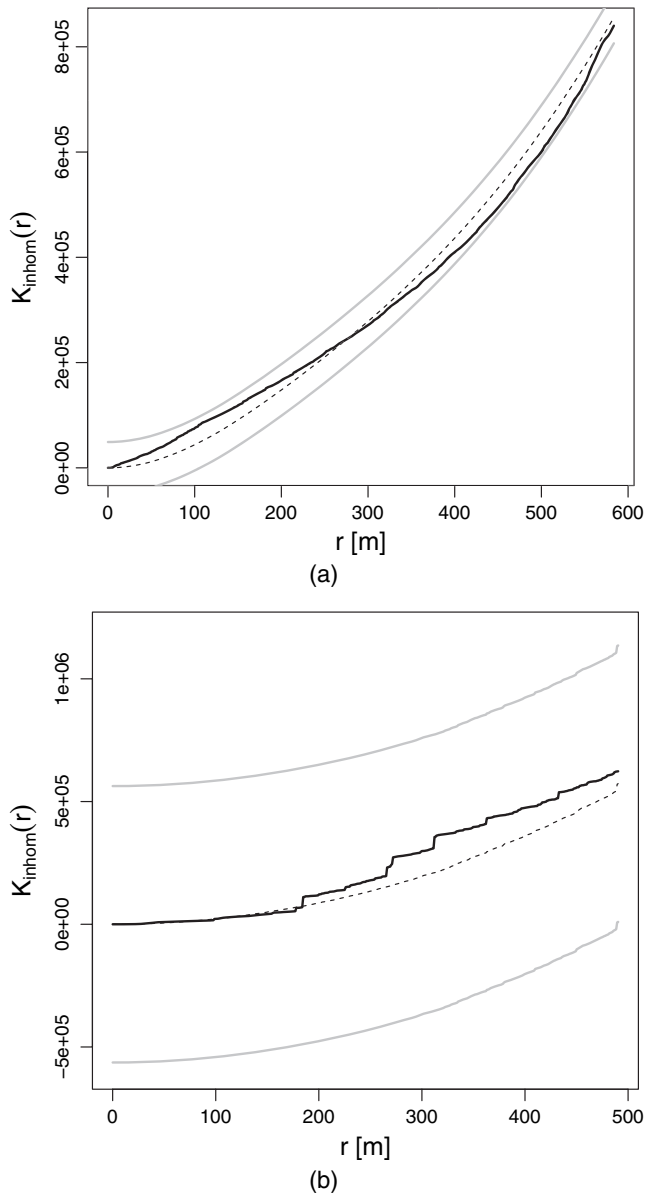

Fig. 5. Inhomogeneous K-function and simultaneous envelopes generated by inhomogeneous Poisson point processes representing the critical points of a Monte Carlo test with significance level 0.05 for (a) example A and (b) example B: ———, observed K-function; - - - - - - -, estimated theoretical K-function; ———, simultaneous 95% envelope

this does not necessarily mean that they will cause damage. To ensure comparability of the quantile-based and the intensity-based construction method, different quantiles were used for examples A and B. We considered three radii for the discs that were used in the standard method: according to Oberfinanzdirektion Niedersachsen, a radius of 50 m was used in the past, whereas 100 m and 150 m are common values nowadays. In each of those settings, 1000 iterations—in which $\tilde{X}$ was thinned and thus different processes $\tilde{Y}$ and $\tilde{Z}$ were obtained—were performed.

We are interested in finding the method which yields the smallest zones covering as many unexploded bombs as possible. Moreover, we investigate whether the parameters α and p of the intensity-based and the quantile-based method are adhered to. The probability that at least one unexploded bomb lies outside the high-risk zone should be α in the intensity-based construction method. We therefore compared the computed fraction p_{out} of generated high-risk zones for which at least one unexploded bomb was outside with α. For the quantile-based construction method, the fraction p_{miss} of unexploded bombs outside the high-risk zone should take a value near $1-p$. Therefore, the fraction p_{miss} was computed in every iteration, as well as the area of the zone.

The results for example A are shown in Table 1, which contains the mean and the standard deviation of p_{miss} and the area, as well as the fraction p_{out}. For the quantile-based method, the mean of p_{miss} of 1000 iterations is close to $1-p$ for all six combinations of parameters, whereas, for the intensity-based construction method, the fraction p_{out} of generated high-risk zones for which at least one unexploded bomb was outside exceeds α in most cases. The relative bias $(p_{\text{out}}-\alpha)/\alpha$ of the intensity-based method is between -0.063 and 1.840; the mean of the relative bias $\{p_{\text{miss}}-(1-p)\}/(1-p)$ of the quantile-based method between -0.021 and 1.448. The results for example B are shown in Table 2 in the on-line appendix, since they are similar to the results for example A. The relative bias $(p_{\text{out}}-\alpha)/\alpha$ of the intensity-based method is between 0.233 and 1.185, whereas the mean of the relative bias $\{p_{\text{miss}}-(1-p)\}/(1-p)$ is between -0.036 and 0.517. The resulting high-risk zones constructed using the standard method are very small for both examples. The mean fraction of unexploded bombs outside the high-risk zone and the fraction of generated high-risk zones for which at least one unexploded bomb was outside are large, especially for example B. This would of course change if larger radii were chosen, but the choice of the radius will always remain arbitrary unless the quantile-based method is used.

The values that were obtained for the intensity-based and the quantile-based construction methods are similar: there are only small differences in p_{out}, the mean of p_{miss} and the mean area.

Table 1. Results of the simulation: mean and standard deviation of the fraction p_{miss} of unexploded bombs outside the high-risk zone from 1000 iterations, fraction p_{out} of generated high-risk zones for which at least one unexploded bomb was outside and mean and standard deviation of the area of the zone, example A, intensity-based method, quantile-based method and standard method

Method		*Results for the following values of q:*					
		0.1	*0.1*	*0.1*	*0.15*	*0.15*	*0.15*
Intensity	α	0.4	0.2	0.1	0.4	0.2	0.1
	Mean p_{miss}	0.011	0.006	0.005	0.010	0.005	0.005
	Standard deviation p_{miss}	0.016	0.011	0.010	0.013	0.008	0.008
	p_{out}	0.375	0.249	0.192	0.498	0.318	0.284
	Mean area (m^2)	2711785	2987652	3186754	2870287	3115226	3295427
Quantile	p	0.99	0.995	0.999	0.99	0.995	0.999
	Mean p_{miss}	0.010	0.005	0.002	0.011	0.006	0.002
	Standard deviation p_{miss}	0.016	0.011	0.007	0.015	0.010	0.007
	p_{out}	0.324	0.196	0.086	0.465	0.285	0.139
	Mean area (m^2)	2663233	3097025	3345761	2701678	3108936	3367984
Standard	Radius (m)	50	100	150	50	100	150
	Mean p_{miss}	0.086	0.019	0.010	0.092	0.022	0.011
	Standard deviation p_{miss}	0.042	0.021	0.015	0.035	0.019	0.013
	p_{out}	0.979	0.575	0.377	1.000	0.751	0.545
	Mean area (m^2)	927549	1873954	2646576	907818	1847601	2622831

An important aspect for the comparison of the construction methods is the way in which the probability of non-explosion q is taken into account. In the standard method, q does not influence the shape of the high-risk zone at all. As the process $\tilde{Y}$ is obtained by independent thinning of $\tilde{X}$, it contains a smaller number of points for $q = 0.15$ than for $q = 0.1$. So the high-risk zones constructed with the standard method are even smaller for a higher probability of non-explosion in our simulation setting. The quantile-based method does not explicitly take into account the probability of non-explosion, q, either. However, nearest neighbour distances will typically be larger for high values of q and so will the radius defined by the p-quantile. In our setting, we observe that the mean area of the high-risk zone increases if $q = 0.15$ is considered, but usually so do the mean of p_{miss} and the fraction p_{out}. In other words, the probability of non-explosion is not taken into account sufficiently. Indeed, the high-risk zones become larger with increasing q, but the failure probability rises nonetheless. The only approach which uses q as a parameter is the intensity-based method. The results show that the high-risk zones have a larger area for $q = 0.15$ than for $q = 0.1$. In some cases, p_{out} and the mean of p_{miss} decreased.

In summary, we cannot recommend the use of the standard method, but at this point we remain indifferent about the quantile-based and the intensity-based method. Furthermore, the simulation revealed that the specified parameters α and p of the intensity-based and quantile-based method are not exactly adhered to.

3.3. High-risk zones with fixed area

Both the intensity-based and the quantile-based construction method can instead be used to construct a zone with a specified area. The standard method does not need to be considered in this context, since it can be interpreted as a simplified version of the quantile-based method. If the area is fixed, it is easier to compare the intensity-based and the quantile-based method, because it is sufficient to consider p_{out} and p_{miss} to decide which method yields the better high-risk zones. In the quantile-based method, fixing the area of the high-risk zone means that the optimal radius must be determined. The union of all respective discs, centred on the observations, must result in the desired area. If this radius is not smaller than the minimum of the nearest-neighbour distances and not larger than their maximum, we can determine which quantile of the nearest-neighbour distances corresponds to this radius. For the intensity-based method, the cut-off value c needs to be determined: the area which results from all points with an intensity of at least c must have the desired area. We can then determine to which value of

$$\alpha = \hat{P}\{N_Z(W \backslash R_c) > 0\} = 1 - \exp\left[-\left\{\frac{q}{1-q}\hat{\Lambda}_Y(W \backslash R_c)\right\}\right] \tag{9}$$

this optimal c corresponds.

The results for example A are depicted in Table 2. The quantile-based method gave better high-risk zones, i.e. its values of p_{miss} and p_{out} were smaller than for the intensity-based method. The comparison of the two methods is also illustrated in Fig. 6, which shows boxplots of p_{miss} for the simulations.

Table 4 in the on-line appendix shows the results for example B. In five of six cases, the high-risk zones determined by using the intensity-based method gave better results. This relationship is also visible in Fig. 6.

In summary, both construction methods yielded comparable results in the simulation. To decide which method should actually be applied, however, we should keep in mind that the more heuristic quantile-based construction method has a vague theoretical justification. The high-risk zones consist of the union of discs around the observations, so their shape is less flexible

Table 2. Results of the simulation for a given area: mean and standard deviation of the fraction p_{miss} of unexploded bombs outside the high-risk zone from 1000 iterations, fraction p_{out} of generated high-risk zones for which at least one unexploded bomb was outside, mean and standard deviation of the area of the zone and the parameter p, example A, intensity-based method and quantile-based method

Method		*Results for the following values of q and area:*					
		q = 0.1, area = 1500000	*q = 0.15, area = 1500000*	*q = 0.1, area = 2000000*	*q = 0.15, area = 2000000*	*q = 0.1, area = 2500000*	*q = 0.15, area = 2500000*
Intensity	Mean p_{miss}	0.076	0.078	0.028	0.032	0.014	0.016
	Standard deviation p_{miss}	0.038	0.031	0.025	0.023	0.018	0.015
	p_{out}	0.976	0.993	0.712	0.875	0.466	0.668
	Mean area (m^2)	1500001	1500001	2000001	2000000	2500001	2500001
	Mean α	0.988	0.999	0.885	0.963	0.573	0.732
Quantile	Mean p_{miss}	0.036	0.038	0.016	0.019	0.012	0.013
	Standard deviation p_{miss}	0.027	0.022	0.019	0.017	0.015	0.013
	p_{out}	0.809	0.938	0.528	0.721	0.446	0.608
	Mean area (m^2)	1500002	1500002	2000001	2000000	2500002	2500002
	Mean p	0.961	0.960	0.982	0.980	0.986	0.986

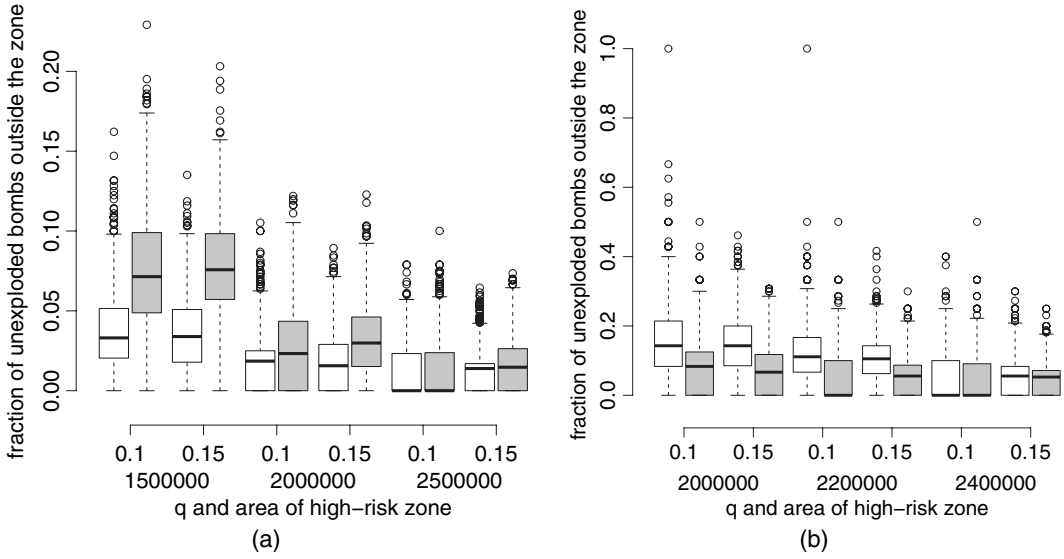

Fig. 6. Fraction p_{miss} of simulated unexploded bombs outside the high-risk zone for a given area for (a) example A and (b) example B: the boxplots compare the quantile-based (□) and intensity-based (■) methods

than the shape of intensity-based high-risk zones. In particular, possible anisotropy cannot be taken into account. Another disadvantage is that it is not possible to construct a quantile-based high-risk zone with an arbitrarily small failure probability. The smallest failure probability is obtained when the maximum of the nearest neighbour distances is used as the radius of the discs. However, the probability that the maximum of the nearest neighbour distances of the full

process exceeds the maximum of those of the thinned process (which is the failure probability) depends on the characteristics of the point pattern and cannot be influenced in any way. For these reasons, we recommend the intensity-based construction method, which is investigated in more detail in the following sections.

4. Simulation procedure based on the estimated intensity function

The results for the intensity-based method in Section 3 showed that the specified parameter α is usually not exactly adhered to. A possible reason is that the estimation of the intensity function on the basis of a small number of observations introduces uncertainty which is not accounted for in the construction method for high-risk zones. The problem is aggravated by thinning in the simulation scenario in Section 3. Therefore, it is not sufficient to choose a reasonable α expressing the intended failure probability, but it is indispensable to investigate the behaviour of the high-risk zone in a setting of realistically high intensity. To do so, we use a simulation procedure where the full process X of exploded as well as unexploded bombs is assumed to be an inhomogeneous Poisson point process with intensity $\lambda_X(\mathbf{s}) = 1/(1-q) \times \lambda_Y(\mathbf{s})$. In the first step, $\lambda_Y(\mathbf{s})$ is estimated by using the kernel method that was described in Section 2.2.1. In each of subsequent 10000 iterations performed, an inhomogeneous Poisson point process X^* with intensity $\lambda_{X^*}(\mathbf{s}) = 1/(1-q) \times \hat{\lambda}_Y(\mathbf{s})$ is simulated. As before, it is partitioned into the process Y^* of exploded and the process Z^* of unexploded bombs by drawing independent Bernoulli-distributed random numbers with probability q, i.e. $P(x^* \in Z^*) = q$. Finally, a high-risk zone based on Y^* is constructed by using the intensity-based method. This procedure can be regarded as a parametric bootstrap. The fraction p_{out} of high-risk zones for which at least one unobserved bomb from X^* is outside the zone expresses the risk. In addition, the fraction p_{miss} of unexploded bombs outside the high-risk zone and the area of the high-risk zone are computed.

Again, a probability of non-explosion q of 0.1 and 0.15 and $\alpha = 0.4$, $\alpha = 0.2$ and $\alpha = 0.1$ were considered. The results for examples A and B are shown in Table 3 and in Table 5 of the on-line appendix respectively. In example A, p_{out} was close to α for $q = 0.1$ and $\alpha = 0.4$, but clearly larger than α for all other combinations. Even larger differences were observed for example B. The relative bias $(p_{\text{out}} - \alpha)/\alpha$ was between 0.0187 and 0.5890 for example A and between 0.1382 and 1.0180 for example B.

Table 3. Bootstrap result: mean and standard deviation of the fraction p_{miss} of unexploded bombs outside the high-risk zone from 10000 iterations, fraction p_{out} of generated high-risk zones for which at least one unexploded bomb was outside and mean and standard deviation of the area of the zone; example A, intensity-based method and oracle estimator

Method		*Results for the following values of q and α:*					
		q = 0.1, α = 0.4	*q = 0.1, α = 0.2*	*q = 0.1, α = 0.1*	*q = 0.15, α = 0.4*	*q = 0.15, α = 0.2*	*q = 0.15, α = 0.1*
Intensity	Mean p_{miss}	0.011	0.006	0.004	0.008	0.004	0.003
	Standard deviation p_{miss}	0.016	0.012	0.009	0.011	0.008	0.007
	p_{out}	0.407	0.243	0.159	0.428	0.274	0.187
	Mean area (m^2)	3041859	3311320	3497413	3201929	3431759	3591073
Oracle	Mean p_{miss}	0.010	0.005	0.002	0.007	0.003	0.001
	Standard deviation p_{miss}	0.014	0.010	0.007	0.009	0.006	0.004
	p_{out}	0.398	0.202	0.100	0.398	0.199	0.097
	Area (m^2)	2752877	3027528	3226426	2914614	3154787	3331556

To investigate the reason for the deviation of p_{out} from α, the simulation procedure was modified. Instead of estimating the underlying intensity of the simulated patterns in every iteration, the high-risk zone was determined on the basis of the intensity which was used for simulation afterwards. We call this an oracle estimator, since it represents the most optimistic performance of our intensity-based method. In all cases, p_{out} was close to α. The relative bias $(p_{\text{out}} - \alpha)/\alpha$ was between -0.0260 and 0.0105 for example A and between 0 and 0.0630 for example B. This indicates that the estimation of the intensity is a crucial issue for the procedure and that failure to meet the specification originates from this step. A reason may be that the number of observations is too sparse to guarantee optimal estimation.

To correct for the bias with respect to α, we propose to use the parametric bootstrap in the following way.

(a) For a given spatial point pattern, the intensity function is estimated and the simulation procedure based on the estimated intensity function is applied as described above.
(b) The fraction p_{out} represents an estimator for the failure probability, so the procedure yields an updated failure probability, which is possibly different from the intended α.

The procedure can be repeated until we find the value for the parameter α which results in the intended failure probability. To find this value more quickly, the search can be realized in an adaptive procedure.

An alternative approach for constructing high-risk zones based on the estimated intensity results if the cut-off value c is specified directly instead of specifying α. The interpretation of a high-risk zone $R_c = \{\mathbf{s} \in W : \hat{\lambda}_Z(\mathbf{s}) \geqslant c\}$ determined in this way is that the expected number of unexploded bombs does not exceed c in any location outside the high-risk zone. Afterwards, it is possible to determine the corresponding α as defined in equation (9) and to assess the risk by using the simulation procedure. High-risk zones obtained for $q = 0.1$ when cutting at $\hat{\lambda}_Z(\mathbf{s}) = 0.000\,001$ are shown in Fig. 8 in the on-line appendix.

5. Spatial clustering

The starting point of the intensity-based method is the assumption of an inhomogeneous Poisson point process. This assumption implies that, beyond spatial variation in the intensity function, there is no stochastic dependence between observations. However, we obtained high values for Ripley's K-function and the pair correlation function in Section 3.1, which may indicate clustering or inhomogeneity. As Diggle *et al.* (2007) stated, there is a 'fundamental ambiguity' between these two phenomena. Both mechanisms generate patterns with aggregation, so they are difficult to distinguish—see also Ripley (1988). An extension of the Poisson process is the Cox process (Cox, 1955), where the intensity $\lambda(\mathbf{s})$ is replaced by a non-negative random field $\Phi(\mathbf{s})$. Conditional on Φ, the Cox process is a Poisson process with intensity function Φ. Because of this relationship, an inhomogeneous Poisson process and a Cox process cannot be distinguished when only one realization is available (Møller and Waagepetersen, 2007).

Therefore, we consider the process model which fits the subject matter theory best (see Schabenberger and Gotway (2005)). According to our co-operation partners Oberfinanzdirektion Niedersachsen, this is rather the case for inhomogeneous Poisson models than for cluster models: the assumption of several targets in a property is not justified for most examples. In some cases, like for example B, the target of the attack even seems to be outside the property to be cleared. However, as clustering cannot be ruled out, it appears advisable to perform a sensitivity analysis and to investigate the behaviour of our intensity-based method in case the bombing point pattern is generated by a cluster process instead of an inhomogeneous Poisson point process.

We used the class of Neyman–Scott processes (Neyman and Scott, 1958), which are special Poisson cluster processes (Bartlett, 1964; Møller and Waagepetersen, 2007; Illian *et al.*, 2008). The (unobserved) cluster centres form a Poisson point process. Each cluster centre has a random number of offspring: the cluster points. The number of cluster points is independent and identically distributed with a discrete probability mass function. The positions of the cluster points relative to the cluster centres are independent and identically distributed according to a bivariate distribution function. The Poisson cluster process consists of the cluster points only (Schabenberger and Gotway, 2005). To obtain simulated processes resembling the patterns observed , we applied the definition of Cressie (1993), which states that the cluster centres may form an inhomogeneous Poisson point process, whereas most other researchers postulate a homogeneous Poisson process (e.g. Stoyan *et al.* (1995)).

Note that classical cluster processes such as the Thomas or the Matérn process, whose parameters can be estimated with the method of minimum contrast by using the K-function (Diggle and Gratton 1984; Møller and Waagepetersen, 2003; Waagepetersen, 2007), cannot describe the observed patterns in an appropriate way. Simulations based on these point process models yielded patterns whose points were concentrated on only a small fraction of the original area. Using a Neyman–Scott point process model results in a modified simulation procedure for X^* compared with Section 4. First, the cluster centres were simulated as an inhomogeneous Poisson point process with intensity function

$$\lambda_{C^*}(\mathbf{s}) = \frac{1}{\nu(1-q)}\hat{\lambda}_Y(\mathbf{s}).$$

The number of points per cluster follows a Po(ν) distribution; the cluster points are placed independently and uniformly inside a disc of radius r centred at the cluster centres.

Here, the parameter $\nu > 0$ determines the extent of clustering. If ν is small, the process of cluster centres will contain almost as many points as the simulated Poisson processes in Section 4. Each of these centres will be replaced by a small number of cluster points. Note that, even for $\nu = 1$, the process is clustered: even though the process of cluster centres is a Poisson process with the same intensity as in the previous section, the resulting process is not, since each cluster centre is replaced by a small number of cluster points and this number may be 0 or larger than 1. If ν is large, the process of cluster centres will consist of a small number of points, whereas the clusters will be big.

The radius r in the sensitivity analysis was chosen to be 80 m, which is a value that is larger than most of the observed nearest neighbour distances in both examples, but sufficiently small to obtain clearly visible clusters (see Fig. 9 in the on-line appendix). Six different values for ν were considered to study the consequences of different extents of clustering. 1000 iterations were performed for each combination of parameters. All other aspects of the simulation setting remained as in Section 4.

To facilitate the comparison between the cluster model and the inhomogeneous Poisson process model, we integrated the results obtained with the latter into the figures. As we can see in Fig. 7, the intensity-based construction method for high-risk zones is conservative if the pattern is a cluster process instead of an inhomogeneous Poisson process. In most cases, the mean fraction p_{out} was smaller than the α-values of 0.4, 0.2 and 0.1. In particular, the numbers are smaller than for the inhomogeneous Poisson point process that was investigated in Section 4. The larger the value of ν, i.e. the more clustered the pattern, the smaller the value of p_{out}. Fig. 10 in the on-line appendix illustrates that smaller high-risk zones result as a consequence of larger ν. For example A, the high-risk zones were usually larger than in an inhomogeneous Poisson process, whereas they were often smaller for example B.

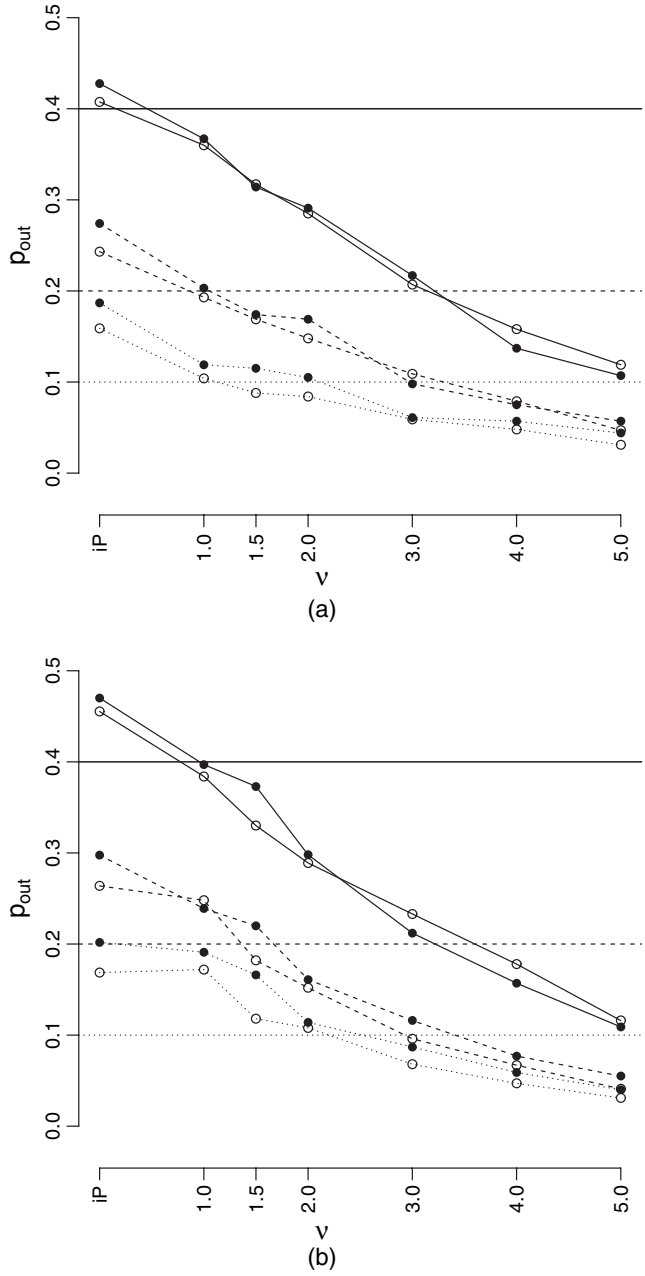

Fig. 7. Fraction of high-risk zones with at least one exploded bomb outside for different values of the cluster parameter ν and for the inhomogeneous Poisson process model for (a) example A and (b) example B: —○—, $\alpha = 0.4$, $q = 0.1$; —●—, $\alpha = 0.4$, $q = 0.15$; - -○- -, $\alpha = 0.2$, $q = 0.1$; - -●- -, $\alpha = 0.2$, $q = 0.15$; ⋯○⋯, $\alpha = 0.1$, $q = 0.1$; ⋯●⋯, $\alpha = 0.1$, $q = 0.15$

6. Discussion

The quantification of risk is an important step when making rational decisions for clearance of unexploded bombs. As considerable costs of clearance must be taken into account, it is rarely possible to investigate the entire property of interest. As a consequence, it is necessary to determine high-risk zones covering a small area and a large fraction of the unexploded bombs at the same time. Our methods described in this paper provide a quantification of risk to guide the decision-making process. We have presented a construction method for high-risk zones based on the intensity of the underlying spatial point process. The high-risk zones consist of those parts of the observation window where the intensity is largest, i.e. the intensity function exceeds a cut-off value c. In general, the heuristic quantile-based construction method yielded comparably good results but has a vague theoretical justification. The shape of these high-risk zones is less flexible than the shape of intensity-based high-risk zones, anisotropy cannot be taken into account and it is not possible to construct a quantile-based high-risk zone with an arbitrarily small failure probability, so we recommend use of the intensity-based method.

As we have demonstrated, the estimation of the point process intensity function is a crucial issue for our intensity-based method. It seems to be the reason why, in many cases, α was not exactly adhered to. However, by using the bootstrap correction, it is possible to decide on acceptable risks and to find the optimal, i.e. smallest, high-risk zone for a fixed probability that not all unexploded bombs are inside the high-risk zone. Other approaches to estimate the intensity, such as the adaptive estimator in the R package `sparr` (Davies *et al.*, 2011), which also corrects for edge effect bias, might help to improve the intensity-based method in general.

Cluster processes and inhomogeneous Poisson point processes cannot be distinguished on the basis of a single realization. From an empirical and subject matter point of view, the inhomogeneous Poisson point process represents a reasonable model. However, the assumption of clustering is clearly relevant. The intensity-based construction method cannot take clustering into account; nor can the degree of clustering be determined. However, the consequences of clustering can be investigated in a sensitivity analysis. Here, we found that the intensity-based method is conservative in an underlying cluster process. If there really is clustering, further developments are needed to extend our method to provide more efficient high-risk zones. If it was possible to estimate the clustering parameter ν, a correction could be used for the intensity-based method. However, given the impossibility of distinguishing inhomogeneous and clustered patterns, this seems to be a difficult task and we have not found any approach which enables us to estimate ν on the basis of a single realization of the process. The correction can be applied, though, if any prior knowledge about the number of bombs in a cluster is available. For example, it may be possible to derive such information from historical records of the allies, for instance the number of bombs dropped simultaneously by the bombing planes.

The intensity-based construction method is flexible and can be applied in different ways and for various situations. If the budget for investigation is fixed, this is usually equivalent to a fixed area of the high-risk zone. In this case, a high-risk zone with the desired area can be determined and the risk can be assessed by using the simulation procedure based on the estimated intensity function. The intensity-based method can also be used to construct high-risk zones such that the expected number of unexploded bombs does not exceed a fixed value c in any location outside the high-risk zone. The corresponding failure probability can then be determined by using the simulation procedure based on the estimated intensity function. For the probability q of non-explosion of a bomb, the values 0.1 and 0.15 have been considered in this paper. This probability is not fixed in general but depends on the subsoil and the year of the attack,

for example. As q is a parameter of the intensity-based construction method, this method can deal with a broad range of situations and its application is not restricted to the cases that were considered in this paper. If more information was available, it would be possible to model q as a function of space.

To demonstrate the applicability of the intensity-based construction method, we used the method on six additional bombing point patterns which partly represent very complex or particular situations from ordnance disposal in Germany (see Figs 11–13 in the on-line appendix). From the point of view of Oberfinanzdirektion Niedersachsen, the resulting high-risk zones clearly differ from the high-risk zones which have been used up to now. They have a plausible shape and a clear risk interpretation.

All the methods presented are also relevant in other applications. One may think of the locations of cases of infectious diseases as a spatial point pattern. If not all cases are reported to the appropriate authorities, only a thinned process is observed. Here, the temporal order of cases will, however, provide valuable extra information. But additional complexity due to inhomogeneous population structure and contact behaviour provides extra challenges.

Acknowledgements

We are grateful to Wilfried Möller and Christian Meinhardt (Oberfinanzdirektion Niedersachsen) and Andreas Bernhardt and Robert Brosy (Mull und Partner Ingenieurgesellschaft mbH, Hannover) who provided the data and additional information on the search for unexploded bombs. We also thank Volker J. Schmid (Ludwig-Maximilians-Universität München) for valuable discussions. The comments of the Joint Editor, as well as the comments of the Associate Editor and two referees, were very helpful in improving the paper substantially.

References

Baddeley, A. (2010) Analysing spatial point patterns in R. *Technical Report*. Commonwealth Scientific and Industrial Research Organisation. (Available from `www.csiro.au/resources/pf16h.html`.)

Baddeley, A., Møller, J. and Waagepetersen, R. (2000) Non- and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statist. Neerland.*, **54**, 329–350.

Baddeley, A. and Turner, R. (2005) Spatstat: an R package for analyzing spatial point patterns. *J. Statist. Softwr.*, **12**, 1–42.

Baddeley, A., Turner, R., Møller, J. and Hazelton, M. (2005) Residual analysis for spatial point processes (with discussion). *J. R. Statist. Soc.* B, **67**, 617–666.

Bartlett, M. S. (1964) The spectral analysis of two-dimensional point processes. *Biometrika*, **51**, 299–311.

Choi, E. and Hall, P. (1999) Nonparametric approach to analysis of space-time data on earthquake occurrences. *J. Computnl Graph. Statist.*, **8**, 733–748.

Cox, D. R. (1955) Some statistical models connected with series of events (with discussion). *J. R. Statist. Soc.* B, **17**, 129–164.

Cressie, N. A. C. (1993) *Statistics for Spatial Data*, 2nd edn. New York: Wiley.

Davies, T. M., Hazelton, M. L. and Marshall, J. C. (2011) sparr: analyzing spatial relative risk using fixed and adaptive kernel density estimation in R. *J. Statist. Softwr.*, **39**, 1–14.

Diggle, P. (1985) A kernel method for smoothing point process data. *Appl. Statist.*, **34**, 138–147.

Diggle, P. J. (2003) *Statistical Analysis of Spatial Point Patterns*, 2nd edn. London: Arnold.

Diggle, P. J., Gómez-Rubio, V., Brown, P. E., Chetwynd, A. G. and Gooding, S. (2007) Second-order analysis of inhomogeneous spatial point processes using case-control data. *Biometrics*, **63**, 550–557.

Diggle, P. J. and Gratton, R. J. (1984) Monte Carlo methods of inference for implicit statistical models. *J. R. Statist. Soc.* B, **46**, 193–212.

Duong, T. (2007) ks: kernel density estimation and kernel discriminant analysis for multivariate data in R. *J. Statist. Softwr.*, **21**, 1–16.

Duong, T. and Hazelton, M. L. (2005) Cross-validation bandwidth matrices for multivariate kernel density estimation. *Scand. J. Statist.*, **32**, 485–506.

Illian, J., Penttinen, A., Stoyan, H. and Stoyan, D. (2008) *Statistical Analysis and Modelling of Spatial Point Patterns*. Chichester: Wiley.

McDonald, J. A. and Small, M. J. (2009) Statistical analysis of metallic anomaly patterns at former air force bombing ranges. *Stoch. Environ. Res. Risk Assessmnt*, **23**, 203–214.
Møller, J. and Waagepetersen, R. P. (2003) *Statistical Inference and Simulation for Spatial Point Processes*. Boca Raton: Chapman and Hall–CRC.
Møller, J. and Waagepetersen, R. P. (2007) Modern statistics for spatial point processes. *Scand. J. Statist.*, **34**, 643–684.
Neyman, J. and Scott, E. L. (1958) Statistical approach to problems of cosmology (with discussion). *J. R. Statist. Soc.* B, **20**, 1–43.
Ripley, B. D. (1977) Modelling spatial patterns (with discussion). *J. R. Statist. Soc.* B, **39**, 172–212.
Ripley, B. D. (1981) *Spatial Statistics*. New York: Wiley.
Ripley, B. D. (1988) *Statistical Inference for Spatial Processes*. Cambridge: Cambridge University Press.
Schabenberger, O. and Gotway, C. (2005) *Statistical Methods for Spatial Data Analysis*. Boca Raton: Chapman and Hall–CRC.
Stoyan, D., Kendall, W. S. and Mecke, J. (1995) *Stochastic Geometry and Its Applications*, 2nd edn. Chichester: Wiley.
Waagepetersen, R. P. (2007) An estimating function approach to inference for inhomogeneous Neyman-Scott processes. *Biometrics*, **63**, 252–258.
Wand, M. P. and Jones, M. C. (1993) Comparison of smoothing parameterizations in bivariate kernel density estimation. *J. Am. Statist. Ass.*, **88**, 520–528.
Wand, M. P. and Jones, M. (1995) *Kernel Smoothing*. London: Chapman and Hall.

Supporting information
Additional 'supporting information' may be found in the on-line version of this article:

'Web Appendix: Determining high-risk zones for unexploded WWII bombs using point process methodology'.