



Improving mine recognition through processing and Dempster–Shafer fusion of ground-penetrating radar data

Nada Milisavljević^{a,b}, Isabelle Bloch^{b,*}, Sebastiaan van den Broek^c, Marc Acheroy^a

^aSignal and Image Centre-Royal Military Academy, Av. de la Renaissance 30, 1000, Brussels, Belgium

^bEcole Nationale Supérieure des Télécommunications, Département TSI-CNRS URA 820, 46 rue Barrault, 75013, Paris, France

^cTNO Physics and Electronics Laboratory, P.O.Box 96864, 2509 JG, The Hague, The Netherlands

Received 18 December 2001; accepted 1 August 2002

Abstract

A method for modeling and combination of measures extracted from a ground-penetrating radar (GPR) in terms of belief functions within the Dempster–Shafer framework is presented and illustrated on a real GPR data set. A starting point in the analysis is a preprocessed C-scan of a sand-lane containing some mines and false alarms. In order to improve the selection of regions of interest on such a preprocessed C-scan, a method for detecting suspected areas is developed, based on region analysis around the local maxima. Once the regions are selected, a detailed analysis of the chosen measures is performed for each of them. Two sets of measures are extracted and modeled in terms of belief functions. Finally, for every suspected region, masses assigned by each of the measures are combined, leading to a first guess on whether there is a mine or a non-dangerous object in the region. The region selection method improves detection, while the combination method results in significant improvements, especially in eliminating most of the false alarms.

© 2002 Pattern Recognition Society. Published by Elsevier Science Ltd. All rights reserved.

Keywords: Humanitarian mine detection; Ground-penetrating radar; Dempster–Shafer framework; Mass assignment; Randomized Hough transform for hyperbola detection

1. Introduction

Despite decades of great efforts of research centers all over the world, the humanitarian mine detection problem is still unsolved, mainly due to the necessarily high detection rate that is requested, as well as a large variety of types of mines and of scenarios where mines can be found. A conventional metal detector (MD) is the oldest mine detection sensor, and in reality, it is still the one that is most often used by the deminers. Unfortunately, there are situations where MD cannot be used, due to either soil type (ferrous soils), metallic

debris, that often remain on old battlefields, or, nowadays, plastic or low-metal content mines, for which MD is practically useless. In these cases, other sensors are preferable, and one most often reaches for a GPR. Namely, GPR detects any object below the soil surface if it differs from the surrounding medium [1,2] in: the conductivity (metallic targets), the permittivity or the dielectric constant (plastic and non-conducting targets), or the permeability (ferrous metals).

Earth materials are mostly non-magnetic, and the change in conductivity mainly affects absorption of the GPR signal by the medium, so it is usually the contrast in the permittivity that leads to a reflection of the electromagnetic waves radiated by the transmit antenna of the GPR [3], and consequently, the detection of backscattered echoes by its receiving antenna. GPR also has two main drawbacks [4]: it cannot see through a water table if the water is highly conductive, and mineralogic clays at frequencies below

* Corresponding author. Tel.: +33-1-45-817585; fax: +33-1-45-813794.

E-mail addresses: nada@elec.rma.ac.be (N. Milisavljević), isabelle.bloch@enst.fr (I. Bloch), s.p.vandenbroek@fel.tno.nl (S. van den Broek), marc.acheroy@elec.rma.ac.be (M. Acheroy).

400 MHz drastically decrease its performance. Otherwise, this sensor is very promising from the mine detection point of view and it attracts a lot of attention within the demining community [3,5–7]. However, a lot of research is also devoted to GPR because:

- the data it provides are often very difficult to interpret,
- various types of clutter can significantly increase false alarms,
- different GPR technologies exist [5,8], and
- the acquired data are strongly scenario dependent.

As a consequence, different data sets need different processing and interpretation methods [9,10].

In this paper, a method for analyzing GPR data is described, as a part of a work done within the Belgian HUDEM project [11]. In Section 2 GPR data presentation in the form of A-, B- and C-scans is briefly described, and preprocessing based on time-varying gain and background removal [5] applied on A-scans is presented. A way for projecting the energy of each preprocessed A-scan that results in one C-scan is introduced. Taking into account some characteristics of the GPR data in general as well as the way that the preprocessing is done, a method for selecting suspected regions is proposed in Section 3 based on the analysis of regions around the local maxima. Once the possibly dangerous regions are selected, a detailed analysis of each of the regions is performed, in order to extract measures that can give information about the true nature of the alarm. Section 4 discusses measures extracted from the preprocessed A- and C-scan data. Measures can also be extracted from B-scans, and more precisely from hyperbolae detected by the randomized Hough transform (RHT) [12] (Section 5). After that, as shown in Section 6 each extracted piece of information is modeled in terms of belief functions and modeled measures are then combined within the Dempster–Shafer (DS) framework [13,14]. The non-probabilistic interpretation [15] of the DS method is chosen, on the one hand, to compensate for the fact that the data are not numerous enough for a reliable statistical learning, and on the other hand, to be able to easily include and model existing partial knowledge about the mines and the GPR sensor itself. A simple way of making decisions or guesses about the true identity of each region is also proposed. Finally, results obtained by applying the proposed method on GPR data acquired in a sand lane at the TNO test facilities [16] within the Dutch HOM-2000 project are given and discussed in Section 7.

2. A-scan preprocessing and resulting C-scan

2.1. Types of GPR data presentations

A common presentation of the signals obtained by GPR is in the form of scans: A-, B- and C-scans (adopted from

acoustic terminology), and data processing can be applied to any of them. A single amplitude-time waveform, with the GPR antennas at a given fixed position, is referred to as an A-scan. A B-scan presents an ensemble of A-scans gathered along one axis, or, in other words, it is a two-dimensional (2D) image representing a vertical slice in the ground. Note that reflections from a point scatterer located below the surface are present in a broad region of a B-scan, due to the poor directivity of the transmitting and the receiving antenna. Finally, a C-scan is a 3D data set resulting from collecting multiple parallel B-scans, hence recording the data over a regular grid in the soil surface plane. Usually, a C-scan is represented as a horizontal slice of this 3D data set by plotting its amplitudes at a given time.

2.2. A-scan preprocessing methods

In order to use the information on energy contained in an A-scan, the weakening of the signal with the depth should be compensated. Another problem that should be overcome is the strong reflection at the air/ground interface, that, besides strongly biasing the energy contained in each A-scan, often hides reflections from objects buried just below or placed just above the surface.

2.2.1. Time-varying gain

While propagating from the transmitter towards a buried object and being scattered back to the receiver, the electromagnetic waves of the GPR are subject to some losses [3,5]. In particular, the deeper the object is buried, the higher the losses introduced by the soil are. In order to compensate for these attenuations in function of R or, more directly, of time t , a TVG is introduced, by which a fixed gain of X dB/s (or /m) is added to the raw signal s , so the amplified signal in the time-domain is

$$s_{TVG}(t) = s(t)10^{Xt/20}. \quad (1)$$

The optimal value of TVG strongly depends on the type of GPR used, on the type and characteristics of the soil, on the level of moisture and on the depth range of interest, so it is usually chosen to meet the operational requirements. Its values can typically vary from 0.1 dB/ns up to 100 dB/ns.

2.2.2. Background removal by finding average values within sliding windows

We discuss here a background removal method which aims at reducing clutter¹ and eliminating strong air/ground interface reflections. Within regions where the soil surface is not rough, where the electromagnetic properties of the soil are unchanged and where the antenna distance from the ground is kept constant, it can be expected that the position at which these strong reflections occur remains constant. In

¹ That is why this method is often referred to as clutter reduction [5].

addition, in such regions it can be assumed that some background disturbances affect the neighboring A-scans approximately equally. That is often the case for smaller regions. A simple way for removing the background consists then in choosing a window within which for each sample number, the mean of neighboring A-scans is found and subtracted from the value of a central A-scan at that sample number. Usually, this sliding window within which the mean is calculated is in the direction of scanning

$$s_{new}(x, y, t) = s_{TVG}(x, y, t) - \frac{1}{2n+1} \sum_{i=-n}^n s_{TVG}(x, y-i, t), \quad (2)$$

where s_{TVG} is a signal after the TVG and before background removal, s_{new} is the same signal after background removal, x is the cross-track coordinate, y is the scanning direction coordinate, and n is the half-width of the sliding window. The choice of n should be a compromise between the expected size of the objects, the distance between two neighboring objects, and the possible change of the height of the antennas. If the range of expected sizes of objects and their distances is not known, if surfaces are rough or if the height of the antennas is not constant, this type of background removal is not really useful. An ideal solution would be to have a “known-to-be-empty” region for the background estimation, for which the GPR distance from the soil surface is somehow (if possible) kept the same as for the measuring region, but this request is rarely met in reality. Finally, if neither of the above requirements for background estimation by finding the mean of several neighboring A-scans is met, more complicated ways for background removal would have to be investigated.

2.3. C-scan containing energy projections of preprocessed A-scans

The advantage of GPR in providing 3D information can be a drawback too, in terms of data quantity. In order to reduce the amount of GPR data, we choose to project the information contained in A-scans in one plane. Another important reason for creating a unique C-scan from 3D data is in using it later both to select possibly dangerous regions as well as to fuse it with other sensors that give 2D images of the regions, such as an imaging metal detector or an infrared camera [17].

Several possibilities for creating such C-scans as projections of A-scans exist, of which we find as the most useful one to sum square values of each A-scan, i.e. to represent it by its real energy. Namely, once TVG and background removal are applied to the data, the signal attenuation effects of the depth and the strength of the air/ground interface reflection are suppressed to some level. It allows to project the energy contained in one A-scan in one point, as well as to put together all such points in one plane and by that indirectly induce comparison between different A-scans, belonging either to objects at various depths or simply to the

background. In other words, each A-scan $s_{new}(x, y, j\Delta T)$, where j is the sample number, and ΔT is the sampling time, is represented by its energy $E_s(x, y)$:

$$E_s(x, y) = \sum_{j=1}^N s_{new}^2(x, y, j\Delta T). \quad (3)$$

Since the peak energy is proportional, amongst other, to the volume of the object seen from that point, the resulting C-scan is a good start for selection of regions possibly containing (dangerous) objects.

3. Region selection by local maxima analysis

A simple way for region selection is to threshold projected A-scans. This gives good results, even without TVG [18]. Still, a potential general problem introduced by this method is that deeper buried objects have weaker signals and, accordingly, weaker energies. Therefore, it can happen that they are not detected if the threshold is not low enough, while a threshold that is too low leads to an increase in the number of false alarms. Even if TVG is applied, such a simple thresholding can rarely work well enough in case of humanitarian mine detection, which demands the highest possible detection rates. As shown in Section 2, the problem of correctly choosing the appropriate TVG is subtle, and strongly depends on a variety of factors. Consequently, it is practically impossible to be sure the chosen TVG is the right one.

In order to preserve weaker signals, possibly belonging to deeper buried objects, without causing a strong increase in the number of false alarms, we propose a simple method for the region selection. The idea is to find local maxima and analyze their neighborhood by grouping together all the points within some window around each of them, the value of which is close to the value of that local maximum. As a result, we get a “blob” within a window around every local maximum. Note that a blob belonging to one local maximum can actually be a group of blobs, and not just one connected blob, since within a window around a local maximum there could be some regions with higher and some with lower energies than the chosen threshold (on percentage of the maximum decrease of a local maximum value) allows.

This means the local maxima method needs three pieces of information as input. The first one is the size of the window around the local maxima within which the points are analyzed. Again (see Section 2), it must be larger than the size of the largest expectable object convolved with the GPR antenna opening and smaller than the minimum expectable distance between two objects in a lane. If this interval does not exist, mistakes can hardly be avoided—either an object that is too large could be detected twice, or a few closely placed objects could be detected as a single object. A way to deal with this problem is shown in Section 7.7. A second input the local maxima method asks for is the minimum value of local maxima that should still be detected

(thr_1). Obviously, this value affects the number of detected regions. Finally, the minimum percentage of the local maximum value so that neighboring pixels are grouped with it (thr_2) has to be chosen too. The choice of this value should not be critical, since it should influence sizes of all obtained blobs similarly.

4. Choice of A- and C-scans measure and their extraction

The selection of regions that possibly contain mines is usually the end of the detection process, resulting in a lot of false alarms in reality. Therefore, we are looking for some useful and reliable measures that can be relatively easily extracted from GPR data, and that should give more information regarding the true identity of objects within each of these regions.

The acquisition step in the cross-track direction is often around the size of a typical AP mine. Consequently, depending on its position in comparison with the gathered B-scans, the same mine can appear on one, two or, in a limit case,² even three successive B-scans. In such cases, both size and shape of the 2D projection are useless. In order to overcome this problem and still extract some information from the pre-processed C-scan, two assumptions have to be made. The first one is that the objects are not prolonged in one direction, so that the dimension in the scanning direction can give an idea about the object size. Taking into account that GPR antennas are usually close to the ground during the acquisition and that analyzed objects are not too deeply buried, it can be also assumed that the antenna opening does not change significantly with the actual object depth (so-called near-field assumption), meaning that the dimension of a region along the scanning direction does not depend on burial depth. Under these two assumptions, a first measure is chosen: the width of each selected region in the scanning direction, y_{size} .

Since the analyzed C-scan is obtained from energy projected A-scans preprocessed by means of TVG and background removal, if the difference in material and shapes is ignored, it can be said that the energy of a local maximum is proportional to the volume of the corresponding object. This is the second chosen measure, E .

The reflection of GPR signal on the interface causes peaks in A-scans, the strength of which depends mainly on the type of material of the object. Once the background is removed, the first peak should correspond to the position of the top surface of the object, i.e. to its depth. That is the third chosen measure, expressed in number of time samples and obtained as the sample number of the first maximum in s_{new} . We convert the depth n_1 , expressed in numbers of samples, to d_1 in centimeters through proportionalities, by

locating an object with a known depth in a calibration area, and assuming that electromagnetic properties of the sand remain constant within the lane, as well as that the soil is homogeneous and isotropic [19]. For such an object buried at the depth of d_{ref} , the first maximum in its A-scan appears at the sample number n_{ref} . Analyzing the raw A-scans, it can be seen that the air/ground interface appears around the sample number n_{surf} . This means that for buried objects ($n_1 > n_{surf}$), one can estimate the depth in centimeters by

$$d_1 = -\frac{d_{ref}(n_1 - n_{surf})}{n_{ref} - n_{surf}}. \quad (4)$$

The soil surface is taken as the reference point for the depth, and the negative sign is chosen to indicate depths in the ground. In case that an object is surface-laid ($n_1 < n_{surf}$), the height of its top above the surface could be estimated taking antennas as the reference object, and knowing its height. Unfortunately, it is not possible to determine which sample number the antenna position corresponds to, since in practice, sample number 0 is chosen in an arbitrary way. This makes distance estimations above the surface unreliable. Still, the real power of GPR is in its subsurface detection abilities, and for surface-laid objects various other sensors could be used in combination with this sensor. Therefore, we choose to assign all positive depths to zero.

5. B-scan hyperbola detection and chosen measures

5.1. Hyperbola detection on B-scans by RHT

Another way for getting useful measures from GPR data is to analyze B-scans obtained from A-scans preprocessed by the TVG and background removal methods. Namely, while A-scans can give local information at the position of a local maximum only, B-scans can provide more global information about its neighborhood. An interesting way to gain information about the size of an object, its 3D position and the propagation velocity of electromagnetic waves above the object is by analyzing characteristic hyperbolic shapes on B-scans [12]. These shapes result from the poor directivity of GPR antennas, due to which reflections of a small object, approximated by a point scatterer, are smeared out [3] over a broad region in B-scans.

The geometry of GPR data acquisition can be presented as given in Fig. 1, for a given displacement of the antennas from the starting position of a B-scan line. Y_t denotes the displacement for the transmitting (T) and Y_r for the receiving (R) antenna, where their difference, $\Delta Y = Y_r - Y_t$ is a characteristic of the GPR. The lateral distance of the observed object from the same starting point is Y_0 . Height of the antennas above the soil surface is H , and it is assumed to be constant. It is assumed that the size of the object, d , is comparable to the wavelength of the GPR signal, λ , so

² The finite antenna opening should be taken into account as well, i.e. the fact that antennas typically have a poor directivity, so objects appear larger than they really are.




Fig. 1. The geometry of GPR data acquisition.

that it can be approximated by a point scatterer.³ Its burial depth is equal to D . The path of electromagnetic waves traveling from T to the object consists of two parts, one through the air (w_1) and another through the soil (w_2). At the interface between the two media waves are refracted so that a change of path angle occurs, from θ_1 to θ_2 , with respect to the soil surface. Similarly, the path of waves reflected from the object and traveling back to the GPR receiver consists of two parts, w_3 through the soil and w_4 through the air, with corresponding angles θ_3 and θ_4 . A detailed analysis of this situation and corresponding complex calculations can be found in Ref. [17]. In the following, we introduce some simplifications of the geometry of GPR data acquisition, leading to simplified calculations that can be found in, e.g., Ref. [12]. There are several arguments in favor of these simplifications:

- the main interest in GPR is for subsurface imaging,
- it is easy to determine the air/ground interface on A- or B-scans, due to the strong reflection at such interfaces,
- GPR antennas often operate very close to the ground,
- the distance between T and R is often negligible in comparison with other distances.

Accordingly, at a first approximation, it can be said that $H = 0$ and $\Delta Y = 0$, and, consequently, $Y_r = Y_t = Y_a$, where Y_a is the distance of the central point between T and R from the starting position in a B-scan line. It further means that $w_1 = w_4 = 0$ and $w_2 = w_3 = w$, with w being the (one-way) wave-path between the antennas and the object:

$$w = \sqrt{D^2 + (Y_a - Y_0)^2}. \quad (5)$$

Finally, the round-trip travel-time or time-of-flight (TOF) can be found as

$$TOF = 2 \frac{w}{v} = 2 \frac{\sqrt{D^2 + (Y_a - Y_0)^2}}{v} \quad (6)$$

with v being the propagation velocity of the electromagnetic waves through the soil (assuming that it is a constant value). TOF is usually expressed in discrete values, as a number of time samples j taken every ΔT :

$$TOF = j\Delta T. \quad (7)$$

The data is stored in such a way that the y -axis is also discretized:

$$Y_a = i_a\Delta Y, \quad (8)$$

$$Y_0 = i_0\Delta Y. \quad (9)$$

Here, ΔY presents the acquisition step in the y -direction, while i is the number of samples, i_a being the sample number in y -direction corresponding to one A-scan collecting position of antennas in a B-scan line, and i_0 being the discretized lateral position (y -coordinate) of the object. Substituting Eqs. (7), (8) and (9) in Eq. (6), we come to the following:

$$j^2 = A + B(i_a - C)^2 \quad (10)$$

with

$$A = \left(\frac{2D}{v\Delta T} \right)^2, \quad (11)$$

$$B = \left(\frac{2\Delta Y}{v\Delta T} \right)^2, \quad (12)$$

$$C = i_0. \quad (13)$$

Eq. (10) is the parametric equation of a hyperbola, showing that indeed, due to a poor directivity of T and R, objects

³ This assumption is valid for AP mines and standard GPR frequencies.

appear as hyperbolae in B-scans: moving the antennas during the acquisition of a B-scan, their position i_a changes, and one small object leaves a hyperbola in the acquired image.

As a next step, a way to detect or extract hyperbolae from B-scans has to be found. For that, the RHT [20,21] for hyperbola detection [12] can be of a great help. In the following, we only briefly describe this idea which is discussed in detail in Ref. [22], and which is based on modifying a previously developed method for ellipse detection by RHT [23,24]. Firstly, we randomly choose three foreground pixels in an image containing edges of a thresholded B-scan image (after TVG and background removal). Then, we find the three parameters (A, B, C) of the hyperbola that contains these points by substituting their coordinates in Eq. (10). If the obtained parameters are realistic [22], they are stored as a potential final solution. This process is repeated a preset number of times, chosen as a compromise between speed and accuracy. Whenever a new hyperbola is found, it has to be checked whether it is already found, and if it is the case, the number of times it was found increases. At the end, two possibilities exist for deciding which hyperbolae are the real solutions, by ranking them based either on the number of times each of them was found, or on the number of foreground pixels each of them contains. If the number of times the RHT is performed is high enough, the two ways should give the same results. If the number of times is not very high, the second way leads to better results.

If the detection is performed in an automated way, another question is how to determine how many highly ranked hyperbolae to preserve. It depends on how many objects can be expected in a scene. While estimating that number, one should not forget that depth is still a free dimension, so that there can be a few objects, placed one below the other. Therefore, in case there is no certain information regarding the possible number of objects in the scene, the safest is to choose several hyperbolae and eliminate some of them later, on the basis of how realistic the measures estimated from each of them are (Section 6). If human interaction is envisaged, the problem can be simply solved by visually estimating how many hyperbolae indeed exist in the scene.

5.2. Chosen B-scan measures

Eqs. (12), (11) and (13) can be rewritten, respectively, in the following forms:

$$v = \frac{2\Delta Y}{\Delta T\sqrt{B}}, \quad (14)$$

$$D = \frac{v\Delta T\sqrt{A}}{2}, \quad (15)$$

$$i_0 = C. \quad (16)$$

Therefore, once the typical hyperbolic shape is extracted from a B-scan, the propagation velocity in the medium above the object, the burial depth of the object and its cross-track

coordinate xy -position (in other words, its 3D position) can be estimated. In addition, the hyperbola opening is proportional to the size of the object d [12]:

$$d = k\frac{A}{B} \quad (17)$$

with k being the characteristic of the scattering function that depends on object shape. Since we do not have information about the shapes of objects, we assume that their scattering functions are approximately the same.

There are two ways of finding the burial depth, not always leading to the same results, so we use two notations in the following, D^* and D . Taking into account that \sqrt{A} is the position of the top of hyperbola, D^* is the depth obtained from it through proportionalities, as explained in Section 4 for depths below the soil surface (see Eq. (4)), i.e. for values $\sqrt{A} > n_{surf}$. If this condition is not satisfied, meaning that the top of such an object is found to be above the surface, D^* is set to 0 (see Section 4). D is the depth that is found from the velocity using Eq. (15), taking into account, once again, where the top of the obtained hyperbola (\sqrt{A}) is in comparison with the position of the air/ground interface. If $\sqrt{A} > n_{surf}$ is satisfied, Eq. (15) is modified to

$$D = \frac{v\Delta T(n_{surf} - \sqrt{A})}{2}. \quad (18)$$

In this case, D presents the depth measured from the air/ground interface and it has a negative sign. If $\sqrt{A} \leq n_{surf}$, Eq. (15) is applied, giving the distance between the top of hyperbola and the sample number 0. As said in Section 4, this number does not say much, except that its positive sign means that such an object is laid above the surface. Consequently, D^* and D must not be compared directly above the surface. Below the soil surface, D^* and D ideally should be equal. In reality, their differences result from the way they are estimated, since D^* assumes that v remains the same everywhere and equal to the one that is found for the reference object. If v estimated from hyperbolae differs significantly from that reference value, the value of D will be affected. The antennas of the GPR are generally not optimized for detection above the surface, which diminishes the need for a precise depth estimation there. The only important information in that case is its positive sign, indicating that such an object is above the surface.

The propagation velocity calculated from Eq. (14) should have a value within a known range of values for that type of soil if the object is buried below the soil surface (if D is negative). If the object is placed above the soil surface (positive D), ideally, v should be equal to the propagation velocity in the air, $c = 3 \times 10^8$ m/s.

Based on the above, we select the following independent measures extracted by hyperbola detection for further modeling:

- depth information given by D^* ;
- propagation velocity v together with the sign of D ;
- ratio between object size and its scattering function, d/k .

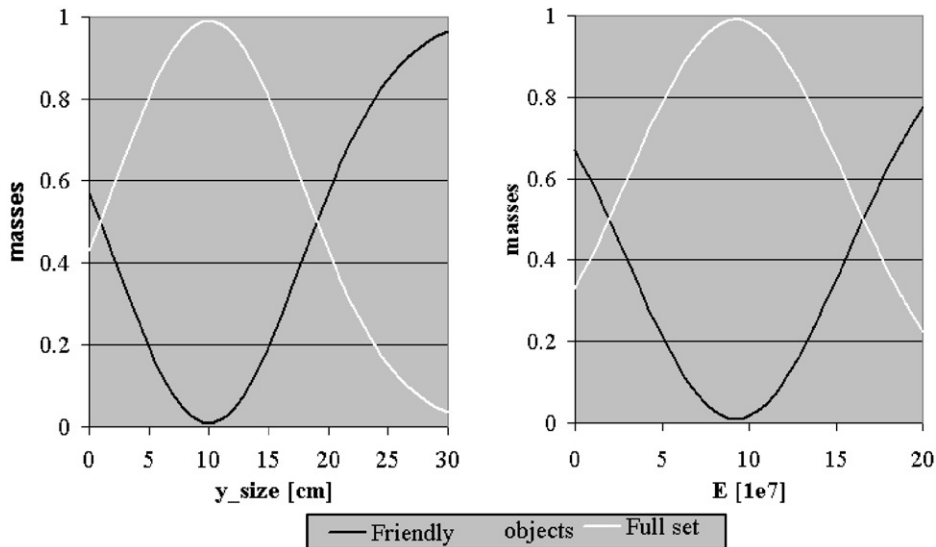


Fig. 2. Masses assigned by y_{size} measure (left) and by E measure (right).

6. Modeling and combination of measures in terms of belief functions

In real mine detection situations, the acquired data are far from numerous enough for reliable statistical learning. Besides, they are highly variable depending on the context and conditions [25]. Furthermore, not every possible object, neither mines nor objects that can be confused with them, can be modeled. On the other hand, some general knowledge exists regarding AP mines, their sizes, shapes, burial depths, etc., as well as regarding GPR detection possibilities. For these reasons, we decide to model and combine the GPR measures in terms of belief functions within the DS theory, since in this framework ignorance, partial knowledge, uncertainty and ambiguity can be appropriately modeled [13,14].

In the following, the frame of discernment Θ consists of two classes of objects: M (mine), and F (friendly, i.e. non-dangerous object, including background). The modeling step aims at defining a mass function for each measure expressing the information provided by this measure on the presence of a mine. Note that the mass function is the distribution of an initial unitary amount of belief among the subsets of Θ [15].

6.1. A-scan and preprocessed C-scan measures

y_{size} mass assignment: This measure is extracted as the width of a region selected by the local maxima method. It cannot provide information about mines alone. Although we know the approximate range of sizes of AP mines, still, whenever an object has a size within that range, it can be something else as well. Therefore, in that range, masses

should be mainly assigned to the full set, Θ . If the object is too large or too small, it is far more likely that it is not a dangerous one, and a large part of mass should be given to friendly objects. Therefore, it makes sense to model masses as given in Fig. 2 (left). The position of the center and the width of the central interval, where masses go mainly to the full set, depend on the available information. If there is no information regarding expectable mine size, or if a wide range of sizes can be expected, these curves should be quite non-informative, so the central interval should be very wide. If it is known which types of mines can be expected in a minefield and their size is similar, the curves, and accordingly, this measure, become very selective (narrow central interval).

E mass assignment: Similarly to y_{size} , whenever the energy is as expected for mines, it can be any other object as well, assigning masses mainly to Θ . Otherwise, it is likely that an object is friendly, so masses can be modeled as shown in Fig. 2 (right). The remaining reasoning is the same as for the previous measure.

d_1 mass assignment: If an object is buried too deep, it is possibly a non-dangerous one. Otherwise, it can be anything. Following this logic, masses can be modeled as shown in Fig. 3 (left).

6.2. B-scan (hyperbolae) measures

D^ mass assignment:* Since its meaning is the same as that of d_1 , D^* is modeled as given in the left side of Fig. 3.

v mass assignment: As said earlier, the value of the propagation velocity depends on the medium, so, in the case of the soil, it should be around the values for this medium, and if it is the air, it should be close to c . In order to decide which

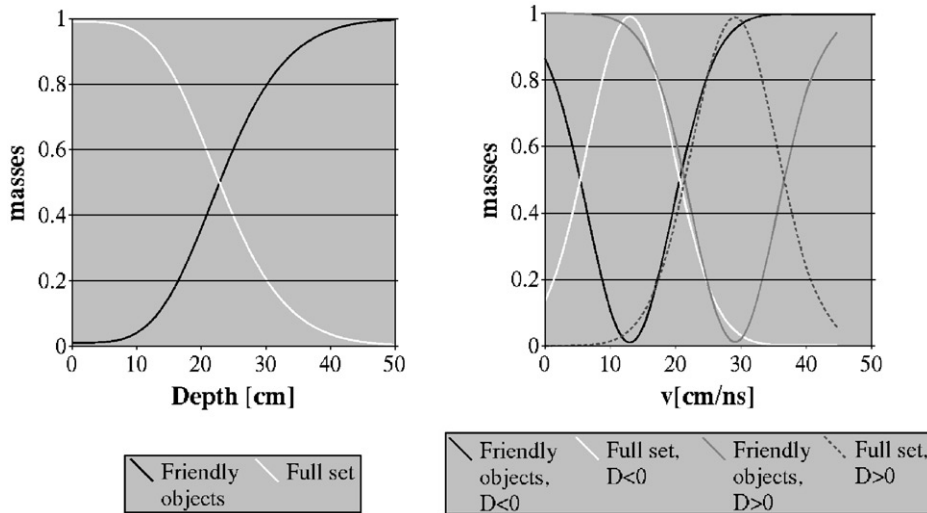


Fig. 3. Masses assigned by the depth measure (left) and by velocity measure (right).

model should be used, the sign of D is used as indicator. If the value of v is expectable for a particular medium, an object that gives that estimation of v can be anything. If v differs significantly from expected values for that medium, it can be expected that it is something friendly or simply background. This reasoning is illustrated in the right side of Fig. 3.

d/k mass assignment: If this measure is within a range of values that can be expected for mines, such an object can be anything from the full set. For very low or high values of this parameter, it is quite certain that the object is non-dangerous. Along this idea, masses are modeled as shown in the left side of Fig. 4.

6.3. Discounting

As pointed out already, the behavior of the GPR is strongly scenario-dependent (see also [26]), referring to:

- the quality of the acquired data,
- GPR reliability/detection ability for a particular type of soil, moisture, depth, etc.,
- types of objects under analysis.

Because of that, we use the possibility offered by the DS framework to give different importances to masses using discounting [13,27,28]. Discounting factors have been proven to be useful in the context of mine detection [26], and we specify them here for GPR measures. For a measure l , the discounting factor d_l consists of three types of parameters:

- g_l —confidence level of GPR in its assessment when judging this measure l (0—not confident at all, 1—completely confident);

- b_l —level of importance of the measure l (1—very low, b_{scale} —very high, where b_{scale} is the scale for the b parameter);
- s —deminer’s confidence into GPR (1—very low, s_{scale} —very high, where s_{scale} is the scale for the s parameter.)⁴

Due to the fact that a main factor that influences performance of a standard GPR is depth, g_l is the same function of depth, g (such as the one shown in the right side of Fig. 4) for all measures. Other factors, such as moisture, type of the soil, etc. are difficult to measure or quantify and for the moment, they are included in the model via the deminer’s confidence in the GPR within a particular scenario. Since these factors influence the maximum depth range of a GPR as well, the shape of g can be adjusted to the scenario. For example, our preliminary analysis shows that for this particular type of GPR and for relatively dry sand, there is no need for heavy discounting. Note that the reliability regarding the depth information depends on the used GPR and the scenario at hand, so in practice, for some GPR, it is known in which situations it makes sense to include the discounting in function of the depth.

Each discounting parameter can simply be used in a successive discounting, and then the global discounting factor is their product:

$$d_l = g \frac{s}{s_{scale}} \frac{b_l}{b_{scale}}. \tag{19}$$

⁴ As discussed in Ref. [26], the choice of b_{scale} and s_{scale} values should be left to the deminer, depending whether he prefers a wider or a narrower scale.




Fig. 4. Masses assigned by the d/k measure (left), discounting factor g as a function of the depth (right).

Using these coefficients, initial masses, m_l , assigned for each measure l are modified into new masses, m_{NEW} :

- $\forall A \subset \Theta, A \neq \Theta$:

$$m_{NEW}(A) = d_l m_l(A); \tag{20}$$

- for the full set:

$$m_{NEW}(\Theta) = d_l m_l(\Theta) + 1 - d_l. \tag{21}$$

In the following, we compare the combination obtained from masses without discounting and from masses with discounting. We assume that the deminer does not intervene so that all measures are equally important and that there is a full confidence in GPR ($s = s_{scale}, b = b_{scale}$). In other words, the only discounting parameter that will be included is g .

6.4. Combination of measures

Masses that are calculated for each measure of GPR can be combined using the well-known Dempster’s rule [13] in unnormalized form [14]:

$$m(A) = \sum_{A_i \cap B_j \dots \cap C_k = A} m_1(A_i) m_2(B_j) \dots m_n(C_k), \tag{22}$$

where m_1, m_2, \dots, m_n are basic mass assignments corresponding to n measures extracted from GPR data, and their focal elements are $A_1, A_2, \dots, A_l, B_1, B_2, \dots, B_p, \dots, C_1, C_2, \dots, C_q$, respectively. A general idea, discussed in Refs. [25,26], for using unnormalized instead of more usual normalized form is to preserve conflict. Here, due to the chosen measures and their modeling linked with some specificities of GPR data, there is no possibility of having conflict, so unnormalized and normalized DS rules, actually, give the same results. Indeed, focal elements are always F and Θ and therefore never contradict.

6.5. Guesses

The last step in the fusion process is decision making. Since the final decision has to be taken by the deminer in this type of application, to avoid confusion, we use the term “guess” for the first decision made automatically from the numerical combination results. This section aims at defining these guesses. The application is very special in comparison to other fusion applications in the sense that the different possible hypotheses do not have the same impact: mines cannot be missed, so false recognition of a mine as a friendly object is incomparably worse than false recognition of a friendly object as a mine. Therefore in case of ambiguity, it is safer to choose M . Here, due to the fact that GPR is an anomaly detector (not restricted to mines), the mass model we proposed focuses on F and Θ only, and M never appears as a focal element. Namely, whenever an object can be a mine, from the point of view of GPR, it can be something else as well. Since mines must not be missed, it is necessary to be cautious [26], so the mass assigned to the full set is treated as the possibility that an object is a mine. On the basis of this reasoning, guesses are determined by simply comparing which mass is larger, i.e. the one assigned to friendly objects (in that case, our guess is “F”, friendly object), or to the full set (the guess is “M”, mine). Note that this could not have been done at the modeling step. For instance, if we decide for each measure that everything assigned to Θ is put on M (for sake of cautiousness), then the fusion will have to face conflicts with F , while keeping the mass on Θ really means both F and M are possible, and another measure decides.

It is possible to derive direct links between the guesses and belief (Bel) and plausibility (Pls) functions [14]. Taking into account that the focal elements are here only F and Θ , we have for each measure and after combination (since no

conflict occurs with this simple model):

$$m(\Theta) = 1 - m(F) = Pls(M), \quad (23)$$

$$Bel(F) = m(F). \quad (24)$$

The general decision rule we proposed [17] is as follows:

$$G(M) = Pls(M), \quad (25)$$

$$G(F) = Bel(F), \quad (26)$$

$$G(\emptyset) = m(\emptyset). \quad (27)$$

The last term aims at dealing with the open-world assumption and is useful when combining several sensors that can for instance focus on different objects [17]. Here only GPR measures are considered, and this case cannot occur. The decision rule is then as follows:

$$\text{if } m(F) > m(\Theta), \text{ then } F, \quad (28)$$

$$\text{if } m(\Theta) \geq m(F), \text{ then } M, \quad (29)$$

which is equivalent to

$$\text{if } Bel(F) > Pls(M), \text{ then } F,$$

$$\text{if } Pls(M) \geq Bel(F), \text{ then } M. \quad (30)$$

The equivalence in terms of plausibility of M for $m(\Theta)$ and of belief of F for $m(F)$ illustrates the cautiousness requirement in this type of application.

7. Results

7.1. Specificities of the acquired data and of the preprocessed C-scan

The used data were acquired in a part of a lane filled with sand at the TNO test facilities [16]. The analyzed part of the lane is actually its upper part (see Figs. 8 and 9). It contains several types of antipersonnel (AP) mines buried at various

depths, as well as some false alarms. The area reserved for one object is approximately $45 \text{ cm} \times 45 \text{ cm}$. Objects are buried in lines. These lines are also the direction of collecting B-scans. The data acquisition step in this direction was around 9 mm, while in the cross-track direction, it was 10 cm. In depth, presenting wave travel-time proportional to it, the sampling time was 50 ps. The non-analyzed part of the lane, containing antitank (AT) mines and some false alarms, was used for testing, TVG estimation and derivation of models.

7.2. Choice of TVG

An example of a raw A-scan belonging to a target buried at around 30 cm below the surface is given in the left side of Fig. 5, where the x -axis is expressed in number of time samples, so directly proportional to t . The position of the first maximum, around time-sample number 38, corresponds to the air/ground interface. The result of applying the TVG of 1 dB/ns is given on the right side of Fig. 5. This value of TVG is our final choice, based on the analysis of energies of one object buried at different depths, aimed to preserve approximately the same energy regardless the depth. To give a more global illustration of influence of TVG on the data, a raw B-scan and the same scan after TVG are shown in Fig. 6, in case of a deeper buried AT mine.

7.3. Results of background removal

The number of targets in our data as well as their separation decrease the usefulness of the method to some extent, so some other, more sophisticated and complex methods for background removal would probably give better results. Nevertheless, we are looking for a realistic, simple, fast and automated method, and since it gives satisfactory results, this method is accepted as our final choice. A result of sliding window mean background removal applied on the A-scan from the right side of Fig. 5 is shown in Fig. 7 (left), while the B-scan from the right side of Fig. 6 after background removal is shown in Fig. 7 (right).




Fig. 5. An example of a raw A-scan (left), and after the TVG of 1 dB/ns (right).




Fig. 6. A raw B-scan of an AT mine (left) and after TVG of 1 dB/ns (right).




Fig. 7. The result of background removal applied to the signal from the right side of Fig. 5 (left) and to the B-scan from the left side of Fig. 6 (right).




Fig. 8. The C-scan containing energy projected A-scans after TVG of 1 dB/ns and background removal; ‘x’ marks represent the ground-truth.

7.4. Obtained C-scan containing energy projections of preprocessed A-scans

In Fig. 8, a C-scan is shown, obtained after applying $TVG = 1$ dB/ns and background removal (Eq. (3)), where $j = 1, 2, \dots, 256$, $\Delta T = 50$ ps) on raw data. Note that, due to the great difference in acquisition steps in the two directions

of the surface plane, the two axes in Fig. 8 are not equal, in order to see them clearer. In all the remaining images presenting processed C-scans of the lane, axes are equal, to give a better view of the lane and of selected regions.

7.5. Selected regions

As an illustration, Fig. 9 (left) shows the result of the local maxima method for $thr_1 = 4 \times 10^7$, $thr_2 = 80\%$, applied on the data shown in Fig. 8, where a gray-scale value represents a selected region. For a comparison, the result of simple thresholding of the same data on value 4×10^7 is presented in Fig. 9 (right). As can be seen, besides the fact that the local maxima method is aimed for detecting regions belonging to objects with weaker signals without increasing clutter detection (Section 3), it differs from simple thresholding in two other points. Namely, size and shape of regions are different, and several selected regions are grouped together as corresponding to the same object. Instead of having 72 thresholded areas or alarms, they are grouped in 26 alarms. The real identity of these 26 regions is given in Table 1, with “real object” being the real object closest to the center of the alarm. The coordinates of the corresponding local maximum are (x_c, y_c) , where x is the cross-track coordinate, and y is the scanning coordinate.




Fig. 9. Ground truth information (x) and the regions selected by: (a) left—the local maxima method applied to the data from the left side of Fig. 8, for $thr_1 = 4 \times 10^7$, $thr_2 = 80\%$, window half-size 20 cm (in the cross-track direction) and 25 cm (in the scanning direction); (b) right—simply thresholding the data from the left side of Fig. 8, for the threshold level equal to 4×10^7 .

Table 1
The real identity of 26 selected regions

Region	x_c (m)	y_c (m)	Real object	Identity
1	1.85	7.1	Waterlevel tube	F
2	0.85	5.27	Iron fragment, 10 cm depth	F
3	1.85	5.59	Al can, 0 depth	F
4	1.45	6.89	NR24, with ring (metallic, small), 0 depth	M
5	1.45	4.8	M14 (low metal, small), 0 depth	M
6	0.55	7.36	Wooden AP (low metal, medium), 10 cm depth	M
7	1.85	6.18	Brick, 10 cm depth	F
8	0.55	5.95	PMN2 (metallic, medium), 10 cm depth	M
9	1.35	6.45	PMN (metallic, medium), 0 depth	M
10	0.45	6.42	PMN, 10 cm depth	M
11	1.55	5.68	NR24, no ring (no metal, small), 0 depth	M
12	0.95	7.08	PFM-1 butterfly (metallic, small), above the surface	M
13	1.35	5.93	PMN2, 0 depth	M
14	1.45	5.07	M14, 0 depth	M
15	1.05	4.89	Nothing	F
16	0.25	7.27	Nothing	F
17	0.65	4.93	M14, 10 cm depth	M
18	0.95	5.84	Iron fragment, 0 depth	F
19	0.25	7.53	Nothing	F
20	1.75	5.09	Iron fragment, above the surface	F
21	1.35	7.36	Wooden AP, 0 depth	M
22	0.35	4.81	Nothing	F
23	0.25	5.25	Nothing	F
24	0.55	7.1	NR24 with ring, 10 cm depth	M
25	1.75	4.72	Nothing	F
26	0.95	6.19	PFM-1 butterfly, 10 cm depth	M

Table 2

Measures extracted from the preprocessed C-scan after local maxima detection (Fig. 9, left) and A-scans after TVG and background removal, for each of the selected regions from Table 1

Region	y_{size} (cm)	$E(10^7)$	n_1 (Sample no.)	d_1 (cm)
1	2.76	17.3	57	-5.7
2	17.48	14.6	49	-3.3
3	4.6	14.6	41	-0.9
4	3.68	13.4	32	0
5	10.12	12.3	29	0
6	3.68	12.2	58	-6
7	3.68	11.4	89	-15.3
8	3.68	11.2	71	-9.9
9	2.76	10.51	52	-4.2
10	3.68	9.41	77	-11.7
11	11.04	8.56	59	-6.3
12	2.76	8.48	30	0
13	2.76	7.81	32	0
14	16.56	7.51	30	0
15	21.16	7.44	46	-2.4
16	25.76	7.39	40	-0.6
17	3.68	0.23	71	-9.9
18	25.76	6.59	55	-5.1
19	2.76	5.70	41	-0.9
20	18.4	5.69	42	-1.2
21	3.68	5.67	47	-2.7
22	18.4	5.49	33	0
23	16.56	5.27	240	-60.6
24	12.88	5.13	32	0
25	14.72	4.67	48	-3
26	12.88	4.22	47	-2.7

Most of these regions contain objects, but some are simply clutter. These alarms are the basis for the extraction, modeling and combination of measures presented in the following. The regions are ordered by the values of their local maxima, with the highest one on the top. Note that a depth of 0 means that an object is buried just below the surface. It can be seen that the regions 5 and 14 possibly refer to the same object, meaning that, although the chosen size of the window around the objects works satisfactory for most of the objects, in this case it was too small.

7.6. Extracted A- and C-scans measure

The three chosen measures, y_{size} , E and d_1 , corresponding to 26 selected regions, are given in Table 2. In our case: $d_{ref} = 30$ cm, $n_{ref} = 138$, $n_{surf} = 38$ (see Section 4).

7.7. Extracted measures from B-scans

Table 3 contains measures that are found from hyperbolae detected in B-scans corresponding to each of 26 selected regions. Here, X_0 presents the x -coordinate of a B-scan, Y_0 is found from the C parameter of a detected hyperbola (through

Eqs. (9) and (13)), while d/k , v , D and D^* are determined as discussed in Section 5.2. Several important conclusions can be made by analyzing Table 3. For example, comparing the parameters of regions 5 and 14 (their xy -coordinates as a first indicator), it can be easily concluded that it is, actually, the same object split in two regions. Therefore, even if the arrangement of objects and their sizes are such that it makes it difficult to determine the optimum size of the window for local maxima analysis, so that one object is detected twice, these errors can be corrected at the moment each of the regions is analyzed in more detail. It can be further concluded that a similar effect would occur if the size of the window is too large, so that at the beginning, one selected region contains a few objects. In that case, a few hyperbolae would be detected in the corresponding B-scan.

Another important piece of information can be obtained by analyzing results for the propagation velocities, v . For sandy soils, v ranges typically from 5.5×10^7 to 1.73×10^8 m/s [12]. Using the information about the true identity of objects in each of the regions given in Table 1, it can be seen that in case of buried objects, v has indeed values within the given range, while for surface-laid objects, it is close to c . Note also that in cases of false alarms, v often reaches unreasonable values, sometimes even higher than in free space. This is a good indicator that there is really nothing in that region, and that the obtained hyperbola is caused by some artifacts or simply clutter.

Finally, the measure d/k also shows that there is a range of values that describes most objects, and that extreme values mainly belong to background.

7.8. Combination of A-scan and preprocessed C-scan measures

Once masses are assigned for each of the measures given in Table 2, they are combined either immediately or after being discounted (See Sections 6.3 and 6.4). The results are given in Table 4. False alarms are written in italic.

Observing the part concerning non-discounted masses, and taking into account real identities of the objects (Table 1), it can be seen that clutter is mainly discarded (except in two cases: regions 19 and 25). Furthermore, for several mines (regions 6, 7, 11, 21, as the most critical ones) the masses assigned to friends and to the full set are very close, meaning that, although correctly classified here, in case of slight changes in the model, they could be easily misclassified. If the discounting is included, it means that some part of the mass initially assigned to friends goes to the full set, so the sensitivity of some masses is decreased. However, some non-dangerous objects, the two masses of which are close to each other and that are well classified without discounting, can also become potential mines (such as regions 20 and 22). The danger of this type of misclassification is in slowing down the mine removal process, since the non-dangerous objects will have to be removed too. However, this is much less important than the danger of

Table 3

Measures extracted from the preprocessed B-scans after hyperbola detection, for each of the selected regions from Table 1

Region	X_0 (m)	Y_0 (m)	D^* (cm)	$(d/k)[(y\text{-dir. sample no.})^2]$	$v(10^8 \text{ m/s})$	D (cm)
1	1.85	7.1	-4.5	486	1.48	-5.54
2	0.85	5.29	-2.4	363	1.53	-3.07
3	1.85	5.62	0	244	1.53	0
4	1.45	6.95	0	768	2.95	25.61
5	1.45	5.02	0	548	3.13	21.6
6	0.55	7.35	-6.3	471	1.44	-6.44
7	1.85	6.17	-13.8	586	10.6	-12.3
8	0.55	5.96	-6	556	1.51	-7.5
9	1.35	6.46	-2.7	371	1.5	-3.54
10	0.45	6.44	-10.8	542	1.17	-10.44
11	1.55	5.75	-0.6	265	1.51	-0.75
12	0.95	7.09	0	469	2.52	20.03
13	1.35	5.95	-2.4	393	1.6	-3.17
14	1.45	5.03	0	562	3.15	21.91
15	1.05	5.07	-0.6	1245	3.26	-1.63
16	0.25	7.18	0	1099	4.37	30.7
17	0.65	4.95	-9.3	839	1.55	-12.07
18	0.95	5.85	-5.4	730	1.78	-7.9
19	0.25	7.4	0	1285	3.8	33.21
20	1.75	5.09	0	493	3.04	20.55
21	1.35	7.37	0	575	2.79	22.15
22	0.35	4.79	0	178	1.59	12.33
23	0.25	5.13	-15.9	15	0.16	-2.09
24	0.55	6.91	-7.5	645	1.49	-9.32
25	1.75	4.69	0	1407	4.48	34.69
26	0.95	6.16	-8.1	362	1.08	-7.3

the opposite situation, where some mines are not removed because they are classified as non-dangerous objects. Therefore, when the chosen measures do not distinguish well between mines and friends, the safest way is to apply discounting, and by that increase confidence in the final guess (although possibly slightly increasing the number of false alarms).

In the discounted case, note that 5 of 12 non-dangerous objects are still discarded, which is a good improvement when compared with the starting situation, meaning that without further analysis of the regions, all 26 regions would be treated as dangerous. Also, it can be seen that the strongest discounting occurs in the case of region 23, due to the fact that the extracted depth was around -60 cm.

Finally, it must be repeated that there is no rejection of mines. Although our goal is to decrease both number of false rejections and of false alarms, it must not be forgotten that, taking into account the problem we are dealing with, false rejections have an incomparably higher importance than false alarms.

Note that discounting parameters s and b are not included. It can be expected that these already good results can be further improved once the deminer's knowledge and experience is included in the reasoning process. It is out of the scope of this paper to discuss the way the deminer would

reason for each detected region, since that process can be quite subjective.

7.9. Combination of B-scan (hyperbolae) measures

After assigning masses to measures extracted from detected hyperbolae (Table 3), they are combined in two ways, without being discounted and after discounting. Masses after their combination are given in Table 5, together with first guesses on object identities. As can be seen, discounting does not have any influence on the final guesses this time, indicating that we can be more confident in obtained results, i.e. that chosen measures and their models are well-suited for the data and robust. Also, looking back to the real identities of the regions, shown in Table 1, it can be seen that all the alarms corresponding to background are classified as non-dangerous (regions 15, 16, 19, 22, 23, 25), and, in addition, with a high confidence in most of the cases.

On the other hand, from the placed false alarms, only one of them is well classified (region 3, aluminum can), but this result is a direct consequence of the fact that this method actually detects whether there is an object, regardless whether it is a mine or not.

Nevertheless, it has to be pointed out that this method is very promising for several reasons. Namely, clutter, which

Table 4

Resulting masses after combination of masses assigned by chosen A-scan and C-scan measures, without and with discounting, for each of the selected regions from Table 1; values of discounting parameters are also given

Region	No discounting			g	With discounting		
	$m(F)$	$m(\Theta_G)$	Guess		$m(F)$	$m(\Theta_G)$	Guess
1	0.78	0.22	F	0.98	0.77	0.23	F
2	0.62	0.38	F	0.91	0.58	0.42	F
3	0.52	0.48	F	0.71	0.39	0.61	M
4	0.48	0.52	M	0.6	0.31	0.69	M
5	0.18	0.82	M	0.6	0.11	0.89	M
6	0.41	0.59	M	0.99	0.4	0.6	M
7	0.46	0.54	M	0.8	0.4	0.6	M
8	0.37	0.63	M	0.96	0.36	0.64	M
9	0.39	0.61	M	0.95	0.37	0.63	M
10	0.35	0.65	M	0.91	0.32	0.68	M
11	0.05	0.95	M	1	0.05	0.95	M
12	0.38	0.62	M	0.6	0.23	0.77	M
13	0.39	0.61	M	0.6	0.24	0.76	M
14	0.34	0.66	M	0.6	0.21	0.79	M
15	0.67	0.33	F	0.85	0.57	0.43	F
16	0.88	0.12	F	0.68	0.61	0.39	F
17	0.36	0.64	M	0.96	0.34	0.66	M
18	0.89	0.11	F	0.97	0.87	0.13	F
19	0.47	0.53	M	0.71	0.35	0.65	M
20	0.54	0.46	F	0.75	0.42	0.58	M
21	0.41	0.59	M	0.87	0.36	0.64	M
22	0.55	0.45	F	0.6	0.35	0.65	M
23	1	1.1e-4	F	0.03	0.01	0.99	M
24	0.28	0.72	M	0.6	0.18	0.82	M
25	0.39	0.61	M	0.89	0.36	0.64	M
26	0.36	0.64	M	0.87	0.32	0.68	M

is typically the most significant cause of false alarms in case of GPR, is almost completely discarded here. In addition, slight modifications of the way measures are modeled and inclusion or not of the discounting factors do not affect final results, meaning that the chosen measures and models allow for a good and reliable discrimination. Finally, there are no false rejections.

Table 6 summarizes all results discussed here and in the previous subsection.

The results obtained by this method and by a voting method are compared in Ref. [29] on a part of the data used here. In that paper, we both analyzed performance of each of the sensors as well as combined the following three sensors: infrared camera, metal detector, and GPR. Regarding GPR results using belief functions, we analyzed only results obtained by B-scan measures. For this sensor, both methods classified well all mines (12 of them). The real power of belief functions was seen in case of clutter-caused alarms, where 6 out of 7 were correctly classified (so rejected) by belief functions, while all 7 were misclassified (so became false alarms) using voting. The conclusion was that without additional knowledge, which is the case of voting, all alarms

are treated as mines, leading to a high false alarm rate, while in the belief function framework, knowledge helps in decreasing the false alarm rate without decreasing the result of mine detection.

7.10. Combination of A-, B- and C-scans measures

Finally, we address the combination of A-, B- and C-scan measures, which may appear as a promising solution. Firstly, it should be noted that it cannot be done straightforward, due to some dependencies between measures (such as two burial depth measures, one per set). Therefore, we performed preliminary tests on combination by excluding one of the two burial depth measures. These tests show that there is no significant improvement in the final results. The combination of all measures is therefore not useful. Moreover, it is highly questionable how realistic this way would be in humanitarian mine detection context, due to the different types of processing that each of them requests. A potential solution, that will be analyzed in detail in our future work, is in finding ways for selecting a subset of the two sets containing most discriminant and cognitively independent measures.

Table 5

Resulting masses after combination of masses assigned by chosen B-scan measures, without and with discounting, for each of the selected regions from Table 1; values of discounting parameters are also given

Region	No discounting			g	With discounting		
	$m(F)$	$m(\Theta_G)$	Guess		$m(F)$	$m(\Theta_G)$	Guess
1	0.14	0.86	<i>M</i>	0.96	0.14	0.86	<i>M</i>
2	0.3	0.7	<i>M</i>	0.85	0.25	0.75	<i>M</i>
3	0.94	0.06	F	0.6	0.65	0.35	F
4	0.07	0.93	M	0.6	0.04	0.96	M
5	0.12	0.88	M	0.6	0.07	0.93	M
6	0.15	0.85	M	0.99	0.14	0.86	M
7	0.2	0.8	<i>M</i>	0.85	0.17	0.83	<i>M</i>
8	0.11	0.89	M	0.99	0.11	0.89	M
9	0.28	0.72	M	0.87	0.24	0.76	M
10	0.12	0.88	M	0.94	0.11	0.89	M
11	0.42	0.58	M	0.68	0.29	0.71	M
12	0.26	0.74	M	0.6	0.16	0.84	M
13	0.29	0.71	M	0.85	0.25	0.75	M
14	0.12	0.88	M	0.6	0.07	0.93	M
15	0.99	0.01	F	0.68	0.82	0.18	F
16	0.96	0.04	F	0.6	0.68	0.32	F
17	0.21	0.79	M	0.98	0.21	0.79	M
18	0.28	0.72	<i>M</i>	0.98	0.27	0.73	<i>M</i>
19	0.89	0.11	F	0.6	0.65	0.35	F
20	0.12	0.88	<i>M</i>	0.6	0.07	0.93	<i>M</i>
21	0.06	0.94	M	0.6	0.04	0.96	M
22	0.94	0.06	F	0.6	0.67	0.33	F
23	0.95	0.05	F	0.77	0.85	0.15	F
24	0.08	0.92	M	0.99	0.08	0.92	M
25	0.99	0.01	F	0.6	0.79	0.21	F
26	0.3	0.7	M	0.99	0.3	0.7	M

Table 6

Number of correct and wrong classifications over total number of mines (14) or friendly objects (6 placed and 6 clutter-caused), per analyzed set of measures, with or without discounting

	A- and C-scans measure		B-scan measure	
	No disc.	With disc.	No disc.	With disc.
Mine detection	14/14	14/14	14/14	14/14
Mine rejection	0/14	0/14	0/14	0/14
Correctly recognized placed F	5/6	3/6	1/6	1/6
Placed F recognized as a mine	1/6	3/6	5/6	5/6
Correctly recognized clutter-caused F	4/6	2/6	6/6	6/6
Clutter-caused F recognized as a mine	2/6	4/6	0/6	0/6

8. Conclusion

In this paper, we investigate a method based on GPR data for humanitarian mine detection. This study is motivated by the fact that GPR has very useful features for this type of application. In particular, it is able to provide 3D information about the subsurface structure, in contrary to most other sensors. The difficulty of interpreting GPR data calls for

specific processing relying on the specificities of the sensor characteristics. Our contribution in this direction is twofold. Firstly, we develop tools for processing A- and C-scans, for extracting useful measures from local maxima, and for detection and measuring hyperbolas in B-scans. Secondly, we propose an appropriate model of the extracted measures in terms of belief functions and discounting factors, which includes knowledge about the sensor and the context of the ap-

plication. The fusion of these measures leads to information on the presence of a friendly object or of any (undifferentiated) object. We derive a decision rule adapted to this result and to the constraint of the application (no mine should be missed).

Normally, after selecting suspicious regions, the process of detection is finished and the mine removal operations begin. Still, both in general as well as in our case, a lot of non-dangerous objects and clutter are selected as well, which means that the removal proceeds slowly. To deal with that, we propose a way for further analysis of suspected regions through choosing discriminative measures, their modeling and combination in terms of belief functions within the DS framework. Two sets of measures are separately analyzed, one extracted from preprocessed A-scans and a C-scan on which regions are selected, and the other extracted from hyperbolae detected on preprocessed B-scans corresponding to the selected regions. For each of these measures, masses are modeled and assigned either to friendly objects or to the full set, i.e. both mines and friendly objects, due to the fact that GPR, as well as a great majority of other mine detection sensors, is not aimed for mine but for anomaly detection.

Reasons for eventual discounting of measures before their combination are also discussed here, and assigned masses are combined both without being discounted as well as after discounting. It is shown that in cases when the measures do not give a good discrimination between friendly objects and the full set, it is safer (taking into account the extreme danger in misclassifying mines) to make final guesses about the true object identity based on results of combining discounted masses. Here, such situations correspond to the case of combining A- and C-scans measures. On the other hand, B-scan measures extracted from detected hyperbolae show a great potential in discarding alarms caused by clutter, which is often the largest problem for GPR. Also, these measures discriminate well between friendly objects and the full set, so there is no need for further discounting. Both sets of measures have a 100% of detection of mines. Note that the analyzed data set contains 14 mines, all of them are small, most have little metal, and some have no metal at all.

A careful analysis and choice of most selective and independent measures from the two sets together, taking into account the processing aspect, could be a good solution for further improvements of the results, and it will certainly be done in future.

The methodology presented here can be applied to other mine detection sensors (where only the choice of measures and their modeling would be different, i.e. would be a function of the operational principles) and their fusion, as shown in Ref. [29]. In addition, various steps of the methodology can be useful themselves, such as A-, B- and C-scans processing, which can be useful to various other areas of GPR application besides mine detection, where visualization of a shallow subsurface using close-range GPR data is important, such as road inspection. Finally, the proposed region selection method can provide an important improvement in

fastening any detection process, where a coarser acquisition step could be used for an initial observation of a whole terrain, to which the region selection can be applied, and then a finer acquisition step can be applied only to the selected regions. This method can be used for other sensors (such as metal detector or infrared camera) and for other applications that involve terrain inspection than mine detection.

Acknowledgements

We are thankful to the TNO Physics and Electronics Laboratory for giving us the permission and great opportunity to work on the data gathered on their test facilities within the Dutch HOM-2000 project. Also, we thank Bart Scheers from the Royal Military Academy for useful comments and suggestions.

References

- [1] C.E. Baum (Ed.), *Detection and Identification of Visually Obscured Targets*. Taylor & Francis, Philadelphia, USA, 1999.
- [2] J.D. Young, L. Peters Jr., A brief history of GPR fundamentals and applications, in: *Proceedings of Sixth International Conference on Ground Penetrating Radar (GPR'96)*, Sendai, Japan, 1996, pp. 5–14.
- [3] B. Scheers, *Ultra-wideband ground penetrating radar, with application to the detection of anti personnel landmines*, Ph.D. Thesis, Université Catholique de Louvain, Louvain-La-Neuve, Belgium, 2001.
- [4] G.R. Olhoeft, Application of ground penetrating radar, in: *Proceedings of Sixth International Conference on Ground Penetrating Radar (GPR'96)*, Sendai, Japan, 1996, pp. 1–3.
- [5] D.J. Daniels, *Surface-Penetrating Radar*, The Institution of Electrical Engineers, London, UK, 1996.
- [6] K.B. Jakobsen, H.B.D. Sorensen, O. Nymann, Stepped-frequency ground-penetrating radar for detection of small non-metallic buried objects, in: *Proceedings of SPIE Conference on Detection Technologies for Mines and Minelike Targets*, Vol. 3079, Orlando, USA, 1997, pp. 538–542.
- [7] G.D. Sower, S.P. Cave, Detection and identification of mines from natural magnetic and electromagnetic resonances, in: *Proceedings of SPIE Conference on Detection Technologies for Mines and Minelike Targets*, Vol. 2496, Orlando, USA, 1995, pp. 1015–1024.
- [8] D.J. Daniels, Ground probing radar techniques for mine detection, in: *Proceedings of Seventh International Conference on Ground Penetrating Radar (GPR'98)*, Vol. 1, Lawrence, Kansas, USA, 1998, pp. 319–323.
- [9] L. van Kempen, H. Sahli, Ground penetrating radar data processing: a selective survey of the state of the art literature, Technical Report IRIS-TR-0060, Vrije Universiteit Brussel, Brussels, Belgium, 1999.
- [10] T.R. Witten, Present state-of-the-art in ground penetrating radars for mine detection, in: *Proceedings of SPIE Conference on Detection Technologies for Mines and Minelike Targets*, Vol. 3392, Orlando, USA, 1998, pp. 576–585.

- [11] N. Milisavljević, M. Acheroy, Overview of mine detection sensors and ideas for their estimation and fusion in the scope of HUDEM project, in: Proceedings of Australian–American Joint Conference on the Technologies of Mines and Mine Countermeasures, Sydney, Australia, 1999.
- [12] L. Capineri, P. Grande, J.A.G. Temple, Advanced image-processing technique for real-time interpretation of ground penetrating radar images, *Int. J. Imaging Systems Technol.* 9 (1998) 51–59.
- [13] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ, 1976.
- [14] P. Smets, What is Dempster–Shafer’s model? in: R.R. Yager, M. Fedrizzi, J. Kacprzyk (Eds.), *Advances in the Dempster–Shafer Theory of Evidence*, Wiley, New York, 1994, pp. 5–34.
- [15] P. Smets, R. Kennes, The transferable belief model, *Artif. Intell.* 66 (1994) 191–234.
- [16] W. de Jong, H.A. Lensen, Y.H.L. Janssen, Sophisticated test facility to detect landmines, in: Proceedings of SPIE Conference on Detection Technologies for Mines and Minelike Targets, Vol. 3710, Orlando, USA, 1999, pp. 1409–1418.
- [17] N. Milisavljević, Analysis and fusion using belief function theory of multisensor data for close-range humanitarian mine detection, Ph.D. Thesis, École Nationale Supérieure des Télécommunications, Paris, France, 2001.
- [18] M.G.J. Breuers, P.B.W. Schwing, S.P. van den Broek, Sensor fusion algorithms for the detection of land mines, in: Proceedings of SPIE Conference on Detection Technologies for Mines and Minelike Targets, Vol. 3710, Orlando, USA, 1999, pp. 1160–1166.
- [19] G.R. Olhoeft, Electrical, magnetic and geometric properties that determine ground penetrating radar performance, in: Proceedings of Seventh International Conference on Ground Penetrating Radar (GPR’98), Vol. 1, Lawrence, Kansas, USA, 1998, pp. 177–182.
- [20] V.F. Leavers, *Shape Detection in Computer Vision Using the Hough Transform*, Springer, London, 1992.
- [21] L. Xu, Randomized Hough transform (RHT): basic mechanisms, algorithms, and computational complexities, *CVGIP: Image Understanding* 57 (2) (1993) 131–154.
- [22] N. Milisavljević, I. Bloch, M. Acheroy, Application of the randomized Hough transform to humanitarian mine detection, in: Proceedings of the Seventh IASTED International Conference on Signal and Image Processing (SIP2001), Honolulu, Hawaii, USA, 2001, pp. 149–154.
- [23] R.A. McLaughlin, Randomized Hough transform: improved ellipse detection with comparison, Technical Report TR97-01, The University of Western Australia, CIIPS, 1997.
- [24] N. Milisavljević, Comparison of three methods for shape recognition in the case of mine detection, *Pattern Recognition Lett.* 20 (11–13) (1999) 1079–1083.
- [25] N. Milisavljević, I. Bloch, M. Acheroy, A first step towards modeling and combining mine detection sensors within Dempster–Shafer framework, in: Proceedings of 2000 International Conference on Artificial Intelligence (IC-AI’2000), Vol. II, Las Vegas, USA, 2000, pp. 745–751.
- [26] N. Milisavljević, I. Bloch, M. Acheroy, Characterization of mine detection sensors in terms of belief functions and their fusion, first results, in: Proceedings of the Third International Conference on Information Fusion (FUSION 2000), Vol. II, Paris, France, 2000, pp. ThC3.15–ThC3.22.
- [27] D. Dubois, M. Grabisch, H. Prade, P. Smets, Assessing the value of a candidate, in: 15th Conference on Uncertainty in Artificial Intelligence (UAI’99), Stockholm, Sweden, Morgan Kaufmann, San Francisco, 1999, pp. 170–177.
- [28] P. Smets, Belief functions: the disjunctive rule of combination and the generalized Bayesian theorem, *Int. J. Approx. Reasoning* (9) (1993) 1–35.
- [29] N. Milisavljević, S.P. van den Broek, I. Bloch, P.B.W. Schwing, H.A. Lensen, M. Acheroy, Comparison of belief functions and voting method for fusion of mine detection sensors, Proceedings of SPIE Conference on Detection and Remediation Technologies for Mines and Minelike Targets VI, Vol. 4394, 2001, pp. 1011–1022.

About the Author—NADA MILISAVLJEVIĆ is a researcher at the Signal and Image Centre of the RMA. She graduated from the University of Novi Sad, Yugoslavia, in Electrical Engineering (Electronics and Telecommunications) in 1992. She received the M.Sc. degree in Optical Electronics and Laser Technics from the University of Belgrade, Yugoslavia, in 1996, and her Ph.D. degree from ENST in Paris, France, in 2001. Her current research interests include data fusion, evidence theory, humanitarian mine detection and pattern recognition.

About the Author—ISABELLE BLOCH is a professor at ENST (Signal and Image Processing Department). She graduated from Ecole des Mines de Paris in 1986, received Ph.D from ENST Paris in 1990, and the “Habilitation Diriger des Recherches” from University Paris 5 in 1995. Her research interests include 3D image and object processing, 3D and fuzzy mathematical morphology, discrete 3D geometry and topology, decision theory, information fusion in image processing, fuzzy set theory, evidence theory, structural pattern recognition, spatial reasoning, medical imaging.

About the Author—SEBASTIAAN VAN DEN BROEK received his Ph.D. in physics in 1997 at the University of Twente (The Netherlands), where he did work on numerical modeling of electromagnetic activity in the human brain. Since 1998 he works as a researcher at TNO Physics and Electronics Laboratory in the field of electro-optics. His research interests are in image processing, sensor fusion and automatic target detection and tracking.

About the Author—MARC ACHEROY is an ordinary professor at the Faculty of Applied Sciences of the RMA in Brussels, Belgium, and founder and director of the Signal and Image Centre of the RMA. He received his degree of civil engineering in mechanics and transport from the RMA, the degree of civil engineering in control and automation (in 1981) and his Ph.D. degree in Applied Sciences (in 1983) from the Université libre de Bruxelles. His main research interests include pattern recognition, humanitarian demining, and image restoration and compression.