Sensor fusion for anti personnel landmine detection, a case study

Eric den Breejen, Klamer Schutte, and Frank Cremer
TNO Physics and Electronics Laboratory,
P.O. BOX 96864, 2509 JG The Hague, The Netherlands

ABSTRACT

In this paper the multi sensor fusion results obtained within the European research project GEODE (Ground Explosive Ordnance Detection system) are presented. The lay out of the test lane and the individual sensors used are described. The implementation of the SCOOP algorithm improves the ROC curves, as the false alarm surface and the number of false alarms both are taken into account. The confidence grids, as produced by the sensor manufacturers, of the sensors are used as input for the different sensor fusion methods implemented. The multi sensor fusion methods implemented are Bayes, Dempster-Shafer, fuzzy probabilities and rules. The mapping of the confidence grids to the input parameters for fusion methods is an important step. Due to limited amount of the available data the entire test lane is used for training and evaluation. All four sensor fusion methods provide better detection results than the individual sensors.

Keywords: Sensor fusion, Multi sensor, Bayes, Dempster-Shafer, Fuzzy probabilities, Rules, Landmine

1. INTRODUCTION

In this paper we present the sensor fusion results obtained in the frame of the European research project GEODE coordinated by THOMSON-CSF DETEXIS. In the consortium Emrad provides an experimental GPR, ELTA provides a COTS GPR, Förster provides a dual frequency metal detector, Marconi communications provides a 3-5 μm and a 8-10 μm wavelength band infrared camera, THOMSON-CSF DETEXIS constructed a test lane facility and TNO-FEL studies the multi sensor fusion methods.

From the data obtained during the tests the sensor manufacturers produce a confidence grid map for the entire test lane. The five confidence maps are the input for the fusion methods. The confidence values are calculated for every grid cell of 2.5*2.5 cm of the 25 meter long and 1 meter wide section of test lane used for measurements.

The SCOOP (Split Clusters On Oversized Patches) algorithm is developed to present results in ROC curves that take both the number of false alarms and the false alarm surface into account. The graphical representation of the detection results allows a visual check of the sensor results.

The sensor fusion methods implemented are Bayes, Dempster-Shafer, fuzzy probabilities and rules. For each of the fusion methods, the mapping of the confidence values to the input parameters of the methods is described. The results of the fusion methods are presented in ROC curves for the different fusion methods and sensor combinations.

2. MULTI SENSOR DATA ACQUISITION

In the THOMSOM-CSF DETEXIS, test lane 26 mine objects are buried or laid on the surface at known locations. Also 6 additional false alarm objects are placed in the test lane. In Table 1, the layout of the mine objects and the false alarm objects in the test lane is given.

The test lane is divided in three parts with different terrain types. The first part is bare agricultural ground, the second part is the vegetation area, the third and last part is the bare sand area. The agricultural part is 15 meters long, the vegetation part 5 meter and the sand area is also 5 meter long. For the measurements, the different sensors were connected to the trolley and moved over the test lane to perform the needed measurements.

The result of the processing for each sensor is a confidence number for each grid cell of the test lane. The sensor fusion algorithms are performed on each single grid cell. For the GEODE project, it was agreed to use a grid cell size of 2.5*2.5 cm, to make detection of the smallest APLs possible. This results in an overall grid dimension for the entire test lane of

Part of the SPIE Conference on Detection and Remediation Technologies
for Mines and Minelike Targets IV • Orlando, Florida • April 1999
SPIE Vol. 3710 • 0277-786X/99/$10.00

40*1000 grid cells. The confidence levels are not the actual probability of a mine. The confidence numbers are used to indicated the order in the probability of a object as detected with the sensor. This means that a higher confidence number implies a higher probability of a mine, but these do not necessarily scale linearly.

lane #	Name	x [m]	y [m]	Metal	Size	Depth
1	MAUS 1	0.50	0.25	2	2	0
2	Mle 59	0.75	1.13	1	1	1
3	Mle 72A	0.25	1.75	2	1	0
4	VS 1.6	0.50	2.63	2	2	2
5	PFM 1	0.75	2.88	3	2	0
6	Foot print	0.25	3.88	1	2	0
7	Cartridge case	0.50	4.50	3	1	1
8	MK 2	0.25	5.38	3	2	2
9	Piquet 62	0.75	6.25	1	1	0
10	Mortar 60	0.50	7.25	3	2	2
11	MD 82B (M14)	0.25	7.63	2	1	0
12	Cylindric print	0.50	8.50	1	2	0
13	Trip wire	0.75	9.50	3	1	0
14	PRB 409	0.50	10.38	2	2	1
15	Mle 51	0.75	11.50	1	1	0
16	Mle 72A	0.25	12.50	2	1	0
17	TS 50	0.50	13.50	2	2	1
18	Mle 59	0.25	14.50	1	1	0
19	PMN	0.75	15.50	3	2	0
20	TS 50	0.50	16.50	2	2	1
21	Stone	0.75	17.25	1	2	2
22	Coca cola can	0.25	17.75	3	2	1
23	Trip wire	0.75	18.13	3	1	0
24	VS 69	0.25	18.50	3	2	0
25	VS 2.2	0.75	19.25	2	2	2
26	HEC3A1	0.50	19.50	2	1	0
27	Mle 51	0.25	20.38	1	1	1
28	Mle 51	0.75	21.25	1	1	1
29	Chew. gum paper	0.25	21.75	2	1	2
30	BLU 62	0.75	22.50	3	1	1
31	Mle 59	0.25	23.38	1	1	1
32	PMN	0.75	24.25	3	2	0

Table 1: The layout of the THOMSOM-CSF DETEXIS test lane in Paris. The legend: for Metal 1=No metal, 2=Low metal and 3=High metal; for Size 1=Small and 2=Large; for Depth 0=Surface, 1=Slightly buried and 2=Buried. Objects 1-18 are laid in the bare agricultural area, objects 19-26 are laid in the vegetation area and objects 27-32 are laid in the bare sand area.

The GPR as supplied by ELTA is a high ground clearance GPR. The centre of the object detected by ELTA is indicated with a circle of 40 cm diameter. The mine confidence for this area is the value calculated by the standard ELTA processing. The second layer is 30 cm wide ring around the first circle, the third layer is a 30 cm wide ring around the second layer (The confidence values for the second and the third layer were set to 35% and 5% of the centre value after optimisation). The levels were optimised by evaluating the detection performance for each combination of levels. As criterion, the area under the ROC was maximised.

The Emrad GPR system is based upon a four channel radar with a bandwidth comparable to the expected radar cross-section of anti-personnel mines. The Emrad philosophy is that for reliable anti-personnel mine detection, it is considered to be necessary to collect radar data on a 50 mm square grid over the area to be searched. The spacing of the antenna elements, in the existing radar, is approximately three times this value, so that to complete a search of the test lane, it was essential to carry out successive scans with the antennas progressively scanned over the width of the test area.

The Förster dual frequency, continuous wave, Minex metal detector. The Förster metal detector is a very sensitive detector for small quantities of metal. The current Förster processing does not detect small metal objects on top of a strong signal produced by a large metal object. The areas where a strong signal is present, the detector blindness or uncertainty is increased. This results in a non zero confidence level for these areas.

Marconi processes the data of two infrared cameras in the wavelength bands 3-5 μm and 8-10 μm. The 3-5 μm camera has a better detection performance than the 8-10 μm camera. This is probably the result of a different SNR between the cameras (the 3-5 μm camera is of a newer generation). In this paper only the results of the 3-5 μm camera are presented.

2.2. GRAPHICAL REPRESENTATION OF DETECTION RESULTS

For the graphical representation of the detection results we use four strip graphs. In these graphs the same test lane is presented four times in strips above each other. In the top strip the confidence level output of the sensor is given. In the second strip from the top the mine objects are represented as 8*8 grid cell large squares. In the third strip the combination of the top strip and the second strip is presented, this indicates whether the sensor did have a confidence value at the position of a mine object. In the fourth and bottom strip a representation of all the objects, including false alarm objects, in the test lane is given. In Figure 1 the four strip graph representation is shown for a random sensor.

To be able to present neutral results, this random sensor is introduced. This random sensor is created by generating an uniformly distributed confidence level grid. This confidence level grid has the same dimension as the other sensors in the GEODE sensor suite. To create objects (clusters with similar confidence level) with higher and lower confidence levels, the random confidence levels are Gaussian filtered with a kernel width of about 3.

The confidence levels of the random detector are coded in grey onto the first strip of Figure 1. The grey coding for the confidence in the top strip ranges from light grey to black, and white for the case the confidence level is zero. In the third strip we see that all mine objects represented in the second strip are detected, as they have a confidence level higher than zero (white).

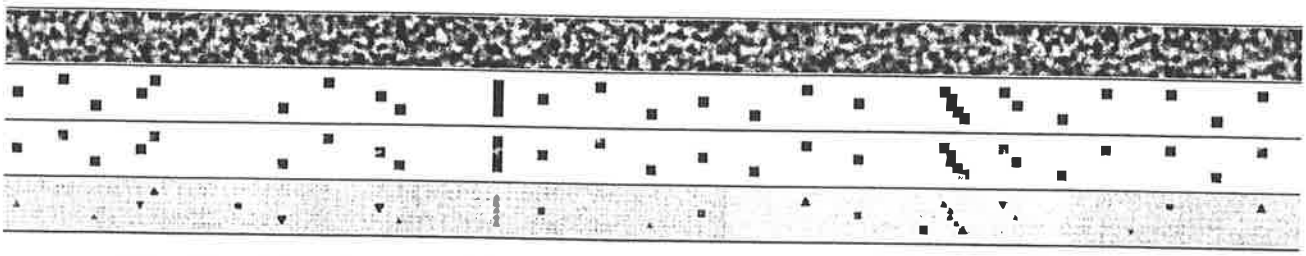

Figure 1: Four strip graph representation of THOMSOM-CSF DETEXIS test lane in Paris for a random detector. The same test lane is presented four times. Top strip is the confidence map of the sensor, Second strip is the location of the mines, Third strip is the combination of the top strip and the second strip indicating what confidence level the sensor produces at the location of a mine. The legend for the bottom strip representing all objects is: Grey coded for metal content, Shape for level of burial and Size indicates large or small objects. Upwards pointing triangles are used to indicate that the object is on top of the test lane; squares indicate slightly buried objects (1-5 cm) and downwards pointing triangles indicate buried objects (>5 cm). The size of these shapes reflects the size of the objects, two classes are defined, larger or smaller than 8 cm. The grey coding of the objects is used for metal content: light grey is used for non metal, dark grey is used for low metal and black for high metal. The bare agricultural area is dark grey, the vegetation area is light grey, and the bare sand area is dark grey at right end.

3. SCOOP ALGORITHM

For the evaluation of the detection results, we implemented a simple ROC curve calculation. The flowchart for this first simple implementation is shown in Figure 2 A, the resulting ROC curve is shown in Figure 2 B. The simple implementation of the ROC curve does not show a continuous rising curve as we would expect for a hypothetical ROC curve. The number of false alarms per square meter decreases for very high detection rates. This is caused by the fact that at high detection rate almost the entire lane is marked as mine. The simple implementation considers a large cluster to be a mine when within the cluster a mine object is present. The surface of the cluster is not taken into account.

At very low confidence thresholds the detected clusters cover a large area. When this is considered as a single object, as is the case in simple implementation, the ROC curve is not as expected continuous increasing. The SCOOP (Split Clusters On Oversized Patches) algorithm splits a large detected cluster in multiple objects of fixed size. Scattered small and not connected clusters of detections produce a high number of false alarms as in the simple implementation of the ROC.

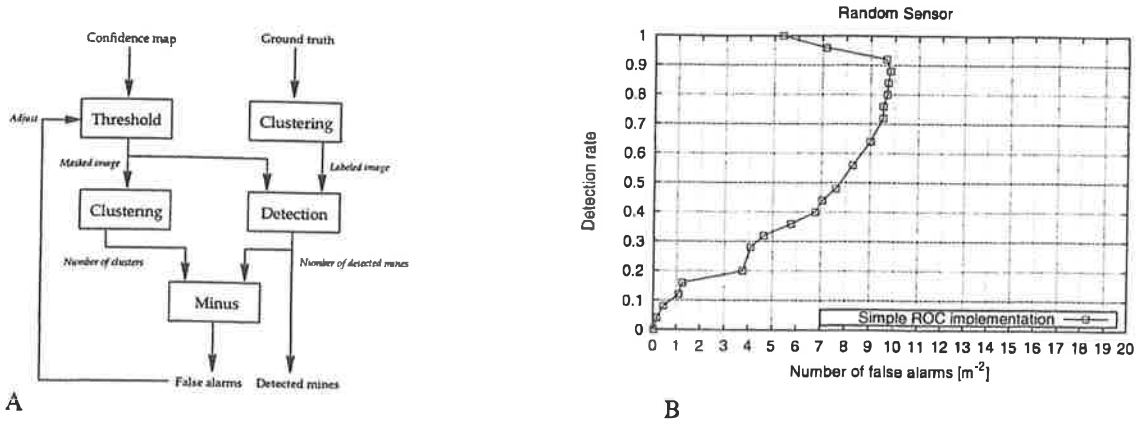

Figure 2: The flowchart for the simple ROC implementation, and the resulting ROC curve for the random sensor.

The SCOOP algorithm is implemented as follows: Each cluster is considered separately: the cluster size (in grid cells) is calculated and the number of mines within the cluster is determined. For every mine within the cluster, the cluster size is reduced with the standard dig area (an area of 20 by 20 cm or 64 grid cells). If the effective cluster size is still above the standard dig area, each dig area which fits in the effective cluster size is counted as one false alarm. The flowchart of the evaluation method with the SCOOP algorithm is presented in Figure 3 A. With the SCOOP algorithm the number of false alarm objects per square meter is 25, if the area is completely covered by a single object.

In Figure 3 B, the previously described random sensor is evaluated using the method with the SCOOP algorithm. In this ROC curve, we see that there is now a functional relationship between the false alarms per square meter and the detection rate as opposed to simple implementation. From now on, we will only use the evaluation method with the SCOOP algorithm, since it produces comparable ROC curves for confidence grids.

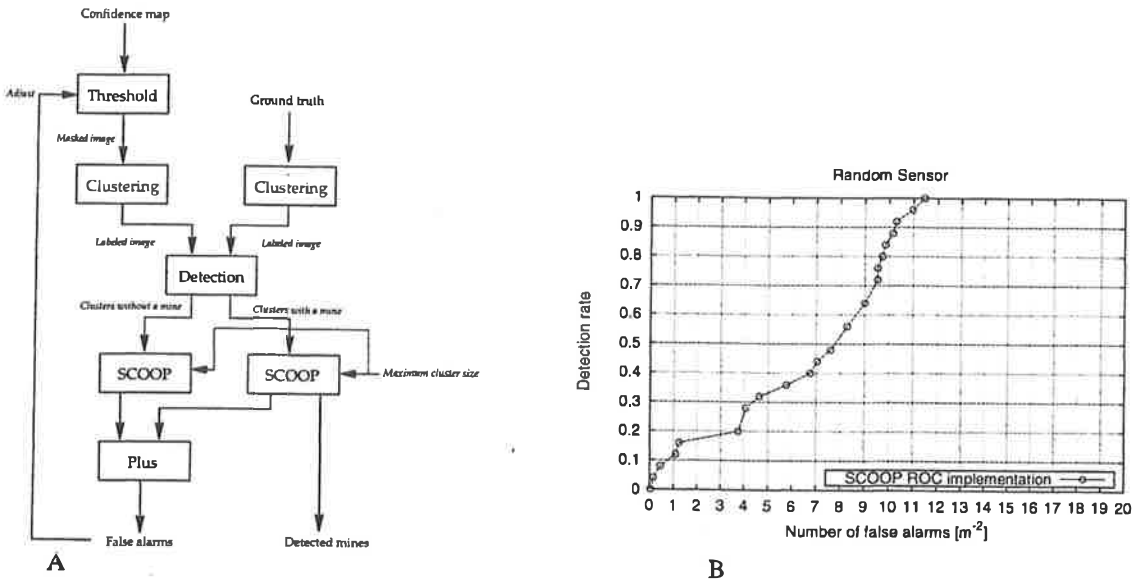

Figure 3: The flowchart for the implementation of ROC with the SCOOP algorithm, and the resulting ROC curve for the random sensor The number of false alarms is with the SCOOP implementation continuously increasing with the detection rate.

4. INDIVIDUAL SENSOR RESULTS

The individual sensor results are presented in 4 strip graphs, Figure 4 to Figure 7, as described in paragraph 2.2. The resulting ROC curves for the individual sensors are included in Figure 12, in which the comparison is made between detection results of fusion methods and the individual sensors.

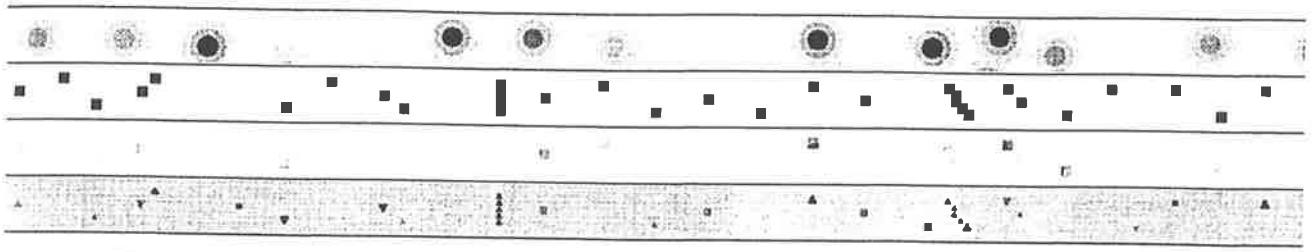

Figure 4: Output of ELTA's processing and the ground truth.

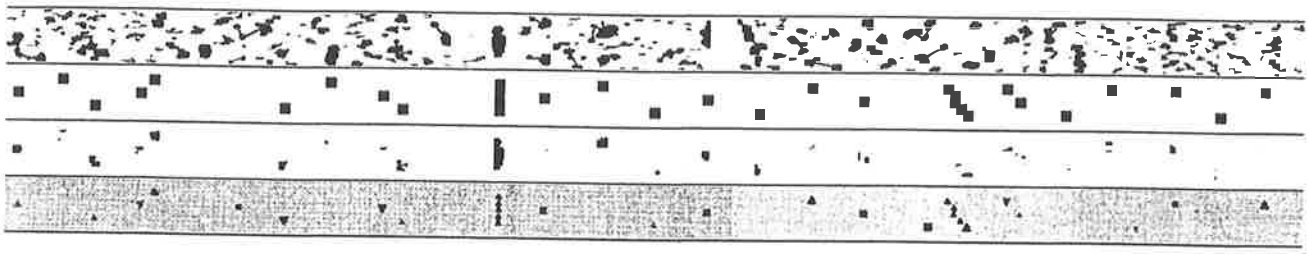

Figure 5: Output of Emrad's GPR after spatial processing by TNO and the ground truth of the test lane.

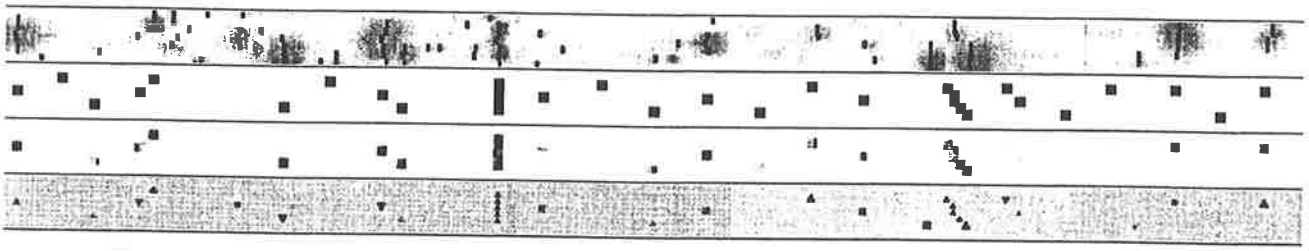

Figure 6: Output of Förster's processing and the ground truth of the test lane.

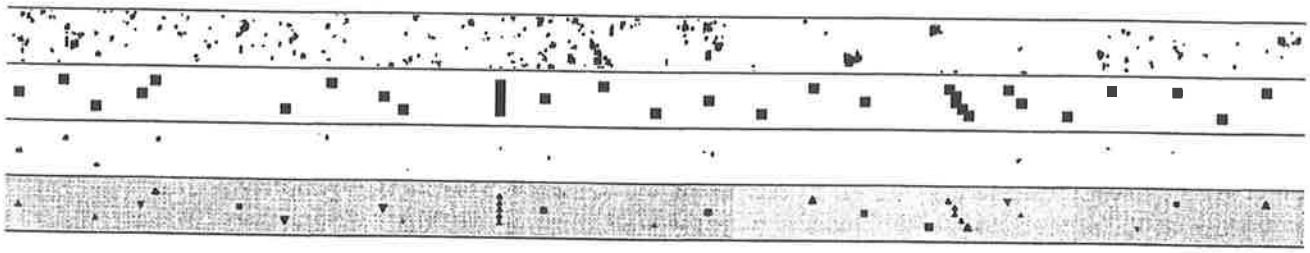

Figure 7: Output of Marconi's processing for the 3-5 µm camera and the ground truth of the test lane.

5. FUSION IMPLEMENTATION

5.1. OPTIMISATION METHOD

To get optimal detection performance, the parameters of Dempster-Shafer, Bayes and fuzzy probabilities has to be set correctly[1]. There is one parameter attached to each of the three sensors, besides the threshold. Each of these three parameters will be set to 11 different levels independently. The threshold will be varied in 101 steps. This gives a total of $11^3*101=134,431$ parameter settings.

For each parameter setting, the detection performance is evaluated using the SCOOP algorithm. This evaluation results into a detection rate and a number of false alarms per square meter, which is a point in the ROC diagram. Out of the 134,431

points in the ROC diagram only the best points are selected. These points are connected together and presented in an ROC curve. Note that to go from one point in the ROC to the other, not only the threshold has to be adapted, but that also the other three parameter settings can be different for the other point.

5.2. DEMPSTER-SHAFER

For Dempster-Shafer[2], three inputs per sensor are needed: the probability mass assigned to a mine M(m), the probability mass assigned to background M(b) and the unassigned probability mass M(m∪b). The sum of these masses is always one, so there are only two independent masses.

For each sensor, there is only one confidence level, so for each sensor this confidence level needs to be mapped onto the three probability masses. In this project, a simple approach is chosen: the unassigned mass will be kept constant for each of the sensors, and the probability mass assigned to a mine is linear with the confidence level. These two conditions fully determine the mapping, see Figure 8.

The probability mass assigned to a mine M(m) is zero at confidence level 0 and the probability mass assigned to background M(b) is zero for confidence level 1. The uncertainty level M(m∪b) used in the Dempster-Shafer implementation is optimised for all sensors, using the optimisation method, described in paragraph 5.1, the value range = [0,1].

The ROC curves for the Dempster-Shafer method for both sensor combinations are presented in Figure 10 A.

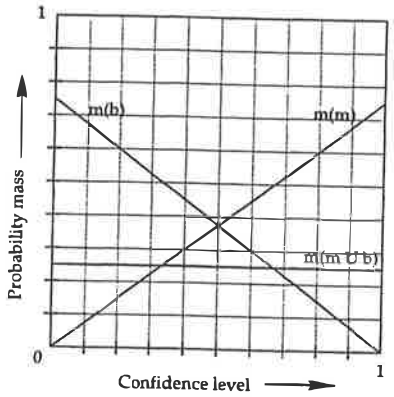

Figure 8: Mapping the confidence level onto the three probability masses: M(m), M(b) and M(m∪b).

5.3. BAYESIAN REASONING

For Bayes[2] we used the assumption of independence of the individual sensors. The conditional probability $P_{sensor}(m|Conf)$ needs to be known for every confidence level (Conf) and each sensor. This conditional probability is very difficult to estimate using the limited amount of data available from the test lane. The conditional probability is given by:

$$P_{sensor}(m|Conf) = (1 - U) \cdot Conf + U/2,$$

with U the level of uncertainty. This uncertainty is varied from U=0 (so $P_{sensor}(m|Conf) = Conf$) to U=1 (so $P_{sensor}(m|Conf)=0.5$), in eleven steps independent for each sensor. These uncertainty levels are the parameters used in the optimisation method described in paragraph 5.1.

The result of the Bayesian fusion is the product of the multiplication of the three sensor probabilities. Each sensor mapping combination results into a joint probability grid. The resulting ROC curve is presented in Figure 10 B.

5.3. FUZZY PROBABILITIES

The confidence level is mapped onto the a fuzzy probabilities[3,4] with a Gaussian membership function. The centre of the Gaussian kernel is set at the confidence level. The width of the kernel is the parameter which has to be optimised, using the

method described in paragraph 5.1. Small width implies that the sensor has a large influence, whereas a large width implies a limited influence. The width is varied between 0.01 and 1 in 11 steps on a logarithmic scale.

The three fuzzy probabilities are combined using the minimum operator and the centre of area (COA) defuzzification method. This means that for every probability level, the minimum of the three membership functions of the fuzzy probabilities is taken. Of the resulting joint fuzzy probability, the probability, for which the area can be divided into two equal areas, is determined. This probability is the output of the fusion process. The resulting ROC curve is presented in Figure 11 A.

5.5. RULES

For both combinations of three sensors we implemented the rules fusion method[5]. The Emrad GPR confidence values range from 0 to 6, ELTA GPR confidence ranges from 0 to 255, Förster confidence ranges from 0 to 255 and Marconi IR confidence ranges from 0 to 100. An example of a rule for Emrad GPR, Förster MD and Marconi 3-5 μm IR camera is:
((GPR >= 3) AND (MD >=255)) OR
((GPR >= 6) AND (MD >= 128) AND (IR >= 74)) OR
((MD >=1) AND (IR > =93))
The optimal rule set found for this sensor combination can be found in Table 2, and the resulting ROC curve in Figure 11 B.

false alarms	detection rate	rule 1	rule 2	rule 3
0.04	0.12	6,255,90	6,191,93	
0.08	0.16	6,255,72	6,191,93	
0.44	0.36	6,255,0		
0.56	0.40	5,255,0		
0.72	0.48	3,255,0		
0.88	0.52	3,255,0	0,1,93	
1.04	0.56	3,255,0	6,128,74	0,1,93
1.16	0.64	0,255,0		
1.92	0.68	0,255,0	0,8,88	
2.52	0.72	6,0,0	0,255,0	
2.96	0.80	5,0,0	0,255,0	
3.52	0.84	5,0,0	0,255,0	0,8,88
3.68	0.92	4,0,0	0,255,0	
5.48	0.96	3,0,0	0,255,0	
8.36	1.00	1,0,0		

Table 2: The number of false alarms and detection rate for different rules. The three values in rules represent the threshold level on the confidence value for the sensors. The sensors used in this rule set are Emrad, Förster and Marconi 3-5 μm.

The rule sets are derived by selecting the minimal rule set for a set of mines[1]. A difference is that in the procedure as described earlier only one value of each sensor corresponded with a mine, where with the real data used now several values in the confidence grid are related to a single mine. For fusion this leads to the problem that the position of the maximum confidence for a given mine for one sensor might not coincide with the maxima for the other sensors.

For fusion this leads to the problem that the position of the maximum confidence for a given mine for one sensor might not coincide with the maxima for the other sensors.

This problem is visualised in Figure 9, with the simplification that only six grid cells for a mine are drawn instead of the 64 (8 by 8) grid cells we use in reality. With the six mine confidence combinations, we will get the rules:
from grid cell 2: (GPR >= 3) AND (IR >= 1)
from grid cell 3: (GPR >= 1) AND (MD >= 1) AND (IR >= 2)
from grid cell 4: (MD >= 4) AND (IR >= 2)

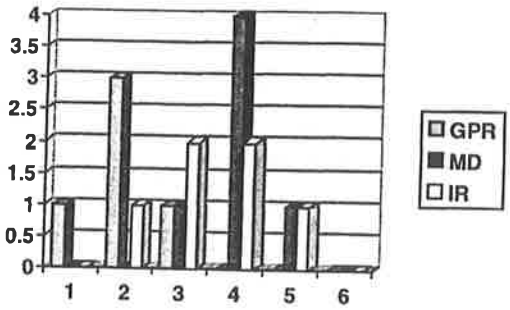

Figure 9: What is the best rule to catch this mine, given these six grid cells over a mine? For each grid cell the confidence result of the three sensors is given and due to not coinciding maxima the best rule is not directly clear.

Grid cells 1, 5 and 6 will lead to less restrictive rules compared to the rules of grid cells 2, 3 or 4, and so are not of interest. Which of the rules should be used is dependent on:

- rules needed for other mines
- the number of false alarms generated for each of the rules.

This is solved by keeping all these rules for each mine, and selecting the optimal rule set in a selection process using all mines and the whole test lane. So for every mine we can have several rules, and we have to select that combination of rules which has the lowest false alarm rate. Full detection rate is guaranteed by keeping at least one rule for every mine. This is implemented by an exhaustive search procedure.

This will lead to the point with 100% detection rate on the ROC curve. Other points on the ROC curve are calculated by excluding combinations of mines such that, for a given detection rate, the false alarm rate is as low as possible. These combinations of mines are the result of a search process with early pruning. The result is that every point on the ROC curve has a different rule set. The resulting ROC curves for both sensor combinations are presented in Figure 11 B.

6. MULTI SENSOR FUSION RESULTS

6.1. SENSOR FUSION METHODS

In this paragraph the resulting ROC curves of the four implemented sensor fusion methods, Bayes, Dempster-Shafer, fuzzy probabilities and rules, are presented for both sensor combinations. The sensor combinations contain either the ELTA GPR or the Emrad GPR.

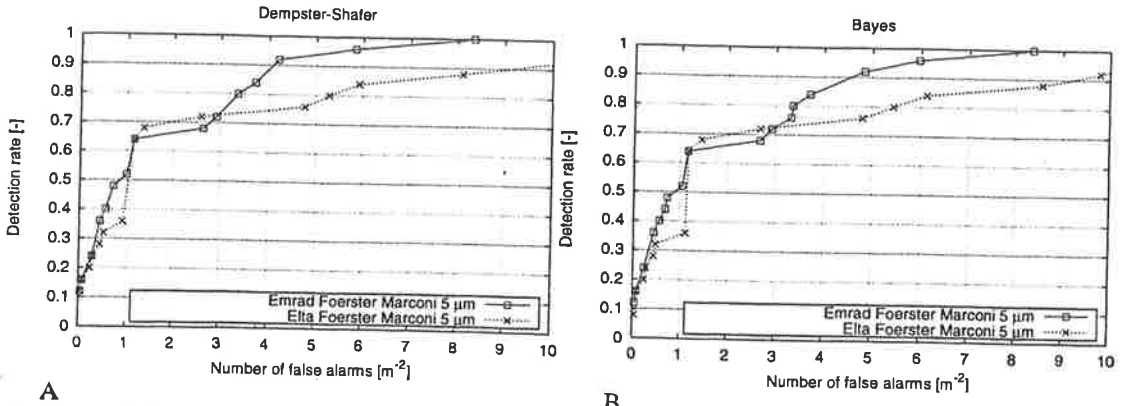

Figure 10: ROC curves for Dempster-Shafer and Bayes fusion methods for both sensor combinations.

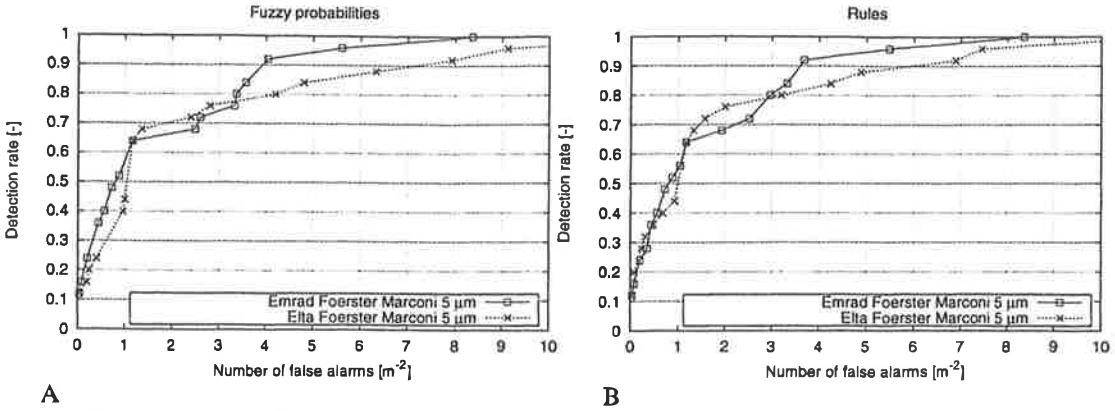

Figure 11: ROC curves for fuzzy probabilities and rules fusion methods for both sensor combinations.

6.2. COMPARISON OF SENSOR FUSION METHODS

In this paragraph the different sensor fusion methods are presented for the two possible sensor combinations with a single GPR, metal detector and a single infra red camera. The ROC curves of the individual sensors used in the combination are also presented in each figure. The fusion method ROC curves have a dotted line connecting the calculated points and the individual sensors have a continuous line connecting the calculated points.

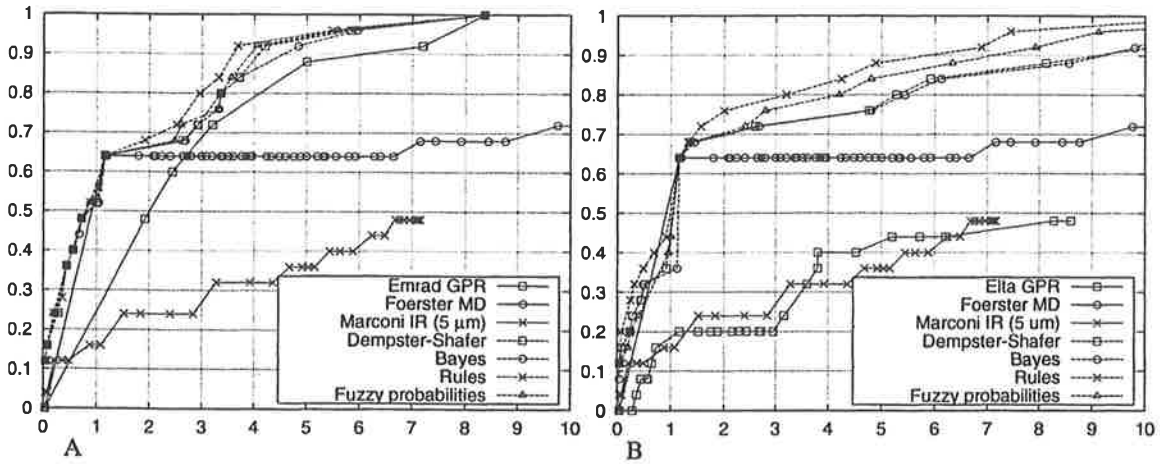

Figure 12: A: ROC curves for the different sensor fusion methods for the sensor combination Emrad GPR, Förster and Marconi 3-5 µm. B: ROC curves for the sensor combination ELTA GPR, Förster and Marconi 3-5 µm. The individual sensor ROC curves for the sensors in the combination are presented with continuous line connecting the calculated points.

7. CONCLUSIONS

All conclusions are based on the measurements of the 25 meter test lane of THOMSOM-CSF DETEXIS in Paris. The number of measurements is limited compared with the number of variables.

The number of variables are: terrain type three (agriculture, vegetation and sand), size two (large and small mines), metal three (no metal, low metal and high metal content mines) and burial depth three (on the surface, slightly buried (1-5 cm) and deep buried (>5 cm) mines). This leads to 54 possible combinations. In the test lane only 26 mines are present and 6 selected false alarm objects.

Due to absence of multiple measurements of all possible combinations the entire test lane is used for training and evaluation.

7.1. CONTRIBUTIONS OF THE INDIVIDUAL SENSORS TO THE SENSOR FUSION PROCESS

Within the GEODE project, every sensor has as primary input to the sensor fusion process a confidence level grid. Every sensor manufacturer strives for a high confidence level for his sensor on a location on the grid corresponding to each mine. However, for each mine the different sensors have their maximum confidence values on different locations on the grid. This means that the sensor fusion processes sometimes can not utilise the optimum results of every sensor as the sensor fusion processes are performed on a single grid cell level.

The ELTA GPR is a COTS single sensor mine detection solution. As sensor used in the sensor fusion system, this has the following drawbacks:

- The spatial accuracy is rather low. In combination with other sensors, it is difficult to make one to one correspondences between the ELTA detections and the detections of the other sensors.
- The system is not capable of generating alarms corresponding to low mine probability. This means that the output of the ELTA GPR can not be used to declare area free of mines.

The Emrad GPR is a product currently still under development. It's current strength is that it is capable of detecting of all mines. However it's capability with regard to clutter rejection is not yet fully developed. As it currently is not capable to generate detections in a range with low false alarm rates (smaller than 2 false alarms per square meter), it's contribution to sensor fusion in the ranges with low false alarm rates is limited.

The metal detector proves to be very valuable in detecting mines with a medium to high metal content. For scenarios where a lot of such mines occur, this sensor will thus be very important.

Due to the low number of surface laid mines present in the THOMSOM-CSF DETEXIS test lane, the contribution of the IR camera is limited. However, especially in the range with low false alarm rate, the IR camera contributes to improve the detection performance of the sensor fusion system.

7.2 MULTI SENSOR FUSION

From the results as given in Figure 10 and Figure 11, we can conclude that:

- For all fusion methods the results as presented in the ROC curves is better than the results of the individual sensors.
- Rules gives for every point on the ROC the best performance.
- The order of average performance of the fusion methods is: rules, fuzzy probabilities, Dempster-Shafer and Bayes. However, the differences in performance between the sensor fusion methods are not very large.

The results given above should be seen in the light that we use the same data set for training and evaluation. As the rule sensor fusion method has considerable more degrees of freedom compared to the other sensor fusion methods, it can be expected that performance of the rules method will deteriorate more when independent training and evaluation data is used compared to the other methods. As such, we currently cannot state which methods should be used in an operational sensor fusion system.

7.3 FUTURE RESEARCH

The work of the GEODE project is continued in the LOTUS (Light Ordnance detection with Teleoperated Unmanned System) project. For the sensor fusion processing measurement campaigns with improved location accuracy should improve the sensor fusion results. The processing of the individual sensors should be more adapted towards to fusion processing.

In future data acquisition campaigns, considerable more data should be collected, to be able to generate independent training and evaluation sets. Only when these independent sets are available, final conclusions can be drawn regarding the best sensor fusion method to be used in an operational system.

Within the GCFF (GEODE Common File Format), object data is stored for all the sensors. Within the current sensor fusion processes this object data is not used. We realise this data could prove useful in the sensor fusion process. Future research should investigate towards the use of the object data into the sensor fusion process.

An operational system foreseen as the final results of a development process of which the GEODE project is a part, will consist of an automated sensor fusion system in combination with a man in the loop to finalise any detection decision. Under the assumptions that such a system will have:

- a forward speed of one meter per second,
- a width of one meter,
- an allowed operator load of one possible detection per second to be classified by the man in the loop,

the automatic sensor fusion system should be capable of reducing the false alarm rate to less then one false alarm per square meter while keeping its detection rate close to 100%. As such, we can conclude that both the individual sensors as well as the sensor fusion system need substantial improvements before an operational system can become reality.

8. ACKNOWLEDGEMENTS

This research is partly funded by the European Union as ESPRIT project GEODE, number 26337. The consortium consisted of the following companies: Emrad limited, United Kingdom, ELTA Electronics Industries Ltd, Israel , Institut Dr. Förster, Germany, Marconi communications, Italy, THOMSOM-CSF DETEXIS, France and TNO Physics and Electronics Laboratory, The Netherlands.

9. REFERENCES

1. Frank Cremer, Eric den Breejen and Klamer Schutte, "Sensor Data Fusion for Anti-Personnel Land-Mine Detection", EuroFusion98, Great Malvern, UK 1998.
2. L.A. Klein, "Sensor and Data Fusion Concepts and Applications", SPIE, Washington, 1993.
3. A.N.S. Freeling, "Possibilities versus fuzzy probabilities - two alternative decision aids" in *Fuzzy Sets and Decision Analysis*, H.-J. Zimmermann, L.A. Zadeh, B.R. Gaines (Eds.), Elsevier Science Publishers, Amsterdam, 1984.
4. H.-J. Zimmermann, *Fuzzy sets, decision making and expert systems*, Kluwer Academic Publishers, Boston, 1987.
5. L.A. Zadeh, "Fuzzy sets as a basis for the theory of possibility" in *Fuzzy Sets and Systems*, Volume 1, North-Holland Publishing Company, 1978