

# Infrared processing and sensor fusion for anti-personnel land-mine detection

John Schavemaker, Frank Cremer, Klammer Schutte and Eric den Breejen

Electro-Optical Systems, TNO Physics and Electronics Laboratory  
P.O. Box 96864, 2509 JG The Hague, The Netherlands  
phone: +31 70 374 0860, fax: +31 70 374 0654, Schavemaker@fel.tno.nl

## Abstract

*In this paper we present the results of infrared processing and sensor fusion obtained within the European research project GEODE (Ground Explosive Ordnance DEtection system) that strives for the realization of a vehicle-mounted, multi-sensor, anti-personnel land-mine detection system for humanitarian demining. The system has three sensor types: an infrared camera, a ground penetrating radar and a metal detector. The output of the sensors is processed to produce confidence values on a virtual grid covering the test bed. A confidence value expresses a confidence or belief in a mine detection on a certain position. The grid with confidence values is the input for the decision-level sensor fusion and provides a co-registration of the sensors. We describe the methods TNO-FEL developed for the processing of infrared (3-5  $\mu\text{m}$ ) data to produce confidence values. We show results of experiments with infrared processing and sensor fusion on real sensor data. The performance of the processing and fusion are measured with the SCOOP evaluation method that yields a less biased probability of false alarm by taking into account the spatial arrangement of false alarms.*

*Keywords: AP land-mine detection, infrared processing, sensor fusion, performance evaluation.*

## 1 Introduction

The existence of (abandoned) land mines in a large number of post-war areas forms a major threat to human lives in these areas. The majority of these minefields is currently cleared by manual prodding, which is a very slow and tedious process and which cannot circumvent that the number of active mines is still worldwide increasing. As such, the detection of land mines by any (technical) means is an important research issue.

Current research focuses on the improvement of existent sensors [6] and the combination of multiple sensors (i.e. sensor fusion) to land-mine detection [2, 3, 8, 4, 1, 12]. The use of one sensor is generally believed to be insufficient for land-mine detection meeting the requirements of humanitarian demining for the reason that a single sensor has a false-alarm rate which is too high or a detection rate which is too low. The aim of sensor fusion is to make higher probabilities of detection ( $P(d)$ ) possible with a lower probability of false alarms ( $P(fa)$ ). In order to combine or fuse sensors, the sensor readings of the different sensors must be converted to a common grid and common confidence values. We show the processing steps taken for the infrared camera.

## 2 Data Acquisition

The experiments are based on data acquisition at the test lane of THOMSON-CSF DETEXIS, Paris, France (see also [8]). This test lane measures 25 by 1 square meter and contains 26 mines that are either buried or laid on the surface. For more information about the mines, see [10]. Additionally, the lane contains six false-alarm objects. Tables 1 and 2 give details on the mines and false-alarm objects.

Object	Name	x [m]	y [m]	Metal	Size	Depth
1	MAUS 1	0.50	0.25	low	large	surface
2	Mle 59	0.75	1.13	no	small	slightly buried
3	Mle 72A	0.25	1.75	low	small	surface
4	VS 1.6	0.50	2.63	low	large	buried
5	PFM 1	0.75	2.88	high	large	surface
6	<i>Foot print</i>	0.25	3.88	no	large	surface
7	<i>Cartridge case</i>	0.50	4.50	high	small	slightly buried
8	MK 2	0.25	5.38	high	small	buried
9	Piquet 62	0.75	6.25	no	small	surface
10	Mortar 60	0.50	7.25	high	large	buried
11	MD 82B (M14)	0.25	7.63	low	small	surface
12	<i>Cylindric print</i>	0.50	8.50	no	large	surface
13	Trip wire	0.75	9.50	high	small	surface
14	PRB 409	0.50	10.38	low	large	slightly buried
15	Mle 51	0.75	11.50	no	small	surface
16	Mle 72A	0.25	12.50	low	small	surface
17	TS 50	0.50	13.50	low	large	slightly buried
18	Mle 59	0.25	14.50	no	small	surface

Table 1. Objects 1-18 in the bare agricultural area in the test lane (continued on Table 2). The items in italics are false-alarm objects.

The test lane is divided into three parts with different types of terrain. The first part is bare agricultural ground, the second part is a vegetation area, and the third part is bare sand. The agricultural part is 15 meter long, the vegetation part five meter, and the sand area is also five meter long. For the measurements, the sensors were one by one attached to a trolley and moved over the test lane. The applied sensors were a dual-frequency metal detector of Förster, a mid wavelength band (3-5  $\mu\text{m}$ ) infrared camera and ground penetrating radar of Emrad. To perform decision-level sensor fusion, the raw sensor data is processed and mapped to obtain decision-level data on a reference grid. The sensor processing results in confidence levels on a grid with grid cells of  $2.5 \times 2.5$  cm. This grid cell size ensures that there are even multiple cells over the smallest land mine.

Images of the infrared camera in the wavelength band 3-5  $\mu\text{m}$  were recorded and pre-processed by Marconi Communications. The mines in these images have a higher apparent temperature (and thus a higher intensity in the image) than their surrounding.

### 3 Performance evaluation

For a comparison of performance, Receiver-Operator Characteristics (ROC) curves are used. In a ROC curve, the detection rate is plotted as a function of the false-alarm rate. The detection rate is defined as the fraction of the detected mines. The number of false alarms per unit area is calculated using the SCOOP (Split Clusters On Oversized Patches) [8] evaluation method. The name SCOOP refers to the scoop size area that need to be checked by the mine-clearance personnel, it is typically set to an area of 20 by 20 cm<sup>2</sup>.

Object	Name	x [m]	y [m]	Metal	Size	Depth
19	PMN	0.75	15.50	high	large	surface
20	TS 50	0.50	16.50	low	large	slightly buried
21	<i>Stone</i>	0.75	17.25	no	large	buried
22	Coca cola can	0.25	17.75	high	large	slightly buried
23	Trip wire	0.75	18.13	high	large	surface
24	VS 69	0.25	18.50	high	large	surface
25	VS 2.2	0.75	19.25	low	large	buried
26	HEC3A1	0.50	19.50	low	small	surface
27	Mle 51	0.25	20.38	no	small	slightly buried
28	Mle 51	0.75	21.25	no	small	slightly buried
29	<i>Chewing gum paper</i>	0.25	21.75	low	small	buried
30	BLU 62	0.75	22.50	high	small	slightly buried
31	Mle 59	0.25	23.38	no	small	slightly buried
32	PMN	0.75	24.25	high	large	surface

Table 2. Objects in the test lane (continued from Table 1). Objects 19–26 are in the vegetation area and objects 27–32 are in the bare sand area. The items in italics are false-alarm objects.

The flowchart of the SCOOP evaluation method is shown in Figure 1. The grid cells which have a confidence value above a threshold are clustered. Every cluster is treated differently depending on whether it contains one or more mines. If a cluster does not contain any mines, it is counted as a number of false alarms equal to the number of scoops it contains. Clusters with one or more mines that are larger in area than the product of the number of mines and the scoop size also contribute to the number of false alarms. In this case, the number of false alarms is set equal to the cluster area divided by the scoop size minus the number of mines. The method is repeated for different threshold values to obtain a number of ROC points. The SCOOP evaluation method results in a percentage of detected mines as a function of the number of false alarms per unit area. This kind of performance measure is requested by the demining experts as an indication of the reduction in workload. It takes into account that more time is needed to reject false alarms occupying large areas (multiple 'scoops') than false alarms occupying small areas.

## 4 Infrared image processing methods

In this section we describe two methods for mine detection using infrared data. The processing methods described here made the infrared camera the best sensor of the GEODE sensor platform, as shown in the experiments section. The methods rely on the principle that the mines have a higher (apparent) temperature than their surrounding. This is reflected in a higher intensity in the infrared image. In Figure 2(a) the infrared image data is shown for the GEODE test lane. One can clearly distinguish some of the mines listed in Tables 1 and 2 (the origin is in the bottom left corner). Note that depending on the demining scenario (type of soil, time of day) the mines can also have a lower apparent temperature. In that case, one has to use negative contrast.

As there is a correlation between apparent temperature and mines an obvious choice for the conversion of sensor data to confidence levels is to use the raw infrared data as confidence levels. Because the spatial resolution of the infrared camera does not match the resolution imposed by the  $2.5 \times 2.5$  cm grid cells for sensor fusion, the infrared data must be resampled. The ROC for mine detection using an infrared camera can then be calculated using the above-mentioned SCOOP algorithm, the result is shown in Figure 4(b). The ROC expresses mine detection as a function of false alarms for different global thresholds on the grid with confidence levels, i.e. infrared image data.

However, global thresholding does not take into account the changing surroundings of mines. For instance, the vegetation area of the test lane is much colder than the agricultural and bare sand area and global thresholding cannot account for that. As such, our proposed methods apply local contrast enhancement to make them invariant for the local background intensity. Furthermore, global thresholding would not reduce the number and size of false alarms. Our proposed methods perform false-alarm reduction by selecting blobs on morphological and size attributes. In order to do so, some a priori knowledge on the scoop size is used.

In the following sections we describe our proposed methods. The methods result in some different operating points (a certain detection rate with a certain number of false alarms) on the ROC.

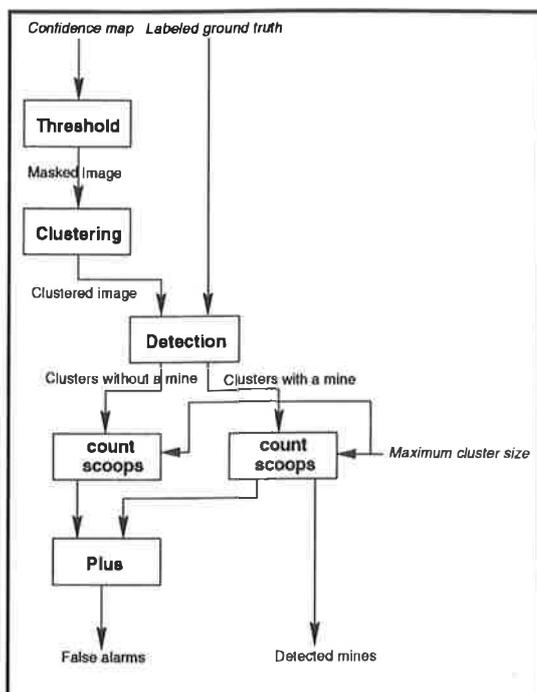


Figure 1. The flowchart for the SCOOP performance evaluation method.

#### 4.1 Method 1

Method 1 is a sequence of resampling, blob search and contrast enhancement. It consists of the following processing steps:

**1. resampling** : reduce the size of the infrared image to GEODE standard grid size by resampling the image. The resampling is done by selecting the local maximum in the image data for each grid cell.

**2. blob search** : search in the resampled image for blobs with a positive contrast. Blobs are generated by applying all possible thresholds to the image. Select those blobs whose size is between  $l \times l$  and  $u \times u$ . Optimal  $l$  and  $u$  on the GEODE dataset are found to be 9 and 9 respectively.

**3. local contrast enhancement** : calculate the local average intensity and local variance in intensity for every grid cell using a window size of  $40 \times 40$  grid cells (one square meter). Normalize the intensity of each grid cell to its local mean and variance.

**4. output** : As output, give the maximum value calculated in step 3 for each blob found in step 2 to all grid cells in that particular blob.

The processing steps are visualized in Figures 2(a) to 2(d). The ground truth (mines) is shown in Figure 2(h).

#### 4.2 Method 2

This method is an adaption of method 1. It outperforms the previous method for lower number of false alarms but is less good for higher number of false alarms. See also Figure 4(b) in Section 6. As such, depending on the required detection rate and number of false alarms, one can choose for one of the methods.

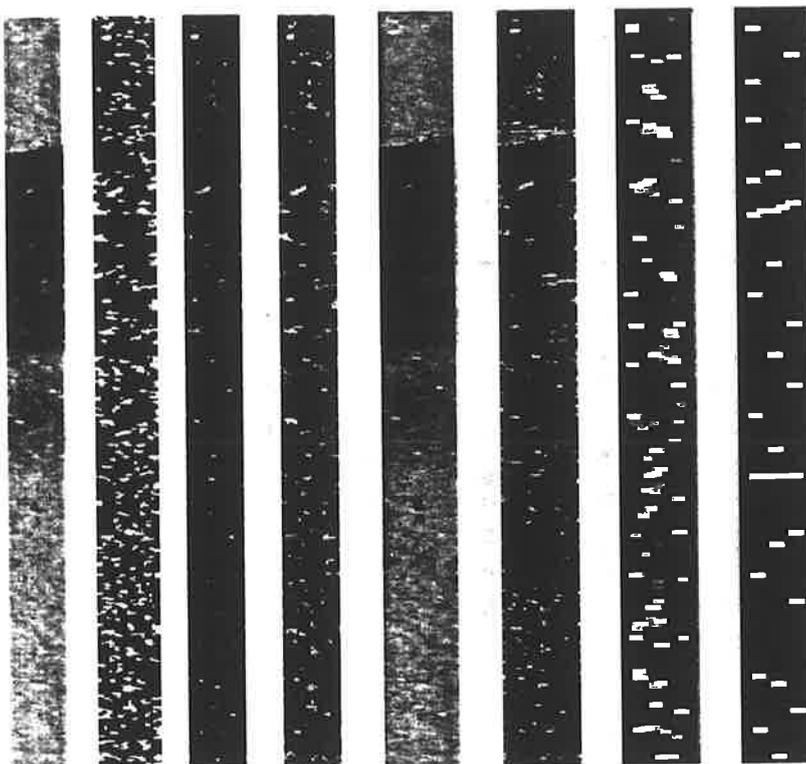


Figure 2. Infrared processing steps on GEODE data for both methods. Method 1: (a) resampling (b) blob search (c) local contrast enhancement (d) maximum. Method 2: (e) resampling (f) normalization (g) projection (h) Ground truth.

1. **resampling** : reduce the size of the infrared image to GEODE standard grid size by resampling the image. The resampling is done by selecting the local maximum in the image data for each grid cell.
2. **local positive contrast enhancement** : normalize the resulting resampled image to its local mean and local standard deviation (positive contrast). The window size used in the mean and standard deviation calculation is  $40 \times 40$  grid cells, corresponding with one square meter.
3. **multi-level thresholding** : create for each threshold  $t \in [0;255]$  a binary image of the resampled image. The result is a stack of binary images.
4. **multi-level opening and closing** : for each stack element (binary image) blobs are removed based on size and shape properties using image operators (binary openings and closings) from mathematical morphology [13]. The remaining blobs are reshaped into squares corresponding with the SCOOP size ( $8 \times 8$  grid cells).
5. **projection of resulting levels on confidence grid** : project all blobs in each stack element onto a confidence grid in which the confidence level is determined by the normalized intensity value of the blob.

The steps taken for method 2 are visualized in Figures 2(e) to 2(g). The false-alarm reduction is clearly visible as only a selected number of blobs remain present on the confidence grid.

## 5 Sensor-fusion methods

In this section we discuss our concept of decision-level sensor fusion. The advantage of decision-level fusion is that all knowledge about the sensors can be applied separately. Each sensor expert knows the most about the capabilities and limitations of their own sensor and can use this information to optimize the detection performance. The availability of this expert knowledge was the reason for choosing decision-level fusion for our application.

In our application, a fusion technique is considered to be a function that separates mines from background on the basis of the output of the different sensors. The output of each sensor is a measure of confidence in the presence of a mine and is called the confidence level.

The general layout of our concept of a sensor-fusion method is shown in Figure 3. The input of each sensor-fusion method is a confidence level per grid cell. A confidence level at a certain location expresses a confidence or belief in a mine detection on that position, but it is not necessarily a probability of detection

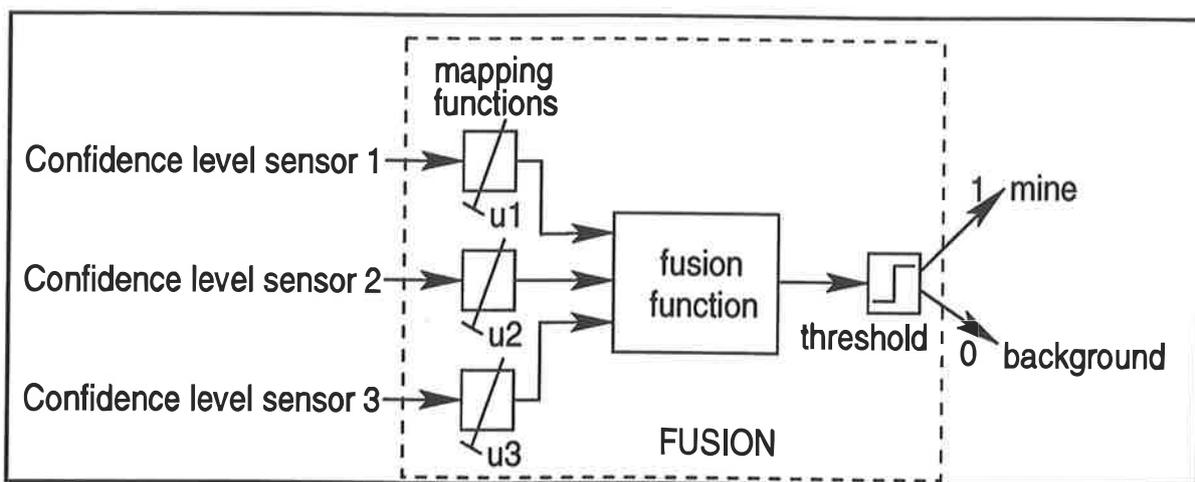


Figure 3. The generic decision-level sensor-fusion layout.

in a statistical sense. The confidence levels are used to indicate an order in probability of a detection of an object given a certain sensor. This means that a higher confidence level implies a higher probability of a mine, but these do not have to scale linearly.

The output of the fusion process is one for a detection and zero for no detection per grid cell. Each of the methods scales the influence of each of the sensors in a different way. This mapping requires one parameter ( $u_1; u_2; u_3$ ) per sensor. This mapping may remove the differences in definitions of the confidence levels. The mapped inputs are combined in a fusion function to acquire a single value per grid cell. The mapping functions and the fusion function are given in Table 3. For a more detailed description of the mapping and fusion functions we refer to [3, 4, 9, 7, 5, 11].

## 6 Experiments

In this section we show results of experiments with both processing methods on the infrared (3-5  $\mu\text{m}$ ) dataset as recorded within the GEODE project. Furthermore, we show results of sensor fusion using the infrared camera, GPR, and metal detector. The results of both processing methods and sensor-fusion algorithms are evaluated with the SCOOP algorithm as described in Section 3 and [8]. The implementation of the SCOOP algorithm improves the ROC curves, as the false alarm surface and the number of false alarms are both taken into account.

method	mapping function	fusion function
best sensor	none	selection
naive Bayes	linear scaling	product
Dempster-Shafer	uncertainty level	Dempster's rule of combination
rules	linear scaling	summation
fuzzy probabilities	fuzzy membership	minimum
voting	threshold	summation

Table 3. The different functions for scaling the input and combining these into a single (fused) result.

### 6.1 Infrared processing

In Figure 4(b) we have shown the ROC curve of processing methods 1 and 2 for the infrared dataset of GEODE. The GEODE dataset has about 50 percent surface mines (see Tables 1 and 2) which can easily be detected without almost any false alarms. Additionally, the ROC curve in case of no processing (confidence levels are raw sensor data) is presented for comparison. Figure 4(a) gives the ROC curves of the other sensors on the same test lane, the ROC curve of processing method 1 is added for comparison.

From Figure 4 we may conclude that the infrared sensor outperforms the other sensors on this particular dataset. Only the metal detector has an operating point at 64% detection with one false alarm that is comparable in performance.

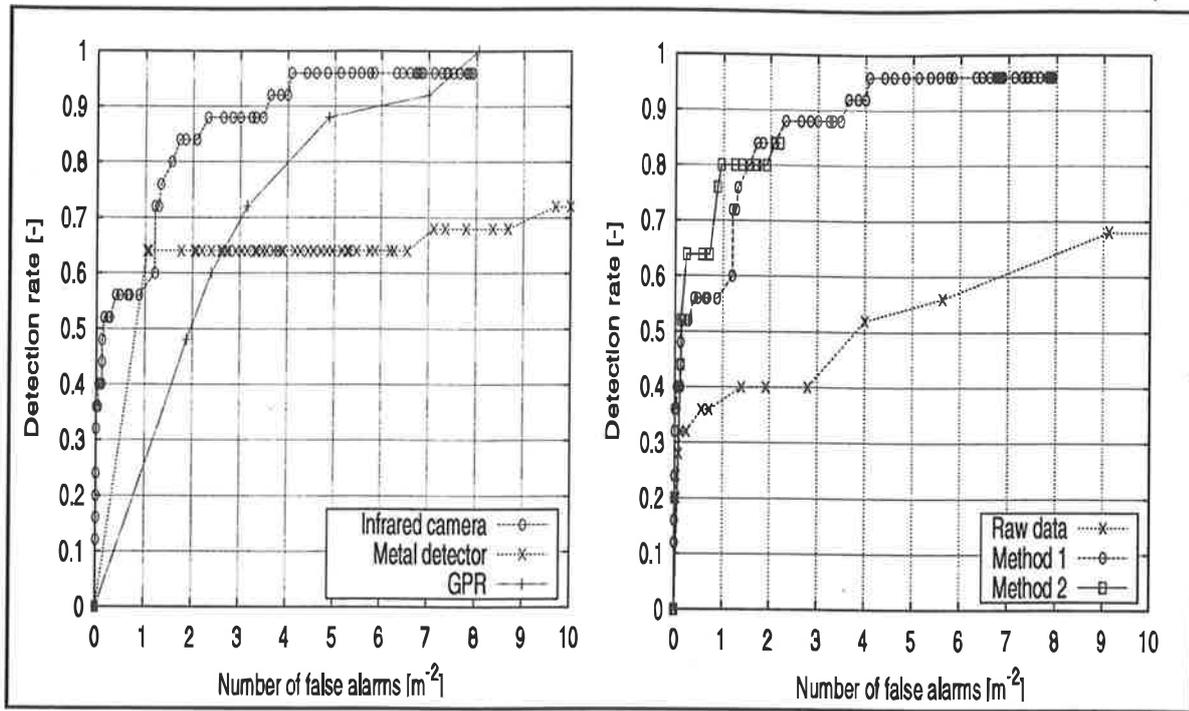


Figure 4. (a) ROC curves of the confidence levels of the infrared camera (processing method 1), metal detector and ground-penetrating radar. (b) ROC curves for infrared image processing of the GEODE dataset for no processing (global thresholding of raw sensory data) and methods 1 and 2.

## 6.2 Sensor fusion

For an unbiased comparison of sensor-fusion methods, with the limited data set we have, a leave-one-out evaluation method is used, see [14]. In the leave-one-out evaluation method, the parameters for each method are acquired on a training set, which contains all but one sample (a region containing one mine and on average 1 m<sup>2</sup>). The acquired parameters from the training set are tested on the single sample left out (the evaluation set). This is repeated for all mines and their surrounding region as evaluation set. The results from the training and evaluation sets are summed and normalized to acquire a detection rate and number of false alarms per m<sup>2</sup>.

The results of this leave-one-out evaluation in Figure 5 show that some of the sensor-fusion methods perform better than the best sensor on the evaluation set. Furthermore the method with most parameters, the rule-based method, performs best on the training set, but has the worst performance on the evaluation set. This performance loss is a confirmation of what we already expected, see [8]. The Dempster-Shafer implementation seems to be the most robust, while Fuzzy Probabilities gives very unpredictable results.

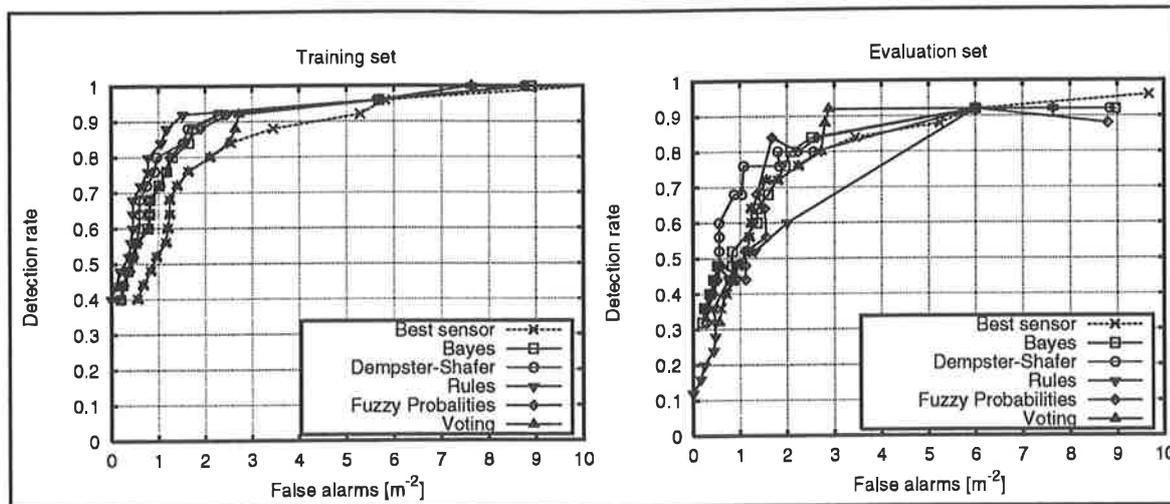


Figure 5. ROC curves for the different sensor-fusion methods evaluated with leave-one-out for the sensor combination low ground clearance GPR, metal detector, and infrared. The training set results are given in (a) and the evaluation set results are given in (b).

## 7 Conclusions

We have shown an increase of performance for the infrared processing for the GEODE dataset when compared with the original infrared processing. Our processing methods made the infrared sensor the best sensor for that particular dataset.

A point of concern is that the current implementations of processing methods are not computationally efficient and cannot, in their current state, perform in real time. We see, however, possibilities to increase processing speed by changing the implementation of the methods to an implementation based on DSP boards.

Another point of concern is the scenarios in which the infrared processing must operate. From the results of the GEODE and LOTUS [4] datasets we may conclude that infrared processing performs better for mines on the surface. We have the opinion that the processing may be optimized for different scenarios. Our design of the test lanes for the LOTUS trials is done in such a way that different scenarios and corresponding sketches can be used to determine the performance of infrared processing under different circumstances.

Concerning sensor fusion, the results of the independent training and evaluation sets, obtained by using a leave-one-out method, show that the Dempster-Shafer implementation performs better than the best sensor. The decrease in performance of the rule-based method is the largest, which is according to our expectations.

This method is clearly overtrained due to the many optimization parameters. The fuzzy probabilities method gives very unpredictable results. The voting fusion-method performs similar compared to the other methods.

The actual performance (if a very large data set is used) of these methods will be somewhere between the results of the training and evaluation sets. Our current experiments on the LOTUS datasets [4] can validate that assumption.

## 8 Acknowledgments

This research is partly funded by the European Union as ESPRIT project LOTUS, number 29812. The consortium consists of the following companies: Emrad limited, United Kingdom, Institute Dr. Förster, Germany,

THOMSOM-CSF DETEXIS, France and TNO Physics and Electronics Laboratory, The Netherlands.

## 9 References

[1] M. G. J. Breuers, P. B. W. Schwering, and S. P. van den Broek. Sensor fusion algorithms for the detection of land mines. In A. C. Dubey and J. F. Harvey, editors, *Proc. SPIE Vol. 3710, Detection and Remediation Technologies for Mines and Minelike Targets IV*, pages 1160–1166, Orlando (FL), USA, Apr. 1999.

[2] F. Cremer, E. den Breejen, and K. Schutte. Sensor data fusion for anti-personnel land-mine detection. In M. Bedworth and J. O'Brien, editors, *Proceedings of EuroFusion98. International Conference on Data Fusion*, pages 55–60, Great Malvern, UK, Oct. 1998.

[3] F. Cremer, J. G. M. Schavemaker, E. den Breejen, and K. Schutte. Detection of anti-personnel landmines using sensor-fusion techniques. In T. Windeatt and J. O'Brien, editors, *Proceedings of EuroFusion99. International Conference on Data Fusion*, pages 159–166, Stratford Upon Avon, UK, Oct. 1999.

[4] F. Cremer, J. G. M. Schavemaker, E. den Breejen, and K. Schutte. Towards an operational sensor fusion system for anti-personnel landmine detection. In A. C. Dubey and J. F. Harvey, editors, *Proc. SPIE Vol. 4038, Detection and Remediation Technologies for Mines and Minelike Targets V*, Orlando (FL), USA Apr. 2000.

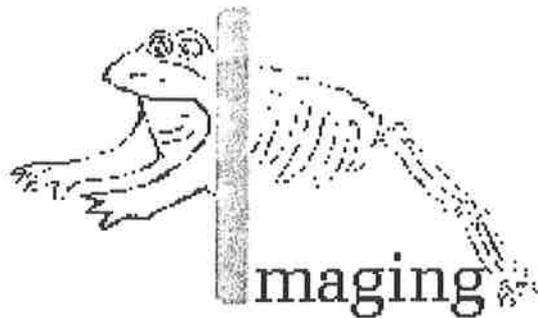
[5] B. V. Dasarathy. Decision fusion benefits assessment in a three-sensor suite framework. *SPIE Optical Engineering*, 37(2):354–369, Feb. 1998.

[6] W. de Jong, F. Cremer, K. Schutte, and J. Storm. Usage of polarisation features of landmines for improved automatic detection. In A. C. Dubey and J. F. Harvey, editors, *Proc. SPIE Vol. 4038, Detection and Remediation Technologies for Mines and Minelike Targets V*, Orlando (FL), USA, Apr. 2000.

[7] A. P. Dempster. Upper and lower probabilities induced by a multi-valued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.

- [8] E. den Breejen, K. Schutte, and F. Cremer. Sensor fusion for anti personnel landmine detection, a case study. In A. C. Dubey and J. F. Harvey, editors, *Proc. SPIE Vol. 3710, Detection and Remediation Technologies for Mines and Minelike Targets IV*, pages 1235–1245, Orlando (FL), USA, Apr. 1999.
- [9] K. Fukunaga. *Introduction to statistical pattern recognition*. Academic press, Inc., Boston, USA, 2<sup>nd</sup> edition, 1990.
- [10] C. King, editor. *Jane's mines and mine clearance*, volume 98-99. Jane's Information Group, Surrey, UK, 3rd edition, 1998.
- [11] L. A. Klein. A boolean algebra approach to multiple sensor voting fusion. *IEEE transactions on Aerospace and Electronic systems*, 29(2):317–327, Apr. 1998.
- [12] P. B. W. Schwing, M. G. J. Breuers, and S. P. van den Broek. Sensor fusion algorithms for the detection of land mines. In *Proc. of Third NATO IRIS Joint Symposium*, pages 365–371, Quebec, Canada, Oct. 1998.
- [13] J. Serra. *Image Analysis and Mathematical Morphology*. Academic Press, 1982.
- [14] S. M. Weiss and C. A. Kulikowski. *Computer systems that learn*. Morgan Kaufmann Publishers, 1991.

Proceedings  
symposium  
Imaging  
*making the invisible visible*



Thursday 18 May 2000  
Technische Universiteit Eindhoven



2000 years

**IEEE** student  
branch  
**EINDHOVEN**