

MARKED POINT PROCESSES FOR THE AUTOMATIC DETECTION OF BOMB CRATERS IN AERIAL WARTIME IMAGES

C. Kruse*, F. Rottensteiner, C. Heipke

Institute of Photogrammetry and GeoInformation, Leibniz Universität Hannover, Germany
(kruse, rottensteiner, heipke)@ipi.uni-hannover.de

Commission II, ICWG II/III

KEY WORDS: Marked Point Processes, RJMCMC, Simulated Annealing, Aerial Wartime Images, Bomb Craters

ABSTRACT:

Many countries were the target of air strikes during the Second World War. The aftermath of such attacks is felt until today, as numerous unexploded bombs or duds still exist in the ground. Typically, such areas are documented in so-called impact maps, which are based on detected bomb craters. This paper proposes a stochastic approach to automatically detect bomb craters in aerial wartime images that were taken during World War II. In this work, one aspect we investigate is the type of object model for the crater: we compare circles with ellipses. The respective models are embedded in the probabilistic framework of marked point processes. By means of stochastic sampling the most likely configuration of objects within the scene is determined. Each configuration is evaluated using an energy function which describes the conformity with a predefined model. High gradient magnitudes along the border of the object are favoured and overlapping objects are penalized. In addition, a term that requires the grey values inside the object to be homogeneous is investigated. Reversible Jump Markov Chain Monte Carlo sampling in combination with simulated annealing provides the global optimum of the energy function. Afterwards, a probability map is generated from the automatic detections via kernel density estimation. By setting a threshold, areas around the detections are classified as contaminated or uncontaminated sites, respectively, which results in an impact map. Our results, based on 22 aerial wartime images, show the general potential of the method for the automated detection of bomb craters and the subsequent automatic generation of an impact map.

1. INTRODUCTION

1.1 Motivation

Although the last combat operations of the Second World War took place more than 70 years ago, their aftermath is still present today. Unexploded ordnance, such as grenades or bombs, remain hidden in the ground. Bombs that were dropped by planes are particularly dangerous due to their high explosive force. Experts of Lower Saxony's explosive ordnance disposal service assume that approx. 10 % - 15 % of all bombs did not detonate. During construction works there is still a real danger today that these duds explode, and there have also been incidents without external influence. Typically, surveillance flights were carried out shortly after an air strike. The resulting images are being used today to find potentially dangerous sites. In Germany, manual interpretation of these images is carried out by the explosive ordnance disposal services of the respective federal states. In this context, a central task is the identification of duds in the images. Usually, such investigations are restricted to particularly endangered or otherwise relevant areas. Nevertheless, the processing effort is still immense. For many applications, it is sufficient to have comprehensive information on the basic occurrence of warlike impacts in the form of "impact maps". An impact map indicates whether areas are likely to be contaminated or not. In this context, contaminated areas are expected to contain one or more duds with a high degree of certainty, whereas uncontaminated areas should not contain any dud. The automatic creation of such a map could accelerate the manual evaluation process and thus save time and money. To do so, an automatic detection of indications of bombing in aerial wartime images, especially of bomb craters, is essential.

The problem we intend to solve is the automatic detection of duds. In order to deduce a probability for their occurrence, bomb craters in aerial wartime images are used as they indicate the areas where unexploded bombs may be located. Hence, this probability can be used to identify contaminated areas to be represented an impact map. The paper focuses on a scenario in which the correctness of the results is important: we are interested in detecting areas that have a very high likelihood of containing a dud so that it makes sense to send a team of experts to that area to probe it, i.e. to visit the site and take measurements using geophysical detectors. As this is expensive, false detections have to be minimised. The main benefit for the explosive ordnance disposal service is that in these areas the aerial images would then no longer have to be inspected manually.

For object detection from images, models are increasingly being used in a probabilistic framework, in which prior knowledge is integrated in the form of probabilities. Well-known approaches are Markov Random Fields (Geman and Geman, 1984) and Conditional Random Fields (Kumar and Hebert, 2006). By modelling context, knowledge about the objects can be integrated into the image. For example, similar classes for pixels in a local neighbourhood can be favoured in this way (Geman and Geman, 1984). However, it is difficult to integrate more global constraints about the objects, for example regarding their shape. This is where marked point processes (MPPs; Descombes and Zerubia, 2002; Daley and Vere-Jones, 2003) come into play. A MPP is a random process whose realizations are configurations of objects. This stochastic approach uses a strong object model which offers a considerable flexibility in integrating knowledge about the objects and their relationship to other objects. Sampling provides

* Corresponding author

the globally optimal configuration for objects of a certain type. Thus, knowledge about the objects can be integrated beyond pixel-based relations. The number of objects in the scene during sampling is variable. MPPs have shown to achieve good results in various object detection problems (e.g. Lafarge et al., 2010; Börns and Benedek, 2015; Benedek, 2017; Schmidt et al., 2017).

The detection probability of individual bomb craters is often rather low and their individual representation is not very descriptive to indicate areas that probably contain duds. Therefore, we are interested to deduce area-based statements about the occurrence of warlike impacts that are based on statistical modelling, e.g. by modelling the probability density function for a dud to occur. In general, parametric approaches may be used, where an analytical model for the probability density function is assumed, and training data provide the parameters of the function. On the other hand, nonparametric approaches estimate the density function directly from the data, which avoids having to select a model for the distribution and to estimate its parameters (Bishop, 2006). Kernel density estimation (Parzen, 1962) is a quite popular nonparametric density estimation technique (e.g. Ruan, 2010; Scott, 2015) and is also employed in our work.

In this work, we use MPPs to detect bomb craters in aerial wartime images, building on (Kruse et al., 2018), where craters were represented as ellipses; the results showed the potential of the method. During sampling, high magnitudes of the grey value gradients at the ellipse borders are favoured and an overlap of ellipses is penalized. In this paper, an additional term that requires the grey values inside the object to be homogeneous is analysed. Homogeneity is measured by the standard deviation of the grey values within the object. Another goal of this paper is to investigate how the use of a circle as an object model instead of an ellipse affects the results, because we think that a model with fewer parameters might be more stable. Finally, besides the general goal in this work to achieve results with a high quality, in our last investigations we want to tune the algorithm's parameters with regard to the proposed application scenario, i.e. the results should have the highest possible correctness, as long as completeness does not suffer too much. As in our previous work, we create a probability map based on the automatically detected bomb craters by kernel density estimation. By applying a threshold, areas around the detections are classified as contaminated or uncontaminated sites, respectively, which results in an impact map.

The remainder of this paper is structured as follows. First, related work on object detection tasks, mainly in connection with MPPs, is given (Section 1.2). The mathematical basics of MPPs and kernel density estimation are described in Section 2. In Section 3, we present the extended model we use for the detection of bomb craters and of the impact map. Experiments and results based on aerial wartime images are presented in Section 4. Finally, a summary and outlook for future work is given in Section 5.

1.2 Related Work

Here, we focus on applications of MPPs and methods for detecting craters, also including planetary craters, because they have a similar appearance as bomb craters.

In the context of MPPs, model knowledge can be integrated in different ways. Typically, simple geometric primitives, which can be described by a small number of parameters, are used to represent the objects to be detected. The extraction of buildings or other man-made objects in the scene is often accomplished

with rectangles (Tournaire et al., 2010; Chai et al., 2012; Brédif et al., 2013). In these papers, the input data are digital surface models. For a rectangle to be included in the object configuration during sampling, high gradient magnitudes of the heights at the rectangle border must be present. Furthermore, rectangles overlapping with each other are penalized. Rectangles were also used to interpret facades of buildings based on rectified images (Wenzel and Förstner, 2016), where statistics of typical configurations of facade objects (windows, entrances) were learned from training data. In addition to rectangles, ellipses, e.g. for the detection of flamingos (Descamps et al., 2008) and seed products (Dubosclard et al., 2014), or circles, e.g. to detect tree crowns (Zhang et al., 2013), are used. Benedek (2017) proposed a method for extracting complex hierarchical object structures from digital images using different types of objects, namely ellipses, rectangles and isosceles triangles, thus considering multiple object models. Inside this MPP framework object-subobject ensembles in parent-child relationship are admitted and corresponding objects may form coherent object groups. Bomb craters have no object-subobject relationships, so complex hierarchical object structures are not necessary in our case. MPPs are also applied to biomedical imagery. Kowal and Korbicz (2018) proposed a method for the detection of cell nuclei in microscopic images. In a Bayesian framework, regions with intensity distributions characteristic for nuclei are detected and approximated by circles. The process of circle generation can be viewed as a MPP that also penalizes an overlap of objects. Descombes (2017) uses a flexible approach based on MPPs to handle biological variabilities in the images. We have another application domain and the structure of the images is different. Cedilnik et al. (2018) propose a method to detect objects (small particles) that are only a few pixels wide by a dictionary of shapes. However, bomb craters are not only a few pixels wide and cannot easily be described by such a dictionary due to their partly variable shape.

In the context of planetary crater detection, we are aware of two contributions dealing with MPPs (Troglio et al., 2010; Solarna et al., 2017). To reduce the computational effort in the optimization process, Solarna et al. (2017) create a birth map from the available contour map via generalized Hough transform and Gaussian filtering. Other approaches, such as unsupervised (e.g. Meng et al., 2009) and supervised (e.g. Wetzler et al., 2005; Urbach and Stepinski, 2009) algorithms, have also been proposed. In the former case, common edge filters are applied to highlight the edges in the image. Afterwards, the Hough transform is used to reconstruct the circular shape of the craters. However, an expert is needed to choose the best filters and the procedure is not very robust to noise. Supervised methods used for crater detection include boosting (Bandeira et al., 2012) and Convolutional Neural Networks (Cohen et al., 2016). To the best of our knowledge, only three papers deal with the detection of bomb craters in aerial wartime images. Jensen et al. (2010) use a two-step approach. First, candidates are searched via cross correlation with representative crater-templates. Afterwards, the candidates are classified by linear discriminant analysis. Merler et al. (2005) use different boosting approaches for the classification of image sections, which requires a high computational effort. Their result is a map of the spatial density of craters, an indicator for the risk of finding duds. Brenner et al. (2018) proposed a method based on Convolutional Neural Networks. Algorithms based on machine learning need a large set of training data to be functional. However, there is currently not enough data available to use such approaches in an appropriate way.

The cited articles and our previous work (Kruse et al., 2018) show the potential of stochastic methods based on MPPs in

different fields. In addition to its applicability for heterogeneous image contents (e.g. varying crater sizes), the discussion of the related work shows that MPPs allow a flexible integration of knowledge about the objects, too. That is why we suppose the procedure of MPPs is suitable for our application scenario as it can be well adapted to our needs. Hence, this paper proposes a special type of MPPs based on another object model (circle) and an enhanced energy function compared to our previous work. In contrast to images of planetary surfaces, objects such as trees or houses make correct detections more challenging in our case. Apart from our previous work, this kind of detecting bomb craters with MPPs and the subsequent inclusion of the results in an impact map are not yet available in the literature.

2. MATHEMATICAL BASICS

2.1 Marked point processes

Marked point processes (Daley and Vere-Jones, 2003; Descombes, 2013) are stochastic processes describing a random configuration of objects of a certain type in a bounded region F , here a digital image. An object u_i is characterized by its position $p_i = (x_i, y_i)$ and a vector m_i (mark) of parameters of an object model. There are different types of point processes. The homogeneous Poisson point process assumes a purely random distribution of objects $u_i = (p_i, m_i)$ in space that are not related to each other. In a Poisson point process the probability $P_\lambda(n)$ for the number of objects n follows a discrete Poisson distribution. The *intensity parameter* λ describes the expected number of objects within F . The object positions are uniformly distributed. In practice, the assumption of complete randomness often does not apply, because there are dependencies between the objects. Models that are more complex are required to measure the quality of the object configuration. To achieve this goal, a probability density $h(\cdot)$ of the MPP can be formulated with respect to a reference point process, which is usually defined as the Poisson point process. We define $h(\cdot)$ by a Gibbs energy $U(\cdot)$ in the form of $h \propto \exp - U(\cdot)$. The Gibbs energy consists of two parts, a data energy $U_D(\cdot)$ and a prior energy $U_P(\cdot)$. The relative influence is weighted by a parameter $\beta \in [0, 1]$.

$$U(\cdot) = \beta \cdot U_D(\cdot) + (1 - \beta) \cdot U_P(\cdot). \quad (1)$$

The conformity of the object configuration with the input data is measured by $U_D(\cdot)$, while interactions between the objects are taken into account by $U_P(\cdot)$. The optimal object configuration $u^* = \{u_1, \dots, u_n\}$ can be determined by maximizing the probability density $h(\cdot)$, i.e. $u^* = \arg \max h(\cdot)$, which is equivalent to minimizing the Gibbs energy $U(\cdot)$, i.e. $u^* = \arg \min U(\cdot)$. The probability density $h(\cdot)$ is usually multimodal and is defined in a configuration space with a variable dimension, because the number of objects can change.

2.2 Reversible Jump Markov Chain Monte Carlo sampling

Markov chains (e.g. Andrieu et al., 2003) can be used to model random state changes of a system. The chains exhibit the Markov property, which implies that in a sequence of states, each sample X_t has a probability distribution that depends only on the previous sample X_{t-1} . Markov Chain Monte Carlo (MCMC) methods (Metropolis et al., 1953; Hastings, 1970) cannot deal with state spaces of different dimensions. However, Reversible Jump MCMC (RJCMC) methods can model scenes with an unknown number of objects. This was first proposed by Green (1995) and is achieved by defining a set of changes (jumps) of the current configuration that are reversible. Reversibility means

that one can always return to a previous state. In each iteration t , the sampler proposes a change of the current object configuration from the predefined set of jumps. For each type of change, there is a density function Q_m (called a kernel). This kernel Q_m leads from an object configuration X_t to a new configuration X_{t+1} according to a probability $Q_m(X_t \rightarrow X_{t+1})$. The new configuration is accepted with a certain acceptance probability α depending on the energy difference of states X_t and X_{t+1}

$$\alpha = \min \left(1, \frac{Q_m(X_{t+1} \rightarrow X_t)}{Q_m(X_t \rightarrow X_{t+1})} \cdot \exp - \left(\frac{U(X_{t+1}) - U(X_t)}{T_t} \right) \right). \quad (2)$$

In (2), the kernel ratio $Q_m(X_{t+1} \rightarrow X_t)/Q_m(X_t \rightarrow X_{t+1})$ describes the ratio of probabilities for changing the configuration from X_{t+1} to X_t and vice versa. The Gibbs energies (1) of the new and current object configuration are represented by $U(X_{t+1})$ and $U(X_t)$, respectively; T_t is the temperature used in simulated annealing at iteration t . To find the optimum of the energy, we combine the RJCMC procedure with simulated annealing (Metropolis et al., 1953; Kirkpatrick et al., 1983). The sequence of temperatures T_t tends towards zero while $t \rightarrow \infty$. A logarithmic cooling schedule guarantees convergence to the global optimum for any initial configuration X_0 . As this leads to high computation times, a faster cooling scheme based on a geometric sequence is typically used instead. It provides an approximate result, which is usually close to the optimum (Van Laarhoven and Aarts, 1987).

2.3 Kernel Density Estimation

In order to estimate the probability density function (pdf) of a random variable, kernel density estimation (e.g. Parzen, 1962) can be used. This nonparametric approach estimates the pdf directly from the data (Scott, 2015). Given a sample (x_1, x_2, \dots, x_n) drawn from a distribution with an unknown density p , an estimate \hat{p} of this density can be calculated via

$$\hat{p}(x) = \frac{1}{nh} \sum_{i=1}^n K \left(\frac{x - x_i}{h} \right). \quad (3)$$

Here, h is the bandwidth parameter and K is a kernel function. Note that this definition of the term *kernel* is not to be confused with the kernels used for RJCMC sampling from Section 2.2. The kernel function $K(k)$ has to be a non-negative function ($K(k) \geq 0$) that integrates to one ($\int K(k) dt = 1$). Equation (3) can be thought of as an estimate of the pdf by averaging the effect of a set of kernel functions centred on each data point. Often, the Gaussian kernel is considered as a good choice for the Kernel function, but other functions (e.g. triangular) can also be used.

3. METHODOLOGY

We use MPPs in combination with RJCMC sampling and simulated annealing to find the optimal configuration of objects in the scene. Bomb craters are modelled as ellipses or circles, respectively (Section 3.1). During sampling, the object configuration changes continuously and a global energy function is minimized. The possible types of change are described in Section 3.2. In the optimization process, we evaluate the object configuration in each iteration. Here, high gradient magnitudes along the objects border are favoured and overlapping objects are penalized. Moreover, we penalize higher standard deviations of the grey values within the object (Section 3.3). Section 3.4 describes how we limit the search space; image pre-processing is described in Section 3.5. On the basis of the final object

configuration, a probability map is created by kernel density estimation which allows to differentiate between potentially contaminated and uncontaminated areas (Section 3.6). Parts of the methodology presented in this paper are based on (Kruse et al., 2018), which provides more information about some topics.

3.1 Object model

In Kruse et al. (2018) the object model used for bomb craters is based on an ellipse. In this paper, an alternative object model is tested. As Zhang et al. (2013), we model the bomb craters as circles. Each circle is described by its position $P(x, y)$ and one mark $M(r)$. The resulting 3-tuple (x, y, r) takes the following values in state space Z

$$Z = P \times M = [0, X] \times [0, Y] \times [r_m, r_M]. \quad (4)$$

Here, (x, y) are the coordinates of the circle centre, X and Y represent the width and height of the image. The minimum and maximum value of the radius is limited by $[r_m, r_M]$. If not otherwise mentioned, the term “object” is used in the following to refer to both object models.

3.2 Changes in the object configuration

In each iteration of the sampling process, the object configuration is changed according to a kernel Q_m with associated proposition probability p_{Q_m} . In total, four types of change (kernels) exist, namely *birth* (Q_B), *death* (Q_D), *translation* (Q_T) and *mark-variation* (Q_M) kernels. The related proposition probabilities are $p_{Q_B}, p_{Q_D}, p_{Q_T}, p_{Q_M}$. On the one hand, an object can be added to the current object configuration by Q_B . The position of a new object is sampled from likely positions for bomb craters detected in the way described in Section 3.4. This procedure provides information about the size of the associated crater, which is used for the initialization of the circles radius or the two semi-axes of the ellipse (the ellipse orientation is drawn randomly from a uniform distribution), respectively. On the other hand, Q_D removes a randomly selected object from the current object configuration. The kernel ratio, described in equation (2), considers the probability of changing the configuration from X_t to X_{t+1} and vice versa. Similar to our previous work, we model the kernel ratio of the birth event by

$$\frac{Q_B(X_{t+1} \rightarrow X_t)}{Q_D(X_t \rightarrow X_{t+1})} = \frac{p_{Q_D}}{p_{Q_B}} \cdot \frac{\lambda}{n}. \quad (5)$$

Here, the Poisson parameter λ describes the expected number of objects while n represents the actual number of objects in the scene. For the death event, the kernel ratio corresponds to the inverse birth rate. For changing the position of an object by Q_T , a randomly chosen object is shifted from its current position by a random (local) displacement vector. The movement is realized in a given interval based on a uniform distribution. The kernel ratio is set to one. Finally, Q_M allows to change the marks of the object. An ellipse of the current configuration is randomly selected, from which the new semi-minor and semi-major axis as well as the orientation are drawn from a uniform distribution within predefined intervals. Analogous changes are applied to a circle and its radius. The kernel ratio is also set to one in both cases.

3.3 Energy function

To evaluate each object configuration, we use the Gibbs energy (1) which describes the consistency of the configuration with our bomb crater model and is minimized during sampling.

3.3.1 Data energy: The data energy $U_D(X_t)$ from (1) checks the consistency of the object configuration with the input data. Bomb craters are characterized by locally darker grey values in comparison to the surrounding area due to the shadow cast by the sun. Its shape is often circular within the bomb craters. Consequently, bomb craters are assumed to have high gradient magnitudes in the transition region from dark to bright and more or less homogenous grey values inside (Fig. 1b). We adopt the data term $U_G(X_t)$ of our previous work (Kruse et al. 2018) and add a new term $U_H(X_t)$. Thus, the data energy is modelled by

$$U_D(X_t) = U_G(X_t) + U_H(X_t). \quad (6)$$

Gradient magnitudes: According to the assumptions made earlier, a newly created or modified object leads to a reduction of the data energy if high gradient magnitudes occur along the edges of the object, i.e. the shape of the circle respectively ellipse fits to the border of the shadow. We determine the gradients along the edge of the object and model the corresponding data term by

$$U_G(X_t) = f_G \cdot \sum_{o_j \in X_t} \left(c - \frac{1}{n_b} \sum_{j=1}^{n_b} \nabla_{IMG_{p_j}}^b \right). \quad (7)$$

In (7), $\nabla_{IMG_{p_j}}^b$ is the component of the grey value gradient at the border pixel p_j in the direction of the normal vector of the object o_j pointing outside. To calculate the sum of the gradients along the border of the object, n_b pixels p_j are used. The edge of the object is approximated by a polygon with a constant number of n_p vertices (set to $n_p = 32$ in our experiments). The term is weighted by a factor f_G and the constant $c \geq 0$ ensures that the energy only decreases if the sum in (7) is larger than c . Without the introduction of c , objects with very small gradient magnitudes at the object border would already reduce the energy, so that the optimal configuration would consist of an extremely large number of (mostly false positive) objects.

Homogeneous grey values: This term requires the grey values inside the object to be homogeneous. Homogeneity is measured by the grey value standard deviation σ within the object. It is assumed that σ is higher for a false positive (FP) object than within a bomb crater (Fig. 1b-d). A modified or newly created object o_j increases the data energy if σ_j is higher than a predefined threshold H_t , which results in

$$U_H(X_t) = f_H \cdot \sum_{o_j \in X_t} \max(0, \sigma_j - H_t). \quad (8)$$

The energy term is weighted by f_H . For the computation of σ_j we consider the grey values of all pixels inside the object o_j . In addition, a certain number of pixels around the border of the object can be excluded from the calculation of σ_j , because the shapes bomb craters may deviate slightly from the geometrical model. This possibility is controlled via a parameter H_e (Fig. 1a) describing the width of the area that is excluded.

3.3.2 Prior energy: Using this energy, certain object configurations can be favoured based on prior knowledge. To avoid the accumulation of objects in regions with low data energy, configurations with overlapping objects are penalized. As in (Kruse et al., 2018), we consider all possible combinations of overlapping object pairs o_i, o_j . The overlapping areas A_{ij} of the objects o_i and o_j as well as the respective relative overlapping areas $A_{ij}/A_i(o_i)$ and $A_{ij}/A_j(o_j)$ are calculated. Here, $A_i(o_i)$ and

$A_j(o_j)$ are the areas of the objects o_i and o_j , respectively. The prior energy with penalization weight f_p becomes

$$U_p(X_t) = f_p \cdot \sum_{o_i, o_j} \max\left(\frac{A_{ij}}{A_i(o_i)}, \frac{A_{ij}}{A_j(o_j)}\right). \quad (9)$$

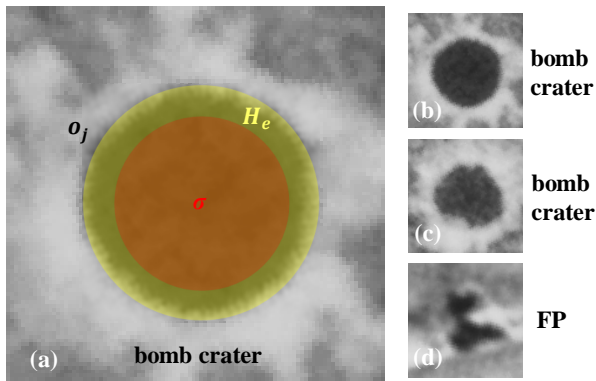


Figure 1: Homogeneity of an object o_j is measured by the grey value standard deviation in the red area; pixels in H_e (yellow) are excluded (a). Bomb craters with a circular (b) and not quite circular shape (c); (d) shows a FP with a high standard deviation within the object.

3.4 Limitation of the search space

To reduce the computational effort for sampling, we restrict the search space for the MPPs in the image. For that purpose, we use the blob detector described in (Mallick, 2015) and implemented in OpenCV. A blob is a group of connected pixels that share common properties (e.g. similar grey values). The aim of the blob detection is to find and mark these pixel regions in the image. Subsequently, in the sampling process these blob locations are used to restrict the search space by only allowing the birth of an object at such a position. In order not to miss any craters, the parameters of the blob detector are selected appropriately (Section 4.1.2). This is important because other image locations are no longer taken into account.

In the first step of blob detection, the image is converted into several binary images by applying different thresholds. Beginning with a minimum threshold B_{T_min} this threshold is increased by B_{T_step} (parameter for the step size) until a maximum threshold B_{T_max} is reached. Subsequently, connected components are extracted from each binary image and their centres are calculated. The centres are grouped according to their coordinates, with neighbouring centres forming a group that corresponds to one blob. In this context, the procedure requires another parameter B_D which ensures that blobs located closer than B_D are merged. The algorithm also provides filter options (circularity, convexity, inertia ratio and size B_N) and allows to detect only dark blobs, bright blobs or both types of blobs.

3.5 Preprocessing

The quality of the images differs considerably due to their age as well as the circumstances during the acquisition. A central aspect is that sometimes areas of the image are underexposed, whereas others are overexposed. To counteract this, a Contrast Limited Adaptive Histogram Equalization (CLAHE; Pizer et al., 1987 – we use the OpenCV implementation) is applied to each image as a pre-processing step. First, the image is divided into small blocks with a size of C_b and each block is histogram equalized. To avoid the amplification of noise, contrast limiting is applied before the

equalization: pixels with a histogram bin above the specified contrast limit C_l are distributed uniformly to other bins. Finally, bilinear interpolation is applied in the borders of the blocks.

In order to make use of the full radiometric resolution in the sampling process, the grey values are normalized locally to the interval $[0, 255]$ in a window centred over the object centre with the size of the object's largest expansion (radius or semi-major axis) and an additional n_p pixels in each direction. We observed empirically that these image processing operations are reasonable, especially regarding the investigations in Section 4.4.

There are objects that have a similar appearance as bomb craters. Among others, shadow casts by houses and trees can lead to false detections. If stereoscopic imagery is available, one way of counteracting these limitations is to integrate height information to find and subsequently exclude such areas from further processing. Similar to our previous work, we instead manually masked out high stationary objects (trees, houses, others) including their shadows. Pixels inside such masked areas are considered neither for the detection nor for the evaluation. We run our experiments (Section 4) on both, the masked and non-masked images to investigate the general influence of the masking procedure.

3.6 Impact map

We use the detected bomb craters to derive a probability for each location that there are duds nearby. Kernel density estimation with the kernel function $K(k) = (1 - |k|)$ is performed to generate the associated probability map from the centres of the detected craters. In this context, the bandwidth h in equation (3) indicates how large the area of influence of a detection is. Using the probability map, the entire scene is classified into potentially contaminated and uncontaminated areas. For that purpose, a threshold p is applied to the probabilities resulting in an impact map. This threshold specifies from which probability within the probability map an area is classified as contaminated.

We are interested in detecting areas that have a very high likelihood of containing a dud so that it makes sense to send a team of experts to that area to probe it. In general, probing is recommended for areas of bomb craters as well as clusters of craters. Thus, the focus of our work is on avoiding false detections, because these cause high unnecessary costs, i.e. areas classified as contaminated should in fact be contaminated.

4. EXPERIMENTS

Our method is evaluated on a total of 22 panchromatic aerial wartime images (Section 4.1.1) for which we empirically set the parameters in our approach (Section 4.1.2). The evaluation of the results is carried out pixel-based (Section 4.1.3). The aims of our experiments are as follows. First, we check whether the circle as an object model leads to similar or even better results compared to the ellipse used in our previous work (Section 4.2). Second, we want to validate the extended data energy by analysing the influence of the newly presented data term on the existing procedure (Section 4.3). Both, in Sections 4.2 and 4.3, the goal is to achieve results with a high quality. In contrast, in Section 4.4, we tune the algorithm with regard to the proposed application scenario, i.e. the results should have a high correctness. For this purpose, the parameter c (equation 7) of the first data term U_G is considered while the parameters of the second data term U_H are kept constant (Section 4.4).

4.1 Test data and test setup

4.1.1 Data: We evaluate our method based on 22 aerial images of scale 1:6000 to 1:10500, acquired by the Allied forces during the Second World War over Lower Saxony, Northern Germany, and scanned with 1200 dpi with 8 bit radiometric resolution, resulting in ground sampling distances (GSD) of approx. 0.13 m to 0.22 m. The images cover areas of about 2 km² to 6 km² each and consist of about 10.000² pixels. Experts of Lower Saxony's explosive ordnance disposal service generated the area-based reference by manual annotation, i.e. for each reference crater we know its position and radius. The degree of impact, i.e. the number of bomb craters in each image, varies between zero and almost 1000. The focus of the investigations and developments is on rural sites. In densely built-up areas, it is not possible to clearly identify craters in the images because they are covered by the rubble of destroyed buildings. Furthermore, the investigated images are representative for certain cases (e.g. different lightning situations and variable image content).

4.1.2 Parameter settings: In our experiments, we set the free parameters to values that were determined empirically. If not specifically indicated, the parameters for all images and tests were set to identical values. Compared to (Kruse et al., 2018), among some other less relevant parameter changes, we chose the same value for c (equation 7) in all images, and except for two images, we did not change the minimal blob size $B_{N,m}$ (Section 3.4). Thus, there are almost no varied parameters, which makes the procedure more relevant for a potential use case.

We select the parameters of the blob detector (Section 3.4) as $B_{T,min} = 10$, $B_{T,max} = 245$, $B_{T,step} = 2$ and $B_D = 5$. As bomb craters or their shadows, respectively, are generally characterized by darker grey values than those in their surroundings, the procedure should only detect dark blobs. The parameters for filtering are set in a loose way, which allows craters to deviate from a circle. Similarly, the selection of B_N in the interval $[B_{N,m}, B_{N,M}]$ makes it possible to detect bomb craters with a different number of pixels and, thus, different sizes. Depending on the GSD of the respective image, we set $B_{N,m}$ and $B_{N,M}$ in a way that blobs in-between diameters of 5 m and 15 m can be detected. For two images with comparably small craters, the lower boundary was reduced to 3 m. Although selecting such loose filter restrictions results in many false detections, experiments have shown that a more restrictive choice excludes the detection of many bomb craters in advance. On the other hand, it is also important that the minimum blob size $B_{N,m}$ is not set too small.

For CLAHE, we set the parameters for the block size to $C_b = 60 \times 60$ pixels and the contrast limit to $C_b = 2$. Here, the selection was performed based on visual inspection in a way that underexposed and overexposed areas, respectively, are removed appropriately. The parameter n_p , which refers to the extended size of the window in which the normalization is performed, is set to $n_p = 20$ pixels to make sure that the complete part of the inside of the crater as well as some surrounding pixels are contained in the window.

We weight the data and prior energy of the marked point processes equally, i.e. β from equation (1) is set to $\beta = 0.5$. Simulated annealing uses a geometric cooling scheme by reducing the temperature T_t from (2) using a factor f_T in the form $T_t = T_0 \cdot f_T^t$. Here, we set the start temperature $T_0 = 100$ and $f_T = 0.9994$. The lower and upper limits of the semi-major (a_m , a_M) and semi-minor (b_m , b_M) axes (ellipse) or radius (circle), respectively (Section 3.1), are derived from the minimum and

maximum blob radius $B_r \in \{B_{r,min}, B_{r,max}\}$ occurring in the image after blob detection. This results in $B_{r,min} = a_m = b_m = r_m$ and $B_{r,max} = a_M = b_M = r_M$. We set the proposition probabilities of the kernels (Section 3.2) to $p_{Q_B} = 0.4$, $p_{Q_D} = 0.4$ and $p_{Q_T} = p_{Q_M} = 0.1$. The probabilities for the translation and mark-variation event are comparatively low, since the objects, in particular the circle, no longer have to be significantly shifted or changed due to the size information provided by the blob detector. In order to avoid manual intervention, λ from equation (5) is set to $\lambda = \#blobs/20$, where $\#blobs$ is the number of blobs. The parameter c of the first data term ($U_D(X_t)$, equation 7) is set to 1100 and the weight of that term is set to $f_G = 1$. The parameters of the second term $U_H(X_t)$ of the data energy in equation 8 are set as $f_H = 5$ and $H_t = 15$. Due to the fact that craters are not always round (e.g. Fig. 1c), we set $H_e = 6$. The weight f_H is set to a relatively low value, because otherwise objects could be generated in homogeneous non-crater areas. Finally, in connection with the prior energy (9), minor overlap of objects is possible with $f_p = 10000$. The initial configuration for the sampling procedure is an empty set of objects.

In connection with the kernel density estimation, we derive an impact map with a radius of 20 m. For this purpose, the bandwidth h from equation (3) is varied based on the image scale and an appropriate threshold p for the probabilities is set in a way that the area around the centre of an object is always classified as contaminated within a radius of 20 m for single detections. As areas of bomb craters and their immediate surroundings are likely to contain duds, the radius of 20 m is set relatively small in order to detect those areas that probably need to be probed by experts.

4.1.3 Evaluation criteria: The pixel-based evaluation of the results is based on the impact map (Section 3.6) generated from the automatically detected bomb craters. The reference centres of the bomb craters are used for the generation of the reference impact map (same parameter setting as for the generation of the impact map from the object centres). The corresponding impact maps from the reference and the automatic detection are compared and each pixel is classified as being either a True Positive (TP), False Negative (FN), False Positive (FP) or True Negative (TN). A TP is a pixel that was correctly classified as contaminated in both, the reference and automatic detection. FN pixels have been classified as uncontaminated by the automatic detection although they are in fact contaminated. FP pixels were falsely classified as contaminated. Finally, a TN pixel was correctly classified as uncontaminated in both cases. The completeness is the percentage of the actually contaminated area found by our method, i.e. $TP / (TP + FN)$. The correctness is the percentage of areas from the automatic detection that lie in areas which are actually contaminated, i.e. $TP / (TP + FP)$. The quality considers both types of errors, FN and FP (Heipke et al., 1997).

4.2 Comparison of two object models

The numerical values of completeness (CP), correctness (CR) and quality (Q) for the comparison of the object models ellipse and circle can be found in Table 1. Here, the quality measures for the total area based on all 22 images, non-masked (TA non-m) and masked (TA m), are shown (more detailed information, e.g. the results for every image, is presented in Table 2, Section 4.3). Our results show that if a circle instead of an ellipse is used as an object model, the quality of the results remains almost the same. The completeness decreases on average by about 1 %, while the correctness increases by about 2 %. Strongly elliptical craters can lead to FNs, whereas FPs that appear similar to elliptical bomb craters in the image are also somewhat reduced. Consequently,

the circle should be preferred as object model, because a model with fewer parameters is considered to be more stable. Thus, in the following experiments the circle is used as object model.

	ellipse			circle		
	CP [%]	CR [%]	Q [%]	CP [%]	CR [%]	Q [%]
TA non-m	56.5	69.5	45.3	56.0	71.2	45.6
TA m	56.2	76.7	48.0	55.3	78.3	48.0

Table 1: Evaluation results for the models ellipse and circle.

4.3 Validation of the expanded data energy

Analogously to Table 1, Table 2 shows in addition to the total areas (non-masked and masked) the quality measures for each of the 22 investigated non-masked images (IMG non-m; details for the masked images are not shown) with the respective number of bomb craters (NC). Note that the abbreviations are the same as the ones in Table 1. The results considering only the first data term (GRAD) are compared to those additionally taking into account the second data term (GRAD + HOM). As mentioned in the previous section, the used object model is a circle.

IMG non-m	NC	GRAD			GRAD + HOM		
		CP [%]	CR [%]	Q [%]	CP [%]	CR [%]	Q [%]
I	0	N/D	0	0	N/D	0	0
II	0	N/D	0	0	N/D	0	0
III	0	N/D	0	0	N/D	0	0
IV	0	N/D	0	0	N/D	0	0
V	9	11	6	4	11	7	4
VI	16	19	18	10	19	20	11
VII	17	23	7	5	17	6	5
VIII	18	31	15	11	30	22	14
IX	24	12	9	5	12	12	6
X	24	19	12	8	19	14	9
XI	33	36	35	22	36	48	26
XII	47	41	55	31	39	62	31
XIII	56	46	36	25	42	41	26
XIV	130	47	45	30	46	51	32
XV	168	35	43	24	35	48	25
XVI	236	59	58	41	56	63	42
XVII	331	36	70	31	35	71	30
XVIII	373	29	71	26	29	75	27
XIX	396	75	93	71	75	96	73
XX	475	66	64	48	59	74	49
XXI	554	76	91	71	75	93	71
XXII	925	78	75	62	72	85	64
TA non-m		59	65	44	56	71	46
TA m		58	73	47	55	78	48

Table 2: Evaluation results for the first data term (U_G) and the combination of both, the first and second data term ($U_G + U_H$); (N/D: not defined).

The results from Table 2 show that an average completeness and correctness of 56 % and 71 %, respectively, can be achieved by considering both data terms (TA non-m, columns 6-8). Compared to the results only based on the first data term (TA non-m, columns 3-5) with a completeness and correctness of 59 % and 65 %, respectively, the quality increases from 44 % to 46 %. In many cases, the correctness increases by several percent while the completeness remains (almost) the same (e.g. VIII, XI, XIV, XV and XVIII). On the other hand, due to the additional data term, (almost) the same number of TPs as FPs are apparently

removed from the object configuration during sampling, which has a larger negative effect on the completeness (e.g. XVI, XX and XXII). This elimination of bomb craters occurs because they do not always have homogeneous dark grey values inside, are not exactly convex, or have an elliptic shape, which can also lead to higher standard deviations. Thus, the assumption that the standard deviation of the grey values for a bomb crater is smaller than for a FP is only partly correct. When using a larger weight and a more restrictive parameter setting (e.g. $f_H = 40$, $H_T = 10$, $H_e = 6$; equation 8) the completeness decreases to 43 % while the correctness increases to 84 % (TA non-m).

4.4 Focus on correctness

In general, the quality measures shown in the previous section are too low to integrate the results into the workflow of the explosive ordnance disposal service (Table 2). With an average correctness of 65 % (using only the first data term) or 71 % (using both data terms) for the non-masked images (Table 2, underlined), too many areas would be unnecessarily probed, resulting in enormous costs. That is why we vary the parameters of the algorithm in a way that the results have a higher correctness at the expense of completeness. To achieve this aim, it is for example possible to vary the parameter c of the data energy U_G . Increasing c will result in more and more objects with smaller gradients at the object border being removed from the object configuration. Similarly, the parameters of the second data term U_H could be varied, but this term has more free parameters. This is why we increase the parameter c starting with $c = 1100$ and keep the parameters of the second data term U_H constant as described in Section 4.1.2 (Fig. 2). Again, the completeness and correctness are based on the total area of all 22 images.

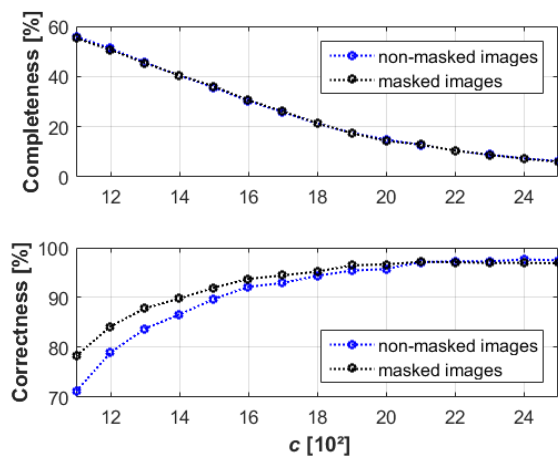


Figure 2: Completeness and correctness as a function of parameter c for the total area of all 22 images.

Figure 2 shows the dependence of completeness and correctness on parameter c of the first data term. It can be seen that the completeness decreases more or less linearly until approx. $c = 1900$, while the correctness is similar to a root function, i.e. initially it increases comparatively more strongly and from approx. $c = 1300$ the loss in completeness is higher than the gain in correctness. However, having in mind the proposed scenario, e.g. for $c = 1400$ a correctness of almost 90 % (non-masked) or 90 % (masked) can be achieved with a remaining completeness of 40 %. From approx. $c = 2100$ onwards, the correctness lies at about 97 %, whereas the completeness still decreases. Furthermore, it should be noted that the curves for completeness (non-masked and masked) overlap (as was to be expected) and the curves for correctness converge with increasing c . The latter

is of interest if it is not possible to automatically derive height information. Finally, Figure 3a shows an example for the resultant optimal object configuration after the sampling procedure for a part of image XXI. Here, there are five false detections (Fig. 3a, cyan circles), and comparably few bomb craters are not detected (Fig. 3a, red arrows). This observation can also be seen in Figure 3b in which the two impact maps generated from the reference centres and the centres of the automatic detection (Fig. 3a) are superimposed. Here, c was set to $c = 1100$. Mainly one larger area that consists of the three undetected craters was falsely classified as uncontaminated (Fig. 3b, red). By increasing c from $c = 1100$ to $c = 1500$ (Fig. 3c) to $c = 1900$ (Fig. 3d), the areas that were wrongly classified as contaminated almost vanish (Fig. 3b-d, pale blue). However, the number of pixels falsely not detected as contaminated increases as well (Fig. 3b-d, red). Consequently, for the proposed scenario, the procedure provides very good results for $c = 1900$ in this example: almost all areas classified as contaminated (Fig. 3d, dark green / pale blue) actually have to be probed.

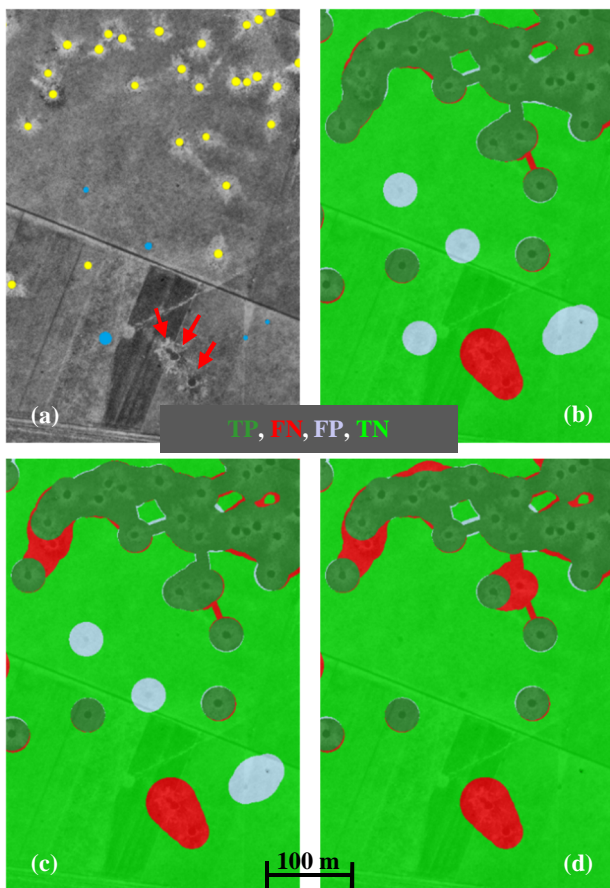


Figure 3: Subset of image XXI with the final object configuration for the MPP of circles for $c = 1100$ (a). Correct detections are shown in yellow, false detections in cyan and missing detections are tagged by red arrows. (b) superimposition of the corresponding impact map and its evaluation with TP-areas in dark green, FN-areas in red, FP-areas in pale blue and TN-areas in lime green. (c) and (d) show the impact maps and the evaluation resulting from $c = 1500$ and $c = 1900$, respectively.

5. CONCLUSION AND OUTLOOK

In this work, we present a stochastic approach based on marked point processes for the automatic detection of bomb craters in

aerial wartime images. The detections are used to generate an impact map that provides a quick overview of contaminated areas for detecting areas that have a very high likelihood of containing a dud so that it makes sense to have them probed by a team of experts. The approach was evaluated on a total of 22 panchromatic images. Experts of Lower Saxony's explosive ordnance disposal service generated the area-based reference by manual annotation. We could show that using a circle instead of an ellipse as object model does not affect the results significantly; the quality stays almost the same. The analysis of the additional data term revealed that the assumption that the standard deviation of the grey values for a bomb crater is smaller than for an FP is often fulfilled. In some cases, it is possible to reduce the number of FPs while maintaining the same level of completeness; in others, TPs are also eliminated. The quality could be slightly increased from 44 % to 46 % by the second data term. In general, the results with a correctness of approx. 65 % (only considering the first data term) or 71 % (considering both data terms) for the non-masked images are not good enough for an integration into the workflow of the explosive ordnance disposal service, because too many areas would have to be probed unnecessarily, resulting in enormous costs. In this context, our final set of experiments shows that the correctness can be increased at the expense of completeness by varying only one parameter of the data energy. Thus, the procedure is more attractive for the proposed use case. On the basis of the investigated images, e.g. a correctness of about 90 % with a remaining completeness of approx. 40 % can be achieved. Furthermore, additional height information would increase the correctness, especially for lower values of the parameter mentioned before.

A problem in connection with false detections arises from objects that appear like bomb craters in the image (e.g. shadows of houses or trees, image errors, others). The former could be tackled by height information (we simulated it). Thus, the integration of height will be considered in the future. FPs stemming from image errors, shadows of non-stationary objects with a certain height or other objects are often partly different from typical craters, especially with regard to their shape. Thus, future experiments, based on the idea behind the second data term, will include investigations regarding the shape of a detection. In this context, we do not consider grey values but determine the deviations from a given shape, e.g. a circle. To achieve this, first, we will determine the contour of a detection and then find a shape which best represents the contour. Deviations should indicate to what extent the detection is actually of the given shape. This could be integrated into another data term in the way that larger deviations increase the energy. Another idea to increase the correctness of the results is to classify an area as contaminated only after several detections (and not already for individual ones). For this purpose, the threshold p (Section 3.6) applied to the probability map created with kernel density estimation (Section 2.3) could be adapted. This, obviously, would reduce the completeness as well (except in strongly bombed areas). Finally, we will consider information from multiple overlapping images for the generation of the impact maps.

ACKNOWLEDGEMENTS

The authors would like to thank the State Office for Geoinformation and Surveying of Lower Saxony and its explosive ordnance disposal service as a department of the regional directorate Hamelin-Hanover for providing the data and the financial support to this project.

REFERENCES

- Andrieu, C., Freitas, N. de, Doucet, A., Jordan, M.I., 2003. An introduction to MCMC for machine learning. *Machine Learning*, 50 (1-2), pp. 5–43.
- Bandeira, L., Ding, W., Stepinski, T.F., 2012. Detection of sub-kilometer craters in high resolution planetary images using shape and texture features. *Advances in Space Research*, 49 (1): 64–74.
- Benedek, C., 2017. An Embedded Marked Point Process Framework for Three-Level Object Population Analysis. *IEEE Transactions on Image Processing: a publication of the IEEE Signal Processing Society*, 26 (9), pp. 4430–4445.
- Bishop, C. 2006. *Pattern Recognition and Machine Learning* (1st edition). Springer, New York, USA.
- Börcs, A., Benedek, C., 2015. Extraction of vehicle groups in airborne LiDAR point clouds with two-level point processes. *IEEE Transactions on Geoscience and Remote Sensing*, 53 (3), pp. 1475–1489.
- Brédif, M., Tournaire, O., Vallet, B., Champion, N., 2013. Extracting polygonal building footprints from digital surface models: A fully-automatic global optimization framework. *ISPRS Journal of Photogrammetry and Remote Sensing*, 77, pp. 57–65.
- Brenner, S., Zambanini, S., Sablatnig, R., 2018. Detection of Bomb Craters in WWII Aerial Images. Proceedings of the OAGM Workshop, pp. 94–97.
- Cedilnik, N., Debreuve, E., De Graeve, F., Besse, F., Descombes, X., 2018. SPADE: Small Particle Detection Method Using A Dictionary Of Shapes Within The Marked Point Process Framework. *IEEE International Symposium on Biomedical Imaging (ISBI)*, Washington, DC, United States.
- Chai, D., Förstner, W., Yang, M.Y., 2012. Combine Markov random fields and marked point processes to extract building from remotely sensed images. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, I-3, pp. 365–370.
- Cohen, J.P., Lo, H.Z., Lu, T., Ding, W., 2016. Crater Detection via Convolutional Neural Networks. <http://arxiv.org/abs/1601.00978> (12 March 2019).
- Daley, D.J., Vere-Jones, D., 2003. *An Introduction to the Theory of Point Processes: Volume 1: Elementary Theory and Methods* (2nd edition). Springer, New York.
- Descamps, S., Descombes, X., Béchet, A., Zerubia, J., 2008. Automatic flamingo detection using a multiple birth and death process. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1113–1116.
- Descombes, X., Zerubia, J., 2002. Marked point process in image analysis. *IEEE Signal Processing Magazine*, 19 (5), pp. 77–84.
- Descombes, X., 2013. *Marked Point Processes for Object Detection*. In: *Stochastic Geometry for Image Analysis*. John Wiley & Sons, Inc, Hoboken, NJ, USA, pp. 11–27.
- Descombes, X., 2017. Multiple objects detection in biological images using a marked point process framework. *Methods*, 115, pp. 2–8.
- Dubosclard, P., Larnier, S., Konik, H., Herbulot, A., Devy, M., 2014. Automatic method for visual grading of seed food products. *International Conference on Image Analysis and Recognition (ICIAR)*, 1, pp. 485–495.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6 (6), pp. 721–741.
- Green, P.J., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82 (4), pp. 711–732.
- Hastings, W.K., 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57 (1), pp. 97–109.
- Heipke, C., Mayer, H., Wiedemann, C., Jamet, O., 1997. Evaluation of automatic road extraction. *International Archives of Photogrammetry and Remote Sensing*, XXXII/3–4W2, pp. 151–160.
- Jensen, L., Drauschke, M., Förstner, W., 2010. Automatische Detektion von Bombentrümmern in digitalisierten Luftbildern des 2. Weltkriegs. In: *DGPF Tagungsband 19/2010*, pp. 228–235.
- Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. *Science*, 220 (4598), pp. 671–680.
- Kowal, M., Korbicz, J., 2018. Marked Point Process for Nuclei Detection in Breast Cancer Microscopic Images. In: Augustyniak, P., Maniewski, R., Tadeusiewicz, R., *Recent Developments and Achievements in Biocybernetics and Biomedical Engineering*. PCBBE 2017. Advances in Intelligent Systems and Computing, vol 647. Springer, Cham, pp. 230–241.
- Kruse, C., Rottensteiner, F., Hoberg, T., Ziemis, M., Rebke, J., Heipke, C., 2018. Generating impact maps from automatically detected bomb craters in aerial wartime images using marked point processes. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-3, pp. 127–134.
- Kumar, S., Hebert, M., 2006. Discriminative random fields. *International Journal of Computer Vision*, 68 (2), pp. 179–201.
- Lafarge, F., Gimel'farb, G., Descombes, X., 2010. Geometric feature extraction by a multimarked point process. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32 (9), pp. 1597–1609.
- Mallick, S., 2015. Blob Detection Using OpenCV (Python, C++). <https://www.learnopencv.com/blob-detection-using-opencv-python-c/> (27 March 2019).
- Meng, D., Yunfeng, C., Qingxian, W., 2009. Method of Passive Image Based Crater Autonomous Detection. *Chinese Journal of Aeronautics*, 22 (3), pp. 301–306.
- Merler, S., Furlanello, C., Jurman, G., 2005. *Machine Learning on Historic Air Photographs for Mapping Risk of Unexploded Bombs*. In: Roli, F., Vitulano, S. (Eds.) *Image Analysis and*

- Processing – ICIAP 2005. Lecture Notes in Computer Science (LNCS)*, Springer, Berlin, Heidelberg, 3617, pp. 735–742.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21 (6), pp. 1087–1092.
- Parzen, E., 1962. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33 (3), pp. 1065–1076.
- Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., Romeny, B.t.H., Zimmerman, J.B., Zuiderveld, K., 1987. Adaptive Histogram Equalization and its Variations. *Computer Vision, Graphics, and Image Processing*, 39 (3), pp. 355–368.
- Ruan, D., 2010. Kernel density estimation-based real-time prediction for respiratory motion. *Physics in Medicine and Biology*, 55 (5), pp. 1311–1326.
- Schmidt, A., Kruse, C., Rottensteiner, F., Sörgel, U., Heipke, C., 2016. Network detection in raster data using marked point processes. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLI-B3, pp. 701–708.
- Schmidt, A., Lafarge, F., Brenner, C., Rottensteiner, F., Heipke, C., 2017. Forest point processes for the automatic extraction of networks in raster data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 126, pp. 38–55.
- Scott, D.W., 2015. *Kernel Density Estimators*. In: *Multivariate Density Estimation*. John Wiley & Sons, Inc, pp. 137–216.
- Solarna, D., Moser, G., Le Moigne, J., Serpico, S.B., 2017. Planetary Crater Detection and Registration Using Marked Point Processes, Multiple Birth and Death Algorithms, and Region-Based Analysis. *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 2337–2340.
- Tournaire, O., Brédif, M., Boldo, D., Durupt, M., 2010. An efficient stochastic approach for building footprint extraction from digital elevation models. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65 (4), pp. 317–327.
- Troglio, G., Benediktsson, J., Moser, G., Serpico, S., 2010. Crater Detection Based on Marked Point Processes. *International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 1378–1381.
- Urbach, E.R., Stepinski, T.F., 2009. Automatic detection of sub-km craters in high resolution planetary images. *Planetary and Space Science*, 57 (7), pp. 880–887.
- Van Laarhoven, P.J., Aarts, E.H., 1987. *Simulated Annealing: Theory and Applications*. Springer Netherlands, Dordrecht.
- Wenzel, S., Förstner, W., 2016. Facade Interpretation Using a Marked Point Process. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, III-3, pp. 363–370.
- Wetzler, P.G., Honda, R., Enke, B., Merline, W.J., Chapman, C.R., Burl, M.C., 2005. Learning to Detect Small Impact Craters. *IEEE Workshops on Applications of Computer Vision (WACV/MOTION)*, pp. 178–184.
- Zhang, J., Sohn, G., Brédif, M., 2013. Single tree detection from airborne laser scanning data using a marked point process based method. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-3/W1, pp. 41–46.