Sensor Fusion in Anti-Personnel Mine Detection Using a Two-Level Belief Function Model

Nada Milisavljević and Isabelle Bloch

***Abstract*—A two-level approach for modeling and fusion of anti-personnel mine detection sensors in terms of belief functions within the Dempster–Shafer framework is presented. Three promising and complementary sensors are considered: a metal detector, an infrared camera, and a ground-penetrating radar. Since the metal detector, the most often used mine detection sensor, provides measures that have different behaviors depending on the metal content of the observed object, the first level aims at identifying this content and at providing a classification into three classes. Depending on the metal content, the object is further analyzed at the second level toward deciding the final object identity. This process can be applied to any problem where one piece of information induces different reasoning schemes depending on its value. A way to include influence of various factors on sensors in the model is also presented, as well as a possibility that not all sensors refer to the same object. An original decision rule adapted to this type of application is proposed, as well as a way for estimating confidence degrees. More generally, this decision rule can be used in any situation where the different types of errors do not have the same importance. Some examples of obtained results are shown on synthetic data mimicking reality and with increasing complexity. Finally, applications on real data show promising results.**

***Index Terms*—Belief functions, confidence degrees, Dempster–Shafer method, discounting factors, humanitarian mine detection, sensor fusion, mass assignment, sensor clustering.**

I. Introduction

DESPITE the great efforts and motivation of research teams around the world, there is no single sensor used for humanitarian mine detection that can reach the necessarily high detection rate in all possible scenarios. As a result, a very attractive approach to finding a solution is in taking the best from several complementary sensors. One of the most promising sensor combinations consists of an imaging metal detector (MD), a ground penetrating radar (GPR), and an infrared camera (IR). Here we propose a method of combination that can be easily adapted for other sensors and their combinations.

Since reliability and detection capabilities of any sensor are strongly scenario-dependent, it is important to characterize each of the sensors under combination. In other words, the ways for modeling the influence of various factors on sensors and on results of their combination have to be investigated, with the aim of obtaining fusion results as good as possible.

Manuscript received January 11, 2002; revised July 7, 2002 and November 26, 2002. This paper was recommended by Associate Editor E. Santos.

N. Milisavljević is with the Signal and Image Centre-Royal Military Academy, Av. de la Renaissance 30, 1000 Brussels, Belgium, and also with the Ecole Nationale Supérieure des Télécommunications-TSI-CNRS URA 820, 75013 Paris, France (e-mail: nada@elec.rma.ac.be).

I. Bloch is with the Ecole Nationale Supérieure des Télécommunications-TSI-CNRS URA 820, 75013 Paris, France (e-mail: Isabelle.Bloch@enst.fr).

Digital Object Identifier 10.1109/TSMCC.2003.814034

The choice of an adequate combination method has to be carefully done. The level of fusion mainly depends on the types and similarity of sensors to be fused. For example, fusion is most often done at a pixel-level if the data to be fused are images acquired by the same camera but in different spectral bands [1], [2], or by passive IR and millimeter-wave cameras [3]. However, there are also some attempts to fuse data from similar sensors (or even the same one) at a higher, feature [4] or decision [5] level. On the other hand, fusion of dissimilar sensors, as it is the case here, has to be done at a higher-level, not because of different physical phenomena they detect, but because of a difference in resolution, problems of registration, etc. Note that only the literature on fusion of mine detection sensors is reviewed here, but a lot of work is also done on each sensor separately, as shown in [6]. It is well-known that there is no universal approach for information fusion and that the choice of a particular one strongly depends on the problem itself [7]–[9]. Most of the efforts made in the field of fusion of dissimilar mine detection sensors are based on statistical approaches [10]–[12]. They provide excellent results for a particular scenario, but ignore, or briefly mention that once more general solutions are looked for, several important problems have to be faced in this domain of application [13]. Namely, the data has the following characteristics:

1) they are not numerous enough to allow for a reliable statistical learning, as shown in [14], [15]; for instance, to establish that the probability of detection is 99.6% (set by the United Nations as the rate at which a mine-free area can be considered safe) with a confidence of 98.1%, more than 1000 samples of each type of mine are needed;
2) they are highly variable depending on the context and conditions;
3) they do not give precise information on the type of mine (ambiguity between several types).

In addition, it is impossible to model every object (neither mines nor objects that could be confused with them).

Therefore, we propose an approach based on belief functions (BFs) in the framework of Dempster–Shafer (DS) theory [16], [17], since ignorance, uncertainty and ambiguity can be appropriately modeled in this framework. Although the requirement of 99.6% detection probability is still often mentioned, it is quite controversial and has been suffering from various critics in the demining community with respect to its justification and how realistic it is. As a result, it disappeared recently from the mine action standards, which now states“no error” until some depth is determined in function of the scenario, future land use, etc. There is no probabilistic interpretation behind this standard, hence, using a nonprobabilistic approach, such as belief function, does not raise any inconsistency. To our knowledge, there are just two attempts for applying DS theory to this problem

1094-6977/03$17.00 © 2003 IEEE

[18], [19]. Both works treat an alarm as a mine, and not as an object that could be a mine, as well as a false alarm. In addition, in [18], beliefs are obtained classically, as probabilities [20], so this method faces the same problem as statistical ones: the lack of a sufficient amount of training data. In [19], the efficiency of the proposed method is estimated using real data containing a few mines and no placed false alarms, hence, the obtained results do not show how the method deals with false alarms.

The lack of training data is partly compensated by available knowledge on mine detection sensors, on well-known mine laying principles, on mines themselves, as well as on objects that each of these sensors confuses with mines. Consequently, our main motivation for exploring possibilities of modeling the mine detection problem in the framework of BFs (with its nonprobabilistic interpretation) is to be able to easily include and model this partial knowledge by exploiting different features of BF theory without any statistical information on the data or on noise. In this domain of application, we have not met any similar approach yet.

Another important aspect of the work presented here is linked to the fact that mistakes are not allowed in case of humanitarian mine detection. In our opinion, it is not possible to reach the highest possible level of detection no matter which fusion method is chosen, simply because it is not possible to predict everything in all real situations where mines can be found. Because of that, our idea is to give to the deminer as much information as possible, starting from processed data of separate sensors up to the possible conclusions, but the final decision has to be left to the deminer. Accordingly, the result of this DS model should be an ordered list of guesses what a currently observed object could be, together with the confidence in these results. Similar ideas of leaving the final decision to the deminer can be found in the literature [18], but the approach presented here goes one step further. Namely, we allow the deminer to influence the modeling and the combination process. Indeed, his knowledge, experience as well as some of his observations that simply cannot be detected by any sensor are precious, and have to be included in the reasoning process in order to obtain significant improvements in detection and false alarm rates.

For the considered sensors, there is no criterion which allows to tell that the object is a mine, it can be just the opposite. Consequently, our results show how expectable it is that an object is not a mine, or that it is either a mine or something else. This point should not be understood as a drawback of the method. It should not be forgotten that mines must not be missed, so detecting that something is not a mine and that it is a mine or something else seems to be the safest approach to this complex problem.

The contribution of this paper is twofold: first, a methodological contribution in the sense that the paper proposes formal models for dealing with important problems in fusion, such as fusion of sensors having different reliabilities, or providing information on different physical objects, which have a more general impact than the treated application; secondly, an adaptive contribution in the field of humanitarian mine detection. In the following, based on characteristics of the three analyzed mine detection sensors and on general directions for applying the DS approach, explained above, the appropriate choice of measures and of respective mass assignment for each of the sensors is discussed. The identity of the object under observation is analyzed at two levels. The idea behind the two-level approach is that it allows to simplify the problem since the sensors behave differently depending on the metal content and, therefore, the pieces of knowledge and information to be modeled are different. The general scheme of this two-level approach is described in Section II. The choice between three types of object corresponding to different degrees of metallic content is done at the first level, explained in Section III. At the second level, discussed in Section IV, the object is analyzed in more detail, taking into account the specificity of the chosen type, with the final aim of classifying the object into a mine or a friendly (nondangerous) object. The masses assigned by the chosen measures are combined using Dempster's rule [16], [17], [21] in unnormalized form, which has the advantage of providing a natural measure of conflict. Two main reasons can explain the conflict: either some sources of information are not completely reliable, or they do not refer to the same object. The first aspect is solved using discounting factors [16], [21], [22], and the second by combining only pieces of information referring to the same object. An original decision rule is proposed, in order to overcome limits of classical rules for this type of problems. Finally, the presented method is illustrated on some examples in Section V and on real data in Section VI.

Although the proposed approach was initially dedicated to a specific application, some questions are raised that are more general, and the methodological solutions we propose to answer these questions can be used in many other applications. This is the case for the two-level approach, for the two ways of solving conflict, for the decision rule, as will be seen mainly in Sections II and IV.

II. Two-Level Approach

The two-level approach is motivated by the characteristics of the sensors and their different behavior depending on the metal content of the observed object. These characteristics are briefly summarized here (more details can be found in [23]). A standard GPR detects any buried object as long as its dielectric constant differs from the one of the soil, IR detects thermal contrast between the object and the soil, while MD detects metal. GPR and IR are strongly influenced by various environmental factors, such as moisture, temperature of the soil, time of the day, etc., and some of these factors are difficult (even impossible) to quantify. The behavior of MD is quite independent of the environment, as long as the soil is not ferrous. Therefore, the large variety of objects that can be detected can be classified on the basis of their metal content. Furthermore, in the scope of humanitarian demining, mines are classified in three types [24]: metallic, low-metal content, and nonmetallic. In this classification, a metallic mine is the one made almost completely of metal, except e.g., handles, while a mine with low metal content has metal only in small parts, e.g., in the fuse. The same type of classification of objects is adopted in this paper, so we subdivide objects into MO (metallic object), LMO (low-metal content object), and NMO (nonmetallic object) types. GPR and IR responses are not significantly influenced by the metal content of an object. On the contrary, the information that can be extracted from the MD regarding the true nature of an object

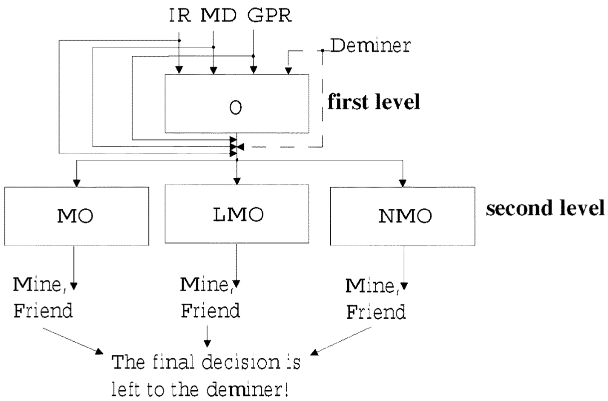

Fig. 1. Two-level approach.

(whether it is a mine or a friendly object) importantly changes between MO, LMO, and NMO cases. That is why we decide to split our analysis into two levels (see Fig. 1).

At the first level, we want to determine whether an object is MO, LMO, or NMO, hence, which kind of information can be expected from MD regarding the true nature of the object. This classification could be done based on MD alone, but we cannot know for sure how much we can rely on it within some scenario, so at this point, the responses of the other two sensors have to be taken into account as well. This is detailed in Section III.

Then, at the second level, the object is further analyzed (this analysis depends on the metal content) in order to decide about the final object identity (a mine or a friendly object). In case of ambiguity between different types, several possibilities among MO, LMO, and NMO are investigated. As an output of this level, an ordered list of guesses about the true object identity is given, together with confidence degrees. The second level is discussed in detail in Section IV.

The idea of having two levels comes from the fact that one type of information induces different types of processing of the other pieces of information depending on its values. Two levels are indeed sufficient if only one type of information leads to different processing, as is the case in this application. Although this idea came from the specificities of the metal detector, it should be noticed that it can be applied to any problem where we have such situations. It is general in this sense. Extensions to more than two levels are possible. For instance, if after the first branching issued from the analysis of this particular information, it appears that another piece of information leads also to several ways of reasoning or to several processing branches, then another level should be added, and so forth. This can be seen as a process similar to decision trees, but in the framework of belief functions.

III. First Level—The Choice Between MO, LMO, and NMO

A. Description

All three sensors considered here give images, so the measures that can tell the most about the absolute and/or relative amount of metal that an object contains are the following:

1) area of the object on the MD image;
2) strength of the MD response (ratio between the maximum value of pixels in the image and the image scale);
3) ratio between the object area on the MD image and on the GPR image (ignored if GPR area is 0);
4) ratio between the object area on the MD image and on the IR image (ignored if IR area is 0).

The interest of the first measure is obvious, but the others deserve some justification. If the area of the object on an MD image is small, the strength of the response could still be high in case of a high metal content object which has, e.g., a long cylindrical shape, of which the top is seen in the image. However, if the strength of the MD response is not high, the area could still be large, indicating that something might be wrong with the amplification of the signal. Then, if the area of the object in the MD image is large and the strength of the MD response is high, this area could still be much smaller than the one seen by GPR or IR, meaning that the percentage of metal is small.

Based on these four measures, we model the problem related to the first level in the BF framework. These measures provide beliefs about the object with respect to the three classes, MO, LMO, and NMO. Note that this assumes that the classes are crisp, which corresponds to the usual way of reasoning in humanitarian demining. Each measure gives rise to a mass function having LMO, NMO, and MO as focal elements. No other disjunction is taken into account in the present version of the model.[1] The full set, $\Theta_1 = (\mathrm{LMO} \cup \mathrm{NMO} \cup \mathrm{MO})$, is taken as a focal element only for the two comparison measures to account for cases in which the area on the MD image is larger than the area in the IR, or the GPR. The numerical representation of mass functions assumes that we can assign numbers that represent degrees of belief. The general shapes and tendencies are derived from the knowledge we have and its modeling. There certainly remain some arbitrary choices, which might appear as a drawback of the method, however, it is not necessary to have precise estimations of these values, and a good robustness is observed experimentally. This can be explained by two reasons: first, the representations are used for rough information, hence do not have to be precise themselves, and secondly, several pieces of information are combined in the whole reasoning process, which decreases the influence of each particular value (of individual information). Therefore, the chosen numbers are not crucial. What is important is that ranking is preserved, as well as the shape of the functions, and these are derived from knowledge. These comments apply to all functions proposed for mass assignments in this paper.

Masses are assigned, as illustrated in Figs. 2 and 3. Namely, if there is no response of the MD, or if the detected area is very small, then the highest mass assigned by the MD area measure (Fig. 2) should be assigned to NMO. If the area is quite large, it is almost certain that the object is MO. For some moderate areas, it can be expected that the object is LMO. The exact range of areas corresponding to each type of objects depends on the particular situation at hand, the types of mines that can be expected, etc. The reasoning is similar for modeling the strength of MD

[1]This simple model was sufficient in our experiments. More complex models where soft boundaries between classes are taken into account could also be developed, but they are not useful in our context.

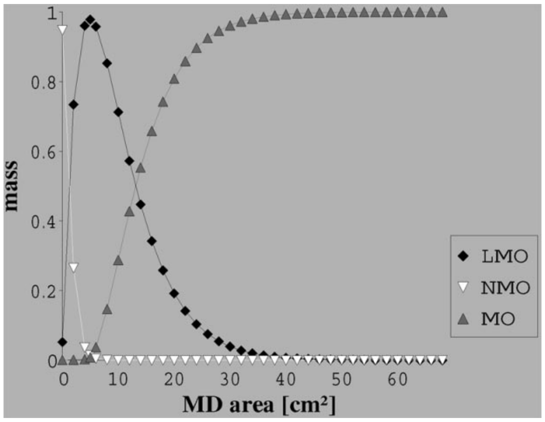

Fig. 2. Masses assigned by MD area measure.

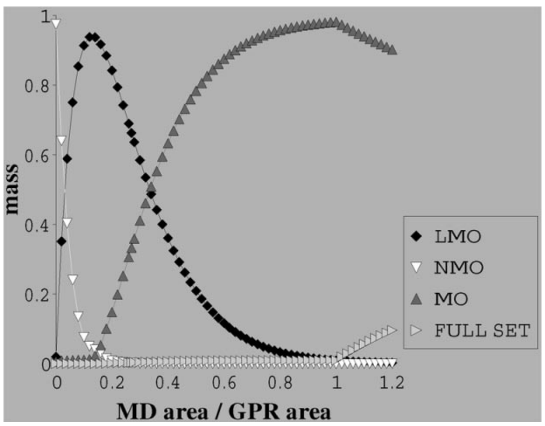

Fig. 3. Assigned masses by comparing MD and GPR areas.

response. Ideas behind modeling the two area comparison measures (shown in Fig. 3 for MD and GPR) are the following:

1) if MD area is negligible when compared to GPR (or IR) area, such an object might be NMO;
2) if MD area is significantly smaller than GPR (or IR) area, the object is likely LMO;
3) if MD and GPR (or IR) areas are similar, the object is MO;
4) if MD area is quite larger than the other two, ignorance about the type of object is high, so the mass should be mainly assigned to Θ_1.

The deminer is included in the reasoning process by expressing his opinion regarding the sensors and, accordingly, the importance of each of the four measures via discounting factors [16], [21], [22]. Since there are only four measures at the first level that all involve either MD alone or comparison of MD data with GPR or IR data, the deminer's opinion regarding the sensors and regarding the importance of measures can be treated as one factor for each measure. The deminer expresses his opinion about the importance of each measure i through a number f_i, on a scale $f_{\rm scale}$ that he predefines. The highest number corresponds to the highest importance of the measure, inducing the smallest discounting. These factors modify masses as follows:

1) for any subset $A \neq \Theta_1$, and any measure i, new masses, $m_i(A)$, are computed from the initial ones, $m_{i\,\rm IN}(A)$, as

$$m_i(A) = \frac{f_i}{f_{\rm scale}} \cdot m_{i\,\rm IN}(A); \tag{1}$$

2) for the full set:

$$m_i(\Theta_1) = 1 - \frac{f_i}{f_{\rm scale}} \cdot [1 - m_{i\,\rm IN}(\Theta_1)]. \tag{2}$$

Masses assigned by the four measures are combined by Dempster's rule in unnormalized form [16], [17], [21]

$$m(A) = \sum_{\substack{k,l \\ A_k \cap B_l = A.}} m_i(A_k) \cdot m_j(B_l) \tag{3}$$

where m_i and m_j are masses assigned by measures i and j after discounting, and their focal elements are $A_1, A_2, \ldots A_m$ and $B_1, B_2, \ldots B_n$, respectively. Since Dempster's rule is commutative and associative, it can be applied repeatedly, until all measures are combined, and the result is independent of the order used in the combination. The measures under combination have to fulfill the condition of being cognitively independent [21], which is weaker than the notion of statistical independence required in the probabilistic case. Measures are cognitively independent if each of them assigns masses without any direct link to the others. This is the case in the proposed model.

If there is no intersection between some focal elements A_k and B_l, the resulting mass is assigned to the empty set. Since the only disjunction among focal elements that is taken into account at this level is the full set, the mass of the empty set indicates that one measure gives masses to one singleton, and some other measure to another singleton. A high mass on the empty set corresponds to a high conflict among measures, and therefore to more difficulty to distinguish between the singletons.

Finally, a good indicator for deciding which path to follow at the second level (see Fig. 1) is the mass of the empty set. If there is a singleton $S \in \Theta_1$ which has a higher mass than the others and than the empty set, i.e.,

$$\exists S \in \Theta_1, \quad S \neq \emptyset : m(S) = \max_{A \in \{\rm NMO, LMO, MO, \emptyset\}} m(A) \tag{4}$$

the path corresponding to S should be followed. Otherwise, we should group singletons. For each of the three possible combinations, masses assigned to corresponding subgroups by measure i are calculated from the masses given to each of the singletons of the subgroup separately

$$m_{i\,\rm NEW}(A \cup B) = m_i(A) + m_i(B), \tag{5}$$

where A and B are singletons that are grouped together. After calculating these new masses, measures are combined using Dempster's rule (3). Then, if for some of these three combinations the mass of the empty set is lower than the other two resulting masses and if the masses of the two focal elements are well-distinguished, the path with the higher mass has to be followed. If the focal element with the higher mass consists of two

TABLE I
STARTING ANALYSIS

	m(MO)	m(LMO)	m(NMO)	$m(\Theta_1)$	$m(\emptyset)$
aM	0.81	0.19	5e-9	0	0
rM	0.64	0.36	8.7e-5	0	0
aMG	0.93	6.2e-2	7.9e-8	9.4e-3	0
aMI	0.96	0.03	7.8e-9	9.7e-3	0
C	0.47	1.9e-4	4e-17	0	0.53

TABLE II
FIRST CASE OF GROUPING, $\{\mathrm{LMO}, \mathrm{MO}\}$ AND $\{\mathrm{NMO}\}$

	$m(\mathrm{MO} \cup \mathrm{LMO})$	m(NMO)	$m(\Theta_1)$	$m(\emptyset)$
aM	1	4.9e-9	0	0
rM	1	8.7e-5	0	0
aMG	0.99	7.9e-8	9.4e-3	0
aMI	0.99	7.8e-9	9.7e-3	0
C	1	3.9e-17	0	8.6e-5

TABLE III
SECOND CASE OF GROUPING, $\{\mathrm{LMO}, \mathrm{NMO}\}$ AND $\{\mathrm{MO}\}$

	m(MO)	$m(\mathrm{LMO} \cup \mathrm{NMO})$	$m(\Theta_1)$	$m(\emptyset)$
aM	0.81	0.19	0	0
rM	0.64	0.36	0	0
aMG	0.93	6.3e-2	9.4e-3	0
aMI	0.96	0.03	9.7e-3	0
C	0.47	1.9e-4	0	0.53

singletons, then these two paths should be followed in parallel. If the empty set has a significant mass in all analyzed cases, then all three paths should be followed in parallel. This idea is illustrated in the following example.

B. Example of Ambiguity Between the Three Types of Objects

It can happen that the data collected by the three sensors result in ambiguity between the three types of objects. A way to deal with such situations is illustrated using the following values for each of the four measures:

1) MD area is 20 cm^2;
2) strength of MD response is 0.4;
3) ratio of MD and GPR areas is 0.7;
4) ratio of MD and IR areas is 0.8.

assuming that the deminer gives equal and maximum importance to each measure (no discounting).

1) Starting Analysis: $\{\mathrm{LMO}\}, \{\mathrm{MO}\}$, *and* $\{\mathrm{NMO}\}$*:* The focal elements are $\{\mathrm{LMO}\}, \{\mathrm{MO}\}$, and $\{\mathrm{NMO}\}$. Masses are assigned for each measure[2] as given in Table I, where measure aM is MD area, rM—MD response, aMG—ratio of MD and GPR areas, and aMI—ratio of MD and IR areas. In the last row of this table, named C, the masses resulting from the combination of the measures by Dempster's rule (3) are shown. The highest mass is assigned to the empty set, indicating that the conflict between measures is too high. Therefore, we group the singletons.

2) First Case of Grouping: $\{\mathrm{LMO}, \mathrm{MO}\}$ *and* $\{\mathrm{NMO}\}$*:* Here, the focal elements are $\{\mathrm{LMO}, \mathrm{MO}\}$ and $\{\mathrm{NMO}\}$. For every measure i, the new mass of subset $\{\mathrm{LMO}, \mathrm{MO}\}$ can be found from (5). Masses assigned for each of the measures are shown in Table II, as well as the masses after combination. In this case, the focal elements are very well distinguished, and the resulting mass of the empty set is very low.

3) Second Case of Grouping: $\{\mathrm{LMO}, \mathrm{NMO}\}$ *and* $\{\mathrm{MO}\}$*:* This time, the focal elements are $\{\mathrm{LMO}, \mathrm{NMO}\}$ and $\{\mathrm{MO}\}$. The new masses of the subset $\{\mathrm{LMO}, \mathrm{NMO}\}$ are calculated using (5) and given in Table III. They do not change much when compared with the starting case, since NMO does not have significant masses. Therefore, we can expect similar combination results as for the starting case. Indeed, the resulting masses, given in the last row of Table III, are the same as the ones shown in Table I.

4) Third Case of Grouping: $\{\mathrm{MO}, \mathrm{NMO}\}$ *and* $\{\mathrm{LMO}\}$*:* Logically, it does not make sense to group MO and NMO together, and LMO separately, but for other applications, where unsupervised clustering may be required, all cluster possibilities should be checked. Nevertheless, as in the previous subsection, this case does not differ significantly from the starting one, since very low masses are assigned to NMO.

C. Final Remarks Regarding the First Level

In the example given above, the lowest mass of the empty set is obtained when LMO and MO are grouped together, so these two paths should be followed.

Note that the values for each measure are chosen so that it is not sure at all if it could be a LMO or MO, so obtained results are in accordance with the analyzed situation. Also, it shows the power of this grouping and of keeping masses unnormalized: it is easy to check that if at the starting case, Dempster's rule in normalized form were applied, the mass of MO would be very high when compared with the other two, so that this case would be chosen, i.e., LMO case would be discarded.

There are two ways to deal with ambiguous situations such as the one shown here, when more than one path is chosen to be followed, i.e., 1) to follow all the chosen paths in parallel, or 2) to have additional paths for ambiguities. Both ways introduce additional problems. The first one calls for a particular procedure for merging the results, while the second means that for each of the possible combinations of paths, a set of measures and adequate modeling have to be carefully chosen. Here, we choose the first solution for the two following reasons; 1) the

[2]Masses are rounded on two digits, and small values are not truncated but preserved in order to avoid multiplications by zero during combination.

TABLE IV
THE FOCAL ELEMENTS IN THE FRAME OF DISCERNMENT AT THE SECOND LEVEL.

	GPR	IR	MD, MO	MD, LMO
area	$(FR \cup FI)$, Θ_2	$(FR \cup FI)$, Θ_2	$(MFR \cup MFI)$, Θ_2	/
elongation	$(MR \cup FR)$, $(MI \cup FI)$, Θ_2	$(MR \cup FR)$, $(MI \cup FI)$, Θ_2	$(MMR \cup MFR)$, $(MMI \cup MFI)$, Θ_2	/
ellipticity	$(MR \cup FR)$, $(MI \cup FI)$, Θ_2	$(MR \cup FR)$, $(MI \cup FI)$, Θ_2	$(MMR \cup MFR)$, $(MMI \cup MFI)$, Θ_2	/
burial depth	$(FR \cup FI)$, Θ_2	/	$(FR \cup FI)$, Θ_2	$(FR \cup FI)$, Θ_2
height	$(FR \cup FI)$, Θ_2	/	/	/
burial depth comparison	/	/	/	$(FR \cup FI)$, Θ_2

problem of merging the results can be simplified by the fact that the real aim is to distinguish between mines and friendly objects (giving much more importance to mines, since they must not be missed); this means that as long as the final choices represent a mine (nonmetallic, metallic, circular, asymmetric, ...), they should receive the highest importance, 2) modeling of additional paths for ambiguities will be much more natural once MO, LMO, and NMO are treated not as crisp, but as fuzzy classes, which is not addressed here.

IV. SECOND LEVEL—CLASSIFICATION INTO MINES OR FRIENDLY OBJECTS

The analysis continues in (at least) one of the three directions, LMO, NMO, or MO. Taking into account that all three sensors give images, that around 90% of the mines have an elliptical (regular) top surface (circular, but seen as elliptical under some burial angle), and that the major goal of our humanitarian demining efforts is to distinguish between a mine and a nondangerous (friendly) object (stones, cans, etc.), the frame of discernment is defined for each case as follows:

1) for MO: $\Theta_{2\mathrm{MO}} = \{\mathrm{MMR}, \mathrm{MMI}, \mathrm{MFR}, \mathrm{MFI}\}$, where MMR denotes metallic mine of regular shape, MMI metallic mine of irregular shape, MFR metallic friendly object of regular shape, and MFI metallic friendly object of irregular shape;
2) for LMO: $\Theta_{2\mathrm{LMO}} = \{\mathrm{LMMR}, \mathrm{LMMI}, \mathrm{LMFR}, \mathrm{LMFI}\}$, with: LMMR (low-metal content mine of regular shape), LMMI (low-metal content mine of irregular shape), etc.;
3) for NMO: $\Theta_{2\mathrm{NMO}} = \{\mathrm{NMMR}, \mathrm{NMMI}, \mathrm{NMFR}, \mathrm{NMFI}\}$.

Some measures are the same for all three cases, so we use simplified notations sometimes, such as

1) FR for denoting any friend (regardless its metallic content) of regular shape (either MFR, LMFR or NMFR, depending whether MO, LMO, or NMO is analyzed);
2) FI for any friend of irregular shape (MFI, LMFI or NMFI);
3) MR for any mine of regular shape (MMR, LMMR or NMMR);
4) MI is any mine of irregular shape (MMI, LMMI or NMMI);
5) Θ_2 for any frame of the discernment at this level.

A. Choice of Measures and Modeling the Resulting Masses

While the measures of MD are function of the metal content of the observed object, the measures of IR and GPR remain the same. These measures are (Table IV)

1) for both GPR and IR:
   i) area, assigning masses to $(\mathrm{FR} \cup \mathrm{FI})$ and the full set, due to the fact that whenever the area is too small or too large, it is quite certain that such an object is not a mine, while whenever the area is within some range corresponding to expected sizes of mines, the object can be a mine or anything else as well;
2) elongation, giving masses to $(\mathrm{MR} \cup \mathrm{FR}), (\mathrm{MI} \cup \mathrm{FI})$ and the full set;
3) ellipticity, that also assigns masses to $(\mathrm{MR} \cup \mathrm{FR}), (\mathrm{MI} \cup \mathrm{FI})$, and Θ_2;
4) for GPR alone:
   i) burial depth, assigning masses to $(\mathrm{FR} \cup \mathrm{FI})$ and Θ_2;
   ii) depth dimension (height) of an object, again giving masses to $(\mathrm{FR} \cup \mathrm{FI})$, and the full set.

Regarding MD, there is no measure when the object is NMO. For the other two cases, its measures are

1) both for LMO and for MO:
   i) burial depth or comparison of burial depth with GPR (which one is taken into account depends on whether this sensor disagrees or agrees with GPR, respectively), assigning masses to $(\mathrm{FR} \cup \mathrm{FI})$ and the full set;
2) for MO alone—same as for both GPR and IR:
   i) area, assigning masses to $(\mathrm{MFR} \cup \mathrm{MFI})$ and the full set;
   ii) elongation, giving masses to $(\mathrm{MMR} \cup \mathrm{MFR}), (\mathrm{MMI} \cup \mathrm{MFI})$ and the full set;
   iii) ellipticity, that assigns masses to $(\mathrm{MMR} \cup \mathrm{MFR}), (\mathrm{MMI} \cup \mathrm{MFI})$ and the full set too.

In the following, we present a model for each of the measures, based on the bibliographic survey as well as on the single-sensor trials within the Belgian HUDEM project. Note that the proposed equations for mass assignments are only examples illustrating the required tendencies. The numbers and parameters in these mass functions are not really important. What is important is the general shape of these functions, as mentioned also in Section III-A.

1) Ellipse Fitting Mass Assignment: We apply an ellipse fitting algorithm using the randomized Hough transform [25], [26] on thresholded images of the sensors. The mass for the subset of regular shapes assigned by this measure indicates how well this shape fits an ellipse:

$$m_f(\mathrm{MR} \cup \mathrm{FR}) = \max\left(0, \min\left\{\frac{A_{\mathrm{oe}}-5}{A_o}, \frac{A_{\mathrm{oe}}-5}{A_e}\right\}\right) \quad (6)$$

where A_{oe} is the part of object area that belongs to the fitted ellipse as well, A_o is the object area, A_e is the ellipse area. The subtraction of 5 pixels is introduced to include the limit case of an ellipse with just 5 pixels, where we cannot judge about the shape at all, so ignorance should be maximum. In case of MD, this equation makes sense only for MO case.

The mass of irregular subset is the larger value of two values, the percentage of ellipse area that does not belong to the object and the percentage of object area that does not belong to the fitted ellipse:

$$m_f(\mathrm{MI} \cup \mathrm{FI}) = \max\left\{\frac{A_e - A_{\mathrm{oe}}}{A_e}, \frac{A_o - A_{\mathrm{oe}}}{A_o}\right\}. \quad (7)$$

Again, this equation is used regardless the metal content for IR and GPR, while for MD it makes sense only when the object is MO.

The full set gets the remaining mass:

$$m_f(\Theta_2) = 1 - m_f(\mathrm{MR} \cup \mathrm{FR}) - m_f(\mathrm{MI} \cup \mathrm{FI}). \quad (8)$$

2) Elongation Mass Assignment: Here, a thresholded image of the object is used as input too. First, the center of gravity of the image is calculated, and then we compute the following quantities; 1) minimum and maximum distance of bordering pixels from the center of gravity, and the ratio between them (ratio1) (if the center of gravity is not inside the boundary, this does not make sense, and ratio1 is set to 0), and 2) second moments, and from them the ratio of minor and major axis of the obtained quadratic form (ratio2). The mass value for the regular subset is the smaller of the two ratios:

$$m_e(\mathrm{MR} \cup \mathrm{FR}) = \min(\mathrm{ratio1}, \mathrm{ratio2}) \quad (9)$$

while the mass of the irregular subset is the absolute value of their difference

$$m_e(\mathrm{MI} \cup \mathrm{FI}) = |\mathrm{ratio1} - \mathrm{ratio2}|. \quad (10)$$

The full set takes the rest

$$m_e(\Theta_2) = 1 - \max(\mathrm{ratio1}, \mathrm{ratio2}). \quad (11)$$

Note that the method could be applied to edges as well, but the influence of noise on the position of the center of gravity would be much larger.

3) Area/Size Mass Assignment: When a preliminary information about the expected size of mines is available, a range of areas of detected object that could be a mine, but something else as well, can always be predicted, taking into account possible deformations because of some burial angle. Outside that range, it is much more expectable that objects are friendly. Therefore, masses are modeled as:

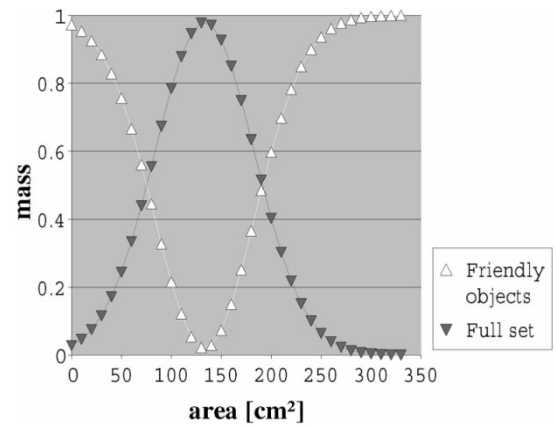

Fig. 4. Illustration of mass assignments for area/size.

$$m_a(\Theta_2) = \lambda \cdot \exp\frac{[A_o - 0.5 \cdot (A_{\min} + A_{\max})]^2}{0.5 \cdot (A_{\max} - A_{\min})^2}, \quad (12)$$

$$m_a(\mathrm{FR} \cup \mathrm{FI}) = 1 - m_a(\Theta_2) \quad (13)$$

where the approximate range of expectable mine areas is from $A_{\min}$ to $A_{\max}$, as illustrated in Fig. 4 for $A_{\min} = 80$ cm^2 and $A_{\max} = 180$ cm^2. The parameter λ is usually set to a value closed to 1 (here we took 0.98). Taking it not exactly equal to 1 allows to have non zero values for the complement function, which avoids to exclude completely solutions (since mass functions are combined using product operators, having zero somewhere is very strong).

4) Burial Depth Mass Assignment: This measure is useful in case of GPR regardless the metallic content of the object. On the contrary, the depth comparison measure has to be ignored if GPR and MD disagree about the depth position of the object, because it means that they do not refer to the same object. If that is the case, in order not to lose the depth information extracted by MD, its own burial depth measure has to be introduced, modeled in the same way as for GPR.

Similarly to area/size of mines, there is a range of burial depths where it is more expectable that mines can be buried. When a GPR starts to sense something on these small depths, mass should be assigned to the full set, since it can be either a mine or something else. At higher depths, it is far more likely that the detected object is something else but mine. Corresponding masses are (see Fig. 5):

$$m_d(\Theta_2) = \frac{\lambda}{\cosh(\mu \cdot \mathrm{depth})^2}, \quad (14)$$

$$m_d(\mathrm{FR} \cup \mathrm{FI}) = 1 - m_d(\Theta_2). \quad (15)$$

The parameter λ has the same meaning as before. The parameter μ defines the speed of decreasingness of the function. It depends on the expected range of the depth, and has to be tune for each application, depending on the historical context, the type of soil, the type of mines expected in the concerned area, etc. (here we have chosen $\mu = 0.04$).

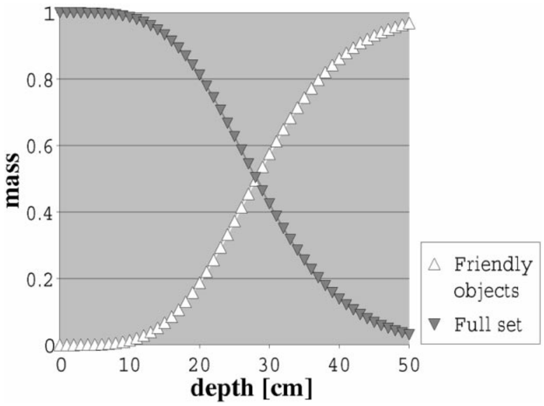

Fig. 5. Example of mass assignments for depth.

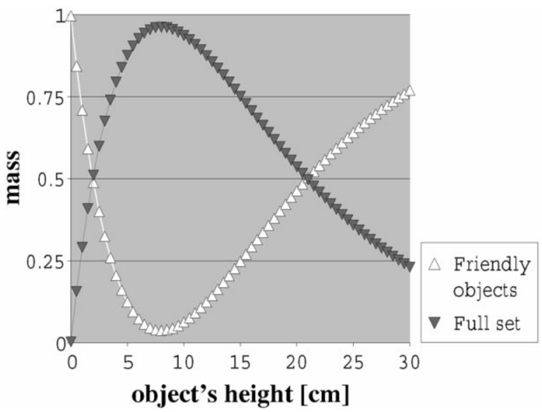

Fig. 6. Masses assigned for depth dimension.

5) Depth Dimension Mass Assignment: Once again, some range of depth dimensions or heights of object exists, where it is more expectable that an object is a mine (but it can be something else as well), while some too small or too large objects are quite surely nondangerous, as shown in Fig. 6.

6) Mass Assignment for Comparison of GPR and MD Depth Information: If there is an agreement between the depth dimension of the object detected by GPR, and the depth position of metal detected by MD, the object can be anything from our frame of discernment. Otherwise, the object should not be a mine. Accordingly, masses are assigned as follows (where x is the depth position detected by MD, measured from the top level detected by GPR, and d is the height of the object detected by GPR):

1) in case of LMO:

i) for $x \in [0, d]$:

$$m_c(\Theta_2) = \lambda, \tag{16}$$

$$m_c(\mathrm{LMFR} \cup \mathrm{LMFI}) = 1 - \lambda; \tag{17}$$

ii) otherwise:

$$m_c(\Theta_2) = \frac{\lambda}{\cosh\left[0.75 \cdot \left(\frac{2 \cdot x}{d} - 1\right)\right]}, \tag{18}$$

$$m_c(\mathrm{LMFR} \cup \mathrm{LMFI}) = 1 - m_c(\Theta_2); \tag{19}$$

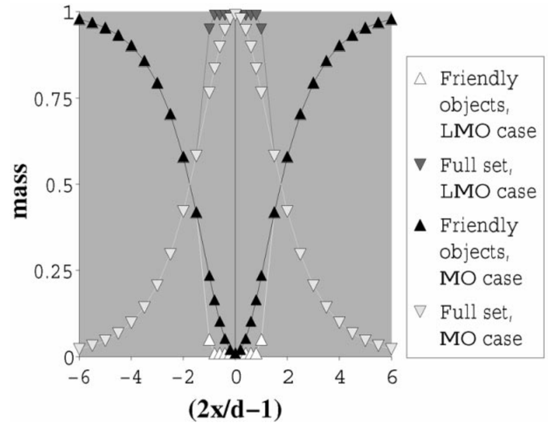

Fig. 7. Assigned masses for MD and GPR depth comparison.

2) in case of MO:

$$m_c(\Theta_2) = \frac{\lambda}{\cosh\left[0.75 \cdot \left(\frac{2 \cdot x}{d} - 1\right)\right]}, \tag{20}$$

$$m_c(\mathrm{MFR} \cup \mathrm{MFI}) = 1 - m_c(\Theta_2). \tag{21}$$

The mass assignment for LMO differs from the one in case of MO within the GPR region because the small metallic piece can be anywhere within the interval detected by GPR. In case of MO, the maximum value is given to the mass of the full set when the position of metal detected by MD is in the middle of the depth interval extracted by GPR, i.e., when $x = d/2$ (see Fig. 7).

B. Discounting Factors

The behavior of each of the three sensors is strongly scenario-dependent, referring to

1) the quality of the acquired data;
2) the reliability of each of the sensors under particular weather conditions, type of soil, etc.;
3) the types of objects under analysis.

These are the main reasons for including discounting factors [16], [21], [22] in the model, and they become very important at the second level. They are presented in detail in [23] for objects with high metal content, so we just briefly describe them here.

Discounting factors consist of three types of parameters:

1) g_{ij}: confidence level of sensor j in its assessment when judging measure i (varying from 0 for not confident at all to 1 for completely confident);
2) b_{ij}: level of importance of measure i of sensor j, varying from 1 to b_{scale}, where b_{scale} is the scale for b parameters;
3) s_j: deminer's confidence into opinion of sensor j, varying from 1 to s_{scale}, where s_{scale} is the scale for s parameters.

Discounting and combination do not commute, so it has to be pointed out when each of the discounting parameters is used. Firstly, g_{ij} and $(b_{ij}/b_{\mathrm{scale}})$ are used to discount masses assigned by measure i of sensor j, and that is done for all measures of that sensor. Then we combine the measures per sensor. After that, the resulting masses are discounted using (s_j/s_{scale}) parameters, before combining the sensors.

Discounting of masses is done as given by (1) and (2), where (f_i/f_{scale}) is replaced by either $g_{ij}, (b_{ij}/b_{\mathrm{scale}})$ or (s_j/s_{scale}).

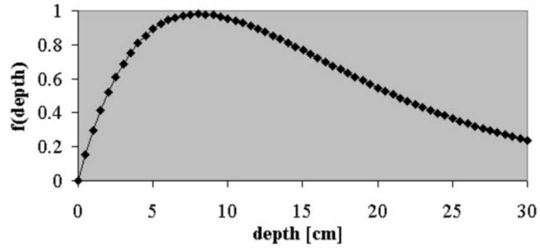

Fig. 8. Confidence in GPR in function of estimated depth.

1) Calculating g_{ij}:

a) Area/Size: Shape is much more sensitive than area to data corruption. Besides, as long as area is moderately corrupted, it does not affect mass assignments significantly. Therefore, it is not needed to apply discounting to this measure, so we do not do it for IR and MD. Still, since GPR does not detect well shallowly buried or very deeply buried objects, we model this confidence level as a function of the extracted depth information (Fig. 8)

$$g_{aG} = f(\text{depth}). \tag{22}$$

b) Ellipse fitting: The confidence of a sensor in masses assigned by ellipse fitting measure is a function of the shape itself. The larger the mass difference between regular and irregular sets, the larger the confidence in the assessment

$$g_{fI} = |m_{fI}(\text{MR} \cup \text{FR}) - m_{fI}(\text{MI} \cup \text{FI})| \tag{23}$$

$$g_{fG} = |m_{fG}(\text{MR} \cup \text{FR}) - m_{fG}(\text{MI} \cup \text{FI})| \tag{24}$$

$$g_{fM} = |m_{fM}(\text{MMR} \cup \text{MFR}) - m_{fM}(\text{MMI} \cup \text{MFI})|. \tag{25}$$

c) Elongation: Similarly, the confidence level in estimation of elongation is

$$g_{eI} = |m_{eI}(\text{MR} \cup \text{FR}) - m_{eI}(\text{MI} \cup \text{FI})| \tag{26}$$

$$g_{eG} = |m_{eG}(\text{MR} \cup \text{FR}) - m_{eG}(\text{MI} \cup \text{FI})| \tag{27}$$

$$g_{eM} = |m_{eM}(\text{MMR} \cup \text{MFR}) - m_{eM}(\text{MMI} \cup \text{MFI})|. \tag{28}$$

d) Comparison of GPR and MD depth information (g_{cG}): This measure is not discounted since it can indicate that the two sensors should be analyzed separately.

e) GPR depth dimension, g_{hG}: This confidence level is defined as a function of the extracted depth (Fig. 8)

$$g_{hG} = f(\text{depth}). \tag{29}$$

f) Burial depth: In case of GPR, this factor also depends on the burial depth itself

$$g_{dG} = f(\text{depth}). \tag{30}$$

We do not discount this factor in case of MD since its reliability is quite independent of environmental conditions.

2) How to Estimate s_j?: The deminer's confidence in a sensor j depends on factors that affect reliability of that sensor, such as environmental conditions, time of the day, moisture, etc. His confidence might be also biased by his trust in each of the sensors.

The idea is that, for each sensor, the deminer gives a number describing his belief in its reliability, where the higher the number, the larger the confidence. The values he gives have to be rescaled, so he has to provide his scale as well (as detailed in [23]).

3) Estimating b_{ij}: There are several open possibilities for these parameters:

1) their choice can be again left to the deminer (how important to him each of the measures is);
2) their values can be preset for each of the predictable cases (e.g., MO, LMO, etc.);
3) they can be at the beginning chosen all the same, and after combination of masses, their values can be tuned depending on which subset has the highest mass [23].

Because of lack of additional information, we adopt the first possibility in this paper.

It should be noted that this way of using discounting factors can be applied to any situation where we have some gradual confidence in the measures and where the user has to give his own input. This happens in many domains, as for instance in medical imaging in diagnosis support systems.

C. Combination

The combination is performed using Dempster's rule without normalization (3). The mass assigned to the empty set expresses the level of conflict between the pieces of information that have been combined. Two main causes of the conflict can be advocated: either there are several objects and the sensors do not provide information on the same object, or some sources of information are not completely reliable. The possibility that sensors do not refer to the same object is addressed here. It is a major improvement over our initial work [23], [27]. Unreliability is modeled as discounting factors, the choice of which has been addressed in Section IV-B.

After assigning masses for each measure of each sensor, a first check is performed, to see whether GPR and MD agree about the depth position of the object:

1) if that is the case, we check if that depth is reachable by IR (because the detectability of IR decreases quite significantly with depth):
   i) if it is, we continue and analyze all the sensors together;
   ii) if it is not, IR is analyzed separately;
2) if that is not the case, we check which one of them detects an object placed closer to the surface, and check whether it is close enough to the surface to be seen by IR;
   i) if it is, IR is grouped together with that sensor;
   ii) if it is not, all sensors refer to different objects and have to be analyzed separately.

Then, the fusion of measures per sensor is done (3), and internal conflicts of each sensor are analyzed. There is no MD measure for NMO, and for LMO just one exists, so for these two types of objects only internal conflicts of GPR and IR can be analyzed.

If the internal conflict of a sensor is high, its data are corrupted by noise, occlusion etc., so they should be discounted. We cannot know for sure which part of information is more affected by this corruption than others. Still, discounting is performed on each measure, since it is modeled in such a way that its influence is proportional to the level of ambiguity and uncertainty of the information, as shown in the previous subsection.

After discounting the sensors that have high internal conflicts, the fusion of sensors (plus the measure on the depth information comparison in cases of LMO and MO) is performed, according to the way the sensors were clustered at the beginning. Note that performing the combination in two steps (first for all information extracted from one sensor and then between sensors) is consistent because of the associativity of Dempster's rule: if there is no discounting, the result is the same as for direct global combination. What still can happen is that even if the sensors were not clustered into subgroups, the conflict after fusion is high. In other words, although it seemed quite possible that they refer to the same object, they actually do not. For example, although IR could reach some depths of 5 cm, where GPR and MD detect something, it actually detects something nonmetallic on the surface. In that case, clustering again takes place, but this time in an unsupervised manner [28], toward finding the combination of the sensors where the conflict among the sensors is the lowest and goes down to the algebraic sum of internal conflicts of the sensors. To illustrate the idea, let i and j be two sensors. The mass of the empty set resulting from their combination (3) is

$$\begin{aligned} m_{ij}(\emptyset) &= \sum_{A\cap B=\emptyset} m_i(A)m_j(B) \\ &= m_i(\emptyset)\sum_B m_j(B) + m_j(\emptyset)\sum_A m_i(A) \\ &\quad - m_i(\emptyset)m_j(\emptyset) + \sum_{A\cap B=\emptyset, A\neq\emptyset, B\neq\emptyset} m_i(A)m_j(B) \\ &= m_i(\emptyset) + m_j(\emptyset) - m_i(\emptyset)m_j(\emptyset) \\ &\quad + \sum_{A\cap B=\emptyset, A\neq\emptyset, B\neq\emptyset} m_i(A)m_j(B) \end{aligned} \tag{31}$$

where the first three terms correspond to the algebraic sum of the internal conflicts. The last term is equal to 0 if and only if all focal elements of source i are consonant with all focal elements of sources j (i.e., they have never an empty intersection).

D. Decision

After clustering and combination of sensors (3), the resulting masses are found for each of the clusters. On the basis of these masses, final conclusion about the true identity of each object under analysis has to be made for every cluster. Usual decision rules rely on beliefs *Bel*, plausibilities *Pl* [16] and pignistic probabilities P [29], defined as follows:

for any subset B:

$$\mathrm{Bel}(B) = \sum_{A\subset B, A\neq\emptyset} m(A), \tag{32}$$

$$\mathrm{Pl}(B) = \sum_{B\cap A\neq\emptyset} m(A), \tag{33}$$

for any singleton C:

$$P(C) = \sum_{A\in 2^{\Theta}, C\in A} \frac{m(A)}{|A|[1-m(\emptyset)]}. \tag{34}$$

Unfortunately, a decision rule based on any of these three functions does not lead to fruitful conclusions, since there are no focal elements containing mines alone [23], [27]. Namely, after combination of the measures at the second level, the resulting focal elements are: $\mathrm{FR}, \mathrm{FI}, (\mathrm{FR}\cup\mathrm{FI}), (\mathrm{MR}\cup\mathrm{FR}), (\mathrm{MI}\cup\mathrm{FI})$ and Θ_2. Beliefs and plausibilities are[3]:

$$\mathrm{Bel}(M) = 0 \tag{35}$$

$$\mathrm{Bel}(F) = m(\mathrm{FR}) + m(\mathrm{FI}) + m(\mathrm{FR}\cup\mathrm{FI}) \tag{36}$$

$$\mathrm{Pl}(M) = m(\mathrm{MR}\cup\mathrm{FR}) + m(\mathrm{MI}\cup\mathrm{FI}) + m(\Theta_2) \tag{37}$$

$$\mathrm{Pl}(F) = 1 \tag{38}$$

where M denotes mines ($M = \mathrm{MR}\cup\mathrm{MI}$) and F denotes friendly objects ($F = \mathrm{FR}\cup\mathrm{FI}$). Pignistic probabilities (34) are computed only for singletons, thus we have to derive probabilities for M and F using standard probabilistic equations. For MR and MI, we find

$$P(\mathrm{MR}) = \frac{m(\mathrm{MR}\cup\mathrm{FR})}{2} + \frac{m(\Theta_2)}{4} \tag{39}$$

$$P(\mathrm{MI}) = \frac{m(\mathrm{MI}\cup\mathrm{FI})}{2} + \frac{m(\Theta_2)}{4}. \tag{40}$$

Since $\mathrm{MR}\cap\mathrm{MI} = \emptyset$, the pignistic probability for M is equal to the sum of $P(\mathrm{MR})$ and $P(\mathrm{MI})$:

$$P(M) = \frac{1}{2}[m(\mathrm{MR}\cup\mathrm{FR}) + m(\mathrm{MI}\cup\mathrm{FI}) + m(\Theta_2)]. \tag{41}$$

The pignistic probability for F is found in the similar way:

$$\begin{aligned} P(\mathrm{FR}) &= m(\mathrm{FR}) + \frac{m(\mathrm{FR}\cup\mathrm{FI})}{2} \\ &\quad + \frac{m(\mathrm{MR}\cup\mathrm{FR})}{2} + \frac{m(\Theta_2)}{4} \end{aligned} \tag{42}$$

$$\begin{aligned} P(\mathrm{FI}) &= m(\mathrm{FI}) + \frac{m(\mathrm{FR}\cup\mathrm{FI})}{2} \\ &\quad + \frac{m(\mathrm{MI}\cup\mathrm{FI})}{2} + \frac{m(\Theta_2)}{4} \end{aligned} \tag{43}$$

hence

$$\begin{aligned} P(F) &= m(\mathrm{FR}) + m(\mathrm{FI}) + m(\mathrm{FR}\cup\mathrm{FI}) \\ &\quad + \frac{1}{2}[m(\mathrm{MR}\cup\mathrm{FR}) + m(\mathrm{MI}\cup\mathrm{FI}) + m(\Theta_2)]. \end{aligned} \tag{44}$$

Note that the following relations are always true:

$$\mathrm{Bel}(M) \le \mathrm{Bel}(F) \tag{45}$$

$$\mathrm{Pl}(M) \le \mathrm{Pl}(F) \tag{46}$$

$$P(M) \le P(F) \tag{47}$$

so the decision would always be made in favor of F.

Since in case of any ambiguity, far more importance has to be given to mines, we propose to define guesses, $G(A)$, where $A \in \{M, F, \emptyset\}$, in the following way:

$$G(M) = \sum_{M\cap B\neq\emptyset} m(B) \tag{48}$$

[3]For the sake of simplicity, we assume that masses are normalized in this illustration.

$$G(F) = \sum_{B \subseteq F, B \neq \emptyset} m(B) \tag{49}$$

$$G(\emptyset) = m(\emptyset). \tag{50}$$

The guess value of a mine is the sum of masses of all the focal elements containing mines, no matter whether they are regular or irregular, the guess of a friend is the sum of masses of all the focal elements containing nothing else but friends (again, either regular or irregular), and the guess of something else is equal to the mass of the empty set. Note that the guess of a mine is equal to its plausibility, while the guess of a friend is equal to its belief, which reflects the fact that we have to be cautious in deciding F. In other words, the introduced guesses are but a cautious way for estimating confidence degrees.

Once guesses for the three types of objects are found, they are ordered and this list, together with confidence degrees, is given as the final result.

If more than one path is followed at the second level, each of the paths results in one ordered list of guesses, for each of its clusters. The problem of merging the results does not have a unique solution. The one that we adopt in this paper is as follows:

1) if all the followed paths result in the same number of clusters, and sensors are clustered in the same way, for each of the clusters, the final guesses are the ones obtained in the path where the guess of a mine is the highest. Let $G_i(M)$ be the guess of M in path i. We choose the path m such that

$$G_m(M) = \max_i G_i(M). \tag{51}$$

   The final guesses are $\{G_m(M), G_m(F), G_m(\emptyset)\}$. This set is then ordered and provided as the final result;
2) if the number of clusters is not the same for all paths, we choose the one with the highest number of clusters as the final result (since it is a result of the finest analysis, and since it is safer to claim, in case of ambiguity, that there is more than one object although there could be actually only one);
3) if the number of clusters for all the paths is the same, but sensors are not clustered in the same way, most likely something went wrong, so the safest is to claim complete ignorance and serve results of analysis of each of the sensors separately.

The decision rule we proposed here can be applied in any situation where the classes have not the same importance, and where the classification errors do not have the same cost depending on their type. For instance, our decision rule applies to any problem where non detection is worse than false alarm (as here), or the contrary. Again this may happen in many problems, like medical diagnosis.

V. First Results

At this time, we do not have real data on which all the features of the proposed approach can be illustrated. We have results on real data, as presented in [30], but only a part of the model applied in that case, due to some particularities of the analyzed data set (they will be presented in Section VI). Since the aim of this paper is to propose a complete model dealing with a large range of cases (that actually occur in reality), we have chosen to illustrate all aspects of the method on synthetic experiments. These experiments are drawn from real observations (measurements) of each of the sensors separately.

For the sake of simplicity, we give similar examples for MO, LMO and NMO cases, clearly stating what applies for each of them.

Only the second level is illustrated here, since the number of interesting combinations of the two levels is very large. The first level has been illustrated in Section III.B. These experiments are ordered by increasing difficulty, in order to introduce progressively the types of problems our approach is able to solve. The first example discusses an ideal case for this problem, i.e., the object is a mine, and it is seen correctly and equally by all three sensors. The next case is more complex: each sensor senses a different physical object, and all objects differ in shape, allowing for a possibility to cluster the sensors. This case leads directly to the third, final example, which is the most difficult one, discussing what happens if the sensors see different objects but of a similar shape. Here we only deal with the classification step. We assume that the preliminary steps of detection of objects in suspect areas is already done. The noise of sensors is involved in these steps only, and is therefore not subject to tests in this work. We only assume that we have objects on which measurements are performed in order to classify them into friend or mine.

A. Ideal Example

We can imagine the situation where an object is buried at a depth where mines can be expected (5 cm), its area is similar to the one of mines (120 cm^2), its depth dimension is as for smaller mines (5 cm). Three different cases are analyzed in parallel, i.e., that this object is:

1) an elliptical MO, seen approximately equally by all three sensors, and MD and GPR agree about its depth position;
2) or an elliptical LMO, seen approximately equally by IR and GPR, while MD detects a small object, of which position is in agreement with depth interval detected by GPR;
3) or an elliptical NMO, seen approximately equally by GPR and IR.

The masses are given in Table V and no discounting is included yet. The measures are indicated according to notations introduced in Section IV.B, while we note $m_R = m(\mathrm{MR} \cup \mathrm{FR}), m_I = m(\mathrm{MI} \cup \mathrm{FI})$, and $m_F = m(\mathrm{FR} \cup \mathrm{FI})$. For the MO case, all measures should be taken into account, for the LMO case the last three should be ignored, and for NMO case the last four should not be taken into account.

The first check, performed on measures extracted from each of the sensors, shows that MD and GPR agree about the depth, and that this depth can be reached by IR. Therefore, the following step consists in combining the measures per sensor. The obtained masses are given in Table VI (the last column is valid just for the MO case).

Since the internal conflicts of the sensors are not significant, there is no need for their discounting, and in the next step, we fuse:

TABLE V
EXAMPLE 1, MASSES ASSIGNED PER MEASURE BEFORE ANY CHECKING AND DISCOUNTING.

	m_R	m_I	m_F	$m(\Theta_2)$
eG	0.72	0.1	0	0.18
fG	0.96	3.9e-2	0	1e-3
aG	0	0	0.06	0.94
dG	0	0	8.7e-4	1
hG	0	0	0.12	0.88
eI	0.72	0.1	0	0.18
fI	0.96	3.9e-2	0	1e-3
aI	0	0	0.06	0.94
cG	0	0	0.01	0.99
eM	0.72	0.1	0	0.18
fM	0.96	3.9e-2	0	1e-3
aM	0	0	0.06	0.94

TABLE VI
RESULTING MASSES FOR EXAMPLE 1, SEPARATE SENSORS.

	GPR	IR	MD
m(FR)	0.14	0.05	0.05
m(FI)	1.8e-3	6e-4	6e-4
m_F	3e-5	9.7e-6	9.7e-6
m_R	0.71	0.82	0.82
m_I	9.1e-3	1e-2	1e-2
$m(\Theta_2)$	1.5e-4	1.7e-4	1.7e-4
$m(\emptyset)$	0.13	0.12	0.12

1) for MO, the three sensors as well as the remaining measure on depth comparison between MD and GPR;
2) in the case of LMO, IR, and GPR as well as the measure on depth comparison between MD and GPR;
3) in the case of NMO, IR, and GPR.

Table VII contains the resulting masses. The resulting conflict does not exceed the algebraic sum of the internal conflicts of the sensors, meaning that these sensors indeed refer to the same object. The first guess is that this object is either MR or FR. This is in accordance with the envisaged situation. Note that the resulting masses for LMO and NMO cases are almost identical, due to the fact that the only differing measure between the two cases, i.e., depth comparison of MD and GPR, is almost completely ignorant about the identity of the object.

Guesses (48), (49), and (50) are similar for all the paths (see Table VIII), and the first guess is that the object is a mine.

TABLE VII
RESULTING MASSES FOR EXAMPLE 1, AFTER SENSOR FUSION

	MO	LMO	NMO
m(FR)	0.17	0.16	0.16
m(FI)	3.7e-7	2.7e-5	2.6e-5
m_F	1.5e-12	7.1e-9	6.8e-9
m_R	0.47	0.58	0.58
m_I	1e-6	9.7e-5	9.8e-5
$m(\Theta_2)$	4.3e-12	2.5e-8	2.5e-8
$m(\emptyset)$	0.36	0.26	0.26

TABLE VIII
GUESSES FOR EXAMPLE 1

	MO	LMO	NMO
G(M)	0.47	0.58	0.58
G(F)	0.17	0.16	0.16
$G(\emptyset)$	0.36	0.26	0.26

B. Example Where Shapes Differ

Here, GPR sees some object from 5 cm to 10 cm of depth. In LMO and MO case, MD detects something at the depth of 30 cm, and for MO case, that is a rectangular shape with area of 10 cm^2. GPR detects an elliptical shape of area equal to 120 cm^2, while IR senses an X-like shape the area of which is 10 cm^2. At first, masses are assigned as given in Table IX.

A strong discrepancy in depth information between MD and GPR is detected during the first check, so these two sensors are grouped separately. Since the GPR response is closer to the soil surface, IR is grouped with it. As a result, we have two clusters, one containing IR and GPR, and another containing MD. Thus, the measure on MD and GPR depth comparison has to be ignored and the burial depth measure is introduced for MD. Measures of this sensor in case of MO assign masses as given in Table X, of which only burial depth measure remains in case of LMO. Therefore, this measure is the only source of information when judging about the identity of this LMO object (see the first column of Table XII). It tells that this object has a metallic part on 30 cm of depth, and the first guess is that it is a friendly object. In MO case, the fusion of MD measures gives results as shown in the second column of Table XII. The internal conflict of MD in this case is high, so the masses are discounted (Table XI). Their fusion leads to the results given in the last column of Table XII. Calculated guesses, given in Table XV, show that this object is most likely friendly.

Regarding IR and GPR cluster, we firstly fuse measures per sensor, as shown in the first two columns of Table XIII. Measures of IR have a relatively strong internal conflict, so we discount the IR measures (Table XIV) and then fuse them again (the third column of Table XIII). The next step is to fuse discounted

TABLE IX
STARTING ASSIGNED MASSES, EXAMPLE 2

	m_R	m_I	m_F	$m(\Theta_2)$
eG	0.72	0.1	0	0.18
fG	0.96	3.9e-2	0	1e-3
aG	0	0	5.4e-2	0.94
dG	0	0	8.7e-4	0.99
hG	0	0	0.12	0.88
eI	0.16	0.69	0	0.15
fI	0.42	0.58	0	5.1e-4
aI	0	0	0.953	4.7e-2
cG	0	0	0.99	0.01
eM	0.64	0.19	0	0.17
fM	0.67	0.33	0	7.3e-4
aM	0	0	0.953	4.7e-2
dM	0	0	*0.57*	*0.43*

TABLE X
MASSES ASSIGNED PER MEASURE FOR MD, EXAMPLE 2, NO DISCOUNTING

	m_R	m_I	m_F	$m(\Theta_2)$
dM	0	0	0.57	0.43
eM	0.64	0.19	0	0.17
fM	0.67	0.33	0	7.3e-4
aM	0	0	0.953	4.7e-2

TABLE XI
DISCOUNTED MASSES FROM TABLE X

	m_R	m_I	m_F	$m(\Theta_2)$
dM	0	0	0.57	0.43
eM	0.29	8.5e-2	0	0.63
fM	0.23	0.11	0	0.66
aM	0	0	0.953	4.7e-2

IR and nondiscounted GPR, and the result of this combination is given in the last column of Table XIII. The mass of the empty set (0.5) is much higher than the algebraic sum of internal conflicts of the two sensors (0.164), hence these sensors do not refer to the same object. As a final result, we conclude that the three sensors refer to three different objects:

- GPR object: 1) MR, 2) FR, 3) MI, 4) FI, 5) something else;
- MD object: 1) F (R or I), 2) M (R or I), 3) something else;
- IR object: 1) FI, 2) FR, 3) MI, 4) MR, 5) something else,

TABLE XII
RESULTING MASSES FOR MD, EXAMPLE 2

	MD, LMO	MD, MO	MD, MO, disc.
m(FR)	0	0.53	0.39
m(FI)	0	0.12	0.13
m_F	0.57	1.2e-4	0.41
m_R	0	1.1e-2	8.1e-3
m_I	0	2.4e-3	2.7e-3
$m(\Theta_2)$	0.43	2.5e-6	8.4e-3
$m(\emptyset)$	0	0.34	4.6e-2

TABLE XIII
RESULTING MASSES FOR GPR AND IR, EXAMPLE 2

	GPR	IR	IR, disc.	IR+GPR
m(FR)	0.14	0.12	0.11	0.47
m(FI)	1.8e-3	0.46	0.37	9e-3
m_F	3e-5	7.3e-5	0.44	7.9e-5
m_R	0.71	6.1e-3	5.4e-3	1.9e-2
m_I	9.1e-3	2.3e-2	1.8e-2	3.6e-4
$m(\Theta_2)$	1.5e-4	3.6e-6	2.2e-2	3.2e-6
$m(\emptyset)$	0.13	0.38	3.9e-2	0.5

or, simpler:

- GPR: 1) mine, 2) friendly object, 3) something else;
- MD: 1) friend, 2) mine, 3) something else;
- IR: 1) friendly object, 2) mine, 3) something else.

Obtained results are in agreement with the imagined situation. The importance of introducing a burial depth measure for MD if it is in disagreement with GPR can be seen: it is the simplest and safest way for treating such cases when MD detects something different than GPR and IR, it is the only source of information in case of LMO, and it can give a complementary information about the true object identity in MO case.

The usefulness of providing confidence values (guesses, Table XV) together with the ordered lists is also shown here: for MD, the difference between first and second guesses for LMO path is much narrower than in the case of IR. Without confidence values these two cases seem the same, which could be a dangerous mistake.

This example proves that any available measure, even not directly involved in the final decision, can provide useful information toward disambiguating some situations. For example, the focal elements for ellipticity and elongation are: $\mathrm{MR} \cup \mathrm{FR}$, $\mathrm{MI} \cup \mathrm{FI}$ and Θ_2, meaning that the two shape measures cannot distinguish between mines and friendly objects. Nevertheless, they help detecting the conflict between sensors, hence determining

TABLE XIV
MASSES ASSIGNED PER IR MEASURE AFTER DISCOUNTING, EXAMPLE 2.

	m_R	m_I	m_F	$m(\Theta_2)$
eI	0.09	0.36	0	0.55
fI	6.6e-2	9.1e-2	0	0.84
aI	0	0	0.953	4.7e-2

TABLE XV
GUESSES FOR EXAMPLE 2.

	GPR	MD, LMO	MD, MO final	IR, final
G(M)	0.72	0.43	0.02	0.04
G(F)	0.15	0.57	0.93	0.92
$G(\emptyset)$	0.13	0	0.05	0.04

that the sensors do not refer to the same object. Namely, if the shape measures are not included in this example, the following GPR and IR measures remain (see Table IX): GPR area, depth and height, and IR area. For any of them, the focal elements are $\mathrm{FR} \cup \mathrm{FI}$ and Θ_2, thus there is no conflict resulting from their combination. In other words, without the shape information, we would not be aware that GPR and IR do not detect the same object, that, again, could have serious consequences in humanitarian demining reality.

C. Example With Several Objects of a Similar Shape

If everything remains the same as in the previous case, except that shapes and areas seen by GPR and IR are similar, it would not be possible to split these two sensors, so it would not be possible to notice that they do not refer to the same object! Unfortunately, this is not a drawback of the model, but of the choice of sensors, so there is nothing that can be done to avoid this misinterpretation. The only way to cope with this problem is that in case that IR and GPR are grouped together and MD separately, we claim that the exact position of the object seen by IR and GPR is somewhere between the soil surface and the lowest surface level detected by GPR. This makes sense since GPR does not see objects on the soil surface and just below it, so we cannot be sure whether the upper level detected by GPR is really the highest level of the object.

VI. PRELIMINARY RESULTS ON REAL DATA

In this section, we present results that have been obtained on real data, provided by TNO Physics and Electronics Laboratory (The Hague, The Netherlands), within the Dutch HOM-2000 project. These data include IR, GPR and MD images, obtained on a sand lane containing 21 mines and 7 friendly objects. After the processing of each type of data, 42 regions are obtained, 28 corresponding to regions containing the actual objects, and 14 for which clutter produced alarms. This means that finally we have to recognize 21 mines and 21 false alarms.

On these data, MD measures are unfortunately not sufficient to allow to use the first level of our approach. Therefore only the second level could be applied. Moreover, shape and area measures are not available on GPR and MD data in these experiments due to a coarse data acquisition step in one of the surface dimensions. However, the results still show the interest of our approach, and the improvement obtained by fusion.

When using A-scan and C-scans measures, the following results are obtained:

1) 19 mines are detected, and two are missed (due to the sensitivity of the used sensors, and not due to the method proposed here);
2) Six placed false alarms are correctly recognized, and 1 is wrongly classified as a mine;
3) Eight clutter alarms are classified as F, and six are wrongly classified as mines.

If we compare these results to the ones obtained on each sensor separately, significant improvement can be observed:

1) from IR only we get six more non detected mines and one more false alarm,
2) from MD only we get three more false alarms and one more non detected mine,
3) from GPR only we get four more false alarms.

This shows that the fusion using the proposed approach allows to improve the mine detection rate, while decreasing the false alarm rate.

When using B-scan measures, the results are exactly the same for the mines and for the false alarms. But the number of false alarms due to objects increases while the one for clutter decreases, keeping the global false alarm rate constant.

VII. CONCLUSION

We presented a method for modeling and fusion of mine detection sensors in terms of belief functions within the Dempster-Shafer framework. A two-level approach is introduced. At the first level, the object under analysis is classified according to its metal content. Based on that classification, at the second level the chosen type is analyzed in detail, with the goal of determining whether the object is a mine or a nondangerous, friendly object. Measures that can be extracted from each of the sensors are presented and modeled. Since importance of each measure, the confidence of the sensors in their assessments regarding each of the measures, as well as the deminer's opinion about the reliability of each of the sensors depend strongly on the scenario, discounting factors are included in the model in order to account for these parameters. Guess functions are introduced as a way of making decisions in this extremely sensitive and dangerous problem of humanitarian demining. It avoids problems encountered with classical decision rules for this application. Several examples are given, based on synthetic data, showing that the proposed model is promising.

The two-level approach described here is a logical basis to explain the complexity of the mine detection problem. The presented theory is general, and can be easily adjusted depending on specificities of particular data sets. This will be tested in detail in our future work, when the proposed approach will be applied on additional sets of real data. First results on some real data have already been obtained [30], but only a part of the model applied in that case.

Furthermore, the paper proposes formal models for dealing with important problems in fusion, such as different reliability among the sensors and possibility that different sensors provide information on different physical objects, which have a more general impact than the humanitarian mine detection application.

Note that there are possibilities to integrate the two levels in one, but such an approach has some drawbacks [31]. In future, we will work on comparing the two approaches with the aim of improving them as well as of developing a model which takes the best of each of them.

Several aspects of our model apply more generally, such as the multi-level approach, in case a piece of information induces different processing (in a decision-tree like manner), the two ways to deal with conflict (discounting and clustering), and the decision rule which allows to take into account the possible different importance of each type of error.

References

[1] S. L. Earp and J. M. Dang, "Maximum likelihood multispectral pixel-level fusion using a linear-quadratic detector," in *Proc, SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 3079, Orlando, FL, 1997, pp. 232–248.

[2] D. Mishra, A. K. Chan, and C. K. Chui, "Histogram equalization, image registration and data fusion for multispectral images," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 2496, Orlando, FL, 1995, pp. 1025–1031.

[3] F. C. Nivelle and P. Lhomme, "Detection of land mines with passive IR and mmW imaging sensors," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 3079, Orlando, FL, 1997, pp. 614–624.

[4] G. A. Clark *et al.*, "Multispectral image fusion for detecting land mines," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 2496, Orlando, FL, 1995, pp. 850–864.

[5] P. D. Gader *et al.*, "Fuzzy logic detection of landmines with ground penetrating radar," *Signal Processing*, vol. 80, pp. 1069–1084, 2000.

[6] N. Milisavljević, "State of the Art of the Sensors for Mine Detection—A Study," Signal and Image Centre, Royal Military Academy, Brussels, Belgium, Tech. Rep., 1998.

[7] I. Bloch, "Incertitude, imprécision et additivité en fusion de données: Point de vue historique," *Traitement du Signal*, vol. 13, no. 4, pp. 267–288, 1996.

[8] ——, "Information combination operators for data fusion: A comparative review with classification," *IEEE Trans. Syst., Man, Cybern., A*, vol. 26, pp. 52–67, Jan. 1996.

[9] P. Smets, "Varieties of ignorance and the need for well-founded theories," *Inform. Sci.*, no. 57–58, pp. 135–144, 1991.

[10] P. Jägerbro *et al.*, "Combination of GPR and metal detector for mine detection," in *Proc. IEE 2nd Int. Conf. Detection Abandoned Landmines*, Edinburgh, UK, 1998, pp. 177–181.

[11] P. Gao, S. Tantum, and L. Collins, "Single sensor processing and sensor fusion of GPR and EMI data for landmine detection," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 3710, Orlando, FL, 1999, pp. 1139–1148.

[12] A. Gunatilaka and B. A. Baertlein, "Comparison of pre-detection and post-detection fusion for mine detection," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 3710, Orlando, FL, 1999, pp. 1212–1223.

[13] N. Milisavljević and M. Acheroy, "An approach to the use of the bayesian rule in decision level fusion for multisensor mine detection," in *Proc. Physics Signal Image Processing (PSIP'99) Conf.*, Paris, France, 1999, pp. 261–266.

[14] M. Roughan and D. McMichael, "Data Fusion for Land-Mine Detection: Data Requirements for Experimental Design," Cooperative Research Centre for Sensor, Signal and Information Processing, Adelaide, Australia, Tech. Rep. LSOD, DSTO, 1996.

[15] R. Voles, "Confidence in the assessment and use of mine detection systems," *Proceedings of IEE Second International Conference on the Detection of Abandoned Landmines*, pp. 28–30, 1998.

[16] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton Univ. Press, 1976.

[17] P. Smets, "What is Dempster-Shafer's model?," in *Advances in the Dempster-Shafer Theory of Evidence*, R. R. Yager, M. Fedrizzi, and J. Kacprzyk, Eds. New York: Wiley, 1994, pp. 5–34.

[18] E. den Breejen, K. Schutte, and F. Cremer, "Sensor fusion for anti personnel landmine detection, a case study," in *Proc. SPIE Conf. Detection Technologies Mines Minelike Targets*, vol. 3710, Orlando, FL, 1999, pp. 1235–1245.

[19] S. Perrin, "Contribution à l'algorithmique multicapteur pour la détection de mines antipersonnel," Ph.D. Thesis, Ecole Centrale de Lille, USTL, France, 2001.

[20] T. Denoeux, "Reasoning With Imprecise Belief Structures," Université de Technologie de Compiègne, France, Tech. Rep. Heudiasys 97/44, 1997.

[21] P. Smets, "Belief functions: The disjunctive rule of combination and the generalized Bayesian theorem," *Int. J. Approx. Reas.*, no. 9, pp. 1–35, 1993.

[22] D. Dubois, M. Grabisch, H. Prade, and P. Smets, "Assessing the value of a candidate," in *15th Conf. Uncertainty Artificial Intelligence (UAI'99)*. San Mateo, CA, 1999, pp. 170–177.

[23] N. Milisavljević, I. Bloch, and M. Acheroy, "Characterization of mine detection sensors in terms of belief functions and their fusion, first results," in *Proc. 3rd Int. Conf. Information Fusion (FUSION 2000)*, vol. II, Paris, France, 2000, pp. ThC3.15–ThC3.22.

[24] E. Banks, Ed., *Antipersonnel Landmines—Recognising and Disarming*, London-Washington: Brassey's, 1997.

[25] R. A. McLaughlin, "Randomized Hough Transform: Improved Ellipse Detection With Comparison," Univ. Western Australia, CIIPS, Tech. Rep. TR97-01, 1997.

[26] N. Milisavljević, "Comparison of three methods for shape recognition in the case of mine detection," *Pattern Recognit. Lett.*, vol. 20, no. 11–13, pp. 1079–1083, 1999.

[27] N. Milisavljević, I. Bloch, and M. Acheroy, "A first step toward modeling and combining mine detection sensors within Dempster-Shafer framework," in *Proccedings of 2000 International Conference on Artificial Intelligence (IC-AI'2000)*, vol. II, Las Vegas, NV, 2000, pp. 745–751.

[28] J. Schubert, "On nonspecific evidence," *Int. J. Intell. Syst.*, vol. 8, pp. 711–725, 1993.

[29] P. Smets, "Constructing the pignistic probability function in a context of uncertainty," in *Uncertainty in Artificial Intelligence 5*, L. N. Kanal, M. Henrion, R. D. Shachter, and J. F. Lemmer, Eds. New York: Elsevier, 1990, pp. 29–39.

[30] N. Milisavljević, S. P. van den Broek, I. Bloch, P. B. W. Schwering, H. A. Lensen, and M. Acheroy, "Comparison of belief functions and voting method for fusion of mine detection sensors," in *Proc. SPIE Conf. Detection Remediation Technologies Mines Minelike Targets VI*, vol. 4394, Orlando, FL, 2001, pp. 1011–1022.

[31] N. Milisavljević and I. Bloch, "A two-level approach for modeling and fusion of humanitarian mine detection sensors within the belief function framework," in *Proc. 10th Int. Symp. Applied Stochastic Models Data Analysis*, vol. 2, France, 2001, pp. 743–748.

Nada Milisavljević received the B.E.E. degree from the University of Novi Sad, Yugoslavia, in 1992, the M.Sc. degree in optical electronics and laser technics from the University of Belgrade, Yugoslavia, in 1996, and the Ph.D. degree from Ecole Nationale Supérieure des Télécommunications (ENST) in Paris, France, in 2001.

Currenly, she is a Researcher at the Signal and Image Centre of the RMA, Brussels, Belgium. Her current research interests include data fusion, evidence theory, humanitarian mine detection and pattern recognition.

Isabelle Bloch received the degree from Ecole des Mines de Paris, France, in engineering and mathematics in 1986, and the Ph.D. degree from ENST Paris in 1990, and the "Habilitation Diriger des Recherches" from University Paris 5 in 1995.

Currenly, she is Professor at ENST, Signal and Image Processing Department. Her research interests include 3-D image and object processing, 3D and fuzzy mathematical morphology, discrete 3-D geometry and topology, decision theory, information fusion in image processing, fuzzy set theory, evidence theory, structural pattern recognition, spatial reasoning, and medical imaging.