Why is Information Fusion Useful for Mined Area Reduction?

Nada Milisavljević
Department of Communication Information Systems & Sensors (CISS)
Royal Military Academy, Brussels, Belgium

Isabelle Bloch
Télécom ParisTech
CNRS LTCI, Paris, France

1   Introduction

Antipersonnel landmines are munitions that are victim-activated and that do not discriminate: they injure or kill whoever triggers them. In addition, they remain active for decades after the end of a conflict during which they were placed. They can be found in many places where civilians carry out their daily activities, such as on roads, agricultural fields, in forests, around and inside various buildings, hence inhibiting freedom of movement, denying access to food, water, etc. According to the Mine Ban Treaty, mine-affected states are required to clear all antipersonnel landmines from mined areas under their jurisdiction or control within 10 years of becoming party to the Mine Ban Treaty. Although first deadlines expired on 1 March 2009, 24 States Parties failed to meet them and were granted extensions, according to the Landmine Monitor of August 2009 (ICBL, 2009). As stated in the Monitor, antipersonnel landmines still affect more than 70 states as well as seven areas not internationally recognized. Thanks to the Mine Ban Treaty, mine clearing operations have been organized in a more controlled and effective way, yet mine clearance remains a slow and resource demanding process. It is estimated that, on average, a deminer clears 10 m2 during a working day using conventional tools such as metal detectors and prodders. Thus, humanitarian mine clearance operations must be understood and designed correctly, providing efficient aid to innocent people who may be severely injured by this dreadful pollution. The recommendations made during the Standing Committee on Mine Clearance, Mine Risk Education and Mine Action Technologies state that: 1) technologists should avoid building technologies based on assumed needs and should work interactively with end-users, 2) appropriate technologies could save human lives and increase mine action efficiency, and 3) nothing is more important than understanding the working environment (Acheroy, M., 2003). Besides the very long time needed to clear polluted terrain, actual demining campaigns show that the false alarm rate is very large, the threat of plastic mines (which cannot be detected by metal detectors) is not negligible and the variety of mine clearance scenarios is high, depending on the country, the region, the climate, etc. These facts prove that the humanitarian mine detection is a very complex and sensitive problem. In addition, the experience shows that it will be a long process to achieve a mine-free world, so the concept of a mine-free world is evolving softly toward the concept of a mine impact-free world, although a mine-free world remains the final goal of the Mine Ban Treaty.

Although the Mine Ban Treaty was signed eleven years ago, a reliable determination of the size of the global landmine problem does not exist yet (ICBL, 2009). As a matter of fact, various landmine impact surveys have overestimated the size of contaminated areas, i.e., mined or mine suspected areas, or more generally, sus-

pected hazardous areas (SHAs), which are areas suspected of having a mine/ERW hazard (ERW – explosive remnants of war, such as grenades, mortars, cluster munition remnants, artillery shells, air-dropped bombs) (IMAS, 2009). The overestimation of the size of contaminated areas leads to a significant waste of financial and material resources as well as to a waste of time, which delays the use of such areas by civilians. Therefore, it is very important to determine techniques that will ensure the efficient release of formerly suspect mined areas (i.e., area reduction consists in identifying the mine-free areas out of the mine-suspected areas). The first priority of mine action becomes in allowing affected regions to reach their level of socio-economic standards. This new vision increases the importance of tools that facilitate prioritization and contribute to a rational and efficient distribution of the available resources.

Several information management systems are developed and used. An example is the Information Management System for Mine Action – IMSMA (IMSMA, 1999), developed thanks to the Geneva International Centre for Humanitarian Demining (GICHD) and in use in more than 40 affected countries. Other examples are systems completing IMSMA, such as the EODIS system (Askelin, J.-I., 1999) developed by SWEDEC in Sweden and the PARADIS system (Delhay et al., 2005) developed by the Royal Military Academy (RMA) in Belgium. Possible entries of such management systems are danger and risk assessment maps provided by the Space and airborne Mined Area Reduction Tools (SMART) project (SMART consortium, 2004; Acheroy, 2005), funded by the European Commission and described in this chapter. The maps, obtained using data fusion, synthesize the knowledge gathered from the existing data. In the framework of SMART, the fusion module, detailed in this chapter, is a very important step, since it allows taking the best benefit from all available data, and from the large efforts made in the scientific community to design detectors and classifiers adapted to these data. It has proven to be a required step before constructing risk maps. This is an improvement in comparison to existing information management systems in this area. In particular, the described approach exploits all available data and knowledge and automatically adapts to the quality of the detectors and classifier outputs.

For both close range detection and area reduction, efficient modeling and fusion of extracted features can improve the reliability and quality of single-sensor based processing (Acheroy, 2003). However, due to a huge variety of scenarios and conditions within a minefield (specific moisture, depth, burial angles) and between different minefields (types of mines, types of soil, minefield structure), a satisfactory performance of humanitarian mine action tools can only be obtained using multi-sensor and data fusion approaches (Keller et al., 2002; Milisavljević & Bloch, 2005). Recommendations for the development of multisensor mine detection systems and fusion of multisensor data have been expressed in (MacDonald & Lockwood, 2003), together with the description of the latest technologies in landmine detection. As the sensors used are typically detectors of different anomalies, combinations of these complementary pieces of information may improve the detection and classification results. Finally, in order to take into account the inter- and intra-minefield variability, uncertainty, ambiguity and partial knowledge, fuzzy set or possibility theory (Dubois & Prade, 1980) as well as belief functions (Smets, 1990) within the framework of the Dempster-Shafer theory (Shafer, 1976) prove to be useful.

In this chapter, we explain how a system based on analysis of modeling and fusion of extracted features in case of mined area reduction can be designed. Three approaches are shown, two of them based on the belief functions and one based on the fuzzy logic. They are also illustrated using real data of synthetic-aperture radar (SAR) and multi spectral sensors, collected within the SMART project. In all cases, importance of context information (knowledge about mine laying principles, mine records, etc.) is demonstrated. More details can be found in (Milisavljević & Bloch, 2003).

In Section 2, the SMART approach to the problem of mined area reduction is summarized, followed by the presentation of the image processing tools. Section 3 presents the available information for fusion in SMART, defines the decision space, and discusses the knowledge modeling issue, using the example of Glinska Poljana site in Croatia. Section 4 gives a brief overview of belief functions and possibility theory for numerical information fusion, while Section 5 presents three main fusion strategies applied in SMART. Section 6 describes additional knowledge inclusion and spatial regularization. Results are shown and discussed in Section 7. Finally, Section 8 contains final SMART steps (formation of the danger maps) and some validation results.
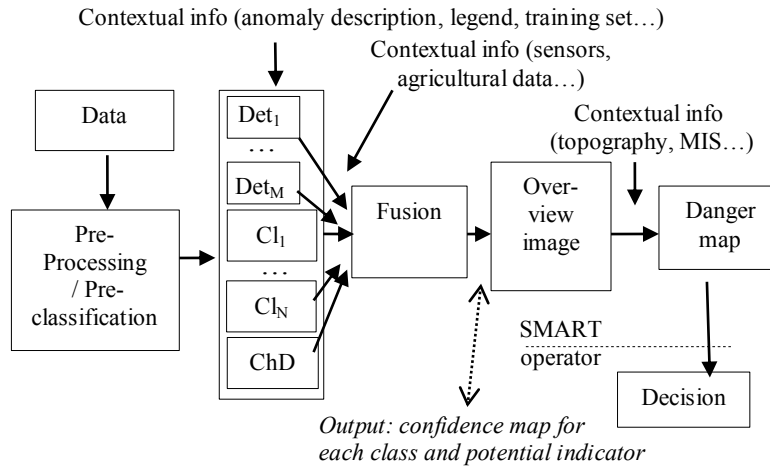
Figure 1: The global SMART approach; Det – detector, Cl - classifier, M – number of detectors, N – number of classifiers, ChD – change detection, MIS - mine information system (Milisavljević & Bloch, 2003).

2   Area Reduction: the SMART Approach

2.1   Overview of the Approach

The aim of area reduction is to find which mine-suspected areas do not contain mines and this task is recognized as a mine action activity that should result in reduction in time and resources. Several well-known methods are in use to perform area reduction, especially using mechanical means. These expensive methods change and damage the environment and the ecosystem most of the time. To avoid this, some approaches have been developed that acquire the necessary information remotely, from air or space, using appropriate sensors associated with context information collected from the field and integrated in a geographical information system (GIS).

The SMART project, funded by the European Commission/DG/INFSO, is among these approaches, and it is applied to Croatia. The goal of this project is to provide the human analyst with the SMART system, i.e. a GIS-based system augmented with dedicated tools and methods designed to use multispectral and radar data in order to assist him in his interpretation of the possibly mined scene during the area reduction process. The usefulness of such image processing tools to help photo-interpretation is, at first place, in the possibility to process automatically a large amount of data and help a visual analysis (SMART consortium, 2004). The use of SMART includes a field survey and an archive analysis in order to collect knowledge about the site, a satellite data collection, a flight campaign to record the data and the exploitation of the SMART tools by an operator to detect indicators of presence or absence of mine-suspected areas. With the help of a data fusion module based on belief functions and fuzzy sets, the operator prepares thematic maps synthesizing all the knowledge gathered with these indicators. These maps of indicators can be transformed into risk maps showing how dangerous an area may be, according to the location of known indicators, and into priority maps indicating which areas to clear first, based on socio-economic impacts and political priorities. These maps are designed to help the mined area reduction process. Results obtained using SMART have shown a reduction rate of 25% (0.98 km2 over analyzed 3.9 km2) and an error rate of 0.1% (Yvinec, 2005) for what SMART considers as not mined while it is actually mined.

2.2   Overview of the Data Processing Tools

Fig. 1 illustrates the global SMART approach. All these tools provide the pieces of information that are then combined in the fusion module. Based on field campaigns and discussions with mine action specialists of the Croatian Mine action Centre (CROMAC), it was clear that the first task is to set up a list of features (indicators) that need to be looked for in the data, and that could be seen in the data and related to the absence or presence of mines or minefields. Indicators of mine absence have the most important role in determining if an area is actually mine-safe. These indicators are not numerous, and the key one is a cultivated field. Most of available indicators are the ones that indicate mine presence. Therefore, SMART has two applications: 1) area reduction, by detecting indicators of mine absence, and 2) suspicion reinforcement, by detecting indicators of mine presence (Yvinec, 2005).

The next step consists in preprocessing, i.e., registering and geocoding the available data as well as restoring the SAR images by a speckle reduction method based on the non-decimated wavelet transform (Duskunovic et al., 2000). This step is followed by the development of methods and tools to detect the indicators, based on two approaches: 1) anomaly detection (detecting specific objects in the data) and 2) classification. The final step fuses the information produced by anomaly detectors and classifiers and builds the so-called danger maps.

The design of the system was based on considerations regarding the data and objectives: preprocessing and extraction of information have to take into account the specificities of each sensor and each type of data, and were therefore designed separately. On the contrary, the fusion step aims at gathering all information and exploits the complementarities between data and between sensors. It has then to be performed globally.

This chapter focuses on the fusion step, which provides an intermediary result in SMART, consisting of improved landcover classification maps, along with confidence values. Thus, it is a very useful result, exploited by the deminers together with the final result. This type of output was designed in agreement with the end-users. The idea is to provide them with summarized information that they can better exploit than large data masses, but including an evaluation of this summary. This is aimed at facilitating their reasoning, while leaving them with the final decision making step for demining and with potential actions to take.

3   Data, Knowledge and their Specificities in SMART

3.1   Available Images

The available images include SAR, multispectral, high resolution optical and satellite data. SAR data were collected with the E-SAR system of the German Aerospace Centre (DLR) in fully polarimetric P- and L-band and in vv-polarization (waves are vertically transmitted and received) X- and C-band. Multispectral Daedalus data were collected with a spatial resolution of 1 m and in 12 channels, ranging from visible blue to thermal infrared. SAR and Daedalus data were geocoded. DLR also provided a complete set of RMK photographic aerial views recorded with a colored infrared film at a resolution of 3 cm (note that RMK is a Zeiss digital aerial camera). This non-geocoded data set is used as evidence to control the processing tools and for qualitative interpretation by photo-interpreters.

Finally, geocoded KVR-1000 (KVR, 2011) black-and-white satellite images with a resolution of 2 m, recorded before the war in Croatia, were purchased in order to assess the changes in the landscape due to the war.

3.2   The Legend

The legend (expected classes in the images), derived based on the existing and gathered knowledge about the mined areas, is given in Table 1. Ground truth was provided as a set of regions (training regions and validation regions). In the fusion module, training regions are used for estimating the parameters of some of the proposed methods; validation regions are used for the evaluation of the results.

Class no.	Legend
1	Abandoned agricultural land
2	Agricultural land in use
3	Asphalted roads
4	Rangeland
5	Residential areas
6	Trees and shrubs
7	Shadow
8	Water

Table 1: Expected classes in the images.

3.3   Input of the Fusion Module

Table 2 summarizes the input of the fusion module.

A logistic regression classification was developed on SAR data at RMA (Borghys et al., 2004). The results consist of confidence images for each class, except for class 4, which is not detected by this approach.

A classification into hedges, trees, shadows, and rivers from SAR data has been developed at DLR (Keller et al., 2002). The method relies on the satisfaction of several criteria. The number of satisfied criteria directly provides the confidence images for hedges and trees (after scaling on [0, 1]). Shadows and rivers, provided as binary images, are "discounted" (work done at RMA based on spectral characteristics of these types of land-cover, and on existing landcover indices and meanings of Daedalus bands). Hedges and trees are then grouped to form class 6 using a maximum operator. Shadows and rivers are directly interpreted as classes 7 and 8.

Data type	Type of result
SAR	Classification with confidence images per class
SAR & Daedalus	Detection of hedges, trees, shadows, rivers, with confidence degrees, sometimes discounted
Daedalus	Supervised classification, result as a decision image
Daedalus	Region-based classification with confidence images per class
Daedalus	Belief function classification with confidence images per class
SAR & Daedalus	Binary detection of roads
SAR	River detection (binary)
Daedalus & KVR	Change detection (binary image)

Table 2: Summary of the input of the fusion module.

Several classifiers have been developed for Daedalus:

- a supervised classification method based on minimum distance has been developed at RMA and a decision image is provided (Keller et al., 2002);
- a region-based classification was performed at the Free University of Brussels (Université Libre de Bruxelles – ULB) and confidence images interpreted as membership degrees to each class are provided;
- a belief function classification was developed at RMA and confidence images per class are provided (Keller et al., 2002).

Road detection was performed at ULB and RMA (Borghys et al., 2002). Linear structures are provided. They are dilated to obtain roads with a width corresponding to the real width.

A tool for river detection previously developed at Telecom ParisTech was used too. It is based on a Markovian approach (Tupin et al., 1998). This is not directly a result of SMART but it is interesting to show how such knowledge can be introduced in the fusion process.

Change detection was obtained at ULB, based on a comparison between older KVR images and images made during the project. It provides mainly information on abandoned regions (class 1). Again, this is an important knowledge that both improves the landcover classification and provides interesting results after the fusion for the construction of danger map.

Other anomaly detection and classification tools developed in SMART were not used in the fusion module, or finally not used at all. For example, detectors of power poles, hilltops and strategic locations are not included in the legend. Therefore, they are not considered as input data for the fusion process, but they are added in the final results (construction of danger maps).

3.3 Knowledge Modeling in the SMART Project

Knowledge inclusion is one of the main features of our algorithms with respect to the commercial ones. This aspect has led to a lot of work in SMART, at different levels.

First, a study of each sensor and on related knowledge (properties and behavior of the sensor, ability to detect or not some classes, confusions between classes, complementarities between bands or channels, etc.) has been performed and used in the design of each classifier and detector.

More interesting for the fusion is the knowledge we have about possible associations between sensors. This is intensively used in the methods we propose. The type of knowledge we use concerns properties of each sensor and each classifier, as well as the complementarities between them. For instance, the knowledge that rivers can be detected by exploiting the distributions of SAR data leads to a specific way of using the result, as will be seen in the following section.

An important piece of knowledge is also provided by the comparison between KVR images and images acquired during the project. In particular this comparison allows us to extract some information about the changes between two different dates. These changes are used as an additional piece of knowledge in our procedure.

The landcover is usually not completely chaotic but there are some uniform regions showing the same landcover. This fact is used at two levels: in the classifier developed at ULB, which is based on a homogeneity criterion in regions, and at higher level, as a final step of fusion, to regularize the results in these homogeneous regions.

Finally, we can group the knowledge sources as follows:

- information provided by CROMAC (landcover labeling, mine laying records, mine accidents, MIS and GIS system,...);
- information provided by ULB (result of the second ground-truth mission, especially on potential anomaly observations);
- information coming from our experience on working with the data (reliability for various channels, various classes, mixing of some classes,... but for this, we again depend on the input regarding existing types of classes and where some representatives of each of them can be found);
- information coming from the sensor principles of operation (physical meaning of the data, of the features, choice of most promising features,...).

Some knowledge can be indirectly integrated (our experience, principles of operation), but for the others, we have to define how to integrate/introduce them in our algorithms (mine laying records, mine accidents... - e.g., as a kind of discounting on how possible it is that an area is contaminated). Some of these pieces of

knowledge are however not directly linked to the classes of interest and are therefore added at a higher level, during the construction of danger map.

The two first items above constitute useful knowledge for testing and validation. From such type of knowledge, training and validation regions could be defined, that are used respectively for parameter estimation in some fusion methods and for evaluation of the results.

The next two items are already used in some of the classifiers, i.e. before the fusion module. Some of them are used in the fusion too, as explained in Section 6. In particular, we have some "sure" detection of roads and rivers (or at least of a part of it), which can be imposed on the fusion result. Change detection constitutes also an important piece of knowledge, that allows improving the results on class 1 (abandoned fields). Also the border (no registered data) can be imposed as prior information. All this is detailed in the following sections.

4   Belief Functions and Possibility Theory for Numerical Information Fusion

4.1   Belief Function Fusion – Overview of its Main Features

Belief function theory or Dempster-Shafer evidence theory (DS) (Shafer, 1976; Smets, 1990) allows representing both imprecision and uncertainty, using plausibility and belief functions derived from a mass function, *m*. In this formalism, any combination of possible decisions from the decision space U can be quantified rather than considering only the singletons of U. This is one of the main advantages of the DS approach. Indeed, it leads to a very flexible and rich modeling, able to fit a very large class of situations, occurring in image fusion in particular. Examples of situations where DS theory may be successfully used include ignorance or partial ignorance, confusion between classes (in one or several sources of information), partial reliability, etc. (van Cleynenbreugel et al., 1991; Mascle et al., 1997; Le Hégarat-Mascle et al., 1998; Tupin et al., 1999; Milisavljević & Bloch, 2003; Milisavljević & Bloch, 2005).

The most common fusion operator, called Dempster rule of combination, has a conjunctive behavior. This means that all imprecision on the data has to be introduced explicitly at the modeling level, in particular in the choice of the focal elements. For instance, ambiguity between two classes in one source of information has to be modeled using a disjunction of hypotheses, so that conflict with other sources can be limited and ambiguity can be possibly solved during the combination. This rule also provides a measure of conflict, useful for evaluating whether the combination is meaningful or not.

For some applications, such as humanitarian demining, it may be necessary to give more importance to some classes (e.g., mines, since they must not be missed) at the decision level. Decision rules in DS theory can then be chosen in an adaptive way. For example, maximum of plausibility can be used for the classes that should not be missed, and maximum of belief for the others (Milisavljević & Bloch, 2001).

4.2   Fuzzy and Possibilistic Fusion – Overview

In the framework of fuzzy sets and possibility theory (Zadeh, 1965; Dubois & Prade, 1980), the modeling step consists in defining a membership function to each class or hypothesis in each source, or a possibility distribution over the set of hypotheses in each source. Such models explicitly represent imprecision in the information, as well as possible ambiguity between classes or decisions.

For the combination step in the fusion process, the advantages of fuzzy sets and possibilities rely in the variety of combination operators, which may deal with heterogeneous information (Dubois & Prade, 1985). Among the main operators, we find t-norms, t-conorms, mean operators, symmetrical sums, and operators taking into account conflict between sources or reliability of the sources. We do not detail all operators in this chapter, but they can be easily found in the literature, with a synthesis in (Bloch, 1996), along with a classification in terms of behavior and properties. Note that, unlike other data fusion theories (e.g., Bayesian or Dempster-Shafer combination), fuzzy sets provide a great flexibility in the choice of the operator, which can be

adapted to any situation at hand. In particular, nothing prevents using different operators for different hypotheses or different sources of information.

An advantage of this approach is that it is able to combine heterogeneous information, which is usually the case in multi-source fusion (as in the examples developed in the next sections), and to avoid to define a more or less arbitrary and questionable metric between pieces of information issued from these images, since each piece of information is converted in membership functions or possibility distributions over the same decision space.

Decision is usually taken from the maximum of membership or possibility values after the combination step. Constraints can be added to this decision, typically for checking for the reliability of the decision (is the obtained value high enough?) or for the discrimination power of the fusion (is the difference between the two highest values high enough?). Local spatial context can be used to reinforce or modify decisions (as shown in Section 7).

5 Fusion Strategies in SMART

In all that follows, the computations are performed at pixel level. A final regularization step is then applied (Section 6). Different fusion strategies have been developed (Bloch et al., 2007) and we present here three most promising ones. The choice of this approach was guided by the availability of training data (hence a supervised approach was privileged), and by the exploitation of the features of the two fusion theories described in the previous section. Note that the primary goal of this chapter is to show why fusion is useful and how it can be done, and that some examples are given here, with details in (Milisavljević & Bloch, 2005; Bloch et al., 2007). These examples provide good results, but other methods could probably be designed as well.

5.1 Adding a Global Discounting Factor (BF1)

Here, we consider each classifier as one information source.

At a very first step, the focal elements can be simply the classes (singletons), and the classifier outputs (confidence values) can be directly used as mass functions. When no confidence values are provided but only a decision image or a binary detection, the mass takes only values 0 or 1. Therefore, such an approach inevitably results in a high mass of the empty set after the combination. Moreover, in this way, only the classes detected by all classifiers can be obtained as resulting focal elements. Hence, no good result can be expected with this approach. It shows the interest of really using belief function theory or any other that takes into account the specificities of the classifiers, disjunctions of classes and ignorance.

Thus, we introduce an additional step where focal elements are still the singletons, as above, but also U, where U is the set of possible classes (also called the full set or the frame of discernment, using terms of the belief function theory) as described in Table 1. The definition of m(U) takes into account both the fact that some classes are not detected - thus it should be equal to 1 at points where 0 is obtained for all detected classes - as well as global errors. So, we propose to use a discounting factor α equal to the sum of the diagonal elements of the confusion matrix, divided by the cardinality of the training areas. This discounting is applied on all masses defined as in the first step (as classifier outputs, i.e., confidence values). Then:

$$m(U) = 1 - \alpha . \tag{1}$$

Note that this approach uses explicitly the confidence matrix, which should be computed on the training areas for each classifier or detector. It results that at each step of the fusion, the focal elements are always singletons and U. Decision rule can be maximum of belief, of mass or of pignistic probability (all being equivalent in this case).

This approach is very easy to implement and models in a simple way the fact that classifiers or detectors may not give any information on some classes and may be imperfect.

5.2 Use of Confusion Matrices for More Specific Discounting (BF2)

The main idea behind this approach is to use the confusion matrices for defining more specific discounting for each class. Namely, each class of each classifier (or detector) is considered as a source and we take into account the behavior of the classifier with respect to the other classes, using the confusion matrices to define discounting for each class. From the confusion matrix computed from the decision made from one classifier and from training areas, we derive a kind of probability that the class is C_i given that the classifier says C_j as:

$$C(i,j) = \frac{conf(i,j)}{\sum_i conf(i,j)} \quad , \qquad (2)$$

where the values $conf(i,j)$ denote the coefficients of the confusion matrix. We can ignore the low values and normalize the others, in order to reduce the number of non-zero coefficients (thus the number of focal elements in the following). We used a threshold value of 0.05.

There are several ways to use this normalized confusion matrix (Eq. 2), for example, by setting $m(C_i)=c(i,j)$ for detected pixels in case of detectors and deriving a more complex method for classifiers. The most interesting way, applying to both classifiers and detectors in a similar way, is as follows. From $v(C_j)$ (denoting the result provided for class C_j by a classifier), we define:

$$m(C_i) = v(C_j) \, c(i,j) \qquad (3)$$

for all classes C_i which are confused with C_j (which provides $\sum_i m(C_i) = v(C_j)$), and:

$$m(\Theta) = 1 - v(C_j). \qquad (4)$$

In comparison to the simplest method, instead of keeping a mass on C_i only (and U), this mass (Eqs. 3 and 4) is spread over all classes potentially confused with C_i, hence better exploiting the richness of the information provided by a classifier.

5.3 Fuzzy Fusion (FUZZY)

In order to compare the previous methods with a fuzzy approach, we test a simple method. Namely, for each class, we choose the best classifiers and combine them with a maximum operator, possibly using some weights. Then decision is made according to a maximum rule. The choice is made based on the confusion matrix for each classifier or detector, by comparing the diagonal elements in all matrices for each class. In Section 7, we detail the best detections in the illustrated example, according to the confusion matrix of each classifier or detector. These best detections provide the inputs of the combination step, and a simple maximum operator performs well for this step.

This approach is very fast. It uses only a part of the information, which could also be a drawback if this part is not chosen appropriately. Some weights have to be tuned, which may need some user interaction in some cases. Although it may appear ad hoc to some extent, it is interesting to show what we can get by exploiting the best parts of all classifiers.

6 Knowledge Introduction and Spatial Regularization

As mentioned in Section 3.4, a large attention was paid to knowledge inclusion in the design of the system. Note that knowledge on the classifiers, their behaviors, etc. is already included in the previous steps. At this step, we use only the pieces of knowledge that directly provide information on the landcover classification. Other pieces of knowledge (such as mine reports) are not directly related to classes of interest, but rather to the dangerous areas, and are thus included in the danger map construction, which follows the fusion.

Several pieces of knowledge proved to be very useful at this step. They concern on the one hand some "sure" detections. Some detectors are available for roads and rivers, which provide areas or lines that surely

belong to these classes. There is almost no confusion, but some parts can be missing. Then these detections can be imposed on the classification results. This is simply achieved by replacing the label of each pixel in the decision image by the label of the detected class if this pixel is actually detected. If not, its label is not changed. As for roads, additional knowledge is used, namely on the width of the roads (based on observations done during the field missions). Since the detectors provide only lines, these are dilated by the appropriate size, taking into account both the actual road width and the resolution of the images.

Another type of knowledge is very useful: the detection of changes between images taken during the project and KVR images obtained earlier. The results of the change detection processing provide mainly information about class 1, since they exhibit the fields which were previously cultivated, and which are now abandoned. These results do not show all regions belonging to class 1, but the detected areas surely belong to that class. Then a similar process can be applied as for the previous detectors.

With the proposed methods, it was difficult to obtain good results on class 2, while preserving the results on class 1 that is crucial since it corresponds to fields no longer in use, hence potentially dangerous. Therefore we use the best detection of class 2 (extracted from region based classification on Daedalus) as an additional source of knowledge. As shown in Section 7, this additional knowledge introduction leads to better results.

The last step is regularization. Indeed, it is very unlikely that isolated pixels of one class can appear in another class. Several local filters were tested, such as a majority filter, a median filter, or morphological filters, applied on the decision image. A Markovian regularization approach on local neighborhoods was tested too. The results are not significantly better. A better approach is to use the segmentation into homogeneous regions provided by ULB. In each of these regions, a majority voting is performed: we count the number of pixels in this region that are assigned to each class and the class with the largest cardinality is chosen for the whole region (all pixels of this region are relabeled and assigned to this class). This type of regularization, which is performed at a regional level rather than at a local one, provides good results, as will be seen in the following.

7   Results of BF1, BF2 and FUZZY

Results shown here are obtained on the Glinska Poljana site in Croatia. In case of BF1, for each classifier, the discounting factor α (Eq. 1) is calculated from the normalized sum of the diagonal elements of the confusion matrix obtained on the training areas (Table 3). After this type of fusion, a lot of confusion occurs between classes 1 and 2, but this is largely improved by knowledge inclusion, while the noisy aspect is suppressed by regularization. In order to assess classification accuracy, we use user's accuracy (UA) and producer's accuracy (PA) measures that can be derived directly from confusion matrices. UA represents the probability that a given pixel will appear on the ground as it is classified. PA is the percentage of a given class that is correctly identified on the map. Table 4 shows some results for a few classes. Note that the most interesting classes for danger map building are 1, 2, 3 and 8, and that, regarding the purpose of the project, PA is important for classes 1 and 8, and UA for classes 2 and 3.

In addition, the "best classifier" (BC) in Table 4 is not always the same one, but the result is the one provided by the classifier that is the best for a particular class.

In order for the reader to have a better visual idea about the images containing the results, Fig. 2 (left) contains the raw image of Glinska Poljana in a visible channel of Daedalus, together with a zoom of the bottom middle part (Fig. 2, right).

Team	Data type	Type of result	α
RMA	SAR	Classification with confidence images per class (except class 4)	0.41
DLR & RMA	SAR & Daedalus	Detection of hedges, trees, shadows, rivers, with confidence degrees for hedges and trees; rivers and shadows discounted based on Daedalus bands	0.11
RMA	Daedalus	Supervised classification, result as a decision image	0.46
ULB	Daedalus	Region based classification with confidence images per class	0.80
RMA	Daedalus	Belief function classification with confidence images per class	0.67

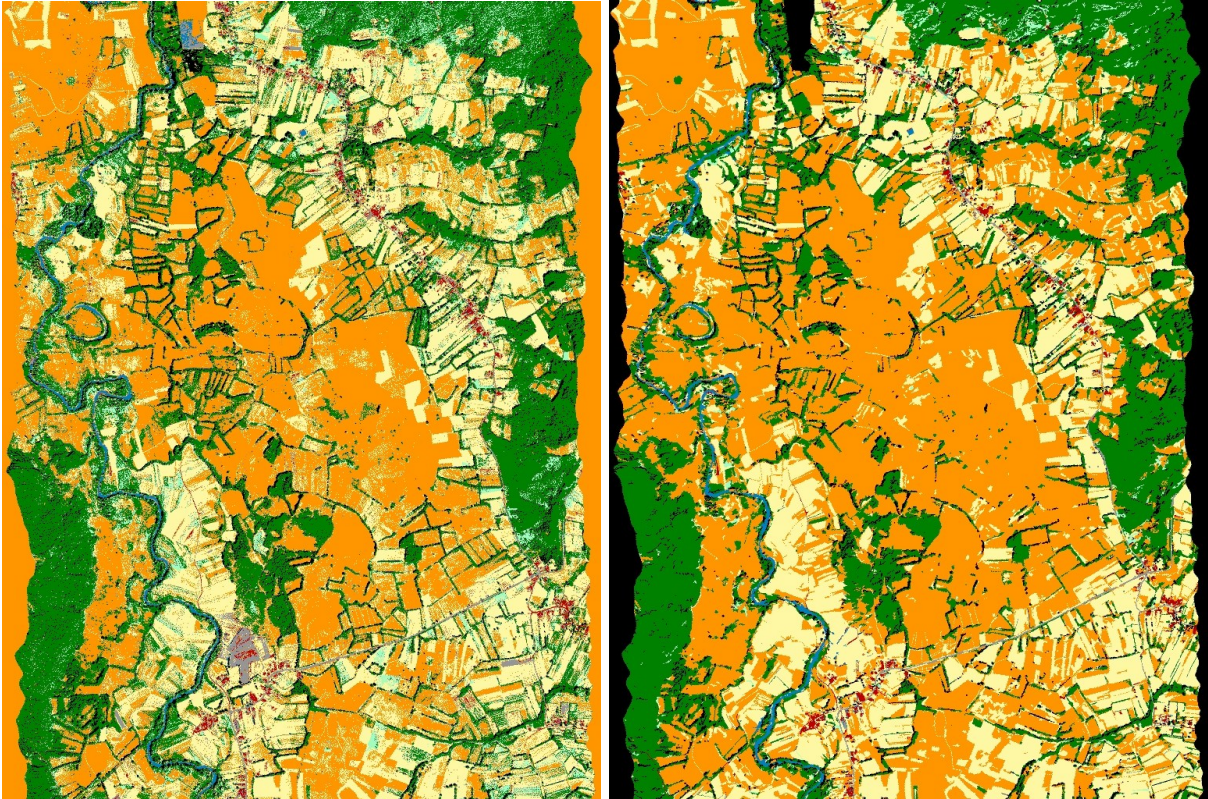**Table 3:** Discounting factors for method BF1.
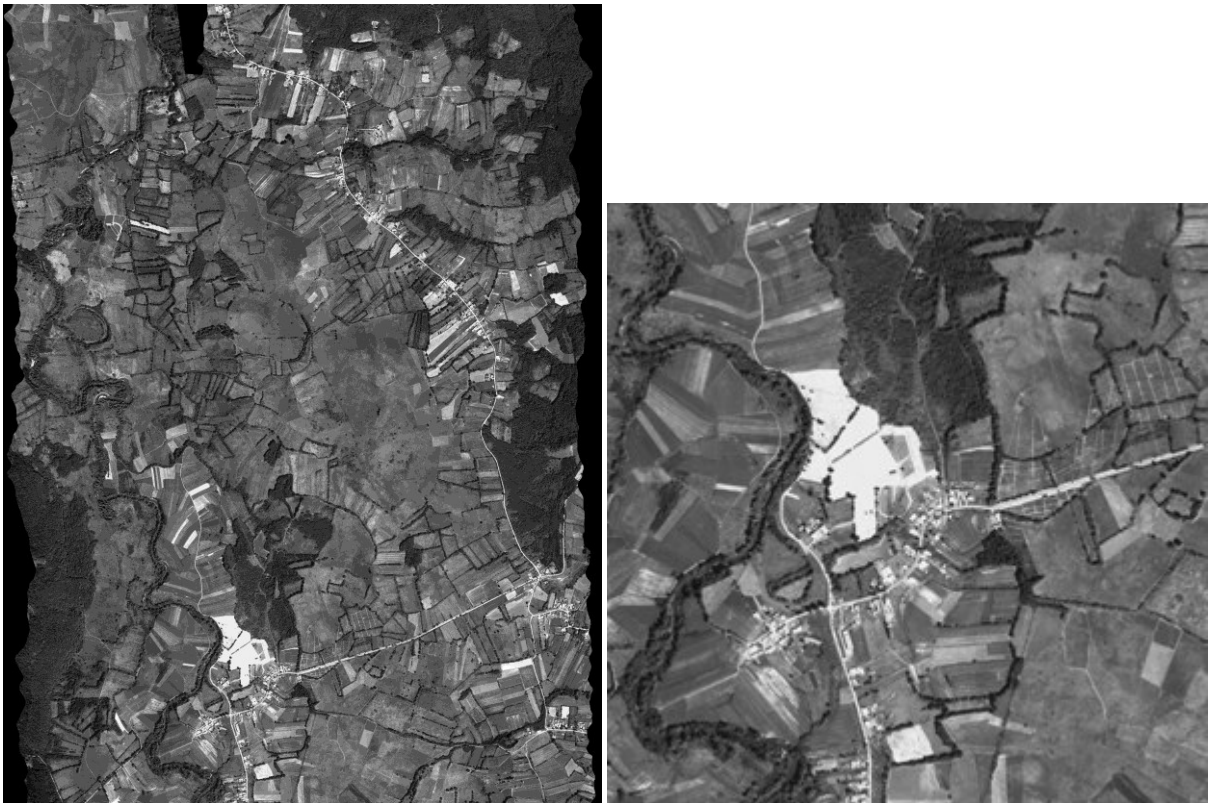
Figure 2: Visible channel of Daedalus (left), zoom of the bottom middle part (right).

Class	BC	BF1	BF2	FUZZY
1 (PA)	0.84	0.81	0.78	0.89
2 (UA)	0.87	0.86	0.81	0.95
3 (UA)	0.88	0.96	0.96	0.98
8 (PA)	0.96	0.97	0.99	0.99

Table 4: UA and PA for all three methods (after knowledge inclusion and spatial regularization) and the best classifier (BC) for each important class.
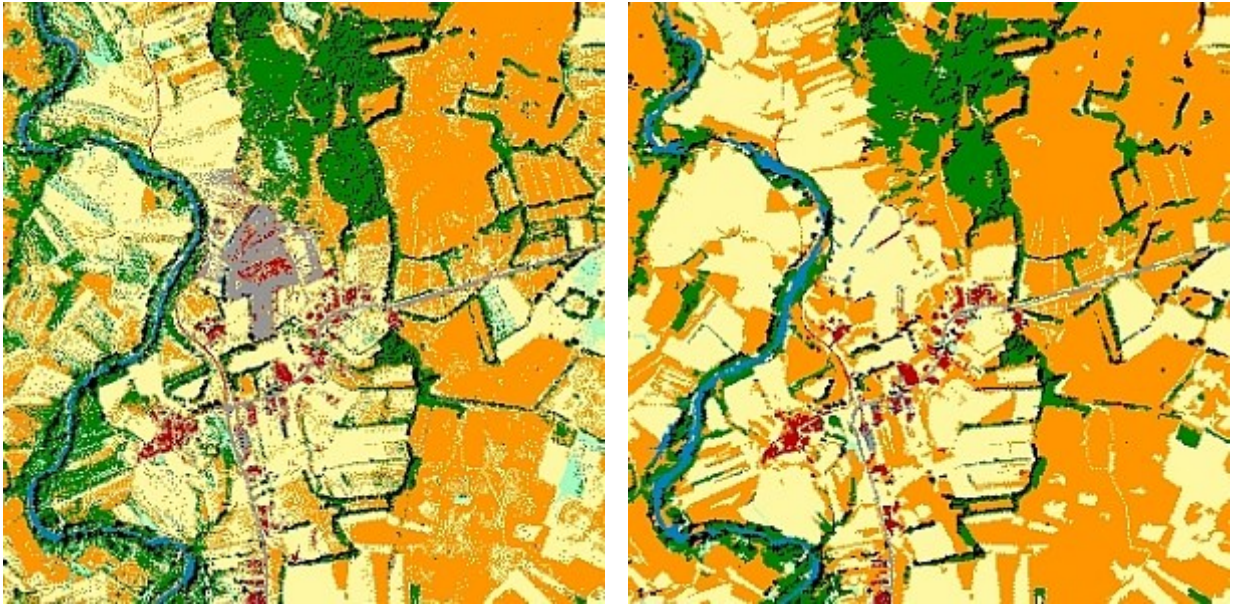

Figure 3: BF1 results: basic (left), after knowledge inclusion and spatial regularization (right).

After classification of this area using BF1 (basic version), we obtain the results given in Fig. 3 (left), while knowledge inclusion and spatial regularization applied to these results lead to Fig. 3 (right). The color code in all classification results is as follows: class 1 – orange; 2 – yellow; 3 – medium grey; 4 – light green; 5 – dark red; 6 – dark green; 7 – brown; 8 – blue. Zoomed parts of these two results that correspond to the region from Figure 2 (right) are given in Figure 4.

The fusion module also provides confidence and stability images. The confidence image represents, at each pixel, the confidence degree of the decided class. The stability image is computed as the difference between the confidence in the decided class and confidence in the second most possible class. If the stability is high, this means that there is no doubt about the decision, and if it is low, the decision should be considered carefully. The confidence image and the stability image can be multiplied to provide a global image evaluating the quality of the classification in each point. In the BF2 method, the confusion matrices for each classifier are

normalized row by row, and the coefficients that are higher than 0.05 are used for discounting the corresponding classes. The results of the basic version of this type of fusion yield a poor detection of class 1 and a lot of confusion between this class and classes 2 and 7. In addition, class 4 is not detected and detection of class 3 is worse than with BF1. However, the results are largely improved by knowledge inclusion and confusions are strongly reduced. Finally, the noisy aspect is suppressed by the regularization, leading to an improved detection, in particular for class 8. Results are given in Fig. 5 left (after knowledge inclusion and spatial regularization), while the zoomed part corresponding to the region from Fig. 2 (right) is shown in Fig. 6 (left). UA and PA are given in Table 4.


Figure 4: Zoomed parts of Fig. 3: basic BF1 results (left), after knowledge inclusion and spatial regularization (right).

For the fuzzy method, the following outputs of classifiers have been used for each class:

1: SAR logistic regression, region-based classification, belief function classification and change detection;
2: region-based classification and belief function classification;
3: region-based classification and road detection;
4: region-based classification, minimum distance classification and belief function classification;
5: region-based classification and belief function classification;
6: region-based classification and SAR trees and hedges detection;
7: SAR logistic regression, SAR shadow detection, minimum distance classification and belief function classification; the maximum is discounted by a factor 0.5, taking into account that this class is not really significant for further processing (shadows "hide" meaningful classes);
8: region-based classification, belief function classification and river detection.
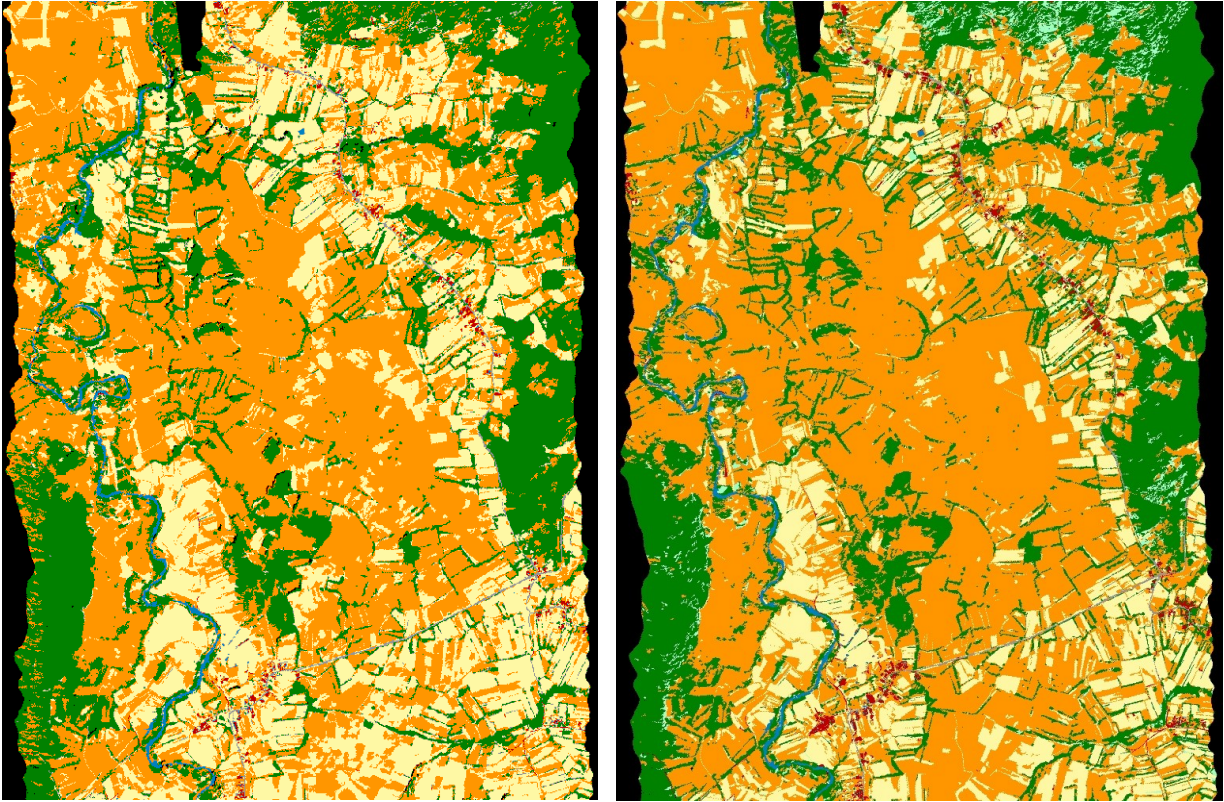
Figure 5: Results with BF2 (left) and FUZZY (right), both after knowledge inclusion and spatial regularization.

The results of this fusion in its basic version are already very good, because of the fact that not all pieces of information provided by the classifiers are used, but only the best part of them. Further improvements are obtained by knowledge inclusion. After the regularization step, class 7 disappears, but this is not a problem since this class is not significant for further analysis.

Results of this method are shown in Fig. 5 right (after knowledge inclusion and spatial regularization). Fig. 6 (right) represents the zoomed region corresponding to the one from Fig. 2 (right). Table 4 contains PA and UA for this type of fusion too.

In order to get a synthetic view of the results obtained by the three methods, the normalized sums of the diagonal elements of the confusion matrices are given in Table 5. The two methods based on belief functions provide similar global results, BF1 being somewhat better. The differences appear mainly when looking at each class individually. The improvement achieved with knowledge inclusion is significant. Regularization provides an additional improvement. The final results are globally better than the ones obtained by each of the initial classifiers, as can be seen by comparing the values with those displayed in Table 3 (the best classifier provides a global accuracy of 0.80). The fuzzy method is the best in its basic version, since it already selects the best inputs, thus the improvement due to the next steps is not as important as for the belief function methods.
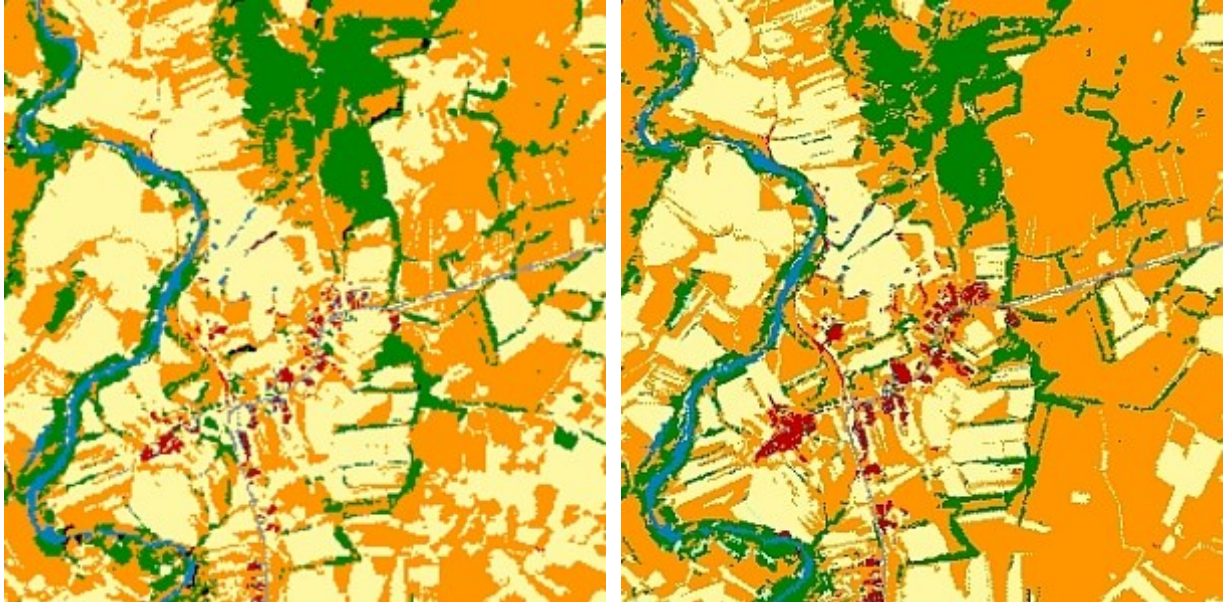
Figure 6: Zoomed parts of Fig. 5: BF2 results (left), FUZZY results (right).

Method	Basic	Knowledge inclusion	Spatial regularization
BF1	0.70	0.81	0.85
BF2	0.65	0.78	0.81
Fuzzy	0.79	0.83	0.84

Table 5: Comparison of the normalized sum of diagonal elements of the confusion matrices for the three tested methods.

8 Danger Maps and First Results of SMART Validation

The danger maps are synthetic documents designed to help the end users in their decision-making process regarding area reduction. They are created from results of all detection and classification tools and methods used in SMART (as well as some other sources such as fieldwork). These maps constitute the final output of the system and represent the basis for proposing areas for area reduction. Note that the results are for decision makers and that the reduction of a suspicious area is not an automatic process.

Four types of danger maps are developed in SMART (SMART consortium, 2004). The most useful continuous location maps, such as the one in Fig. 7, are obtained as a weighted sum of factors derived from the number of indicators of mine presence at each point (IMP), with a superimposition of vectors having a see-through inside, representing the number of indicators of mine absence (IMA) at each point.

During the process of area reduction, the decision makers can refer to information relating to the IMA and the associated confidence values. The other key element is the information that concerns the IMP and the associated confidence values. As pointed out by the end users, this information can be of use for prioritizing the mine clearance operations.
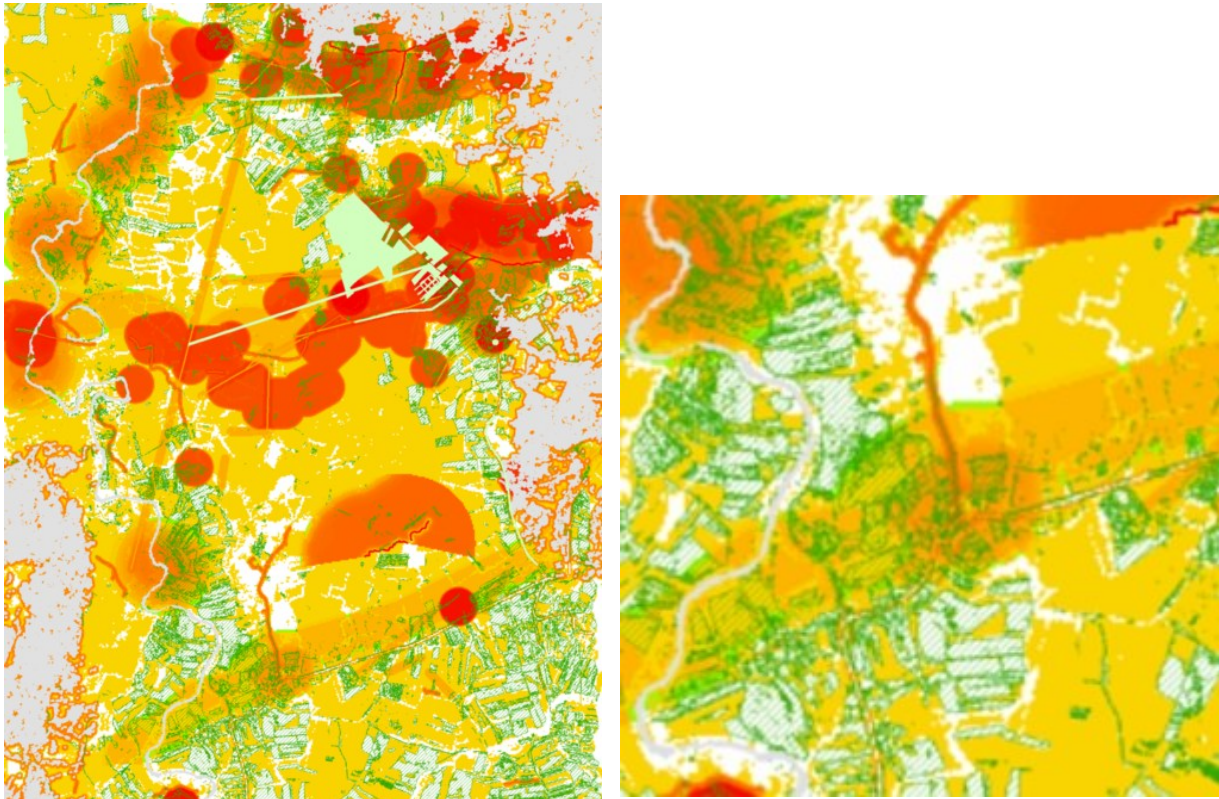
Figure 7: Left: Continuous location map (SMART consortium, 2004). Grey areas are outside of the scope of SMART, while no data exists for white areas. Demined areas are light green. IMAs are superimposed as parallel white and green lines. The degree of danger is on the scale from green (low) via yellow (intermediate) to red (high). Right: Zoomed part of the continuous location map (left) corresponding to the region from Fig. 2 (right).

Validation was done by blind tests at three sites in Croatia (Yvinec, 2005) having 3.9 km2 in total: Glinska Poljana (0.63 km2, fertile valley surrounded by hills), Pristeg (1.5 km2, rocky, Mediterranean area) and Čeretinci (1.7 km2, flat agricultural area). In each site clearing was performed after the flight campaign in order to have the true status of the mine presence, but this information was not available before the validation of produced danger maps. From these maps, a selection of areas proposed for area reduction was done, and areas considered as suspect were selected too. In average 25% of the mine-free area has been proposed for reduction: Glinska Poljana – 7.7%, Pristeg – 9.0% and Čeretinci – 47%. The error rate of 0.1% is relatively constant for all three sites. In addition to this technical evaluation, a panel of independent mine action experts working in Croatia has evaluated the SMART method and danger maps. They recognized SMART as a successful project that solved several crucial problems of the aerial survey of the suspected areas, especially by approved indicators of mine presence, efficient use of very different sensor techniques, data fusion and danger map functionalities. It has been found that it might be even more suited for risk assessment.

The results presented in this chapter show moreover the usefulness of information fusion for this type of application. Combining various detection or classification results and introducing contextual and domain knowledge contribute to the improvement of the interpretation of the images and the derived danger maps.

Finally, note that such results can be integrated in robotics systems for demining applications, by providing exploitable and useful inputs to the system.

Acknowledgements

References

Acheroy, M. (2003). Mine Action Technologies: Problems and Recommendations. *Journal for Mine Action*, 7(3).

Acheroy, M. (2005). Spaceborne and Airborne Mined Area Reduction Tools. In *Workshop on Inventory and analysis of operationally validated results related to mine action space- and airborne surveys*, Zagreb, Croatia, Scientific Council of CROMAC (invited paper).

Askelin, J.-I. (1999). EOD IS - a Loaded Laptop. *FOA magazine, Swedish Defense Research*, 1.

Bloch, I. (1996). Information combination operators for data fusion: a comparative review with classification. *IEEE Transactions on Systems, Man and Cybernetics*, 26, 52-67.

Bloch, I., Milisavljević, N. & Acheroy, M. (2007). Multisensor Data Fusion for Spaceborne and Airborne Reduction of Mine Suspected Areas. *XI International Journal of Advanced Robotic Systems*, 4(2), 173-186.

Borghys, D.; Lacroix, V. & Perneel, C. (2002). Edge and Line Detection in Polarimetric SAR Images. In *International Conference on Pattern Recognition*, 2 (pp. 921-924), Quebec, Canada.

Borghys, D.; Perneel, C.; Keller, M.; Pizurica, A. & Philips, W. (2004). Supervised Feature-Based Classification of Multichannel SAR Images Using Logistic Regression. In *European Conference on Synthetic Aperture Radar (EUSAR)* 2 (pp. 949-952), Ulm, Germany.

Delhay, S., Lacroix, V. & Idrissa, M. (2005). PARADIS: GIS Tools for Humanitarian Demining. In *ISCRAM 2005: Information Systems for Crisis Response And Management Conference*.

Denœux, T. (1995). A k-nearest Neighbor Classification Rule based on Dempster-Shafer Theory. *IEEE Transactions on Systems, Man and Cybernetics*, 25 (5), 804-813.

Dubois, D. & Prade, H. (1980). *Fuzzy sets and systems: Theory and applications*. Collection: Mathematics in science and engineering 144, New York, London: Academic Press.

Dubois, D. & Prade, H. (1985). A review of fuzzy set aggregation connectives. *Information Sciences*, 36, 85-121.

Dubois, D. & Prade, H. (1988). Representation and Combination of Uncertainty with Belief Functions and Possibility Measures. *Computational Intelligence,* 4, 244-264.

Duskunovic, I.; Stippel, G.; Pizurica, A.; Philips, W. & Lemahieu, I. (2000). A New Restoration Method and its Application to Speckle Images. In *IEEE Int'l Conf. on Image Processing (ICIP 2000)* (pp. 273-276), Vancouver, Canada.

ICBL (2009) - International Campaign to Ban Landmines. *Landmine Monitor 2009: Toward a Mine-Free World*. Mines Action Canada, Handicap International, Human Rights Watch, Landmine Action, Norwegian People's Aid.

IMAS (2009) - International Mine Action Standards. *IMAS 08.20: Land Release*. First Edition, United Nations Mine Action Service (UNMAS), New York, USA.

IMSMA (1999). http://www.imsma.org, Information Management System for Mine Action.

Keller, M.; Milisavljević, N.; Suess, H. & Acheroy, M. (2002). Reduction of Mine Suspected Areas by Multisensor Airborne Measurements: First Results. In *SPIE Conference on Detection Technologies for Mines and Minelike Targets*, 4742 (pp. 857-871), Orlando, USA.

KVR (2011). http://www.spotimage.com/web/en/186-kvr-1000.php, KVR-1000 images.

Le Hégarat-Mascle, S.; Bloch, I. & Vidal-Madjar, D. (1998). Introduction of Neighborhood Information in Evidence Theory and Application to Data Fusion of Radar and Optical Images with Partial Cloud Cover. *Pattern Recognition*, 31(11), 1811-1823.

MacDonald, J. & Lockwood, J.R. (2003). *Alternatives for Landmine Detection*. RAND Corporation.

Mascle, S.; Bloch, I. & Vidal-Madjar, D. (1997). Application of Dempster-Shafer Evidence Theory to Unsupervised Classification in Multisource Remote Sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 35(4), 1018-1031.

Milisavljević, N. & Bloch, I. (2001). A Two-Level Approach for Modeling and Fusion of Humanitarian Mine Detection Sensors within the Belief Function Framework. In *Applied Stochastic Models and Data Analysis*, 2 (pp. 743-748), Compiègne, France.

Milisavljević, N. & Bloch, I. (2003). Sensor fusion in anti-personnel mine detection using a two-level belief function model. *IEEE Transactions on Systems, Man and Cybernetics*, Part C, 33(2), 269-283.

Milisavljević, N. & Bloch, I. (2005). *Improving Mine Recognition through Processing and Dempster-Shafer Fusion of Multisensor Data*. In: Computer-Aided Intelligent Recognition, Techniques and Applications, M. Sarfraz (Ed.), ch. 17, pp. 319-343, J. Wiley, New York, ISBN: 0-470-09414-1.

Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton Univ. Press.

SMART consortium (2004). SMART - final report.

Smets, P. (1990). The Combination of Evidence in the Transferable Belief Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5), 447-458.

Smets, P. (1993). Belief Functions: The Disjunctive Rule of Combination and the Generalized Bayesian Theorem. *International Journal of Approximate Reasoning*, 9, 1-35.

Tupin, F.; Maître, H.; Mangin, J. F.; Nicolas, J. M. & Pechersky, E. (1998). Linear Feature Detection on SAR Images: Application to the Road Network. *IEEE Transactions on Geoscience and Remote Sensing*, 36(2), 434-453.

Tupin, F.; Bloch, I. & Maître, H. (1999). A First Step Towards Automatic Interpretation of SAR Images using Evidential Fusion of Several Structure Detectors. *IEEE Transactions on Geoscience and Remote Sensing*, 37(3), 1327-1343.

van Cleynenbreugel, I.; Osinga, S.A.; Fierens, F.; Suetens, P. & Oosterlinck, A. (1991). Road Extraction from Multi-temporal Satellite Images by an Evidential Reasoning Approach. *Pattern Recognition Letters*, 12, 371-380.

Yvinec, Y. (2005). A Validated Method to Help Area Reduction in Mine Action with Remote Sensing Data. In *The 4th Int'l Symposium on Image and Signal Processing and Analysis (ISPA 2005)*, Zagreb, Croatia.

Zadeh, L. (1965). Fuzzy sets. *Information and Control*, 8(3), 338-353.